

Characterization and Control in Large Hilbert Spaces

by

Colm A. Ryan

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics

Waterloo, Ontario, Canada, 2008

© Colm A. Ryan 2008

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Computational devices built on and exploiting quantum phenomena have the potential to revolutionize our understanding of computational complexity by being able to solve certain problems faster than the best known classical algorithms. Unfortunately, unlike the digital computers quantum information processing devices hope to replace, quantum information is fragile by nature and lacks the inherent robustness of digital logic. Indeed, for whatever we can do to control the evolution, nature can also do in some random and unknown fashion ruining the computation. This thesis explores the task of building the classical control architecture to control a large quantum system and how to go about characterizing the behaviour of the system to determine the level of control reached. Both these tasks appear to require an exponential amount of resources as the size of the system grows. The inability to efficiently control and characterize large scale quantum systems will certainly militate against their potential computational usefulness making these important problems to solve. The solutions presented in this thesis are all tested for their practical usefulness by implementing them in either liquid- or solid-state nuclear magnetic resonance.

Acknowledgements

I would like to thank my parents for their constant advice (asked for or not), for striking that balance between expecting much from me and yet letting me make my own mistakes, and for impressing the importance of education but also the need for other interests.

I certainly could not have gained all that I did in my doctoral work without the positive and exciting atmosphere created at the Institute for Quantum Computing. A huge debt goes to my supervisor Dr Raymond Laflamme. Dr Laflamme continually impressed upon this engineering graduate the importance of connecting to theory and where to find ideas for experiments. By helping build the IQC Dr Laflamme greatly enriched my graduate experience. In the same vein I wish to thank Dr David Cory for the many discussions about how to connect one's work to the broader community and his practical experimental advice.

Working alone on a project can be miserable and frustrating so I wish to thank the other students and post-docs I have been able to bounce ideas (scientific and otherwise) off: Dr Camille Negrevergne, Dr Jonathan Baugh, Dr Martin Laforest, Osama Moussa, Jeremy Chamilliard and Adam Hubbard. Although not at IQC, Dr Jonathan Hodges from the Cory group was also a wonderful source of ideas, and discussions.

It would not have been possible to implement all these experiments without the technical expertise of Mike Ditty. Not only did he always help fix the spectrometers I had broken but also was always keen to help fill the huge gaps in my chemistry knowledge.

I would never have had the interest in physics I have if it wasn't for my high-school physics and chemistry teacher Mr. Ertman who took the time and effort to teach the class as much as it was willing to learn and not simply the curriculum.

And finally and most importantly I would like to thank my wife Veronica Kos for believing that I am clever enough to do good science and letting me fight late into the night when the spectrometer is not behaving.

Contents

List of Tables	vii
List of Figures	viii
1 Introduction	1
2 Quantum control in large Hilbert Spaces	4
2.1 Motivation	4
2.2 Establishing controllability	5
2.2.1 NMR Hamiltonians	6
2.3 Quantum Optimal Control	9
2.4 Scalable Quantum Control	18
2.5 Achieving scalable optimal control	32
2.6 Future Improvements	37
3 Characterizing quantum dynamics	39
3.1 The curse of dimensionality in characterization	39
3.2 Characterizing complex quantum dynamics	40
3.3 Characterizing time suspension sequences in solid-state NMR	48
3.4 Randomized benchmarking of quantum control sequences	56
3.4.1 Single-Qubit Benchmarking	59
3.4.2 Multi-Qubit Benchmarking	69
3.4.3 Future Improvements	74
4 Creating pure states out of mixed ones: Algorithmic Cooling	76
4.1 Theory	76
4.2 Experiment	80
4.3 Future directions for algorithmic cooling	94

5 Conclusion	97
Appendices	
A Molecular Hamiltonians	98
A.1 Crotonic acid	98
A.2 TMSS	98
B GRAPE Optimal Control Code	101
B.1 Usage	102
B.2 Code Details	103
C Pulse Implementation Correction	110
D Papers with significant contribution	114
D.1 Benchmarking Quantum Control Methods on a 12-Qubit System . .	114
D.2 Low temperature probe for dynamic nuclear polarization and multiple- pulse solid-state NMR.	120
D.3 Experimental Implementation of a Quantum Walk	130
E Publications	139
Bibliography	141

List of Tables

2.1	Decomposition of pulse with pre- and post-error terms	27
2.2	Subsystem optimal control pulses compared with standard shaped pulses	36
3.1	Simulated and experimental results from three twirled maps.	55
4.1	Table of polarization results from algorithmic cooling protocol.	93

List of Figures

2.1	Discretized pulse	10
2.2	r.f. inhomogeneity profile	14
2.3	Robustness profile for crotonic acid GRAPE pulse.	15
2.4	Robustness profile for malonic acid GRAPE pulse.	16
2.5	Simple shaped pulse on crotonic acid	20
2.6	Efficient pulse decomposition.	24
2.7	Virtual 180 cancellation.	29
2.8	Example of subsystem optimal control	34
3.1	DQC1 quantum circuit for measuring fidelity decay.	42
3.2	Robustness profile of composite 180 pulses.	44
3.3	Fidelity decay curves for regular and complex environments.	46
3.4	Pulse sequence for creating a two-qubit labelled pseudo-pure state.	47
3.5	Quadrature components of hard pulse.	52
3.6	Subsection of Cory48 sequence.	53
3.7	Quantum circuit for symmetrized characterization of quantum processes.	54
3.8	Quantum circuit implementing single qubit benchmarking.	59
3.9	Single qubit fidelity decay with no r.f. selection.	61
3.10	Single qubit benchmarking with r.f. selection.	63
3.11	Signal fluctuations due to r.f. selection.	67
3.12	Distribution of measured error rates in multiqubit benchmarking simulations.	71
3.13	Multi-qubit benchmarking with r.f. selection.	74

4.1	Performance of algorithmic cooling with depolarizing errors.	81
4.2	Number of steps to saturate algorithmic cooling protocol.	82
4.3	Fitted malonic acid spectrum	89
4.4	Quantum circuit and NMR pulses for algorithmic cooling	91
4.5	Results from algorithmic cooling	92
4.6	Electron polarization versus temperature for a variety of fields . . .	95
A.1	Crotonic acid natural Hamiltonian parameters.	99
A.2	TMSS natural Hamiltonian parameters	100
C.1	Schematic of spectrometer setup for pulse correction.	111

Chapter 1

Introduction

Quantum information processing can solve certain problems, probably most importantly simulation of quantum systems, exponentially faster than the best known classical algorithms [1]. However, to realize the benefits of quantum information processing, exquisite control over the evolution of the system is necessary. Otherwise, the errors introduced by either imperfect control or unknown and unwanted interactions with the environment surrounding the system, will quickly reduce the potentially highly entangled coherent quantum state, full of superpositions of classical states, into a more mundane (and computationally less powerful) mixture of classical states. Fortunately, things are not as bad as initially might be feared, and quite surprisingly, arbitrarily good control is not needed for arbitrarily long computation. It is possible to encode the quantum information into a non-local state spread out over many qubits so that the original state can be recovered by detecting and correcting local errors. Indeed, if the error rate is below a so-called fault-tolerance threshold, then the encoding and correction procedure can be concatenated to suppress the error to arbitrarily low values with only a polynomial resource overhead [2]. Thus, when exploring how to improve coherent control of quantum systems we have a specific target to meet. Subject to the uncertainties of architecture and error model assumptions, that number is around 10^{-4} computationally relevant errors per gate. That doesn't mean that once we reach a fault-tolerant level of control our job is finished. The resource overheads near the threshold are extreme and every improvement in control reduces that overhead to a more manageable level and makes a large scale quantum information processor (QIP) more of a possibility.

Although every proposed implementation must have a way of performing one and two qubit gates, and nearly all of them should theoretically be able to achieve fault-tolerant control levels when the many engineering and materials issues have been overcome, often little attention is paid to the question of how the control

schemes scale up with the size of the QIP. A particularly glaring example comes from quantum optimal control where very high fidelity pulses, robust to many experimental imperfections, are possible with a few qubits. This method has had great success on the small toy systems currently explored in the lab including nuclear magnetic resonance (NMR), ion traps and superconducting qubits. But the scheme requires a full simulation of the quantum system on a classical computer, and so is explicitly not efficient and in practice is limited to controlling less than 10 qubits. Clearly some form of modularization and sub-system approach is needed. These control issues and techniques are discussed in Chapter 2.

In parallel with working on controlling a large QIP, we also need tools to characterize and diagnose the evolution of the QIP. Unfortunately, even characterizing the state of a quantum system, let alone its evolution, is an exponentially hard problem that requires an exponential number of experiments. There are $\mathcal{O}(2^{4n})$ number of parameters to characterize a quantum process or map for n qubits. Protocols and algorithms are needed which can coarse-grain this huge number of parameters down to a more manageable number of average values which are experimentally feasible to extract and are meaningful in terms of characterizing the level of control or some other quantity of interest: for example, the spatial distribution or correlation of the errors affecting the QIP. Three approaches to different facets of this problem are discussed in Chapter 3.

Finally, one of the useful benefits of establishing high fidelity coherent control of a quantum system is that the control can be leveraged to improve other operations such as state preparation. So for example, while an error (the particulars of the error, whether it is worst case or some average, depend on the particular assumptions of the threshold proof or simulation) rate of 10^{-4} is needed for gates, an error rate on state preparation or measurement of 10^{-2} is reasonable for fault-tolerant computation. This is particularly relevant to a NMR QIP where the highly mixed thermal state gives us very noisy state preparation. With good control, many noisy qubits can be combined to produce one less noisy qubit and with the ability to reset a qubit to some heat bath polarization (or many qubits initialized at the beginning) then the process can be repeated many times to reduce the state preparation error even further in a process known as algorithmic cooling [3]. As discussed in Chapter 4 we were able to control our system well enough to implement this protocol in the solid-state and purify one qubit above the Shannon bound - the polarization limit for closed systems.

All of the presented protocols were tested in either liquid-state or solid-state NMR QIP. Relying on the many decades of experience in controlling nuclei through radio-frequency fields, and the fortunate availability of molecules with several coupled spin- $\frac{1}{2}$ nuclei with long coherence times, high fidelity control with average error rates of less than 1% is possible over several qubits. The commercially available

control hardware and software is at a sufficient level that attention can be concentrated on the high-level control software and algorithmic ideas rather than technical details specific to a particular system. This makes NMR an ideal test-bed for ideas of quantum control and characterization.

Chapter 2

Quantum control in large Hilbert Spaces

2.1 Motivation

It is usually relatively straightforward to establish control with only one or two qubits. One qubit gates can be effected by turning on a control field and inducing Rabi oscillations. For a two qubit gate letting a coupling Hamiltonian evolve for the correct amount of time and potentially adding some one qubits gates is sufficient. Indeed in most cases the control sequences for these small systems can be written down by hand and optimized in a heuristic manner. However, as the number of qubits grows the control schemes become more complicated, particularly in the case of always on couplings which must be refocussed as in NMR QIP [4]. Even for single qubit rotations there is usually an issue of addressability: it is not possible to effect a rotation of a single qubit without having an effect on other qubits. In most situations we have a good model for the system and apparatus so these unitary errors coming from imperfect control could conceivably be simulated and corrected at a later time. However, if we are to realize the reputed power of a QIP, the effort and complexity of our classical control schemes used to drive the quantum computation must not grow exponentially with the number of qubits. Thus it is necessary to explore methods for scalable high fidelity quantum control in large Hilbert spaces.

Nuclear magnetic resonance techniques have a long and illustrious history and represent a paradigm for quantum control. Indeed much of the language from NMR has been adopted for other spin and pseudo-spin systems. The availability of long lasting coherence and a clean, well understood, system offers a test-bed for developing and benchmarking quantum control techniques. In this chapter I explore methods for quantum control based in liquid-state NMR. However, the techniques

developed should find application beyond magnetic resonance, as many other systems, such as electron dots and superconducting qubits, have similar Hamiltonians. The work in this chapter was published in Ref. [5].

2.2 Establishing controllability

The time evolution of a closed quantum system can be described by the differential Schrödinger equation. In all cases I assume that the system is governed by a time-independent natural Hamiltonian \mathcal{H}_{nat} (known as the drift term in optimal control) and the time-dependent control Hamiltonian $\mathcal{H}_C(t)$ whose time dependence is under the experimenter's control. In theoretical work this is often assumed to be arbitrary control, but in practice there are finite power and bandwidth limits imposed on the control fields. Then, given a density matrix description of the state, $\rho(t)$, the time evolution is given by:

$$\dot{\rho}(t) = -\frac{i}{\hbar} [\mathcal{H}_{nat} + \mathcal{H}_C(t), \rho(t)]. \quad (2.1)$$

The differential equation is integrated to give the unitary form of the dynamics,

$$\rho(t) = U(t) \rho(0) U^\dagger(t). \quad (2.2)$$

With the unitary time evolution operator calculated by integrating the instantaneous propagator,

$$U(t) = \mathcal{T} \exp \left[\frac{-i}{\hbar} \int (\mathcal{H}_{nat} + \mathcal{H}_C(t)) \right]. \quad (2.3)$$

The Dyson time ordering operator \mathcal{T} tells us that when we calculate the integral in the exponential to be careful to get the order of the non-commuting terms right. In practice, time is discretized at a suitable rate such that $\mathcal{H}_C(t)$ is approximately piecewise constant. Then the propagator for timestep j with length Δt is,

$$U_j = \exp \left[\frac{-i}{\hbar} \Delta t (\mathcal{H}_{nat} + \mathcal{H}_C(t_j)) \right], \quad (2.4)$$

and the total unitary is calculated by multiplying all the U_j together $U(t) = \prod_j U_j$. Thus simulation of a quantum evolution reduces to calculating matrix exponentials whose complexity and calculation time increases by a factor of approximately 8 for each qubit added. As an aside, we currently rely on Matlab's built in `expm` which in turn is using scaling and squaring and Pade's approximant [6]. Other groups [7] have reported for quantum simulations that other approximations such as a Tchebychev series are faster and are worth exploring.

The first relevant question in quantum control is whether it is even possible to control the system: whether the full Hilbert space can be explored. This question is answered in computer science with universal gate sets: a small set of, say, one and two qubit gates which when put together appropriately can approximate, arbitrarily accurately, any unitary [8]. For physical implementations the question is more appropriately asked in the Hamiltonian language: whether the evolution under the natural (or drift) Hamiltonian combined with the time dependent control Hamiltonians can approximate any unitary (or average Hamiltonian) eventually. Note that this will not tell us how long it might take or how strong a control field is required to achieve the unitary; nevertheless, it is a starting point. The question is formally answered by determining whether the control and natural Hamiltonians, and commutators of them, generate the entire Lie algebra [9]. There are standard numerical techniques for determining whether a system is fully controllable given a set of natural and control Hamiltonians [10]. There are also intuitive results about controllability. Any system where arbitrary single qubit rotations are possible (the ability to generate rotations about two non-parallel axes) and a two body coupling Hamiltonian of any form that connects all qubits will be sufficient to generate complete control [11].

2.2.1 NMR Hamiltonians

For NMR in the liquid-state and solid-state the Hamiltonian may provide universal control in certain situations. Since the majority of NMR QIP is performed on the archetypal qubit, the spin- $\frac{1}{2}$ nucleus, this thesis will not consider quadrupolar nuclei. Nuclei with greater than spin- $\frac{1}{2}$ have a quadrupolar moment which couples the spin to electric field gradients. They tend to have fast decoherence rates as a result and are usually unsuitable for QIP purposes. The magnetic dipole of the nuclei interacts with the static magnetic field, which produces a Zeeman splitting between the two energy levels (spin aligned or anti-aligned with the static field), giving the Hamiltonian

$$\mathcal{H}_{Zeeman} = \gamma_i B_o (1 + \delta) Z = \omega_i Z, \quad (2.5)$$

where γ_i is the gyromagnetic ratio of nucleus i , B_o is the strength of the magnetic field and Z is the Pauli operator σ_z . In a similar manner I will use X/Y for σ_x/σ_y and I for the identity matrix. A subscript indicates that Pauli operator acting on the subscripted spin with the identity operator acting on all other spins; for example, in a three qubit system $X_2 = I \otimes X \otimes I = IXI$. Furthermore, I will always assume that the static magnetic field is along the z direction. This Pauli operator notation will also be used to refer to quantum states in NMR experiments. The thermal state in NMR is highly mixed and is mainly the maximally mixed state

with a small deviation called the deviation density matrix. So, for example, the thermal state for two spins is,

$$\rho_{thermal} = \frac{1}{4} \mathbb{1} + \gamma_1 \epsilon ZI + \gamma_2 \epsilon IZ, \quad (2.6)$$

where epsilon is some constant factor depending on the field and temperature and is $\approx 10^{-5}$ in typical fields at room temperature. Because the identity term is unobservable in NMR and unchanging under the dynamics it is typically dropped from the description and only the second two terms, the deviation density matrix is written down. This is known as the product-operator notation [12].

It is convenient to use a semi-classical picture of the spins precessing about the magnetic field with an angular frequency ω_i . The Larmor frequency is $\omega_i/2\pi$ and is on the order of tens to hundreds of megahertz in today's superconducting magnets where the field ranges from 2-20T. In the absence of symmetry of the molecule, each nuclear site in the molecule experiences a different electronic environment. The electron cloud partially shields the nuclei from the applied magnetic field so that the local magnetic field experienced by each nucleus is slightly different. This gives each nucleus a slightly different Larmor frequency called the chemical shift, denoted by $\delta \ll 1$ in Eq. 2.5. This separation in frequency space (of a few kilohertz) is what allows different nuclei of the same nuclear species to be distinguished and selectively addressed. Different nuclear species have different gyromagnetic ratios, so their Larmor frequencies are usually widely separated by tens or hundreds of megahertz in frequency in typical fields. The contrast between the homonuclear and heteronuclear frequency ranges leads to huge differences in the timescales for homonuclear and heteronuclear control. It is only when a purely heteronuclear system is considered that single-qubit rotations are much faster (approximately three orders of magnitude) than two-qubit coupling gates. It is commonly assumed that the requirement of different chemical shifts will limit the number of qubits in molecular system where the resonance frequencies cannot be engineered. However, there are numerous schemes where the full power of computation can be realized with global control [13]. These schemes do however rely on a unique addressable qubit at the end of the chain. In practice there are few molecules which satisfy the chemical and symmetry requirements to have sufficiently resolved peak groups to define individual qubits and have long coherence times so that potential molecules with more than 7 qubits are always of interest.

The nuclear spins interact with each other by two mechanisms: a direct spin-spin coupling via the dipolar interaction, and an indirect electron-mediated interaction, the J coupling. In a secular approximation, the dipolar coupling between spins i and j takes the anisotropic form,

$$\mathcal{H}_D = D_{ij} (2Z_i Z_j - X_i X_j - Y_i Y_j). \quad (2.7)$$

The dipolar coupling constant depends on the orientation of the two molecules. If θ is the angle between the static field and the vector joining i and j is \mathbf{r}_{ij} then,

$$D_{ij} = \hbar \left(\frac{\mu_0}{4\pi} \right) \frac{\gamma_i \gamma_j}{r_{ij}^3} (3 \cos^2 \theta - 1). \quad (2.8)$$

In the liquid state¹, the rapid tumbling motion of the molecules averages to zero the pseudo-dipolar or anisotropic portion of the J coupling and the dipolar coupling, both within the same molecule and (to a good approximation) between molecules [14]. This leaves only the weaker and isotropic J coupling. The coupling Hamiltonian between spins i and j is of an exchange form,

$$\mathcal{H}_J = \frac{\pi}{2} J_{ij} (X_i X_j + Y_i Y_j + Z_i Z_j) = \frac{\pi}{2} J_{ij} \sigma_i \cdot \sigma_j. \quad (2.9)$$

Since this coupling is mediated by the electrons binding the two nuclei, there is no J coupling between different molecules. If the difference in the resonant frequencies of the two nuclei involved is much greater than the coupling strength, then a secular approximation is valid. The off-diagonal terms can be ignored and the coupling is reduced to the weak-coupling Ising form where the $X_i X_j$ and $Y_i Y_j$ terms are dropped:

$$\mathcal{H}_{J_{weak}} = \frac{\pi}{2} J_{ij} Z_i Z_j. \quad (2.10)$$

This approximation is always valid for heteronuclear systems and for the majority of the homonuclear systems considered for a QIP. With this approximation, all the terms in the internal Hamiltonian commute, and tracking the system during periods of free evolution becomes particularly straightforward. This can be used to greatly simplify the control schemes.

The internal Hamiltonian enables one-qubit rotations about the Z axis and two-qubit controlled- Z gates. However, for universal control, one-qubit rotations about another axis are needed. These are implemented by the control-field Hamiltonians. By applying a radio-frequency (r.f.) field, we can induce transitions between energy levels whose energy difference is resonant with the field. The applied field is linearly polarized perpendicular to the static field and it can be decomposed into two components rotating in opposite directions at the r.f. frequency. Since the spin is precessing about the static field at its Larmour frequency, only the component that is rotating in the same direction as the spin and at a comparable frequency will have an effect. Indeed it is convenient to move into a frame rotating at the pulsing frequency: the so called rotating frame [14]. In the rotating frame the control field Hamiltonian is a constant magnetic field and can be described as,

$$\mathcal{H}_{control} = \omega_{nut}(t) (\cos(\phi(t))X + \sin(\phi(t))Y). \quad (2.11)$$

¹Although there are quasi-liquid states such as liquid-crystals where the partial alignment of the molecules with the external magnetic field leads to a residual dipolar coupling between spins on the same molecule.

Up to the hardware limitations of the spectrometer, it is possible to arbitrarily control the amplitude (ω_{nut} , the nutation or Rabi frequency) and phase (ϕ) as a function of time. With a standard liquid-state probe, nutation frequencies of up to 30 kHz are feasible and for solid state probes up to 300kHz is easily attainable. This order of magnitude difference is due to the use of saddle coils in the liquid state for easy top loading sample changing; whereas in the solid state, more efficient solenoid coils are usually used. If the r.f. frequency is on resonance with the spins' Larmor frequency, then in the rotating frame the contribution to the Hamiltonian of the static magnetic field along Z vanishes and the r.f. field looks like a constant field about an axis ϕ from the x-axis. Hence, the spins precess about this axis at a frequency ω_{nut} , and we can induce rotations about any axis in the X - Y plane in the rotating frame. The counter-rotating component will look like a magnetic field rotating at twice the pulsing frequency and its effect will quickly average to zero. In the rotating reference frame at the r.f. frequency, spins whose transition frequencies are off-resonant have an additional Z component in their Hamiltonian, and the effective rotation axis will be the vector sum of the r.f. field and the off-resonant Z field. If the spins are far off-resonance then this rotation axis will be close to the z axis and the spins will not be rotated into the plane. The combination of the internal Hamiltonian and the control fields gives universal control. The challenge is to find the control fields as a function of time that drive the system through the desired unitary evolution.

2.3 Quantum Optimal Control

Although somewhat surprising to traditional NMR spectroscopists, there is a very powerful approach to control which relies on straightforward numerical optimization. Traditionally, NMR has been used to probe an unknown system and so pulse design focussed on designing a huge variety of pulses with different frequency domain characteristics, say narrow band or broad band pulses [15]. Numerical optimization and optimal control was often used for this purpose [16]. In an NMR QIP the situation is turned around. The system is already well known and so pulse design can give pulses that effect the desired gates on the specific system. For small systems, or under unrealistic assumptions about the control available, such as instantaneously fast single qubit gates, it is possible to show analytical results for the time-optimal evolution for certain quantum gates [17]. However, for larger systems we resort to numerical methods. As above, the system is driven by a natural Hamiltonian and the control Hamiltonians. The control Hamiltonian is broken up into a sum of all the control knobs we have,

$$\mathcal{H}_C(t) = \sum_k u_k(t) \mathcal{H}_k. \quad (2.12)$$

In NMR the amplitude and phase (or the quadratures X and Y) Hamiltonians are controllable so that for example $u_k(t) \mathcal{H}_k = u_k(t) \sum_m X_m$ where index m specifies which spins are sensitive to the control field, e.g. all carbon nuclear spins. So, more generally for NMR QIP,

$$\mathcal{H}_C(t) = \sum_k \left(x_k(t) \sum_m X_m + y_k(t) \sum_m Y_m \right), \quad (2.13)$$

where k indexes the different isotopes (proton, carbon, nitrogen, etc.) and m indexes the nuclear spins of each species.

The pulse is first discretized into suitable time steps so that the control parameters become a series u_{k_j} (see Figure 2.1). If the system is fully controllable then we know given enough timesteps there is some sequence u_{k_j} that will create any desired unitary. The question of optimal control then becomes how to find such a sequence. If it is an optimal sequence then it should be the shortest, lowest power etc. depending on what the metric for optimal is.




Figure 2.1: A pulse is discretized into timesteps t_j . The optimal control algorithm then pushes the control parameters u_{k_j} in the correct direction such that the unitary implemented by the pulse is as close to the desired unitary as possible.

Given the series u_{k_j} and a good model for the system and apparatus it is possible to simulate the unitary on a classical computer as described above to obtain U_{sim} . This can then be compared to some desired target unitary U_{goal} . Since global phases do not matter (i.e. we do not have to create the exact unitary U_{goal} but $e^{i\phi}U_{goal}$ is sufficient), the global phase invariant Hilbert-Schmidt (HS) norm is a good choice for the fidelity function,

$$F_{HS} = \frac{\left| \text{tr} \left(U_{goal}^\dagger U_{sim} \right) \right|^2}{N^2}, \quad (2.14)$$

where N is the dimension of the Hilbert space. This fidelity measure can be seen as an imperfect motion reversal, and is linearly related to the average fidelity (the

squared state overlap between the desired and simulated output state averaged (Haar measure) over all input states [18]). It is also equal to the average attenuated correlation between ideal and simulated output density matrices [19].

With this fitness function the performance of any pulse can be calculated. The most naive approach is then to start with some random pulse, calculate its fidelity with the desired unitary and then to vary the controls such that the simulated unitary has a greater overlap with the goal unitary. That, in essence, is quantum optimal control, a numerical search for the best pulse. The clever part of optimal control comes in applying standard optimization strategies to the problem. Unfortunately, it is difficult to guarantee global optimality from any search method. It may be that the best pulse found from the numerical search has a fidelity of 99.9% and takes $1ms$ but without an exhaustive search that does not mean that there is not another pulse that would give 99.99% fidelity and take only $900\mu s$. This makes the “optimal” part of optimal control a bit of a misnomer as the guarantee of optimality is only as good as the search method is at finding global maxima. It should be noted that there are suggestions that in the absence of constraints on the controls, the fitness function landscape is relatively well behaved [20]: that it has no local maximum. This may explain some of the success of quantum optimal control.

The optimizer’s task is a difficult one because of the extreme cost of function evaluations and the “credit assignment problem”. Because each function evaluation requires simulating the quantum evolution (perhaps multiple times if the pulse has to be robust against inhomogeneties in the system) it is intrinsically unscalable and is limited to 10 qubits even with large cluster computers [21]. The credit assignment problem is to decide how much the control \mathcal{H}_k at timestep u_j contributes to creating the desired unitary and hence how to change it. This can be solved of course by taking the derivative of the fitness function with respect to u_{k_j} . However, if done naively numerically this requires an enormous number of function evaluations. Parameter u_{k_j} is increased slightly, the entire simulated unitary is calculated and the forward finite derivative of the fitness function calculated. For example, with a pulse defined on 500 timesteps and four control parameters (X and Y for proton and carbon say) then this will require 2001 simulations of the system. For seven qubits each simulation will take approximately a minute on a standard desktop computer (3GHz, 1GB RAM) so it will take over 30 hours just to make one move in the parameter space. A standard solution to this problem is back-propagation. The unitary is simulated both forwards starting from the identity and backwards from the goal unitary and the unitaries at each point are stored in memory. From these two sets, instantaneous derivatives of the fidelity function can be calculated with a few matrix multiplications rather than simulation of the entire system.

The unitary for each timestep is,

$$U_j = \exp \left[-i\Delta t \left(\mathcal{H}_{int} + \sum_i u_i(j) \mathcal{H}_i \right) \right], \quad (2.15)$$

where $\mathcal{H}_x = \sum_k X_k$ for example and Δt is the length of the timestep. To first order, the derivative of the timestep's unitary propagator with respect to the control fields can be evaluated without finite differencing or another matrix exponential as,

$$\frac{\delta U_j}{\delta u_i(j)} \approx -i\Delta t \mathcal{H}_i U_j \quad (2.16)$$

where we require $|\Delta t \mathcal{H}| \ll 1$ for the approximate derivative to be accurate. The total unitary for the pulse can be calculated as the product of each timestep's unitary,

$$U_{tot} = U_M U_{M-1} U_{M-2} \dots U_3 U_2 U_1. \quad (2.17)$$

The gradient of the fitness function Eq. 2.14 can now be explicitly calculated as

$$\frac{\delta \Phi}{\delta u_i(j)} = \frac{1}{N^2} \times \left[\text{tr} \left(\left(U_{j+1}^\dagger \dots U_M^\dagger U_{goal} \right)^\dagger \frac{\delta U_j}{\delta u_i(j)} U_{j-1} \dots U_1 \right) + c.c. \right]. \quad (2.18)$$

On either side of the derivative are the partial propagators in the forward direction in time (starting from $\mathbb{1}$) and backwards in time (starting from U_{goal}). Thus by storing the forwards and backwards propagation of the unitary and substituting Eq. 2.16 into Eq. 2.18, gradient information about the fitness function can be obtained without finite differencing or recalculation of the entire propagator. The gradient information leads to a much more efficient search determining the direction in which the control parameters should be moved to improve the fitness function. This information can be used by a simple steepest-ascent hill-climbing algorithm to optimize the pulse. For faster optimization, conjugate-gradient techniques dramatically speed up the convergence.

This application of standard optimal control to quantum pulse engineering was codified in the GRadiant Ascent Pulse Engineering algorithm [22]. The code I wrote is closely based on the published algorithm with some modifications discussed below that were important for high fidelity control experimentally. Further technical details of the code can be found in Appendix B.

The numerically optimized sequences from optimal control techniques drive the system through a potentially complicated and non-intuitive path where small errors in the model of the system and apparatus might lead to a large drop in the fidelity of the pulse. An optimal control pulse is fortuitously robust to some noise in the pulse. At the point of convergence the pulse is at a local maximum in the parameter space and so the derivative of the fidelity with respect to any control parameter

is zero. Thus any small fluctuations in the control parameter due to noise in the pulse generation will have no effect to first order. However the pulse can be made explicitly robust to known uncertainties by defining a new fidelity function which is an average fidelity from the pulse simulated over a range of parameters in the model. The two most relevant to NMR QIP are inhomogeneities across the sample in both the r.f. field and the static field. In liquid-state NMR, good shimming can reduce the static field variation to less than 1Hz. However, in the solid-state, susceptibility mismatch between the sample and air in addition to imperfect crystals lead to much larger static field inhomogeneities of up to 300Hz. In both the liquid and solid state the r.f. field varies across the sample (see Figure 2.2 for a typical distribution). The distribution plotted is for the proton coil in a TCI cryogenic probe. I expect the proton coil to be the worst as it is the closest to the sample. The carbon coil is further away and so is relatively more homogeneous over the sample volume. It has also been noted that the inhomogeneity is particularly bad in cryogenic probe heads [23]. An example of a 7 qubit pulse robust to r.f. and static field variation is shown in Figure 2.3 and a pulse for the solid-state malonic acid system is shown in Figure 2.4. It should be noted that the usefulness of the robust pulses extends far beyond ensemble systems. They will be useful in other systems to improve the fidelity of control against for example, control field miscalibration or slow drift in the natural Hamiltonian.

Even with having to make only two simulations per step, finding pulses for large systems is tediously slow - particularly when looking for robust pulses which must be simulated over a range of parameters. The code implements a subsystem style approach to speed up finding pulses. Although this does not directly address the scalability issues with optimal control, it provides a practical method to find pulses on larger spin systems. A similarly motivated approach has been used for typical NMR experiments to allow the simulation of up to 200 spins [24]. If the system can be decomposed into subsystems (possibly overlapping), such that for each subsystem the desired unitary operator factors into one acting on the subsystem and another acting on the complement, then we can find pulses by defining the fitness function as a weighted sum of fitness functions for each subsystem. The subsystems must be defined so that the dominant dynamics of the system is captured. In particular, every strong coupling must be internal to at least one of the subsystems and strongly coupled groups must be simulated together to capture multi-body dynamics. We simulate each subsystem separately to reduce the simulation complexity. The hope is that if the pulse performs the desired unitary on each of the sub-systems then the pulse will most likely perform the desired unitary on the full system because higher order coherences are unlikely to be excited during the pulse.

Take as an example the crotonic acid system with seven qubits (see Appendix A for details of the molecule). Pulses implementing one-qubit rotations have the de-

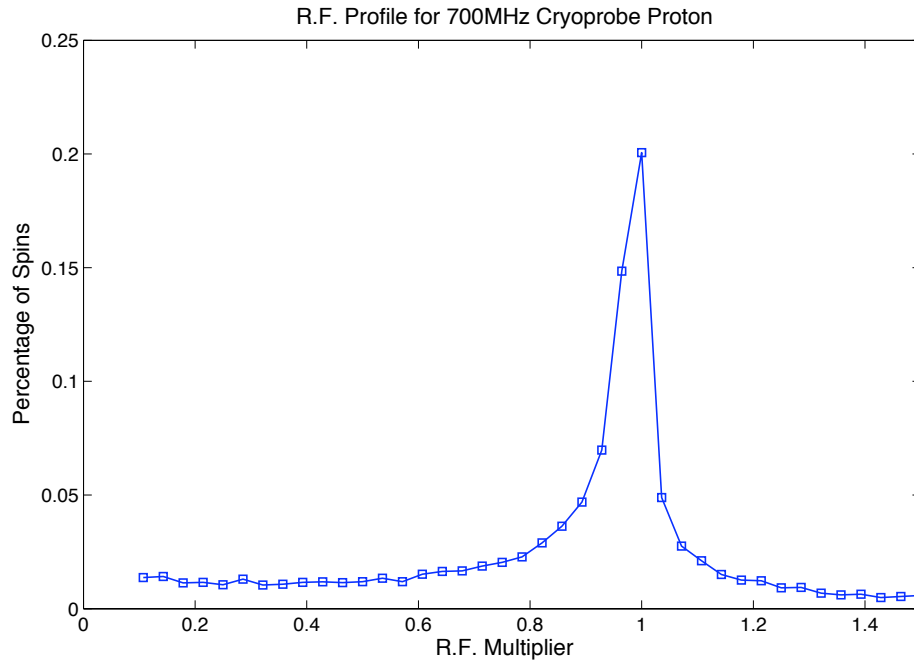


Figure 2.2: r.f. inhomogeneity profile on a Bruker TXI 700MHz cryogenic probe. The profile is measured with a nutation experiment: the integrated intensity of the NMR signal of a clean well isolated peak (in this case chloroform dissolved in deuterated acetone) is measured as a function of pulse length. A nutation field of approximately 12.5 kHz was incremented in time by $5\mu\text{s}$ from $5\mu\text{s}$ to 2ms. The integral of the NMR signal was evaluated at each point. For an ideal qubit we would expect to see Rabi oscillations with perfect contrast. However, due to the r.f. inhomogeneity, different parts of the sample experience different nutation rates and the signal decays. Fourier transforming the signal reveals the distribution of nutation frequencies. The ideal r.f. multiplier of unity is assumed to correspond to the peak of the distribution. Assuming all spins are visible in the experiment then the area under the curve must be normalized which allows us to scale the y-axis appropriately. As intuitively expected, the distribution is skewed towards lower fields. We would expect that the coils will create an volume of high homogeneity field in center of the sample and then towards the edges of the sample the field strength will drop off. Thus, there will be a larger volume at the edges of the sample versus the centre giving a distribution skewed towards lower field. There is also a substantial tail at lower fields which must be suppressed for high fidelity control.

