UNIVERSITY OF CALIFORNIA
Los Angeles

Quantum Information with Semiconductor Electron Spins

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Electrical Engineering

by

Thomas Szkopek

2006

The dissertation of Thomas Szkopek is approved.

Hong-Wen Jiang

Vwani P. Roychowdhury

Kang L. Wang

Eli Yablonovitch, Committee Chair

University of California, Los Angeles

2006

Od wieków wszystkim wiadomo...

kwitnie paproć...

TABLE OF CONTENTS

List of Figures

List of Tables

ACKNOWLEDGMENTS

First and foremost, I am indebted to Prof. Yablonovitch for the scientific direction and financial support required to undertake this work. Above all else, I have learned from him how easy it is to focus in on the fundamental physics in a situation where obfuscating details loom. I have not only gained scientific knowledge, but I have also gained insight into the ways and means of scientific endeavour.

I extend my gratitude to Prof. Roychowdhury for his Aristotelian and Sophoclean teachings on information theory, in the formal setting of a class room and the informal setting of office "chats". The theoretical work presented herein could not have come to fruition without his counsel and criticism.

It is my pleasure to thank Prof. Jiang, who has shown me the beauty of experimental physics below 100μeV temperatures. I am indebted to him for providing essential cryogenic facilities during the last months of my research, and even more so for sharing his experimental expertise and enthusiasm.

I kindly thank Prof. Wang for serving on my PhD committee.

There exist collaborators who have contributed to the work presented here. Hans D. Robinson and Deepak S. Rao refined many semiconductor fabrication procedures and were responsible for much of the initial experimental effort. All semiconductor fabrication was done at the UCLA Nanoelectronics Research Facility where Steve Franz, Hoc Ngo, Huynh Do, and Ivan Alvarado-Rodriguez were conspicuously helpful. Additionally, P. Oscar Boykin, Heng Fan, Daniel Abrams, and Mark Gyure were excellent collaborators.

I have been the beneficiary of interaction at various levels with Seth Lloyd, Stephen A. Lyon, Daniel Gottesman, Karoly Holczer, David Goldhaber-Gordon,

Sir John Pendry and Andy Sachrajda.

Last but certainly not least, it is my pleasure to acknowledge the direct and indirect financial support from the Army Research Office, the Defense Advanced Research Projects Agency, the Defense Microelectronics Activity, the Centre for Nanoscience Innovation for Defense, the Microelectronics Advanced Research Corporation (Materials, Structures, Devices Focus Center), the California NanoSystems Institute and the UC Regents.

VITA

1976	Born, Warszawa, Poland.
1999	B.A.Sc. (Engineering Science), University of Toronto.
2001	M.A.Sc. (Electrical Engineering), University of Toronto.
2001–2006	Research Assistant, Dept. of Electrical Engineering UCLA.

PUBLICATIONS

T. Szkopek, P.O. Boykin, H. Fan, V. Roychowdhury, E. Yablonovitch, G. Simms, M. Gyure, B. Fong, “Threshold error penalty for fault tolerant computation with nearest neighbour communication”, IEEE Trans. Nanotech., vol. **5**, no.1, Jan 2006, pp. 42-49 [quant-ph/0411111].

D.S. Rao, T. Szkopek, H.D. Robinson, E. Yablonovitch, and H.W. Jiang, “Single photo-electron trapping, storage, and detection in a one-electron quantum dot”, J. Appl. Phys. **98**, 114507 (2005) [quant-ph/0410094].

T. Szkopek, V. Roychowdhury, E. Yablonovitch and D.S. Abrams, “Eigenvalue estimation of differential operators with a quantum algorithm”, Phys. Rev. A. **72**, 062318 (2005) [quant-ph/0408137].

H. Fan, V. Roychowdhury, and T. Szkopek, “Optimal two-qubit quantum circuits using exchange interactions”, Phys. Rev. A. **72**, 052323 (2005) [quant-ph/0410001].

E. Yablonovitch, H.W. Jiang, H. Kosaka, H.D. Robinson, D.S. Rao, and T. Szkopek, “Optoelectronic quantum telecommunications based on spins in semiconductors”, Proc. of the IEEE, vol. **91**, no.5, May 2003, pp.761-80.

Y. Sun, T. Szkopek, and P.W.E. Smith, “Demonstration of narrowband high-reflectivity Bragg gratings in a novel multimode fiber”, Optics Comm., vol. **223**, no. 1-3, July 2003, pp. 91-5.

T. Szkopek, V. Pasupathy, J.E. Sipe, and P.W.E. Smith, “Novel multimode fiber for narrow-band Bragg gratings”, IEEE. J. Sel. Top. Quantum Electron., vol. **7**, no. 3, May 2001, pp. 425-33.

Abstract of the Dissertation

Quantum Information with Semiconductor Electron Spins

by

Thomas Szkopek

Doctor of Philosophy in Electrical Engineering
University of California, Los Angeles, 2006
Professor Eli Yablonovitch, Chair

The study of quantum information science has brought a wealth of open questions into the view of the scientific community: it is the aim of this work to address three of these questions.

The first question is: can quantum algorithms be extended for efficient simulation of “classical” physics, such as electromagnetics simulations? Existing quantum algorithms, particularly Shor’s number factorization and Lloyd’s quantum many-body emulator, make use of efficient eigenvalue estimation. We show that a quantum algorithm can be extended to eigenvalue estimation of linear, partial differential operators. It is found that scaling better than existing classical algorithms can be achieved only for problems defined over domains of high dimensionality (>4 for second order differentials). The answer is thus negative for electromagnetics; the fundamental reason is the unitarity of quantum mechanics and hence the inability to “erase” errors during convergence.

The second question is: what are the tolerable error rates for a quantum computer constructed from qubits that only exhibit nearest neighbour interactions?

In many solid state qubit proposals, including electron spins in semiconductor quantum dots, the communication scheme may very well be limited to nearest neighbour interactions. We find that in the worst case scenario - with regards to communication - of a linear stripe of qubits of negligible width, the impact on the error correction threshold for the concatenated [[7,1,3]] Calderbank-Shor-Steane code is merely one order of magnitude in the amplitude or timing of control signals compared to idealized error-free communication models. The overhead associated with error correction itself is substantial enough to limit the impact of additional nearest-neighbour communication operations.

The third and final question is: is it experimentally feasible to implement long distance communication with semiconductor electron spin qubits. A critical component for long distance communication of qubits is a spin coherent photodetector that maps the arbitrary spin state of an incident photon to the spin state of a photoelectron in a fully quantum coherent fashion. The trapping and storage of single photoelectrons is reported here, a significant step towards realizing spin coherent photodetection.

CHAPTER 1

Motivation

Interest in quantum information arose from several theoretical developments, but that which garnered the most attention was Shor's discovery of a quantum algorithm for factoring large numbers [1] more efficiently than the best known classical algorithm - the number field sieve [2] - although it remains to be proved that a more efficient classical algorithm does not exist. The technological significance of Shor's algorithm arises from its potential to allow the breaking of the widely used Rivest-Shamir-Adelmann (RSA) cryptosystem whose security relies upon our current inability to factor large numbers efficiently, as detailed in reviews on the subject [3, 4, 5]. The asymptotic complexity of various implementations of Shor's algorithm are summarized in Table 1.1 for factoring an N digit number. The $O[\text{poly}(N)]$ scaling of Shor's algorithm is clearly advantageous compared with the $O[\exp(\sqrt[3]{N})]$ scaling of the classical number field sieve. The desire to construct specialized hardware to implement Shor's algorithm is the prime motivating factor for applied research in quantum information science, including the work presented here.

1.1 "The" Quantum Algorithm

All quantum algorithms discovered thus far, including Shor's and Lloyd's [9, 10, 11] as well as those of purely theoretical interest (including Grover's [12] and

Table 1.1: Asymptotic scaling of complexity of various implementations of Shor's algorithm for factoring an N digit number under non-local communication. Each is asymptotically superior to the classical number field sieve number, which requires $O[\exp(\sqrt[3]{N})]$ gate operations. The Cleve & Watrous implementation requires $O(N)$ classical operations for pre and post processing.

Reference	Qubits	Operations	Depth
Vedral et al. [6]	$O(N)$	$O(N^3)$	$O(N^3)$
Gossett [7]	$O(N^2)$	$O(N^3)$	$O(N \log N)$
Cleve & Watrous [8]	$O(N^3)$	$O(N^3)$	$O[(\log N)^2]$

Deutsch-Jozsa's [13]) can be viewed as different instances of a single algorithm. Cleve et al. [14] demonstrated that quantum algorithms can all be cast in the form of a multiple-particle interferometer for estimating the eigenvalues of unitary operators. The core of the quantum algorithm is illustrated in Figure 1.1, after [14]. It is in the choice of unitary operator that one specifies the problem that one wishes to solve, which may require great ingenuity.

The estimation of eigenvalues is of course of great practical importance in engineering practice, including the determination of resonance frequencies in continuum mechanics and electromagnetism. It is natural to ask whether the quantum algorithm can be applied to any advantage to these common engineering problems. This question is investigated in Chapter 2 of this dissertation, where the quantum algorithm is applied to eigenvalue estimation of linear partial differential equations. It was found that for differential equations of order $2S$ with eigenfunctions $\psi(x_1, x_2, \ldots, x_D)$ of D arguments, the computational cost required to estimate a low order eigenvalue to accuracy $\Theta(1/N^2)$ is $\Theta((2(S+1)(1+1/\nu)+D)\log N)$ qubits and $O(N^{2(S+1)(1+1/\nu)} \log^c N^D)$ gate oper-

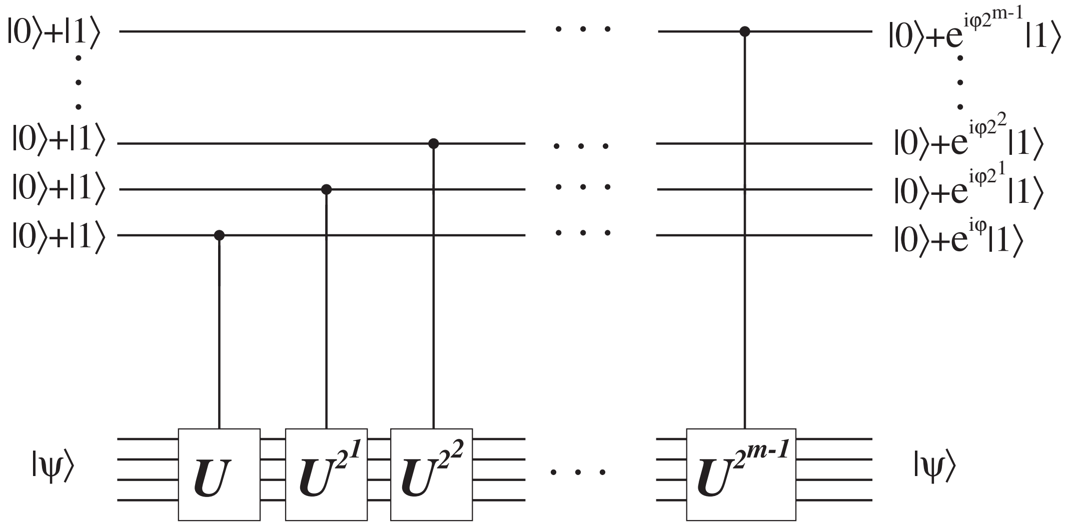

Figure 1.1: The heart of the quantum algorithm is the accumulation of phase increments φ from a unitary operator U acting upon some initially prepared eigenstate $|\psi\rangle$ of U. A computational problem (such as Shor's factorization) is solved by a selection of U and/or $|\psi\rangle$ such that a useful computational result is given by φ. Index register qubits prepared in superposition states $|0\rangle + |1\rangle$ are used to pick up phase differences through conditional applications of U, which can finally be measured by Fourier analysis.

ations, where N is the number of points to which each argument is discretized, ν and c are implementation dependent constants of $O(1)$. Optimal classical methods require $\Theta(N^D)$ bits and $\Omega(N^D)$ gate operations, implying that a quantum algorithm can give improved convergence if the simulation domain has sufficiently large dimension $D > 2(S+1)(1+1/\nu)$.

In the of case second-order differential equations,$2S = 2$, appropriate to resonance problems in electromagnetics and classical continuum mechanics, a quantum algorithm improvement over existing classical methods requires a domain of dimension $D > 4$. Of course, the domain of electromagnetics and continuum mechanics problems is 3+1 (space+time) dimensions, ruling out an advantage over existing classical computational techniques. Thus, to date, Shor's algorithm remains the prime motivating factor for the development of hardware for quantum information processing. Concomitantly, the search for new applications of the quantum algorithm, and more generally new applications of quantum information, remains an open problem.

1.2 Quantum Computer Architecture

In order to develop hardware for quantum information processing, one must cope with the inevitable noise and decoherence arising from interaction with the environment [15]. The means to achieve reliable computation is through fault-tolerant architecture [16] wherein errors induced by noise and decoherence are corrected, and the probability of a logical error is bounded by a constant. This is completely analogous to the bounded error in digital computation where the probability of error is bounded by a constant rather than growing with successive operations as in analogue computation. *Unlike* classical digital computation, it is unknown how to use inherent physical redundancy (such as the multiple electrons in a transistor circuit) to create reliable quantum circuits (with the possible exception of recently discovered topological quantum computation [17]). Rather, the most studied approach to fault-tolerant quantum computation is through the use of error correcting codes, an important class being the Calderbank-Shor-Steane (CSS) family of codes [18, 19]. Logical qubits are encoded with multiple physical qubits in such a manner that errors can be detected and corrected without destruction of the quantum coherence of the logical qubit necessary for quantum computation.

An important theoretical finding for fault-tolerant quantum computation is the discovery that an error threshold exists [20]. If the probability of a physical error during a gate operation, during storage, or during communication of a qubit is below some threshold value, then fault-tolerant computation with any constant error probability is possible by using a sufficiently redundant concatenated encoding (where concatenation is simply the recursive encoding of qubits).

Unlike classical computation, where the reliable communication of information across an integrated circuit is taken for granted, reliable communication of quantum information is much more difficult due to the need to preserve quantum coherence. Generally, the probability of a communication error increases with distance. In the case of solid state quantum computing proposals, such as electron or nuclear spins in semiconductors [21, 22, 23], the only means for communication of quantum information is via nearest-neighbour interactions (pending the development of new hardware such as a quantum bus [24]). It was shown by Gottesman [25] that even in the presence of communication error that linearly increases with distance, an error threshold still exists, although the value of the error threshold was not determined.

In Chapter 3, the dependence of the error threshold on communication error is investigated for the situation of a narrow stripe of qubits (the stripe width being no greater than the logarithm of the stripe length). The error thresholds are numerically estimated for concatenation of the [[7,1,3]] CSS code (ie. a code of 7 qubits representing 1 logical qubit with a minimum of 3 qubit errors required for logical error), under assumptions of error-free communication, communication via teleportation, and communication by nearest neighbour swapping. It was found that the error threshold of 2.1×10^{-5} without any communication error is reduced to 1.2×10^{-7} with nearest neighbour communication errors. This $\sim$ 175X penalty in error probability threshold translates to an $\sim$ 13X penalty in qubit in timing and amplitude accuracy of qubit control pulses. It is found therefore that nearest-neighbour communication only has a "moderate" effect upon the error threshold. This finding has been corroborated by independent investigations [26, 27] of error scaling in two dimensional arrays of qubits restricted to nearest neighbour communication.

1.3 Semiconductor Spin Qubits

The spins of electrons confined to semiconductor quantum dots have been proposed as quantum hardware [21, 22, 23]. Despite the variation in details between proposals, the common idea is to use the spin-1/2 degree of freedom of localized electrons as qubits. The appeal of a semiconductor solution is founded on the hope that the historically successful scaling of semiconductor transistor technology for classical computing, embodied by Moore's Law [28], would apply to the scaling of a quantum computer as well.

The criteria that any qubit technology must satisfy were outlined by DiVincenzo [29]:

- scalable physical system with well defined qubits
- ability to initialize to a fiducial state
- decoherence times long compared to gate operations
- universal set of quantum gates
- measurement capability of specific qubits
- ability to interconvert “stationary” and “flying” qubits
- ability to transmit “flying” qubits faithfully

These are of course bare minimum requirements, and other practical issues must be considered, such as the absolute speed of gate operations for example. General questions regarding the suitability of electron spins in semiconductors as qubits are beyond the scope of this dissertation [30, 31].

A number of important experimental advances have been made towards realizing a logical qubit from electron spins in semiconductors: single-shot measurement of an electron spin state [32], averaged measurement of single electron spin resonance [33], and averaged measurement of Heisenberg exchange interaction between neighbouring electron spins [34]. Another area of active research is that of quantum communication, namely the interconversion of “flying” qubits (photon spins) to “stationary” qubits (electron spins). Research in this area is motivated by anticipated communication needs for distributed quantum computation [35] as well as long distance quantum cryptography [36].

One of the key processes required for long distance quantum communication is the ability to transfer information from photon spins to electron spins. It was proposed [37] that semiconductor optoelectronics could be used to transfer quantum information from the spin (polarization) of a single photon to the spin of a single localized electron in a fully quantum coherent manner (meaning without spin measurement taking place at any point in the process). Specifically, the desired transformation is one that leaves photoelectron and photohole unentangled,

$$\alpha|\sigma_+\rangle + \beta|\sigma_-\rangle \rightarrow \left(\alpha|+\rangle + \beta|-\rangle\right)_{\mathrm{e}} |\mathrm{arb}\rangle_{\mathrm{h}} \tag{1.1}$$

where α and β are probability amplitudes of the incident photon in a circular basis $\{|\sigma_+\rangle, |\sigma_-\rangle\}$, while the electron spin basis is $\{|+\rangle, |-\rangle\}$ and the hole state $|\mathrm{arb}\rangle$ is arbitrary.

The energy states and associated selection rules allowing such a transfer to take place are illustrated in Figure 1.2, based on the work of Vrijen & Yablonovitch [37]. Preliminary studies indicate that it is theoretically possible to obtain the desired energy states and selection rules in a strained quantum well heterostructure with the InGaAsP family of materials [38]. One might equally well consider the transfer of quantum information from photon spin to photoholes, but it is expected that the spin coherence lifetime of electrons exceeds that of holes [31] on account of the lesser spin-orbit coupling in the s-orbital like conduction band compared to the p-orbital like valence band.

In Chapters 4 and 5 we report experimental work towards realizing spin coherent photodetection. In particular, we describe the experimental observation of the trapping, storage and detection of single photoelectrons in gate electrode defined quantum dots [39]. Recent work towards demonstration of spin transfer on a single shot basis is described in Chapter 5. Concluding remarks on future research directions for quantum information in semiconductor spins are collected in Chapter 6.

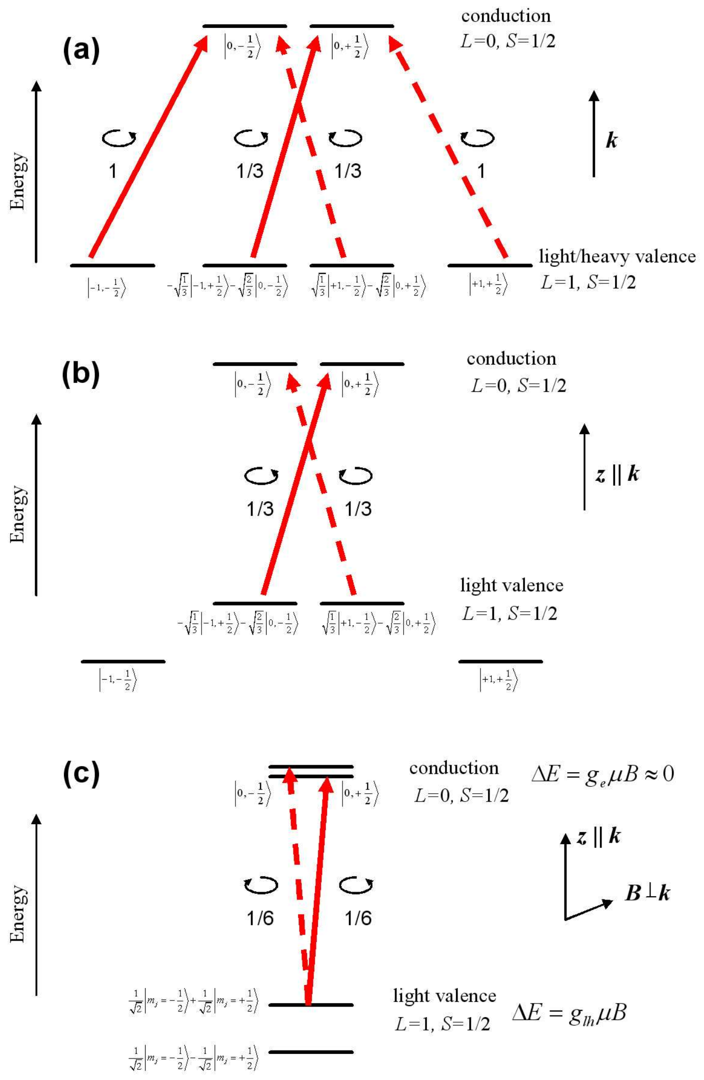

Figure 1.2: The selection rules for band edge transitions in a direct gap semiconductor that will allow spin coherent photodetection. (a) The bulk band edge states and the optical transitions permitted for optical excitation with wavevector k. (b) By the application of strain along the axis z colinear with optical axis k, the heavy hole states can be shifted and the light hole states can be spectroscopically selected. (c) The application of a large magnetic field B normal to the strain axis / optical axis will mix the light hole states. Selecting semiconductors with zero electron Landé g-factor and as large a hole Landé g-factor will allow both electron spin states to be accessed from a single.

CHAPTER 2

Eigenvalue Estimation of Differential Operators with a Quantum Algorithm

We demonstrate how linear differential operators could be emulated by a quantum processor, should one ever be built, using the Abrams-Lloyd algorithm. Given a linear differential operator of order $2S$, acting on functions $\psi(x_1, x_2, \ldots, x_D)$ with D arguments, the computational cost required to estimate a low order eigenvalue to accuracy $\Theta(1/N^2)$ is $\Theta((2(S+1)(1+1/\nu)+D)\log N)$ qubits and $O(N^{2(S+1)(1+1/\nu)}\log^c N^D)$ gate operations, where N is the number of points to which each argument is discretized, ν and c are implementation dependent constants of $O(1)$. Optimal classical methods require $\Theta(N^D)$ bits and $\Omega(N^D)$ gate operations to perform the same eigenvalue estimation. The Abrams-Lloyd algorithm thereby leads to exponential reduction in memory and polynomial reduction in gate operations, provided the domain has sufficiently large dimension $D > 2(S+1)(1+1/\nu)$. In the case of Schrödinger's equation, ground state energy estimation of two or more particles can in principle be performed with fewer quantum mechanical gates than classical gates.

2.1 Introduction

An early motivation for research in quantum information processing has been the simulation of quantum mechanical systems [40]. The Abrams-Lloyd algorithm [9, 10, 11] is an instance of quantum mechanical simulation (followed by variations [41], [42], [43]). We describe in this paper the application of the Abrams-Lloyd algorithm to estimating low order eigenvalues of linear partial differential equations with homogeneous boundary conditions (more precisely, Hermitian boundary value problems). The significance of our analysis is two fold. First, we generalize the Abrams-Lloyd algorithm to boundary value problems other than Schrödinger's equation, which may find application to classical problems. Secondly, we quantify computational cost and determine under what conditions we may expect the Abrams-Lloyd algorithm to give a reduction in computational work compared to optimal classical techniques in order to achieve the same eigenvalue accuracy.

Very briefly, the Abrams-Lloyd algorithm as originally envisaged for the many-body Schrödinger equation is structured as follows. An initial estimate $|\psi(0)\rangle$ of the target eigenstate is loaded into a multiple qubit register. Controlled application of a unitary operation, chosen to correspond to the time evolution operator $\exp(-i\mathcal{H}\tau)$ of the many-body Hamiltonian $\mathcal{H}$ under study for time step τ, allows one to generate a sequence of time evolved states originating from the initial guess, $\{|\psi(0)\rangle, |\psi(\tau)\rangle, |\psi(2\tau)\rangle, \ldots\}$. A spectral analysis of the sequence of time evolved states recovers the frequency (energy) of the target eigenstate (provided the initial guess was "close enough"). The Abrams-Lloyd algorithm is akin to a stroboscope for quantum states evolving under a many-body Hamiltonian. If the total time of evolution is sufficiently long, while the individual time steps are sufficiently small, a high frequency (energy) resolution can be achieved. Following the determination of the eigenvalue, the corresponding eigenstate remains in the qubit register. Although the full amplitude description of an eigenstate is inaccessible, some information about the state can be extracted to a precision ultimately limited by the number of qubits used to represent the eigenstate (so for instance, one can test symmetries of the eigenstate).

The algorithm can be extended to more general partial differential equations rather easily. So long as the boundary value problem is Hermitian, we can map our mathematical problem to a fictional quantum system and apply the algorithm without change. The partial differential operator, $\mathcal{L}$, will correspond to a (possibly) fictional Hamiltonian $\mathcal{H}$, and an initial guess $\psi(0)$ will correspond to an initial wavefunction $|\psi(0)\rangle$. In other words, quantum mechanical amplitudes represent function values. Less obviously, the sequence of time evolved states, $\{|\psi(0)\rangle, |\psi(\tau)\rangle, |\psi(2\tau)\rangle, \ldots\}$ has a mathematical analogue of great use in classical matrix eigenvalue analysis, known as the Krylov subspace: $\text{span}\{\psi(0), \exp(-i\mathcal{L}\tau)\psi(0), \exp(-i2\mathcal{L}\tau)\psi(0), \ldots\}$. The subspace is generated by repeated application of $\exp(-i\mathcal{L}\tau)$, although in classical techniques one more typically uses rational functions of $\mathcal{L}$. Here, τ no longer has the physical meaning of time. Rather, τ sets the scale for how much phase is applied per application of $\exp(-i\mathcal{L}\tau)$. As in the quantum simulation, a large total phase applied one small phase step at a time allows a high resolution estimate of eigenvalues. We quantify these notions now.