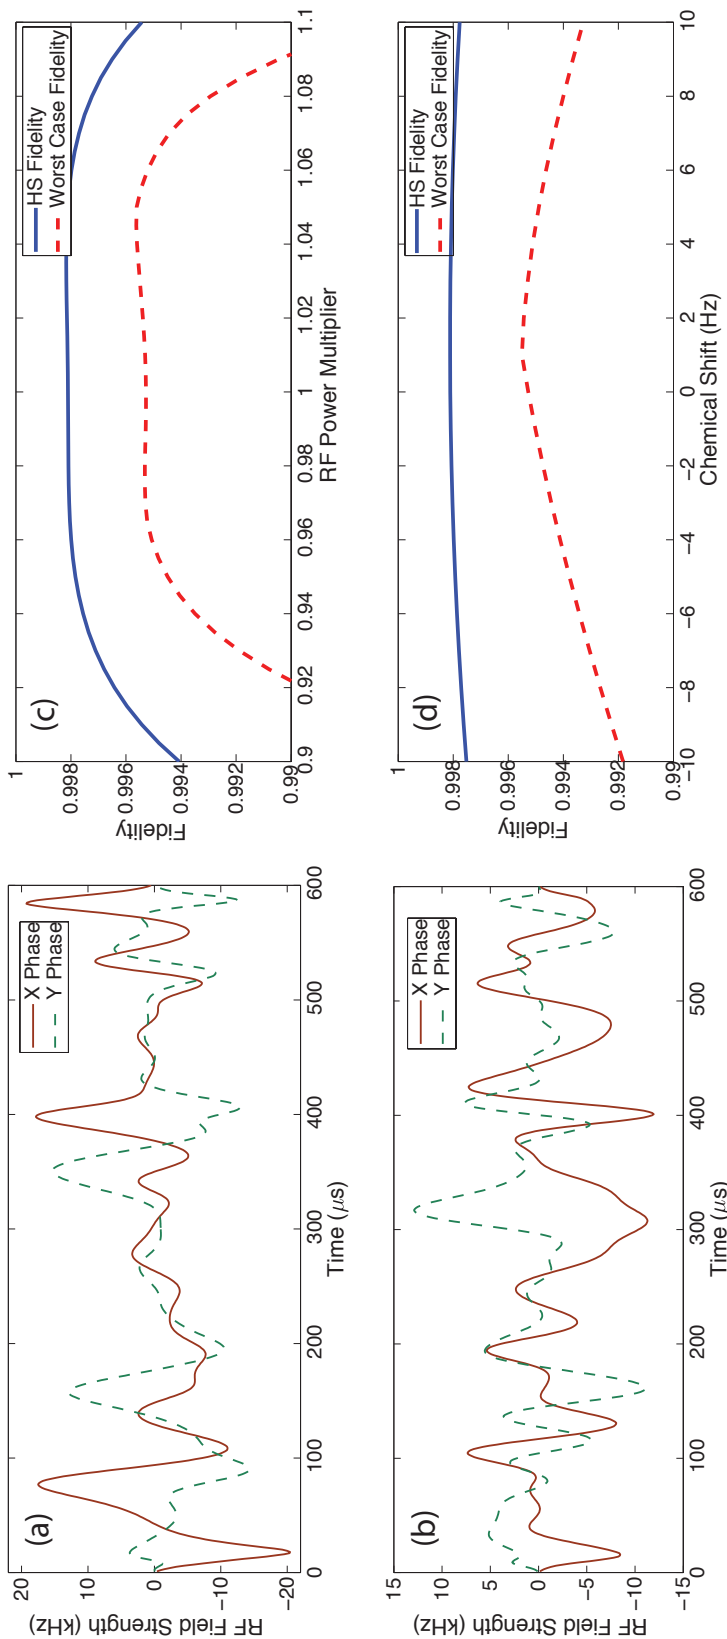


Figure 2.3: Plot showing the quadrature components of the proton (a) and carbon (b) control fields for a GRAPE pulse that implements a 90 degree rotation on H_1 in the crotonic acid molecule. This pulse is one of the more difficult to find one-qubit gates on the crotonic acid molecule (see Appendix A for details of the molecule), as H_1 and H_2 are only 788 Hz apart on a 700 MHz spectrometer. The pulse shape is smoothly varying and goes to zero at the start and end points so as to avoid any transient effects in implementation. The total duration of the pulse is only 600 μs and should be compared with the 3 – 7 ms required for standard soft pulse shapes. This pulse is not time optimal, but shorter pulses tend to require experimentally unfeasible high power. The right-hand-side plots show the robustness of the pulse to static and r.f. field inhomogeneities. Both the HS fidelity (Eq. 2.14 - solid line) and the worst case fidelity (the overlap of the output state from the simulated unitary and the output from the ideal unitary minimized over all input states - dashed line) are shown.




Figure 2.4: Plot showing the quadrature components for a GRAPE pulse that implements a complicated three qubit gate, the controlled-NOT-NOT gate in malonic acid molecule. The C-NOT-NOT gate is two two-qubit controlled-NOT gates combined: if the state of the control bit is 1 then both the target bits are flipped. As discussed in the text, extreme robustness to static field inhomogeneities are necessary for high fidelity control in solid-state NMR and the HS and worst-case fidelity are shown for versus r.f. and static field variation.

sired factorization property with respect to any subsystem. To capture the strong couplings, we use the following four overlapping subsystems: $\{M, C_1\}$, $\{C_1, C_2\}$, $\{H_1, C_2, C_3, H_2\}$, and $\{C_3, C_4\}$. Qubit C_1 is then covered by two subsystems capturing its strongest couplings. If the pulse refocusses these strong couplings to M and C_2 it is likely to also refocus weaker couplings to say C_4 without additional optimization. In general though, even if the pulse performs the desired subsystem unitaries with high fidelity, there is no guarantee that the smaller unconsidered couplings and/or many-body effects will not give a much lower fidelity implementation on the full system. Empirically however, we have found the constructed pulses work well on the full system, both for short one-qubit gates and for longer coupling gates. Furthermore, it substantially speeds up finding pulses for the seven-qubit crotonic acid system. The pulse shown in Figure 2.3 was found with this method. It has 99.9 % fidelity averaged over the above subsystems and showed a 99.7 % fidelity on the entire seven qubit system.

Up to this point it has been assumed that the controls can be varied arbitrarily and that unlimited r.f. power is available: neither is true experimentally. R.F. power limits are easily imposed by resetting the power when it exceeds a preset value. This however, renders the conjugate gradient approach ineffective and more elegant solutions are possible [25]. In practice it is found that enforcing lower power constraints leads to longer pulses. Discontinuities in the amplitude and/or phase of the r.f. will lead to switching transients due to the finite bandwidth of the probe circuitry and pulse generation [26]. There are a few solutions to this problem. Firstly the resonant probe circuit can be modeled with linear elements (capacitors, inductors, resistors), an impulse response calculated, and the field seen by the sample for a given input pulse calculated. Assuming a finite impulse response, modified derivatives in the GRAPE algorithm can be calculated [27]. Intuitively because of the transient effects, the level of the control amplitude at timestep t_j of the input pulse will affect the level seen at the sample at some later time-step t_{j+l} . Thus the derivative of the fitness function at time-step t_j must take this into account. Using the product rule, the derivative becomes a sum of as many terms as are in the impulse response. This solution is elegant but requires an accurate model of the system and will not capture transient effects with frequencies beyond the sampling frequency dictated by the pulse timestep length. A second solution [25], similar to the one for hard-limiting the r.f. power is to apply a low-pass filter to the derivatives to filter out high frequency components outside the bandwidth of the system. This however, may affect the direction of the derivatives and so affect the very useful conjugate gradient techniques. The code implements a third solution which has the side benefit of speeding up pulse finding. The idea is that even though the spectrometer can discretize the points on a very fast timescale (100ns), it is rare that that many points are actually needed to find a high fidelity

pulse. This leads to the following strategy:

1. Find a high-fidelity pulse using relatively long timesteps, e.g. $20 \mu\text{s}$. The length of the timestep is constrained so that approximate derivative (Eq. 2.16) is still accurate and so that there are sufficient number of timesteps to find a high fidelity unitary.
2. Digitally smooth this pulse with shorter timesteps, e.g. $1 \mu\text{s}$, making sure that the pulse bandwidth is within the system limitations.
3. Use this smoothed version of the pulse as an initial guess for the numerical optimization. There will be a small loss of fidelity from the smoothing. Nevertheless, it will be a good starting point (usually in the 95-99% range) and will not take long to re-optimize.
4. If necessary, repeat the smoothing and optimization procedure.

Empirically we have found that the re-optimization of the pulse in the third step changes the pulse very little, and once the new optimized pulse is determined, it is still sufficiently smooth. It is also necessary to ensure that the beginning and end of the pulse go to zero to ensure a smooth experimental turn-on/off. This can be achieved with a penalty function that penalizes high powers at the beginning and the end.

Another improvement in our code is the addition of timestep derivatives. In conventional GRAPE the length of each timestep and the total number of timesteps is fixed. To determine the time optimal pulse an arbitrary pulse length is chosen and many attempts to find a pulse are made. If one is found that has an acceptably high fidelity, then the time is shortened. The process is iterated until it is no longer possible find a high enough fidelity pulse and that point is considered time optimal. This is a long process and given an arbitrary gate and molecule it is difficult to determine a priori what a reasonable time is. Adding derivatives with respect to the length of each timestep and allowing the timesteps to vary circumvents this problem. A penalty function is used to prevent excessively long pulses or to squeeze the time down to the time optimal point and once a timestep length drops below a threshold that timestep is dropped from the pulse.

2.4 Scalable Quantum Control

In larger systems an alternative approach to control must be taken because the classical cost of simulating the full quantum dynamics grows exponentially with the

system size. This prohibits the application of optimal control techniques to the full system. Some well-chosen simplifications and approximations must be made to the model. The control sequence is constructed from simple predefined but imperfect building blocks (for example standard frequency selective pulses). The blocks are then systematically put together to form a pulse sequence ensuring that the errors in the building blocks do not build up excessively as the sequence progresses. It is possible to efficiently design such sequences only if we judiciously simplify the model and take into account only the largest and first order errors in the blocks.

Typical building blocks are one-qubit unitaries that involve selectively rotating one spin. If the spins have distinct resonant frequencies, this corresponds to frequency selective pulses. The goal is to obtain a flat inversion or excitation profile over the range of frequencies for the peak group corresponding to a particular spin, while at the same time having negligible excitation effects outside this region. The problem of obtaining such pulses has of course had much attention in the long history of NMR. The most successful approach has been to use shaped pulses, and a huge variety of increasingly complicated shaped pulses have been developed with various bandwidths, excitation profiles (tipping angles as a function of chemical shift) and robustness to experimental imperfections [15]. In the linear regime the excitation profile in frequency space of a pulse is related to the Fourier transform of its time domain profile². As one would expect, the longer the pulse, the more selective it is in frequency space. Furthermore, one can tailor the excitation profile by shaping the pulse. For example, a Gaussian shaped pulse has a Gaussian shaped excitation profile. Thus, given the internal Hamiltonian of the molecule, it is straightforward to design a set of pulses for single spin rotations. An example of a pulse and its excitation profile is shown in Figure 2.5.

For many NMR spectroscopy experiments, it is sufficient to consider only the excitation profile, but for quantum computing purposes we need to accurately keep track of all the effects of the pulse. In particular there are off-resonant and coupling effects. Although the resonant frequency of a particular spin may be far outside the excitation bandwidth of a selective pulse intended for another spin (so that there is no rotation about an axis in the plane), there still is a substantial phase evolution from the off-resonant or transient Bloch-Siegert effect [29]. The original Bloch-Siegert effect paper [30] dealt with the effect from the counter rotating component of the linearly polarized r.f. field. For modern high-field spectrometers this effect is negligible. However, the counter rotating field is simply an off-resonance field (off-resonant by $2\omega_{rf}$) and so the same calculations apply to any off-resonant field. The effect can be seen by calculation of the Magnus expansion of the average Hamiltonian [31]. Average Hamiltonian theory states that the unitary propagator

²The linear regime applies only to small angle rotations but it still gives good intuition for larger rotation angles.




Figure 2.5: A portion of a single-scan carbon thermal spectrum of the fully labeled molecule trans-crotonic acid [28] at 16.4T, showing three qubits, C_2, C_3 and C_4 . See Appendix A for details of the molecule. By modulating the amplitude of a pulse sent at a frequency resonant with the chemical shift of C_2 we can selectively rotate C_2 . Shown in the inset is a truncated inverse secant pulse of 700 μs . Overlain on the spectrum is the excitation profile of a 90 degree pulse and the inversion profile of a 180 degree pulse of this length. The profiles are determined from a simulation of the spin dynamics of a single qubit as the resonance frequency is varied. The qubit starts along the $+Z$ axis of the Bloch sphere and the plots show the projection onto the $X - Y$ plane for the excitation profile and the $-Z$ axis for the profile. It is clear that the pulse affects C_2 but does not significantly excite C_3 or C_4 . As discussed in the text, C_3 and C_4 experience a transient Bloch-Siegert effect, which must be accounted for in a quantum computing experiment.

for some time-dependent Hamiltonian can be represented as the time evolution of some constant average Hamiltonian:

$$U(t_c) = T \exp \left[-i \int_0^{t_c} \mathcal{H}(t') dt' \right] \quad (\text{definition})$$

$$= \exp \left[-i \overline{\mathcal{H}} t_c \right]. \quad (\text{ansatz})$$

The Magnus expansion provides a way to calculate $\overline{\mathcal{H}}$:

$$\overline{\mathcal{H}}^{(0)} = \frac{1}{t_c} \int_0^{t_c} dt_1 \mathcal{H}(t_1); \quad (2.19)$$

$$\overline{\mathcal{H}}^{(1)} = \frac{-i}{2t_c} \int_0^{t_c} dt_2 \int_0^{t_2} dt_1 [\mathcal{H}(t_2), \mathcal{H}(t_1)]; \quad (2.20)$$

$$\overline{\mathcal{H}}^{(2)} = \frac{-1}{6t_c} \int_0^{t_c} dt_3 \int_0^{t_3} dt_2 \int_0^{t_2} dt_1 \{ [\mathcal{H}(t_3), [\mathcal{H}(t_2), \mathcal{H}(t_1)]] + [\mathcal{H}(t_1), [\mathcal{H}(t_2), \mathcal{H}(t_3)]] \}; \quad (2.21)$$

In the frame rotating at the spin's Larmor frequency, any off-resonant field will have the time-dependent Hamiltonian:

$$\mathcal{H}_R = I_x \cos(\Omega t) + I_y \sin(\Omega t), \quad (2.22)$$

where Ω is the frequency difference between the rotating frame and the pulsing frequency. We can now calculate the average Hamiltonian for any time multiple of the period $t_c = \frac{2\pi}{\Omega}$. Clearly the zeroth order average Hamiltonian vanishes. However, to first order, the effect of an off-resonant pulse is to shift the resonant frequency of the spins. The commutator of X and Y is Z and in particular:

$$\begin{aligned} \overline{\mathcal{H}}^{(1)} &= \frac{-i}{2t_c} \int_0^{t_c} dt_2 \int_0^{t_2} dt_1 [\mathcal{H}(t_2), \mathcal{H}(t_1)] \\ &= \frac{-i\omega_1^2}{2\frac{2\pi}{\Omega}} \int_0^{\frac{2\pi}{\Omega}} dt_2 \int_0^{t_2} dt_1 \{ (I_x \cos \Omega t_2 + I_y \sin \Omega t_2) (I_x \cos \Omega t_1 + I_y \sin \Omega t_1) \\ &\quad - (I_x \cos \Omega t_1 + I_y \sin \Omega t_1) (I_x \cos \Omega t_2 + I_y \sin \Omega t_2) \} \\ &= \frac{\omega_1^2}{2\frac{2\pi}{\Omega}} I_z \int_0^{\frac{2\pi}{\Omega}} dt_2 \int_0^{t_2} dt_1 (\cos \Omega t_2 \sin \Omega t_1 - \cos \Omega t_1 \sin \Omega t_2) \\ &\quad (I_x I_y = -I_y I_x = \frac{i}{2} I_z) \\ &= \frac{\omega_1^2}{2\frac{2\pi}{\Omega}} I_z \int_0^{\frac{2\pi}{\Omega}} dt_2 \left(\cos \Omega t_2 \left(-\frac{1}{\Omega} \cos \Omega t_1 \Big|_0^{t_2} \right) - \sin \Omega t_2 \left(\frac{1}{\Omega} \sin \Omega t_2 \Big|_0^{t_2} \right) \right) \\ &= \frac{\omega_1^2}{2\frac{2\pi}{\Omega}} I_z \frac{1}{\Omega} \int_0^{\frac{2\pi}{\Omega}} dt_2 (-\cos^2 \Omega t_2 + \cos \Omega t_2 - \sin^2 \Omega t_2) \\ &= -\frac{\omega_1^2}{2\frac{2\pi}{\Omega}} I_z \frac{1}{\Omega} \frac{2\pi}{\Omega} \\ &= -I_z \frac{1}{2} \frac{\omega_1^2}{\Omega} \end{aligned} \quad (2.23)$$

where ω_1 is the nutation frequency of the r.f. field and Ω is the offset. Obviously this is a perturbative calculation and higher order terms from the Magnus expansion are necessary when $\Omega < \omega_1$. A few relevant examples are worth considering. For example, a spin 3 kHz off-resonant from a 1 ms pulse that performs a 180 degree rotation on an on-resonance spin will experience an additional phase shift of ~ 15 degrees. This is roughly the situation for a selective pulse on $C_{2/3/4}$ in crotonic acid. This is non-negligible and must be accounted for. The original Bloch-Siegert shift comes from the counter rotating component at $2\omega_{rf}$. In solid state systems we can achieve nutation frequencies of 300kHz at 300MHz Larmour frequency. This will give a shift of 75Hz which is substantial, although in practice, the lines are broadened by dipolar interactions much stronger than this so that linewidths of 10's of kHz are common and the 75Hz is inconsequential. For the liquid-state situation where nutation frequencies of 16.7kHz are possible at 175MHz carbon Larmour frequencies, this gives a shift of approximately 0.5Hz. In practice, fields of that strength are kept on for less than $100\mu s$ and so the effect is minimal. For most optimal control pulses r.f. fields average approximately 10% of this and so the effect is much smaller than the linewidth or other experimental uncertainties.

The second non-ideal effect occurring during a pulse is that the couplings still evolve during the pulse (which may have a duration comparable to $1/2J$), and the coupling Hamiltonians do not commute with the r.f. Hamiltonian. Particularly for long pulses on spins that have strong couplings, there are substantial deviations from the ideal action of the pulse. These can be partially countered by self-refocussing pulses [15], such as the Hermite pulses (see Table 2.2).

Both the off-resonant and the coupling effects can be accounted for through a decomposition scheme where the imperfections are unravelled from the simulated pulse and represented as phase and coupling errors before and after the ideal pulse [28]. These errors are then corrected as part of a larger pulse sequence. A similar method with symmetric negative time evolution was presented in Ref. [32]. The method used here is more accurate and general by not constraining the errors to be of the form of natural Hamiltonian evolution, but requires more complicated refocussing schemes. The effects of the pulse on the entire system is modeled by the following decomposition:

$$\begin{aligned}
 U_{sim} &= e^{-i(H_{rf} + \sum_i (\omega_i \mathbf{Z}_i) + \sum_{i < j} (J_{ij} \mathbf{Z}_i \mathbf{Z}_j)) \Delta t} \\
 &\simeq \prod_i e^{-i\alpha_i^{post} \mathbf{Z}_i} \prod_{i < j} e^{-i(\beta_{ij}^{post} \mathbf{Z}_i \mathbf{Z}_j)} U_{ideal} \times \\
 &\quad \prod_i e^{-i\alpha_i^{pre} \mathbf{Z}_i} \prod_{i < j} e^{-i(\beta_{ij}^{pre} \mathbf{Z}_i \mathbf{Z}_j)}. \tag{2.24}
 \end{aligned}$$

This decomposition resolves the simulated pulse into the ideal unitary operator sandwiched on either side by Z rotations (to account for the transient Bloch-Seigert

shift and the chemical shift evolution during the pulse) and ZZ couplings to account for the couplings that occurred during the pulse (see Fig. 2.6). This model for the pre- and post-error terms does not work for arbitrary pulses, but in the relevant case of spin selective pulses, it will capture most of the first order dynamics. All couplings that do not involve the target qubit are trivial to extract and can be perfectly reversed. For couplings involving the target qubit, the control fields modulate the coupling necessitating a numerical optimization of the error terms. For 90 degree pulses ZZ pre- and post-error terms can represent coupling evolution during the pulse. However, for 180 degree pulses, because the ideal pulse should refocus the couplings, ZZ error terms are not sufficient, and self-refocussing pulses are desirable. Since the pre- and post-error terms do not commute with the control Hamiltonian, the decomposition faithfully represents the true dynamics only when $|\Delta t J| \ll 1$ where Δt is the length of the pulse and J is the strongest coupling.

This efficient pre- and post-error term analysis requires two steps: the first is to *efficiently* compute the relevant dynamics of the system under the r.f. pulse; the second is to extract optimal pre- and post-error terms that are correctable and give a good representation of the simulated pulse. Simulation of the full spin system will of course reveal everything about the pulse and its associated unitary errors. However, this is intrinsically unscalable and so a method that allows us to capture the main dynamics of the pulse without simulating the full system must be used. The errors must be represented in a way that is faithful to the true dynamics but at the same time is amenable to compilation and correction of error terms.

Consider an illustrative three spin system. Spins 1 and 2 are of the same species, whereas spin 3 is different. The propagator U of a pulse affecting spins 1 and 2 can be written as

$$U = \exp\{-i\Delta t(\omega_{nut}(\cos\phi(X_1 + X_2) + \sin\phi(Y_1 + Y_2)) + \omega_1 Z_1 + \omega_2 Z_2 + \omega_3 Z_3 + J_{12}Z_1 Z_2 + J_{23}Z_2 Z_3 + J_{13}Z_1 Z_3)\}, \quad (2.25)$$

where the control and internal Hamiltonian terms are defined as above. If the control fields are functions of time, the pulse must be discretized and integrated over many time steps.

The dynamics of the system can be expanded into contributions of different coupling order. If coupling effects are ignored, then the zeroth-order U_0 can be computed by three independent single spin simulations as

$$U_0 = e^{-i\Delta t(\omega_{nut}(\cos\phi X_1 + \sin\phi Y_1) + c_1 Z_1)} e^{-i\Delta t(\omega_{nut}(\cos\phi X_2 + \sin\phi Y_2) + c_2 Z_2)} e^{-i\Delta t(c_3 Z_3)}. \quad (2.26)$$

The effects of the coupling can then be added in a perturbative manner. The first step is to consider the effect of the couplings between pairs of spins or subsystems:

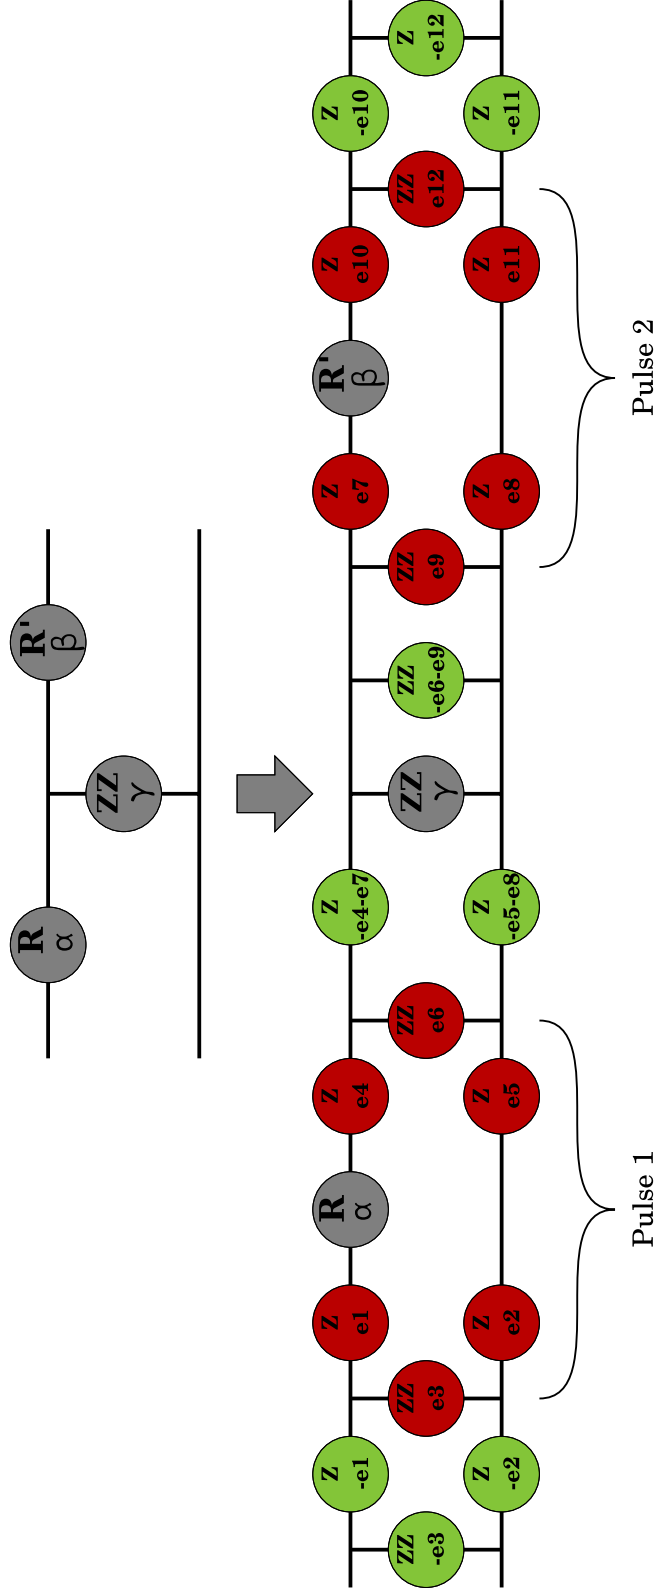


Figure 2.6: Sequence including pre- and post-error terms and corrections. The circles denote rotations and couplings. The axis and angle of the rotation are given one above the other. A rotation around the ZZ axis by θ is the evolution $e^{-iZ_j Z_k \theta/2}$ for the target spins j and k . Grey circles are for the intended sequence. The red circles represent the pulse pre- and post-error terms, while the green circles represent the evolutions that are implemented by the sequence compiler in order to cancel the pulse error terms. The ZZ-corrections are implemented by adjusting the free-evolution periods between pulses and the Z-corrections are performed with phase tracking calculations.

couplings 1-2, 2-3 and 1-3 are considered. By considering only pairs of spins, coherent three-body coupling effects are lost. For example, there is an extra effect on the dynamics of qubit 1 from the coupling to qubit 2 because qubit 2 is in turn affected by the coupling to qubit 3. These higher order coupling effects can be neglected when coupling effects during the pulse are small. Moreover, for soft selective pulses in the weak coupling approximation, many of the second-order contributions vanish because all the internal Hamiltonians of the off-resonant spins commute with each other, so that in most cases, stopping at first order is sufficient to encompass all the relevant dynamics. Indeed it is difficult to devise a situation where these second order coupling effects are relevant. Under simple single spin rotations they will have no effect. A pulse on spin 1 will rotate spin 1 about an axis in the plane and may have a transient Bloch-Siegert effect on spins 2 and 3. However, under a weak coupling ZZ approximation spin 3 will have no effect on spin 1 through spin 2. It is only when one considers pulses which rotate all the spins (hard pulses or optimal control pulses) that these indirect couplings matter. For example, consider starting with polarization on spin 3 or the state IIZ using the product operator notation of NMR [12]. During a “hard” pulse which affects all the spins, this will be rotated into IYY which will then couple to spin 2 turning into IZX but since spin 2 is also being rotated by the hard pulse this will turn into IYX which will in turn couple to spin 1 turning into ZXX . This effect will not be captured by pairwise simulations (or commuting pairwise error terms).

The simulation protocol is then to first simulate all single spins and then to simulate all pairs of spins. Even this may be unnecessary if the couplings vanish beyond a certain radius as might reasonably be expected in a physical architecture. The simulations are of a fixed one or two qubit size and there will be $\mathcal{O}(n^2)$ to perform so the complexity of the simulations scales only polynomially with the number of qubits. This of course could be extended to larger numbers of spins (say consider all groups of three spins) in a cluster-state style expansion [33].

From the efficient simulations of the previous paragraph, we need to determine the pre-and post-error terms surrounding the ideal operation. This is achieved through a simple numerical search optimizing the following fitness function:

$$\Phi = \text{Real} \left\{ \text{tr} \left(U_{sim}^\dagger (U_{post} U_{ideal} U_{pre}) \right) \right\} / N. \quad (2.27)$$

This compares the simulated unitary to its representation as the ideal unitary, restricted to the relevant qubits, sandwiched between the pre- and post-error terms. This fitness function is sensitive to global phase because we are considering only a subsystem, and so to be consistent when representing the gates on the full system, local phases matter. To extract the optimal error terms one can use a classical search algorithm. Note that this search is performed on a relatively low-dimensional parameter space. For pulses designed for one qubit, the single spin simulations

require only two parameters for the Z error terms on this spin. In the pairwise simulations there are six parameters for the Z and ZZ error terms on the coupled spins.

For standard shaped pulses designed to act independently on one or more spins, the following procedure allows one to compute error terms for the pulse representation.

1. Perform single spin simulations including the r.f. control fields for each spin and optimize the Z error terms (with respect to the intended evolution on the spin) to capture the effects of the chemical shift and the transient Bloch-Siegert effect.
2. Perform simulations of each pair of spins, including the pairwise coupling and the r.f. control fields. Optimize the Z and ZZ error terms with respect to the intended evolution on both spins. Remove the contribution of the Z error terms found in the first step from the ones found here (by dividing the terms or subtracting the exponents) to account for the fact that the simulation also accounts for the effects modeled in the first step. Call these the “incremental Z error terms”.
3. Determine the Z error terms for each spin as the product of the terms found in the first step and the incremental terms from each pair-wise simulation involving this spin. The ZZ error terms are the ones found in the second step.

An example of the representation of 90 degree rotation on C2 is shown in Table 2.1. The representation has a fidelity of 99.96% with respect to a simulation of the full system dynamics showing that pairwise simulations and Z and ZZ error terms are a suitable representation.

The error-term representation of Z and ZZ is chosen both so that it faithfully represents the dynamics of the spin system (dictated by the form of the natural Hamiltonian) and so that error terms can be corrected as part of a larger sequence. These pulse representations give us the building blocks for one-qubit gates. In a pulse sequence, two-qubit gates are achieved by periods of free evolution during which the coupling terms in the Hamiltonian evolve. If a J-coupling term is allowed to evolve for a time $1/2J$, then a controlled- Z gate (up to one-qubit Z rotations) is implemented between the coupled spins. All undesired couplings must be refocussed with 180 degree pulses [4]. The error terms can be corrected as part of the sequence. Z error terms are corrected with Z rotations. Z rotations can be considered a reference frame change and can be implemented instantaneously and with high precision by suitably modifying the phase of subsequent pulses and perhaps the observation phase. To be able to implement single spin Z rotations,

	M	H_1	H_2	C_1	C_2	C_3	C_4
M	0.0	0.41	-0.04	7.95	-0.41	0.41	-0.06
H_1	0.41	0.0	0.98	0.18	11.32	-0.11	0.41
H_2	-0.04	0.98	0.0	0.34	0.0	10.07	0.24
C_1	7.95	0.18	0.34	1.72	2.98	0.09	0.44
C_2	-0.41	11.32	0.0	3.00	0.0	4.95	0.09
C_3	0.41	-0.11	10.07	0.09	4.95	7.12	4.40
C_4	-0.06	0.41	0.24	0.44	0.09	4.40	-6.92

Table 2.1: Table showing the decomposition of the single spin and pairwise simulations for the pulse performing a 90 degree rotation on C_2 in the crotonic acid molecule. The diagonal gives the Bloch-Seigert shift in degrees with respect to each nucleus’ rotating frame, although in practice everything is calculated with respect to a single reference frame. The off-diagonal elements give the pre- (above diagonal) and post- (below diagonal) ZZ error terms in degrees. Because the pulse is symmetrical, so are the error terms, but the method is general enough to handle arbitrary pulses.

abstract individual rotating frames are defined. Just as the rotating frame for a single spin eliminates the Zeeman term in the Hamiltonian, individual rotating frames eliminate all chemical shift terms from the internal Hamiltonian. This could be implemented by a transmitter dedicated to each spin and rotating at that spin’s chemical shift, but this is experimentally and financially prohibitive. The phase evolution caused by the chemical shift term in the Hamiltonian is simply a Z rotation, and this evolution can be tracked with respect to some reference frame. Similarly the evolution of the transmitter at its frequency can be tracked, and when a pulse is required, the pulse phase is adjusted to obtain the correct phase in the spin’s rotating frame. Similarly the observation phase must also be calculated and adjusted. This tracking greatly simplifies the pulse sequences by eliminating the need to refocus the chemical shift evolution. It is also worth noting that because Z rotations are “free” and accurate it is beneficial to transform the quantum circuits to use Z rotations preferentially. A drawback of individual reference frames is that any given pulse has one phase and so it is not always possible to have the correct phase in two different reference frames. Therefore, global pulses that nontrivially affect multiple spins are not always feasible, and hence they must be decomposed into a sequence of simpler pulses.

The ZZ error terms are absorbed into the coupling goals. So for example if a ZZ evolution of 0.25 (in units of $\frac{2}{j}$ so that a ZZ evolution of 1 is the identity operator), is desired but the pulses before and after the coupling delay have ZZ

errors of 0.01 then a time delay that gives a ZZ evolution of 0.23 is necessary. Similar calculations apply to refocussing unwanted couplings.

Any unitary can be decomposed into a series of single qubit rotation events and ZZ coupling goals. The job then is to optimize the delays between the events to minimize the distance between the desired and actual couplings. This task is handled by a pulse sequence compiler. Refocussing pulses which change the direction of the coupling events can be used to assist this task. There are standard efficient algorithms for achieving this [4, 34]. However, it is not necessary to refocus all unwanted couplings during every evolution period. It is sufficient to track the evolution, and refocus it only when needed [35]. Single qubit gates and the Z and ZZ error terms in their representation commute with the terms of the internal Hamiltonian, except those involving the target nucleus. Therefore, the couplings can be let evolve through many gates and need to arrive at the goal coupling only when a noncommuting pulse affects either member of the coupling pair. The refocussing scheme can be made more efficient by the use of “virtual 180s”. Conventionally, every time a refocussing pulse is used, a second refocussing pulse must be applied at the end of the period to cancel the first one and ensure that all the spins return to their initial state, so that the computation is unaffected by the refocussing. While the first refocussing pulse is needed to physically refocus the coupling, the second one can be made virtual and not physically applied. The *virtual180* can be implemented by pushing it forwards or backwards (the *virtual180* can be created before or after the physically applied refocussing pulse which may help the refocussing) through the pulse sequence until it can be merged with another pulse (see Figure 2.7). If it is another refocussing pulse, then it can cancel with the virtual 180 created at that refocussing event. Or, if it is a computational pulse, the virtual 180 can be absorbed by modifying the phase of the pulse and introducing a frame change. Using the notation $R_\phi(\theta)$ to denote a rotation of θ about the axis an angle ϕ away from the x-axis in the $x - y$ plane, the following rotations should be equivalent:

$$R_\alpha\left(\frac{\pi}{2}\right)R_\beta(\pi) = R_z(\gamma)R_\delta\left(\frac{\pi}{2}\right) \quad (2.28)$$

using matrix multiplication and trigonometric identities it is easy to solve for γ and δ so that:

$$\gamma = 2(\alpha - \beta) \quad (2.29)$$

$$\delta = 2\beta - \alpha - \pi \quad (2.30)$$

Or in words: a 180 degree pulse at phase β followed by a 90 degree pulse at phase α is equivalent to a 90 degree pulse at phase $2\beta - \alpha - \pi$ followed by a rotation about the z axis of $2(\alpha - \beta)$. Since the Z rotation comes for free as a frame change we have eliminated the need to do the second 180. The sequence compiler keeps track of these *virtual180s* and it considerably simplifies writing pulse programs.




Figure 2.7: (a) A conventional refocussing scheme for a fragment of a pulse sequence. Shaded rectangles are 180 degree refocussing pulses. Unshaded rectangles are 90 degree rotations about an axis specified in the plane. To turn off the coupling between the two qubits during the delay between the two computational pulses, a refocussing 180 degree pulse is placed in the middle of the delay. A second 180 degree pulse is placed at the end of the period to cancel the refocussing pulse so that the refocussing scheme does not affect the computational frame. (b) Figure showing how the second 180 degree pulse can be made “virtual”. Rather than being implemented it is pushed forward to the next pulse on that qubit. If for example the next pulse is a 90 degree pulse, the virtual 180 can be combined to give a 90 degree pulse at a different angle and a Z rotation (dashed rectangle) which is implemented as a frame change (see text for details).

The sequence compiler was built to automate, as much as possible, the design of pulse sequences for an NMR QIP. The compiler takes as one input the fixed information about the internal Hamiltonian, namely the chemical shifts and J coupling values, and the error term information for the pulses obtained by optimal control and implementing the set of one-qubit rotations needed. The second input is a representation of a goal sequence of quantum evolutions that implements the desired algorithm or quantum network. The sequence compiler then determines delays between pulses so as to minimize the total error of the implementation compared to the intended evolution. The goal sequence is described with a purpose built language. Here is an example of such a sequence:

```
;pulse C190 0 @C1:X+
;zz 0.25 C1 C2
;refocus C3180 0.25
;pulse C290 0.75 @C2:0+
;z 0.5 C3
```

This sequence requests a 90 degree rotation about the x-axis on C_1 followed by a ZZ90 (equivalent to a controlled Z gate up to single qubit Z rotations) coupling gate between C_1 and C_2 followed by a 90 degree one-qubit gate on C_2 and a final Z rotation on C_3 . During the coupling period a “floating” refocussing pulse of 180 degrees about the y-axis is executed on C3. The @ instructions specify state assumptions that may simplify the optimization. The output of the compiler is a pulse sequence that can be directly executed by the spectrometer. The compiler automates a number of the tasks discussed above.