The computational cost of the Abrams-Lloyd algorithm for a specified eigenvalue accuracy is limited as a consequence of three sources of error, expressed here in the language of quantum mechanical simulation:

I *truncation error*: Discretization is necessary for a computational model based on qubits. However, discretization of the continuous problem to N

points per coordinate results in $\Theta(1/N^2)$ relative error in low order energy eigenvalues due to truncation of high spatial frequency contributions. The choice of N must be made appropriate to the accuracy that is desired.

II *splitting error*: The full many-body evolution $\exp(-i\mathcal{H}\tau)$ over time step τ can be implemented with universal gates by splitting the full evolution into a sequence of efficiently implementable unitaries $\exp(-i\mathcal{H}^k\tau)$, where $\mathcal{H} = \sum_k \mathcal{H}^k$. The approximation results in an absolute eigenvalue error $O(\|\mathcal{H}\|_2^{\nu+1}\,\tau^\nu)$, where $\|\mathcal{H}\|_2$ is the maximum eigenvalue of the discretized Hamiltonian $\mathcal{H}$ and ν is a constant of $O(1)$ determined by the precise sequence of local operators chosen. Splitting error requires us to use small time steps τ.

III *frequency resolution*: A quantum Fourier transform, like any discrete Fourier transform, can resolve absolute phase to accuracy at best $\pm\pi$. For a sequence of M samples, the relative error in an energy eigenvalue E will be $\pm\pi/(ME\tau)$. Frequency resolution requires us to simulate over a large total time $M\tau$.

The optimal way to balance these errors is as follows. Since we are interested in the continuous problem, we first choose a discretization of N points per sample so that the discrete problem eigenvalue approximates the continuous eigenvalue problem to some desired accuracy $\Theta(1/N^2)$. We wish to solve the discretized problem to an accuracy determined by the truncation error limit; solving the discrete problem to greater accuracy leads to wasted effort since we are interested in the continuous problem, while solving the discrete problem to lesser accuracy implies we have wasted effort by choosing too many discrete points N per coordinate. We can thereby determine the maximum time step τ to keep splitting error no greater than truncation error. Next, we can determine the number of time steps M required to resolve eigenvalues with the quantum Fourier transform at the truncation error limit. In the case of Hermitian boundary value problems, We show in this paper that the resulting computational cost is $\Theta(D \log N)$ qubits and $O(N^{2(S+1)(1+1/\nu)} \log^c N)$ gate operations, where $2S$ is the differential order of $\mathcal{L}$ and c is a constant $O(1)$. This can be compared with the optimal classical cost of $\Theta(N^D)$ bits and $\Omega(N^D)$ gate operations. Near optimal classical methods approaching these costs do in fact exist [1].

We emphasize that in our analysis, we take a constructive approach wherein we account for all the logical operations required to implement the algorithm without recourse to oracles that may or may not have physically efficient implementations. This is in contrast to previous work including the simulation of spin glass physics [47], and Sturm-Liouville problems ([48] and references therin). As stated, our motivation is to compare the computational cost of eigenvalue estimation by the Abrams-Lloyd algorithm and optimal classical methods.

Our paper is organized as follows. In section 2.2, we introduce the one dimensional eigenvalue problem, which will serve as a useful example with which the principles of the algorithm can be illustrated. We derive the truncation error for low order eigenvalues in a way suitable for extension to higher dimensional problems. The algorithm itself is described in section 2.3, followed by an analysis of computational cost as it is applied to the one dimensional problem in section 2.4. A circuit suitable for a 2nd order differential equation is given as a con-

---

[1] A near optimal classical method can be constructed using a combination of Krylov subspace iteration, matrix preconditioning and multigrid solutions, as in [44]. See [45],[46] for a sampling of the vast array of classical numerical techniques available

crete example. Generalization of the algorithm to higher dimensional problems is given in section 2.5 along with an analysis of computational cost, where we show a reduction in computational work polynomial in N over classical techniques. Concluding remarks about the computational efficiency of the Abrams-Lloyd algorithm are given in section 2.6.

2.2 One-Dimensional Problem

To illustrate the essential features of eigenvalue estimation of differential operators, it's instructive to consider a Hermitian one-dimensional problem, which we introduce here in some detail. The primary result of this section is a derivation of truncation error in low order eigenvalues as a function of discretization. Much of the notation used throughout this paper are defined in this section. We begin with a linear, $2S$-order differential operator $\mathcal{D}$ that maps a complex valued function $\psi(x)$, $x \in [0,1]$ to a new function according to the rule,

$$\begin{aligned} \mathcal{D}\psi(x) &= \sum_{s=0}^{S} \frac{\partial^s}{\partial x^s}\left(a_s(x)\frac{\partial^s \psi(x)}{\partial x^s}\right) \\ &= a_0(x)\psi(x) + \frac{\partial}{\partial x}\left(a_1(x)\frac{\partial \psi(x)}{\partial x}\right) + \dots \\ &\quad + \frac{\partial^S}{\partial x^S}\left(a_S \frac{\partial^S \psi(x)}{\partial x^S}\right), \end{aligned} \tag{2.1}$$

where we assume $\psi(x)$ has finite derivatives up to order $2S$. The coefficients $a_s(x)$, $s = 0, 1, 2, \dots, S$ are finite, real valued functions on the domain $x \in [0,1]$ with finite derivatives to order s and satisfy periodic boundary conditions,

$$\frac{\partial^t a_s}{\partial x^t}(0) = \frac{\partial^t a_s}{\partial x^t}(1) \qquad t = 0, 1, \dots, s \tag{2.2}$$

The minimal smoothness assumed of $a_0(x)$ is continuity on $x \in [0,1]$. For concreteness, we impose periodic boundary conditions upon $\psi(x)$ itself,

$$\frac{\partial^t \psi}{\partial x^t}(0) = \frac{\partial^t \psi}{\partial x^t}(1) \qquad t = 0, 1, \ldots, 2S \tag{2.3}$$

although more general homogeneous boundary conditions could be insisted upon. Given the above definitions, a set of eigenfunctions $\phi_f(x)$ with corresponding real eigenvalues λ_f is defined through,

$$\mathcal{D}\phi_f(x) = \lambda_f \phi_f(x), \tag{2.4}$$

and we order the eigenvalues λ_f, $f = 1, 2, 3, \ldots$ in ascending order $\lambda_1 \leq \lambda_2 \leq \lambda_2 \leq \ldots$. The definition and boundary conditions in Eqs. 2.1-2.3 guarantee a Hermitian $\mathcal{D}$, meaning $\int_0^1 dx(\phi_f \mathcal{D}\phi_{f'} - \phi_{f'}\mathcal{D}\phi_f) = 0$ for any pair of eigenfunctions $\phi_f, \phi_{f'}$. All the usual eigenvalue/eigenfunction properties of Hermitian operators follow. Our task is to estimate a low order ($f = O(1)$) eigenvalue λ_f.

The most useful expression of the eigenvalue is the Rayleigh quotient,

$$\lambda_f = \int_0^1 dx \phi_f^*(x) \mathcal{D} \phi_f(x) = \int_0^1 dx \phi_f^*(x) \mathcal{L} \phi_f(x) \tag{2.5}$$

where we impose unity L_2 norm on the eigenfunctions,

$$\|\phi_f\|_{L_2} = \left(\int_0^1 dx \phi_f^*(x) \phi_f(x) \right)^{1/2} \tag{2.6}$$

in anticipation of the quantum algorithm and the operator $\mathcal{L}$, derived from $\mathcal{D}$ by simple integration by parts, is a more convenient (bilinear) operator to work with due to its symmetric form,

$$\varphi^*(x) \mathcal{L} \psi(x) = \sum_{s=0}^{S} \frac{\partial^s \varphi^*(x)}{\partial x^s} a_s(x) \frac{\partial^s \psi(x)}{\partial x^s} \tag{2.7}$$

for any two functions $\varphi(x)$ and $\psi(x)$.

It is useful to work not only in the "space" domain $x \in [0,1]$, but also in the "reciprocal space" domain of integers, $k \in \mathbb{Z}$. The connection between the two representations is defined by the Fourier transforms,

$$\begin{aligned} \widetilde{\psi}_k &= \int_0^1 dx \exp(-2\pi i k x)\psi(x), \\ \psi(x) &= \sum_{k=-\infty}^{\infty} \exp(2\pi i k x)\widetilde{\psi}_k, \end{aligned} \tag{2.8}$$

where tilde will indicate a reciprocal space representation throughout the paper. Our eigenvalue Eq. 2.4 is Fourier transformed to

$$\sum_{k'=-\infty}^{\infty} \widetilde{\mathcal{L}}_{k,k'}\widetilde{\phi}_{f,k'} = \lambda_f \widetilde{\phi}_{f,k}, \tag{2.9}$$

where,

$$\widetilde{\mathcal{L}}_{k,k'} = \sum_{s=0}^{S} (2\pi i k)^s \widetilde{a}_{s,k-k'} (2\pi i k')^s \tag{2.10}$$

is the reciprocal space matrix representation of the operator $\mathcal{L}$ (and $\mathcal{D}$). The Rayleigh quotient of Eq. 2.5 is Fourier transformed to,

$$\lambda_f = \sum_{k,k'=-\infty}^{\infty} \widetilde{\phi}_{f,k} \widetilde{\mathcal{L}}_{k,k'} \widetilde{\phi}_{f,k'}, \tag{2.11}$$

where we now have the Euclidean normalization,

$$\left\| \widetilde{\phi}_f \right\|_2 = \left(\sum_{k=-\infty}^{\infty} \widetilde{\phi}^*_{f,k} \widetilde{\phi}_{f,k} \right)^{1/2} = 1, \tag{2.12}$$

consistent with $\|\phi_f\|_{L_2} = 1$ and our Fourier transform definition.

In a classical digital computer, discretization of the domain $x \in [0,1]$ is required so that values x can be represented with a finite number of bits. For the quantum algorithm we'll be discussing, discretization of the domain will also be required so that values x can be identified with a finite number of qubits. We can then sample the spatial domain at the points $x = 0, 1/N, 2/N, \ldots, (N-1)/N$, where $N = 2^n$ requires n qubits. It will be more convenient to work with the integers $\overline{x} = Nx = 0, 1, 2, \ldots, N-1$. The discrete spatial domain of N points allows us to approximate a function $\psi(x)$ by a vector,

$$\psi^{(N)} = \left(\psi_0^{(N)}, \psi_1^{(N)}, \ldots, \psi_{N-1}^{(N)}\right) \tag{2.13}$$

for computational purposes, where we shall impose Euclidean norm $\|\psi^{(N)}\|_2 = 1$. In particular, we wish to generate discretized approximations $\phi_f^{(N)}$ that approach the continuous problem eigenvector $\phi_f(x)$ such that taking an ever greater number of discretization points N gives us the limit $\lim_{N\to\infty} \sqrt{N}\phi_{f,\overline{x}}^{(N)} = \phi_f(x)$, the factor $\sqrt{N}$ accounting for Euclidean normalization of the vector $\phi_f^{(N)}$ and L_2 normalization of the function $\phi_f(x)$. We discuss how we generate $\phi_f^{(N)}$ and how quantify the quality of our discrete approximations as a function of N further below.

In addition to having a discrete approximation to functions $\psi(x)$, we shall require discrete approximations of the differential operator $\mathcal{L}$ of Eq. 2.7 in the form of an $N \times N$ matrix $\mathcal{L}^{(N)}$ acting on vectors $\psi^{(N)}$. Hence, we'll need $N \times N$ finite difference matrices, which we shall denote $\Delta_{(s)}^{(N)}$, to approximate derivatives $\partial^s/\partial x^s$. There is freedom in choosing finite differences to approximate derivatives, here we (arbitrarily) choose the forward difference for a concrete example,

$$\left(\Delta_{(1)}^{(N)}\psi^{(N)}\right)_{\overline{x}} = N\left(\psi_{\overline{x}+1}^{(N)} - \psi_{\overline{x}}^{(N)}\right), \tag{2.14}$$

and higher order finite differences can be generated by $\Delta_{(s)}^{(N)} = \left(\Delta_{(1)}^{(N)}\right)^s$ for integer s. From the very definition of derivatives, we have $\lim_{N\to\infty} \sqrt{N}(\Delta_{(s)}^{(N)}\psi^{(N)})_{\overline{x}} = \partial^s\psi(x)/\partial x^s$ if $\lim_{N\to\infty} \sqrt{N}\psi_{\overline{x}}^{(N)} = \psi(x)$, the factor $\sqrt{N}$ again accounting for Euclidean normalization of the vector $\psi^{(N)}$ and L_2 normalization of the function $\psi(x)$. The subscript arithmetic $\overline{x} \pm 1$ in the definition of finite differences is to be performed modulo-N, consistent with the boundary conditions of Eqs. 2.2,2.3.

The resulting matrix operator $\mathcal{L}^{(N)}$ is,

$$\mathcal{L}^{(N)} = \sum_{s=0}^{S} \left(\Delta_{(s)}^{(N)}\right)^T \cdot \mathrm{Diag}(a_s^{(N)}) \cdot \Delta_{(s)}^{(N)}. \tag{2.15}$$

where $(\cdot)^T$ indicates matrix transpose and $\mathrm{Diag}(\cdot)$ indicates a diagonal matrix with the vector argument along the diagonal. With the above construction for $\mathcal{L}^{(N)}$, we can now pose a Hermitian matrix eigenvalue problem,

$$\sum_{\overline{x}'=0}^{N-1} \mathcal{L}_{\overline{x},\overline{x}'}^{(N)} \phi_{f,\overline{x}'}^{(N)} = \lambda_f^{(N)} \phi_{f,\overline{x}}^{(N)} \tag{2.16}$$

whose solutions will have the desired properties $\lim_{N\to\infty} \sqrt{N}\phi_{f,\overline{x}}^{(N)} = \phi_f(x)$ and $\lim_{N\to\infty} \lambda_f^{(N)} = \lambda_f$, with the obvious restriction $f \leq N$. A reciprocal space description is useful, for which we introduce the discrete Fourier transforms,

$$\begin{aligned} \widetilde{\psi}_k^{(N)} &= \frac{1}{\sqrt{N}} \sum_{\overline{x}=0}^{N-1} \omega^{-k\overline{x}} \psi_{\overline{x}}^{(N)}, \\ \psi_{\overline{x}}^{(N)} &= \frac{1}{\sqrt{N}} \sum_{k=-N/2}^{N/2-1} \omega^{k\overline{x}} \widetilde{\psi}_k^{(N)}, \end{aligned} \tag{2.17}$$

where $\omega = \exp(2\pi i/N)$ and reciprocal space has been truncated to the set $\mathcal{N} = \{k \in \mathbb{Z} : -N/2 \leq k \leq N/2+1\}$. The eigenvectors are assigned unit Euclidean norm in both $\overline{x}$ and k space representations,

$$\begin{aligned} \left\|\phi_f^{(N)}\right\|_2 &= \left(\sum_{\overline{x}=0}^{N-1} \phi_{f,\overline{x}}^{*(N)} \phi_{f,\overline{x}}^{(N)}\right)^{1/2} \\ &= \left(\sum_{k\in\mathcal{N}} \phi_{f,k}^{*(N)} \phi_{f,k}^{(N)}\right)^{1/2} = 1, \end{aligned} \tag{2.18}$$

so that the discrete analogs of Eqs. 2.5, 2.11 are

$$\begin{aligned} \lambda_f^{(N)} &= \sum_{\overline{x},\overline{x}'=0}^{N-1} \phi_{f,\overline{x}'}^{*(N)} \mathcal{L}_{\overline{x}',\overline{x}}^{(N)} \phi_{f,\overline{x}}^{(N)} \\ &= \sum_{k,k'\in\mathcal{N}} \widetilde{\phi}_{f,k'}^{*(N)} \widetilde{\mathcal{L}}_{k',k}^{(N)} \widetilde{\phi}_{f,k}^{(N)} \end{aligned} \tag{2.19}$$

which we shall find useful below.

We shall call $|\lambda_f^{(N)} - \lambda_f|$ the *truncation* error, alluding to the fact that we wish to approximate λ_f with $\lambda_f^{(N)}$ while truncating reciprocal space from all integers $\mathbb{Z}$ to the subset $\mathcal{N}$. We now proceed to show the well known fact that replacing derivatives by finite differences ultimately limits the convergence of $\lambda_f^{(N)}$ to λ_f as the number of sampling points N increases. Straightforward application of previously stated definitions gives,

$$\begin{aligned}\left(\widetilde{\Delta_{(1)}^{(N)}\psi^{(N)}}\right)_k &= N\left(\exp\left(2\pi ik/N\right)-1\right)\widetilde{\psi}_k^{(N)} \\ &= 2\pi ik\widetilde{\psi}_k^{(N)}\left(1+\Theta\left(\frac{k^2}{N^2}\right)\right) \qquad (2.20)\end{aligned}$$

where we have made use of series expansions and the fact that $|k| \leq N/2$ to arrive at the contribution $\Theta(k^2/N^2)$. The result holds for higher order derivatives.

An important parameter in characterizing truncation error is a reciprocal space cut-off $k(\phi_f)$, which can be defined for every ϕ_f. There will always exist a number $k(\phi_f)$ such that $|\widetilde{\phi}_{f,k}|^2 = O(k^{-(4S+1+\epsilon)})$ for all $|k| > k(\phi_f)$ and some infinitesimal ϵ. This follows simply because ϕ_f must be differentiable up to order $2S$, and therefore the series $\sum_k (2\pi k)^{4S}|\widetilde{\phi}_{f,k}|^2$ giving the norm of the $2S^{\text{th}}$ derivative of ϕ_f must converge. The eigenvalue spectrum of the continuous domain operator $\mathcal{L}$ is unbounded, and it can be shown that $\sup\{k(\phi_f)\}$ does not exist. However, since we restrict ourselves to $f = O(1)$, we can specify a finite $k(\phi_f)$ independent of N. For $N/2 > k(\phi_f)$, a reciprocal space cut-off $k(\phi_f^{(N)})$ must also exist since $\lim_{N\to\infty} \widetilde{\phi}_{f,k}^N = \widetilde{\phi}_{f,k}$. From here on, we shall not distinguish between $k(\phi_f^{(N)})$ and $k(\phi_f)$ as the precise value of the reciprocal space cut-off is not needed, but simply its existence. We thus define another subset of reciprocal space $\mathcal{M} = \{k \in \mathbb{Z} : |k| < k(\phi_f)\}$.

We have now collected enough ingredients to find the truncation error $|\lambda_f^{(N)} -$

$\lambda_f|$. We assume that $N/2 > k(\phi_f)$, so that a "reasonable" representation of ϕ_f can be made on the discretized domain. By "reasonable", we mean the eigenvalue λ_f can be estimated using Eq. 2.11 and the truncated reciprocal space $\mathcal{N}$ to give,

$$\lambda_f = \sum_{k,k' \in \mathcal{N}} \widetilde{\phi}^*_{f,k'} \widetilde{\mathcal{L}}_{k',k} \widetilde{\phi}_{f,k} + O\left(N^{-(2S+\epsilon)}\right) \tag{2.21}$$

where the above result arises from the least convergent (highest order derivative) contribution to λ_f in the region $k, k' \notin \mathcal{N}$,

$$\begin{aligned} &\sum_{k,k' \notin \mathcal{N}} \widetilde{\phi}^*_{f,k'} (2\pi i k)^S \widetilde{a}_{S,k-k'} (2\pi i k')^S \widetilde{\phi}_{f,k} \\ &= \sum_{k,k' \notin \mathcal{N}} O\left(k'^{-(S+1/2+\epsilon)} k^{-(S+1/2+\epsilon)}\right) \\ &\qquad = O\left(N^{-(2S+\epsilon)}\right) \end{aligned} \tag{2.22}$$

where we have made use of $|\widetilde{\phi}_{f,k}| = O(k^{-(2S+1/2+\epsilon)})$ for $k > k(\phi_f)$. Thus, for $N/2 > k(\phi_f)$, truncation of the reciprocal space sum in Eq. 2.11 gives $O(N^{-(2S+\epsilon)})$ error.

Using the finite difference error of Eq. 2.20, the reciprocal space matrix $\widetilde{\mathcal{L}}^{(N)}$ can be written,

$$\widetilde{\mathcal{L}}^{(N)}_{k,k'} = \sum_{s=0}^{S} \left[(2\pi i k)^s \widetilde{a}_{s,k-k'} (2\pi i k')^s \left(1 + \Theta\left(\frac{k^2 + k'^2}{N^2}\right)\right)\right], \tag{2.23}$$

where we have used the fact that there is some freedom in approximating $a_s(x)$ by $a^{(N)}_{s,\overline{x}}$. We choose to match spectral components $\widetilde{a}^{(N)}_{s,k} = \widetilde{a}_{s,k}$, and accept that $a^{(N)}_{s,\overline{x}}$ may exhibit oscillation artifacts (Gibb's phenomenon) due to discarding the contributions $\widetilde{a}_{s,k}$ for $k \in \mathbb{Z} - \mathcal{N}$. Note that the smoothness of $a_s(x)$, meaning continuity and finite s order derivatives for $x \in [0, 1]$, implies the existence of reciprocal space cut-offs $k(a_s)$. We use Eq. 2.19 to decompose,

$$\lambda^{(N)}_f = \sum_{k,k' \in \mathcal{M}} \widetilde{\phi}^{*(N)}_{f,k'} \widetilde{\mathcal{L}}_{k',k} \widetilde{\phi}^{(N)}_{f,k} \left(1 + \Theta\left(\frac{k(\phi_f)^2}{N^2}\right)\right)$$

$$+\sum_{k,k'\in(\mathcal{N}-\mathcal{M})} \widetilde{\phi}^{*(N)}_{f,k'}\widetilde{\mathcal{L}}_{k',k}\widetilde{\phi}^{(N)}_{f,k}\left(1+\Theta\left(\frac{k^2+k'^2}{N^2}\right)\right) \tag{2.24}$$

where the error summed over $\mathcal{N}-\mathcal{M}$ is of order,

$$\begin{aligned}\sum_{k,k'\in(\mathcal{N}-\mathcal{M})} \Theta\left(\frac{1}{N^2}\right) O\left(\frac{k^2+k'^2}{k'^{S+1/2+\epsilon}k^{S+1/2+\epsilon}}\right)& \\ = \Theta\left(\frac{1}{N^2}\right) O\left(\frac{1}{k(\phi_f)^{2S-3}}\right)&\end{aligned} \tag{2.25}$$

The diminishing contribution of the region $\mathcal{N}-\mathcal{M}$ to $\lambda_f^{(N)}$ ensures that the relative finite difference error $\Theta(k^2/N^2)$ does not approach unity but remains $\Theta(1/N^2)$. Collecting the results of Eqs. 2.21, 2.24, 2.25, we can express $\lambda_f^{(N)}-\lambda_f$ as,

$$\lambda_f^{(N)}-\lambda_f = \sum_{k,k'\in\mathcal{N}} \widetilde{\phi}^{*(N)}_{f,k'}\widetilde{\mathcal{L}}_{k',k}\widetilde{\phi}^{(N)}_{f,k}\left(1+\Theta\left(\frac{1}{N^2}\right)\right) - \sum_{k,k'\in\mathcal{N}} \widetilde{\phi}^{*}_{f,k'}\widetilde{\mathcal{L}}_{k',k}\widetilde{\phi}_{f,k} \tag{2.26}$$

where we have dropped the dependence upon $k(\phi_f)$ as it shall be of no further use. We note that $\widetilde{\mathcal{L}}^{(N)}-\widetilde{\mathcal{L}}=\Theta(1/N^2)$ in the reciprocal space $\mathcal{M}$, so we can consider $1/N^2$ a parameter of expansion in perturbation theory. The lowest order perturbation gives $\|\delta\widetilde{\phi}_f^{(N)}\|_2 = \|\widetilde{\phi}_f^{(N)}-\widetilde{\phi}_f\|_2 = O(1/N^2)$ for a non-degenerate ϕ_f. Degenerate eigenvectors might be perturbed substantially, but this is merely the result of there being no preferred basis for the span of the degenerate eigenvectors. The same bounds on truncation error can be shown to apply to the degenerate case. Noting that Eq. 2.26 is second order in eigenvector and $\delta\widetilde{\phi}_f$ is orthogonal to $\widetilde{\phi}_f$, the contribution of $\delta\widetilde{\phi}_f^{(N)}$ to the eigenvalue error is $O(1/N^4)$ and can therefore be ignored. The relative truncation error is,

$$\left|\frac{\lambda_f^{(N)}-\lambda_f}{\lambda_f}\right| = \Theta\left(\frac{1}{N^2}\right) \tag{2.27}$$

which is the final result of this section. We emphasize that truncation error arises solely from the uniform discretization of the domain $x\in[0,1]$.

2.3 Quantum Algorithm - One Dimension

We present now the quantum algorithm as it applies to the Hermitian, one-dimensional boundary value problem discussed in the previous section. We will show the various computational steps, and the rationale behind them.

First, we set forth some preliminaries. We will represent a vector $\psi^{(N)}$ with a quantum state composed of $n = \log_2 N$ qubits whose probability amplitudes are encoded as follows,

$$|\psi^{(N)}\rangle = \frac{1}{\sqrt{N}} \sum_{\overline{x}=0}^{N-1} \psi_{\overline{x}}^{(N)} |\overline{x}\rangle, \tag{2.28}$$

where $|\overline{x}\rangle$ is an n qubit state storing the binary representation of $\overline{x}$. Similarly, the finite difference matrix $\mathcal{L}^{(N)}$ is mapped to an operator,

$$\Lambda^{(N)} = \sum_{\overline{x},\overline{x}'=0}^{N-1} |\overline{x}\rangle \mathcal{L}_{\overline{x},\overline{x}'}^{(N)} \langle \overline{x}'|. \tag{2.29}$$

and we define the unitary exponential,

$$U = \exp(i\Lambda^{(N)}\tau) = \sum_{q=0}^{\infty} \frac{(i\Lambda^{(N)}\tau)^q}{q!}. \tag{2.30}$$

where τ is a dimensionless constant whose value is chosen in advance of the simulation and where the unitarity of U follows from the Hermitian nature of $\Lambda^{(N)}$. The constant τ must be carefully chosen to arrive at a desired accuracy in eigenvalue $\lambda_f^{(N)}$ without an unnecessarily large number of operations. The prescription for choosing τ is described further below in section 2.4. Note that τ is now an abstract scaling parameter rather than the time step of a quantum simulation.