- Phase tracking: To avoid having to use one physical spectrometer channel per qubit, the compiler tracks the evolution of the nuclear rotating frames and the spectrometer channel frames throughout the computation. When a pulse about a certain axis is required for a nucleus, a simple calculation of the phase difference between the channel frame and the nuclear frame determines the phase at which the pulse should be sent in the channel frame to achieve a specific phase in the target nuclear frame. As noted above this freedom allows us to avoid having to refocus the chemical shifts at every step. It also allows free (pulse-less) implementation of Z rotations by executing a frame change on the target nuclei and updating the phases of all subsequent pulses. In addition, the relative phase evolution of the nuclei and the observation channel are tracked and the observation phase is adjusted appropriately. The frequency changes can be implemented by either changing the transmitter frequency or linearly ramping the phase of the pulse. The phase tracking calculations assume that the frequency change is phase coherent, but experimentally this

is not always the case (particularly for large frequency changes ($> 5\text{kHz}$) we observe random phase errors), and we have achieved more consistent results with phase ramping on our Bruker Avance spectrometers.

- **State assumptions:** In some algorithms the state of the system might be known at certain steps. For example, at the beginning of the algorithm we may know that a particular qubit is in the state $|0\rangle$ or the maximally mixed state $\mathbb{1}$. This can simplify the refocussing and phase tracking calculations. For example, if two spins are in the maximally mixed state, the coupling between them has no effect and so does not need to be refocussed. If a qubit is in a pseudo-pure state then its coupling effect with the other spins is reduced to an additional Z rotation, which can be accounted for with the phase tracking calculation. State assumptions are implemented by specifying the nuclear states after pulse commands, as shown in the example. The compiler also implements some elementary state update rules, e.g. a qubit in the state $|0\rangle$ will be updated to $|1\rangle$ after a 180 degree pulse.
- **Cancellation of *virtual180s*:** As described above the refocussing scheme is simplified by absorbing some 180 degree pulses into 90 degree pulses or canceling them with other 180 degree pulses.
- **Modifying coupling goals:** In some cases it may be expedient to modify the coupling goals by 0.5 using the following identity:

$$\exp\left(-i\frac{\pi}{2}Z \otimes Z\right) \simeq \exp\left(-i\frac{\pi}{2}Z\right) \otimes \exp\left(-i\frac{\pi}{2}Z\right) \quad (2.31)$$

Or that a ZZ coupling of 0.5 is equivalent to single qubit Z flips on both qubits (up to a global phase). For example, we may wish to have a coupling of 0.25 between two spins. However, because of previous refocussing pulses the coupling may be evolving in the negative direction and it might be easier to achieve -0.25. This physical coupling can be made equivalent to the desired coupling by performing phase flips on both qubits. In another situation with a strong coupling it might be easier to achieve 0.75 than 0.25 and again the 0.75 can be made equivalent by adding two phase flips. The compiler can in some instances recognize these situations and add the phase flips automatically.

- **Optimization of delays:** The pulse sequence can be considered as a sequence of events (computation and refocussing pulses) with delays in between. The delays between pulses serve the dual purposes of allowing desired coupling gates to occur and also allowing time for refocussing unwanted coupling effects. The refocussing pulses change the direction in which the couplings are evolving, which can help reach the coupling goals. At the beginning of each event certain couplings must be at their goal values. In most cases, only

those couplings that do not commute with the pulse’s intended effect have a fixed goal target. Other couplings are simply tracked until a fixed goal is required [35]. The coupling evolution for each pair of spins is calculated at each event from both the coupling evolution during the delays and the coupling errors terms in the pulse representations. For those pairs that have a fixed target at this event the calculated coupling is compared to the goal. A euclidian distance function is defined as the sum squared error between the goal and actual couplings and is related to an estimate of fidelity loss for the whole sequence. Optimizing the pulse sequence is now reduced to the task of optimizing the delays between each period. This is handled by a simple iterative optimization to minimize the total contribution to the distance function of the events bounding the delays. That is, the delays are individually optimized one by one, starting at the last one. After the first one is optimized, the sequential optimization starts again at the last one, repeating the process until the improvement is smaller than a threshold, or a goal distance is achieved. Although effective, this optimization strategy is simplistic and easily gets trapped in local minima. It would be useful to develop strategies that optimize all delays together with a non-linear least-squares optimization and consider different distance functions such as maximal error.

The sequence compiler can optimize only the delays between events and it must be given a suitable sequence of refocussing pulses to start with. Designing an exact refocussing scheme may require many refocussing pulses, and each refocussing pulse takes a finite duration and introduces its own errors. In some cases, the theoretical control accuracy gained by bringing the calculated ZZ coupling evolution closer to the goal is lost due to decoherence and pulse imperfections. There is therefore a trade-off between the theoretical accuracy of the control scheme and its duration and number of refocussing pulses. In practice this entails designing a good refocussing scheme and optimizing it. The total error from unrefocussed ZZ couplings is then calculated and judged whether it is acceptably low. If not, additional refocussing pulses are added to correct the errors and improve the optimization. In addition, a penalty function prevents the optimization from using excessive time for refocussing schemes.

2.5 Achieving scalable optimal control

Each method described above has its strengths and limitations. Optimal control theory can give robust and time-optimal control sequences for strongly coupled systems, where conventional pulse design fails. However, the method is intrinsically unscalable and is limited in the number of spins it can handle in practice.

The number of qubits currently available in NMR already pushes these methods to their limit. The pre- and post-error method with pairwise simulations provides a scalable, efficient solution to the design of control sequences, but is limited by the properties of the pre-defined pulses, particularly the fidelity of the error-term representation. On long pulses targeting spins with strong couplings, the decomposition of Eq. 2.24 may fail to give a high fidelity representation of the pulse, and better pulse engineering is needed. Here we show how the two methods can be combined to allow optimal control techniques to be applied to larger systems. A more computer science based approach to scaling up optimal control techniques also utilizing subsystems has also been considered in Ref. [7].

The idea is to consider only a subsystem of the QIP's qubits in designing the optimal control pulse. The subsystem should be chosen such that it encompasses all the qubits affected by the control fields being optimized. In NMR, particularly relevant subsystems consist of the spins of the same nuclear type, for example, one subsystem is all protons and another, all carbons. A pulse that is designed on the subsystem without consideration for the other spins does not implement the desired unitary on the whole system. To determine the effect of the pulse on the entire system, pair-wise simulations between the subsystem and the other spins are performed. These simulations capture and track both the evolution of the internal Hamiltonian on the other spins, and deviations due to couplings between the subsystems. The pulse on the full system can then be represented in the same manner as described above by adding pre- and post-error terms, which can be accounted for during the optimization of the refocusing scheme and the phase tracking calculations as part of a larger sequence.

There is no guarantee that the pulse decomposition with the error terms will give a high fidelity representation. The optimal control pulse drives the subsystem through some complicated trajectory, and it may not be possible to extract the effect of the couplings to other systems as simple ZZ errors before and after the ideal gate. In general this decomposition works well for short one-qubit unitaries but breaks down for two-qubit gates taking more time. However, whereas the pulse is optimized considering only one subsystem, it can be made robust to the effect of couplings to other subsystems by incoherently averaging over the states of the other subsystems as described in Ref. [36]. This can also be seen as ensuring the pulse decouples the considered spins from the rest of the system. This is equivalent to making the pulse robust against Larmor frequency variations. Thus, averaging over all possible states is not necessary, and making the pulse robust to frequency shifts from the sum of the J couplings is sufficient. This will make the pulse more difficult to find, but the resulting pulse will have a higher fidelity representation on the full system.

To illustrate the basic ideas, consider a register of qubits organized into sub-

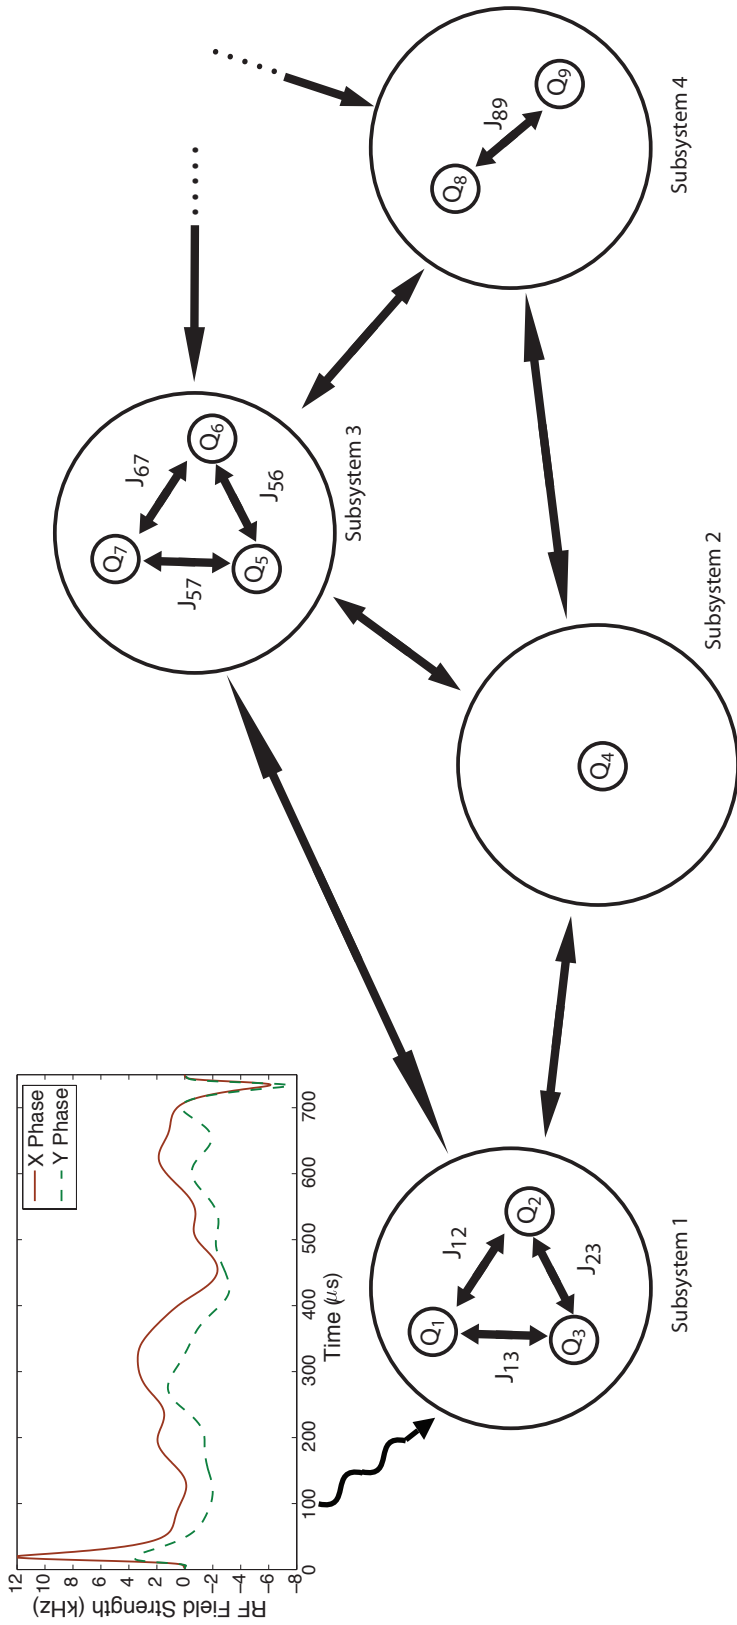


Figure 2.8: Example of a block architecture: For clarity, couplings between qubits of blocks have been contracted into single arrows. It is assumed that the longer range couplings, such as those between subsystem 1 and 4, vanish. For example, in crotonic acid, subsystem 1 could be the proton spins, so that Q_1, Q_2, Q_3 are M, H_1, H_2 , and subsystem 3 the carbon spins and there would be no further subsystems. In the inset an example of a sub-system optimal control pulse, in this case the pulse implements a 180 degree rotation on the methyl qubit of the crotonic acid molecule, is found considering only subsystem 1. The effects of the other subsystems are then characterized using a set of “subsystem” pairwise simulations. As previously, such simulations can be used to compute optimal pre- and post-error terms that enable control of the couplings by optimization of the refocusing scheme and the phase tracking calculations.

system blocks, as illustrated in Figure 2.8. Take for example the first subsystem of Fig. 2.8 consisting of three proton spins, where the other subsystems consist of other nuclear types. It may be that to effect a one-qubit rotation on, say, qubit Q_1 by a simple shaped pulse method requires a very long selective pulse, because the chemical shift difference to the nearest spin is small. Then, the effect of the couplings cannot be taken into account by only Z and ZZ pre- and post-error terms, because $\Delta t J$ is too large. However, one can find an optimal control pulse that considers only subsystem 1 and that refocusses the intra-subsystem couplings while implementing the ideal gate. Perturbations due to inter-subsystem couplings and the internal evolution of the other subsystems during the pulse can then be taken into account by pre- and post-error terms.

The utility of such a scheme can be demonstrated on the crotonic acid molecule where the natural subsystems are the three proton qubits and the four carbon qubits. Because of their refocussing properties, 180 degree rotations are more difficult to find and to represent. However, 180 degree rotations can be found easily for all the spins by use of this subsystem technique and then represented through subsystem pair-wise simulations. For example, a pulse on the methyl qubit is found by considering only the proton subsystem, and then, to determine the dynamics on the full system, simulations are performed on individual and pairs of the five subsystems, $\{M, H_1, H_2\}$, $\{C_1\}$; $\{C_2\}$; $\{C_3\}$; and $\{C_4\}$. The results of the fidelities of the representation to the full seven spin dynamics are summarized in Table 2.2 and compared with the fidelities of the pulse representations of standard pulse designs. The crotonic acid molecule is not ideally suited to this approach because of the large couplings between subsystems; nevertheless, the subsystem GRAPE pulses have similar or better fidelities than standard pulses of a similar or longer length. In particular the GRAPE pulses have consistently high fidelities, even where other pulses, such as selective pulses on H_1 and H_2 , break down. Furthermore, since the clock of decoherence is always ticking, shorter pulses usually perform better.

Finding an optimal control pulse on the entire system such as in Figure 2.3 may give shorter, higher fidelity pulses than the subsystem with error term strategy described above. This is because the full optimal control method is able to exploit more control handles. However, optimization over the full system is too difficult for large systems, whereas the methods used to find the pulses in Table 2.2 are scalable if the subsystems used are kept small. The combination of optimal control and use of pre- and post-error terms is therefore well suited for designing control sequences for the class of QIPs made of subsystems with strong internal couplings but weaker couplings between subsystems.

	isech		Hermite180		GRAPE	
	Length	Fidelity	Length	Fidelity	Length	Fidelity
M	896	99.50	2000	99.96	750	99.92
H_1	3300	91.71	7000	97.71	3000	99.79
H_2	3300	97.42	7000	97.42	3000	99.79
C_1	128	99.92	300	99.97	60	99.95
C_2	700	99.45	1400	99.93	700	99.92
C_3	700	99.39	1400	99.94	700	99.85
C_4	700	99.86	1400	99.96	700	99.88

Table 2.2: Table comparing the performance of optimal control pulses constructed and corrected by using the subsystem approach (last column) to more conventional shaped pulses (first two columns). The pulse lengths are given in μs and the percent fidelities are the Hilbert-Schmidt norm of the simulation of the pulse on the full system including pre- and post-error terms compared to the ideal pulse (see Eq. 2.24). All pulses implement one-qubit 180 degree rotations (see text). The isech shape is an inverse secant and is gaussian like in shape. The Hermite180 is designed to have more of a top hat inversion profile and is closer to a sinc shape [15]. The crotonic acid molecule is not ideally suited to the subsystem approach because of the large couplings between subsystems. This required the effects of the other subsystems be suppressed through incoherent averaging over chemical shifts of the J coupling strength (see text and Ref. [36]) in order to obtain high fidelities.

2.6 Future Improvements

There are several logical improvements to our control schemes mainly of a practical nature. The successful numerical optimization of GRAPE pulses is highly dependent on the initial starting point and it can take many tries to find a high fidelity pulse which is experimentally feasible. Currently this requires an heuristic approach relying on experience and anecdotal results, which is tedious and not always repeatable. Improvements to the code to allow multiple optimizations to be tried in parallel would return a more definitive answer as to whether a certain combination of pulse length and number of timesteps will work and make pulsefinding a more systematic and foolproof process. In addition, the current code can handle up to about 7 spins but parallelized code running on large clusters has been able to find pulses for up to 10 spins [21] and presumably would make finding pulses for seven qubits a much more routine task.

The optimization step in the sequence compiler could also be improved in a number of areas:

1. Designing refocussing schemes is currently done by hand and is tedious. It should be possible to automate this step. However, it would require a reasonably clever algorithm. The couplings are allowed to evolve over many events, and because the refocussing pulses themselves take time and introduce errors, it is a difficult optimization task.
2. The current optimization is local in the pulse sequence, i.e., it only considers the effects of changing a delay on a limited number of events whereas in longer pulse sequences the length of a delay at the beginning of sequence may affect the ability to reach coupling goals near the end of the sequence. In practice this situation is rare because coupling events are usually terminated fairly regularly by pulses.
3. The current optimization is unable to recognize the modular nature of J-couplings in all situations. Because the J-coupling evolution is periodic a ZZ coupling of 0.25 is equivalent to 1.25. Furthermore, as discussed above,

$$e^{-i\frac{\pi}{2}Z_1}e^{-i\frac{\pi}{2}Z_2} = e^{-i\frac{\pi}{2}Z_1Z_2} \quad (2.32)$$

or a ZZ coupling of 0.5 is equivalent to single qubit Z gates on both spins. For example in some situations if the goal coupling is 0.25 it may give a better optimization to let the coupling go to 0.75 and add two Z gates. Not recognizing this makes the optimization task unnecessarily difficult. However, incorporating this into the optimization is difficult as it makes the distance function non-linear and creates local minimum in the parameter space.

I have rewritten the optimization sub-routine to use a non-linear least squares solver which takes into account the modular nature of the goal couplings and implements a penalty function to prevent excessively long sequences. However, it is much slower to run and can become trapped in local minima and so in practice it currently offers little benefit.

Chapter 3

Characterizing quantum dynamics

3.1 The curse of dimensionality in characterization

For many of the same reasons a QIP is reputedly so powerful, it is also difficult to characterize and diagnose what is going wrong when the results are not as expected. The exponentially large Hilbert space in which the QIP operates implies an exponentially large number of experiments are necessary to fully characterize the system. Thus an important avenue of research is to be able to “coarse-grain” or average the information and to efficiently extract fewer relevant parameters. This is of course not an exclusively quantum phenomenon and it is similarly difficult to characterize a large classical process.

The first task in characterizing a quantum system is to determine its state. The quantum state of n qubits can be defined by a density operator: a positive operator with trace equal to unity and of size $2^n \times 2^n$. Any such operator can be decomposed in a basis and a convenient basis is the Pauli basis: the tensor product of n single qubit Pauli operators. Since there are 4 single qubit Pauli operators there are 4^n Pauli operators that need to be specified. Because of the trace constraint there are actually only $4^n - 1$ independent parameters: the identity component is always fixed at $\frac{1}{2^n}$. Quantum state tomography then becomes simply running all the experiments necessary to measure the expectation values of all the n qubit Pauli operators. The state can then be reconstructed through a maximum likelihood algorithm.

In NMR the measurement reveals the simple single minus 1 coherences [14]. Because of phase sensitive detection this can be thought of as measuring $\langle \sigma_x \rangle$ and $\langle \sigma_y \rangle$ for any one spin. Any other Pauli operator can be transformed into one of these terms by a unitary gate and so in a series of experiments each Pauli term can be measured. In some situations it may be possible to reduce the num-

ber of experiments (although it will still grow exponentially with n). In particular for NMR, all homonuclear spins have sufficiently close enough resonant frequencies so that they can all be observed in the same experiment (although if using individual rotating frames appropriate phase correction must be applied). Secondly, if couplings to other spins are resolvable then terms such as XZ as well as XI can be measured in the same experiment. Obtaining more than one term of the density matrix per experiment allows for substantially fewer experiments. For example, with three homonuclear spins, the set of seven readout pulses $Y_{90}II, IYY_{90}, IIX_{90}, Y_{90}Y_{90}I, Y_{90}Y_{90}Y_{90}, X_{90}X_{90}X_{90}$ are sufficient to observe all the terms of a three qubit density matrix, as opposed to the naive 63 experiments. This can be generalized and an algorithm has been presented for determining the necessary readout pulses [37].

To characterize a quantum process the situation is even worse. There is a prescribed procedure for quantum process tomography [38] but it is obvious what the task is. The quantum process can be characterized in its Liouville representation. Given an orthonormal basis for n qubit density matrices, P_i (again the Pauli basis is common), then the map Λ can be represented as a matrix $\Lambda_{\alpha\beta}$,

$$\Lambda_{\alpha\beta} = \text{Tr} (P_{\alpha}^{\dagger} \Lambda (P_{\beta})) . \quad (3.1)$$

This representation can then be converted to any other desired [39]. The experimental procedure then is to input all 4^n density matrices into the process and then perform state tomography on the output. The quantum process can then be reconstructed through a maximum likelihood algorithm. This reconstruction can be tricky in the presence of imperfect readout pulses and in many cases appears to be numerically ill-conditioned [40]. Although there will be some constraints if the process is limited to be completely positive (CP), there will still be $\mathcal{O}(2^{4n})$ parameters to be determined. It should be noted that given other resources, such as a perfect side channel to allow some qubits to be unaffected by the noise, a quadratic reduction in the number of experiments is possible [41].

3.2 Characterizing complex quantum dynamics

One such coarse-grained measurement which is useful is the distinction between chaotic and regular dynamics. Presumably it would make work unnecessarily difficult to try and build a QIP in a chaotic system. Of course there is no such thing as chaotic evolution in a closed quantum system: the Schrodinger equation governing its evolution perfectly preserves the distance between input states and so there is no equivalent of the classical idea of two points in phase space being driven exponentially far apart by chaotic evolution. Nevertheless, one can ask whether

systems that are chaotic in their classical limit exhibit any signs of this behaviour as a quantum system. This is the search for quantum chaos.

Fidelity decay under perturbation, initially proposed by Peres [42], is an important analogue to the classical notion of instability, and has received much study. Rather than considering the overlap of two nearby states evolving under the same evolution, this method considers the overlap of the same initial state evolved under two nearby unitary maps. If the discrete time evolution is given by the map U , then we also consider the map $U_p = UP$ where $P = \exp(-iV)$ for some hermitian matrix V . Then, the fidelity decay after n steps is given by,

$$F_n(\psi) = \left| \langle \psi | (U^n)^\dagger U_p^n | \psi \rangle \right|^2. \quad (3.2)$$

The analogy with classical mechanics suggests that chaotic systems might show an exponential fidelity decay, while regular systems would decay at a slower (e.g. polynomial) rate. However, numerical studies of small systems show a much more complicated situation, with different behaviours depending on the form and strength of the perturbation [43, 44, 45]. Yet, for a range of sufficiently strong perturbations, chaotic systems show a universal exponential fidelity decay at an average rate which is given by the Fermi Golden Rule (FGR) [45, 46]. The decay rate depends on the statistics of the system eigenvectors in the eigenbasis of the perturbation, or vice-versa how random the perturbation looks in the eigenbasis of the system. For a random map, almost all perturbations will look random and thus the average decay is universal, in the sense that it is independent of the system dynamics and depends only on the strength of the perturbation. Specifically, the rate of the universal exponential decay depends only on the variance of the eigenvalues of the perturbation (i.e., the 2-norm of the perturbation, $\|V\|_2$) and the exponential decay continues until a saturation level $\mathcal{O}(\frac{1}{N})$, due to the finite system dimension N . However, for a specific integrable system, a well-chosen perturbation will show structure in the system's eigenbasis, leading to large fluctuations and recurrences in the fidelity decay. Hence, direct measurement of the fidelity decay under different applied perturbations can provide important information about the complexity of the system dynamics.

Given some dynamical system under study, the fidelity decay measured for any one particular initial state will fluctuate from the average - an effect particularly pronounced with small systems. Hence estimating the average decay can require averaging over many different initial states. However, Poulin et al. [47] worked out a quantum circuit (see Fig. 3.1) for directly measuring the average fidelity decay within the deterministic quantum computation with a single bit (DQC1) model [48]. Using the identity that,

$$\overline{\langle A \rangle_\psi \langle B \rangle_\psi} = \frac{\text{Tr}(A)\text{Tr}(B) + \text{Tr}(AB)}{N^2 + N}, \quad (3.3)$$

where $\langle A \rangle_\psi$ denotes $\langle \psi | A | \psi \rangle$ and the overline an average (Haar measure) over states $|\psi\rangle$. Then the fidelity decay after n steps, averaged over a uniform (Haar) measure of the initial states $|\psi\rangle$ (see Eq. 3.2) for an arbitrary system of dimension N takes the general form,

$$\overline{F(n)} = \frac{\left| \text{Tr} \left\{ (U^n)^\dagger (PU)^n \right\} \right|^2 + N}{N^2 + N}. \quad (3.4)$$

Therefore, experimental determination of the average fidelity decay for a particular U and P requires measuring the trace, for which there exists an efficient DQC1 circuit [49].




Figure 3.1: Ideal quantum circuit for measuring fidelity decay after n steps. The top qubit in a pure (or pseudo-pure) state can be considered either as a probe of the bottom system of K qubits which starts in the maximally mixed state or as a toy system being decohered by the maximally mixed environment below. Making the final rotation about the x (y) basis gives the real (imaginary) part of the trace in Eq. 3.4. Taken from Ref. [47].

The circuit shown in Figure 3.1 gives the necessary steps. The circuit has been simplified with a few steps that may not be obvious at first. Initially everything in the trace of equation 3.4 is controlled on the state of the first qubit. However, the controls on all the U gates (U and U^\dagger) can be removed because in the case that the control is in the state $|0\rangle$ where we do not want anything executed on the bottom registered, the still controlled P 's will not run and the U 's and U^\dagger 's will cancel each other. In the case of a $|1\rangle$ on the first qubit everything will be executed as desired. However, the final n U^\dagger gates are not necessary as they will have no effect on the state of the first qubit.

For this project I implemented the fidelity decay measurement in liquid state NMR. We used our liquid-state work-horse crotonic acid molecule. The solvent was deuterated acetone. The spectrometer was a 700MHz Bruker Avance II system with a TXI probe.

Since very good decoupling of the probe qubit from the rest of the system is needed, I used the hetero-nuclear system with the methyl qubit as the probe and

the four carbon backbone as the system being investigated. The probe could then be refocussed with very good robust composite 180 pulses designed by Dr Knill. Each 180 is replaced with five 180's so that a 180 degree rotation about the X axis is replaced with:

$$R_0(180) = R_{\frac{\pi}{6}}(180)R_0(180)R_{\frac{\pi}{2}}(180)R_0(180)R_{\frac{\pi}{6}}(180). \quad (3.5)$$

There is also an additional 60 degree Z rotation associated with the composite pulse which can be handled with phase tracking. The pulse is somewhat robust to calibration errors and is particularly robust to off-resonant errors so that it is good for applying a high fidelity 180 degree pulse to all the protons simultaneously. Comparisons of the performance of the pulse to simple rectangular pulse and the BB1 composite pulse are shown in Figure 3.2. The other two protons H_1 and H_2 were placed in a pseudo pure state at the beginning of the computation so that they had no effect on the rest of the molecule (besides an additional chemical shift which can be tracked).

Because the natural Hamiltonian in NMR is simple and regular we had to simulate complex evolution using our universal control to generate arbitrary dynamics. We investigated the difference in fidelity decay response for regular evolution (natural Hamiltonian) and complex evolution which is simulated by a pseudo-random map. There is strong numerical evidence for the ability of random maps to match the statistics of chaotic maps [50]. Generating a truly random map is exponentially hard in the number of qubits, but pseudo-random operators are able to approximate a fully random map well. We used a decomposition for pseudo-random operators based on repeated applications of a two step process: (1) Individual random rotations are applied to each qubit; (2) Simultaneous two-body interactions between neighbours given by the unitary, $U = \exp \left[i \left(\frac{\pi}{4} \right) \sum_{j=1}^{n-1} \sigma_z^j \otimes \sigma_z^{j+1} \right]$ [51]. For a finite number of repetitions of these two steps, the resulting distribution of maps is biased with respect to the uniform measure; however, it converges exponentially to the uniform measure [52]. For our experiments we used four rounds of this two step process which numerically was found to generate a sufficiently random map for seeing an average fidelity decay.

In practice performing individual single qubit rotations for the random maps would require several long “soft” pulses and it would be difficult to refocus all the other qubits. Fortunately, I found that because we are trying to implement a random map, extra coupling dynamics are not an issue and the refocussing is not necessary. Furthermore, exactly implementing the single qubit rotations is not necessary. Because the carbon spectrum is so wide, a hard pulse of random length between $10 - 34\mu\text{s}$ followed by a short random delay of 1-3ms to introduce a random Z rotation from the chemical shift followed by another random hard pulse is sufficient to perform random single qubit rotations on all the carbons. The




Figure 3.2: The robustness of the BB1 (dotted-green) and Dr Knill’s (dashed-red) composite 180 pulses to both calibration errors (a) and off-resonance errors (b). The plots show the average fidelity (HS norm) with the ideal operation (with a 60 degree Z pulse correction for Dr Knill’s 180) from simulating the single spin dynamics. Also shown is a simple hard rectangular pulse (solid-blue) for comparison purposes. The off-resonance x-axis is plotted in units of the nutation frequency of the pulse so it is independent of the pulse length. The simulations assume no delays between the pulses of the composite pulses.

couplings within the carbon sub-system are not refocussed but instead contribute to the second step of the random map iteration. Secondly exactly implementing a ZZ coupling of 0.25 between all nearest neighbours is too demanding and not necessary. Instead, one of the pairs was chosen at random and a ZZ 0.25 coupling implemented for that pair. The other couplings will still sufficiently entangle the qubits to generate a random map.

The controlled perturbation was implemented using the natural Hamiltonian of the molecule that provides a coupling between the probe and the system qubits of the form: $\sum_{j=1}^4 J_{MC_i} \sigma_z^M \otimes \sigma_z^{C_i}$. By changing reference frames (or applying a Z-rotation) to one of the qubits, this can be seen as a controlled rotation of the system qubits about the z axis depending on the state of the probe qubit. Therefore, a controlled operation on the system qubits can be implemented by allowing the natural evolution for a time proportional to the desired strength of the perturbation. This perturbation can be transformed into a rotation about another axis, by sandwiching the evolution with two rotations of the target qubits. This will allow us to test how the fidelity decay response depends on the relationship between the system Hamiltonian and the perturbation.

Different pseudo-random maps can be generated by varying the parameters of the random single qubit pulses and delay times which implement the coupling gate. This allows an averaging over random maps which is important because of the fluctuations in the fidelity decay due to the finite size of the system. The average fidelity decay depends only on the relative randomness between the system and perturbation. But, the fidelity decay for any particular map will fluctuate (an effect independent of the initial state which is already averaged over by this circuit implementation). These fluctuations disappear as the size of the system increases beyond a few qubits, but for small systems, this effect is pronounced and requires an averaging over different random maps to obtain good statistics. To quantify the fluctuations, numerical simulations of the fidelity decay experiment were performed over 1000 different random maps created assuming perfect control. Averages and deviations were then calculated to compare with the experimental results (see Figure 3.3). This allows us to say whether the deviation from the average exponential decay we see is a signature of regular dynamics or is simply an expected fluctuation.

The experiment started with polarization on the methyl qubit. The thermal polarization on the other qubits was dephased using gradient pulses. The spin- $\frac{1}{2}$ subspace of the three equivalent protons was then selected using a specialized pulse sequence [53]. A three qubit labelled pseudo-pure state was then created using two repeats of the following steps to create a two-qubit labelled pseudo-pure state between methyl and C_1 . Using the product operator notation and starting with the polarization on the first qubit:

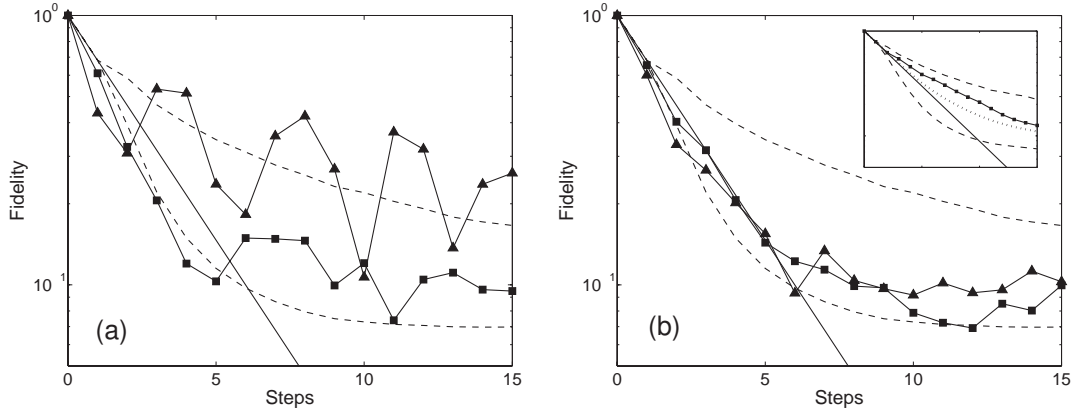


Figure 3.3: Comparison of fidelity decay curves for regular (a) versus chaotic (b) evolution under different forms of perturbation. The natural evolution of the molecule provided the regular system and a pseudo-random map was used to model chaotic evolution. The perturbation form was controlled by the presence or absence of conjugate rotations on either side of the coupling interaction between the probe and system. In the case of regular evolution, when the perturbation is a rotation about the z-axis (\blacktriangle curves), which commutes with a system coordinate, the fidelity decay shows substantial fluctuations away from the average. Changing the perturbation to a rotation about the x-axis (\blacksquare curves) substantially alters the form of the decay. In the chaotic case, the evolution looks random in the eigenbasis of either perturbation and the decays exhibit the universality of the fidelity decay for complex dynamics. The FGR result (solid curves) and standard deviations (dashed lines) calculated from 1000 numerical simulations assuming perfect control are also shown. The inset of figure (b) shows the experimental average fidelity decay from twenty different random maps. The dotted line show the corresponding average from the numerical simulations and we can see that deviation from the exponential decay is also present in the numerical results and well explained by the finite size saturation level.

$$\begin{aligned}
& ZI \\
& \downarrow R_{\pi/4}(\pi/2)^1 \\
& \frac{1}{\sqrt{2}}(XI - YI) \\
& \downarrow R_{\pi}(\pi/2)^1 \\
& \frac{1}{\sqrt{2}}(XI + ZI) \\
& \downarrow J \text{ coupling for time } \frac{1}{2J} \\
& \frac{1}{\sqrt{2}}(YZ + ZI) \\
& \downarrow R_{\pi/2}(\pi/2)^1 \\
& \frac{1}{\sqrt{2}}(YZ + XI) \\
& \downarrow R_{-\pi/4}(\pi/2)^1 \\
& \frac{1}{\sqrt{2}} \left(\frac{1}{\sqrt{2}} \left(ZZ + \frac{1}{\sqrt{2}}(YZ - XZ) \right) + \frac{1}{\sqrt{2}} \left(ZI + \frac{1}{\sqrt{2}}(XI + YI) \right) \right) \\
& \downarrow \text{crusher gradient} \\
& \frac{1}{2}(ZI + ZZ) = \frac{1}{2}Z(I + Z) = \frac{1}{2}Z0
\end{aligned}$$

Figure 3.4: Pulse sequence for creating a two-qubit labelled pseudo-pure state. The notation $R_{\phi}(\theta)^k$ specifies a rotation of angle θ about an axis ϕ from the x -axis on qubit k .

After the pseudo-pure was created on C_1 , it was then moved out to either H_1 or H_2 using state specific swap gates. The fidelity decay protocol was then run and the fidelity decay was measured from the real and imaginary part of the integral of the methyl group peaks. The baseline for the integral is distorted from a large impurity peak only a few hundred Hz away. This was somewhat suppressed by phase cycling the experiment.

The results of Fig. 3.3 also show how the form of the perturbation effects the decay. The pseudo-random map shows a universal response, in that the decays are identical under different perturbations. On the other hand, the natural evolution's decay varies wildly. This result demonstrates the necessity of choosing the perturbation carefully in the context of distinguishing between regular and chaotic evolution: fidelity decay will provide useful information only if the applied perturbation commutes with the system's coordinate [45] i.e. the dominant term in the Hamiltonian.

These experiments also highlight the relevance of these techniques to decoherence studies. For this purpose, we consider the pseudo-pure qubit to be the system we are interested in and the maximally mixed bottom register, the environment.

The decoherence rate of the system is typically governed by some *macroscopic* parameters of the environment, such as its temperature, cutoff frequency, etc., see e.g. Refs. [54, 55]. Recently, the importance of the *dynamics* of the environment has been expressed by a few authors [47, 56], and these experiments constitute a direct demonstration of these effects because we can directly control the complexity of the environment’s dynamics. This is because the environment is no longer an uncontrolled degree of freedom but is rather the bottom register of qubits. We can explicitly tune the dynamics of these qubits between regular and complex dynamics using our universal control as demonstrated above.

The work in this section was published in Ref. [57].