We shall call the register of $n = \log_2 N$ qubits the *accumulator* register. In addition, a register of $m = \log_2 M$ qubits will be required to count phase steps, which we shall call the *index* register. Several ancilla qubits will be required, their number depending on the desired precision for the coefficients $a_s^{(N)}$ that specify Λ. The first steps are to load an initial state $\psi^{(N)}\rangle$ into the accumulator and to form an equal superposition of all index qubit states, giving a complete state,

$$|\Psi\rangle = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} |\psi^{(N)}\rangle|j\rangle, \tag{2.31}$$

The state $|\psi^{(N)}\rangle$ is an initial estimate of the field eigenvector of interest. To determine the required computational work to arrive at a suitable initial estimate $|\psi^{(N)}\rangle$, it is useful to decompose the accumulator state in terms of the initially unknown eigenstates $|\phi_f^{(N)}\rangle$,

$$|\Psi\rangle = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} \sum_{x=0}^{N-1} \alpha_f |\phi_f^{(N)}\rangle|j\rangle = \sum_{f=0}^{N-1} \alpha_f |\Psi_f\rangle, \tag{2.32}$$

where $\alpha_f = \langle\phi_f^{(N)}|\psi^{(N)}\rangle$. As will be shown, the probability the Abrams-Lloyd algorithm will give an estimate of eigenvalue λ_f in a single iteration is $|\alpha_f|^2$. To obtain an estimate of λ_f with probability approaching unity, approximately $1/|\alpha_f|^2$ iterations will be required. It is thus necessary for $|\psi^{(N)}\rangle$ to have a large overlap with $|\phi_f^{(N)}\rangle$ in order to avoid numerous iterations of the algorithm. The best technique proposed thus far is that of Jaksch and Papageorgiou [49], where a more coarsely defined $\phi_f^{(N_0)}$ is determined first (ie. $N_0 < N$). According to the analysis of the previous section, a coarse approximation limited by truncation error will allow one to achieve,

$$|\alpha_f|^2 \leq \left\| \widetilde{\phi}_f - \widetilde{\phi}_f^{(N_0)} \right\|_2^2 = \left\| \delta\widetilde{\phi}_f^{(N_0)} \right\|_2^2 = 1 - O(1/N_0^2) \tag{2.33}$$

for $f = O(1)$. Thus, one might solve for a desired $\phi_f^{(N_0)}$ classically (with cost that we will discuss later), and load the state $|\phi^{(N_0)}\rangle$ into the accumulator with $\Theta(N_0)$ operations.

We shall now follow the linear portion of the algorithm as it operates on a particular component $|\Psi_f\rangle$, reintroducing the full superposition over all f in the final (nonlinear) measurement step. The next stage of the algorithm is to apply the unitary U to the accumulator conditional upon the index to produce the superposition,

$$\begin{aligned} |\Psi'_f\rangle &= \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} U^j |\phi_f^{(N)}\rangle |j\rangle \\ &= \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} \exp(ij\lambda_f^{(N)}\tau) |\phi_f^{(N)}\rangle |j\rangle. \end{aligned} \tag{2.34}$$

Only M conditional applications of U are in fact required to form $|\Psi'_f\rangle$ from $|\Psi_f\rangle$. One applies U conditional on $j > 1$, then one applies U conditional on $j > 2$ and so forth until the $(M-1)^{\text{th}}$ conditional U is applied for $j = M - 1$. The conditional applications of U can be performed with a single additional ancilla qubit as follows. With at most $\log M$ logical operations, one can entangle the index register with an ancilla to form the state $\sum_j |j\rangle|C_{j,j'}\rangle$ where the ancilla $C_{j,j'} = 1$ for $j \geq j'$ and $C_{j,j'} = 0$ otherwise. The j'^{th} application of U can be implemented as a U conditional on the ancilla $C_{j,j'}$. The ancilla is then disentangled from the quantum register by running the initial entangling operation once again.

The operator U acts in the full n qubit Hilbert space of $|\psi^{(N)}\rangle$, which will in general be prohibitively large, but it is nevertheless possible to efficiently generate an approximation to U using operations in a few qubit Hilbert space. The structure of $\Lambda^{(N)}$ is a band diagonal matrix resulting from local operations, and thus it has a block diagonal representation *in the qubit basis* of the accumulator.

To illustrate explicitly some of the key features of the algorithm at work, it's useful to consider the simple example where $\mathcal{D} = \partial/\partial x\{a(x)(\partial/\ \partial x)\}$. The following decomposition is appropriate,

$$\Lambda^{(N)} = \; N^2 \begin{pmatrix} d_0 & & & & \\ & d_1 & & & \\ & & d_2 & & \\ & & & d_3 & \\ & & & & \ddots \\ & & & & \end{pmatrix} - N^2 \begin{pmatrix} 0 & a_1 & & & & \\ a_1 & 0 & & & & \\ & & 0 & a_3 & & \\ & & a_3 & 0 & & \\ & & & & \ddots & \\ & & & & & \end{pmatrix}$$

$$-N^2 \begin{pmatrix} 0 & & & & & & a_0 \\ & 0 & a_2 & & & & \\ & a_2 & 0 & & & & \\ & & & 0 & a_4 & & \\ & & & a_4 & 0 & & \\ a_0 & & & & & \ddots & \end{pmatrix}$$

where $d_{\overline{x}} = a_{\overline{x}} + a_{\overline{x}+1}$. The operators $\Lambda^{(N,p)}$ can be written more compactly,

$$\begin{aligned} \Lambda^{(N,1)} &= N^2 \sum_{\overline{x}} d_{\overline{x}} |\overline{x}\rangle\langle\overline{x}| \\ \Lambda^{(N,2)} &= -N^2 \sum_{\overline{x} \text{ even}} a_{\overline{x}+1} \left\{ |\overline{x}\rangle\langle\overline{x}+1| + |\overline{x}+1\rangle\langle\overline{x}| \right\} \\ \Lambda^{(N,3)} &= -N^2 \sum_{\overline{x} \text{ odd}} a_{\overline{x}+1} \left\{ |\overline{x}\rangle\langle\overline{x}+1| + |\overline{x}+1\rangle\langle\overline{x}| \right\}, \end{aligned}$$

where $\Lambda^{(N,1)}$ is diagonal, and $\Lambda^{(N,2)}$,$\Lambda^{(N,3)}$ act in one qubit subspaces (conditional upon $\lfloor \overline{x}/2 \rfloor$) in lieu of the full Hilbert space of $\Lambda^{(N)}$.

The unitary U can be approximated to take advantage of the above decomposition in several ways. For a general decomposition,

$$\Lambda^{(N)} = \sum_{p=1}^{R} \Lambda^{(N,p)}, \tag{2.35}$$

where for our simple example $R = 3$, the Baker-Campbell-Hausdorff formulae can be used to show,

$$\begin{aligned} U_\Pi &= \prod_{p=1}^{R} \exp\left(i\Lambda^{(N,p)}\tau/2\right) \prod_{p=R}^{1} \exp\left(i\Lambda^{(N,p)}\tau/2\right) \\ &= \exp\left(i\Lambda^{(N)}\tau - \frac{i}{3!}\sum_{p,q=1}^{R} \left[\Lambda^{(N,p)}, \left[\Lambda^{(N,q)}, \Lambda^{(N,R)}\right]\right]\tau^3 + O\left(\left\|\Lambda^{(N)}\right\|_2^4 \tau^4\right)\right) \end{aligned} \tag{2.36}$$

where terms bilinear in $\Lambda^{(N,p)}\tau$ are suppressed by the symmetry of the product formula shown. One may approximate U by U_Π to take advantage of the efficient implementation of $\exp(i\Lambda^{(N,p)}\tau)$ at the cost of introducing error.

The quantum circuit for implementing U_Π for our simple example $\mathcal{D} = \partial/\partial x\{a(x)(\partial/\ \partial x)\}$ is shown in Fig. 2.1 for the particular case of an accumulator with $N = 2^4$. The ancillae initialized to state $|0\rangle$ are used to store the coefficients $d_{\overline{x}}$ or $a_{\overline{x}+1}$, to three bit precision with resolution δ: $d_{\overline{x}} = (\mathbf{d}_{\overline{x},2}2^2 + \mathbf{d}_{\overline{x},1}2^1 + \mathbf{d}_{\overline{x},0}2^0) \times \delta$ with $\mathbf{d}_{\overline{x},j} \in \{0,1\}$, and a similar binary description for $a_{\overline{x}+1}$. The circuit "$d_{\overline{x}}$" maps $|\overline{x}\rangle|\text{anc}\rangle \to |\overline{x}\rangle|\mathbf{d}_{\overline{x}} \oplus \text{anc}\rangle$, while "$a_{\overline{x}+1}$(even/odd)" maps $|\overline{x}\rangle|\text{anc}\rangle \to |\overline{x}\rangle|\mathbf{a}_{\overline{x}+1} \oplus \text{anc}\rangle$ for even/odd $\overline{x}$. Single qubit rotations $R'_p = |0\rangle\langle 0| + \exp(i(N^2\delta\tau/2)2^p)|1\rangle\langle 1|$ give the desired diagonal phase shifts for $\exp(i\Lambda^{(N,1)}\tau/2)$. Single qubit rotations $X'_p = \exp(i(-N^2\delta\tau/2)2^p\sigma_x)$, where $\sigma_x = |0\rangle\langle 1| + |1\rangle\langle 0|$, to implement the desired off diagonal couplings of $\exp(i\Lambda^{(N,2)}\tau/2)$ and $\exp(i\Lambda^{(N,3)}\tau/2)$. The values of τ and δ can be inferred from the restriction that the operator *splitting* error is comparable to truncation error, described in section 2.4. The parity shift operators, defined $D^{\pm}|\overline{x}\rangle = |\overline{x} \pm 1\rangle$, are required to shift the block structure of $\Lambda^{(N,3)}$ so that only operations on the least significant qubit $\overline{x}_0$ need be performed. The $D^{\pm}$ can be implemented using quantum Fourier transforms (at cost of $O(\log^2 N)$ operations) and single qubit rotations. Final disentanglement of ancillae is achieved by a second application of "$d_{\overline{x}}$" or "$a_{\overline{x}+1}$".

The reason for the ease of implementing U_Π is apparent in Fig. 2.1, one

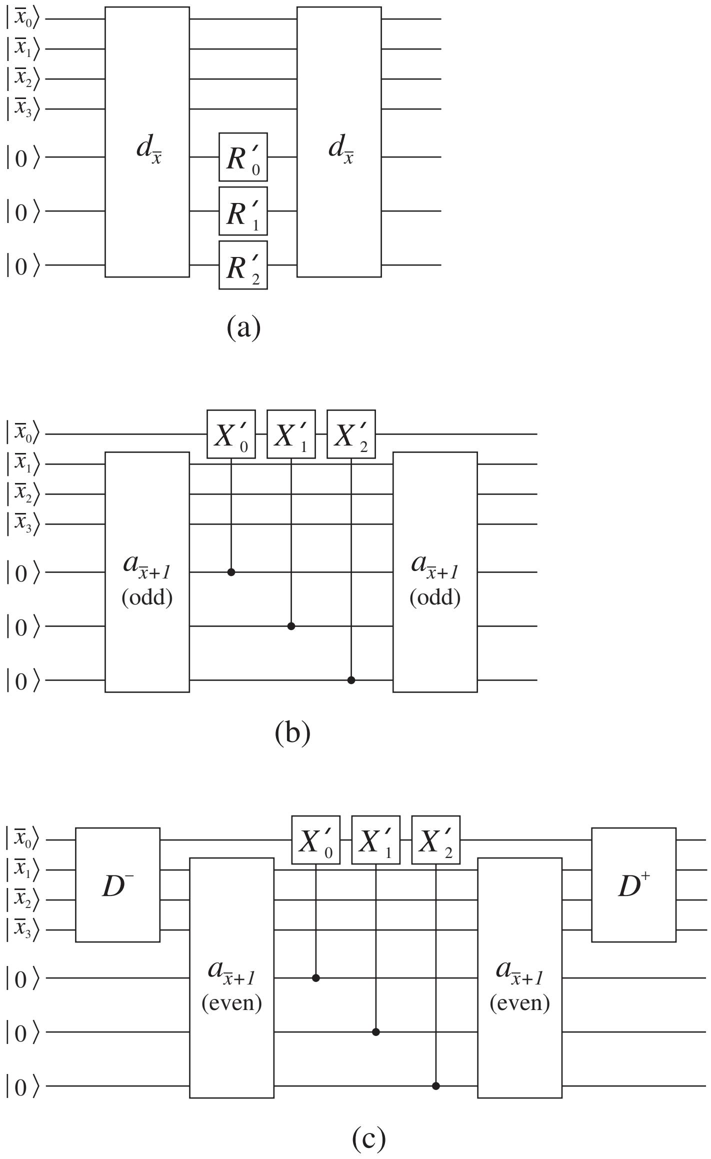

Figure 2.1: The quantum circuits for applying: (a) $\exp(i\Lambda^{(N,1)}\tau/2)$, (b) $\exp(i\Lambda^{(N,2)}\tau)$, and (c) $\exp(i\Lambda^{(N,3)}\tau/2)$, to the accumulator qubits $|\overline{x}\rangle = |\overline{x}_3\overline{x}_2\overline{x}_1\overline{x}_0\rangle$ for the decomposition of Eq. (2.35) with $N = 2^4$.

applies single qubit unitaries conditional upon the evaluation of $a_{\overline{x}}$. The Solovay-Kitaev theorem guarantees that the single qubit unitaries can be implemented to an accuracy $\Theta(1/N^2)$ with $\Theta(\log^c N)$ universal quantum gates [50]. We also assume that evaluation of $a_{\overline{x}}$ requires $O(\log N)$ operations. Roughly speaking, the differentiability of $a(x)$ rules out pathological functions that have greater complexity.

Approximating U by U_Π implies that the algorithm will give an estimate of the eigenvalue $\lambda_{f,\Pi}^{(N)}$ of the operator,

$$\Lambda_\Pi^{(N)} = \Lambda^{(N)} + O\left(\left\|\Lambda^{(N)}\right\|_2^3 \tau^2\right) \tag{2.37}$$

instead of the desired eigenvalue $\lambda_f^{(N)}$. We call the error introduced by using U_Π the *splitting* error, which has value $O\left(\left\|\Lambda^{(N)}\right\|_2^3 \tau^2\right)$ provided $\|\Lambda^{(N)}\|_2 \tau < 1$. The splitting error will be shown in the next section to limit the computational efficiency of estimating eigenvalues.

Applying U_Π rather than U, Eq. (2.34) takes the form

$$|\Psi'_{f,\Pi}\rangle = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} \exp(ij\lambda_{f,\Pi}^{(N)}\tau)|\phi_{f,\Pi}^{(N)}\rangle|j\rangle. \tag{2.38}$$

The eigenvalue is encoded in the phase periodicity of $|\phi_{f,\Pi}\rangle|j\rangle$, and can be determined to at most the $\pm\pi/M$ precision allowed by a $\log_2 M$ bit representation of a full 2π radians. We briefly review the procedure for retrieving the phase [14], beginning with the application of the quantum Fourier transform,

$$\mathrm{QFT} = \frac{1}{\sqrt{M}} \sum_{l=0}^{M-1} \sum_{m=0}^{M-1} \exp\left(-2\pi i l m/M\right) |l\rangle\langle m|, \tag{2.39}$$

to the index qubits. The resulting state is,

$$\mathrm{QFT}|\Psi'_{f,\Pi}\rangle = \sum_{l=0}^{M-1} b_{l,f}|\phi_{f,\Pi}^{(N)}\rangle|l\rangle. \tag{2.40}$$

with the coefficients,

$$b_{l,f} = \frac{1}{M} \sum_{j=0}^{M-1} \exp\left(ij\left(\lambda_{f,\Pi}^{(N)}\tau - 2\pi l/M\right)\right), \tag{2.41}$$

which have square modulus,

$$|b_{l,f}|^2 = \frac{\sin^2\left[M\left(2\pi l/M - \lambda_{f,\Pi}^{(N)}\tau\right)/2\right]}{M^2 \sin^2\left[\left(2\pi l/M - \lambda_{f,\Pi}^{(N)}\tau\right)/2\right]}. \tag{2.42}$$

A projective measurement of the index produces $|l'\rangle$ where $|\lambda_{f,\Pi}^{(N)}\tau/2\pi - l'/M| < 1/2M$ with a probability $|b_{l',f}|^2 \geq (M^2\sin^2(\pi/2M))^{-1}$. All eigenvalues will satisfy $|\lambda_{f,\Pi}^{(N)}| < \pi/\tau$ since we will impose $\|\Lambda^{(N)}\|_2\tau \ll 1$ (to be made precise in the next section), so identification of l' will determine an eigenvalue $\lambda_{f,\Pi}^{(N)}$ uniquely to a precision $\pm\pi/M\tau$.

Since we began not with the desired state alone, but with a superposition $|\Psi\rangle = \sum_{k=0}^{N-1}\alpha_f|\Psi_f\rangle$, measurement of the index will determine a particular $\lambda_{f,\Pi}^{(N)}$ with relative probability $|\alpha_f|^2$. It is the initial trial wavefunction $|\psi^{(N_0)}\rangle$ that determines the probability $|\alpha_f|^2$ of the eigenvalue/eigenvector pair being selected by a projective measurement.

Upon completion of the eigenvalue readout (via index bits l'), the accumulator is left in the eigenstate $|\phi_f^{(N)}\rangle$. This is useful since it allows further information to be extracted. For instance, one can efficiently test whether $|\phi_f^{(N)}\rangle$ has a particular symmetry, such as inversion symmetry about a particular point $\overline{x}$ in the domain. This can serve as a partial check as to whether the desired $|\phi_f^{(N)}\rangle$ was indeed selected by the projective measurement.

2.4 Computational Cost - One Dimension

We now analyze the computational cost for implementing the Abrams-Lloyd algorithm for the one dimensional Hermitian problem described in the preceding sections. As pointed out, there are three sources of error that must be considered to determine the required number of operations for a given accuracy in eigenvalue estimation.

First, uniform discretization of the continuous problem to $N = 2^n$ points on the spatial domain introduces a truncation error,

$$\left| \frac{\lambda_f^{(N)} - \lambda_f}{\lambda_f} \right| = \Theta\left(\frac{1}{N^2}\right) \tag{2.43}$$

The truncation error quantifies the accuracy with which the discrete problem represents the continuous problem for low order (ie. $f = O(1)$) eigenfunctions ϕ_f. To compare algorithms, classical or quantum, we may ask how many operations are required to achieve the $\Theta(1/N^2)$ accuracy in the solution of the discrete eigenvalue problem.

Second, splitting $\Lambda^{(N)}$ into parts so as to approximate U with a product U_Π of local operators results in what we have termed splitting error. From Eq. 2.37 the eigenvalue $\lambda_{f,\Pi}^{(N)}$ of U_Π is,

$$\lambda_{f,\Pi}^{(N)} \quad = \quad \lambda_f^{(N)} + O\left(\left\|\Lambda^{(N)}\right\|_2^3 \tau^2\right) \tag{2.44}$$

where we choose τ such that $\|\Lambda^{(N)}\|_2 \tau = \|\mathcal{L}^{(N)}\|_2 \tau < 1$. However, from the the finite difference formula Eq. 2.14 and the form of $\mathcal{L}^{(N)}$ in Eq. 2.15, the spectral radius $\|\mathcal{L}^{(N)}\|_2 = \Theta(N^{2S})$. Hence, the splitting error is,

$$\lambda_{f,\Pi}^{(N)} \quad = \quad \lambda_f^{(N)} + O\left(N^{6S}\tau^2\right) \tag{2.45}$$

which, unlike truncation error, *increases* polynomially with an increase in the number of discretization points N. The splitting error results from the fact that the product U_Π creates deviations from the true advancement in phase at high spatial frequencies. For example, in the system described in Eq. 2.35, it is the non-commuting nature of advancing even pairings of points and odd pairings of points that generates an error with spatial frequency $N/2$.

Third, the measurement of phase $\lambda_{f,\Pi}^{(N)}\tau$ via the quantum Fourier transform is limited by the uniform discretization of 2π radians into $M = 2^m$ intervals. The limited phase resolution allows us to specify $\lambda_{f,\Pi}^{(N)}$ upon completion of the algorithm to a precision $2\pi/M\tau$.

The three sources of error allow us to determine the optimal number of index bits $m = \log_2 M$, the value of the constant τ, and thus the complexity of the algorithm. Obviously, there is nothing gained in solving the discretized problem to an accuracy greater than the truncation error $\Theta(1/N^2)$ if the goal is to study the continuous problem. We can thus allow the splitting error $O(N^{6S}\tau^2)$ to be of the same order as the truncation error,

$$\Theta\left(\frac{1}{N^2}\right) \geq O\left(N^{6S}\tau^2\right) \quad \rightarrow \quad \tau \leq \Omega\left(\frac{1}{N^{3S+1}}\right) \tag{2.46}$$

Since $\lambda_{f,\Pi}^{(N)} = O(1)$ for our low order eigenvalue with $f = O(1)$, the phase advancement $\lambda_{f,\Pi}^{(N)}\tau \leq \Omega(1/N^{3S+1})$ for the low order eigenfunction becomes exceedingly small. In order to resolve this phase so that our final eigenvalue uncertainty does not exceed the truncation error, we require

$$\frac{2\pi}{M\tau} \leq \Theta\left(\frac{1}{N^2}\right) \quad \rightarrow \quad M \geq O\left(N^{3(S+1)}\right) \tag{2.47}$$

thus prescribing the number $m = \log_2 M$ of index register qubits.

The complexity of the eigenvalue estimation can now be stated. The determination of a suitable initial guess eigenstate $\phi_f^{(N_0)}$ requires the determination of an eigenvector of an $N_0 \times N_0$ problem. This can be done classically in $\Omega(N_0)$ steps, since each of N_0 points in the spatial domain description of $\mathcal{L}^{(N_0)}$ must contribute to the eigenvalue. Near optimal classical methods are in fact known. In the case of a tridiagonal $\mathcal{L}^{(N_0)}$, bisection gives an eigenvalue to $\Theta(1/N_0^2)$ precision with $\Theta(N_0 \log N_0)$ operations [45]. Low order eigenvalues of wider bandwidth $\mathcal{L}^{(N_0)}$ matrices can be determined to the same precision with the same order of operations using more complex classical techniques 1. Only a modest N_0 is required for the probability of a successful iteration of the quantum algorithm, $1 - O(1/N_0^2)$, to be comparable to unity. Following the construction of an initial eigenstate estimate, this estimate must be loaded into the accumulator register, which can be done in $\Theta(N_0)$ steps. We suppose that N will exceed N_0 by a substantial factor, so that the initial state preparation is a negligible cost compared to the remainder of the algorithm. The majority of the computational steps in the quantum algorithm are accounted for by the $M \geq O\left(N^{3(S+1)}\right)$ applications of U_Π, each of which requires $O(\log^c N)$ gate operations for some constant $c = O(1)$. The final quantum Fourier transform requires $\Theta(\log^2 M)$ gate operations, a negligible $\log(N)$ contribution compared to the M applications of U_Π.

Thus, to achieve $\Theta(1/N^2)$ accuracy in the final eigenvalue, at least $O(N^{3(S+1)} \log^c N)$ operations and $\Theta(D \log N)$ qubits are required. In contrast, an eigenvalue can be found using classical techniques to $\Theta(1/N^2)$ accuracy using $\Theta(N \log N)$ operations. The quantum algorithm requires significantly more work than classical algorithms for the one dimensional problem. Nonetheless, we show in the next section that the quantum algorithm is easily extended to higher dimensional problems where increased efficiency over classical techniques is indeed possible.

2.5 Higher Dimensional Problems

Here we will generalize the results of the one dimensional problem to the multidimensional problem. Many of the arguments presented in the earlier sections are not specific to the single dimension domain, and in many cases we can simply replace scalars with vectors. The continuous problem we wish to solve involves an operator $\mathcal{D}$ mapping functions $\psi(\mathbf{x})$, defined over a D-dimensional cubic domain $\mathbf{x} \in \mathcal{S} = [0,1]^{\otimes D}$, to functions $\mathcal{D}\psi(\mathbf{x})$. Rather than explicitly writing out the general form of a multidimensional Hermitian operator $\mathcal{D}$ analagous to the single dimensional operator of Eq. 2.1, we simply state that $\mathcal{D}$ must satisfy,

$$\int_0^1 dx_1 \cdots \int_0^1 dx_D \left(\phi_f^* \mathcal{D}\phi_{f'} - \phi_{f'}^* \mathcal{D}\phi_f\right) = 0 \tag{2.48}$$

for any eigenfunctions ϕ_f satisfying $\mathcal{D}\phi_f = \lambda_f \phi_f$. We can then define an "equivalent" bilinear operator $\mathcal{L}$ that maps any two vector functions $\psi^*(\mathbf{x})$ and $\varphi(\mathbf{x})$ to a scalar function $\psi^* \mathcal{L}\varphi$. This can be done by using the higher dimensional forms of integration by parts, which in one dimension allowed us to relate $\mathcal{D}$ to $\mathcal{L}$. We exclude "trivial" problems that are readily expressed as a tensor product of single dimensional problems, $\mathcal{L} = \mathcal{L}_1 \otimes \mathcal{L}_2 \otimes \ldots \otimes \mathcal{L}_D$. We are therefore considering problems whose structure is instead a sum of tensor product terms,

$$\mathcal{L} = \sum_{\beta=1}^{B} \mathcal{L}_{\beta,1} \otimes \mathcal{L}_{\beta,2} \otimes \ldots \otimes \mathcal{L}_{\beta,D} \tag{2.49}$$

for some constant $B > 1$, and where the differential order of each one dimensional $\mathcal{L}_{\beta,\alpha}$ is $2S_{\beta,\alpha}$. The differential order of $\mathcal{L}$ is then $2S = \max_\beta\{\sum_\alpha 2S_{\beta,\alpha}\}$. Of course, we retain the Hermitian property

$$\int_0^1 dx_1 \cdots \int_0^1 dx_D \left(\phi_f^* \mathcal{L}\phi_{f'} - \phi_{f'}^* \mathcal{L}\phi_f\right) = 0 \tag{2.50}$$

and the associated eigenvalue/eigenvector properties. Normalizing the eigenfunctions allows us to write,

$$\lambda_f = \int_0^1 dx_1 \cdots \int_0^1 dx_D \phi_f^* \mathcal{L} \phi_f \tag{2.51}$$

which is simply the Rayleigh quotient.