3.3 Characterizing time suspension sequences in solid-state NMR

Randomization has been proposed as a useful protocol in noise characterization [51, 46]. Through appropriate randomizations the noise can be symmetrized and coarse-grained quantities of the noise extracted. The protocols are based on the ideas of twirling. The quantum process to be characterized is conjugated by a random unitary. The original noise map maps a density matrix ρ to $\Lambda(\rho)$. Under a S -twirl, a unitary U from the group S is chosen, applied to ρ , then the noise map is applied and then U is undone with U^\dagger :

$$\rho \rightarrow \int_S U^\dagger \Lambda(U\rho U^\dagger) U. \quad (3.6)$$

Λ can be written in a Kraus decomposition as $\Lambda(\rho) = \sum_k A_k \rho A_k^\dagger$. Then the twirled map can be written as (also exchanging the order of integration and summation):

$$\rho \rightarrow \sum_k \int_S U^\dagger A_k U \rho U^\dagger A_k^\dagger U. \quad (3.7)$$

The A_k can be decomposed in the Pauli basis,

$$\rho \rightarrow \sum_{i,j} \alpha_{ij} \int_S U^\dagger P_i U \rho U^\dagger P_j^\dagger U. \quad (3.8)$$

The chi matrix α_{ij} describes the superoperator and the effect of the twirling depends on what group S is and can be stated in terms of how it affects the chi matrix. If S are fully Haar distributed random unitaries then the effect is to completely depolarize the noise, i.e., all terms where $i \neq j$ will vanish and all diagonal terms for traceless Pauli’s will have equal weight. Because for any ρ , $\frac{1}{4^n} \sum_j P_j \rho P_j = \frac{1}{2^n} \mathbb{1}$,

twirling with Haar random unitaries will turn any map into a depolarizing map with a single parameter d describing the noise strength $\rho \rightarrow (1 - d)\rho + \frac{d}{2^n} \mathbb{1}$ [46]. Haar random unitaries are powerful depolarizers and gain from the concentration of measure in large Hilbert spaces. However, they are also inefficient to create in that it takes an exponential number of parameters to specify and thus create a truly Haar random unitary. One work around is to use pseudo-random unitaries from the previous section, where with a polynomial number of simple one and two qubit gates, a distribution of unitaries exponentially close to the Haar measure can be achieved [51]. Fortunately, decomposing the A_k as Pauli's makes another elegant solution clear. A random element from the n qubit Clifford group will map any n qubit Pauli to another evenly distributed random n qubit Pauli. Thus it is clear from Eq. 3.8 twirling with the full n qubit Clifford group will also depolarize the noise map [58]. Full Clifford twirling is still difficult as the size of the Clifford group grows exponentially with n . However, full Clifford twirling is not necessary as a Chernoff-Hoeffding bound argument [59] shows that sampling from the Clifford group is sufficient to efficiently estimate many parameters of the twirled map to any desired accuracy.

Another interesting twirl is a $C^{\otimes n}$ -twirl - a tensor product of n single qubit Cliffords. The effect of this twirl [60] is to average α_{ij} to zero for $i \neq j$ and when combined with random qubit permutations to equalize all the α_{ii} of the same Pauli weight, i.e., the terms which have the same number of non-identity Pauli's. Since this Pauli weight can be associated with the error weight (i.e. how many qubits are affected by the error) it is an interesting and useful thing to measure the probability of different error weights p_w . For example, most quantum error correction codes are formulated in the stabilizer formalism [61]. An important metric is the distance of the code which determines up to what weight Pauli error the code can correct. Thus measuring the probabilities of Pauli error weight would in principle determine the failure probability of the code. If the code had distance $2t + 1$ so that it could correct errors of weight t , then the failure probability for that code would be $\sum_{w>t} p_w$. Furthermore, measuring the p_w would allow some tests of the independence of the errors. Under an independent error model we would expect $p_w \propto p_1^w$. Although confirming this proportionality experimentally does not confirm this error model, deviations from it confirm the error model is not independent.

The twirling procedure performs the necessary coarse-graining of the map reducing the number of parameters depending on the form of the twirl. An experimental procedure to measure these reduced number of parameters in the case of the $C^{\otimes n}$ -twirl is detailed below:

1. Prepare eigenstates of Pauli operators. Because of the Clifford gates used for twirling, the particular direction of the Pauli chosen does not matter i.e.

the states could all be eigenstates of σ_z . However, the weights of the Pauli operators and the locations of the non-identity terms should be randomly chosen.

2. Send these states through a $C^{\otimes n}$ twirled version of the channel. That is randomly choose a set of single qubit Cliffords, apply them, apply the channel to be characterized, and then apply the inverse of the Cliffords. Every time a state is sent through, a new set of random Cliffords is chosen.
3. Measure the correlation of the output state with the initial input state.

After repeating this procedure sufficiently many times, sufficient statistics can be accumulated to get estimates of c_w , where c_w is the average correlation between the input and output state for a Pauli operator of weight w . The accuracy of the estimate can be determined from a Chernoff-Hoeffding bound but the number of experiments necessary to estimate c_w to an accuracy ϵ is only $\mathcal{O}(\epsilon^2)$. The p_w are linearly related to the c_w . The relationship can be worked out by exhaustively counting the effects of all possible errors. As a relevant example consider a three qubit system. When inputting a eigenstate of a weight 1 Pauli operator, say ZII , then to work out c_1 we have to work out which error terms commute and anti-commute with the state. There are 9 weight 1 error terms $\{XII, YII, ZII, IXI, \dots, IIZ\}$. After the twirling they all have the same probability $\frac{1}{9}p_1$. Only two of these errors (XII, YII) will invert the sign this input state (and indeed for any weight 1 input state only two weight 1 errors will flip the sign of the state). Similarly, there are 27 error terms of weight 2, and 12 of them will flip the sign of the state ZII . And finally of the 27 error terms of weight 3, 18 will flip the sign of ZII . Thus the correlation of the output state with the input state for eigenstates of Pauli terms of weight 1 will be,

$$\begin{aligned} c_1 &= p_0 + p_1 \frac{1}{9} (-2 + 7) + p_2 \frac{1}{27} (-12 + 15) + p_3 \frac{1}{27} (-18 + 9) \\ &= p_0 + \frac{5}{9}p_1 + \frac{1}{9}p_2 - \frac{1}{3}p_3. \end{aligned} \tag{3.9}$$

The equations for c_2 and c_3 can be worked out in a similar manner. Under unital dynamics (and the twirled map is always unital) the identity input state is mapped to the identity by definition, so it is assumed that $c_0 = 1$. Thus,

$$\begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & \frac{5}{9} & \frac{1}{9} & -\frac{1}{3} \\ 1 & \frac{1}{9} & -\frac{5}{27} & \frac{1}{9} \\ 1 & -\frac{1}{3} & \frac{1}{9} & -\frac{1}{27} \end{bmatrix} \begin{bmatrix} p_0 \\ p_1 \\ p_2 \\ p_3 \end{bmatrix}. \tag{3.10}$$

This equation can be inverted to determine the p 's after measuring the c 's. Unfortunately, the matrix relating the p 's and c 's is not well conditioned and uncertainties

in the c 's from both the incomplete twirling and experimental uncertainty can blow up into exponentially large errors in the p 's for large w . This limits the weight of the Pauli errors which can be efficiently estimated. Nevertheless, on physical grounds we expect high-weight errors to decay beyond some level.

The scheme is ideally suited to measuring the identity map or say a quantum communication channel where nothing is supposed to happen. In NMR QIP the Hamiltonian contains always-on couplings which need to be refocused or decoupled during wait steps. In solid-state NMR these couplings are dominated by the dipolar couplings which have the form $\mathcal{H}_{DD} = \sum_{i < j} D_{ij} (2Z_i Z_j - X_i X_j - Y_i Y_j)$. Because of the asymmetry and different signs in this Hamiltonian, it can be averaged to zero through an appropriate pulse sequence. A powerful body of techniques under the terms average Hamiltonian theory and the Magnus expansion have been developed to analyse such sequences [31, 62].

A family of sequences are known as time suspension sequences as they suppress both the chemical shift and the dipolar terms and so aim for an average Hamiltonian of zero. Despite the impressive theoretical tools available, it is still difficult to determine in a real experiment what the form of the errors will be in a given time suspension sequence, and so it is useful to characterize these sequences experimentally. Although finite pulse effects and some models of phase transients can be compensated for, Figure 3.5 shows that the real pulse is far from an ideal rectangular shape. These deviations can somewhat be nulled out through standard tune-up procedures to adjust the probe matching [63]. Even after doing this though deviations remain. The phase transients are difficult to model as well. Furthermore, the model for the system used in the average Hamiltonian calculation may not capture the real dynamics due to an inaccurate model or parts explicitly ignored, e.g. decoupled spins. In practical terms an NMR spectroscopist would ask questions like: “How much signal is left after the sequence?” or “How narrow is the line under stroboscopic observation?” However, using the symmetrized characterization approach more detailed and generalizable information is available than these state to state fidelity measures.

The Cory48 sequence is a time suspension sequence with not surprisingly 48 pulses [64]. The sequence suppresses the dipolar term to second order in the Magnus expansion, the chemical shift to zeroth order and the cross terms between the dipolar and the chemical shift to first order. So, we would expect that it is quite effective at suppressing weight two errors from dipolar couplings but may be less effective at suppressing weight 1 errors from chemical shifts. The sequence consists of 48 pulses separated by delays τ or 2τ as a series of two-pulse solid echos (see Figure 3.6). The pulse spacing τ defines the averaging rate. Stronger couplings must be averaged at a faster rate for effective suppression. Thus we would also expect an error rate that depends on τ . For a given amount of time as τ decreases




Figure 3.5: The envelope of a nominal $1.3\mu\text{s}$, 75MHz pulse measured at the sample. The r.f. field at the sample is measured via a pickup coil and digitized on a high-speed oscilloscope. The r.f. is then digitally demodulated and in-phase and quadrature components determined. The phase is adjusted to maximize the integral of the in-phase component. The pulse was recorded on a home-built Q-spoiled probe and the distortions would be worse on a higher-Q probe. The Q or quality factor of the probe measures the losses in the resonant circuit. It can be interpreted in a number of ways but a relevant measure is the resonant frequency divided by full-width half-max (linewidth) of the resonance. The higher the Q the narrower the linewidth and the narrower the bandwidth of the probe which makes transient effects worse. The probe was tuned and matched in a standard manner to minimize the reflected power. However, different probe configurations have an effect on the form and magnitude of the phase transients and the probe can be tuned to minimize them by tuning to a point off-resonance (at the expense of coupling efficiency and therefore nutation frequency).

the error rate should decrease as the averaging improves, until these improvements are overcome from the errors from finite pulse widths and pulse imperfections from the increased number of pulses [65]. More interestingly, the error rate can increase as τ decreases if another process on a faster time scale starts to interfere with the averaging of the time suspension sequence.




Figure 3.6: The first four pulses of the Cory48 time-suspension sequence. The pulses are 90 degree pulses about the phase specified. The actual delays between the pulses are reduced to take into account the finite width of the pulses. Also shown are the phases of the full 48 pulses of the sequence.

The experiment is shown in circuit form in Figure 3.7. It was performed on our three qubit solid-state molecule of choice - malonic acid (see Chapter 4 for details of the typical Hamiltonian). The first step is to depolarize the carbons and then apply a reset sequence which transfers polarization from the protons to one of the carbons (see Chapter 4). The signal is then “cleaned-up” by rotating the polarization on the other two carbons into the plane and allowing the uncoupled protons dipolar fields to dephase them. This gives an approximate initial state of IIZ ($C_1C_2C_m$). IIZ is not an eigenstate of the carbon Hamiltonian and the dipolar flip-flop terms will also create small amounts of ZII and IZI . One of three states $\{IIZ, IZZ, ZZZ\}$ is then created at random via a state-to-state optimized transformation. These are the eigenstates of Pauli operators of weights 1, 2 and 3. Then a random permutation of the qubits is applied via unitary SWAP gates (the SWAP gate simply exchanges the state of two qubits; if followed by numbers it specifies which qubits are exchanged, e.g. SWAP13 exchanges the state of qubits 1 and 3). Then 3 randomly chosen single qubit Cliffords are applied in series. This eases the experimental burden by requiring only $3 \times 12 = 36$ pulses, rather than $12^3 = 1728$ pulses, to be found and corrected for implementation errors (see Appendix C). The expense is the additional errors from doing three gates versus one parallel gate. The channel to be characterized is then applied. The inverse of the Cliffords, SWAPs and state transformations are applied and finally a readout pulse is applied to determine the fidelity of the output state by both scaling and integration and comparison to a reference spectrum.




Figure 3.7: The circuit for the symmetrized characterization of a quantum process. The initial state is the weight one deviation density matrix IIZ . This is randomly transformed into weight 2 (IZZ) or weight 3 (ZZZ) state. The location of the non-identity terms is then randomized via a randomly chosen SWAP gate. Three single qubit Cliffords (C_{rand}) are applied in series. The channel to be characterized Λ is then applied. Finally all the steps are reversed and the polarization is read out on the first qubit.

The choice of the reference spectrum allows us to normalize out errors in the state preparation and twirling sequences. There are two steps to the normalization procedure. Preparing the weight 2 and 3 Pauli operators require different preparation sequences whereas the weight 1 IIZ requires no additional state preparation. Therefore individual reference spectra are used for each Pauli weight by running the state preparation and its inverse for each weight. Secondly, by running the twirling protocol without any channel, we obtain three spectra for the different Pauli weight input states. One can imagine commuting all the errors in the SWAP and Clifford pulses into the middle which then becomes the channel being characterized. This is not entirely accurate because the errors will be pulse dependent. However, this gives us some information about the level of control and we can use the spectra obtained as references for characterizing other channels. Indeed, the fact that the protocol even works in NMR with its highly mixed input state (which can be interpreted as huge errors on the preparation step) shows how robust this normalization feature is. The no-channel gave c 's of $\{c_0 = 1; c_1 = 0.96; c_2 = 0.93; c_3 = 0.92\}$ which corresponds to a channel with p 's of $\{p_0 = 0.93; p_1 = 0.04; p_2 = 0.03; p_3 = 0\}$. Since there are 8 gates (2 SWAPS and 6 Clifford gates) it is reasonable to claim an error per gate of approximately 1%. All subsequent experiments are then normalized with respect to these spectra to determine the additional error coming from the channel.

The SWAP and Clifford gates can be chosen at random. However, because the input state is known, only certain Clifford combinations will have a useful impact on the twirling for particular input states. There are 1728 possible Clifford gate combinations which make exhaustive averaging too time consuming. It therefore makes sense to choose the random gate combinations so that they implement the

most effective twirl. Consider the weight 1 input state ZII . Only the SWAP12 and SWAP13 perform a non-trivial operation and are thus worth doing. The other SWAP can be used as a dummy gate for when no SWAP is called for to keep the number of gates consistent. If the polarization is left on the first qubit (state ZII) then changing the Clifford gates of qubit 2 and 3 have no effect and so can safely be left out of the averaging. Thus each of the possible weight 1 Pauli input states (ZII, IZI, IIZ) have only 12 Clifford gate combinations to be averaged over for a total of 36 experiments. Similarly the weight 2 Pauli inputs have 144 possible combinations that may have an effect on the twirling. We choose rather arbitrarily to perform 432 experiments, all 144 useful experiments for each of the three weight two Pauli inputs.

I characterized three channels. Firstly, I engineered noise consisting of a single phase flip on a random qubit by modifying the phase of subsequent pulses. This a weight 1 error happening with 100% probability so that we expect the measurement to return p 's $\{p_0 = 0; p_1 = 1; p_2 = 0; p_3 = 0\}$. Secondly we characterized two time suspension sequences. The first was 48 pulses of the Cory48 sequence with a $10\mu s$ pulse spacing. The second was 96 pulses of the Cory48 sequence with $5\mu s$ pulse spacing. The total time of the two sequences is the same so that the comparison reveals information about the performance of the sequence as a function of pulse spacing. The results from all three experiments are summarized in Table 3.1 along with results from simulations. The simulations are of the carbon sub-system only and assume unitary evolution. They reveal expected deviations due to incomplete twirling, imperfect time-suspension suppression of evolution or pulse dependent unitary errors.

	p_0		p_1		p_2		p_3	
Map	Sim.	Exp.	Sim.	Exp.	Sim.	Exp.	Sim.	Exp.
Engineered	0	$0.01^{+0.01}_{-0.01}$	1	$0.99^{+0.01}_{-0.03}$	0	$0.01^{+0.02}_{-0.01}$	0	$0.00^{+0.01}$
Cory48 $10\mu s$	0.47	$0.44^{+0.01}_{-0.02}$	0.42	$0.45^{+0.03}_{-0.03}$	0.12	$0.10^{+0.04}_{-0.08}$	0.00	$0.01^{+0.03}_{-0.01}$
Cory48 $5\mu s$	0.92	$0.84^{+0.01}_{-0.01}$	0.07	$0.15^{0.02}_{-0.03}$	0.01	$0.01^{+0.03}_{-0.01}$	0.01	$0.00^{+0.02}$

Table 3.1: Simulated and experimentally measured p values for the three maps implemented. Uncertainties for the experimentally measured values come from a maximum likelihood algorithm.

The experimental results match the simulations well for the engineered error and the Cory48 sequence with $10\mu s$ spacing. There is a larger discrepancy though between the simulated and experimental results for the Cory48 with $5\mu s$ pulse spacing which warrants more explanation. Fortunately, because the input state for each experiment is known a more detailed analysis of the location of the errors

is possible. For the weight 1 Pauli input state the state is one of ZII, IZI, IIZ after the SWAP gate. By adding up all the experiments corresponding to the polarization being on each qubit separately we can determine the location of the errors. In our case, the results for ZII and IZI were nearly perfect where as the result for IIZ (polarization on the methyl carbon) was approximately $2/3$ of the expected amount. Similarly for the weight two Pauli inputs, when there was polarization on C_m the signal was much less than expected. This suggested that all the errors were occurring on the C_m qubit. The physical reason behind this is a breakdown of the proton decoupling. C_m has approximately an order of magnitude stronger coupling to the protons than the other two carbons (see Section 4 and Figure 4.3 for details of the proton-carbon Hamiltonian) and so we would expect it to be most adversely affected by any proton decoupling inadequacies. After the initial polarization transfer from the reset step, the protons are decoupled with the efficient SPINAL64 sequence [66]. The SPINAL64 is a windowless phase modulated sequence which is a series of rotations of $\approx 165^\circ$. With the available proton power this corresponds to a pulse length of $1.5\mu s$. With the $5\mu s$ pulse spacing there is actually only $3.7\mu s$ between carbon pulses due to the finite pulse width. The proton decoupling works by modulating the proton dipolar fields the carbons experience at a rate much faster than the coupling to produce zero average Hamiltonian. However, if the carbons themselves are being modulated at a comparable rate, then the designed average Hamiltonian from the decoupling no longer works and the decoupling breaks down. Further confirmation of this explanation is found from the fact that the experiment improves p_0 from 0.76 to 0.84 by decreasing the proton pulse time for the $\approx 165^\circ$ pulse from $2.4\mu s$ to the $1.5\mu s$ used in the final experiment. Further increases in decoupling power are prevented because of the onset of arcing in the probe capacitors.

These experiments confirm that symmetrized characterization techniques can extract useful information about a quantum process. In such a small and well characterized system the results are perhaps not terribly surprising and unexpected. Nevertheless the technique provides a systematic approach which will scale to many qubits. The work on this symmetrized characterization was published in Ref. [60].

3.4 Randomized benchmarking of quantum control sequences

The above protocol is very effective for characterizing the identity or do-nothing gate; it is best suited to characterizing quantum channels where the ideal transformation is perfect unadulterated transmission rather than quantum gates where we need something to happen. To overcome this one could imagine doing and then

undoing the gate and then using the above protocol to characterize the combined error. This will work well in the case of a pulse independent error, say in the case where the errors are decoherence limited. However, in most realistic scenarios the control fields themselves will introduce errors and the errors will be coherent unitary rotations that depend on the gate being implemented. An obvious example is a miscalibration error. If the power was consistently too high then doing and undoing the gate will still do a perfect identity gate even though the fidelity of the gates themselves is very low.

Benchmarking the level of control for gates is important for several reasons. It is useful to establish whether fault-tolerant control can be achieved. Using the tools of quantum error correction and certain assumptions about the resources available and the error model, it is possible to show that if the error in the computation is sufficiently low then arbitrarily long computation is possible [2]: that is the error rate is low enough that the gain from performing error correction is not overcome by the increased number of gates. The “error sufficiently low” is usually reported as a error per gate and is reasonably thought to be around 10^{-4} , although with some perhaps unrealistic assumptions about the amount of resources available, it could be as high as 10^{-2} [67]. An experimental protocol is needed to efficiently determine whether a given device reaches that level of control.

Agreeing on a standard benchmarking of quantum control would allow a fair comparison between different devices and different technologies. The protocol should return a single (or very few) numbers which accurately represent the computationally relevant error per gate and are useful in the context of fault-tolerant computation. Additionally, an experimental procedure for measuring the relevant error would be useful in optimizing the design of the control hardware and software. It is easy to imagine a closed loop optimization procedure where the quantum computer optimizes its own gates by running the benchmarking procedure as a fitness function. One could start with reasonably good guesses for the pulses that implement the desired gates and then run the benchmarking protocol to obtain a number for the error rate. Then, in a similar manner to the optimal control algorithms, the pulse shapes could be varied by some optimization program to minimize the error rate. This self-optimization procedure would have a couple key benefits. Firstly, the simulation of the quantum system’s evolution, which is explicitly inefficient on the classical computer, is now efficiently computed on the quantum computer itself. This removes a barrier to finding control sequences for large systems. Secondly, because we are using the quantum system to simulate itself, we are sure to capture all the dynamics, some of which might be missed in the model which is used to simulate the system on a classical computer. Work along this lines has previously been reported where improvements have been seen by optimizing decoupling schemes on the system itself versus simulation of a model on a classical computer [68]. Further-

more, in the experiments reported below it was used to optimize the calibration for the single-qubit pulses.

As seen in the above characterization protocols, randomization is a useful tool for coarse graining the information needed. In the present task it has been proposed that randomizing over computational gate sequences and measuring the fidelity decay (overlap between the measured output state and ideal error free output state) as a function of increasing number of gates would be a useful measure of control [69]. It is important to carefully choose the gate set from which the randomized sequence is chosen. It has to meet two criteria: it must depolarize the noise so that the noise can be represented with one parameter; and it should also be possible to efficiently track and measure the fidelity of the final output state. Depolarizing the noise means that any channel should be turned into an average channel which is a depolarizing channel: with probability d the maximally mixed identity state is returned and with probability $(1 - d)$ the original density matrix is returned:

$$\bar{\Lambda}(\rho) = (1 - d)\rho + \frac{d}{D}\mathbb{1}, \quad (3.11)$$

where D is the dimension of the Hilbert space. The depolarizing parameter d then gives some information about the strength of the noise and is related to the average fidelity of the channel.

In the unrealistic scenario of pulse independent errors, then fully Haar random unitaries are efficient depolarizers and a self-inverting sequence of l random unitaries will give an exponential fidelity decay [46]. The density matrix after l applications of the average noise (Eq. 3.11) is

$$\rho_l = (\Lambda^{average})^l(\rho) = (1 - d)^l \rho + \frac{1 - (1 - d)^l}{D}\mathbb{1}. \quad (3.12)$$

And so the fidelity with the initial density matrix ($Tr(\rho_l^\dagger \rho)$) is

$$F_l = (1 - d)^l Tr(\rho^2) + \frac{1 - (1 - d)^l}{D}. \quad (3.13)$$

Thus, by measuring the fidelity as a function of the number of gates, and fitting the curve to an exponential, gives a measure of the strength of the noise.

Unfortunately, fully Haar random unitaries are exponentially hard to generate and it would also be exponentially hard to track the state through the computation to either determine the final fidelity or a recovery gate. Fortunately, gate fidelity requires only a 2-design (gate fidelity is a polynomial function of order two in the elements of the unitary) and there exist more accessible gate sets which satisfy the 2-design requirements. The n -qubit Clifford gate set is one such 2-design [70]. The

Clifford gate set is attractive for a number of reasons. They are computationally relevant. Many fault-tolerant constructions [67], rely on Clifford group gates. By themselves, they are not a universal gate set, but they can be bootstrapped with the preparation of so called “magic states” which furnish universality. Furthermore, high-fidelity Clifford group operations can be used to boost the fidelity of other operations such as state preparation, measurement and even to purify the magic states [71]. Finally, certain relevant input states can be tracked efficiently (by keeping track of the stabilizers of the state [61]) through the computation, allowing the ideal output state or a recovery gate to be efficiently calculated.

3.4.1 Single-Qubit Benchmarking

This leads to the protocol introduced in Ref. [69] for randomized benchmarking of quantum gates. The circuit implementing this protocol for a single qubit is shown in Fig 3.8. A series of random computational gates is chosen from a generating set for the Clifford group, for example, 90 degree rotations about the x and y axis. The series is truncated at different lengths to measure the fidelity decay curve. Random Pauli gates are inserted between every computational gate. These Pauli gates do not advance the computation but rather change the Pauli frame which helps to depolarize the noise. These gates would occur naturally in fault-tolerant constructions based on teleportation [72]. A fiducial initial state is prepared, say $|0\rangle$, and tracked through the computational gates. At each truncation a recovery gate is chosen at random which will return the state to the initial fiducial state. For example, if the state is in an eigenstate of σ_x , the recovery gate should be chosen from Y_{90} or Y_{-90} . Finally, many runs are made randomizing over different Pauli frames and different computational gate sequences to induce an effective twirling of the noise.



Figure 3.8: Quantum circuit implementing single qubit benchmarking. A fiducial state is prepared and a sequence of computational gates G is applied. The recovery gate R is chosen to return the state to fixed state. The Pauli gates P interleaved with the computational gates induce a Pauli randomization.

This protocol was implemented in liquid state NMR. A single qubit molecule was chosen to be $\approx 0.3\%$ solution of unlabeled chloroform in a d_6 -acetone solvent¹. The sample was not freeze-pumped to avoid excessively long T_1 recovery times.

¹This is also the official Bruker lineshape sample to determine if the static magnetic field

The measured T_1 was 7 seconds and the T_2 was measured to be 4.5 seconds with a standard CPMG refocussing sequence. The unrefocussed T_2^* was 0.45 seconds measured from the linewidth.

The experiments were performed on a 700MHz Bruker Avance spectrometer using a TCI cryogenic probe. The cryo-probe provides enhanced sensitivity and associated improved signal to noise ratio but the high quality factor of the probe resonant circuit leads to phase-transient and radiation damping effects [26]. To address the amplitude and phase transient issues with the high Q cryoprobe, $24\mu s$ gaussian shaped $\frac{\pi}{2}$ pulses were used. These avoid the unwanted effects due to their more slowly varying amplitude profile. Radiation damping effects were suppressed with a very low concentration sample. Since the errors are expected to be dominated by pulse miscalibration, amplifier drift and r.f. inhomogeneity, composite pulses robust to r.f. field variation were also tested. The BB1 family of pulses from Wimperis et al. [73] are robust to pulse length (calibration) errors ϵ up to at minimum order ϵ^6 (see Figure 3.2) and are universally compensating in that they are robust unitary operations rather than robust for a particular state to state transformation. The pulses consist of a compensating block followed by the desired pulse so that a rotation by an angle θ about the x axis can be replaced as follows:

$$R_x(\theta) = (180)_{\phi_1} (360)_{\phi_2} (180)_{\phi_1} R_x(\theta), \quad (3.14)$$

where, ϕ_1 and ϕ_2 depend on the pulse flip angle according to,

$$\phi_1 = \frac{1}{3}\phi_2 = \arccos\left(\frac{-\theta}{4\pi}\right). \quad (3.15)$$

The location of the compensating block is not important and it can be placed before or after the pulse. The pulse can even be symmeterized by placing the compensating block between two halves of the pulse [74].

The results of the single qubit benchmarking with BB1 composite pulses are shown in Figure 3.9. It is clear that the pulse fidelity is not good and furthermore that the curve does not fit an exponential decay well indicating that the noise has not been completely depolarized or that error per gate depends on the sequence length. However, this can be understood from the fact that the r.f. field strength varies across the sample and that this r.f. inhomogeneity is particularly bad in cryogenic probes [23]. Indeed, by measuring the r.f. inhomogeneity profile and simulating the experiment across that variation, we were able to reproduce both quantitatively and qualitatively the result showing we understand well the error model. The result intuitively makes sense: we would expect spins which see an r.f. field very different to the ideal field to very quickly end up at some random point

homogeneity meets specifications.

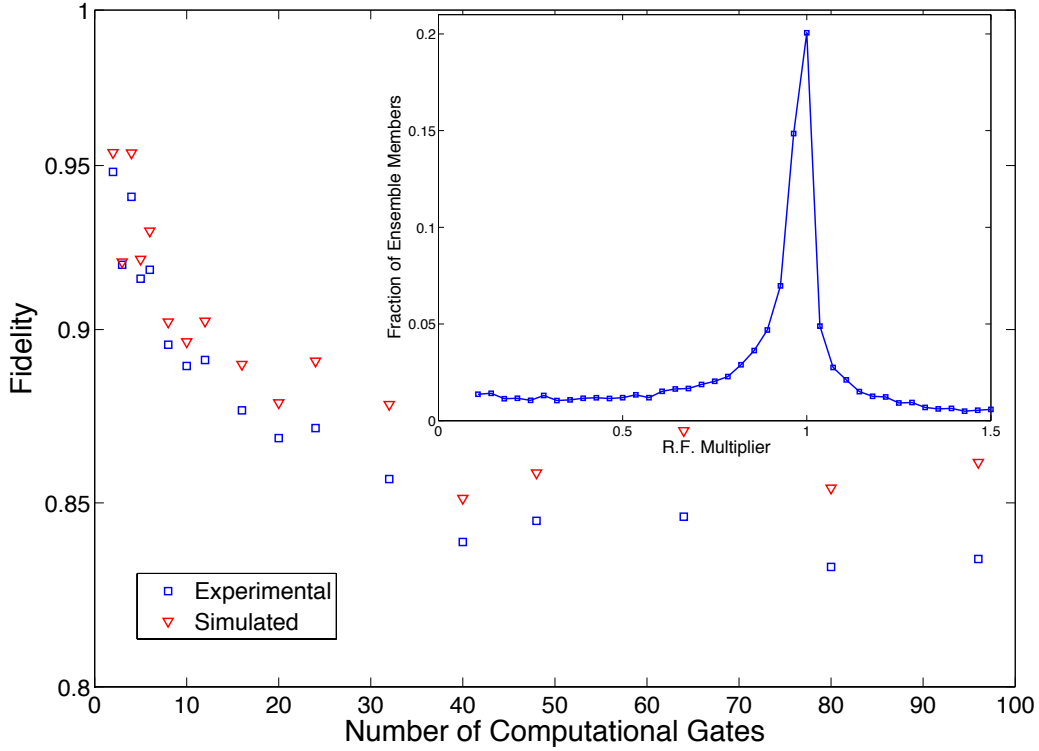


Figure 3.9: Experimental (\square) fidelity as a function of number of randomized gates for a single qubit using BB1 composite pulses plotted on a semi-log plot. The fidelity decay is clearly non-exponential indicating coherent pulse dependent errors. This effect is caused by the large distribution of r.f. field strengths across the sample. Also shown are the results from simulations of the pulse sequence (∇) averaged over the measured r.f. profile. The simulations match the experimental results well demonstrating the error model captures the problem.

on the Bloch sphere, whereas those close to the ideal field strength will closely track the ideal evolution for many gates. Thus we expect the fidelity to quickly decay (with large fluctuations) as the spins at the edge of the r.f. profile are depolarized and then for the fidelity to level off and decay much more slowly. This intuitive picture can be confirmed in a more detailed analytical analysis of the coherent pulse dependent errors introduced by this model.

The issues arising from r.f. inhomogeneity can be largely eliminated by running a r.f. selection sequence. This is a sequence of pulses and gradients that leaves polarization on only a subset of ensemble of processors that experience an r.f. field within a certain range, say $\pm 2\%$ of the ideal field strength[28]. For calibration purposes and again to avoid the sharp transitions of hard pulses I developed an

optimal control pulse which implemented the r.f. selection. The pulse was designed to rotate spins outside the $\pm 2\%$ range of desired powers to the $X - Y$ plane while leaving the calibrated spins along the z-axis. The unwanted spins are then dephased using gradient techniques. The pulse was found by modifying the GRAPE code to have different goal states for different r.f. powers and then defining a global fitness function as the average of the fitness function at each r.f. power from a range of 0.4 to 1.3 spaced every 0.02. The spacing had to be this tight otherwise the pulses found would not behave as expected in between the defined points. For r.f. powers outside the $\pm 2\%$ range the goal state was I^+ which conveniently asks for the state to be in the $X - Y$ plane (but does not care where in the plane), whereas within the $\pm 2\%$ range the goal state was Z . Two rounds of the r.f. selection sequence were used to create an even sharper profile and reduced the signal by $\approx 70\%$. Using the r.f. selection dramatically improves the results and gives an exponential decay which we fit to give an error per randomized computation gate of $1.3 \pm 0.1 \times 10^{-4}$ (see Figure 3.10).

Because the r.f. selection sequence is so selective, the amount of signal returned after the sequence (or in other words the fraction of the sample seeing the correct power) is very sensitive to the power calibration of the r.f. selection pulse. Small fluctuations in this power (of less than 1%) due to random fluctuations in the signal generation and/or amplification, or small fluctuations in the resonance frequency of the high Q probe lead to large fluctuations ($\pm 5\%$) in the signal returned from the r.f. selection sequence. These must be eliminated to measure an accurate fidelity decay without excessive averaging. A solution is to use stroboscopic observation of the signal after the r.f. selection sequence. Because the measurement in NMR only very slightly affects the state of the ensemble system [75], it is possible to observe the amount of signal returned from the r.f. selection sequence and then proceed with the benchmarking experiment. The fidelity result after the benchmarking sequence can then be normalized by the initial input signal. This requires more post-processing and filtering of the signal: performing ourselves much of what is normally done automatically by the spectrometer hardware and software. However, the benefits are worthwhile because the normalized signal fluctuates less than 0.5% between experiments.

There are other proton signals within the bandwidth of preamplifier and receiver. When looking at a narrow spectrum, these signals will be aliased into the spectral window. In conventional digital NMR the spectrometer oversamples the FID at a rate sufficient to encompass all these proton signals. The FID is then digitally filtered to remove all signal outside the specified spectral window. This filtering leads to a group delay or a transient effect at the start of the signal for a number of points which depends on the length of the digital filter. When we take a stroboscopic reference after the r.f. selection sequence there is an insufficient




Figure 3.10: Semi-log plot of the average fidelity as a function of the number of randomized gates for a single qubit using BB1 composite pulses after a r.f. selection sequence. The error bars (68% confidence) indicate the uncertainty from randomization (i.e. different computational sequences and Pauli randomizations give different fidelities due to coherent errors). The uncertainty in each measurement due to signal to noise and fluctuations in the amount of signal from the r.f. selection sequence is less than 0.5%. The fidelity decay is a good fit to a single exponential shown in red with 68% confidence fits.

number of points for the standard spectrometer filtering so we must acquire the data in analog mode and perform our own digital post-processing in MATLAB.

The pulse sequence commands for stroboscopic observation are not standard for liquid-state NMR so I include some relevant parts here for reference. We use the *AvanceSolids.incl* include file for some useful macros for spectrometer machine level commands such as *setfrtp7-1*. When performing stroboscopic observation it is important to set the dwell time and number of time domain points correctly on the spectrometer because it is no longer handled automatically and it is easy to overrun the receiver memory. In general it is easier to have slightly too many points. It is also important to open up the spectrometer filter (FW Bruker parameter) to the maximum allowed (1.25MHz). We used a dwell time of $20\mu\text{s}$ per complex point to cover the entire proton spectrum.

```
;;; Pulse sequence for the single qubit benchmarking

#include <Avancesolids.incl>
#include <Grad.incl>

;;;Declarations:start

define pulse comp90
"comp90 = 235.2u"
define pulse comp180
"comp180 = 261.6u"
define pulse rfsel
"rfsel = 6250u"
define pulse gradpul
"gradpul = 2m"
define delay aqreal
"aqreal = aq-10m"

;;;Declarations:stop

1 ze
2 d1

10u reset:f1 reset:f2

50u UNBLKGRAD

;;;;;;;;;;;;; Water/Solvent signal suppression ;;;;;;;;;;;;;;
```

```

;;; pulses not shown

100u reset:f1 reset:f2

;;;;;; RF selection
( 6u rf sel:sp10 ph0 6u ):f2
gradpul:gp4
1m

;;; Reference taking
;; Receiver start up
STARTADC      ;prepare adc for sampling, set reference frequency
RESETPHASE    ;reset reference phase
RGP_ADC_ON    ;open ADC gate
10u REC_BLK   ;blank the receiving path

;;Composite Y90 readout pulse
( 5u comp90:sp2 ph1 5u ):f2

25u REC_UNBLK ;unblank the reciever
10m DWELL_GEN ;10ms delay while generating dwell pulses
10u do:f1
10u REC_BLK

;; Flip back with final gradient cleaning
100u
( 5u comp90:sp2 ph3 5u ):f2
gradpul:gp6
1m

;;; Main benchmarking sequence goes here

;;;Benchmarking control pulses.....

;;Composite Y90 readout pulse
( 5u comp90:sp2 ph1 5u ):f2

;;;Pulses:stop

;; Final acquisition
25u REC_UNBLK

```

```

aqreal DWELL_GEN

rcyc=2
100m mc #0 to 2 F0(zd)
20u BLKGRAD
exit

;;;Phase definitions....

ph0 = (4) 0
ph1 = (4) 1

```

The FID is loaded into MATLAB and split into the first 500 points and the rest. The first 500 points are the reference which contains signal from the desired chloroform signal as well as the other acetone/solvent signal. The first 500 points are filtered with a fifth order Butterworth low-pass filter with a cut-off frequency at $1/32$ of the sampling frequency. The second half of the filtered signal (points 250-500) is then used as the reference value to avoid transient effects. The rest of the FID is then decimated by a factor of 32 to return the standard FID with all other signals filtered out. The effects of this stroboscopic referencing scheme can be seen by comparing a series of the same experiment repeated many times and observing the fluctuations in signal intensity. The experiment is simply the r.f. selection sequence followed by a readout pulse. As shown in Figure 3.11, because of the r.f. selection sequence, the signal intensity fluctuates in a bi-modal pattern between 1.04 and 0.96 giving an unacceptable variation for determining the state fidelity. However, the intensity of the stroboscopic reference is almost perfectly correlated with the signal intensity fluctuations and so when normalized the fluctuations are suppressed to less than 0.2%.