Discretization proceeds as in the single dimensional case, with each domain coordinate $x_i \in [0, 1]$ discretized to N points. Functions $\psi(\mathbf{x})$ are represented by rank D tensors $\psi_{\overline{\mathbf{x}}}^{(N)}$, ie. for $D = 2$ dimensions $\psi_{\overline{\mathbf{x}}}^{(N)}$ is a matrix of numbers, for $D = 3$ dimensions $\psi_{\overline{\mathbf{x}}}^{(N)}$ is a "cube" of numbers and so forth. Partial derivatives are converted to finite differences as in the one dimensional case. The operator $\mathcal{L}$ is can thus be discretized to a tensor $\mathcal{L}_{\overline{\mathbf{x}},\overline{\mathbf{x}}'}^{(N)}$. The truncation error in the multidimensional problem is,

$$\left| \frac{\lambda_f^{(N)} - \lambda_f}{\lambda_f} \right| = \Theta\left(\frac{1}{N^2} \right) \tag{2.52}$$

which is identical to the one dimensional case because the relative finite difference errors on each coordinate are $\Theta(1/N^2)$.

The implementation of the Abrams-Lloyd algorithm for the multidimensional problem proceeds in a completely analagous fashion to the one dimensional case, with the number of accumulator qubits $D \log_2 N = Dn$ so as to represent a volume $V = N^D$. As before, an initial estimate of the desired eigenvector $\phi_f^{(N)}$ is required. A coarse classical simulation can produce an eigenvector $\phi_f^{(N_0)}$ with $N_0 < N$. Since truncation error scales as $\Theta(1/N^2)$, the required value of N_0 is such that the probability of a successful iteration of the algorithm, $1 - O(1/N_0^2)$, approaches unity. The computational cost is $\Theta(N_0^D \log N_0)$ classical gate operations for generating the initial eigenstate and $\Theta(N_0^D)$ gate operations to load the state into the accumulator.

The heart of the algorithm is the controlled application of the unitary $U = \exp(i\Lambda^{(N)}\tau)$ where $\Lambda^{(N)} = \sum |\overline{\mathbf{x}}\rangle \mathcal{L}^{(N)}_{\overline{\mathbf{x}},\overline{\mathbf{x}}'} \langle \overline{\mathbf{x}}'|$. As before, U acts within a large Hilbert space, so an approximating operator U_Π is applied instead. The operator U_Π is a sequence of operations acting conditionally upon a much smaller Hilbert space than the full $D \log_2 N$ qubits. We quantify the size of this Hilbert space now. The multidimensional $\Lambda^{(N)}$ is no longer represented by a band diagonal matrix, but has the structure of a sum of tensor products as in Eq. 2.49,

$$\Lambda^{(N)} = \sum_{\beta=1}^{B} \Lambda^{(N)}_{\beta,1} \otimes \Lambda^{(N)}_{\beta,2} \otimes \ldots \otimes \Lambda^{(N)}_{\beta,D} \tag{2.53}$$

The local nature of $\Lambda^{(N)}$ is quantified by the maximum number of states $|\mathbf{x}'\rangle$ for which $\langle \overline{\mathbf{x}}'|\Lambda^{(N)}|\overline{\mathbf{x}}\rangle$ is not zero (maximizing over all possible $|\overline{\mathbf{x}}\rangle$). This volume, v, is the maximum product of matrix bandwidths,

$$\begin{aligned} v &= \max_\beta \{(2S_{\beta,1}+1)(2S_{\beta,2}+1)\ldots(2S_{\beta,D}+1)\} \\ &\le \left(1 + \frac{2S}{D}\right)^D \end{aligned} \tag{2.54}$$

where we have used the restriction $2S = \max_\beta \sum_\alpha \{2S_{\beta,\alpha}\}$ to arrive at the bound on v. It follows that we can split $\Lambda^{(N)} = \sum_{p=1}^{R} \Lambda^{(N,p)}$ where the $\Lambda^{(N,p)}$ act conditionally upon a Hilbert space of $r = \lceil \log_2 v \rceil$ qubits. The size of this reduced Hilbert space is independent of domain size N^D, so that $\exp(i\Lambda^{(N,p)}\tau)$ can be applied to the requisite accuracy (polynomial in $1/N$) with only $\Theta(\log^c N)$ universal gates for some constant c. As before, we assume that the function evaluations required for conditional action upon the r-qubit subspace entails at most $O(\log N)$ universal gates. The total number of the split up operators $\Lambda^{(N,p)}$ is bounded $R \le v = (1 + 2S/D)^D$ independently of the domain size N^D. Thus, U_Π can be applied with $O(\log^c N)$ work for some $c = O(1)$.

The approximation U_Π can be composed by using a symmetric product as in

Eq. 2.36 so that the splitting error is,

$$\lambda_{f,\Pi}^{(N)} = \lambda_f^{(N)} + O\left(\left\|\Lambda^{(N)}\right\|_2^3 \tau^2\right) \tag{2.55}$$

More generally, an approximation of U correct to higher order in $\|\Lambda^{(N)}\|_2\tau$ can be implemented [51, 52], [53]. For the sake of generality, we assume we have a product operator U_Π correct to order $\|\Lambda^{(N)}\|_2^\nu\tau^\nu$, and set $\nu = 2$ to recover the simple symmetric product results. In practice, one can not take ν arbitrarily large since the number of terms in U_Π grows exponentially in ν. The optimal choice of ν is that which minimizes the overall computational cost.

Using the fact that $2S$ is the differential order of $\mathcal{L}$, the splitting error becomes,

$$\lambda_{f,\Pi}^{(N)} = \lambda_f^{(N)} + O\left(N^{2S(\nu+1)}\tau^\nu\right) \tag{2.56}$$

The final phase measurement through a quantum Fourier transform proceeds as in the one dimensional case, with the same precision of $\pm\pi/(M\tau)$ in determining $\lambda_{f,\Pi}^{(N)}$, where $m = \log_2 M$ is the number of index qubits. Requiring that the final eigenvalue be determined to the truncation error limit as in the one dimensional case, the same line of reasoning as in the previous section leads to,

$$\begin{aligned} \tau \leq \Omega\left(\frac{1}{N^{2(S(1+1/\nu)+1/\nu)}}\right) \\ M \geq O\left(N^{2(S+1)(1+1/\nu)}\right) \end{aligned} \tag{2.57}$$

The computational cost of the algorithm is dominated by the M applications of U_Π, each application of U_Π requiring $O(\log^c N)$ number of operations. The computational cost for the quantum algorithm is,

$$\aleph_Q = O(M \log^c N) = O(N^{2(S+1)(1+1/\nu)} \log^c N) \tag{2.58}$$

in addition to the cost for finding and loading an eigenstate with coarse discretization N_0 along each axis. We assume $O(N_0)$. The number of qubits required by the quantum algorithm is $\Theta(\log N)$.

We now consider classical costs associated with the multidimensional eigenvalue equation. Discretization and reduction of the continuous problem to a matrix equation results in a sparse $N^D \times N^D$ matrix with a number of bands depending on the spatial derivatives and dimensions in the continuous problem. The most efficient and near optimal classical method requires

$$\aleph_C = O(N^D \log N) \tag{2.59}$$

operations in order to attain a low order eigenvalue with truncation error accuracy $\Theta(1/N^2)$. The method is near optimal in the classical case since the computational cost per each of $V = N^D$ points in the domain is merely $O(\log N)$. Any classical method must "visit" each point in the simulation domain in order for that point to influence the outcome of the classical calculation, hence the classical computation cost is $\Omega(N^D)$. Of course, the number of bits required is $\Theta(N^D)$.

The maximum improvement in computational efficiency provided by the quantum algorithm presented is,

$$\max\left\{\frac{\aleph_C}{\aleph_Q}\right\} = O\left(\frac{N^{D-2(S+1)(1+1/\nu)}}{\log^{c-1} N}\right) \tag{2.60}$$

with respect to the best known (near optimal) classical algorithm. From the above, we see that the domain dimension must satisfy $D > 2(S+1)(1+1/\nu)$ in order to see any improvement using the Abrams-Lloyd algorithm. In particular, we have $S = 1$ for Schrödinger's equation and we can identify $D/3$ as the number of particles in space (3 degrees of freedom per particle, neglecting spin). A many-body eigenvalue calculation is more efficient than classical simulation for particle number $D/3 > (4/3)(1 + 1/\nu)$. For the case where U_Π is a simple symmetric product, $\nu = 2$ and we require $D/3 > 2$ in order to see improved computational efficiency. Higher order approximations, $\nu > 2$ will result in two (spinless) particle calculations already being done more efficiently using the Abrams-Lloyd algorithm.

We now discuss the generality of the results for domains other than the simple hypercube $\mathcal{S} = [0,1]^{\otimes D}$ discretized to $V = N^D$ points. A more complex domain $\mathcal{S}'$ can be had by deleting regions from $\mathcal{S}$ along planes defined by the uniform discretization scheme. The computational cost incurred is that required to ensure the probability amplitudes in $|\psi\rangle$ do not "spill" into the deleted regions $\mathcal{S} - \mathcal{S}'$. This is easily done by circuits such as those in Fig. 2.1, wherein quantum gates can be used to determine the conditional application of a few-qubit operator through out the simulation domain. The computational cost is therefore proportional to the classical cost of determining whether a point $\overline{x}$ is in or out of the specified domain $\mathcal{S}'$ subtended by the hypercube $\mathcal{S}$. As an explicit example, the subcircuit $a_{\overline{x}+1}$ of Fig. 2.1 for applying $\exp(i\Lambda^{(N,2)}\tau)$ of Eq. 2.35 can be made to compute $|\overline{x}\rangle|\mathrm{anc}\rangle \rightarrow |\overline{x}\rangle|\mathrm{anc}\rangle$ for $\overline{x} \in \mathcal{S} - \mathcal{S}'$ and $|\overline{x}\rangle|\mathrm{anc}\rangle \rightarrow |\overline{x}\rangle|a_{\overline{x}+1} \oplus \mathrm{anc}\rangle$ for $\overline{x} \in \mathcal{S}'$. The effect of this operation is to conditionally apply $\exp(i\Lambda^{(N,2)}\tau)$ to those points $\overline{x} \in \mathcal{S}'$. Clearly, $\mathcal{S}' = \mathcal{S}$ is the simplest domain to consider as there is no added computational cost, but more complex domains are accessible at only the modest cost of describing the domain with a Boolean function.

2.6 Conclusion

Our analysis of the Abrams-Lloyd algorithm raises several questions. Firstly, it is natural to ask what sort of qubit phase rotation accuracy is required during the application of U_Π to the initial guess eigenstate. The phase that is applied to qubits by the operator U_Π during the computation is of the same order as the phase applied to the highest order eigenvector: $\lambda^{(N)}_{N,\Pi}\tau$ where the eigenvalue $\lambda^{(N)}_{N,\Pi} = \Theta(N^{2S})$ for a differential operator of order $2S$ and $\tau = \Omega(1/N^{3S+1})$ for a second order splitting formula. The magnitude of the phase rotations applied to qubits is therefore $\Omega(1/N^{S+1})$. The relative accuracy with which the phase must be applied is $\Theta(1/N^2)$ if the final eigenvalue estimation is to be accurate to the truncation error limit of $\Theta(1/N^2)$. Thus the absolute accuracy required from single qubit rotations is $\Omega(1/N^{S+3})$, independent of the number of dimensions D. The absolute accuracy is a small quantity for very modest values of $N = 100$ (representing a relative eigenvalue accuracy of 10^{-4}) with a second order operator ($2S = 2$). Angular resolution of 10^{-8} in the control of qubits represents a technical feat, but thankfully the principles of fault tolerant quantum computation [20, 54] can be applied here to lessen the accuracy requirements for *physical qubit* operations.

Secondly, it is tempting to compare the quantum and classical algorithms for the simulation of dynamical evolution. The Abrams-Lloyd algorithm simulates the dynamics of the Schrödinger equation $\partial\psi/\partial t = \mathcal{D}\psi$ for some (possibly fictitious) Hamiltonian represented by $\mathcal{D}$, but only limited detail of the dynamics in a quantum simulation are accessible. The probability amplitudes characterizing a register of $D \log N + \log M$ qubits can result in at most $D \log N + \log M$ classical bits of information being extracted by measurement (by the Holevo bound). For instance, in order to obtain the eigenvector coefficients $\phi_f^{(N)}$, at least $\Theta(N^D/\log N)$ iterations of the algorithm would be required. This is in contrast to a classical simulation of dynamical evolution where $\Theta(N)$ bits would be required to store a state at a single dynamical step, and $\Theta(NM)$ bits are required to store the entire evolution of an initial state over M dynamical steps. We emphasize that the strength of the Abrams-Lloyd algorithm is not in its ability to provide great detail into dynamical evolution but rather in extracting useful classical information (such as eigenvalues) from a very compact representation of that dynamical evolution.

Finally, the analysis of the Abrams-Lloyd algorithm raises the question as to why the eigenvalue convergence for low dimensional problems (ie. small D) is less than that of optimal classical approaches. Part of the answer lies in the classical theory of matrix eigenvalue calculation. An important tool for numerical estimation of eigenvalues is the Krylov subspace, which is defined to be the span of the set $\{\psi, A\psi, A^2\psi, \ldots, A^{M'-1}\psi\}$ for some initial guess vector ψ, some hopefully small constant $M' < N^D$, and some $N^D \times N^D$ matrix A of which we seek several low order eigenvalues. The Krylov subspace is spanned by at most M' vectors, rather than the full N^D vector space of A, and so projecting onto the Krylov subspace gives an efficient means of estimating eigenvalues/eigenvectors of A. If the matrix whose lowest eigenvalue is sought is $\Lambda^{(N)}$, then we might choose $A = (\Lambda^{(N)} - \mu I)^{-1}$ where μ is an initial estimate of the eigenvalue sought (the eigenvalues of $\Lambda^{(N)}$ being simply related to those of A). With $A = (\Lambda^{(N)} - \mu I)^{-1}$, the vector $A^j\psi$ converges exponentially towards the eigenvector $\phi_f^{(N)}$ whose eigenvalue minimizes $|\lambda_f^{(N)} - \mu|$. In contrast, if $A = \exp(i\Lambda^{(N)})$ as in the Abrams-Lloyd algorithm, there is no such convergence towards a target eigenvector since the eigenvalues of A are of unit norm. The unitarity of quantum gates restricts eigenvalues to lie on the unit circle in the complex plane, which is a poor eigenvalue distribution from the perspective of estimating a target eigenvalue [46]. This leads to the question of whether controlled decoherence can be used to produce non-unitarity evolution to accelerate the selection of a target eigenvector *with a net reduction in gate operations/delay.*

CHAPTER 3

Threshold Error Penalty for Fault Tolerant Quantum Computation with Nearest Neighbour Communication

The error threshold for fault tolerant quantum computation with concatenated encoding of qubits is penalized by internal communication overhead. Many quantum computation proposals rely on nearest-neighbour communication, which requires excess gate operations. For a qubit stripe with a width of $L + 1$ physical qubits implementing L levels of concatenation, we find that the error threshold of 2.1×10^{-5} without any communication burden is reduced to 1.2×10^{-7} when gate errors are the dominant source of error. This $\sim$ 175X penalty in error threshold translates to an $\sim$ 13X penalty in the amplitude and timing of gate operation control pulses.

3.1 Introduction

A critical architectural issue for quantum computation is the internal communication of quantum information within the processor. There are a variety of proposed quantum processor implementations with different mechanisms for internal communication. For instance, the linear ion trap proposal of Cirac and Zoller [55] involves physical motion of massive ions for internal communication, as do proposals using more complex ion trap structures [56]. Alternative proposals involve using photons and cavity QED for communication [57]. The cavity QED approach has been extended to the solid state [58, 24]. Even direct transport of information carrying electrons has been suggested for the solid state [59, 60].

Our paper is motivated by another class of quantum computation proposals that rely upon local communication through nearest neighbour interactions [22, 21, 23]. For instance, communication among electron spins in semiconductors can be performed with sequential SWAP gate operations, generated by a controlled Heisenberg exchange between adjacent electrons. An appealing feature of the SWAP operation is that it is generated by the very same two-qubit interaction used for computational operations. Also, a substantial degree of parallelism can be employed. However, the protection of qubits with concatenated error correction requires communication between a number of physical qubits that grows exponentially with concatenation level. This exponential increase in SWAP operations might suggest that concatenated error correction will fail to reduce the logical qubit error rate. Gottesman [25], and Aharanov and Ben-Or [61] have pointed out that a threshold error exists despite an exponential increase in logical gate count with concatenation level L, although no attempt was made to quantify what that threshold might be. In this paper, we estimate that threshold.

The main result we report here is that the number of nearest neighbour communication operations is merely a constant factor over and above the necessary logical operations for error correction at each concatenation level L. Our estimated error thresholds are summarized in Table 3.1. We analyzed in detail fault-tolerant error correction with a concatenated [[7,1,3]] Calderbank-Shor-Steane (CSS) code [18, 19] on a linear qubit stripe with a width of $L+1$ physical qubits for L levels of concatenation, and find an ~ 175 fold reduction in threshold gate operation error due to nearest neighbour communication overhead. This translates to an ~ 13 fold increase in accuracy of control pulse amplitude and timing in gate operations. Although nearest neighbour communication incurs a significant penalty in the requisite experimental accuracy of qubit gate operations, it is not a fundamental obstacle to fault-tolerant computation in the solid-state. Our analysis is in general agreement with the recent work of Svore *et al.* [26], who also show that internal communication with local interactions incurs an error threshold penalty, although they do not fully account for all communication steps.

Our paper is organized as follows. In the first section, we describe the underlying architecture of a quantum processor composed of electron spin qubits, including a description of the physical layout of electron spin qubits and their grouping into concatenated CSS logical qubits. We describe a fault-tolerant error correction protocol in the second section. Our protocol implements error recovery without direct measurement. In the third section, we calculate the threshold error for gate operations under our error correction protocol, with various assumptions about available resources. The fourth section considers the relation between control pulse accuracy and gate error thresholds.

3.2 Layout Architecture

Given the problem of internal communication in a quantum processor, a higher dimensional architecture is preferred because it would allow qubits to be as close as possible. However, there must be access by control wires, thus limiting the packing geometry. Fig. 3.1 shows a schematic cross-section of a 2-D semiconductor qubit array controlled by gate electrodes accessing qubits from the side. The number of vertical stacked control electrodes is limited to twice the number of

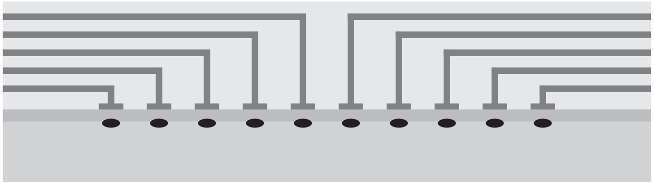

Figure 3.1: A schematic representation showing how the number of available metal wire layers limits the width of a 2-D qubit array to only about 10-20 qubits.

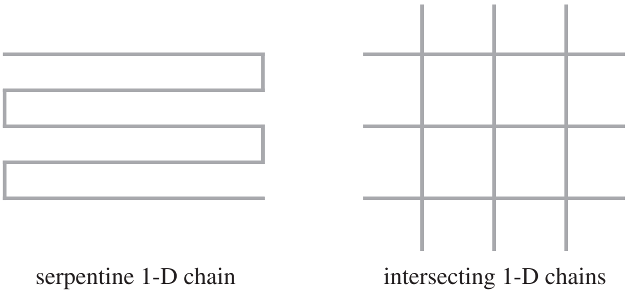

Figure 3.2: The requirement for gate electrode access to qubits restricts the layout to stripes of either serpentine or intersecting geometry.

metal wiring layers in the integrated circuit technology. The need for a reasonable fabrication yield limits the number of metallization layers to ~ 10, which means that the 2-D array can be at most 20 qubits wide. Fig. 3.1 illustrates the case for 5 metallization layers. In this respect, we agree with Copsey *et al.* [62], who pointed out this restriction specifically in the context of semiconductor qubits. Thus, while the qubit array might be locally 2 dimensional, the overall architecture will consist of 1-D stripes of moderate width, as illustrated in Fig. 3.2.

The lowest level of concatenated qubit encoding, $L = 1$, can be laid out along stripe width, but all higher concatenation levels must be laid out along the stripe length, and are effectively 1 dimensional. We are thus led to an essentially 1-D concatenation hierarchy, the most challenging for internal quantum communication.

Universal sets of fault-tolerant operations are known only for CSS error correcting codes of various size [16, 63, 61, 64]. In our work, we shall consider the $[[7,1,3]]$ CSS code. Concatenation [65], where each logical qubit is composed of encoded qubits, which are in turn composed of encoded qubits and so on, can suppress logical error rate to arbitrary degree, provided the physical error rates remain below a threshold value. The self-similarity of concatenation naturally leads to the self-similar logical structure illustrated in Fig. 3.3. There are 7 level $L-1$ logical qubits forming the CSS codeword that represents a single level L logical qubit $|\psi\rangle_L$. A minimum of two logical zeros, $|0\rangle_L$, and six initially arbitrary ancillae, $|a\rangle_{L-1}$, are required to perform error correction on $|\psi\rangle_L$. We consider $L+1$ parallel lines of physical qubits to implement error correction and computation with L levels of concatenation. The error correction protocol is described in detail in the next section. An important feature of the self-similar hierarchy is that at each concatenation level, the same qubit protection block is employed (for ancillae as well as information bearing qubits). Error correction can thus take place at any logical level within an appropriate logical qubit protection block.

3.3 Error Correction Protocol

For estimating error thresholds, we consider an aggressive error correction scheme where every unitary operation U_L at concatenation level L is followed by error correction E_L at level L, as illustrated in Fig. 3.4.

The error correction operation, E_L, can be implemented in a fault-tolerant manner with a Steane error correction circuit [66], slightly modified to that shown

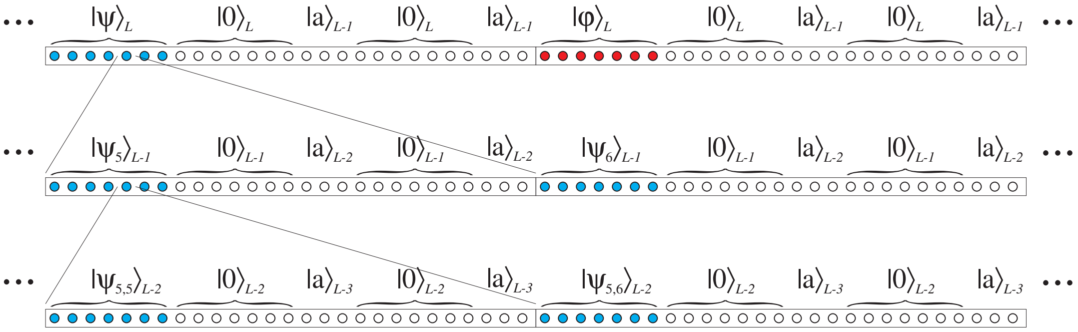

Figure 3.3: A self-similar concatenated hierarchy of logical qubits on a linear array, with concatenation level L down to $L-2$ shown. Error correction requires a minimum of two logical zeros, $|0\rangle_L$, and six ancillae, $|a\rangle_{L-1}$. Altogether, 27 level $L-1$ qubits are minimally required to protect a single level L qubit $|\psi\rangle_L$. The exponential growth with concatenation level L of *physical* nearest-neighbour operations to interact $|\psi\rangle_L$ and $|\phi\rangle_L$ is apparent. We consider a layout with $L+1$ adjacent linear arrays of qubits each organized according to the illustrated logical heirarchy.

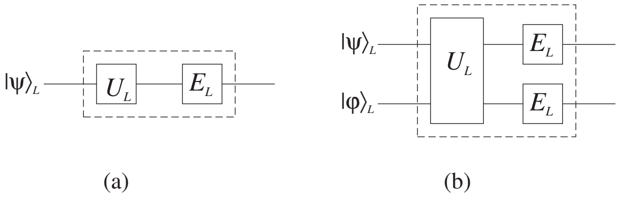

Figure 3.4: Each unitary operation U_L at logical level L is followed by error correction E_L at error correction level L.

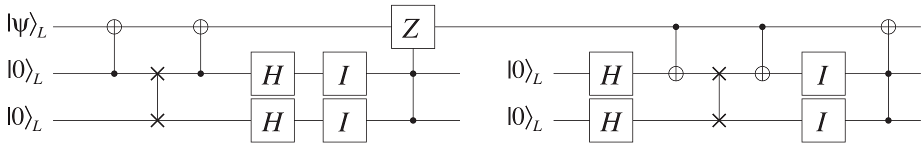

Figure 3.5: A modified Steane error correction circuit (E_L). The indicator block I computes an error syndrome, and decodes the syndrome into a bit-wise error indicator used for error recovery. The logical SWAP gate, as well as the CNOT gates, requires shuffling of the constituent $L-1$ qubits (see Fig. 3.8). We allow only nearest neighbour operations at all logical levels in adherence to self-similarity.

in Fig. 3.5. Error correction takes place within an error correction block, with the logical qubit $|\psi\rangle_L$ and logical zero states $|0\rangle_L$ explicitly shown. The two groups of three $L-1$ ancillae, $|a\rangle_{L-1}$, are made use of within the bit-flip *indicator* circuit, denoted by I. As can be seen in Fig. 3.5, the Steane error correction circuit is particularly parsimonious in its use of gate operations, and leads to particularly favorable error thresholds. The bit-flip indicator block I is essential, where for each logical zero $|0\rangle_L$ it computes a bit-flip error syndrome into three ancillae qubits $|a\rangle_{L-1}$. The syndrome is then decoded within the indicator block I into a bit-wise error indicator that can be directly used for error recovery. Note also that only nearest-neighbour operations at logic level L are employed, in strict adherence to self-similarity from the physical layer up to concatenation level L.

The key point about the bit-flip indicator block I is that it operates on logical zeros that have effectively measured the logical qubit error, but not the logical qubit itself, by virtue of a logical CNOT gate. As was pointed out by Boykin *et al.* [67], the identification of which operations require full quantum coherence and which operations do not is important since "quantum" operations require full protection against both phase-flip and bit-flip errors, while "classical" operations require protection against bit-flip errors only. Note from Fig. 3.5 that the outputs of indicator block I are used only as control bits for the error recovery operations acting upon the logical qubit. Arbitrary phase flips in the output of I have no effect on the logical qubit. Likewise, phase flips on the input of I have no effect on the logical qubit since the syndrome is encoded as bit-flips on the input to I. We need only protect against bit-flip errors in I, so that the operations within I can be thought of as essentially "classical" in nature, even though they are executed by physical qubit gates. Thus, I can in principle be protected with classical fault tolerance, which has been shown to be much more efficient than quantum fault tolerance [68], to ensure that the operations within I will contribute negligibly to the quantum error threshold.

Of course, the requisite logical zeros, $|0\rangle_L$, that allow for efficient fault-tolerant error correction are complex entangled states which must be created with low error probability to begin with. One approach to this problem is to dedicate adjacent quantum circuitry whose sole function is to prepare and purify logical zeros, providing a steady supply at various concatenation levels specifically for this purpose. Alternatively, the preparation of logical zeros can be performed directly within the qubit error protection block. The full error correction circuit is illustrated in Fig. 3.6. Purification of three $|0\rangle_L$'s, prepared by the 0_L block, results in a single $|0\rangle_L$ state for use in error correction. The 0_L zero preparation block is given in Fig. 3.7. Bit-flip errors are corrected with a modified indicator block I_P, which also corrects for a possible parity flip error corresponding to the logical zero being in the state $|1\rangle_L$ (and thus requiring a minimum of 4 ancillae). The qubit protection block must increase in size to accommodate $|0\rangle_L$ preparation in this case. A total of 46 qubits would be required, arranged in the following sequence of $L-1$ qubits (compare with Fig. 3.6): 7 qubits for storing $|\psi\rangle_L$, 7 qubits for storing a $|0\rangle_L$, 3 ancillae $|a\rangle_{L-1}$ for I, 7 qubits for storing a $|0\rangle_L$, 4

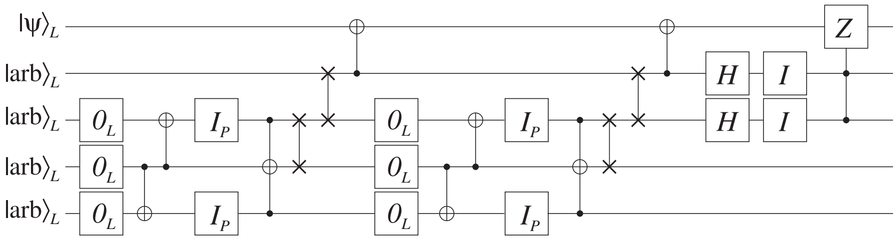

Figure 3.6: Error correction circuit (phase-error portion only) directly incorporating the preparation of requisite logical zeros. Ancillae begin in arbitrary states $|\text{arb}\rangle$. Three 0_L blocks prepare logical zeros that are purified into a single $|0\rangle_L$ state for use in error correction. A modified indicator block I_P corrects for possible parity errors in the raw $|0\rangle_L$'s.

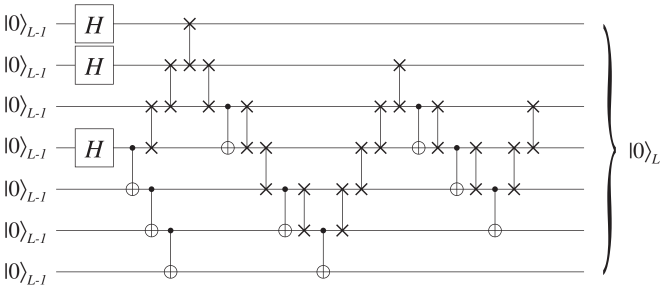

Figure 3.7: Circuit 0_L for preparation of a single logical zero $|0\rangle_L$ from lower level $|0\rangle_{L-1}$'s. Only nearest neighbour operations are employed.

ancillae $|a\rangle_{L-1}$ for I_P, 7 qubits for storing a $|0\rangle_L$, 7 qubits for storing a $|0\rangle_L$ and 4 ancillae $|a\rangle_{L-1}$ for I_P.

3.4 Error Threshold Penalty

The number of physical qubits for our concatenated CSS encoding required to store and protect one logical qubit is 27^L (or 46^L including logical zero preparation). Several levels of concatenation already leads to a large number of physical qubits (although the width of the qubit stripe grows only as $L+1$). Likewise, the number of physical gate operations grows exponentially, N^L, where N is approximately the number of logical operations required at level $L-1$ in order to implement a single logical function at level L. For example, with a single level of encoding, N is simply the number of physical gate operations required to perform some function on our 7-qubit CSS code word (or multiple code words in the case of a multi-qubit logical function).

The number of gate operations N will depend on the function being performed. We consider implementing a simple two-qubit unitary, U_L, followed by error correction, E_L, as illustrated in Fig. 3.4(b). Error correction might require $N = N_E$ logical gate operations at level $L-1$. There will be additional logical SWAP operations at level $L-1$ required to move qubits around since only nearest-neighbour interactions are permitted. We let N_{Ec} be the number of required nearest-neighbour SWAP communication operations, bringing the total number of level $L-1$ operations to $N = N_E + N_{Ec}$. Of course, the unitary U_L will require N_U operations at level $L-1$, as well as N_{Uc} additional communication operations at level $L-1$. The total gate operation count at level $L-1$ to implement U_L followed by E_L is simply $N = N_U + N_{Uc} + N_E + N_{Ec}$. The total *physical* gate count is again approximately $N^L = (N_U + N_{Uc} + N_E + N_{Ec})^L$ because each of the N operations at $L-1$ is simply a unitary U_{L-1} followed by error correction E_{L-1}. The self-similar hierarchy requires that N operations at $L-2$ are required for each operation at $L-1$ and so forth, including communication.

In reality, the gate count $N_U + N_{Uc}$ varies among the various logical qubit operations possible. For instance, Hadamard at level L requires $N_U = 7$ Hadamard gates at level $L-1$ and $N_{Uc} = 0$ communication gates. In contrast, the gate operations $N_U + N_{Uc} = 7 + 42$ involved in a logical SWAP on the same qubit line are illustrated in Fig. 3.8 for adjacent logical qubits. Clearly the number N^L can be very large, although a substantial fraction of operations at each logical level can be performed in parallel. Note the fault tolerance of the logical SWAP gate: a single swap gate failure induces one error in each logical qubit, which can be recovered independently by error correction. Of course, the extra qubits involved in a qubit protection block increases the number of communication swaps N_{Uc}. As a final example, we show the partial sequence of gate operations required for the logical CNOT gate in Fig. 3.9. It is in implementing the CNOT gate that an additional line of qubits is used for every concatenation level, resulting in a total of $L + 1$ lines of qubits. Similar sequences are used for the SWAP and CNOT gates required for the error correction operation E_L, contributing to $N_E + N_{Ec}$.

Despite the exponential increase in physical qubits and physical gate operations with concatenation level (while the width of the stripe merely grows linearly in concatenation level), logical errors are suppressed double-exponentially with concatenation level. We let P_1 be the logical error probability on a first level encoded state, $|\psi\rangle_1$, after a two qubit unitary followed by a single error correction cycle. By the fault tolerant construction of U_L and E_L, the probability of a logical error is bounded above by the probability that two gate operations fail,

$$P_1 \leq \left(\begin{array}{c} N \\ 2 \end{array} \right) \epsilon^2 \simeq \frac{N^2}{2} \epsilon^2, \tag{3.1}$$

where ϵ is the probability of physical gate error, assumed to be equal for all gates, and $N = N_U + N_{Uc} + N_E + N_{Ec}$ as before. While logical error rates shall vary slightly due to differences in $N_U + N_{Uc}$ amongst the logical gate operations with the dominant $N_E + N_{Ec}$ remaining fixed, a conservative estimate can be had by taking the gate counts for the logical CNOT gate as representative. The criterion for error correction to reduce the likelihood of qubit error is $P_1 < \epsilon$.

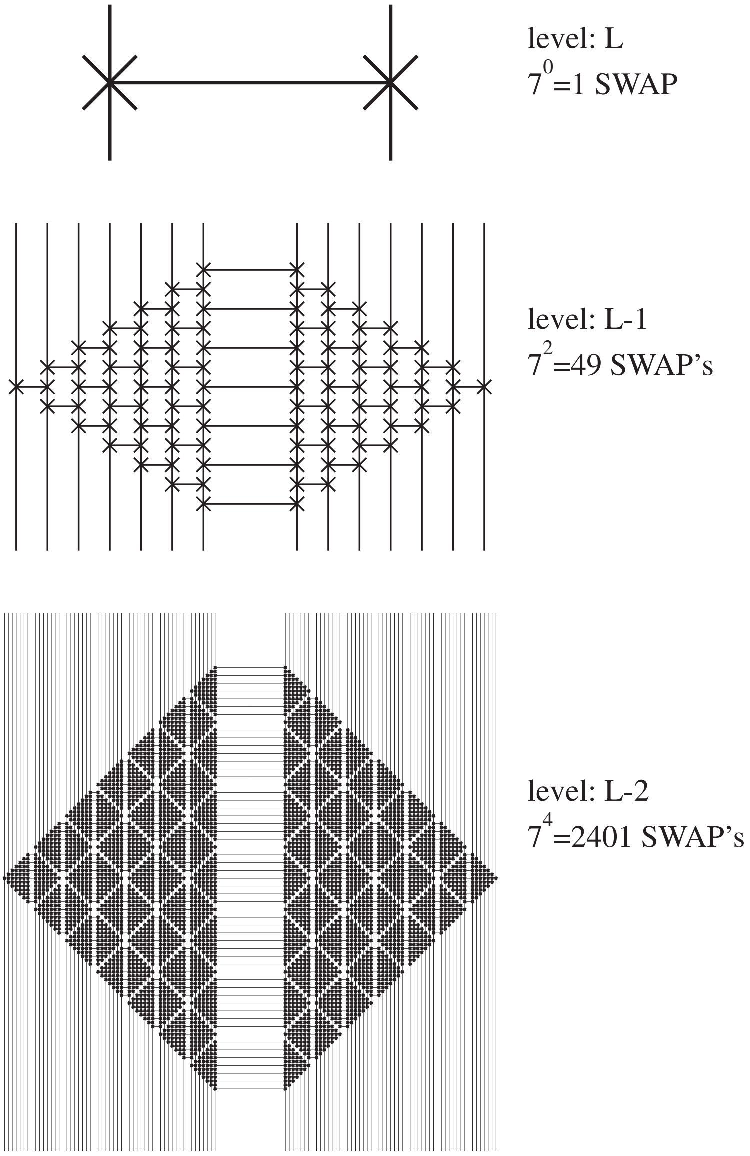

Figure 3.8: A logical SWAP operation illustrated at concatenation levels L through $L-2$ with nearest neighbour interactions only. The number of level $L-1$ SWAPS required to implement a single level L SWAP between adjacent logical qubits is $N_U + N_{Uc} = 7 + 42$. There are 21 level $L-1$ SWAPs to interleave the qubits, 7 level $L-1$ qubit-wise SWAPs, and 21 level $L-1$ SWAPs to undo the interleaving. Note that a single gate failure does not produce correlated errors within a logical qubit. Error correction, and swapping through the additional qubits in a qubit protection block, are omitted here for clarity.

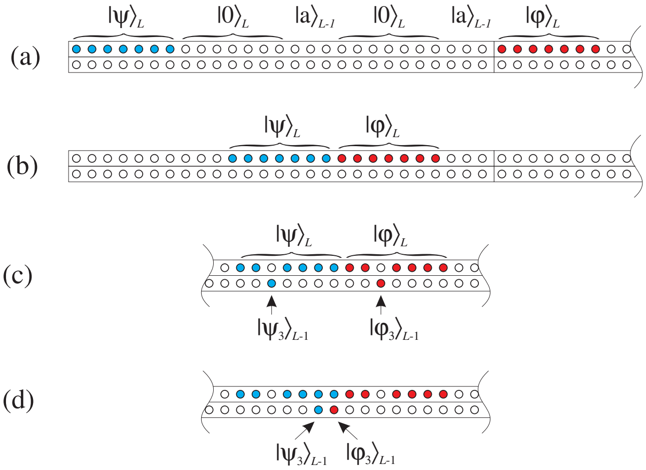

Figure 3.9: Partial sequence for a logical level L CNOT operation illustrated at concatenation level $L-1$ with nearest neighbour interactions only. The (a) logical code words $|\psi\rangle_L$ and $|\varphi\rangle_L$ are (b) first brought into adjacent positions, then (c) each of the 7 constituent $L-1$ qubits are moved into an adjacent qubit row to be (d) brought together for qubit wise interaction (only the third qubits $|\psi_3\rangle_{L-1}$ and $|\varphi_3\rangle_{L-1}$ are shown interacting). The logical qubits are brought back to their original positions for error correction after the logical CNOT. The scheme is applied recursively until physical CNOT gates are performed in the $L+1^{st}$ row. The CNOT gates for the error correction circuit are similarly implemented. Note that a single gate failure does not produce multiple errors within a logical qubit.

Table 3.1: The gate count for error correction, $N_E + N_{Ec}$, and for logical CNOT operations, $N_U + N_{Uc}$, under different assumptions of internal communication resources and ancilla preparation. Approximate threshold gate error probabilities are given, as well as control pulse accuracy thresholds (see text for details).

		Error Correction Gate Count $N_E + N_{Ec}$	Two-Qubit Unitary Gate Count $N_U + N_{Uc}$	Error Probability Threshold $P_{\text{th}} = 2/N^2$	Gate Accuracy Threshold ϕ_{th}
nonlocal	no $\|0\rangle_L$	70	7	3.4×10^{-4}	2.1°
	$\|0\rangle_L$	298	7	2.1×10^{-5}	0.52°
re-CNOT	no $\|0\rangle_L$	238	35	2.7×10^{-5}	0.60°
	$\|0\rangle_L$	1090	35	1.6×10^{-6}	0.14°
SWAP	no $\|0\rangle_L$	1008	203	1.4×10^{-6}	0.13°
	$\|0\rangle_L$	3754	343	1.2×10^{-7}	0.034°

This leads to the threshold error condition $\epsilon < 2/N^2$. Likewise, at higher levels of concatenation,

$$P_L \leq \begin{pmatrix} N \\ 2 \end{pmatrix} P_{L-1}^2 \simeq \frac{N^2}{2} P_{L-1}^2, \tag{3.2}$$

leading to $P_{L-1} < 2/N^2 = P_{\text{th}}$ being the error threshold condition for all L. The corresponding required phase accuracy for gate operations, as described in section 3.6, is $\phi = 2\sqrt{2}/N$. From the above relations, we arrive at the standard logical error probability for concatenated error correction,

$$P_L \leq P_{\text{th}} \left(\frac{\epsilon}{P_{\text{th}}} \right)^{2^L} \tag{3.3}$$

but where N now includes the nearest neighbour communication overhead at a particular concatenation level. The exponent 2^L results in an overwhelming, *super*-exponential in L suppression of logical errors while the number of qubits and gate operations increase only exponentially in L.

Suppose that a quantum computation requires a sequence of T logical gate operations, then a logical error probability $P_L = 1/T$ will give the correct result with only several trials of the computation. The relation between the number, T, of operations in a calculation and concatenation level L can be written,

$$T \geq \frac{1}{P_{\text{th}}} \left(\frac{P_{\text{th}}}{\epsilon} \right)^{2^L} \tag{3.4}$$

or alternatively,

$$L \leq \log_2 \left(\frac{\log_2(T P_{\text{th}})}{\log_2(P_{\text{th}}/\epsilon)} \right) \tag{3.5}$$

For instance, the error threshold might be $P_{\text{th}} = 10^{-6}$ while the physical gate operation error is an order of magnitude better, $\epsilon = P_{\text{th}}/10 = 10^{-7}$. We then have an accessible computation length $T = 10^6 \times 10^{2^L}$, which for $L = 3$ gives

$T \geq 10^{14}$. It follows that interesting calculations can be performed with only a few layers of concatenation (ie. a qubit stripe with a width of only a few qubits) if physical error probabilities well below the error threshold can be achieved.

The problem of estimating error threshold has been reduced to counting gate operations, for which our numerical results are summarized in Table 3.1. Note that we have neglected storage errors in our present analysis since the coherence times of electron spins in semiconductors [69] exceed the expected gate operation times by at least ~ 8 orders of magnitude, with further improvement expected. The top row of Table 3.1 gives the most favourable error thresholds where any qubit can interact with any other qubit without any extra communication operations. The bottom row is the least favourable case where nearest neighbour SWAP operations are used on a linear qubit array to implement all operations. The middle row represents an intermediate case, where the remote-CNOT is used to perform a CNOT gate between distant qubits [70, 54]. The remote-CNOT requires a shared EPR pair, a resource that might be generated by independent hardware with sufficient purity that the EPR error rate contributes negligibly to the overall error rate of the remote-CNOT and the error threshold. Measurement and classical communication are also required for the remote-CNOT (see appendix).

For all three communication schemes, the gate count is given in Table 3.1 for sub-cases where $|0\rangle_L$'s are supplied by adjacent circuitry (a parallel qubit stripe, for instance); or where the $|0\rangle_L$'s are prepared directly within the error correction circuit itself (as in Fig. 3.6) thus burdening the error threshold. In the former case, we assume that the adjacent circuitry can prepare and purify logical zeros to reach an error probability much less than the preparation circuit of the former case, thereby contributing to the error threshold negligibly. This might be achieved by successive rounds of purification.

In all cases, we assume that those portions of the circuit that can be implemented with classical fault-tolerant logic [67], albeit with qubit gates, take advantage of the greater efficiency of classical coding. The threshold error for classical fault-tolerant circuits has been estimated to be between $\sim 1/100$ to $\sim 1/3000$ depending on topology and communication resources [68], we therefore assume the error rates in the classical circuits are negligible compared to the quantum circuits, so that in counting the gate operations we can neglect the operations in I and I_P. Furthermore, the dual-control phase-flip ($\Lambda_2(Z)$) and dual-control bit-flip ($\Lambda_2(X)$ =Toffoli) are assumed to count merely as two-qubit interactions, since fault-tolerant classical logic can be used to generate a single classical control bit. The remaining sundry details involved in counting gate operations are left to the appendix.

Observing the gate error thresholds in Table 3.1, we see that SWAP communication incurs a penalty of $\sim$ 175X compared to the case of free communication. Communication through the remote-CNOT incurs a penalty of $\sim$ 12X compared to the free communication case. The improvement associated with remote-CNOT communication is not as much as one might expect, since the remote-CNOT requires multiple operations proportional to the size of the logical qubits. Thus, internal quantum communication reduces gate error thresholds for fault tolerant computation by a substantial factor that we estimate to be from $\sim$ 12X to $\sim$ 175X. While this certainly increases the difficulty in experimentally realizing fault tolerant gate operations, it is by no means an impasse for solid state quantum computation.

3.5 Threshold Error Calculations

We provide a brief summary here of the counting of gate operations, which then leads to the threshold error. Error correction at concatenation level L with the circuit E_L requires the use of both single qubit unitaries and two qubit unitaries at levels L down to the physical layer. Interestingly, the quantum portions of the circuit E_L (see Figs. 3.5 or 3.6) consists of gate operations that are directly fault tolerant, where qubit-wise (or transversal) operations are sufficient. These operations include CNOT, SWAP, and H (Hadamard rotation). The control bits of the dual control gates are classical, so a full quantum Toffoli is not required. Of course, indirectly fault tolerant gates such as the Toffoli ($\Lambda_2(X)$) or $\pi/8$ rotation ($Z^{1/4}$) are required for universal computation. We do not calculate the error threshold for indirectly fault tolerant gates here.

3.5.1 Free Communication

First, we consider the idealized case where communication is achieved without any extra operations, in other words, any two-qubits can interact directly at any time. In this case, $N_{Uc} = N_{Ec} = 0$ and we need only count the number of computationally useful gates. A directly fault tolerant two-qubit unitary will require $N_U = 7$ operations. The error correction gate count without logical zero preparation is,

$$N_E = 4 \times 7\text{CNOT} + 4 \times 7\text{H} + 7\Lambda_2(X) + 7\Lambda_2(Z) = 70 \tag{3.6}$$

where the $L-1$ gate type and count are indicated. With logical zero preparation, we have,

$$N_E \;\; = \;\; 70 + 12 \times 0_L + 4 \times 7\Lambda_2(X) + 8 \times 7\text{CNOT}$$

$$\begin{aligned} &= 70 + 12 \times (3\text{H} + 9\text{CNOT}) + 84 \\ &= 298 \end{aligned} \tag{3.7}$$

where again $L-1$ gate type and count was indicated.

3.5.2 remote-CNOT communication

Next, we consider the intermediate communication case involving remote-CNOT operation, which we abbreviate as reCNOT. The reCNOT circuit is indicated in Fig. 3.10. For simplicity, we assume that the classical communication and EPR preparation introduce negligible errors compared to the other gate operations involved. We see that a reCNOT between two level $L-1$ qubits requires 5 level $L-1$ operations, so that a reCNOT between two level L qubits requires $N_U + N_{Uc} = 5 \times 7$ level $L-1$ operations. The error correction gate count without logical zero preparation becomes,

$$\begin{aligned} N_E &= 4 \times 7\text{reCNOT} + 4 \times 7\text{H} + 7\Lambda_2(X) + 7\Lambda_2(Z) \\ &= 140 + 28 + 35 + 35 \\ &= 238 \end{aligned} \tag{3.8}$$

where $\Lambda_2(X)$ and $\Lambda_2(Z)$ are counted as reCNOT operations (recall they can be implemented with single classical control bits). With logical zero preparation, we have,

$$\begin{aligned} N_E &= 238 + 12 \times 0_L + 4 \times 7\Lambda_2(X) + 8 \times 7\text{reCNOT} \\ &= 238 + 12 \times (3\text{H} + 3\text{CNOT} + 6\text{reCNOT}) \\ &\quad +140 + 280 \\ &= 238 + 432 + 140 + 280 = 1090 \end{aligned} \tag{3.9}$$

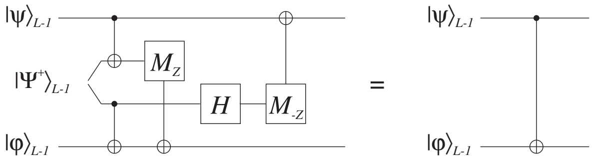

Figure 3.10: The remote CNOT gate requires a shared EPR pair, $|\Psi^+\rangle = (|01\rangle + |10\rangle)/\sqrt{2}$, measurement, M_Z, and classical communication to implement a CNOT operation between distant qubits.

where we have made use of both nearest neighbour CNOT and reCNOT in the logical zero preparation.

3.5.3 SWAP communication

Finally, we consider communication by SWAP gates. Without logical zero preparation, a level L qubit protection block is 27 $L-1$ qubits long. Applying CNOT between two level L qubits as in Fig. 3.9 requires $N_U + N_{Uc} = 203$ level $L-1$ operations on each logical qubit argument. The error correction operation requires,

$$\begin{aligned} N_E &= 4 \times (7\text{CNOT} + 112\text{SWAP}) + \\ &\quad 4 \times 7\text{H} + 2 \times (7\text{SWAP} + 84\text{SWAP}) + \\ &\quad (7\Lambda_2(X) + 154\text{SWAP}) + (7\Lambda_2(Z) + 154\text{SWAP}) \\ &= 1008 \qquad (3.10) \end{aligned}$$

where we note that 112 communication SWAPs are required for applying CNOT between $|\psi\rangle_L$ with an adjacent $|0\rangle_L$, and 84 communication SWAPs are required for logical swapping of a $|0\rangle_L$ with another $|0\rangle_L$ taking account of the extra ancillae $|a\rangle_{L-1}$ in the way.

When logical zero generation is included, the qubit protection block increases in size to 46 qubits. Applying CNOT between two level L qubits now requires

$N_U + N_{Uc} = 343$ level $L-1$ operations because of the increased size of the qubit protection block. The error correction operation requires,

$$\begin{aligned} N_E &= 1008 + 12 \times 0_L + 2 \times (7\text{SWAP} + 84\text{SWAP}) + \\ &\quad 4 \times (7\text{SWAP} + 98\text{SWAP}) + 4 \times (7\text{CNOT} + \\ &\quad 112\text{SWAP}) + 4 \times (7\text{CNOT} + 168\text{SWAP}) + \\ &\quad 4 \times (7\Lambda_2(X) + 154\text{SWAP}) \\ &= 3754 \end{aligned} \tag{3.11}$$

where we note that each logical $|0\rangle_L$ generation requires 27 level $L-1$ operations (Fig. 3.7), and the SWAP communication accounts for all extra ancillae $|a\rangle_{L-1}$ in the way.