An estimate of the expected error rate due to intrinsic decoherence can be made from the measured T_1 and T_2 values. The combined time for a randomized computational gate using BB1 composite pulses is $516.8\mu s$ (including delays between pulses to avoid probe overheating). The measured T_2 consists of contributions from both the purely dephasing T_2 and the dephasing induced by the T_1 process:

$$\frac{1}{T_2^{measured}} = \frac{1}{T_2} + \frac{1}{2T_1} \quad (3.16)$$

Knowing both the measured T_1 and T_2 , the intrinsic value for T_2 can be extracted. Then, the Kraus operators for the a map consisting of purely T_1 and T_2 decoherence acting for the gate time can be written down [76]:




Figure 3.11: The integral of the signal (+) for 100 repeated experiments showing the large fluctuations that the r.f. selection sequence introduces. For reference 80 experiments without the r.f. selection sequence are also plotted (circles). However, the intensity of the stroboscopic reference (*) is perfectly correlated with the signal intensity and can be used to normalize the integral to obtain less uncertainty in the results. A series of 80 experiments with no r.f. selection sequence is also shown for comparison to show that indeed it is the r.f. selection sequence causing the variation.

$$A_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\lambda-\gamma} \end{pmatrix}; A_1 = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{\gamma} \end{pmatrix}; A_2 = \begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix} \quad (3.17)$$

where,

$$\gamma = 1 - e^{-\Gamma_1 t} \quad (3.18)$$

$$\lambda = e^{-\Gamma_1 t} (1 - e^{\Gamma_2 t}) \quad (3.19)$$

and $\Gamma_1 = 1/T_1$ and $\Gamma_2 = 1/T_2$.

From the Krauss operators, the average fidelity can be calculated [46]:

$$\overline{F}_g(\Lambda) = \frac{\sum_k |Tr(A_k)|^2 + D}{D^2 + D} \quad (3.20)$$

Equation 3.20 also provides the relation between the HS fidelity quoted above and the true average fidelity. The difference is small for larger systems. Finally, it also provides the relation between the averaged fidelity decay rate and the average gate fidelity. For the measured T_1 and T_2 this would imply an error per randomized gate of 5×10^{-5} . This represents a lower bound on the expected error rate which we should be able to reach with hardware and software improvements. If the T_2^* , rather than the T_2 , is used in the decoherence model, the estimated error per gate climbs to 4×10^{-4} . We would expect the experimental value to fall between these two numbers and indeed that is what we measured. The randomized gate sequence will somewhat refocus the static field inhomogenities contributing to T_2^* , but they are not explicitly refocussed. The remaining impediments of incoherence across the ensemble members and the fluctuations in power from the amplifier could be overcome with even more robust and compensated pulses (although at some point there is a tradeoff between more highly compensated pulses and the increased losses due to intrinsic decoherence because of the longer pulse times).

For comparison purposes, we also tested other pulse types with the same protocol. Using only simple uncompensated pulses we obtain an error rate of $2.1 \pm 0.2 \times 10^{-4}$ and using GRAPE numerically optimized pulses, an error rate of $1.8 \pm 0.2 \times 10^{-4}$. The GRAPE pulses were numerically optimized to 99.999% fidelity (Hilbert-Schmidt norm) over a range of r.f. powers $\pm 3\%$ from the ideal power. They were $100\mu s$ in length and discretized at $1\mu s$. It is somewhat surprising that the numerically optimized pulses cannot match the performance of the BB1 pulses. However, the BB1 pulses are well suited to compensating for systematic deviations from the ideal pulse shape which manifest themselves as calibration errors. Numerically optimized pulses are somewhat robust to noise in the pulse generation: because the controls are at a local maximum of fidelity, any deviation gives no change in the fidelity to first order. However, numerically optimized pulses are still more sensitive

to other imperfections in the implementation. For example, the optimization and robustness assumes the control fields are constant at each time step. In the experiment, finite bandwidth effects and noise prevent exact implementation of this and lead to a loss of fidelity.

3.4.2 Multi-Qubit Benchmarking

When trying to benchmark the level of control in multi-qubit systems, it is much more difficult to say what quantity the measured error per gate refers to, or even what it should be. In order to make the quoted results relevant to fault-tolerance threshold proofs and calculations, it would be ideal to quantify the error per gate for one and two qubit gates and also storage errors for wait steps. However, it is difficult to isolate the errors for only these gates if the error model does not satisfy the independent error model - that each gate's errors are described by a quantum operation acting only on qubits which the gate affects. In realistic situations it is most likely that applying a gate to qubit a could induce an error on qubit b . The gates should depolarize the noise but at the same time the error per gate should be meaningful in relation to the fault-tolerant thresholds.

One possibility is to choose a generating gate set consisting of single qubit Clifford generators (say the Hadamard and phase gates) and controlled-NOT gates between pairs of qubits. This will generate the multi-qubit Clifford group and indeed after only a polynomial number of gates will approximate a 2-design necessary for depolarizing the noise [77]. The multi-qubit protocol then becomes:

1. Choose a series of lengths of computational gates from a generating set for the Clifford group which to measure the fidelity decay at. The number of random gates necessary to achieve depolarization of the noise depends on the number of qubits and may be large. Thus we only expect the asymptotic error rate to be meaningful.
2. For each truncation length choose n_g random sequences of computational gates.
3. Pick a fiducial input state and track that state through the gate sequence. Determine a recovery sequence which will return the state to one with a known definite output upon measurement. Because the Clifford group operations can be efficiently tracked this is possible to do efficiently on a classical computer [78].
4. Apply some parallelization routine to the random sequence of gates to ensure that the number of wait steps does not grow with the size of the computer.

If only one computational gate is performed per time step the wait steps on all the other qubits will come to dominate the errors as the size of the system grows. Fault-tolerant proofs require some form of parallelization to overcome this. In a similar manner some parallelization is needed in the benchmarking to extract the relevant average error rate. This parallelization step allows a fair comparison between different size qubit registers, say one with 5 and one with 50 qubits. The error per time step may be larger in the 50 qubit computer but many more gates are possible in each timestep.

5. Measure the fidelity decay as in the single qubit case. An exponential fit to the fidelity decay will reveal the average error per one and two qubit gate. It is possible that the average error could mask a distribution of error rates such that for example all single qubit gates are perfect but the two qubit gates are much worse. However, more detailed, but still coarse grained information is available by doing more experiments (see below).

Numerical simulations have confirmed that this protocol will return the correct error rate for a variety of error models such as dephasing and pulse dependent unitary errors. Not surprisingly, larger amounts of randomization are required compared with the single qubit protocol. Even when averaged over many computational gate sequences, each time the protocol is run a different error strength is calculated due to varying degrees of twirling and the random nature of the protocol. The protocol was tested by simulating 256 runs of the protocol with each run averaged over 48 different computational sequences as would be performed in the actual experiments. The error model tested is a qubit independent dephasing model but similar results are obtained for other error models. There is a spectrum of errors from fully depolarizing to fully random pulse-dependent unitaries. The closer the error model is to pulse-dependent unitary, the more randomization is needed. The error rate is then measured by fitting the asymptotic fidelity decay curve (beyond 30 gates). The measured error rate is then compared to the theoretical one and the distribution plotted (see Figure 3.12). The distribution is reasonably tight, varying only $\pm 10\%$ and so we claim that this protocol reveals the error per gate. Higher accuracy can be obtained with further averaging or averaging a few runs of the protocol.

More detailed information about the errors can be obtained by combining the ideas of previous randomization protocols [60, 79] with the randomized computational sequences. The idea is to run localized versions of the benchmarking protocol on different sub-systems in the QIP. The steps of the proposed protocol are as follows:

1. Perform the single qubit benchmarking procedure on each qubit individually. These numbers will give an estimate of the error per gate for single qubit




Figure 3.12: Distribution of measured error rates in multiqubit benchmarking simulations as percent difference from theoretical error rate for 256 simulated experiments. The error model is qubit-independent dephasing with an expected depolarizing parameter of 0.01 (see equation 3.11). The spread of measured error rates is comparable to the statistical and experimental uncertainty in fitting any one run.

gates. Unfortunately, as discussed above, the error model is unlikely to follow an independent error model and the possibility that performing the single qubit gates induces errors on non-target qubits needs to be checked. This can be achieved by measuring the fidelity of the identity operation on the other $n - 1$ qubits. Efficient procedures exist for this measurement. For example, as discussed in section 3.3, performing single qubit Clifford gates at the beginning and their inverses at the end of the sequence, and then randomizing, allows an estimation of the fidelity of the channel with a small number of experiments [60]. A possible concern is that the error model on either the single qubit or the remaining $n - 1$ qubits might be highly non-Markovian. However, the benchmarking procedure should effectively act as a randomized dynamical decoupling sequence [80] which should limit these effects.

2. Perform two qubit benchmarking: the procedure described for multiple qubits above can then be performed on all pairs of qubits (or at least on all pairs of qubits that would have two qubit gates between them given the constraints of the architecture). The other $n - 2$ qubits should be twirled as in the first step to assess the fidelity of the wait steps.
3. This procedure can be iterated to all groups of 3 qubits and so on but because most fault-tolerant constructions are specified in terms of one and two qubits gates going as far as all pairs should be sufficient.

The multi-qubit benchmarking was performed experimentally in liquid-state NMR on a 700MHz spectrometer with a TCI cryoprobe. A three qubit sample was made from selectively labelled ^{13}C tris-(trimethylsilyl)-silane-acetylene dissolved in deuterated chloroform [81]. A table of the natural Hamiltonian parameters is shown in Table A.2. All T_1 's and T_2 's are greater than 2 seconds.

The carbon sub-system is strongly coupled and the conventional NMR control strategies of soft selective pulses and refocussing networks do not work well. Control was instead achieved through the GRAPE optimal control technique [22]. Using contained optimal control pulses also makes it easier to concatenate them together in a pulse sequence as a new refocussing network does not need to be found for every sequence. The pulses were optimized to above 99.95% average fidelity over a range of r.f. powers $\pm 3\%$ from the ideal power. The pulses were discretized at $2 \mu\text{s}$ as a balance between smoothness and limiting the number of points to define due to finite spectrometer memory. Single qubit pulses were 1.2ms long; CNOT gates between H and C_1 (and any single qubit gate on C_2) were 2.4ms ; and CNOT gates between C_1 and C_2 (with any single qubit gate on H) were 4ms . These pulses are not time-optimal but have low enough r.f. powers for

experimental implementation. Shorter pulses tended to require unfeasible high power levels which led to probe heating during long computational sequences. Non-linearities in the pulse generation and transient effects from the probe’s resonant circuit lead to distortions in the implementation of shaped pulses. These can lead to unexpected deviations from ideal performance. To avoid this, the r.f. field at the sample was detected through a pickup coil and corrected through a simple feedback loop. Technical details of this correction procedure are found in Appendix C. This correction procedure was only applied to individual pulses and the longer term power inverse droop we observed² was not corrected but should instead be handled by engineering robust pulses. Due to finite spectrometer memory we are limited to 120 gates in a computational sequence. The same numerically optimized GRAPE r.f. selection sequence as was used in the single qubit experiment was applied before each experiment to the proton nuclei. Polarization on the carbon nuclei was dephased with gradient techniques giving the starting deviation density matrix $ZII (HC_1C_2)$.

The sequence of random gates was chosen in the following manner [78]: with 1/3 probability each; a Hadamard gate on qubit a ; a PHP gate (Hadamard conjugated by a phase gates) on qubit a ; and a CNOT with qubit a as control and qubit b as target. For the CNOT gates, a and b are restricted to be nearest neighbours. The resulting state was then tracked and a recovery sequence to return the state to ZII calculated. The entire sequence was then parallelized with a simplistic iterative scheme of repeatedly checking whether gates in series could be compressed into a single gate. For example, a CNOT gate between qubits 2 and 3 followed by a Hadamard gate on qubit 1 would be compressed to a single timestep which implements both gates in parallel. The fidelity of the state was then measured through a readout pulse on the proton spin.

The results are shown in Figure 3.13. The results fit an exponential decay well and give an error per gate of $4.7 \pm 0.3 \times 10^{-3}$, approximately an order of magnitude worse than the single qubit results. Again, an estimate of the lower bound on the error rate can be obtained from the measured T_1 ’s and T_2 ’s. Assuming an independent and uncorrelated error model (which is unlikely but does not significantly affect the result) gives an average error per gate of 1.3×10^{-3} . Moreover, from the design of the pulses, we would expect an error of 4.4×10^{-4} , which is an order of magnitude smaller than the experimentally measured error rate. This leads us to suspect that there are still errors in the implementation of the pulses and/or knowledge of the chemical properties of the molecule that are not currently handled by

²The r.f. power from our signal generation and amplifier combination increases several percent over a timescale of 10’s of ms . This could also be due to the resonant condition of the probe changing due to heating. See Ref. [82] for another observation of this effect in the context of optimal control.




Figure 3.13: Semi-log plot of the average fidelity as a function of the number of randomized gates for 3 qubit benchmarking using GRAPE pulses after a r.f. selection sequence. The error bars (68% confidence) indicate the uncertainty from randomization (i.e. different computational sequences give different fidelities due to coherent or biased errors). The fidelity decay is a good fit to a single exponential shown with 68% confidence fits.

our pulse design.

The work in this section was published in Ref. [83]

3.4.3 Future Improvements

In both the single and multi-qubit setting we were unable to establish decoherence limited control. The control was most likely limited by inhomogeneities across the sample and imperfect implementation of the numerically optimized pulses. There are two approaches to this problem: software and hardware. In the software the pulses can be made more robust to inhomogeneities by improving the pulse optimization. This will be mainly a brute force effort by parallelizing the code and running it on large clusters. A more involved improvement to the software is to

investigate the robustness of the pulses to noise, particularly noise on a timescale less than the timestep. On the hardware side, further improvements to the pulse generation (more memory and faster discretization) would require new spectrometer purchases. A more useful and less expensive avenue would be to concentrate on reducing the noise in the pulse feedback loop to improve the correction of the pulses and be more certain about what fields the spins are seeing. This is particularly difficult at low pulse powers where the pickup coil signal-to-noise is worse. Finally the ability to perform feedback on multiple pulses in a sequence would allow the correction of slower noise in the signal generation and amplification.

Another line of further work is to explore some of the extensions to the multi-qubit benchmarking. This would be useful not only for further development of the protocols but also for us to better understand the noise and errors in our NMR system. For example, are the errors localized to any one qubit? Or, is there a difference in error rates between the one and two qubit gates?

Chapter 4

Creating pure states out of mixed ones: Algorithmic Cooling

4.1 Theory

Nearly all models of quantum computing (DQC1 is one clear exception - see Section 3.2 for an example implementation) require close to pure quantum states. These are needed not only as an initial input state into the algorithm but also, and more importantly, they are needed throughout the computation for error correction purposes. Quantum error correction can be viewed in a thermodynamic view as follows. The initial pure state has zero entropy and any errors during the computation due to an interaction with an environment at some finite temperature cause the system to become a mixed state and raise the entropy of the system. This entropy can be pumped out of the system of interest into pre-cooled ancilla qubits via quantum error correction but the ancillas themselves must continually be refreshed through some cooling procedure or they would heat during the computation and be ineffective in their error correction role.

There are standard physical approaches to the problem of creating pure states. For example, cooling the qubits' environment to a temperature well below the energy gap to the first excited state and simply waiting for thermal equilibrium to cool the qubits into the ground state with high probability. This works well in qubits that are already in cold environments such as quantum dots or superconducting qubits. Another approach is to drive certain cycling transitions with radiation to pump the population into a certain state. This approach works well on qubits with an optical transition (and the right level structure) and is used with trapped ions and nitrogen vacancy centres in diamond.

In all cases, the initialization is still imperfect. Polarization (the probability of being in the ground state minus the probability of being in the excited state) on

the order of 99% sounds very impressive, but for fault-tolerant computation purposes two or more orders of magnitude smaller deviation from perfect polarization is needed without one of the cooling or purification techniques presented here. However, if the control techniques and technology are good enough to have achieved a fault-tolerant level of control, then the control can be leveraged to algorithmically cool the qubits beyond the physical initialization. The ideas behind algorithmic cooling are very similar to the ideas of classical error correction and are intuitively easy to understand.

There are two related classes of cooling schemes: closed and open system, derived from their thermodynamical namesake. In a closed system qubits cannot be reinitialized and the entropy of the system is fixed. The question is then: how to best move the entropy around the system to maximize the polarization of certain qubits, or in other words, how to compress the uncertainty in a portion of the system making the state of another portion of qubits more certain? These steps will always be limited by two physical bounds. Firstly, as stated above, the system is closed and the entropy fixed. The uncertainty and entropy can be shuffled around the system but cannot be reduced. This is known as the Shannon bound. For a density matrix description of the state the Shannon bound implies that the sum diagonal elements squared cannot change. For quantum systems there is a further bound if the compression step is limited to a unitary transformation. This is the Sørensen bound [84]. In the density matrix description this means that the eigenvalues (diagonal elements for a classical state) cannot be changed; they can only be permuted. For the experimental implementation described below, a relevant situation is for three qubits. For all three qubits at the same polarization ϵ , the Shannon and Sørensen bounds coincide and state that one qubit cannot be polarized to more than 1.5ϵ .

In an open system further steps can be taken after the optimal compression step because entropy can be pumped out of the system and qubits reinitialized. It is convenient to frame the discussion in terms of a heat bath with a polarization bias ϵ corresponding to its temperature. The qubits are initialized by bringing them in contact with the heat bath. Initially all the qubits are brought to the heat bath temperature with a bias ϵ . The optimal compression step is then applied which will boost the polarization of some of the qubits and reduce the polarization on the rest to below that of the heat bath. These hot qubits can then be brought back in contact with the heat bath to re-cool them and pump the entropy out of the system. These rounds of cooling and compression can then be repeated.

It is assumed that the qubit can be initialized to the thermal state $(1 - \epsilon)|0\rangle\langle 0| + \epsilon|1\rangle\langle 1|$, or in words it is initialized to the ground state with a polarization 2ϵ . If we initialized three qubits and then took a majority vote with the output on one of the qubits then the probability that the output qubit is in the wrong state ($|1\rangle\langle 1|$) is the

probability that two (or three) of the qubits were initialized to the wrong state. This will go as $\mathcal{O}(\epsilon^2)$ (or more precisely ${}_3C_2\epsilon^2 + \epsilon^3$). This process can be concatenated, feeding the output qubits of three majority gates into another majority gate which would suppress the error to ϵ^4 . The procedure can be repeated until the desired level of purity is achieved.

An alternative approach is to bring more than three qubits up to the reset polarization and then perform a more complicated multi-qubit gate to boost the probability of the lower energy states. This has been codified for the situation where there are many qubits and one resettable qubit. This is a physically relevant situation for many electron-nuclear systems. A single electron, which can be reset through its fast T_1 and strong magnetic moment or perhaps through an optical handle, is coupled to many nuclei with long T_1 's to store the polarization and on which to perform the compression gates. In this situation an optimal algorithm has been worked out: the partner pairing algorithm (PPA) [85]. In the PPA at each step the reset qubit is reset and then a gate is performed which permutes the diagonal elements of the density matrix so that they are in non-increasing order. The schemes do not polarize a single qubit. In particular in the case of the PPA, the polarization is boosted in a conditional cascade. At each step, the polarization of the first qubit is maximized. Given that has been achieved the next qubit's polarization is maximized, and so on. More specifically the probabilities for the different states decrease in lexicographical order after each refresh and permutation step. $|000\dots\rangle\langle 000\dots|$ has the highest probability $|000\dots 1\rangle\langle 000\dots 1|$ the next highest and so on till $|111\dots\rangle\langle 111\dots|$ has the lowest probability.

Using the optimality of the algorithm, bounds on the purification possible can be worked out. In particular there is a threshold result that if the heat bath or reset polarization $\epsilon < \frac{1}{2^n}$ then pure qubits are impossible to achieve. On the other hand for reset polarizations above the threshold $\epsilon > \frac{1}{2^n}$ arbitrarily pure qubits are achievable.

In the low polarization regime, the maximum polarization achievable with n qubits, including the reset qubit is $2^{n-2}\epsilon$. This has been shown numerically in Ref. [86] but it can also be shown analytically from the results of Ref [85]. There it is claimed that the maximum achievable population (starting from the maximally mixed state where all populations are 2^{-n}) is: $\min(2^{-n}e^{\epsilon 2^{n-1}}, 1)$. In the low polarization regime we can approximate $e^{\epsilon 2^{n-1}} \approx 1 + 2^{n-1}\epsilon$ and ignore all terms $\mathcal{O}(\epsilon^2)$ and higher. The PPA algorithm will reach a steady state when all the terms in lexicographical order have decreasing probability and the reset step does not have any effect (the polarization of the reset qubit is the same as the heat bath). In terms of the density matrix this implies that the diagonal terms in the density matrix are sorted and the difference between i^{th} and $(i+1)^{th}$ diagonal elements is 2ϵ for odd i . The differences do not have to all be the same for every i but on average they must

be and so it does not affect the following argument. Thus, the diagonal elements of the steady state density matrix can be written down as:

$$\frac{1}{2^n} \begin{bmatrix} 1 + 2^{n-1}\epsilon \\ 1 + 2^{n-1}\epsilon - 2\epsilon \\ 1 + 2^{n-1}\epsilon - 2\epsilon \\ 1 + 2^{n-1}\epsilon - 4\epsilon \\ 1 + 2^{n-1}\epsilon - (2^{n-1} - 2)\epsilon \\ \vdots \\ 1 \\ 1 \\ \vdots \\ 1 - 2^{n-1}\epsilon + 2\epsilon \\ 1 - 2^{n-1}\epsilon \end{bmatrix}. \quad (4.1)$$

The polarization of the first qubit, which is maximized in this case, is the sum of the top half minus the sum of the bottom half. There are 2^{n-1} terms in each half and the bottom and top half are symmetrical so:

$$\begin{aligned} P_1 &= \frac{2}{2^n} (2^{n-1}2^{n-1}\epsilon - 0\epsilon - 2\epsilon - 2\epsilon - 4\epsilon \cdots - (2^{n-1} - 2)\epsilon - 2^{n-1}\epsilon) \\ &= 2^{n-1}\epsilon - \frac{\epsilon}{2^{n-1}} (0 + 2 + 2 + 4 \cdots + (2^{n-1} - 2) + (2^{n-1} - 2) + 2^{n-1} + 2^{n-1} - 2^{n-1}) \\ &= 2^{n-1}\epsilon - \frac{\epsilon}{2^{n-1}} \left(\left(2 \sum_{k=1}^{2^{n-2}} 2k \right) - 2^{n-1} \right) \\ &= 2^{n-1}\epsilon - \frac{\epsilon}{2^{n-1}} \left(4 \frac{2^{n-2}(2^{n-2} + 1)}{2} - 2^{n-1} \right) \\ &= \epsilon 2^{n-2}. \end{aligned} \quad (4.2)$$

Thus the maximum polarization boost in the low polarization regime is 2^{n-2} . Note that this shows an error in Ref. [85] where it is claimed with 5 qubits a boost of 16 times is achievable where in fact the maximum boost is 8.

Errors during the algorithmic cooling procedure limit the error level achievable in an unsurprising way [87]. A straightforward error model is to assume a completely depolarizing error after every gate. Even in the case of a biased error model it can be transformed into a depolarizing model through randomization. Under errors and without modifying the protocol the final polarization reached is limited by the depolarization error rate. This is not surprising because one can consider the final

compression gate producing a pure qubit but then being depolarized by the error in the gate. In this situation with control errors, adding more qubits or more steps of purification does not necessarily improve the purification level because more qubits requires more gates. Thus the purification quickly saturates after only a few steps. The behaviour of the cooling algorithm over a range of polarizations and for a couple different depolarization rates is shown in Figure 4.1. The number of steps required so that the improvement in polarization has dropped below an arbitrary bound of 10^{-6} is shown in Figure 4.2. It shows that adding a depolarizing error causes the purification to saturate at a fewer number of steps.

4.2 Experiment

Experimentally demonstrating algorithmic cooling is a demanding task requiring high fidelity control and the ability to reset qubits. The protocol (and indeed any purification protocol) relies on leveraging control quality for purification, much like composite pulses can leverage high fidelity control over the phase of the pulse to improve the performance with respect to another parameter. NMR QIP is probably the only system where sufficient quantum control has thus far been realized. However, in liquid state NMR there is no known access to a dynamical reset step. The only opportunity is to rely on differences in T_1 . We can imagine a molecule where one nucleus has a very fast T_1 and so the reset can be achieved simply by waiting: for example waiting $5T_1$ to achieve 99% of the thermal polarization. However, a very large difference in T_1 's is needed because while waiting for the reset qubit the polarization on the other qubits will decay at their T_1 rate. Furthermore, the short T_1 on the reset qubit will limit its T_2 and the fidelity of the control. Despite these limitations, some improvement from algorithmic cooling was shown in a three qubit liquid-state system [88].

Solid state NMR offers the requisite level of control and an approximation of a reset step which does not require a relaxation process in the system of interest. The reset step relies on the idea of a spin bath: a network of dipolar coupled spins with a certain polarization leading to the concept of a spin temperature. Other spins can be brought in thermal contact with the bath (*vide infra*) and will reach thermal equilibrium at the spin bath temperature. If the number of spins in the bath is large enough compared with the spins being reset (in other words, that the processor molecules are a dilute fraction of the crystal, while the majority are unlabelled molecules with no processor spins), then the bath temperature will be negligibly affected by the exchange of polarization and the process can be repeated. The experiment presented here used a single crystal of malonic acid. A schematic of the molecule and the final parameters of the Hamiltonian are shown in Figure

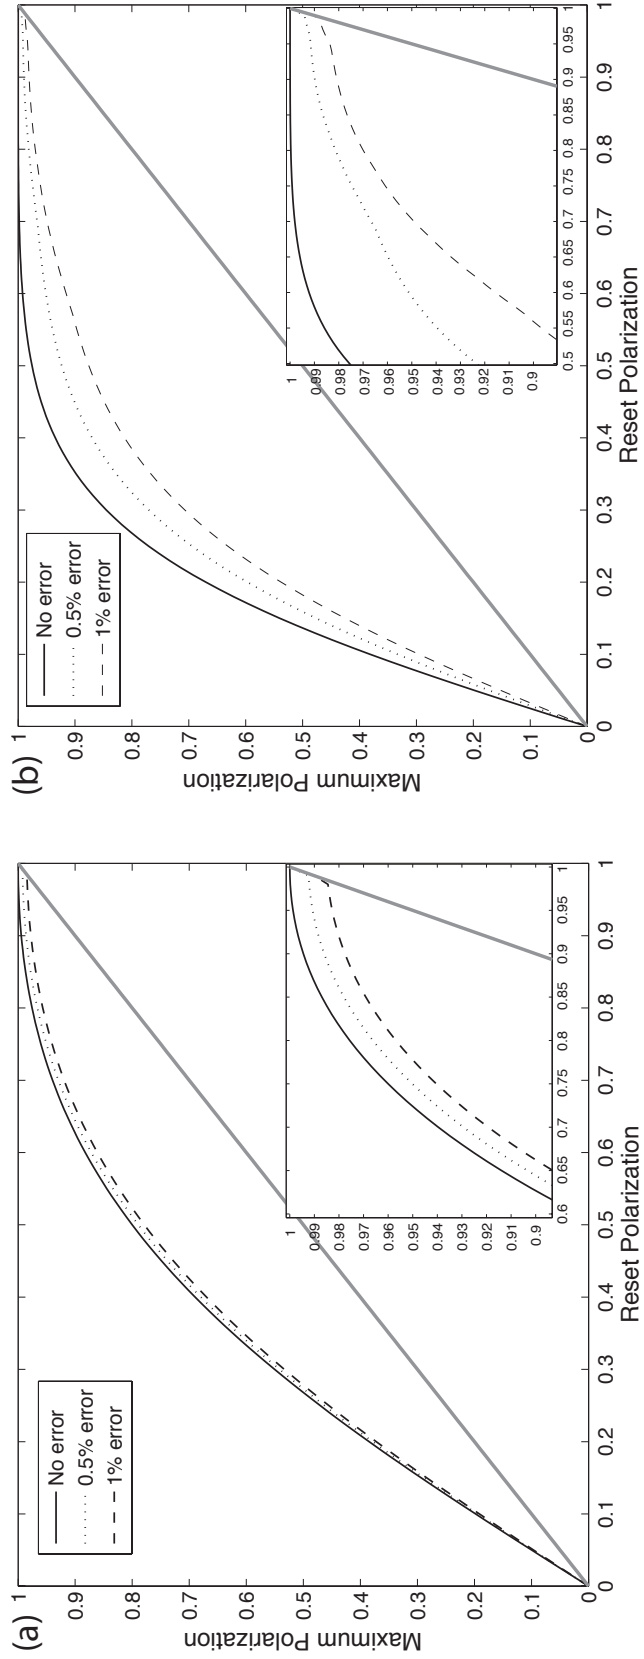


Figure 4.1: Performance of the algorithmic cooling algorithm for three (a) and five (b) qubits with a depolarizing error model for each control step versus the reset polarization. The purification curve with no depolarizing error (solid) and with 0.5% (dotted) and 1% (dashed) depolarization probability are shown (see Equation 3.11).




Figure 4.2: Number of steps to saturate the algorithmic cooling protocol (arbitrarily defined to be improving the polarization of the first qubit less than 10^{-6} per step) versus the reset polarization. The plots is shown for both three (lower curves) and five qubits (upper curves) with no errors and a depolarization error of 1% per gate. The depolarizing error causes the algorithm to saturate at a fewer number of steps.

4.3. The three carbon nuclei will form the processor molecules. The C_m which has a strong coupling to the two methylene protons will be the reset qubit through its interaction with the proton spin bath.

The standard technique in solid-state NMR for achieving the thermal contact between hetero-nuclear spin species is cross-polarization (CP) [89]. The driving force behind the technique is the desire to transfer polarization from a more polarized (and usually abundant) spin species to a less polarized (and usually rare) spin species. A typical situation would be to transfer polarization from proton to natural abundance ^{13}C . Ideally this should lead to a ≈ 4 ($\gamma_H/\gamma_C \approx 3.98$) times boost in polarization of the carbon and reduction in experiment time by a factor of 16. The process is usually explained in thermodynamical terms. The off-diagonal flip-flop terms in the heteronuclear dipolar Hamiltonian are suppressed because of the energy mismatch between the Zeeman energies (assuming the nuclei have sufficiently different gyromagnetic ratios. A way around this mismatch is to match the energies in the doubly rotating frame. In the rotating frame, the energy of the spins is set by the strength of the r.f. field and so by suitably adjusting the nutation frequencies of the two spin species, their rotating frame energies can be made equal and they can exchange energy. This point is called the Hartman-Hahn matching condition. In the proton-carbon example, we would need to drive the carbons with a r.f. field 4 times as strong so that they are both nutating at the same frequency. This standard explanation is not very rigorous. However, it becomes clear when applying average Hamiltonian theory to the situation that the dipolar Hamiltonian is averaged to an exchange interaction for the relevant input states and so polarization will be exchanged [90]. In the doubly rotating frame for a heteronuclear system the dipolar Hamiltonian is reduced to a ZZ form by the large difference in Zeeman energies and so under a driving field ω_1 for species 1 and ω_2 for species 2 the doubly-rotating frame Hamiltonian is:

$$\mathcal{H} = \omega_1 X_1 + \omega_2 X_2 + D_{12} Z_1 Z_2. \quad (4.3)$$

The average Hamiltonian can be worked out by moving to another frame induced by the r.f. driving field:

$$\mathcal{H}' = \exp[it(\omega_1 X_1 + \omega_2 X_2)] D_{12} Z_1 Z_2 \exp[-it(\omega_1 X_1 + \omega_2 X_2)] \quad (4.4)$$

$$= D_{12} [(\cos(\omega_1 t) Z_1 + \sin(\omega_1 t) Y_1) (\cos(\omega_2 t) Z_2 + \sin(\omega_2 t) Y_2)] \quad (4.5)$$

At the Hartman-Hahn matching condition, $\omega_1 = \omega_2$ and the zeroth order average Hamiltonian (see section 2.4) will be non-zero and the \cos^2 and \sin^2 terms average to $\frac{1}{2}$:

$$\overline{\mathcal{H}}^0 = \frac{1}{2} (Z_1 Z_2 + Y_1 Y_2) \quad (4.6)$$

For cross-polarization the polarization is prepared along the direction of the CP field or IX . For this input state the $ZZ + YY$ average Hamiltonian acts as an

exchange Hamiltonian and will swap the polarization back and forth between the two spins.

Under the exchange interaction the polarization should oscillate back and forth indefinitely between a carbon and a single proton spin. However, the carbon is most likely coupled to many proton spins and the proton spins themselves are coupled to each other so that the oscillations will quickly die away leading to a steady-state polarization [90]. However, the initial transient oscillation can be used to achieve close to a full polarization transfer on a much shorter time scale than conventional CP. This is the method we choose to achieve the reset gate. The Hartman-Hahn condition was established for a short time to achieve a polarization swap between C_m and $H_{m1/2}$. Because this proton/carbon coupling is approximately an order of magnitude stronger than any other, the swap is close to selective. After the swap gate the proton field is returned to a high power spin-locking field which decouples the protons yet allows the proton-proton dipolar couplings to refresh the $H_{m1/2}$ by equalizing its polarization with the rest of the proton spins. One can intuitively see that the dipolar couplings will still work under the spin-locking field by moving to the rotating frame where the spin-locking field is a constant field along say the x axis. If the spin-locking field is much larger than the dipolar couplings then this is now analogous to the situation for the static field in the laboratory frame and so we expect the same secular approximation to the homonuclear dipolar coupling to hold (except in this case with the X axis being preferred) yielding

$$\mathcal{H}_{dipolar}^{spin-lock} \propto 2XX - ZZ - YY. \quad (4.7)$$

This can be seen more formally, and the constant of proportionality determined, by an application of average Hamiltonian theory similar to the one above. Again moving to the frame defined by the spin locking field,

$$\mathcal{H}' = \exp(it\omega X_1 + \omega X_2) D_{12} (2Z_1 Z_2 - X_1 X_2 - Y_1 Y_2) \exp(-it\omega X_1 + \omega X_2) \quad (4.8)$$

$$= 2D_{12} [(\cos(\omega t) Z_1 + \sin(\omega t) Y_1) (\cos(\omega t) Z_2 + \sin(\omega t) Y_2)] - \quad (4.9)$$

$$X_1 X_2 - D_{12} [(\cos(\omega t) Y_1 - \sin(\omega t) Z_1) (\cos(\omega t) Y_2 - \sin(\omega t) Z_2)]. \quad (4.10)$$

Again the \cos^2 and \sin^2 will survive the averaging to zeroth order leaving:

$$\overline{\mathcal{H}_{dipolar}^{spin-lock}} = -X_1 X_2 + \frac{1}{2} (Z_1 Z_2 + Y_1 Y_2). \quad (4.11)$$

So that under a spin locking field the dipolar coupling is scaled by 1/2 but the flip-flop terms remain to allow polarization to diffuse through the spin-system.

Other methods can be used to achieve an exchange interaction from a heteronuclear dipolar Hamiltonian. Most notable are multiple pulse techniques. The same sequence of pulses that averages a dipolar Hamiltonian to zero will also average a

ZZ coupling to an isotropic exchange interaction $XX + YY + ZZ$. Thus if the dipolar decoupling sequences are applied synchronously to both spins, this will lead to a swap gate for any input state. This has the added advantage of suppressing the homo-nuclear dipolar couplings which might otherwise lead to errors on the processor spins. Almost any dipolar decoupling sequence will work. However, for algorithmic cooling purposes, where the processor spins have unique non-zero chemical shifts, the sequence must also suppress the chemical shift terms: a true time-suspension sequence. Some possibilities are the Cory48 sequence (or a 12 pulse sub-sequence) [64] or heteronuclear mixing sequences designed to measure heteronuclear dipolar couplings in the solid state such as the SHRIMP16 [91] or WIM24 [92, 93].