3.6 Error Probability and Gate Operation Accuracy

So far, we have worked entirely with error probabilities. In practice, experimental gate accuracy is more naturally specified in terms of control pulse amplitude. Consider the spin (or a qubit pseudo-spin), illustrated in Fig. 3.11. Suppose a control pulse, as used in spin resonance, was to bring the spin into alignment with the x-axis. However, an error in pulse area, phase, or timing may cause a misalignment by some small angle ϕ. The probability of error, ϵ, is then the probability that the spin is not projected into the +x direction when a measurement is performed along the x-axis. The probability of projection along the +x direction is $\cos^2(\phi/2)$, so that the error probability is,

$$\epsilon = \sin^2(\phi/2) \approx (\phi/2)^2. \tag{3.12}$$

The required gate timing and amplitude accuracy is $\phi = 2\sqrt{\epsilon}$, specified as a phase angle, is proportional to the *square root* of the threshold error probability. The

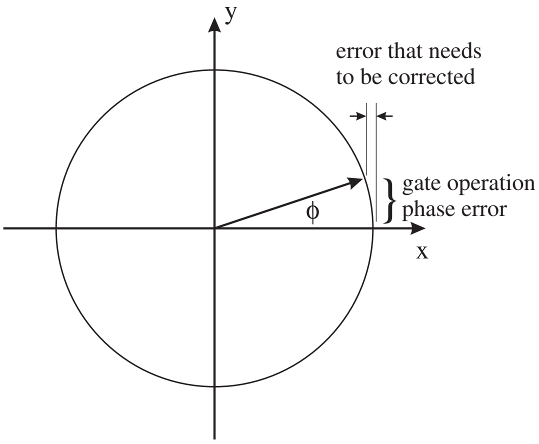

Figure 3.11: A conceptual illustration of a qubit pseudo-spin that might miss a target x-axis by an angle ϕ due to a control pulse error. The resulting probability of qubit error is $\epsilon \approx (\phi/2)^2$.

gate accuracy thresholds are given in degrees in Table 3.1. Of course, the $\sim$ 12X to $\sim$ 175X penalty in error probability threshold becomes only a $\sim$ 3.5X to $\sim$ 13X penalty in control pulse accuracy. In order to achieve an error probability of 10^{-7}, one would require about 1/30 of a degree accuracy in control pulse timing, which is not entirely infeasible since it would require about 1 picosecond phase accuracy in a clock period of about 10 nanoseconds. Recall that an error probability of 10^{-7} for a quantum processor with threshold error probability 10^{-6} and 3 levels of concatenation will allow a computation with $\geq 10^{14}$ operations. Thus, thinking about gate errors in terms of phase angle makes it clear that very small error probabilities are achievable.

3.7 Conclusions

Internal quantum communication remains a challenging architectural problem that impacts the threshold error for fault-tolerant computation with encoded logical qubits. The communication operation overhead required to distribute information among a number of qubits that grows exponentially with concatenation level can be a significant burden. Whether one is limited to nearest-neighbour communication, a communication bus (as in the original Cirac-Zoller ion trap proposal [55]), or communication by modified teleportation schemes such as the remote-CNOT, there is always a communication penalty in error threshold. The minimum communication overhead cost is associated with a communication bus, where a single operation for "transmitting" and a single operation for "receiving" is possible in principle. The question of whether a sufficiently robust communication bus is available for solid state qubits remains open. Ballistic transport of electron spins through mesoscopic wires is predicted to give error rates of ~ 0.6 for GaAs [60], far above our stated threshold requirements even for the free communication case. Much more promising is the combination of cavity QED techniques with confined electron spins [58] or superconducting circuits [24], where an electromagnetic bus can couple a number of qubits. The error rates of such a bus, the reconfigurability of its links, and its parallelism (ie. how many qubits can be transported simultaneously? through the same link?) must all be carefully considered in determining what benefits, if any, we can expect over nearest neighbour architectures. Nonetheless, we expect that communication overhead can be mitigated to a large extent by circuit optimization. Recent work [71] on laying out Shor's factorization algorithm on a linear chain of qubits under the restriction of nearest neighbour interaction has shown that circuit optimization can greatly reduce the number of logical qubit SWAPs required.

CHAPTER 4

Photoelectron Trapping, Detection and Storage

We report the filling of a quantum dot - defined by gate electrodes on a GaAs/Al$_{0.3}$Ga$_{0.7}$As modulation doped heterostructure - by single photoelectrons, originating from weak photon pulses. The gate defined quantum dot could be emptied in a controlled fashion before the arrival of each weak photon pulse. Trapped photoelectrons were detected by a quantum point contact transistor integrated adjacent to the electrostatic potential trap. Each stored photoelectron caused a persistent negative step in the transistor channel current. Such a controllable, non-invasive, single photoelectron detector could allow for information transfer between "flying" photon spin qubits and "stationary" electron spin qubits.

4.1 Introduction

In order to implement spin coherent photodetection [37], wherein the spin of an incident photon is transferred to the spin of a localized photoelectron, one must be able to trap, store and detect individual photoelectrons in quantum dot structures suitable for quantum computing. The most common means to detect single photons, and thus single photoelectrons, is by an avalanche multiplication mechanism such as that found in photomultiplier tubes and avalanche photodiodes [72]. The need to store individually generated photoelectrons prohibits the use of avalanche photomultiplication, but can instead be achieved by the conductive gain of a transistor, which can easily give single electron charge sensitivity [73, 74]. In the context of detecting photo generated electrons, this same mechanism is known as photoconductive gain [75].

The trapping, storage and detection of individual photogenerated carriers has been reported under a variety of conditions. The most spectacular has been the detection of single photoelectrons and photoholes generated by far-infrared ($\lambda \approx 200\mu$m) photons in a GaAs/AlGaAs heterostructure [76]. Incident far-infrared photons were used to excite either a hole or electron in a quantum Hall droplet, resulting in a persistent photoconductance change in an adjacent quantum Hall edge state defined by gate electrodes. This detector holds the record for sensitivity in the far-infrared [77] and has been applied for cryogenic imaging [78]. In other work, photoholes created by direct bandedge absorption have been trapped and detected in self-assembled InAs quantum dots [79] and even Si dopant related traps [80] located adjacent to field effect transistors in AlGaAs/GaAs heterostructures. In the case of quantum communication, it is desired to trap photoelectrons rather than photoholes due to the increased spin coherence lifetime of electrons versus holes. A potential well that binds electrons rather than holes is thus desired, requiring a new device structure.

4.2 Device Structure

The device was based upon an AlGaAs/GaAs heterostructure as illustrated in Fig. 4.1. The quantum well formed at the interface accumulates electrons that form a 2-dimensional electron gas (2DEG). The semiconductor material was grown by commercial (IQEP, Bethlehem Pennsylvania) molecular beam epitaxy on a semi-insulating GaAs substrate. At a temperature of 4.2K, the 2DEG

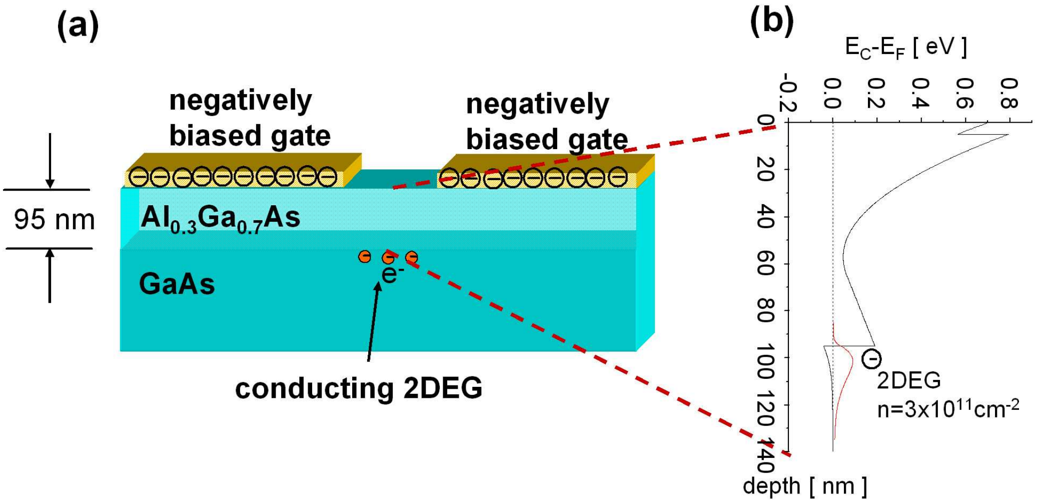

Figure 4.1: (a) A 2DEG at the heterointerface between AlGaAs and GaAs can be selectively depleted into pools and channels with negatively biased surface electrodes. (b) The conduction band edge as calculated from a self-consistent Schrödinger-Poisson equation. The heterolayers are: a 5 nm Si-doped ($1 \times 10^{18}/\mathrm{cm}^3$) GaAs cap layer, a 60 nm Si-doped ($1 \times 10^{18}/\mathrm{cm}^3$) n-$\mathrm{Al}_{0.3}\mathrm{Ga}_{0.7}\mathrm{As}$ layer, a 30 nm i-$\mathrm{Al}_{0.3}\mathrm{Ga}_{0.7}\mathrm{As}$ spacer layer, atop an undoped GaAs buffer of several hundred nm thickness.

in an unprocessed sample was determined by quantum Hall measurements to have a density of $3 \times 10^{11}/\mathrm{cm}^2$, although this could be increased with exposure to light. The sample was measured to have a low mobility of approximately $\mu \sim 200000\mathrm{cm}^2/\mathrm{Vs}$ at 4.2K.

The device structure employed for photoelectron trapping, storage and detection is illustrated in Fig. 4.2. The gate elecrodes depicted in Fig. 4.2(a) were fabricated by electron beam lithography and electron beam evaporation of Ti/Pt/Au. Gates G1 and G2 define a quantum point contact (QPC) when biased with a negative voltage. Current could be passed through the QPC via the source and drain Ohmic contacts, S_{QPC} and D_{QPC}. Adjacent to the QPC, a

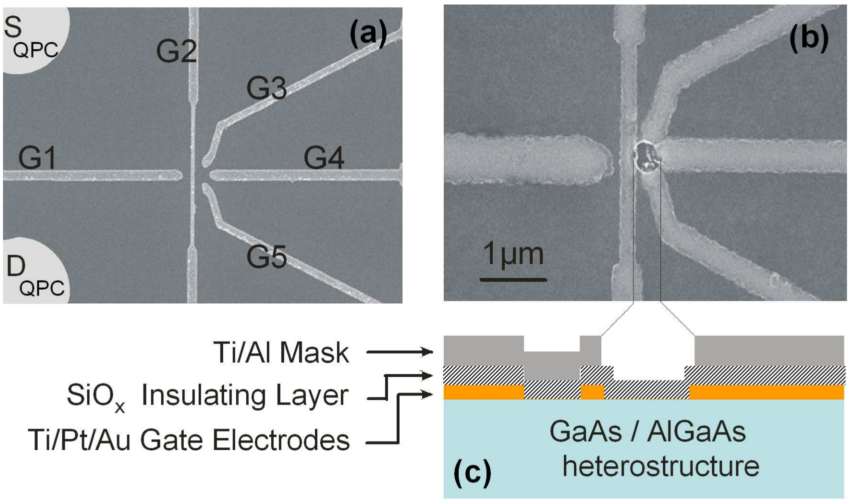

Figure 4.2: (a) Scanning electron graph (SEG) of the surface metallic gates defining a quantum point contact between the source and drain Ohmic contacts. (b) SEG of pinhole aperture etched in an opaque Al/Ti layer, 150 nm thick, acting as a shadow mask to allow illumination of the quantum dot region only. (c) Cross–sectional view of the device structure showing gates buried under Al/Ti/SiO$_x$ layers.

circular quantum dot with a lithographic radius of ∼200 nm was defined by gates G3-G5. A negative voltage applied to G3-G5 would deplete the 2DEG in such a way as to isolate a pool of electrons.

Similar structures employing the integration of a quantum dot adjacent to a quantum point contact have been used to study the distribution of charge in quantum dots [81, 82, 83, 84]. In such structures, the proximity of the quantum dot to the QPC results in sufficient Coulomb interaction between charge on the quantum dot and the electron current in the QPC so that the QPC can function as a near optimal probe of the charge on the quantum dot [85].

Photogeneration of electron-hole pairs was suppressed over the bulk of the sample by a 150 nm thick Al layer deposited as a mask over the entire device area, except for a pinhole aperture directly above the quantum dot, as shown in the scanning electron graph of Fig. 4.2(b). Suppressing bulk carrier generation was essential due to the prevalence of hole trapping at Si donor related traps in the AlGaAs layers, resulting in modulation of 2DEG density and mobility [86, 87] that can overwhelm the desired photoconductivity signal. An $\mathrm{SiO_x}$ layer and a thin adhesion layer of Ti deposited by electron beam evaporation insulated the metal gate electrodes from the Al mask layer.

4.3 Electrical Characterization

Prior to exposure to photons, the sensitivity of the QPC current to charge on the quantum dot was characterized. All experiments were performed with the device mounted on the tail of a top-loading probe in a ^{3}He cryostat (CryoIndustries of America). The sample was immersed in ^{3}He at a base temperature of 430mK - the temperature being limited by poor cryostat design and manufacture. Gates were biased with tuneable voltage sources and the QPC conductance was measured with a low noise transimpedance amplifier (FEMTO LCA-400K-10M with $65\text{fA}/\sqrt{\text{Hz}}$ equivalent input noise current) feeding an analog-digital convertor.

Figure 4.3 plots the current through the QPC versus the voltage on the plunger gate, G4. The plunger was swept at a rate of ∼4 mV/s to push electrons out of the dot, one at a time, into the adjacent 2DEG reservoirs. The charge state of the quantum dot, at the start of the scan in Fig. 4.3, is the same as that at time t_0 (or equivalently t_6) in Fig. 4.4. Upon formation of the dot, a few excess electrons remained trapped in the dot in long-lived metastable states, prior to being pushed out by an increasingly negative plunger gate potential. The QPC current varied in a sawtooth fashion with a small discrete positive step for each electron ejected, as seen in Fig. 4.3. The last electron ejection occurred on curve (c) at a plunger (G4) voltage of ∼-2.75V. In order to ensure that the absence of further steps is not due to very slow tunnelling times, the barrier gate voltage G3 was raised to transparency after the last detected step to allow any remaining electrons to escape. Only a smooth increase in the QPC current was observed due to the capacitive coupling between the QPC and the tunnel barrier gate, with no evidence of remaining electrons (in the form of discrete QPC current changes). The lower inset to Fig. 4.3 shows the steps corresponding to the last two electrons after subtracting out the background slope due to capacitive coupling from the plunger gate to the QPC. The observed single electron step size of about 500 pA per electron on a background of $\sim$ 50nA represents a QPC conductance change of the order of 1%.

Upon sweeping the plunger gate G4 from -4.0V back to -1.5V, at a scan rate of 4mV/s, no electrons were observed to reenter the dot. This came about because the last few electrons remained trapped at a potential energy lying well above

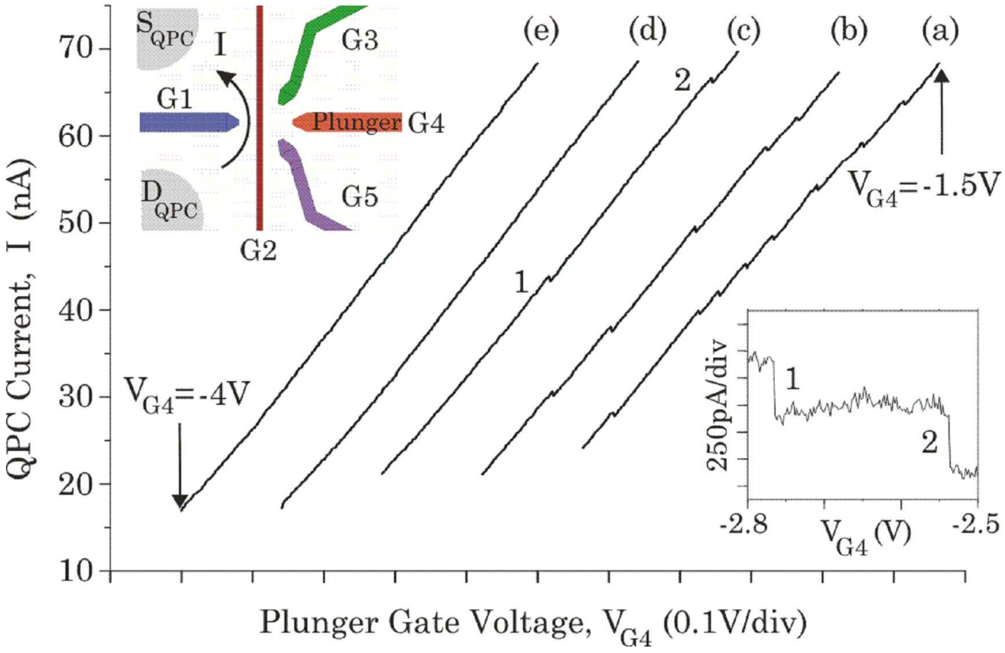

Figure 4.3: Single electron escape from the dot detected by the QPC transistor. The plunger gate, G4, was swept from -1.5V to -4V with a scan rate of -4mV/s starting at curve marked (a) and ending at (e) with each curve spanning 0.5V. Gates G2, G3, and G5 were fixed at -0.9V while G1 was adjusted prior to each curve for optimum QPC sensitivity. The curves are compressed along the G4 voltage axis for compactness. The inset shows the current step sizes of the last two electrons observed in curve (c) after subtraction of the background slope ($V_{SD,QPC}$=3.25mV, $G_{QPC} = 0.35e^2/h$ at the last electron step).

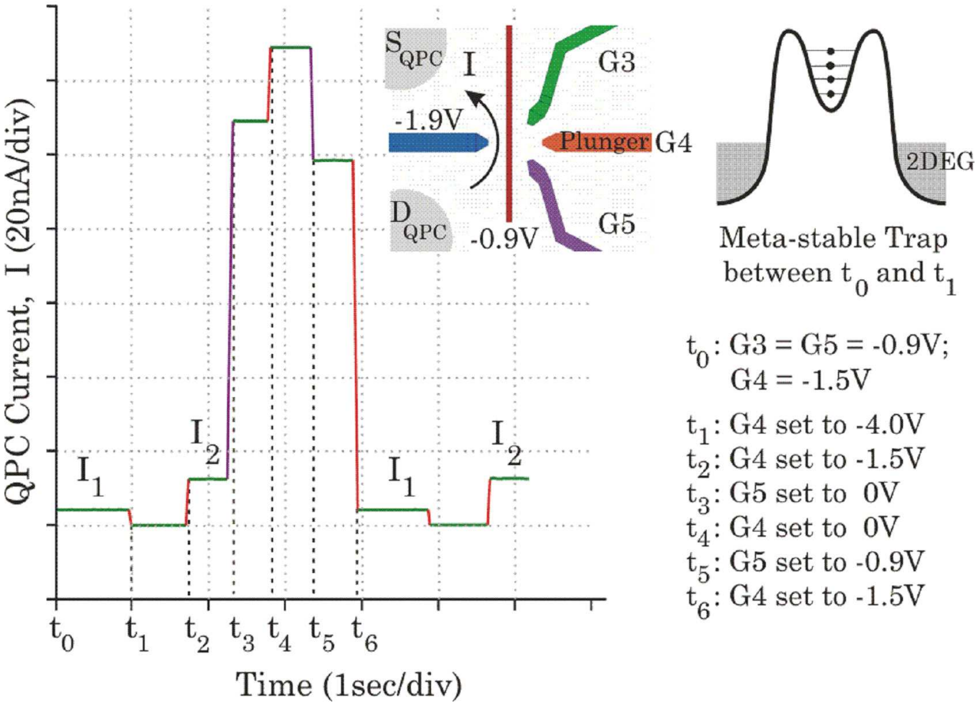

Figure 4.4: Hysteresis measured in the current through the QPC transistor, associated with the transition of the dot from the metastable filled state to the equilibrium empty state. See text for details.

the Fermi level in the surrounding 2DEG, as shown schematically in the right inset to Fig. 4.4. Quantum dots strongly isolated from adjacent reservoirs by opaque tunneling barriers have been previously observed in metastable states for durations exceeding tens of minutes [88].

Figure 4.4 illustrates the metastable/hysteretic behavior in the QPC current associated with the trapping and ejection of electrons from the quantum dot. Immediately following time t_0 (equivalent to time t_6), the quantum dot was prepared in a metastable state with excess trapped electrons. This was achieved by simply biasing the gates G1-G5 from 0V to the respective negative values indicated in Fig. 4.4. No electrons were observed to escape in the interval between times t_0 and t_1. At t_1, the plunger gate G4 was set to -4.0V, overwhelming the quantum dot barriers and forcing all electrons out of the quantum dot. The plunger gate voltage was sufficiently negative to pinch off the QPC channel. At the time t_2, the plunger potential was retracted to -1.5V. The difference in QPC current I_1 (during the interval $t_1 - t_0$) and the QPC current I_2 (during the interval $t_3 - t_2$) corresponds to the effect of the change in quantum dot charge state from filled to empty. Opening and closing the barriers in the interval $t_6 - t_3$ allows the dot to become repopulated with electrons. However, it is with the quantum dot in an empty charge state that photoelectrons can be most readily captured. With an empty quantum dot, the binding (anti-binding) potential for an injected photoelectron (photohole) is greatest. Optical excitation of the sample was thus performed with the dot prepared as in the interval $t_3 - t_2$.

4.4 Photoelectron Trapping, Detection and Storage

Highly attenuated optical pulses at a vacuum wavelength λ=760nm from a continuous wave Ti:Sapph laser were externally modulated by a Pockels cell to create photocarriers at selected time intervals in the sample. The light was delivered to the sample via a single mode optical fiber coupled to bulk optics at the tail of the ^{3}He cryogenic probe. The pulses were focused to a spot size of about 100μm diameter on the sample (at room temperature). The Al mask blocked most of the incident photon flux - as evidenced by the near absence of 2DEG density modulation - with the exception of the are area directly above the dot, where the 200 nm radius pinhole aperture was etched. Assuming a Gaussian profile for the incident spot over the illumination area of radius 50μm, and given the 200nm radius of the electrostatic dot, the photon flux impinging on the dot is a fraction $\sim 10^{-5}$ of the total incident flux.

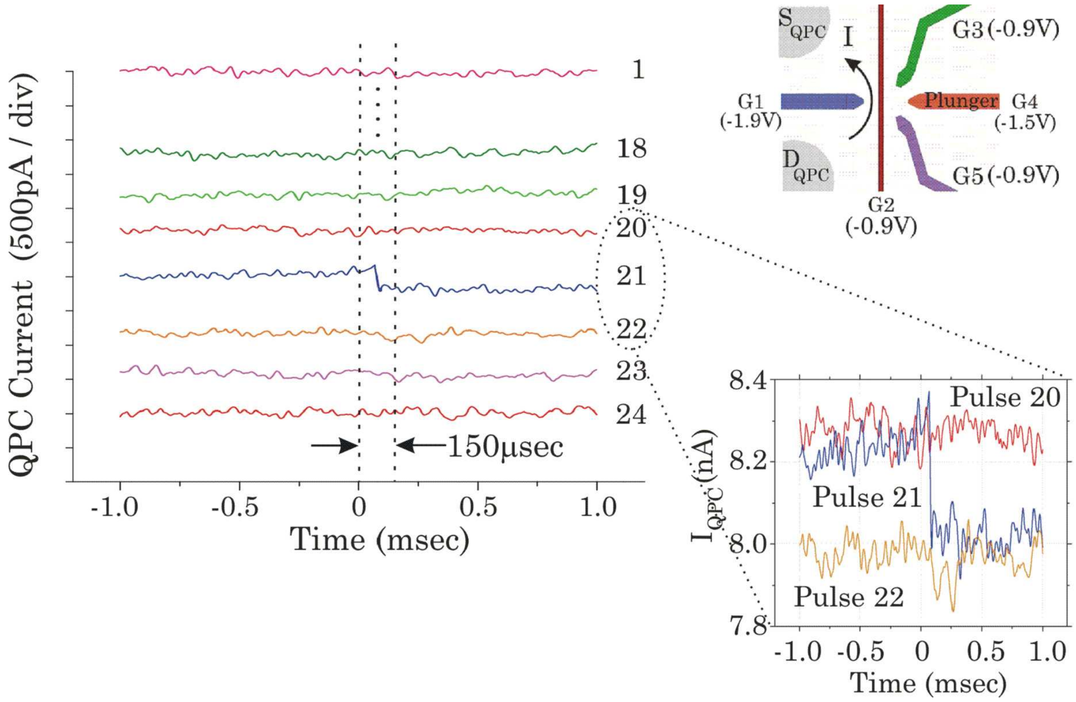

Figure 4.5: (a) Photoelectron trapping in the quantum dot detected by adjacent point contact transistor. The dot was emptied of charge prior to exposure to λ=760nm optical pulses, at a flux of 0.1 photons/pulse through the aperture within a 150μs time window. The QPC current versus time traces are depicted with a vertical offset for clarity. (b) An expanded view of QPC current for pulses 20, 21, and 22 without offset. The charge sensitivity is $\sim 10^{-3}e/\sqrt{\text{Hz}}$.

The QPC current versus time is plotted in Fig. 4.5, illustrating the device response to a series of consecutive optical pulses (with the dot emptied prior to the first exposure). The data of Fig. 4.5 was obtained with a mean incident photon flux of 0.1 photons/pulse through the aperture. Time $t = 0$ marks the time at which the Pockels cell was opened for a duration of 150μs. Photon absorption resulted in electron-hole pair creation, after which the quantum dot potential well served to attract the photoelectron and repel the photohole. Capture of a photoelectron was detected during pulse sequence 21 as a persistent negative photoconductivity step in the QPC current during the optical exposure window. The persistent photoconductivity change of several percent is consistent with the electrical characterization of Fig. 4.3.

Upon emptying the dot with the plunger gate G4, persistent negative photoconductivity was not observed if any one of the gates G3-G5 were not biased negatively to form a quantum dot. We thus rule out the possibility of photoelectron capture in traps other than that formed by negatively biased gate electrodes. The fall time associated with the single electron signal is 20μs, from Fig. 4.5(b). The signal-to-noise ratio in Fig. 4(b) corresponds to a QPC sensitivity to quantum dot charge of about $10^{-3}e/\sqrt{\mathrm{Hz}}$.

Increasing the photon flux through the aperture increases the frequency of occurrence of negative photoconductivity steps. A series of QPC current versus time traces with an increased photon flux of 1.2 photons/pulse through the aperture is illustrated in Figure 4.6. Based on the frequency of occurrence of photodetection events, we estimate the photoelectron trapping quantum efficiency to be approximately 10%, in fair agreement with the ratio of quantum dot volume ($\sim 0.03\mu\mathrm{m}^3$) to effective absorption volume ($\sim 0.5\mu\mathrm{m}^3$ at 4.2K, $\lambda = 760$nm).