Although the multiple pulse technique is probably the most elegant it suffers from several experimental drawbacks for this specific situation. The proton-carbon coupling we are using is quite strong - and we want it that way in order to perform a selective swap gate. This implies that a full cycle of the multiple pulse sequence must be completed in the exchange time. This demands very short pulse and delay times and in some cases this is simply not feasible due to bandwidth and power limitations. Furthermore, the dipolar coupling between the two protons is also quite strong (although somewhat tunable via crystal orientation) which strains the ability of the pulse sequences to suppress homonuclear dipolar interactions and interferes with the heteronuclear exchange. Finally, the multiple pulse sequences require short hard pulses which contain inevitable transient effects. Despite our best efforts to tune these out [63] we were unable to obtain the high fidelity control needed for algorithmic cooling. Specifically, although both the WIM24 and SHRIMP16 worked well as the transfer sequences they were designed to be, they performed poorly at preserving the proton polarization which was not their designed intention.

Optimizing the natural Hamiltonian of the malonic acid was crucial to the success of this experiment. The ability to tune the Hamiltonian through the orientation of the molecule is recognized as a benefit of solid-state NMR. The dipolar coupling depends on the angle between the vector joining the nuclei and the static field as a 2^{nd} -order Legendre polynomial $P_2(\theta) = 3\cos^2\theta - 1$, and the chemical shifts have a tensorial dependence on the orientation of the crystal with respect to the static field. However, given our limited ability to orient and rotate the crystal, and the constraints of the molecular geometry, there is not complete freedom with every term in the Hamiltonian. Furthermore, as with nearly all experimental parameters, there are design tradeoffs which need to be optimized. Our goals and constraints for the Hamiltonian can be listed:

1. Large chemical shift differences between the carbon processor qubits for faster selective control. However, we want small chemical shift differences to make

it easier to refocus the chemical shift terms during the transfer sequence and to minimize the power requirements of the GRAPE pulses.

2. Small proton chemical shifts to improve the ability to preserve the heat-bath (1H) polarization during the decoupling and the transfer sequence.
3. Small dipolar coupling between the methylene protons to make the proton decoupling and transfer sequences more efficient. However, we do want the methylene protons to have a strong dipolar coupling to the rest of the proton spin bath so that they rethermalize quickly after a transfer sequence.
4. Large difference in proton-carbon dipolar coupling between the reset qubit and the other ones to maximize the selectivity of the transfer sequence. Polarization leaking onto other processor qubits is irrelevant during the first steps of the algorithm. However, after one of the qubits has been boosted above the heat bath polarization by the compression step, any contact with the heat bath during the refresh step will lead to a loss of polarization. For the application of multiple pulse transfer sequences we do not want the proton-carbon qubit coupling too large so that the transfer happens before one complete sequence. Furthermore, a large proton-carbon coupling also makes the proton-decoupling difficult and is one of the main factors reducing the fidelity of our carbon control operations.

The Hamiltonian for the entire proton-carbon system is worked out in a two step process. A carbon spectrum is taken after cross-polarization. A high quality fit of the spectrum reveals the full carbon Hamiltonian including chemical shifts, dipolar couplings and J-couplings. From other sources (see below) the positions of the atoms in the unit cell should be known with high precision. These two sets of information constrain the orientation of the crystal with respect to the static magnetic field. The orientation can be further pinned down using additional information such as the methylene proton dipolar coupling from the splitting in the proton spectrum or proton-carbon coupling information from observing coherent oscillations during the transfer sequence. This heteronuclear coupling data is less accurate because the presence of two protons complicates the spin dynamics. Once the orientation of the crystal is obtained, the rest of the proton-carbon and proton-proton dipolar couplings can be determined from the structural data.

Unfortunately, there are inconsistencies with the available data on the molecular structure. X-ray data [94] gives the distance between the two methylene protons as $1.52(1)\text{\AA}$ and the distance between C_m and the methylene protons as $0.97(2)\text{\AA}$ and 0.852\AA . It should be noted this data was used in the original algorithmic cooling experiment [95] where it was difficult to get a consistent Hamiltonian for the system. This X-ray data is inconsistent with unpublished neutron scattering data reported

in an ENDOR study of malonic acid [96], and with NMR data from Ref. [97] and our own experimental work. The neutron scattering data gives the proton-proton distance as $1.756(9)\text{\AA}$ and the distance between C_m and the methylene protons as $1.084(6)\text{\AA}$ and $1.084(9)\text{\AA}$. The X-ray and neutron scattering data agree on other interatomic distances between carbons and oxygen nuclei to better than 1%. Because of the r^{-3} dependance of the dipolar coupling these small differences in distances for the methylene protons lead to large differences in coupling values. The NMR study [97] and our own observations have found a maximum dipolar coupling between the two methylene protons of $66 \pm 1\text{kHz}$. Since the dipolar splitting between two equivalent nuclei is $3D_{ij}$ this implies a maximum dipolar coupling of 22kHz . Indeed a more careful analysis over a range of orientations gives proton-proton distance as $1.76 \pm 0.02\text{\AA}$ - consistent with the neutron scattering data. If isotope effects can be ignored, then further confirmation is obtained from deuterium NMR work [98]. There the $H_1 - C_m - H_2$ bond angle is determined to be 108° (versus 112° from X-ray) and the proton carbon distance as $1.09 \pm 0.025\text{\AA}$. Given this overwhelming evidence in favour of the neutron scattering data my work used that data for the Hamiltonian determination and found more consistent reasonable results. Although X-ray diffraction is not well suited to determining proton coordinates it is possible that this discrepancy is related to the temperature the experiments were performed at. Malonic acid undergoes a phase-transition from the room temperature phase below $47\text{-}60\text{K}$ [99] and to another phase above 352K [100]. However, the X-ray data was taken at room temperature where we performed our experiments. There is no temperature reported for the neutron-scattering data as it is unpublished and was only reported indirectly as a personal communication.

The spectral fitting software to determine the carbon Hamiltonian was continually improved as improvements in other areas of control revealed that for high fidelity control a more accurate Hamiltonian was needed. The Hamiltonian is fit by a model Hamiltonian with chemical shifts, dipolar couplings and isotropic J-couplings.

$$\mathcal{H}_{fit} = \sum_i \omega_i Z_i + \sum_{i<j} D_{ij} (2Z_i Z_j - X_i X_j - Y_i Y_j) + \sum_{i<j} J_{ij} (Z_i Z_j + X_i X_j + Y_i Y_j) \quad (4.12)$$

The electron mediated J-coupling is usually ignored in the solid state. It is usually an order of magnitude or more weaker than the direct dipolar couplings and to first order, the ZZ component can be somewhat absorbed in the dipolar term. However, we found higher fidelity control by including it in the fit. It is manifested in the different peak heights of the tented peaks. There is also an anisotropic component to the J-coupling which is identical in form to the dipolar (and is in fact sometimes called the pseudo-dipolar) and is absorbed by the dipolar term in the fit. Note however, this will affect how well the dipolar couplings obtained from

the fit match the geometry of the molecule. This correction was not taken into account. Reasonable values for the J-couplings can be obtained easily from liquid state experiments with the caveat that in the liquid state the two carboxyl carbons are equivalent.

Given the model Hamiltonian an FID is then simulated. The FID is relatively fast to create: the complete evolution is not simulated, rather only the amplitude of the observable terms (in the eigenbasis) are calculated and the difference in eigenfrequencies. Both Lorentzian and Gaussian line broadening are applied for each nucleus, although it should be noted that defining a T_2 for individual spins does not make sense in a strongly coupled system where the eigenstates are superpositions of computational basis states. A possible consequence of this was observed in some orientations where different line-widths within a peak group were seen. However, this may also be explained by incomplete proton decoupling as the proton decoupling power also affects the line-widths. Finally since the crystal also contains ^{13}C from natural abundance at $\approx 1\%$ which is at a comparable level to our dilution of processor molecules, natural abundance peaks were also added to the FID with their own amplitudes and T_2 's but constrained to the same ω_i . The parameters in the Hamiltonian, the amplitudes and the T_2 's are then optimized by a MATLAB non-linear least squares fitter using three possible fitness functions. The best fits are obtained in the frequency domain by fourier transforming and minimizing a least squares difference for the real part of the spectra. However, this fitness function has many local minima and performs best with a good initial guess. When one peak of the simulated spectrum lines up with the wrong peak of the experimental spectrum it is difficult to move out of that local minimum with a least-squares optimizer. The FID's themselves can be compared and do not suffer as badly from this issue. However, we obtained the best success for the initial fit (and also for fitting the forest of peaks in liquid crystal spectra), with integral curve methods [101]. In this approach a cumulative sum of the spectra is calculated and used for a least-squares fit. Because with each peak the cumulative sum monotonically increases, local minima are avoided. However, good baseline correction is needed for this method. A final fit of a spectrum is shown in Figure 4.3.

Given our limited ability to orient the crystal in the probe it was difficult to satisfy all the constraints on the Hamiltonian listed above. In the end a Goldilocks orientation was found which had just the right balance of everything: carbons were far enough apart in frequency space to be controllable, the coupling between $H_{m_1,2}$ was large enough for the transfer to be selective, yet not too large as to prevent efficient decoupling; the proton spectrum was sufficiently narrow to decouple, the coupling between C_2 and the protons sufficiently weak such that it could store polarization. Many otherwise fine orientations were ruined by short carbon T_2 's. These were caused by either susceptibility mismatch at the air/crystal interface or

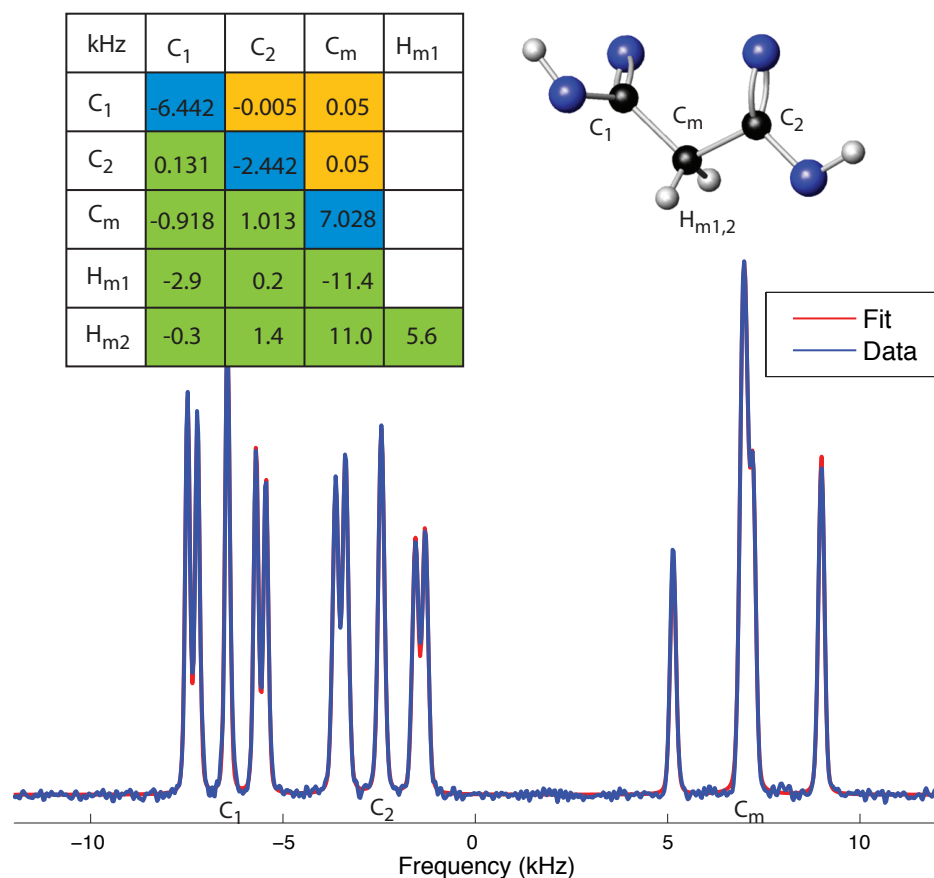


Figure 4.3: A proton-decoupled ^{13}C cross polarization spectrum and associated fit of malonic acid in the orientation used for the algorithmic cooling experiment. The field is 7.1T. The table gives the Hamiltonian values in kHz: Chemical shifts along diagonal, dipolar couplings below and J-couplings above. The chemical shifts are given with respect to the transmitter frequency chosen for the experiment: 75.438255MHz. To minimize the power requirements for the control we want the pulsing frequency to be centered in the spectrum. The relative intensity of the signal from the processor molecules and natural abundance peaks gives us a rough measure of the dilution of the processor molecules by assuming that the natural abundance is 1%. Since the processor molecules are split 4 times, the fact that they are 75% of the natural abundance intensity implies the processor molecules are 3% of the crystal.

dipolar interactions between processor molecules that were particularly strong at that orientation. The first issue could be overcome by the difficult task of shaping the crystal into a sphere [65]. The second leads to an intrinsic T_2 and is overcome by going to lower dilutions - limited by patience to overcome the reduced signal to noise and finally the natural abundance nuclei.

A quantum circuit description of the experiment is shown in Figure 4.4. Up to four rounds of cooling and compression could be run before the polarization of C_2 no longer increased. The carbons were initialized to infinite spin-temperature by rotating the carbon polarization into the plane and letting it dephase under the proton dipolar fields. The proton polarization was then rotated into the plane and held in place with a spin-locking pulse for the duration of the experiment. During carbon control operations, the spin-locking pulse serves the dual-purpose of preserving the proton magnetization and decoupling the protons. When a refresh sequence was required the proton B_1 field was smoothly reduced to the Hartman-Hahn matching condition as the carbon B_1 field was increased to the matching condition at a nutation frequency of $\approx 80\text{kHz}$. The matching condition was held for $30\mu\text{s}$ which empirically gave the maximum polarization transfer for the experimental crystal orientation. The amount of polarization transferred in the initial step determines the effective heat bath polarization. It can be compared to a thermal carbon spectrum to determine the amount transferred. The reset steps boosts the polarization by 3.3X which is less than the ideal 3.98X but the same as achieved with conventional CP. After the transfer sequence the proton polarization is smoothly returned to the high power spin-locking level of $\omega_{r.f.} \approx 250\text{kHz}$. Although higher proton powers are available and would improve the decoupling, they heat the probe changing the resonance condition and periodically produce arcing and capacitor breakdown.

The carbon control sequences are GRAPE optimal control pulses optimized numerically using the GRAPE code I wrote described in Appendix B. The most difficult part of the control is to achieve high fidelity over the wide range of inhomogeneities in the solid state. Without the availability of r.f. selection sequences (the protons are strongly coupled making r.f. selection sequences designed for a single spin useless), the pulses must be robust to the range of r.f. powers over the sample. Also, due to limited shimming ability and susceptibility mismatch due to a non-spherical sample, the T_2^* 's ($\approx 2\text{ms}$) are approximately two orders of magnitude shorter than the intrinsic T_2 's ($> 100\text{ms}$) [102]. By making the pulses robust to chemical shift variations, high fidelity control is available for times much longer than the T_2^* . After the reset sequences, carbon control is used to swap polarization from C_m to C_1 or C_2 and perform the three bit compression gate to boost the polarization on C_2 . The entire pulse sequence is shown in Figure 4.4.

The results of the experiment are shown in Figure 4.5. Because the CP reset sequence produces polarization along the spin-locking field direction from the CP

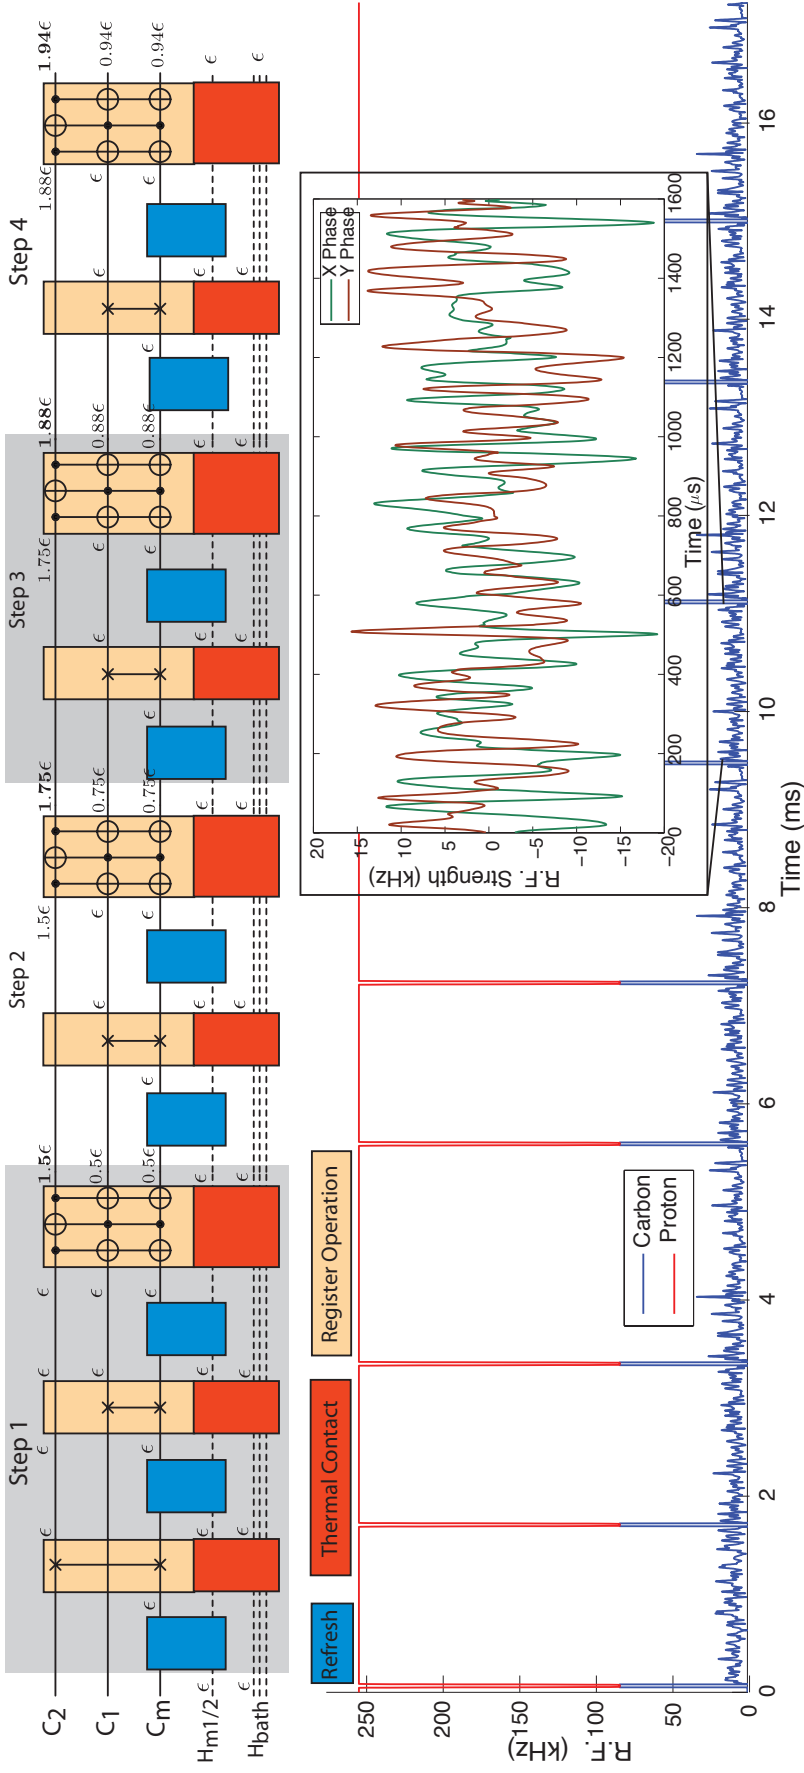


Figure 4.4: The quantum circuit and NMR pulses to implement four rounds of cooling and compression in the PPA algorithmic cooling protocol. The swap gates are 1.6ms long and the compression gate, 2.2ms . The compression gate is equivalent to a permutation of the diagonal elements of the density matrix and one possible implementation is shown decomposed as C-NOT-NOT and Toffoli gates [86]; however, here it was implemented as a single GRAPE pulse. The bottom trace shows the amplitude of the r.f. control fields for the pulse sequence. The inset shows in detail the two quadrature components of one of the GRAPE control pulses which implements a unitary swap gate between qubits 1 and 3. Each set of swaps and a compression gate is considered a step and the ideal polarization on the target qubit C_2 should increase as $1.5, 1.75, 1.88, 1.94$ in steps 1 through 4, respectively

field it was easier to build up polarization along the X direction in the rotating frame. This saved rotating the spins back and forth between the Z and X axes but demanded more rigorous timing of the pulses because the state was always evolving under the natural Hamiltonian. The unitary SWAP gates were used unchanged but the three qubit compression gate had to be modified by conjugating it with global 90 degree rotations (incorporated into the GRAPE pulse). The results were then easily obtained by simply switching the decoupling to the more efficient SPINAL64 sequence [66] and observing the proton-decoupled carbon spectrum. After the final reset step, the heat-bath polarization no longer needs to be preserved and the proton decoupling can be switched to the much more effective SPINAL64 decoupling sequence [66] for the final carbon compression control sequence. With the more efficient decoupling it is possible to boost the polarization of C_2 to $1.69 \pm 0.02X$ the heat-bath polarization - far beyond the 1.5 Shannon bound limit. The full results for each reset and boosting step of the protocol are shown in Table 4.1.




Figure 4.5: The resulting spectrum after four rounds of the algorithmic cooling protocol compared with the initial reset polarization spectrum. The initial reset is shown after being swapped to C_2 for comparison purposes. The polarizations in the table are given with respect to the initial polarization transfer or effective heat bath temperature. They were determined through a combination of integrals and scaling and all results have an uncertainty of ± 0.02 . Because C_1 and C_m have small couplings to the proton spin bath, the natural abundance spins on these nuclei also gain polarization during the transfer sequence. These natural abundance peaks undergo some unspecified unitary during the carbon control sequences and so give a randomly phased signal which interferes with the desired signal.

Step	C_1	C_2	C_m
1	0.10	0.04	1.00
2	0.11	0.98	0.05
3	0.21	0.98	0.98
4	0.96	0.94	0.22
5	0.97	0.93	0.98
6	0.47	1.39	0.49
7	0.51	1.37	0.97
8	0.96	1.34	0.50
9	0.96	1.32	0.97
10	0.68	1.56	0.71
11	0.70	1.53	0.97
12	0.95	1.53	0.72
13	0.97	1.50	0.97
14	0.76	1.64	0.79
15	0.78	1.63	0.96
16	0.92	1.62	0.79
17	0.94	1.57	0.95
18	0.78	1.67	0.85

Table 4.1: Table of polarization results from the three qubit algorithmic cooling protocol for each step. Each step refers to a reset or carbon control operation. The horizontal lines correspond to the “compression steps” which boost the polarization on C_2 . The polarizations were determined from an average of scaling and integration of the spectrum and comparison to the initial reset polarization and are uncertain to ± 0.02 . The results are shown for continuous spin-locking through the final step. As noted in the text, the decoupling during the final boost can be improved using the SPINAL64 sequence giving 1.67X boost for three rounds of cooling and compression and 1.69X boost for four rounds.

The protocol can now handily boost the polarization beyond the Shannon bound; however, it is still limited by issues specific to the malonic-acid system and not the control techniques or cooling algorithm. There are two main error sources in malonic acid. The first is that incomplete proton decoupling during the carbon control operations limits their fidelity. It is well-known that cw decoupling (the spin-locking pulses) provides a poor bandwidth to power ratio. It is particularly bad when the decoupled spins have strong dipolar couplings, as is the case here. Much more efficient decoupling sequences for the solid state such as TPPM or SPINAL64 have been developed, but unfortunately do not preserve the proton polarization. Without switching to SPINAL64 in the final compression step, C_2 is boosted to only 1.67X the heat-bath polarization. This is within the error of the 1.69X boost but is still indicative of the losses from incomplete decoupling. The second limitation is the non-ideality of the heat-bath. The proton spin bath is a finite size and so each polarization transfer reduces its polarization. This loss can be initially estimated as the ratio of carbons to protons. Given the $\approx 3\%$ triply labeled concentration this loss is $\approx 1\%$ per step. As the carbons are polarized this loss will decrease. However, the heat bath is continually heating during the course of the experiment due to relaxation under the spin-locking pulse from another relaxation process known as T_1 in the rotating frame or $T_{1\rho}$ [103]. This leads to a loss of heat-bath (1H) polarization. The combination of these effects in the experiment means the final refresh returns only 95% of the first refresh step.

4.3 Future directions for algorithmic cooling

The future directions for algorithmic cooling are to more qubits and higher reset polarizations. The increased polarization will almost certainly come from using electron spins: either because of its larger magnetic moment or optical handle. It is useful to consider the thermal polarization of an electron at different fields. This is shown in Figure 4.6. Also shown on the plot is the minimum polarization necessary to reach pure qubits (arbitrarily defined as 99% polarization from fault-tolerant threshold expectations) with algorithmic cooling. It is clear that with only 3 qubits, high fields and/or low temperature will be required. Given that coherent control becomes more difficult at these associated higher frequencies, it seems that relying on thermal polarization in bulk ensemble systems is an unlikely route. By increasing the number of qubits by only 2, to five, pure qubits are achievable with much lower reset polarizations. With any number of qubits it would be useful to investigate the possibility of a reset through optical pumping to potentially avoid the difficulty of bulk magnetic resonance at cryogenic temperatures. In nitrogen vacancy centers the ground state can be pumped to 80% probability through optical means [104].

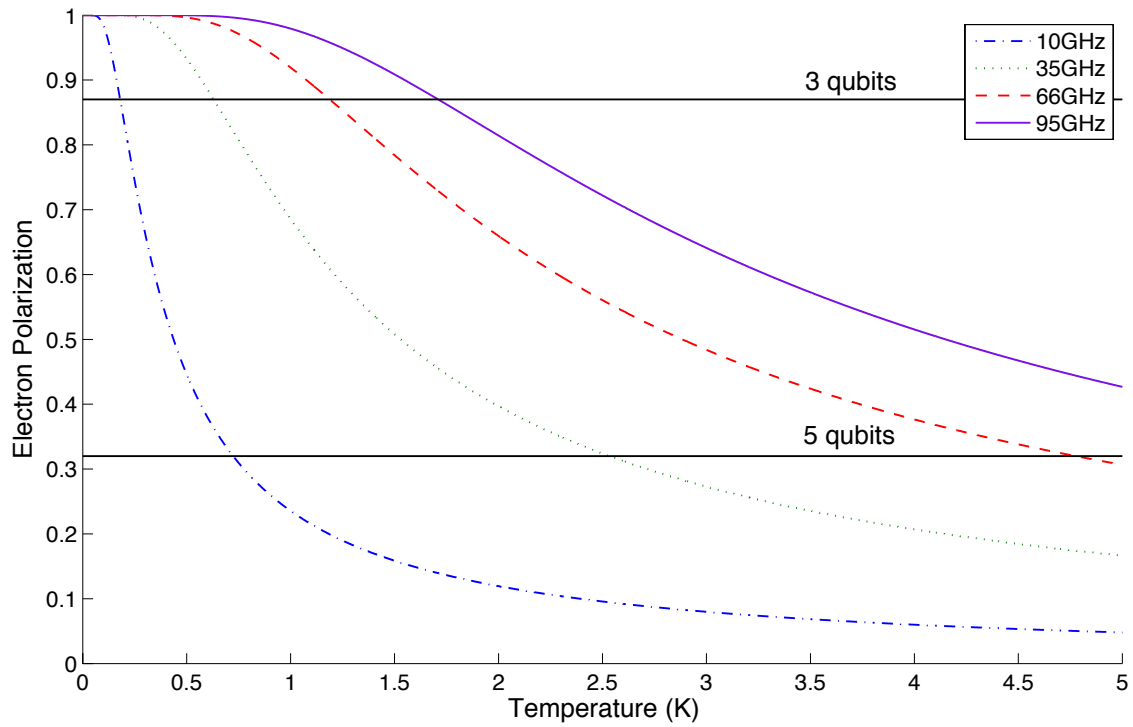


Figure 4.6: Electron polarization at different fields for temperatures from 0-5K. The fields are given in terms of typical ESR frequencies for $g = 2$ electrons. Also shown as black horizontal lines are the minimum reset polarizations necessary to achieve 99% polarization via algorithmic cooling assuming perfect control with 3 or 5 qubits. They give a sense of the reset polarization needed to achieve close to pure qubits using algorithmic cooling. As expected, with 5 qubits a much lower and more experimentally accessible polarization is possible.

In all of these electron-nuclear systems the issue of control must also be addressed. An attractive route is to rely on the electron as a control actuator i.e. control of the entire system will be found using only electron transitions. This greatly simplifies the control hardware and also allows much faster gates [105, 106]. However, universal control over the full Hilbert space is possible only when the nuclei have unique anisotropic hyperfine couplings to the electron. Furthermore, because the electron is coupled to the nuclei the decoherence properties of the nuclei are set by that of the electron. In particular, in the case of a weak ZZ coupling the nuclear T_2 will be set by the electron T_1 . Since any reset, optical or thermal is essentially a T_1 process, in this situation nuclear coherence cannot be preserved during an electron reset. Nuclear polarization though can be preserved. However, when the nuclei are controlled through the electron an anisotropic coupling is required in which case neither coherence nor polarization of the nucleus can be preserved perfectly under an electron T_1 process. Thus, for algorithmic cooling purposes, one may be forced into using both electron and nuclear control fields and a weak coupling regime.

Chapter 5

Conclusion

This thesis work has shown that in addition to the difficult technical and material science issues surrounding building a large-scale QIP, there are also fundamental and more abstract issues related to control and characterization that must be solved. The work in this thesis has experimentally shown that there are potentially efficient approaches to solving both the control and characterization problems. Ideas from randomization will allow us to overcome some of the characterization issues and extract relevant coarse-grained parameters from the system with a reasonable number of experiments. In particular, we can twirl a channel with random gates to obtain information about the probability of different Pauli weight errors occurring in the channel. This was applied to characterizing time-suspension sequences in solid-state NMR. Or, it was shown how applying random sequences of computational gates reveals the average error per gate and the highest fidelity control of any QIP implementation was demonstrated in liquid-state NMR. It has also shown in liquid-state NMR that by applying a subsystem style approach ideas from optimal control theory may be applied to larger systems in a scalable manner. Finally it was also shown that by obtaining high-fidelity control over a quantum system, it may be used to implement ideas about quantum algorithms and in particular one can use an algorithmic approach to overcome the state preparation problem in NMR and boost the polarization of a qubit far beyond its thermal equilibrium polarization.

The work in this thesis has also highlighted the immense technical challenge in integrating the control technologies to build a large-scale QIP with high enough fidelity control to demonstrate a quantum advantage. This is a large engineering challenge which may not be best addressed by small-scale research groups in a university setting. Nevertheless, seeing the great progress over the past five years, particularly in solid state devices such as superconducting qubits, one can only be optimistic that these challenges will be overcome in the near future.

Appendix A

Molecular Hamiltonians

Throughout this work, I have used numerous different molecules for both liquid- and solid-state NMR QIP. Below are two of the more important ones for liquid-state NMR. In the solid state the Hamiltonian depends on the orientation of the crystal so that situation is discussed in Chapter 4. In each case the Hamiltonian was obtained via spectral fitting of thermal spectra using homebuilt software.

A.1 Crotonic acid

Offering up to 7 qubits, good coherence times and good couplings, crotonic acid is our workhorse molecule for liquid state NMR experiments. There is a four chain carbon backbone, two proton spins and a methyl group. The three proton spins of the methyl group are magnetically equivalent and so form a spin $\frac{3}{2}$ and a spin $\frac{1}{2}$ subspace. Through an appropriate pulse sequence, the spin $\frac{1}{2}$ portion can be selected with high fidelity and the spin $\frac{3}{2}$ portion dephased [53]. From that point on the methyl group can be treated like another spin $\frac{1}{2}$ nucleus. Although this procedure gives another qubit it does constrain us to start with polarization on the methyl spin. One caveat to note with this molecule is that the hydroxyl group attached to C_4 usually exchanges with the deuterated solvent. This will lead to a different chemical shift and usually worse T_2 for C_4 .

A.2 TMSS

For the benchmarking protocol we were looking for a three qubit molecule with long T_2 's but not too long T_1 's and that would not require decoupling. This eliminated many candidates with quadrupolar spins such as chloroform because the quadrupolar's rapid T_1 tends to induce a rapid dephasing on nearby, coupled spins. However

	C_1	C_2	C_3	C_4	M	H_1	H_2
C_1	-3010						
C_2	41.6	-25630					
C_3	1.5	69.6	-21541				
C_4	7.0	1.2	72.3	-29552			
M	127.2	-7.1	6.6	-0.9	-1317		
H_1	3.9	155.6	-1.8	6.5	6.8	-4897	
H_2	6.3	-0.7	161.5	3.3	-1.7	15.5	-4101




Figure A.1: Hamiltonian parameters (Hz) for the crotonic acid molecule in a 700MHz spectrometer. Diagonal elements are chemical shifts with respect the base frequency for that nuclear type (700.13MHz for proton and 176.047829MHz for carbon) while off-diagonal elements give the J-coupling constants. Note some of the J-couplings are negative which is often ignored in the literature but is important for high fidelity control. The yellow coloured nuclei are oxygen which have almost 100% natural abundance spin 0 and so do not explicitly appear in the spin Hamiltonian.

we fortunately found the molecule selectively labelled ^{13}C tris(trimethylsilyl)silane-acetylene from previous work in the Cory group [81]. The relevant part of the molecule for us is the proton and two ^{13}C -labelled carbons joined by a triple bond which gives us three qubits. The rest of the molecule can essentially be ignored: the methyl protons are too far away to have a resolvable coupling to the system of interest. The natural abundance ^{29}Si , occurring at about 4%, have a resolvable coupling to the system. However, the parts that can be spectrally resolved can be ignored by not measuring the integral in that region and the unresolved couplings can be handled by treating the Si incoherently and ensuring the control pulses are robust to small chemical shifts. One quirk of this molecule is that due to the large size of the silicon section of the molecule it does not rotate and tumble freely. This, combined with the large CSA of the ethylene group leads to a decoherence rate on C_1 which depends on the state of H . This gives different line-widths in the C_1 peak group depending on the state of H .

	H	C_1	C_2
H	-0.1		
C_1	236.4	-1035.0	
C_2	42.2	132.5	1040.5
T_1	2.4	5.3	5.3
T_2	2.0	2.0	2.4

Figure A.2: Table of natural Hamiltonian parameters (Hz) obtained from spectral fitting on a 700MHz magnet. The diagonal elements give the chemical shifts with respect to the transmitter frequencies while the off-diagonal elements give the J-couplings. The transmitter frequencies were chosen in the center of the spectrum for control convenience and are -1630Hz and -15832Hz from the proton and carbon transmitter base frequencies specified in Figure A.1. T_1 's and T_2 's (seconds) are measured from standard inversion recovery and CPMG echo sequences respectively.

Appendix B

GRAPE Optimal Control Code

An important component of implementing high fidelity control was the ability to find robust high fidelity control pulses via optimal control theory. Our previous approach was borrowed from the strongly modulating composite pulses (SMP) technique [19]. The MATLAB code from MIT for finding SMP's was improved in our lab but had trouble finding high fidelity control pulses for systems larger than three qubits and even for three qubits in the solid state it had difficulty finding pulses above 99% fidelity. This was not due to a limitation of the strongly modulating pulse model; indeed, for some purposes it is necessary to strongly modulate the system to refocus unknown inhomogeneities and for high fidelity implementation. However, the SMP code relied on finite differences to calculate gradients and a simplex search. The application of standard optimal control ideas to the problem of synthesizing quantum control operations in NMR from Khaneja et al. [22] was a great advance and allowed us to expand our high fidelity control up seven qubits with fidelities that are limited by experimental constraints on how well we can implement the control field and natural decoherence rather than the control sequence itself.

Numerical optimization has a few key steps:

1. Start with a random or well-chosen guess for the pulse shape.
2. Calculate the value of the fitness function (overlap of the simulated unitary with the desired unitary) for the pulse shape.
3. Make a small change to the pulse shape that improves the fitness function.
4. Repeat steps 2 and 3 until no more improvements are possible.
5. If the best value of the fitness function found is good enough then the pulse is accepted otherwise start again with a new guess at step 1.

For numerical optimization of a pulse sequence the key information we need at each step is which direction to push the control fields to improve the fidelity. This could be answered with a finite difference gradient but, as discussed in the text (see section 2.3), the GRAPE algorithm’s key improvement is to take advantage of the fact that the simulated unitary is a product of the unitary for each timestep. Thus the derivative can be moved inside the product. If the cumulative propagator is then stored for each time step, then calculating the derivative for each time step (with either the first-order approximate derivative discussed in the text or a simple finite difference) does not require resimulating the entire quantum evolution.