Interspersed among the negative photoconductivity steps, positive photocon-

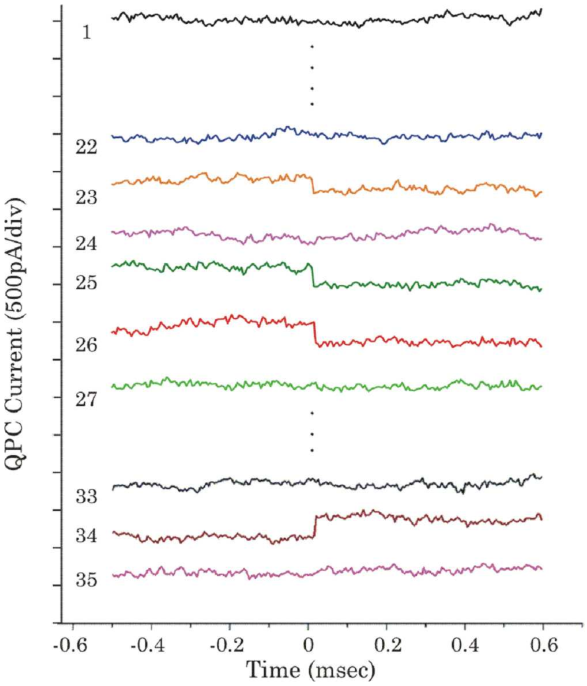

Figure 4.6: An optical pulse series with an average flux of 1.2 photons/pulse within the dot area. Occasional positive steps, of which one is observed here, could be attributed to photohole trapping at a Si donor related defect or the photoionization of the gate defined quantum dot.

ductivity steps were occasionally seen, as in the 34th pulse in Fig. 4.6. Such positive steps occurred even in the absence of a negative voltage bias on gates G3-G5. These persistent positive photoconductivity steps could be attributed to the photohole trapping at Si donor related related defects in close proximity to the QPC, as observed with a bare QPC exposed to light [80]. It is also possible that the persistent positive photoconductivity steps were due to photoionization of the gate electrode defined quantum dot.

4.5 Conclusion

In conclusion, we have demonstrated single photoelectron trapping, storage and detection in a gate electrode defined quantum dot integrated adjacently with a quantum point contact. The photoelectrons were generated by interband transitions, so that the experimental technique demonstrated could be extended to other III-V materials whose bandgaps match photon energies in the low-loss windows of conventional optical fiber ($\lambda \approx 1300$nm or 1500nm). The successful trapping and detection of photoelectrons reported here, in spite of the usually dominant photohole trapping, is a significant step towards the development of spin coherent photodetectors required for the long distance teleportation of "stationary" electron spin qubits via "flying" photon spin qubits.

CHAPTER 5

Towards Spin Coherent Photodetection

We briefly summarize experimental efforts towards the demonstration of single-shot transfer of spin states from weak photon pulses to photoelectrons trapped and stored within gate defined quantum dots in GaAs/AlGaAs heterostructures.

An intermediate step between the trapping, storage and detection of photoelectrons [39] and the single shot quantum coherent transfer of spin from photon to electron [37] is the demonstration of single shot classical transfer of spin from photon to electron. The term “classical” is used here to indicate that measurement of spin is permitted during the photodetection process so that only classical bits can be transferred rather than qubits. Specifically, the photohole is allowed to become entangled with the photoelectron so that the photocarrier pair creation process is,

$$\alpha|\sigma_+\rangle + \beta|\sigma_-\rangle \rightarrow \alpha|+\rangle_{\mathrm{e}}|+\rangle_{\mathrm{h}} + \beta|-\rangle_{\mathrm{e}}|-\rangle_{\mathrm{h}} \tag{5.1}$$

rather than the fully quantum coherent transfer,

$$\alpha|\sigma_+\rangle + \beta|\sigma_-\rangle \rightarrow \left(\alpha|+\rangle + \beta|-\rangle\right)_{\mathrm{e}} |\mathrm{arb}\rangle_{\mathrm{h}} \tag{5.2}$$

that is required for quantum communication.

This intermediate goal was chosen primarily for two reasons; firstly, InGaAsP heterostructures required to implement the selection rules permitting quantum communication [38] have yet to become readily available of sufficient quality for gate electrode confinement of single electrons. Secondly, the GaAs/AlGaAs

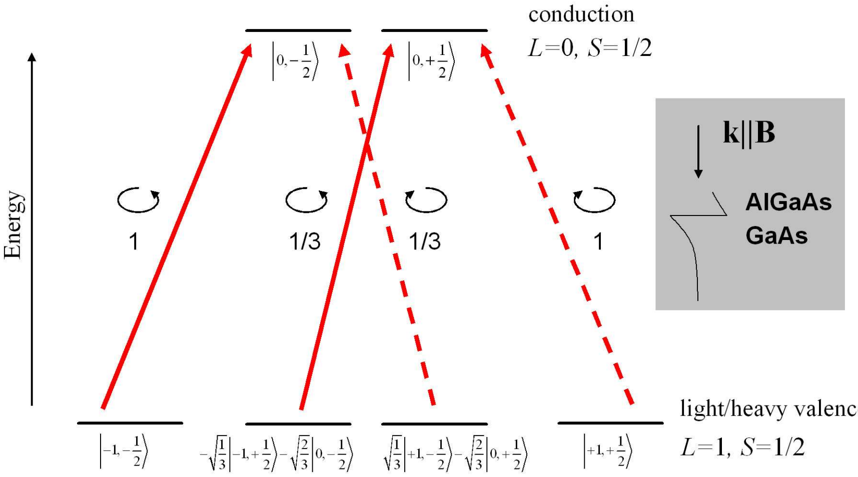

Figure 5.1: The optical selection rules for incident photons creating electron-hole pairs at band edge (Γ point). Without resolving the Zeeman splitting of conduction band or valence band states, a 75% fidelity in spin transfer from photon to electron and hole can be expected in the common quantization axis defined by magnetic field, optical wavevector and Stark splitting field of the heterostructure.

heterostructures already used for photoelectron detection [39] and single shot spin measurement [32]. Classical spin transfer with 75% fidelity, as illustrated in Fig. 5.1, is in fact possible in GaAs/AlGaAs heterostructures.

The experimental demonstration of single shot transfer of spin can be implemented by the combination of single photoelectron detection and single shot spin measurement. The experimental protocol is illustrated schematically in Fig. 5.2, and is based largely upon the single shot spin measurement protocol of Elzerman [32]. Briefly, a photoelectron is trapped by the quantum dot in a spin eigenstate (determined by an applied magnetic field), and detected by the QPC. Following adjustment of the quantum dot potential so that the Zeeman split levels straddle the Fermi level of an adjacent reservoir, electron escape and recapture in the quantum dot will be spin dependent. Monitoring the quantum dot charge state with the QPC will reveal the change in quantum dot charge state. In effect, Pauli exclusion can be used to convert the measurement of spin (magnetic dipole moment) to a much easier measurement of charge (electric monopole moment), as first demonstrated by Xiao *et al.* [89].

Minimal experimental requirements for the protocol described is as follows:

- the Zeeman energy splitting $g\mu B$ must exceed the thermal broadening k_BT of the electron energies in the reservoirs; for GaAs/AlGaAs structures we thus require $T < 2.4$K at a typical laboratory magnetic field of 8T
- the transition rate $\Gamma_{0\leftrightarrow 1}$ between the empty and singly occupied quantum dot must exceed the spin relaxation rate $1/T_1$, where T_1 is the longitudinal spin relaxation rate; for GaAs/AlGaAs heterostructures T_1 ~0.8ms [32]
- the QPC must be able to measure quantum dot charge at a rate exceeding the transition rate $\Gamma_{0\leftrightarrow 1} > 1/T_1$, implying a minimum charge sensitivity

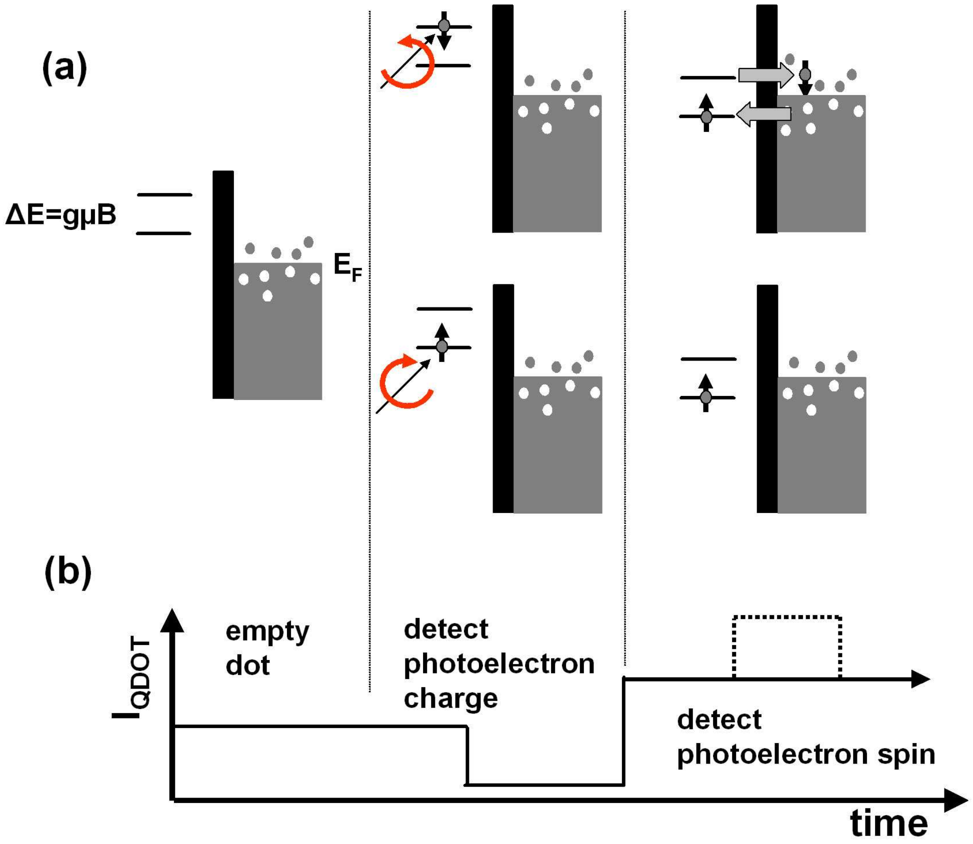

Figure 5.2: (a) the proposed protocol for verification of single shot spin transfer from photon spin to electron spin, along with (b) the expected QPC current. An initially empty dot traps an injected photoelectron in one of two Zeeman split ground states, and the photoelectron charge is detected by the QPC. The Zeeman split ground state energy levels are tuned via a gate electrode bias so as to straddle the Fermi level of the adjacent reservoir. With a sufficiently transparent tunneling barrier, a photoelectron in the higher energy spin state will tunnel out of the quantum dot and will be replaced by a lower energy spin state. The transient change in quantum dot charge state (dotted line) can be detected by the QPC from which the spin can be inferred. The entire process must take place within a T_1 spin flip lifetime.

$\delta q < e/\sqrt{\Gamma_{0\leftrightarrow 1}} < e\sqrt{T_1}$; for GaAs/AlGaAs heterostructures the charge sensitivity must then be $0.03e/\sqrt{\mathrm{Hz}}$

In the following, we describe the experimental progress towards realizing these criteria for single shot spin transfer.

5.1 Device Structure

The device employed for realizing single shot classical spin transfer is illustrated in Fig. 5.3. The heterolayers are: a 10 nm i-GaAs cap layer, a 50 nm Si-doped i-$\mathrm{Al_{0.3}Ga_{0.7}As}$ layer, a 20 nm Si doped ($1\times10^{18}/\mathrm{cm}^3$) n-$\mathrm{Al_{0.3}Ga_{0.7}As}$ layer, a 15 nm i-$\mathrm{Al_{0.3}Ga_{0.7}As}$ spacer layer, atop an undoped GaAs buffer of 500 nm thickness. Atop the GaAs/AlGaAs heterostructure are gate electrodes defined by electron beam lithography and electron beam evaporation. When biased at negative voltages, the gate electrodes define a quantum dot adjacent to a quantum point contact (QPC) in a geometry previously used for the single shot measurement of electron spins [32]. To suppress photohole trapping and the resultant modulation of 2DEG density, an Al mask layer was deposited atop the entire structure with an aperture etched out over the quantum dot region. In contrast with the device previously employed for photoelectron trapping, detection and storage **??**, the insulating layer is 300 monolayers of atomic layer deposited (ALD) $\mathrm{Al_2O_3}$. The advantages of ALD alumina over electron beam evaporated $\mathrm{SiO_x}$ are: increased adhesion to the GaAs substrate and a highly reproducible 5% modulation of 2DEG density in contrast with the tendency of electron beam evaporated $\mathrm{SiO_x}$ to deplete the 2DEG entirely (resulting in poor reproducibility and low device yield).

The gate electrode geometry of the device illustrated in Fig. 5.3 differs from

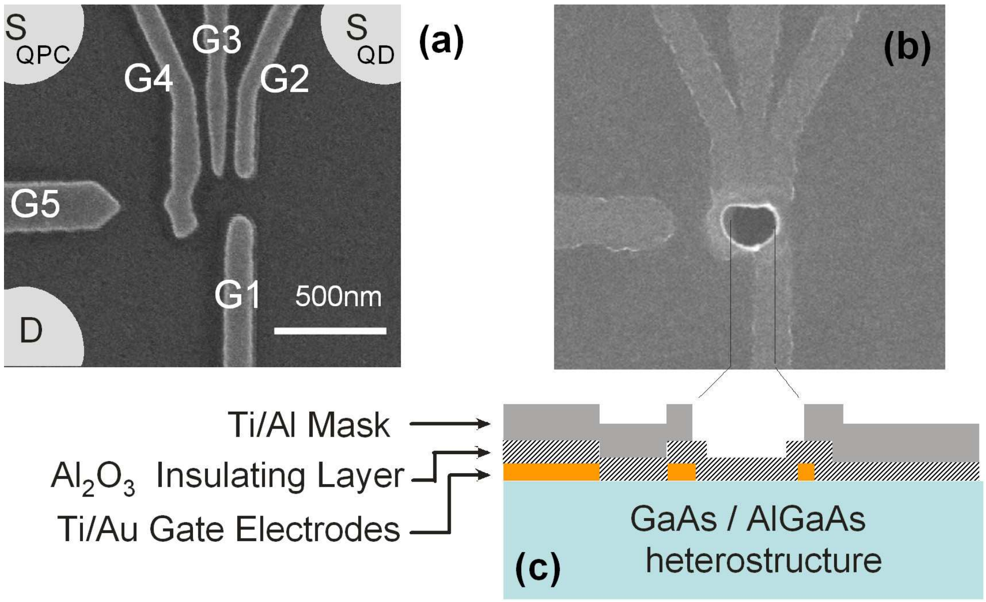

Figure 5.3: (a) Scanning electron graph (SEG) of the surface metallic gates defining a quantum point contact between the source and drain Ohmic contacts. (b) SEG of pinhole aperture etched in an opaque Al/Ti layer, 150 nm thick, acting as a shadow mask to allow illumination of the quantum dot region only. (c) Cross–sectional view of the device structure showing gates buried under Al/Ti/Al_2O_3 layers.

that described in Chapter 4 in two ways: the quantum dot is smaller with a diameter of $\sim$ 220nm as opposed to a $\sim$ 400nm diameter and the tunnel barriers are shorter ($\sim$ 80nm lithographic length). The purpose of these changes in geometry are to allow greater tuneability of the quantum mechanical tunneling of the last electron on the quantum dot and reservoir. Recall that the quantum dot described in Chapter 4 could only be brought into single electron occupation when the tunneling barriers were opaque and the dot existed in a metastable state. This was the result of a large lithographic area of the quantum dot and relatively long tunneling barriers (since tunneling is suppressed exponentially with distance, long barriers behave more "classically").

5.2 Preliminary Results

The sample described above was mounted in an optical microscope on the cold finger of a ^{3}He cryostat insert (Janis Research Company, HE-3-SSXGAS) in a dewar with a superconducting magnet ($\sim$ 8T peak field) as illustrated in Fig. 5.4. The base temperature of all experiments reported herein are 290mK-310mK. Gate electrodes were biased with tuneable voltage sources while source-drain currents were measured with a low noise transimpedance amplifier (FEMTO LCA-400K-10M). A polarization maintaining fiber brought light in from a laser diode source that could be directly modulated.

The dependence of the number of electrons in the quantum dot potential well on the applied gate voltages was mapped out using the QPC. By measuring the differential transconductance $dI_{QPC}/dVG3$ of the QPC with respect to the plunger gate voltage VG3, a stability diagram can be generated indicating the voltage biases at which the quantum dot holds a well defined number of electrons [82, 84]. Measurements of the differential transconductance $dI_{QPC}/dVG3$ as a

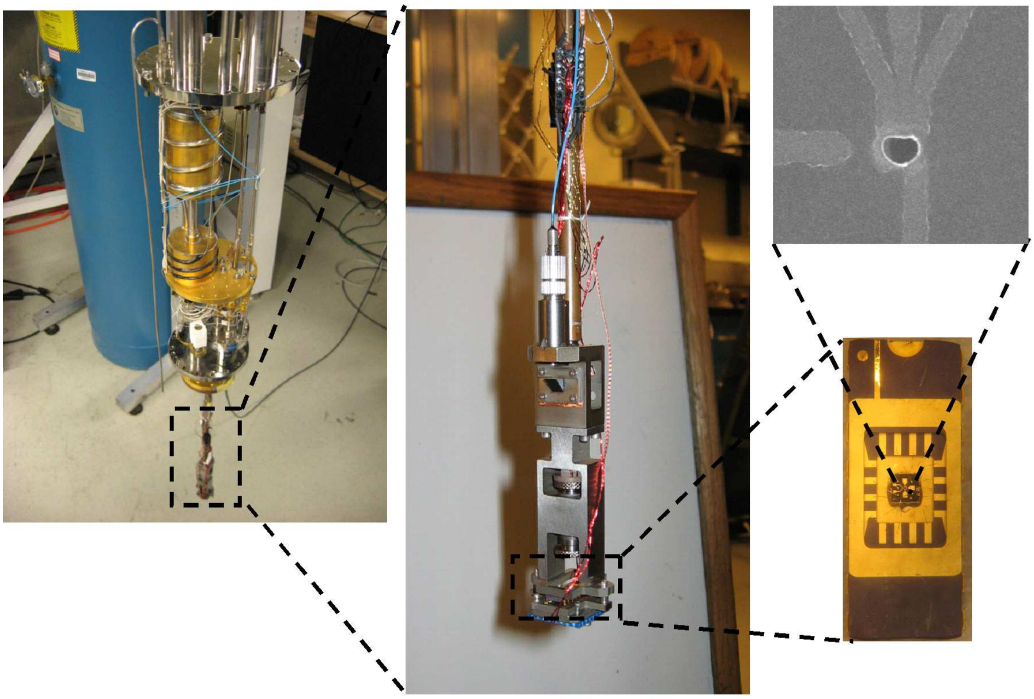

Figure 5.4: The AlGaAs/GaAs device was thermally anchored to a chip carrier that was mounted in an optical microscope on the tail of a ^{3}He cryostat. A polarization maintaining optical fiber was used for delivering light into the microscope from a room temperature source.

function of the bias voltages on G1 and G3 are illustrated in grey scale in Fig. 5.5 under (a) zero magnetic field and (b) high magnetic field $B = 7.55$T normal to sample surface, corresponding to a $\nu = 3/2$ quantum Hall filling factor in the bulk 2DEG.

Regions of nearly constant conductance correspond to a well defined number of electrons on the quantum dot. These are regions of Coulomb blockade, where it is energetically unfavourable for the quantum dot to give up or take up an electron along with its discrete electrostatic addition energy $e^2/2C$. The regions of Coulomb blockade are separated by sharp dips in the conductance, corresponding to bias voltages at which the quantum dot can give up or take up an electron with very small changes in bias. In effect, an electron can enter or exit the dot to screen small gate potential changes. This screening can be detected as a dip in QPC transconductance. One can thus count the number of electrons in the quantum dot.

The stability diagram for a single quantum dot is expected to consist of regions bounded by straight lines, corresponding to a fixed capacitance between quantum dot and gate electrode. The curving nature of the observed lines indicates a capacitance change as a function of gate bias, or in other words, a redistribution of charge as a function of gate voltage. Both the presence of local defects in the semiconductor [90] and tight magnetic confinement at high fields [91, 92] can lead to charge redistribution amongst multiple local minima in the effective potential. Both effects are seen to be present in Fig. 5.5, but the empty and singly occupied quantum dot states are still clearly visible.

The stability diagram also gives information on the electron tunneling rate. The transconductance $dI_{QPC}/dVG3$ in Fig. 5.5 was measured with VG3 modulated at a frequency of 3.381kHz. If the electron tunneling rate between quantum

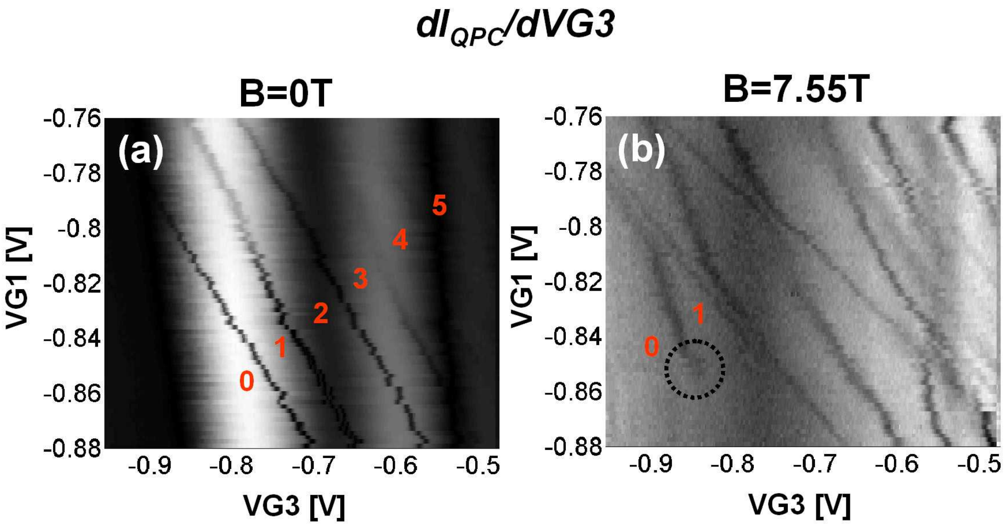

Figure 5.5: The differential conductance $dI_{QPC}/dVG3$ is plotted in (a) zero magnetic field and (b) $B = 7.55$T magnetic field corresponding to quantum Hall filling factor $\nu = 3/2$. White indicates high conductance and black indicates low conductance. A QPC bias of $I_{QPC} = 10$nA and $V_{SD} = 750\mu$V was used. The AC excitation on VG3 was 1.6mV rms at a frequency of 3.381kHz and observed with lock-in amplifier with a 30ms integration time. The broad bands in (a) are resonances in the QPC due to the local dopant/defect potential. Electron number on the quantum dot is indicated in red. The circled region in (b) is the gate voltage bias condition in which random telegraphing was observed as the charge on the quantum dot fluctuated by one electron.

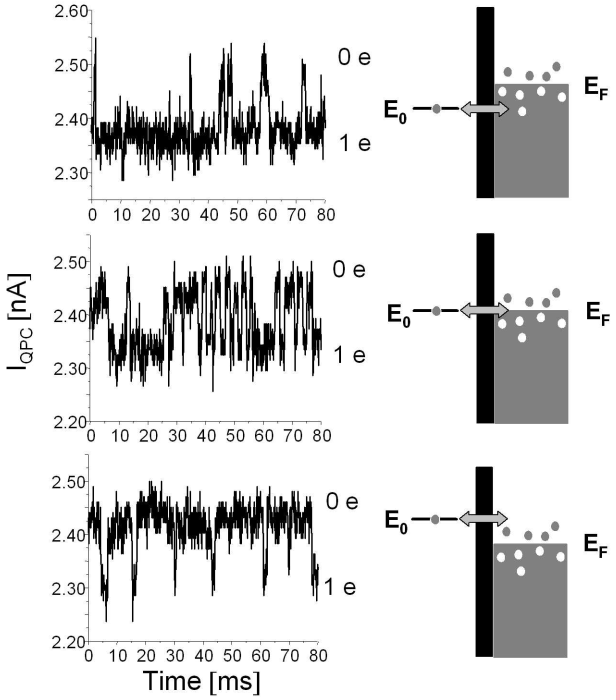

Figure 5.6: Discrete changes in quantum dot charge state are detected by the QPC at bias points indicated by the circled region in Fig. 5.5(b). Varying the potential on VG3 varies the mean occupation of the quantum dot from 0 electrons to 1 electrons. The charge sensitivity here is $0.006e/\sqrt{\text{Hz}}$.

dot and reservoir is below the modulation frequency, screening of the modulated gate potential is suppressed and the transconductance dip will disappear as indicated by the dotted circle in Fig. 5.5(b). With the gate electrodes biased in this region of moderate tunneling rate, the fluctuation in electron number on the quantum dot can be measured directly by the QPC. The quantum dot charge fluctuation manifests itself as a random telegraph signal in the QPC current as shown in Fig. 5.6. The sensitivity of the QPC current to quantum dot charge is $0.006e/\sqrt{\text{Hz}}$, corresponding to a single electron being resolved in 36μs, well below the $T_1 = 800\mu$s spin-flip lifetime of GaAs.

In addition to the high speed measurement of quantum dot charge and the control of tunneling rate through the quantum dot, single shot spin measurement also requires the precise control of the quantum dot energy levels with respect to the Fermi level of the adjacent reservoirs. The tuning of quantum dot energy levels with respect to the reservoir Fermi level can be directly probed by measuring current flow through the quantum dot. The typical differential conductance $dI_{\mathrm{QPC}}/dV_{\mathrm{SDQPC}}$ measured under zero applied magnetic field is plotted in Fig. 5.7. Bias regions of current flow suppressed by Coulomb blockade ideally take the form of diamonds, in which the number of electrons within the quantum dot is well defined. The Coulomb diamonds are bounded by differential conductance peaks corresponding to electron tunneling through the quantum dot ground state. Differential conductance peaks at source drain biases beyond the ground state tunneling peaks correspond to electron tunneling through excited states.