As part of my thesis work I wrote an implementation of the GRAPE code in MATLAB. The full code is available on request from the author (c4ryan@iqc.ca) or author’s supervisor (laflamme@iqc.ca). MATLAB is an interpreted language and therefore might not be ideally suited for numerical work. However, careful use of vectorized code and MATLAB’s “Just-In-Time Compiler” which precompiles loops of code (see MATLAB documentation for more information on both topics) makes the MATLAB code fast enough for our current purposes. The code is closely based on the algorithm presented in Ref. [22]. In this appendix I present some of the extensions and improvements I made.

B.1 Usage

The code requires several auxiliary functions for setting up the Hamiltonian, taking the partial trace, etc. These are either loaded as part the Laflamme group’s NMR simulator package or can be separately loaded by adding the auxiliary_files subfolder to the MATLAB path. It is also useful to set two global variables *pulses* and *pulseguess*. *pulses* is used to access the pulses found if the program is stopped before it finishes and *pulseguess* is used to feed in initial starting points for the search. If no starting point is known then *pulseguess* should be set to empty. The code is then called with `foundpulses = pulsefinder_subsystem('paramsfile');`. The *paramsfile* is the name of a parameter file where all the options for the program are specified (see below for an example). The program will return a cell array of pulse structures which contain the following fields: *initialguess* is the starting point for the search; *params*, the parameters for the search; *pulse*, the last pulse used in the search; *zangles*, the pre and post single qubit Z rotations if the *Zfreedom* flag is turned on (see below).

B.2 Code Details

The params file is a separate .m file which sets all the necessary options for the pulse finding program. It sets all the fields of a global variable *params*. The file is evaluated in the main function and the global params variable is accessed to load the options. An example file for the crotonic acid system is shown below:

```
global params pulseguess

%Reset the params structure
params = [];

%Location of the nuclei file (full path name)
params.nucleifile = '-/Programs/pulsefinder/molecule.def';

%Number of timesteps in the pulse
params.plength = 100;

%Length of each time step in seconds
params.timestep = 1e-6;

%Initial stepsize (this is reasonably important - run some tries
%and choose a value which is close to what the program is choosing
%after 50 or so iterations)
params.stepsize = 1e-1;

%Desired unitary
%H1U
params.Uwant = expm(-i*(pi/2)*(full(mkstate('+1IXIIIII',0))));

params.subsystem{1} = [1 4 5];
params.subsystem{2} = [2 3 5 6];
params.subsystem{3} = [6 7];
params.subsys.weight = [3 4 2];

%params.subsys.weight = [1];
%params.subsystem{1} = [1 2 3 4 5 6 7];

%Desired fidelity for the unitary
% (this is the trace squared fidelity  $F = \text{abs}(U_{\text{goal}}^{\dagger} U_{\text{sim}})^2 / N^2$ )
params.fidelity = 0.99925;

%RF distribution to optimize over
% (will slow down search and convergence dramatically)
%Two dimensional array
% first column is percentage of sample; second
%column is percentage of rf strength it sees.
params.rfdist = [0.3 0.97;0.4 1.00;0.3 1.03];

%params.rfdist = [1 1];

%Hamiltonian distribution to optimize over

params.Hamdistr = [1];
params.Hammatts{1} = 0;

%params.Hamdistr = (1/4)*ones(1,4);
```



```

%Matrices for robustness to Hamiltonian distributions. These will be
%multiplied by 2PI and added to the natural Hamiltonian.
% params.Hammatts{1} = 75*(1/2)*full(mkstate('+0.8ZI+1IZ',0));
% params.Hammatts{2} = 75*(1/2)*full(mkstate('+0.8ZI-1IZ',0));
% params.Hammatts{3} = 75*(1/2)*full(mkstate('-0.8ZI+1IZ',0));
% params.Hammatts{4} = 75*(1/2)*full(mkstate('-0.8ZI-1IZ',0));

%Spins in nuclei file to ignore in search
params.spins.ignore = {'RH','RC'};

%RF control fields (3D array of control Hamiltonians – as many as you like)
params.RFmatts(:, :, 3) = ...
(1/2)*full(mkstate('+1IIIXIII+1IIIIIXII+1IIIIIXI+1IIIIIX',0));
params.RFmatts(:, :, 4) = ...
(1/2)*full(mkstate('+1IIIIYIII+1IIIIYII+1IIIIYI+1IIIIY',0));
params.RFmatts(:, :, 1) = (1/2)*full(mkstate('+1XIIIIII+1XIIIIII+1IIXIIII',0));
params.RFmatts(:, :, 2) = (1/2)*full(mkstate('+1YIIIIII+1YIIIIII+1IYYIIII',0));

%The maximum rf power for each rf field in rad/s
params.rfmax = 2*pi*[25e3 25e3 16.7e3 16.7e3 16.7e3 20e3];

%Some parameters for the random guess
%Scale of the random guess for each RFmatt (between 0 and 1)
params.randscale = [0.05 0.05 0.05 0.05 0.1 0.2];

%Choose every randevery points at random (rest will be fit to cubic spline)
params.randevery = 25;

%Tolerance for improving
% i.e. if over 20 tries we are not improving by an average of at least this,
%we will try a new random starting point
params.improvechk = 1e-8;

%Minimum stepsize (if we're not moving anywhere we should stop searching)
params.minstepsize = 1e-8;

%Number of random guesses to try before giving up
params.numtry = 5;

%A vector of the pulsing frequency for each spin (this is defined
%with respect to the same frequency your nuclei file is)
params.pulsefreq = [-3000 -3000 -3000 -16000 -16000 -16000 -16000];

%Soft pulse buffering delay (delays required before and after soft pulses.
% (Since our time periods are usually greater than 350ns we could probably
%try to use fast shapes)
params.softpulsebuffer = 10e-6;

%If there is a starting guess for the pulse load it in here (should
%be a structure
params.pulseguess = [pulseguess];

%The type of pulse we are searching for (1 for unitary, 2 for state
%to state)
params.searchtype = 1;

%Input and goal states for state to state
params.rhoin = mkstate('+1XII',1);
params.rhogoal = mkstate('+1YII',1);

%Flag for whether you want to allow the time steps to vary
params.tstepflag = 0;

```

```

%Maximum length of pulse (rather soft boundary)
params.tpulsemax = 20e-3;

%Allow Zfreedom or not
params.Zfreedomflag = 1;

%Flag to force beginning and end of pulse to zero
params.onofframps = 1;

% How many seconds to force a display output
params.dispevery = 60;

```

Most of the parameters are self-explanatory from the comments. There is very little error checking or default values built into the program so nearly all the fields have to be properly defined whether they are used or not.

- The *subsystem* field is a cell array denoting which nuclei belong to each sub-system. The desired unitary will automatically be broken up into the sub-system unitaries and a new fidelity function is defined as the sum of the fidelities of the pulse on each sub-system weighted by the *subsys_weight* field. See below for implementation details.
- **mkstate** is a separate function which interprets a text string of Pauli matrices in the product-operator notation and converts it to a numerical matrix.
- The Hamiltonian distribution can be used for adding robustness to inhomogeneities or uncertainties in the natural Hamiltonian. For example in the solid state it can be used to ensure robustness to static field inhomogeneities. The commented out section shows using it to provide robustness to J-couplings to other spins treated incoherently by averaging over all possible states of the other two spins. The Hamiltonian distribution is evaluated for each r.f. distribution point so when the two inhomogeneities are used in combination it can substantially increase the search time.
- The *softpulsebuffer* is used to allow for the delays that the spectrometer requires around shaped pulses to load and process the pulse. This delay is approximately $4\mu\text{s}$ per pulse per channel before and after the pulses, i.e. for parallel pulses at least $8\mu\text{s}$ is needed before and after. The spectrometer will sometimes complain about this delay being missing but note that it will also sometimes fail silently and simply not implement the pulse. Because the timesteps of the pulses are usually greater than $1\mu\text{s}$ it should be possible to use so-called ‘fast-shapes’ which do not require this delay. However, I have found that there are timing issues with ‘fast-shapes’ and they are implemented too quickly. For example, $1\mu\text{s}$ timesteps becomes $0.95\mu\text{s}$ timesteps.

- The *searchtype* field is a flag specifying which fitness function to use: state to state or unitary. Sometimes we know the initial state and desired final state (for say state preparation sequences). In the case of state to state, it is a much easier task to drive the system for only one input state then to worry that all possible input states are correctly evolved as in the case of a unitary goal. If using state to state then the fitness function maximizes the absolute value of the correlation between the simulated output state and the *rhogoal* field starting from the *rhoiin* field input state. This can mean that the negative of the state is achieved. For example, we may want *ZII* to go to *ZZZ* but the code may find a pulse which gives $-ZZZ$.
- The *tstepflag* field is a flag specifying whether to allow the time for each time step to vary. This can be useful when the total length of time needed for the pulse is unknown. However, it slows down the search from the extra derivative evaluation and larger search space. The maximum pulse time is set by the *tpulsemax* field.
- The *Zfreedomflag* allows individual *Z* rotations to be applied for free before and after the pulse. When working with the sequence compiler individual *Z* rotations can be applied instantaneously and accurately with an abstract frame change. Giving the search this freedom makes certain pulses much easier to find because the chemical shifts do not have to be refocussed by the control field, but instead can be tracked and accounted for by the *Z* rotations.
- The *onofframps* determines whether the beginning and end of the pulse are forced to zero with a penalty function. This is useful for an accurate implementation because the control field cannot be turned on/off instantly.

The main file of the code is divided into several subfunctions as follows:

1. Main **pulsefinder**. This function evaluates the parameters file to load all the options, sets up the initial guess, and runs the main optimization loop
 - The system is broken down into its subsystems here. The natural Hamiltonian and control Hamiltonians for the subsystems are obtained by taking the partial trace over the complement of the subsystem spins. The partial trace approach does not work for obtaining the subsystem unitaries as the traced over spins may have a traceless unitary. For example, consider the desired unitary being a 180 degree pulse about the *x*-axis on spin 1. For any subsystem that does not include spin 1 the partial trace will return zero because *X* is traceless. A solution is apply random single-qubit unitaries to all the spins that will be traced out. This

ensures that it is very unlikely that the traced out spins have a traceless unitary and the partial trace can then be taken. The answer must then be normalized to return a valid unitary. The process does not employ any error checking to ensure that the desired unitary is compatible with the sub-system decomposition.

- The initial guess is either loaded from the *pulseguess* field of the params input or is set to a random guess determined by the *randscale* and *randevery* fields. To obtain smooth low-power pulses which are desirable for experimental implementation it is important to start with a smooth low-power guess. The guess is determined by choosing random points between $\pm \text{randscale} * r_{\text{fmax}}$. These are set at every *randevery* number of points of the starting pulse along with 0 as the start and end points. The rest of the points are interpolated using a cubic spline.
- The main optimization loop will keep running optimizations with new random guesses until a pulse is found with fidelity greater than the *fidelity* field in params, or the number of optimizations tried exceeds the *numtry* field.
- The main optimization loop calls the **evalpulse** subfunction which returns the current value of the fitness function and derivatives of all the control variables. The pulse is then moved in the direction of the gradient. For fastest convergence and proper use of the conjugate gradients then the pulse should be moved to the highest fidelity point along the search direction. This is a line search problem for which there are standard solutions but it is difficult because each function evaluation requires a simulation of the quantum system and so is expensive. The best move is determined from a quadratic interpolation of the fitness function along the search direction. The fitness function is evaluated at two further points: one times the current step size along the gradient and twice the step size. A quadratic polynomial is fit to the three points and the maximum of the quadratic is determined. If the maximum is beyond two times the step size the fitness function is evaluated at four times the stepsize and a second quadratic fit to the points. If the maximum is still beyond 4 times the step size then the pulse is moved to 4 times the step size. If the quadratic produces a minimum instead of a maximum, or the maximum step size is less than 0.1, then the pulse is moved by 0.1 times the step size. The step size is dynamically adjusted during the search by updating it by $stepsize_{\text{new}} = \sqrt{\text{currentmove}} \times stepsize_{\text{old}}$ so that the step size is adjusted slowly to the correct range but does not wildly jump around.
- The best current direction of the search is not always simply the gradient

and so the code employs the conjugate-gradient technique [107]. The best direction is a linear combination of the gradient (steepest ascent) and a linear sum of all the previous search directions. One can visually imagine this as climbing a hill along a ridge. If one is not perfectly on the crest of the ridge then the steepest ascent will lead to one to zig-zag back and forth across the crest of the ridge, slowly making one's way to the top. However, by looking at all previous moves the climber will quickly realize that the average direction is simply up along the ridge. In general, the conjugate-gradient converges much more quickly around the maximum where the fitness function surface is approximately quadratic. Further away from the maximum if the surface is not quadratic, the search method may lose conjugacy and break down, i.e. the previous moves are not useful. In these cases the search must be restarted along the gradient direction. The code implements a version of the Polak-Ribiere formula. If λ_k is the search direction for step k and ∇_k is the gradient, then

$$\lambda_k = \nabla_k + \beta_k * \lambda_{k-1}. \quad (\text{B.1})$$

The Polak-Ribiere formula for β_k is

$$\beta_k^{PR} = \frac{\nabla_k^T (\nabla_k - \nabla_{k-1})}{\|\nabla_{k+1}\|^2} \quad (\text{B.2})$$

The formula is modified if β is less than zero where it is assumed that conjugacy has been lost and the algorithm is reset by setting $\beta = 0$. The conjugate gradient technique provided approximately an order of magnitude speed-up in convergence over a simple steepest ascent algorithm.

- Finally the main loop prints out information about the current value of the fitness function, step size and improvement rate at a regular interval. The interval is set but the *dispevery* field in params or when the difference between the current fitness function and the goal has been reduced by 20% since the last information was printed.
- The current pulse is stored in the global variable *pulses* so that it can be accessed even when the pulsefinding process is killed due to impatience at the rate of convergence.

2. **evalpulse**. This subfunction evaluates the current fidelity of the pulse and calls **calcdervs** to obtain the derivative information.

- The buffer delays for the soft pulse are handled by modifying the goal and starting unitary (or goal and starting state) by applying negative time evolution of the natural Hamiltonian for the soft pulse buffer time. This ensures that the pulse itself will compensate for the evolution during

that time so that the overall unitary (or state transfer) for the total time is as desired.

- The fitness function and derivatives are evaluated in a double loop over the r.f. and Hamiltonian distributions and the returned values are an appropriately weighted sum.
 - If activated in the params structure the penalty functions which limit the total time of the pulse and the ensure the beginning and end of the pulse are low power are implemented here. Both are implemented as cosh functions with some heuristic parameters. The on/off ramps can be further fine tuned with three hard-coded parameters: *param1*, *param2* and *numpts*. The *numpts* specifies how many points at the beginning and end of the pulse are affected by the penalty function. The effect of the penalty function decreases exponentially away from the endpoints. *param1* sets the overall strength of the penalty function: values from 1e-4-1e-3 are reasonable. *param2* sets how rapidly the penalty functions decays away from the endpoints: values of 0.1-0.25 are reasonable.
3. **calcderviv**. This subfunction evaluates the derivatives of the control fields or time step lengths for whichever fitness function is specified. Currently it uses only the approximate derivative discussed in the text which limits the length of the time steps such that the condition $|\mathcal{H}\Delta t| \ll 1$ is satisfied. However, it would be useful to add the ability to calculate finite difference derivatives. This would be slightly slower (equivalent to doing three simulations versus two per iteration) but would allow long time steps which would be useful for longer pulses implementing coupling gates.
 4. **some_extra_stuff**. This subfunction provides some extra code for debugging purposes. In particular it calculates the derivatives via finite differencing so that the approximate derivatives can be checked. It also plots the fitness function along the search direction so that the quadratic interpolation and step size can be checked.

Appendix C

Pulse Implementation Correction

For high fidelity control and a good model for simulations of our NMR systems it is important that the r.f. control fields at the sample match what we think they are. Unfortunately, non-linearities in the pulse generation and amplification, and bandwidth constraints of the probe-resonant circuit, mean that there can be large discrepancies between the pulse sequence that we ask the spectrometer to implement and the fields that the spins in the sample actually see. This is solved by measuring the field at the sample and closing a feedback loop which will iteratively adjust the control pulses so that the field at the sample matches the simulated pulse. This appendix gives a few technical details of this unconventional experiment.

The field is detected with a pick-up coil placed close to the sample r.f. coils. A wire is wound into a few (2-4) turn coil approximately 3mm in diameter. The ends of the coil are connected to short the inner and outer conductor of a coaxial cable. The coil has a flat transmission characteristic across the relevant r.f. frequencies. In the solid state, the coaxial cable is then lowered into the bore of the magnet and visually placed within a few centimeters of the sample coil. In the liquid state we have to be more precise. We use a thin coaxial cable which will fit into a 5mm sample tube and then use epoxy to hold the cable in place such that the pickup coil is a few mm above the r.f. sample coils when the sample tube is loaded into the probe via the normal pneumatic lift. The sample tube may be filled with solvent to mimic the experimental conditions; however, I have found that the solvents degrade the coaxial cable and electric tape and that having an empty tube is sufficient. In either case, since the transmission between the r.f. coils and the pickup coil is extremely sensitive to position, it may require some trial and error to find the ideal position. Furthermore, for consistency, the cable should be fixed in place with tape to ensure the coil does not move during the feedback process.

Once the pickup coil is in place we use the spectrometer to detect the r.f. pulse. This avoids having to synchronize the spectrometer and a high-speed oscilloscope,

but will hide certain timing and frequency issues because we are using the spectrometer to test itself. We are also relying on the linearity of the receiver train in the spectrometer. A schematic of the experimental setup is shown in Figure C.1. Because of some unexplained issues in the Bruker pulse generation and routing, the behaviour of the of pulse depends strongly on which signal generation unit (SGU) is used for the pulse and what other SGU's are connected to the same amplifier. Therefore, it is important to use the same SGU/router combination that will be used in the experiment. For example, for proton pulses we can send the pulse via logical channel :f2 through SGU1 and then the proton amplifier. We must then use SGU1 for generating the proton pulses in the actual experiment. Unfortunately, the dwell pulses from the SGU cause the pre-amplifier stage connected to the SGU to switch to receive mode which will interfere with the forward power of the pulse. To prevent this, the SGU generating the dwell pulses for the receiver should be connected to a different pre-amplifier slice. This does not have to be one of the same nuclear type. For example, on the 700, logical channel :f1 is set to proton to receive using SGU4 for generating the dwell pulses and SGU4 is connected to the nitrogen preamplifier slice. In the solid state, with specialized high speed preamplifiers made for stroboscopic observation, this is not a concern.

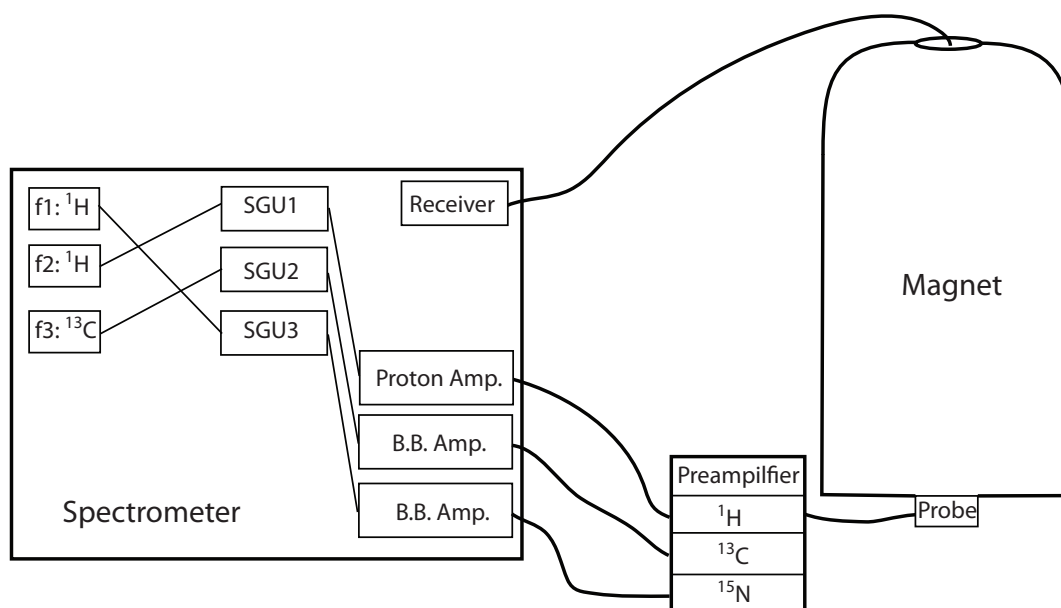


Figure C.1: Schematic of spectrometer setup for pulse correction of a proton pulse. The offset frequencies should be set to the same value for the two proton channels. The B.B. amplifiers are the broadband amplifiers.

The pick-up coil is connected directly to the r.f. input of the receiver. To avoid

wear and tear on the SMA connector on the receiver we connect a short length of SMA coaxial cable to the receiver and swap between the normal and pickup coil connections on the cable rather than the receiver directly.

I have written Bruker AU software as well as MATLAB code to automate the “pulsefixing” procedure. A dataset must be setup for stroboscopic observation of the pulse and the pulse power and pulse program setup correctly. Since we are feeding r.f. power directly into the receiver without protection it is important to be careful and start with a low power and then slowly bring the power up to the calibrated pulsing power. Since the pulse generation non-linearities depend on pulse power, the pulses must be fixed at the correctly calibrated powers that they will be used with in the experiment. The calibration may not be known in advance, and so this may require several iterations of pulsefixing and calibration. For calibration purposes I recommend pulsefixing a short ($\approx 100 - 150\mu\text{s}$ low power ($\approx 20\%$ maximum power) simple (e.g. truncated Gaussian) pulse shape that will implement a 180 degree pulse when properly calibrated.

The Bruker AU program **pulsefixer_grape** prompts the user for the pulse name, the pulse time and the number of feedback loops to perform (usually 4-8 is sufficient). The pulse fixer AU program implements several steps:

1. It first runs a simple rectangular pulse as an amplitude and phase reference. Ideally this means that all the pulses will be consistently fixed no matter how the pick-up coil is setup. The FID is sent to a separate remote computer which has MATLAB installed and a MATLAB program **fitref.m** fits the middle portion of the pulse (to avoid transients) and saves the data so that all future feedback iterations for this pulse are scaled and phased according to this reference.
2. The pulse is then changed to the pulse to be corrected, the acquisition time is set correctly and the pulse itself is then run.
3. The FID file is again sent to the remote computer where the MATLAB program **pulsefixer_grape.m** is run. The matlab code compares the scaled and phased measured pulse with the desired one and creates a new pulse to try and compensate for the pulse generation non-linearities. The new pulse can be generated in one of two ways. A simple point to point comparison can be made between the measured and ideal pulses and the new pulse is generated as the tried pulse minus some damping constant times the difference. This method can lead to correction of noise and introduction of noise into the pulse. However, this can be solved by smoothing the output pulse. Another approach is to fit both the measured and ideal pulse to overlapping cubic polynomials. The number of points to fit and the overlap are parameters that

can be set. The output pulse is then the tried pulse minus some damping constant times a polynomial generated from the difference in the polynomial parameters. The overlapping points are simply averaged. This method is more robust to noise but is much slower. In either case the damping is some constant less than one to ensure the feedback loop is stable (i.e. the pulses do not oscillate around the ideal set point).

4. The AU program then retrieves the new pulse and runs it and then repeats the feedback loop the set number of times.

For some systems we have observed strong signal at exactly the spectrometer proton frequency (e.g. 700MHz or 600MHz). The signal appears to originate within the spectrometer itself in the receive train because it appears when nothing is connected to the receiver input port. Therefore it is filtered out with a notch filter in the **pulsefixer_grape** program.

Appendix D

Papers with significant contribution

This appendix contains papers to which I contributed significantly but were not directly discussed in this thesis.

D.1 Benchmarking Quantum Control Methods on a 12-Qubit System

This work was a collaboration between the Waterloo and MIT groups to explore how far we could push our control techniques. Both groups applied the pseudo-pure state benchmark [28] to the same molecule histidine, which offers 13 possible qubits. The MIT group used an early form of numerically optimized pulses known as strongly modulating pulses [19] whereas the Waterloo group used only simple shaped pulses and the sequence compiler. The MIT group was able to create and observe a twelve qubit pseudo pure state whereas we at Waterloo were only able to achieve a 10 qubit pseudo pure state before control errors, decoherence and the exponential signal loss overcame the signal. This would seem to validate the success of numerically optimized pulses and was the instigation of our investigations of how to combine the two control techniques. A solution to this was presented in Chapter 2 of the thesis.

My work involved accurately working out the natural Hamiltonian (J-coupling signs) with a series of targeted small (2-4 qubits) pseudo-pure states. Because each qubit in a pseudo-pure state reduces the number of peaks by a factor of two, creating pseudo-pure state can dramatically clear up complicated spectra making the Hamiltonian assignment problem easier. I helped work on optimizing pulse

sequences for the pseudo-pure states. I also helped design the protocol for accurate measurement of the signal recovered at each step (complicated by orders of magnitude difference in signal intensity between the reference and final state) and comparison to expected values. Furthermore, we were able to show that with a simple independent dephasing model we could qualitatively explain how the signal loss behaved with the number of qubits.

Benchmarking Quantum Control Methods on a 12-Qubit System

C. Negrevergne,¹ T. S. Mahesh,² C. A. Ryan,¹ M. Ditty,¹ F. Cyr-Racine,¹ W. Power,¹ N. Boulant,²
T. Havel,² D. G. Cory,² and R. Laflamme^{1,3}

¹*Institute for Quantum Computing, University of Waterloo, Waterloo, ON, N2L 3G1, Canada*

²*Department of Nuclear Engineering, MIT, Cambridge, Massachusetts 02139, USA*

³*Perimeter Institute for Theoretical Physics, Waterloo, ON, N2J 2W9, Canada*

(Received 17 December 2005; published 1 May 2006)

In this Letter, we present an experimental benchmark of operational control methods in quantum information processors extended up to 12 qubits. We implement universal control of this large Hilbert space using two complementary approaches and discuss their accuracy and scalability. Despite decoherence, we were able to reach a 12-coherence state (or a 12-qubit pseudopure cat state) and decode it into an 11 qubit plus one qutrit pseudopure state using liquid state nuclear magnetic resonance quantum information processors.

DOI: [10.1103/PhysRevLett.96.170501](https://doi.org/10.1103/PhysRevLett.96.170501)

PACS numbers: 03.67.Lx, 76.60.Pc

Quantum mechanics promises information processors that are more efficient than any known classical devices. However, to bring this potential to reality, we must learn how to control large quantum systems in a scalable fashion. Scalability has at least two components: the complexity of the methods used to obtain coherent control must grow only polynomially with the number of qubits involved, and the errors occurring during the implementation of the control sequence must be small enough to be correctable. These errors can be split in two classes: first, the operational errors due to imperfections in the control procedure and, second, intrinsic errors due to decoherence and relaxation processes. Benchmarking small quantum information processor (QIP) prototypes [1–4] is therefore crucial to characterizing the errors in a physical system and developing general quantum control methods. In a physical system well suited for implementing a QIP, once we have reached an optimal operational control, we will need to take care of intrinsic errors using quantum error correction procedures [5].

Because they have the ability to run nontrivial quantum algorithms, liquid state nuclear magnetic resonance (NMR) based QIPs [6,7] can be used as benchmark systems [1,8,9]. In the present work, we are interested in optimizing operational control strategies in terms of accuracy and the amount of required classical resources. To do so, we have chosen to extend the benchmarking algorithm previously used on a 7-qubit liquid state NMR register [1] to 11 qubits plus one qutrit. This algorithm consists of preparing mixtures of generalized Greenberger-Horne-Zeilinger (GHZ) states of the form: $\rho_{\text{GHZ}} = \mathbf{I}^{\otimes n} + \mathbf{X}^{\otimes n}$ (\mathbf{I} is the identity matrix, \mathbf{X} is the σ_x Pauli matrix, and n is the number of qubits involved in the GHZ states). This state preparation is very similar to the generation of stabilizer operators [10], which are building blocks for quantum error correction codes. Furthermore, this algorithm takes the state of the quantum system to the most fragile reaches of the Hilbert space we are operating in and, therefore,

clearly demonstrates coherent control. Previous work has demonstrated a 12-spin pseudocat state [11] and multiple quantum coherences of much higher order [12]. However, these exploited symmetries in the systems which limited them to a much smaller symmetric subspace of the full Hilbert space. In the present Letter, we benchmark *universal* control methods that allow us to access the full Hilbert space of our system.

In a liquid state NMR QIP, universal control is achieved through the application of a coordinated sequence of radio-frequency (rf) pulses and periods of free evolution. The resulting 1- and 2-qubit gates allow us, in principle, to implement any unitary transformation [13]. The challenge is to efficiently design such pulse sequences to be as short as possible and robust against experimental imperfections in order to minimize systematic error and decoherence [14,15]. In a 3-qubit experiment [16], it is possible to write the pulse sequences by hand and compensate for experimental errors with a few optimization parameters. Moving to larger registers [1,8] requires more complex control schemes that necessitate systematic numerical optimization in the design of the pulse sequence. Going to 12 qubits represents a substantial step forward in the number of quantum degrees of freedom that are controlled.

We will approach coherent control over this large Hilbert space system from two complementary points of view. First, to demonstrate that control methods of high precision are available and experimentally realizable, we build a detailed model of the experimental QIP and, for each desired unitary operation, search for an optimal control sequence based on strongly modulating pulses [17]. Applied over the entire Hilbert space, this approach is not scalable. The amount of classical resources required to search for control sequences grows as the size of the Hilbert space—i.e., exponentially with the number of qubits. However, this approach returns control sequences of high fidelity and with small, known errors, provided our system model is accurate. Because of the exponential

computing cost of determining a suitable pulse, when dealing with a large Hilbert space, we have to lower the dimension of the space over which we search. This can be achieved by searching for pulses only on a subsystem of the spins of the register (in the present case, the carbon nuclear spins) and check that it leads to sufficiently high fidelity control by simulating the effect of rest of the spins as described in Ref. [18].

A second approach to control such a large Hilbert space is to make a series of well-constructed simplifications to the model, in order to permit control sequences to be developed with a complexity that grows only polynomially with the number of qubits. We therefore based the design of our pulse sequences on the method developed in Ref. [1]. Indeed, by using only simple pulses (broadband rectangular hard pulses and selective soft Gaussian shaped pulses) and performing a series of simulations on pairs of spins with significant couplings for each pulse, it is possible to efficiently determine first-order deviations from the ideal pulse. For each of the pulses, these control errors can be represented as phase shifts and spin-spin coupling effects occurring before and after an ideal pulse. One can then modify the phase of each pulse to correct for the phase shifts. Assuming that long range couplings between the spins vanish, the timing between pulses can be efficiently numerically optimized in order to absorb the coupling effects into the refocusing scheme [19]. This design does not take higher-order coupling and off-resonant effects into account and leaves some small couplings unrefocused to minimize the pulse sequence length. These approximations lead to errors in the control. A crucial point of this experimental work was to verify that these approximations hold for larger Hilbert spaces, i.e., that we could find a suitable refocusing scheme that, once optimized, still provides reliable control.

In liquid state NMR, the thermal equilibrium state is almost completely mixed. Therefore, instead of preparing ρ_{GHZ} , we actually prepare the following state:

$$\rho_{\text{cs}} \approx \frac{\mathbf{I}^{\otimes N}}{2^N} + \epsilon \mathbf{X}^{\otimes n} \mathbf{I}^{\otimes (N-n)}, \quad (1)$$

$N = 14$ is the total number of spins-1/2 in the register, and n is the number of qubits involved in the GHZ state. The factor $\epsilon \approx 10^{-5}$ is related to the thermal polarization of the system. The second term of ρ_{cs} , called the deviation density matrix, contains the n -coherence term $|00\dots 0\rangle \langle 11\dots 1| + |11\dots 1\rangle \langle 00\dots 0|$ corresponding to a n -qubit cat state $|00\dots 0\rangle + |11\dots 1\rangle$, as well as lower coherence terms corresponding to the other n -qubit GHZ states. This state preparation (called the encoding of the pseudocat state) is done by propagating the polarization of the two equivalent protons through the chain of nuclei by a sequence of 1- and 2-qubit quantum gates (see Fig. 1). In NMR, only single coherence terms are observable [20]. Therefore, to see the signature of the GHZ state, we need to transform the n -coherence term into a n -qubit labeled

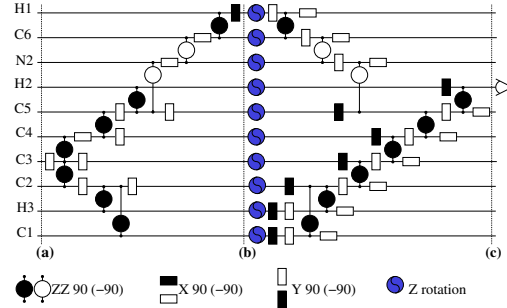


FIG. 1 (color online). Sequence of gates for the 10-qubit pseudocat-state preparation followed by its decoding into a 10-qubit pseudopure state. The initial preparation of the qutrit into a pseudopure state, as well as the refocusing gates, are not shown. Proper cycling of the Z rotations and the phase of observation filter act as a coherence filter. The qubit names correspond to the histidine molecule nuclei (see Fig. 2). (a) After the qutrit pseudopure preparation, the state of the register is $\mathbf{0}^{\text{H}_{4/5}} \mathbf{1}^{\text{H}_1} \mathbf{1}^{\text{C}_6} \mathbf{1}^{\text{N}_2} \mathbf{1}^{\text{H}_2} \mathbf{1}^{\text{C}_5} \mathbf{1}^{\text{C}_4} \mathbf{1}^{\text{C}_3} \mathbf{1}^{\text{C}_2} \mathbf{1}^{\text{H}_3} \mathbf{1}^{\text{C}_1}$. At the end of the encoding in (b), it is $\mathbf{0XXXXXXX}$, and, after filtering, the decoded state in (c) is $\mathbf{0000X00000}$.

pseudopure state of the form $\mathbf{X00}\dots\mathbf{0}$ [where $\mathbf{0} = (\mathbf{I} + \mathbf{Z})/2$]. This step of the algorithm is called the decoding. To average away the other lower coherence order terms present in the $\mathbf{X}^{\otimes n}$ operator, we used two types of coherence filters: gradient and phase cycling techniques. Proof that we have actually created the pseudopure and accompanying pseudocat state by determining the final state through tomography would require an impractically large number of experiments ($\sim 4^{12}$). Nevertheless, since the averaging procedure filters out the signal coming from every term but the desired one (i.e., the highest coherence order term), a single observation of the “read out” nucleus in the resulting NMR spectrum (see Fig. 2) indicates that we have indeed reached the desired coherence.

We applied both methods to design two series of pulse sequences that implement the encoding-decoding procedure, with n going from 1 to 12, on a liquid state NMR QIP, based on uniformly ^{13}C ; ^{15}N labeled *l*-histidine (see Fig. 2). Two different samples were used. The one used for designing strongly modulating pulses was made of 16.7 mg of histidine and 15.9 mg of deuterated phosphoric acid in 1 ml of deuterated water. To design simplified pulse sequences, we prepared a second sample by dissolving 35.3 mg of histidine, 12.5 mg of glycine-2- ^{13}C ; ^{15}N , and 3.4 mg of deuterated phosphoric acid in 1 ml of deuterated water. The labeled glycine molecule has a simple spectrum which allowed us to perform accurate calibrations of the selective pulses on isolated NMR peaks *in situ*. The experiments based on the strongly modulating pulses and the simplified design were, respectively, performed on Avance-600 and Avance-700 Bruker spectrometers at Massachusetts Institute of Technology and Institute for Quantum Computing.




FIG. 2 (color). This *l*-histidine molecule has 14 spin-1/2 nuclei: five ^1H , six ^{13}C , and three ^{15}N . See Ref. [24] for a more complete description of the molecule. The two protons H_4 and H_5 , are chemically equivalent and indistinguishable. As such, they can be seen as a composite particle with a spin-1 and a spin-0 part. We considered only the spin-1 subspace (qutrit) since the spin-0 does not interact with the other spin-1/2. This molecule is therefore a 12-qubits plus one qutrit quantum register. However, the N_3 nuclear spin has a particularly weak coupling with the rest of the molecule; thus, we did not use it. On this plot, we have shown a reference spectrum of H_2 (gray plot) and of the pseudopure state obtained after decoding a 10-qubit cat state onto H_2 (red line). They are arbitrarily scaled for clarity. The reference spectrum was obtained with 2 scans and the pseudopure, with 4000 scans, in order to improve the signal to noise ratio. We also show simulated spectra of the expected reference (yellow plot) and pseudopure state (blue line) for which the amplitude is matched to the experimental data to evaluate the signal loss.