From the typical Coulomb diamonds of Fig. 5.7(b), the width of the tunneling peaks through the quantum dot is a rather large 300μeV, corresponding to an effective electron temperature of 3.4K. This magnitude of broadening presently prohibits resolution of Zeeman splitting $\Delta E = g\mu B = 200\mu$eV in the quantum dot ground state at readily available magnetic fields of $B = 8$T. The fluctuation in reservoir electron energy is almost certainly the result of "technical" noise: capacitive transduction of acoustic noise from cryostat capillary flow lines and inadequate electrical filtering of high frequency noise (10GHz - 1THz). Improved filtering and shielding developed for exquisite low temperature measurements [93] are expected to overcome these "technical" limitations.

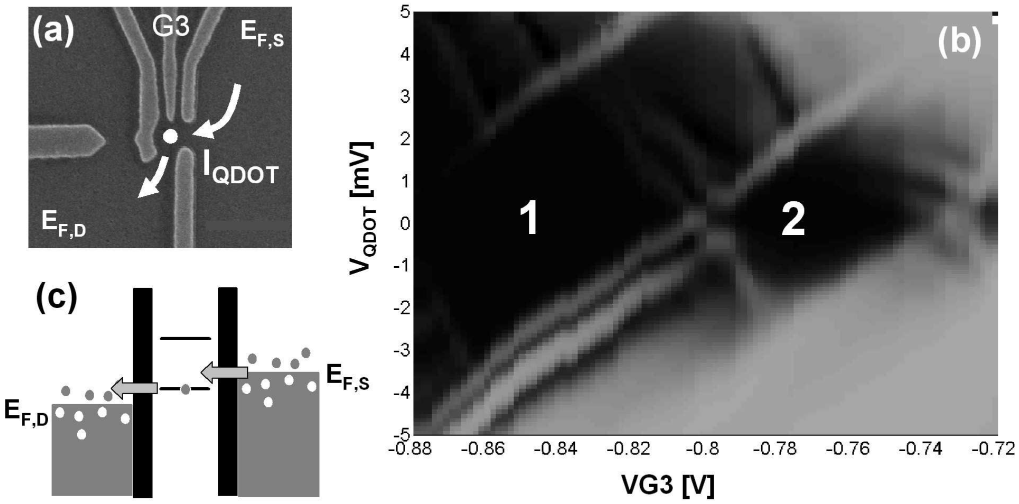

Figure 5.7: (a) The current I_{QDOT} through the quantum dot was measured under varying bias conditions. A typical grey scale plot (b) of the measured differential conductance dI_{QPC}/dV_{QDOT}, where $V_{QDOT} = (E_{FS} - E_{FD})/e$ is the potential bias across the dot, reveals diamond shaped regions where Coulomb blockade suppresses current through the quantum dot. The electron occupancy of the quantum dot is labeled in Coulomb blockade regions; a charging energy of $\sim$ 3.5meV is determined from the half width of the single electron Coulomb diamond. The differential conductance is a sensitive probe of (c) the relative alignment between quantum dot energy levels and the adjacent reservoirs.

5.3 Conclusions

Significant progress has been made towards realizing single shot spin transfer, particularly in the control over charge in quantum dots, but experimental demonstration of photon to electron spin transfer awaits the resolution of two "technical" problems. Firstly, the reduction in effective electron reservoir temperature (noise) is essential. Previously reported single shot spin measurement in AlGaAs/GaAs heterostructures was performed with electron temperatures of the order of 0.29K [32], well below the measured effective temperature of ~ 3.4K reported herein.

Secondly, the tunneling rate $\Gamma_{0\leftrightarrow 1}$ of an electron on and off the quantum dot has been found to be difficult to control to the precision that would allow single shot spin measurement. The limited ability to control the tunneling rate is likely due to excess effective electron temperature with the electron reservoirs. The tunneling rate $\Gamma_{0\leftrightarrow 1}$ is a function of quantum dot energy level alignment with the Fermi level, and as such is averaged over an energy range equal to the effective electron temperature in the reservoirs. Reducing the "technical" noise that is likely producing the elevated effective electron temperature will improve control over tunneling rate.

CHAPTER 6

Conclusions

6.1 Summary

The contributions reported in this dissertation are:

- the means by which a quantum algorithm can be applied to finding eigenvalues of partial differential equations was described; it was found that domains must be of more than 4 dimensions to give any computational advantage over existing classical schemes for electromagnetics eigenproblems
- the penalty in the error correction threshold for fault tolerant quantum computation on a one-dimensional array of qubits was numerically estimated; it was found that error correction circuits contain sufficient complexity that additional communication operations only have a moderate impact on error thresholds even in one dimension
- the trapping, storage, and detection of individual photoelectrons was experimentally demonstrated, which is an important step towards realizing spin coherent photodetection for quantum communication applications

In addition, recent progress towards the demonstration of single shot classical spin transfer from photons to photoelectrons has been detailed. This is the next logical step towards realizing spin coherent photodetection.

6.2 Future Work

Open questions for future research raised by the work reported in this dissertation are briefly discussed.

Perhaps the most immediate theoretical problem is in the advancement of quantum computational architectures, where there are primarily two concerns. The first and foremost is the error threshold. The experimental difficulty in controlling analog quantities such as spin orientation (equivalently qubit location on the Bloch sphere) to fractions of a degree necessitates a search for a fault tolerant scheme with the highest tolerance for errors in spin manipulation. Of course, this can not come at an unreasonable cost in the number of qubits (and the resulting tolerance requirements on communication that will ensue). The second concern is that of quantum circuit latency or depth. Although concatenation is known to be effective at suppressing error, the self-similar structure of concatenation leads to exponential growth of latency with concatenation level. Thus, circuit optimization for minimum error rate and minimum depth is a concern for all proposed implementations of quantum computational hardware.

Experimental implementation of quantum hardware in semiconductors requires a number of problems to be resolved as well. Most importantly is the development of valence IV materials such as the SiGe alloy system to produce reliable gate electrode confinement of electrons. The primary reason is the reduced nuclear spin density, and thus a reduced magnitude of the fluctuating effective field from nuclear spins (Overhauser shift) and enhanced lifetime of $T_2 = 60$ms [69] in isotopically purified valence IV material compared to a $T_2 \sim 10$ns lifetime in GaAs [94](in which there are no spin-free isotopes and a large Overhauser shift ensues). There has already been recent progress in the development of SiGe quantum dot technology [95]. Secondly, the stability of modulation doped heterostructures poses a concern. It is well known that the AlGaAs/GaAs heterostructures used for most work to date are often times eratic in background charge noise [96, 97]. While there has been some experimental effort towards reducing background charge fluctuation [98], the development of material systems based strictly on surface charge rather than partially ionized dopant layers is conjectured here as a possible means to reduce background charge noise. This would improve charge sensitivity (limited by charge noise rather than intrinsic bandwidth) and would also reduce error in coherent charge manipulation.

References

[1] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proc. 35th Annual Symposium on Foundations of Computer Science*, 1994, pp. 124–134.

[2] E. W. Weisstein, "Number field sieve." [Online]. Available: http://mathworld.wolfram.com/NumberFieldSieve.html

[3] A. Ekert and R. Josza, "Quantum computation and Shor's factoring algorithm," *Rev. Mod. Phys.*, vol. 68, pp. 733–753, 1996.

[4] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM J. Comput.*, vol. 26, pp. 1484–1509, 1997.

[5] D. Bouwmeester, A. K. Ekert, and A. Zeilinger, Eds., *The physics of quantum information: quantum cryptography, quantum teleportation and quantum computation.* Springer, 1999.

[6] V. Vedral, A. Barenco, and A. Ekert, "Quantum networks for elementary arithmetic operations," *Phys. Rev. A.*, vol. 54, pp. 147–153, 1996.

[7] P. Gossett, "Quantum carry-save arithmetic." [Online]. Available: xxx.lanl.gov/abs/quant-ph/9808061

[8] R. Cleve and J. Watrous, "Fast parallel circuits for quantum fourier transform," in *Proc. 41st Annual Symposium on Foundations of Computer Science*, 2000, pp. 526–536.

[9] S. Lloyd, "Universal quantum simulators," *Science*, vol. 273, pp. 1073–1078, 1996.

[10] D. S. Abrams and S. Lloyd, "Simulation of many-body Fermi systems on a universal quantum computer," *Phys. Rev. Lett.*, vol. 79, pp. 2586–2589, 1997.

[11] ——, "Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors," *Phys. Rev. Lett.*, vol. 83, pp. 5162–5165, 1999.

[12] L. K. Grover, "Quantum mechanics helps in searching for a needle in a haystack," *Phys. Rev. A*, vol. 79, pp. 325–328, 1997.

[13] D. Deutsch and R. Jozsa, "Rapid solution of problems by quantum computation," *Proc. R. Soc. Lond. A*, vol. 439, pp. 553–558, 1992.

[14] R. Cleve, A. Ekert, C. Machiavello, and M. Mosca, “Quantum algorithms revisited,” *Proc. R. Soc. Lond. A*, vol. 454, pp. 339–354, 1998.

[15] W. Zurek, “Decoherence, einselection, and the quantum origins of the classical,” *Rev. Mod. Phys.*, vol. 75, pp. 715–775, 2003.

[16] P. W. Shor, “Fault-tolerant quantum computation,” in *Proc. 37th Annual Symposium on Foundations of Computer Science*, 1996, pp. 56–65.

[17] M. Freedman, A. Kitaev, M. Larsen, and Z. Wang, “Topological quantum computation.” [Online]. Available: xxx.lanl.gov/abs/quant-ph/0101025

[18] A. R. Calderbank and P. W. Shor, “Good quantum error-correcting codes exist,” *Phys. Rev. A.*, vol. 54, pp. 1098–105, 1996.

[19] A. M. Steane, “Error correcting codes in quantum theory,” *Phys. Rev. Lett.*, vol. 77, pp. 793–7, 1996.

[20] J. Preskill, “Reliable quantum computers,” *Proc. R. Soc. Lond. A*, vol. 454, pp. 385–410, 1998.

[21] B. E. Kane, “A silicon-based nuclear spin quantum computer,” *Nature*, vol. 393, pp. 133–7, 1998.

[22] D. Loss and D. P. DiVincenzo, “Quantum computation with quantum dots,” *Phys. Rev. A*, vol. 57, pp. 120–6, 1998.

[23] R. Vrijen, E. Yablonovitch, K. Wang, H. W. Jiang, A. Balandin, V. Roychowdhury, T. Mor, and D. DiVincenzo, “Electron spin resonance transistors for quantum computing in silicon-germanium heterostructures,” *Phys. Rev. A*, vol. 62, p. 12306, 2000.

[24] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R. S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, “Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics,” *Nature*, vol. 431, p. 162, 2004.

[25] D. Gottesman, “Fault-tolerant quantum computation with local gates,” *J. of Mod. Optics*, vol. 47, p. 333, 2000.

[26] K. M. Svore, B. M. Terhal, and D. P. DiVincenzo, “Local fault-tolerant quantum computation,” *Phys. Rev. A*, vol. 72, p. 022317, 2005.

[27] K. M. Svore, D. P. DiVincenzo, and B. M. Terhal, “Noise threshold for a fault-tolerant two-dimensional lattice architecture.” [Online]. Available: xxx.lanl.gov/abs/quant-ph/0604090

[28] G. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, pp. 114–117, 1965.

[29] D. P. DiVincenzo, "The physical implementation of quantum computation," *Fortschritte der Physik*, vol. 48, pp. 771–783, 2000.

[30] T. Szkopek, H. Robinson, and E. Yablonovitch, "Dynamic quantum infidelity by interaction with the environment," *unpublished*, 2002.

[31] E. Yablonovitch, H. W. Jiang, H. Kosaka, H. D. Robinson, D. S. Rao, and T. Szkopek, "Optoelectronic quantum telecommunications based on spins in semiconductors," *Proc. of the IEEE*, pp. 761–780, 2003.

[32] J. M. Elzerman, R. Hanson, L. H. van Beveren, B. Witkamp, L. M. K. Vandersypen, and L. P. Kouwenhoven, "Single-shot read-out of an individual electron spin in a quantum dot," *Nature*, vol. 430, pp. 431–435, 2004.

[33] M. Xiao, I. Martin, E. Yablonovitch, and H. W. Jiang, "Electrical detection of the spin resonance of a single electron in a silicon field-effect transistor," *Nature*, vol. 430, pp. 435–439, 2004.

[34] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard, "Coherent manipulation of coupled electron spins in semiconductor quantum dots," *Science*, vol. 309, pp. 2180–2184, 2005.

[35] S. D. Barrett and P. Kok, "Efficient high-fidelity quantum computation using matter qubits and linear optics," *Phys. Rev. A*, vol. 71, p. 060310(R), 2005.

[36] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, "Quantum cryptography," *Rev. Mod. Phys.*, vol. 74, pp. 145–195, 2002.

[37] R. Vrijen and E. Yablonovitch, "A spin-coherent semiconductor photodetector for quantum communication," *Physica E*, vol. 10, pp. 569–575, 2001.

[38] A. A. Kiselev, K. W. Kim, and E. Yablonovitch, "Designing a heterostructure for the quantum receiver," *Appl. Phys. Lett.*, vol. 80, pp. 2857–2859, 2002.

[39] D. S. Rao, T. Szkopek, H. D. Robinson, E. Yablonovitch, and H. Jiang, "Single photoelectron trapping, storage, and detection in a one-electron quantum dot," *J. Appl. Phys.*, vol. 98, p. 114507, 2005.

[40] R. P. Feynman, *The Feynman lectures on computation.* Perseus Publishing, 2000.

[41] B. Boghosian and W. Taylor, “Simulating quantum mechanics on a quantum computer.” [Online]. Available: xxx.lanl.gov/abs/quant-ph/9701019

[42] C. Zalka, “Simulating quantum systems on a quantum computer,” *Proc. Roy. Soc. London A.*, vol. 454, pp. 313–322, 1998.

[43] D. A. Meyer, “Quantum computing classical physics.” [Online]. Available: xxx.lanl.gov/abs/quant-ph/0111069

[44] A. Brandt, S. McCormick, and J. Ruge, “Multi-grid methods for differential eigenproblems,” *SIAM J. Sci. Stat. Comput.*, vol. 4, pp. 244–260, 1983.

[45] J. W. Demmel, *Applied Numerical Linear Algebra.* SIAM, 1997.

[46] Y. Saad, *Numerical Methods for Large Eigenvalue Problems.* Halsted Press, 1992.

[47] D. A. Lidar and O. Biham, “Simulating Ising spin glasses on a quantum computer,” *Phys. Rev. E*, vol. 56, pp. 3661–3681, 1997.

[48] A. Papageorgiou and H. Woźniakowski, “Classical and quantum complexity of the Sturm-Liouville eigenvalue problem,” *Quantum Inf. Process.*, vol. 4, pp. 87–127, 2005.

[49] P. Jaksch and A. Papageorgiou, “Eigenvector approximation leading to an exponential speedup of quantum eigenvalue calculation,” *Phys. Rev. Lett.*, vol. 91, p. 257902, 2003.

[50] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information.* Cambridge Univ. Press, 2000, appendix 3.

[51] H. Yoshida, “Construction of higher-order symplectic integrators,” *Phys. Lett. A*, vol. 150, pp. 262–268, 1990.

[52] N. Hatano and M. Suzuki, “Finding exponential product formulas of higher orders.” [Online]. Available: xxx.lanl.gov/abs/quant-ph/0506007

[53] A. Sornborger and E. Stewart, “Higher-order methods for simulations on quantum computers,” *Phys. Rev. A*, vol. 60, pp. 1956–1965, 1999.

[54] D. Gottesman, “Theory of fault-tolerant quantum computation,” *Phys. Rev. A*, vol. 57, pp. 127–137, 1998.

[55] J. I. Cirac and P. Zoller, “Quantum computations with cold trapped ions,” *Phys. Rev. Lett.*, vol. 74, pp. 4091–4, 1995.

[56] D. Kielpinski, C. R. Monroe, and D. J. Wineland, “Architecture for a large-scale ion-trap quantum computer,” *Nature*, vol. 417, pp. 709–11, 2002.

[57] J. I. Cirac, P. Zoller, H. Kimble, and H. Mabuchi, “Quantum state transfer and entanglement distribution among distant nodes in a quantum network,” *Phys. Rev. Lett.*, vol. 78, pp. 3221–4, 1997.

[58] A. Imamoglu, D. D. Awschalom, G. Burkard, D. P. DiVinceno, D. Loss, M. Sherwin, and A. Small, “Quantum information processing using quantum dot spins and cavity QED,” *Phys. Rev. Lett.*, vol. 83, p. 4204, 1999.

[59] G. Burkard, D. Loss, and E. V. Sukhorukov, “Noise of entangled electrons: Bunching and antibunching,” *Phys. Rev. B.*, vol. 61, p. R16303, 2000.

[60] P. Recher, E. V. Sukhorukov, and D. Loss, “Quantum dot as spin filter and spin memory,” *Phys. Rev. Lett.*, vol. 85, p. 1962, 2000.

[61] D. Aharanov and M. Ben-Or, “Fault-tolerant quantum computation with constant error rate.” [Online]. Available: xxx.lanl.gov/abs/quant-ph/9906129

[62] D. Copsey, M. Oskin, F. Impens, T. Metodiev, A. Cross, F. T. Chong, I. L. Chuang, and J. Kubiatowicz, “Toward a scalable, silicon-based quantum computing architecture,” *IEEE J. Select. Topics Quantum Electron.*, vol. 9, no. 6, pp. 1552–69, 2003.

[63] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, “On universal fault-tolerant quantum computing: a novel basis and a new constructive proof of universality for shor’s basis,” in *Proc. 40th Annual Symposium on Foundations of Computer Science*, 1999, p. 486.

[64] A. Steane and B. Ibinson, “Fault-tolerant logical gate networks for css codes.” [Online]. Available: xxx.lanl.gov/abs/quant-ph/0311014

[65] E. Knill and R. Laflamme, “Concatenated quantum codes.” [Online]. Available: xxx.lanl.gov/abs/quant-ph/9608012

[66] A. M. Steane, “Active stabilization, quantum computation, and quantum state synthesis,” *Phys. Rev. Lett.*, vol. 78, pp. 2252–5, 1997.

[67] P. O. Boykin, T. Mor, V. Roychowdhury, and F. Vatan, “Fault tolerant computation on ensemble quantum computers,” in *Proc. of the 2004 International Conf. on Dependable Systems and Networks*, 2004, p. 157.

[68] P. O. Boykin and V. P. Roychowdhury, to appear in Proc. of the 2005 International Conf. on Dependable Systems and Networks (DSN2005). [Online]. Available: xxx.lanl.gov/abs/cs/0504010

[69] A. M. Tyryshkin, S. A. Lyon, A. V. Astashkin, and A. M. Raitsimring, "Electron spin relaxation times of phosphorus donors in silicon," *Phys. Rev. B.*, vol. 68, p. 193207, 2003.

[70] A. Sørenson and K. Mølmer, "Error-free quantum communication through noisy channels," *Phys. Rev. A.*, vol. 58, pp. 2745–9, 1998.

[71] A. G. Fowler, S. J. Devitt, and L. C. L. Hollenberg, "Implementation of shor's algorithm on a linear nearest neighbour qubit array," *Quant. Info. Comput.*, vol. 4, pp. 237–251, 2004.

[72] A. Yariv, *Optical Electronics in Modern Communications.* Oxford University Press, New York NY, 1997.

[73] K. Kandiah and F. B. Whiting, "Low frequency noise in junction field effect transistors," *Solid State Electron.*, vol. 21, pp. 1079–1088, 1978.

[74] M. J. Kirton and M. J. Uren, "Noise in solid state microstructures: a new perspective on individual defects, interface states and low-frequency (1/f) noise," *Advan. Phys.*, vol. 38, pp. 367–468, 1989.

[75] A. Rose, *Concepts in photoconductivity and allied problems.* Krieger, Huntington NY, 1978.

[76] S. Komiyama, O. Astafiev, V. Antonov, T. Kutsuwa, and H. Hirai, "A single-photon detector in the far-infrared range," *Nature*, vol. 403, pp. 405–407, 2000.

[77] P. H. Siegel, "THz technology," *IEEE Trans. MTT*, vol. 50, pp. 910–928, 2002.

[78] K. Ikushima, Y. Yoshimura, T. Hasegawa, S. Komiyama, T. Ueda, and K. Hirakawa, "Photon-counting microscopy of TeraHertz radiation," *Appl. Phys. Lett.*, vol. 88, p. 152110, 2006.

[79] A. J. Shields, M. P. O'Sullivan, I. Farrer, D. A. Ritchie, R. A. Hogg, M. L. Leadbeater, C. E. Norman, and M. Pepper, "Detection of single photons using a field-effect transistor gated by a layer of quantum dots," *Appl. Phys. Lett.*, vol. 76, pp. 3673–3675, 2000.

[80] H. Kosaka, D. S. Rao, H. D. Robinson, P. Bandaru, T. Sakamoto, and E. Yablonocitch, "Photoconductance quantization in a single-photon detector," *Phys. Rev. B*, vol. 65, p. 201307, 2002.

[81] M. Field, C. G. Smith, D. A. Ritchie, J. E. F. Frost, G. A. C. Jones, and D. G. Hasko, "Measurements of coulomb blockade with a noninvasive voltage probe," *Phys. Rev. Lett.*, vol. 70, pp. 1311–1314, 1993.

[82] M. Ciorga, A. S. Sachrajda, P. Hawrylak, C. Gould, P. Zawadzki, S. Jullian, Y. Feng, and Z. Wasilewski, "Addition spectrum of a lateral quantum dot from coulomb and spin-blockade spectrsocopy," *Phys. Rev. B*, vol. 61, pp. R16 315–R16 318, 2000.

[83] D. Sprinzak, Y. Ji, M. Heiblum, D. Mahalu, and H. Shtrikman, "Charge distribution in a kondo-correlated quantum dot," *Phys. Rev. Lett.*, vol. 88, p. 176805, 2002.

[84] J. M. Elzerman, R. Hanson, J. S. Greidanus, L. H. W. van Beveren, S. D. Franceschi, L. M. K. Vandersypen, S. Tarucha, and L. Kouwenhoven, "Few-electron quantum dot circuit with integrated charge read out," *Phys. Rev. B*, vol. 67, p. 161308, 2003.

[85] D. S. Rao, T. Szkopek, H. W. Jiang, and E. Yablonovitch, "What is the fastest speed at which a single electron can be detected?" [Online]. Available: xxx.lanl.gov/abs/quant-ph/0507224

[86] R. J. Nelson, "Long-lifetime photoconductivity effect in n-type GaAlAs," *Appl. Phys. Lett.*, vol. 31, pp. 351–353, 1977.

[87] A. Kastalsky and J. C. M. Hwang, "Study of persistent photoconductivity effect in n-type selectively doped AlGaAs/GaAs heterojunction," *Solid State Comm.*, vol. 51, pp. 317–322, 1984.

[88] J. Cooper, C. G. Smith, D. A. Ritchie, E. H. Linfield, Y. Jin, and H. Launois, "Direct observation of single-electron decay from an artificial nucleus," *Physica E*, vol. 6, pp. 457–460, 2000.

[89] M. Xiao, I. Martin, and H. W. Jiang, "Probing the spin state of a single electron trap by random telegraph signal," *Phys. Rev. Lett.*, vol. 91, p. 078301, 2003.

[90] L. Gaudreau, S. A. Studenikin, A. S. Sachrajda, P. Zawadzki, A. Kam, J. Lapointe, M. Korkusinski, and P. Hawrylak, "Stability diagram of a few-electron triple dot," *Phys. Rev. Lett.*, vol. 97, p. 036807, 2006.

[91] M. Ciorga, M. Pioro-Ladriere, P. Zawadzki, P. Hawrylak, and A. S. Sachrajda, “The break-up of a lateral quantum dot into multiple dots in high magnetic fields,” in *Proc. 26th International Conference on the Physics of Semiconductors*, 2002.

[92] T. H. Oosterkamp, J. W. Janssen, L. P. Kouwenhoven, D. G. Austing, T. Honda, and S. Tarucha, “Maximum-density droplet and the charge redistributions in quantum dots at high magnetic fields,” *Phys. Rev. Lett.*, vol. 82, pp. 2931–2933, 1999.

[93] J. M. Elzerman, R. Hanson, L. H. W. van Beveren, L. M. K. Vandersypen, and L. P. Kouwenhoven, “Semiconductor few-electron quantum dots as spin qubits,” in *Quantum computation:solid state systems*, P. Delsing, C. Granata, Y. Pashkin, B. Ruggiero, and P. Silvestnini, Eds. Kluwer Academic Plenum, 2005.

[94] J. M. Kikkawa and D. D. Awschalom, “Resonant spin amplification in n-type GaAs,” *Phys. Rev. Lett.*, vol. 80, pp. 4313–4315, 1998.

[95] M. R. Sakr, H. W. Jiang, E. Yablonovitch, and E. T. Croke, “Fabrication and characterization of electrostatic Si/SiGe quantum dots with an integrated read-out channel,” *Appl. Phys. Lett.*, vol. 87, p. 223104, 2005.

[96] J. M. Elzerman, “Electron spin and charge in semiconductor quantum dots,” Ph.D. dissertation, Technische Universiteit Delft, 2004.

[97] R. Hanson, “Electron spins in semiconductor quantum dots,” Ph.D. dissertation, Technische Universiteit Delft, 2004.

[98] M. Pioro-Ladrière, J. H. Davies, A. R. Long, A. S. Sachrajda, L. Gaudreau, P. Zawadzki, J. Lapointe, J. Gupta, Z. Wasilewski, and S. Studenikin, “Origin of switching noise in GaAs/Al$_x$Ga$_{1-x}$As lateral gated devices,” *Phys. Rev. B*, vol. 72, p. 115331, 2005.