We based the design and the interpretation of the experiments on a model of the system and the apparatus [18] which includes the following attributes: (i) The Hamiltonian of the system in the static magnetic field of the

spectrometer. The chemical shifts as well as the scalar-coupling strengths and relatives signs were experimentally determined by fits of reference spectra and small targeted multiple-coherence experiments. (ii) Knowledge of T_2 and T_2^* [20] relaxation times of the system. (iii) rf field inhomogeneities were mapped and used in the design of the strongly modulating pulses [21].

This type of experiment comes with a predicted exponential decay of signal as we increase the number of correlated qubits. We also expect high decoherence rates [22,23] and therefore a strong signal attenuation, as it is reasonable to assume that the relaxation rate for each spin included in the multiple quantum coherence add. To evaluate the quality of the control we could reach, the relevant quantity to measure is the amount of signal obtained experimentally with respect to the expected one assuming perfect control. Figure 3 shows how much signal we were able to retain after decoding the highest coherence order cat state into a pseudopure state for each experiment. We could reach a 12-coherence state using strongly modulating pulses and a 10-coherence state with selective pulses. Indeed, the sequences obtained through the simplified design were slightly longer, leading to more decoherence. Moreover, the transverse relaxation times were not the same in both sets of experiments. To distinguish between operational errors and relaxation loss, both decay times (T_2^* and T_2) were used to estimate the signal loss due to transversal relaxation during the pulse sequences (see Fig. 4). It showed that decoherence is the main source of signal loss and, therefore, indicates that we have good operational control.

Strongly modulating pulses			Simple pulses				
Coh. #	Retained signal (%)		Coherence path			Retained signal (%)	Coh. #
2	100	$\text{H}_{4/5}^x \text{C}_3^x$	C_3^x	C_1^x	$\text{H}_{4/5}^0 \text{C}_3^x$	100	1
3	46	$\text{H}_{4/5}^x \text{C}_3^x \text{C}_2^x$	C_2^x	H_3^x	$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x$	91	2
5	27	$\text{H}_{4/5}^x \text{C}_3^x \text{C}_2^x \text{C}_1^x \text{C}_4^x$	C_1^x	C_4^x	$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x \text{H}_3^x \text{C}_1^x \text{C}_4^x$	82	3
6	16	$\text{H}_{4/5}^x \text{C}_3^x \text{C}_2^x \text{C}_1^x \text{C}_4^x \text{C}_5^x$	C_5^x		$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x \text{H}_3^x \text{C}_1^x \text{C}_4^x \text{C}_5^x$	74	4
8	20	$\text{H}_{4/5}^x \text{C}_3^x \text{C}_2^x \text{C}_1^x \text{C}_4^x \text{C}_5^x \text{H}_2^x \text{H}_3^x$	H_2^x	H_3^x	$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x \text{H}_3^x \text{C}_1^x \text{C}_4^x \text{C}_5^x$	39	5
8	6	$\text{H}_{4/5}^x \text{C}_3^x \text{C}_2^x \text{C}_1^x \text{C}_4^x \text{C}_5^x \text{N}_1^x \text{N}_2^x$	N_1^x	N_2^x	$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x \text{H}_3^x \text{C}_1^x \text{C}_4^x \text{C}_5^x$	21	6
10	9	$\text{H}_{4/5}^x \text{C}_3^x \text{C}_2^x \text{C}_1^x \text{C}_4^x \text{C}_5^x \text{N}_1^x \text{N}_2^x \text{H}_2^x \text{H}_3^x$	N_1^x	N_2^x	$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x \text{H}_3^x \text{C}_1^x \text{C}_4^x \text{C}_5^x \text{H}_2^x$	20	7
12	3	$\text{H}_{4/5}^x \text{C}_3^x \text{C}_2^x \text{C}_1^x \text{C}_4^x \text{C}_5^x \text{N}_1^x \text{N}_2^x \text{H}_2^x \text{H}_3^x \text{C}_6^x \text{H}_1^x$	C_6^x	H_1^x	$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x \text{H}_3^x \text{C}_1^x \text{C}_4^x \text{C}_5^x \text{H}_2^x \text{N}_2^x$	0.5	8
					$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x \text{H}_3^x \text{C}_1^x \text{C}_4^x \text{C}_5^x \text{H}_2^x \text{N}_2^x \text{C}_6^x$	0.8	9
					$\text{H}_{4/5}^0 \text{C}_3^x \text{C}_2^x \text{H}_3^x \text{C}_1^x \text{C}_4^x \text{C}_5^x \text{H}_2^x \text{N}_2^x \text{C}_6^x \text{H}_1^x$	1.5	10

FIG. 3 (color online). Description of each of the multiple-coherence experiments. The picture shows how the polarization of the $\text{H}_{4/5}$ is propagated through the nuclei chain to create the cat state. In the series of experiments using simple pulses, we first prepared the qutrit made of the two equivalent protons into a pseudopure state $\mathbf{0} = (\mathbf{I} + \mathbf{Z})/2$ and left it as such for the rest of the experiments. When using strongly modulating pulses, we included the qutrit into the multiple-coherence state. The two first and two last columns show the coherence number we reached and how much signal we were able to retain after decoding the cat state into a observable pseudopure state. Results are shown in percentage of the expected signal assuming perfect control, normalized to the first experiment.




FIG. 4 (color online). Expected decay of the pseudopure state signal due to transversal relaxation for the series of experiments done with simple pulses. Each point corresponds to a different coherence experiment. The length of the pulse sequence increases with the coherence order. Most experimental points are above the estimates given by T_2^* and even T_2 . Indeed, to predict the transversal relaxation during the pulse sequence, we used only a very simple model of decay for multiple coherences that gives an upper bound of the signal loss. Nevertheless, the experimental curve and predicted ones show the same decay pattern. Thus, it is reasonable to say that most of the signal loss comes from decoherence and, therefore, that we have a good operational control over the system. For the coherences 8, 9, and 10 (last three points), the experimental curve goes below the T_2 curve. It reflects a loss of accuracy in our operational control. Indeed, for these experiments, we are controlling the nitrogen nuclei through very small couplings. We are, therefore, using long 2-qubits gates that are sensitive to any small inaccuracy in the values of the Hamiltonian parameters.

In summary, we have reported an algorithmic benchmark performed on the largest quantum information processor to date. Despite the decoherence during the computation, we have been able to demonstrate universal control on up to 11 qubits and one qutrit. This work shows that liquid state NMR allows us to develop operational control methods that can be used to control a large number of quantum degrees of freedom. These methods provide a systematic and efficient way of programming liquid state NMR QIPs. However, the approaches and control techniques behind these methods could also be used to design control sequences in more scalable implementations where the intrinsic errors are smaller.

This work was supported by ARDA, ARO, LPS, NSERC, and by the Cambridge-MIT Institute.

[1] E. Knill, R. Laflamme, R. Martinez, and C.-H. Tseng, *Nature (London)* **404**, 368 (2000).

- [2] A. Furusawa, J. Sorensen, S.L. Braunstein, C. Fuchs, H.J. Kimble, and E.S. Polzik, *Science* **282**, 706 (1998).
- [3] M. Riebe *et al.*, *Nature (London)* **429**, 734 (2004).
- [4] M.D. Barret *et al.*, *Nature (London)* **429**, 737 (2004).
- [5] D.P. DiVincenzo and P.W. Shor, *Phys. Rev. Lett.* **77**, 3260 (1996).
- [6] N. Gershenfeld and I.L. Chuang, *Science* **275**, 350 (1997).
- [7] D.G. Cory, A.F. Fahmy, and T.F. Havel, *Proc. Natl. Acad. Sci. U.S.A.* **94**, 1634 (1997).
- [8] L.M.K. Vandersypen, M. Steffen, G. Breyta, C.S. Yannoni, M.H. Sherwood, and I.L. Chuang, *Nature (London)* **414**, 883 (2001).
- [9] E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne, *Phys. Rev. Lett.* **86**, 5811 (2001).
- [10] D. Gottesman, Ph.D. thesis, California Institute of Technology, Pasadena, CA, 1997, quant-ph/9705052.
- [11] J.-S. Lee and A.K. Khitrin, *J. Chem. Phys.* **122**, 041101 (2005).
- [12] W.S. Warren, D.P. Weitekamp, and A. Pines, *J. Chem. Phys.* **73**, 2084 (1980).
- [13] A. Barenco, C.H. Bennett, R. Cleve, D.P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter, *Phys. Rev. A* **52**, 3457 (1995).
- [14] N. Khaneja, R. Brockett, and S.J. Glaser, *Phys. Rev. A* **63**, 032308 (2001).
- [15] J.A. Jones, *Phys. Rev. A* **67**, 012317 (2003).
- [16] D. Cory, M.D. Price, W. Maas, E. Knill, R. Laflamme, W. Zurek, T.F. Havel, and S.S. Somaroo, *Phys. Rev. Lett.* **81**, 2152 (1998).
- [17] E.M. Fortunato, M.A. Pravia, N. Boulant, G. Teklemariam, T.F. Havel, and D.G. Cory, *J. Chem. Phys.* **116**, 7599 (2002).
- [18] Y.S. Weinstein, T.F. Havel, J. Emerson, N. Boulant, M. Saraceno, S. Lloyd, and D.G. Cory, *J. Chem. Phys.* **121**, 6117 (2004).
- [19] M.D. Bowdrey, J. Jones, R. Laflamme, and M. Knill, *Phys. Rev. A* **72**, 032315 (2005).
- [20] R.R. Ernst, G. Bodenhausen, and A. Wokaun, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions* (Oxford University Press, Oxford, 1994).
- [21] M. Pravia, N. Boulant, J. Emerson, A. Farid, E. Fortunato, T. Havel, R. Martinez, and D. Cory, *J. Chem. Phys.* **119**, 9993 (2003).
- [22] H. Krojanski and D. Suter, *Phys. Rev. Lett.* **93**, 090501 (2004).
- [23] J.-S. Lee and A.K. Khitrin, quant-ph/0508052.
- [24] See EPAPS Document No. E-PRLTAO-96-009617 for a figure that describes in more detail the histidine molecule used in the experiments. The table includes chemical shifts, J -coupling values, and relaxation times (T_2 and T_2^*). For more information on EPAPS, see <http://www.aip.org/pubservs/epaps.html>.

D.2 Low temperature probe for dynamic nuclear polarization and multiple-pulse solid-state NMR.

A major goal of the Waterloo group was to achieve high polarization in bulk ensemble NMR for implementing multiple rounds of error correction. Our initial work in this direction focussed on dynamic nuclear polarization (DNP) again in collaboration with the MIT group. Because the electron spin has a much larger magnetic moment (≈ 660 times greater than proton), for any given field and temperature the electrons will have a much greater polarization than the nuclei. Typical values of the polarization at cryogenic temperatures are shown in Figure 4.6. Using microwaves, the electron polarization can be transferred to the nuclei by driving forbidden flip-flop transitions of the hyperfine interaction. The goal was to use the electrons to enhance the nuclear polarization via DNP and then to turn off the microwaves and perform quantum information processing tasks on the polarized nuclei using the same control techniques as in the room temperature experiments. The probe was designed and built in collaboration with the MIT group. At Waterloo we added a double resonance circuit to control both proton and carbon in a field of 2.1T.

Performing high power NMR experiments at cryogenic temperatures proved to be a challenge with the probe resonance conditions changing, tuning and matching capacitors freezing up and arcing due to the helium gas environment. Nevertheless, we were able to show DNP enhancement of the proton polarization in a frozen glycerol/water solution with the TEMPO free radical. We were also able to take stroboscopic observation of the signal under a dipolar decoupling sequence showing line narrowing and control. Furthermore, we were able to transfer the DNP enhanced proton polarization to the natural abundance carbon nuclei via cross-polarization. Although this experiment achieved the initial goals, it highlighted the difficulty of getting high polarization with bulk-ensemble systems and incoherent polarization transfer. This provided the impetus to move systems where we can coherently control the electron spin with pulsed ESR.



Low temperature probe for dynamic nuclear polarization and multiple-pulse solid-state NMR

HyungJoon Cho ^{a,1}, Jonathan Baugh ^b, Colm A. Ryan ^b, David G. Cory ^a, Chandrasekhar Ramanathan ^{a,*}

^a Department of Nuclear Science and Engineering, Massachusetts Institute of Technology, Cambridge MA 02139, USA

^b Institute for Quantum Computing, University of Waterloo, Waterloo, Ont., Canada, N2L 3G1

Received 8 January 2007; revised 4 April 2007

Available online 30 April 2007

Abstract

Here, we describe the design and performance characteristics of a low temperature probe for dynamic nuclear polarization (DNP) experiments, which is compatible with demanding multiple-pulse experiments. The competing goals of a high-Q microwave cavity to achieve large DNP enhancements and a high efficiency NMR circuit for multiple-pulse control lead to inevitable engineering tradeoffs. We have designed two probes—one with a single-resonance RF circuit and a horn-mirror cavity configuration for the microwaves and a second with a double-resonance RF circuit and a double-horn cavity configuration. The advantage of the design is that the sample is in vacuum, the RF circuits are locally tuned, and the microwave resonator has a large internal volume that is compatible with the use of RF and gradient coils.

© 2007 Elsevier Inc. All rights reserved.

Keywords: Dynamic nuclear polarization; Multiple-pulse NMR; Coherent averaging; Low temperature NMR

1. Introduction

NMR approaches to quantum information processing (QIP) have received much attention over the last decade. Liquid state NMR studies have explored the limits to controlling small quantum systems, and enabled the systematic study of open quantum systems. Solid-state NMR QIP approaches allow us to obtain control over a larger Hilbert space [1–6], and hold promise for the study of many body dynamics [7–11] and quantum simulations. Moreover, in the solid-state, the spins can be highly polarized by dynamic nuclear polarization techniques. The increased polarization allows an exploration of systems with a larger number of qubits, and also allows preparation of the system close to a pure state.

A number of solid-state NMR QIP experiments have recently been published [12–15], demonstrating that the coherent control necessary for QIP can be implemented in single crystal solids at room temperature. The next step is to take advantage of the high polarizations that are accessible in the solid-state. In order to reach high polarization via dynamic nuclear polarization (DNP) it is necessary to irradiate the system with microwaves, and to cool the sample down to liquid helium temperatures. The challenge here is to engineer the NMR probe so that these can be achieved without sacrificing high-fidelity control of the spins. In this paper we describe the design and performance characteristics of a low temperature probe for dynamic nuclear polarization experiments, which is also compatible with demanding multiple-pulse experiments.

Historically, DNP investigations have concentrated on two main areas: increased detection sensitivity of rare spins; and the creation of spin-polarized targets. The groups of Wind and Yannoni have explored a number of microwave resonator designs including horn-reflector [16],

* Corresponding author. Fax: +1 617 253 0760.

E-mail address: sekhar@mit.edu (C. Ramanathan).

¹ Present address: Schlumberger Doll Research, Cambridge, MA, USA.

Fabry–Perot [17], and cylindrical resonators [18] for DNP signal enhancement at room temperature. At lower temperatures Griffin and co-workers have implemented a novel design in which the cylindrical walls of a high-Q microwave cavity formed the RF coil for the NMR [19]. At even lower temperatures, the highest absolute polarizations reported were obtained using CW techniques in a multimode cavity in a dilution refrigerator [20].

The probe discussed here is designed for a top-loading continuous flow cryostat (spectrostatCF, Oxford Instruments) that fits in the bore of a 2.35 T superconducting magnet (89 mm bore diameter). The design specifications comply with the dimensional and functional aspects of this particular cryostat, but could easily be adapted to many

other cryostats used in low temperature NMR applications. This field corresponds to an electron Larmor frequency of 66 GHz for $g=2$. It is possible to obtain relatively inexpensive solid-state microwave sources with up to a Watt of power at this frequency. Fig. 1 shows a three-dimensional solid model of the main section of the probehead, as well as a cut-away view identifying the different components.

2. RF design

There is a wealth of experience in low temperature NMR [21–23]. There are two main challenges in the low temperature design of the tuned RF probe:

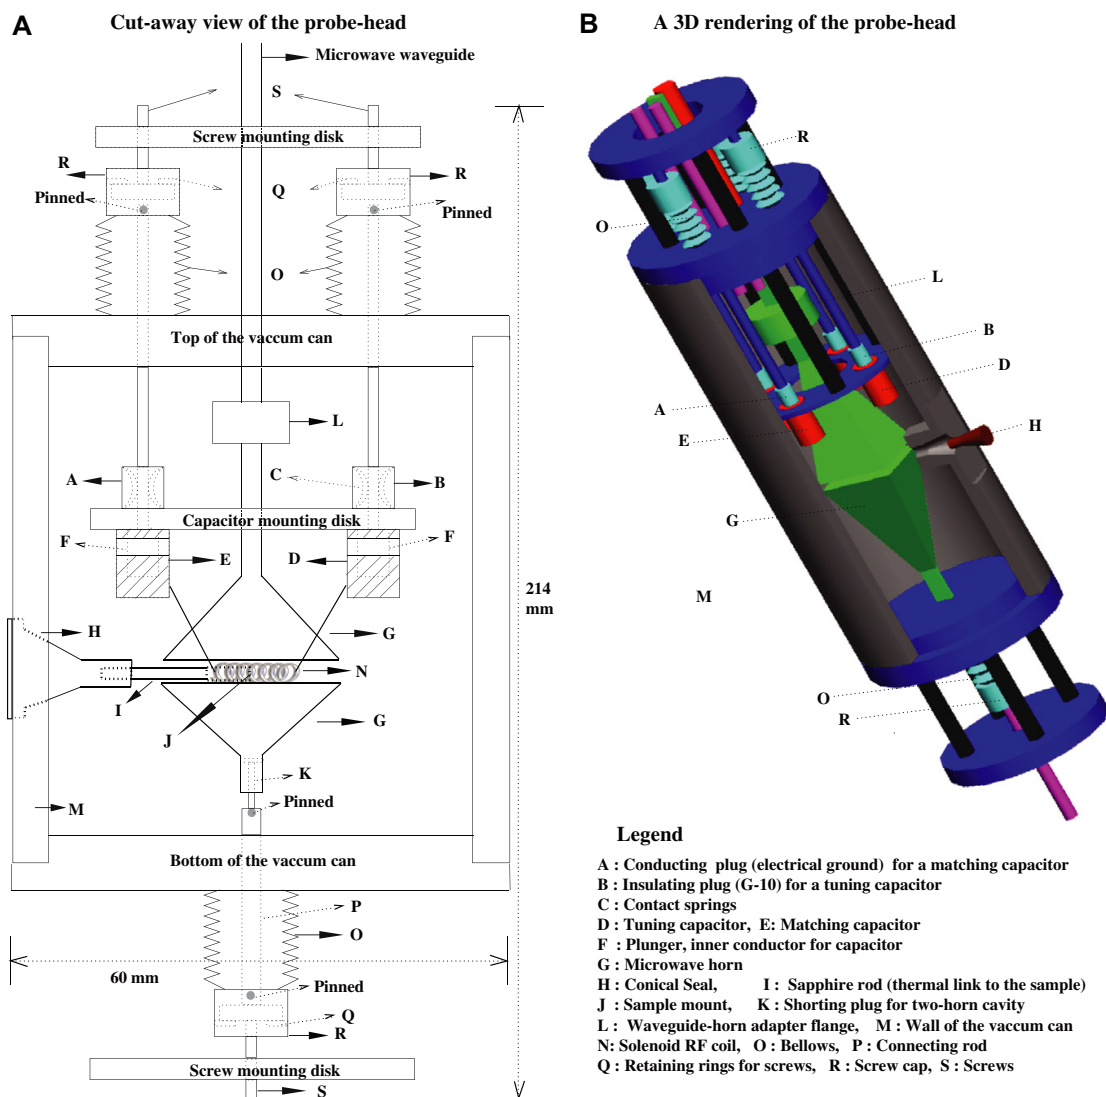


Fig. 1. Probe assembly: (A) cut-away view; and (B) three-dimensional CAD drawing.

- Maintaining the ability to tune and match the resonant circuit at low temperature.
- Ensuring that the sample remains at the desired operating temperature during the experiment.

As the probe is cooled down, the resonance invariably shifts due to contraction of the probe materials, and the Q increases as the coil resistance is reduced. It is therefore necessary to re-tune and match the circuit once the operating temperature has been established. It is difficult to operate variable capacitors at low temperature as they have a tendency to seize if they have not been thoroughly cleaned. Although a simple approach to circuit tunability is to locate the tuning and matching elements outside the cryostat at room temperature, this entails a substantial loss of efficiency since a lossy transmission line becomes part of the resonant circuit [24–26]. An alternative approach that requires no tuning parts uses a large diameter coaxial line to reduce the losses [27]. However, space constraints and the need to reduce thermal losses up the coax limit the applicability of this solution, especially if multiple RF channels are needed. Our approach is to keep the tuning and matching capacitors at low temperature in close proximity to the sample coil.

The next challenge is to cool the sample itself down to the desired temperature. If the sample is in direct contact with the helium gas or liquid, the only requirement is that the sample temperature equilibrate to the bath. However, a tuned circuit in a gaseous ^4He environment at low temperature is known to be susceptible to electric breakdown even at modest power levels. Although this may not be troublesome for a one-pulse experiment, high duty-cycle excitations such as spin-locking or multiple-pulse irradiation are likely to cause breakdown. The alternative is to locate the tuning circuit and RF coil, and usually the sample in a vacuum can evacuated to $\sim 10^{-5}$ torr. The sample is then mounted on a sapphire rod that is heat sunk at the other end to a copper plate. In this configuration we have found it difficult to cool the sample below about 6 K. The challenges are twofold. First we need to find a material that

has high thermal conductivity as well as good dielectric properties at these temperatures. We used sapphire, which is a very good thermal conductor at low temperatures, though the conductivity peaks around 30 K and drops sharply below that [28]. Secondly we need to be able to place the sample in good thermal contact with this material. In our experiments, the quartz tube containing the frozen TEMPO solution was attached to the sapphire rod with Dow Corning vacuum grease.

An alternative solution is to seal the tuning circuit and the coil in a vacuum space, while the sample is immersed in helium. While this design permits both high RF powers and low sample temperatures, the tradeoff is a significantly smaller filling factor, resulting in lower B_1 fields compromising the resultant control [22,29].

2.1. RF transmission

In order to minimize conductive heat losses down the coaxial cables, we used UT-141C-SS semi-rigid coaxial cable (Micro-Coax), which has a silver-plated copper inner conductor and a stainless steel outer conductor with a PTFE dielectric. A series of brass baffles were used to minimize radiative losses to the room temperature flange. We have built two versions of the probe, a single-resonance probe tuned to protons, and a double-resonance probe tuned to protons and carbon. The standard circuit configurations used are shown in Fig. 2. We used standard tune-up sequences to calibrate the probes for multiple-pulse experiments. In order to minimize phase transients, we reduced the Q of the circuits used. In the single-resonance probe we overcoupled the circuit, and in the double-resonance probe we added a 1.2 Ω metal oxide resistor, yielding a Q of around 25.

A 3 mm diameter solenoidal RF coil was used in both probes, a 7-turn coil in the single-resonance circuit and a 9-turn coil in the double-resonance circuit. The wire was slightly flattened and then wound with a non-uniform pitch so as to optimize RF inhomogeneity [30]. The free-standing RF coil was supported by the leads as shown in Fig. 1.

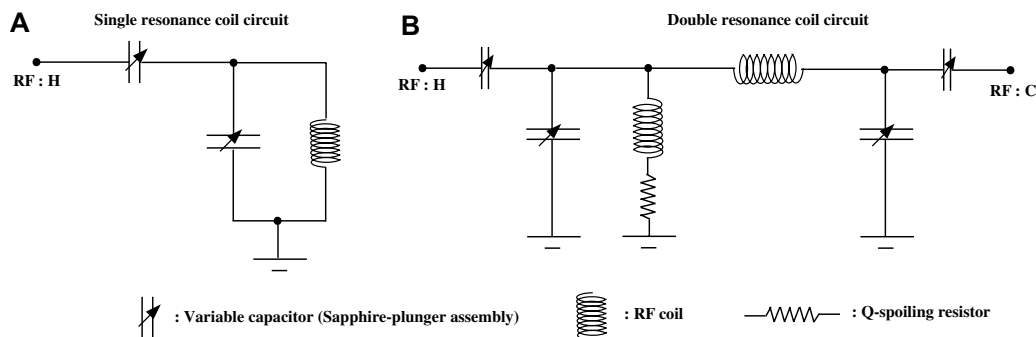


Fig. 2. Circuit schematics for the single- and double-resonance probes.

2.2. Tuning and matching

Fig. 3 shows a schematic diagram of the variable capacitor assembly. In order to achieve variable tuning and matching at low temperatures, we designed coaxial cylindrical capacitors using copper-plated sapphire tubes (7.8 mm OD, 4.8 mm ID, 14 mm in length) with a piston mechanism to obtain a capacitance range of 1–12 pF. The wall thickness of the sapphire dielectric is 1.5 mm, giving us a breakdown voltage of at least 25 kV. The capacitance range is obtained using $C = \epsilon\pi R_i^2 l/d \approx 10l$ pF (using $R_i = 4.8$ mm (ID of the sapphire tube), $d = 1.5$ mm (wall thickness of the tube), and where l is the depth of the plunger in cm). The thickness of the copper plating is 50 μm , which is more than five times the skin depth of ~ 7 μm for copper at 100 MHz. We used a bellows mechanism (beryllium copper bellows, Mini-flex Corporation) to couple the room temperature rotary motion of the tuning rods to linear motion of the piston inside the low temperature vacuum can. A retaining ring (Rotor Clip, Inc.) was mounted inside the adapter on top of the bellows so that the bellows can be either extended or contracted by rotating a screw. Tuning rods made of G-10 extended to the room temperature flange at the top of the cryostat, where they passed through a set of Goddard quick-connect valves (Rego Products).

2.3. Sample cooling and orientation

To ease sample insertion and positioning, we machined a tapered conical seal (copper or brass) that doubles as a sample mount. Dow Corning silicone grease applied to the seal was found to be effective from room temperature to below 4 K. The advantage of using this approach is that one can easily change samples without opening a flange, and the (single crystal) sample orientation can be systematically varied by rotating the conical seal with respect to the vacuum can wall. A dial was inscribed onto the wall of the vacuum can to facilitate reproducible sample orientation. The conical seal also acts as a pressure release valve for the vacuum can, in the event of helium getting trapped inside the can. Fig. 4 shows a diagram of this conical seal. The sapphire rod attached to the conical seal extends to the center of the RF coil and the sample is mounted on the end of this sapphire rod. Clearly, efficient cooling of the sample by thermal contact is a critical issue for our vacuum can approach. Special care must be taken to ensure good thermal contact between the sample and the sapphire, and between the sapphire and the conical seal into which it is inserted. The external side of the conical seal makes contact with the gaseous ^4He environment providing a cooling pathway for the sample.

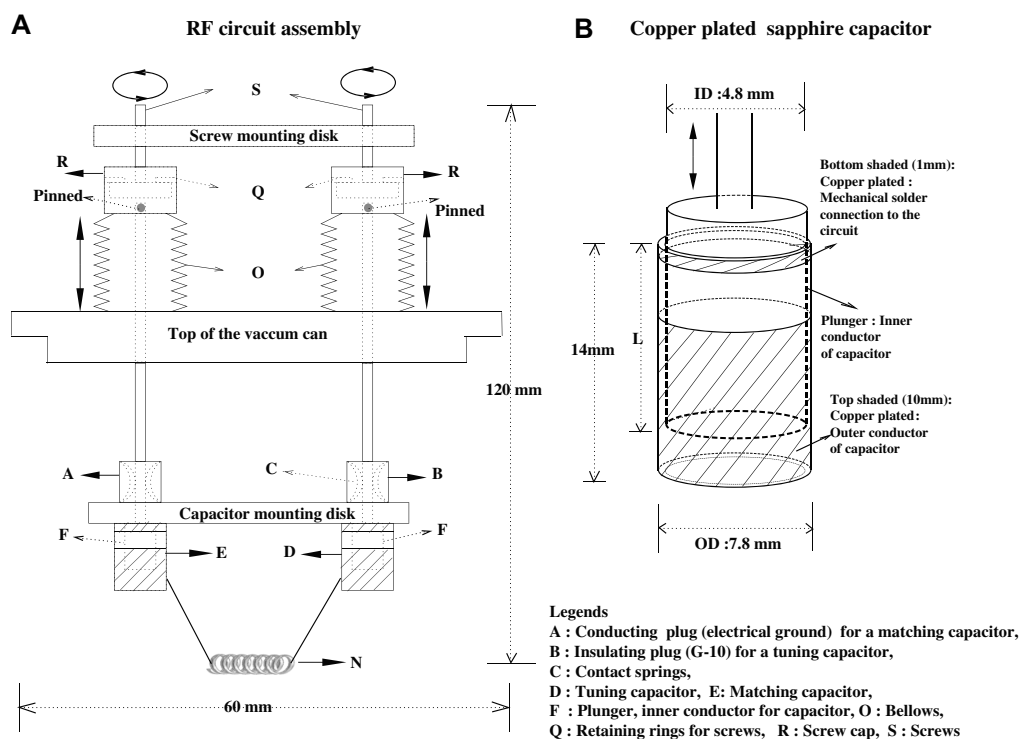


Fig. 3. (A) Schematic diagram of the variable capacitor assembly showing the copper-plated sapphire capacitors and the bellows arrangement that permits low temperature tuning inside the vacuum can. Rotary motion outside the can is converted to linear motion inside the can. (B) Detail of a single copper-plated sapphire capacitor.




Fig. 4. (A) Side view of the wall of the vacuum can, showing the location where the conical seal sits, as well as the angular markings used to orient the sample. (B) Front and side views of the conical seal.

3. Microwave design

Microwave irradiation of the sample was achieved using microwave horns. A horn-mirror configuration was used in the single-resonance probe, and the double-horn configuration [31,32] was used with the double-resonance probe (Fig. 5). Although horn arrangements do not have particularly good microwave quality factors ($Q \sim 10$ – 100), they provide ample space to accommodate extra components such as RF and gradient coils. In addition they do not need any additional impedance matching arrangement. In the double-horn arrangement, the distance between the two horns was set equal to one wavelength of the microwave source so that the $3\lambda/4$ point (shown in Fig. 5) would have a minimum of the E -field on the horn surface. This allowed us to drill holes in the walls of the horn at this location, and pass the leads of the RF coil through the holes, without distorting the microwave mode structure. The shorter leads allow for greater RF efficiency.

Fig. 6 shows reflected power versus the relative distance of the mirror from the horn at 66 GHz microwave frequency, and similar data for the double-horn geometry (gap between the two horns is set to 4.5 mm, one wavelength at 66 GHz) as a function of a sliding shorting plug that is placed within the neck of the end horn. In the horn-mirror geometry, a loss of microwave power is observed as the gap between the horn and the mirror increases. The double-horn setup has a significantly higher quality factor ($Q \sim 100$) and shows reduced loss compared to the horn-mirror setup.

To effectively map out the H field profile in the horn-mirror system, we performed HFSS (High Frequency

Structure Simulator, AnSoft) simulations with the different horn configurations. Fig. 7 shows the simulated H field profiles with the horn-mirror and double-horn geometries and a 66 GHz microwave source. We found that the horn-mirror and double-horn configurations have 4.5 dB and 7 dB gains in H field strength, respectively, compared to a single horn.

3.1. Experimental setup

Fig. 8 shows a schematic of the experimental setup and the estimated insertion losses of the different components. We used a standard fundamental mode (TE_{10}) rectangular WR-15 waveguide to transmit the microwaves from the source to the cavity. The source used was a backshort-tuned Gunn oscillator (Millitech, LLC) with an attached isolator. Its center frequency is 66.2 ± 0.1 GHz with mechanical tuning ability of ± 2.0 GHz. The available output power of this Gunn oscillator is 40–65 mW. We also performed experiments with two other sources, a 60 mW Gunn diode source from Millitech, LLC with a mechanical tuning range of ± 1.5 GHz, and a 1 W source from Quinstar Technologies. The 1 W source was obtained by power combining the output of two impatt diode sources that are injection-locked with a Gunn to reduce phase noise. To connect the probe assembly and microwave source while maintaining vacuum inside the probe assembly, we used a bulkhead flange unit (Aerowave, Inc.) as a vacuum window on the top of low temperature NMR probe with mica and rubber O-ring seals. A mica window serves as a vacuum feed-through for microwave and has low insertion losses. To reduce microwave losses while maintaining thermal iso-




Fig. 5. (A) The horn-mirror arrangement showing the use of a bellows arrangement to move the mirror and tune the cavity from outside the vacuum can. (B) The double-horn arrangement showing the use of a shorting plug inserted into the second horn, that permits tuning. A pair of holes are placed in the upper horn, at a location where the *E*-field is minimum in order to minimize field distortions. Using these holes for the leads of the RF coil, minimizes the length and inductance of the leads, increasing coil efficiency.




Fig. 6. Reflected power measurements for the two cavity geometries. In the lower figure, the *x*-axis shows the location of the shorting plug inside the neck of the horn.

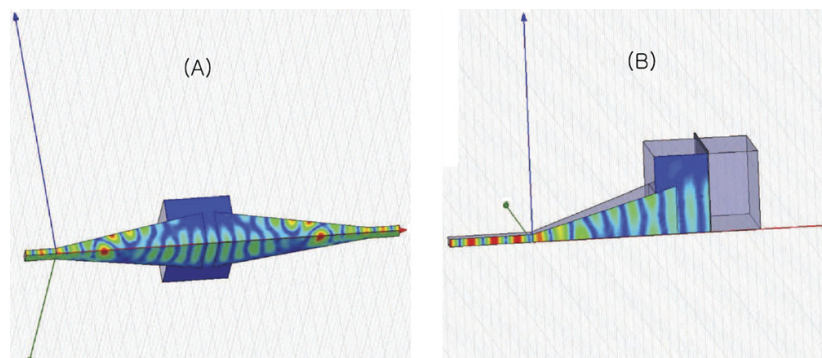


Fig. 7. HFSS simulations showing the H field profiles in (A) the double-horn and (B) the horn-mirror cavity configurations.

lation between the sample area and the waveguide, a 4" section of stainless steel waveguide was brazed into a 40" long coin-silver waveguide. The stainless steel section acts as a heat switch limiting conduction up the waveguide. The total loss of microwave power was ~ 10.6 dB, resulting in ~ 5 mW output microwave power at the horn input. We do not always need to have the directional coupler in place, so we can save the 1.43 dB loss associated there.

In the single-resonance probe, the proton $\pi/2$ pulse was 1 μ s with using a 300 W Bruker BLAX amplifier with 0 dBm input power. For the double-resonance probe, the measured $\pi/2$ pulses were 2 μ s on both proton and carbon

channels, using 300 W Bruker amplifiers on both channels, with 0 dBm input power.

4. Experimental results

DNP experiments were performed at 4 K and 2.35 T (100 MHz ^1H , 25 MHz ^{13}C), using a BRUKER Avance spectrometer and the home-built probes. A sample of 40 mM 4-amino-TEMPO radical (4-amino-2,2,6,6-tetramethylpiperidine 1-oxyl) in a 60/40 glycerol/water solution was used to test DNP. At this field strength and concentration, DNP is mediated by cross-relaxation between electron

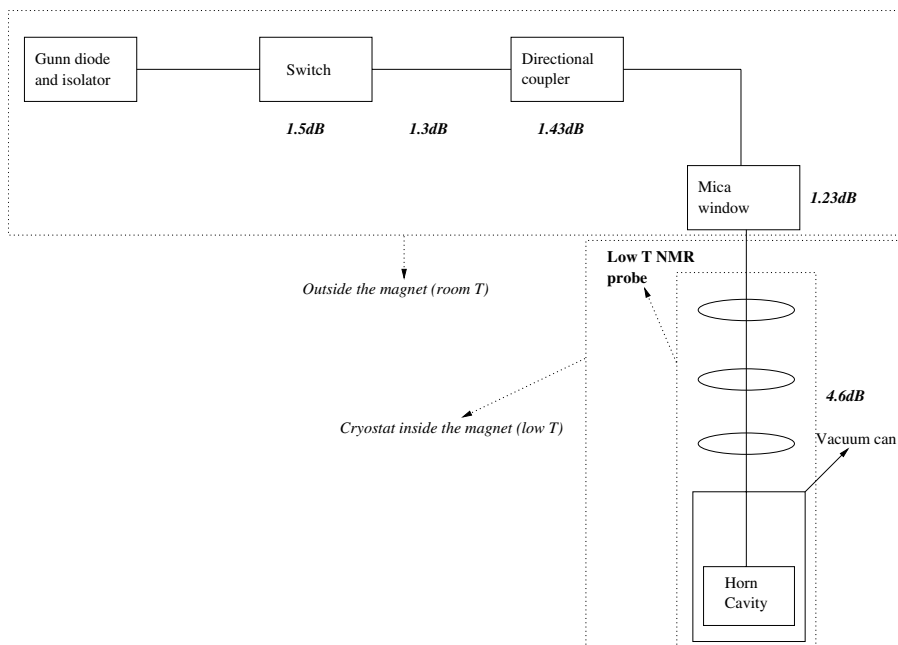


Fig. 8. A schematic drawing of the experimental setup. Estimated losses of microwave power due to waveguide sections and microwave elements are indicated in dB.




Fig. 9. (A) DNP enhancement of glycerol/water doped with 40 mM TEMPO as a function of microwave frequency, as obtained with a horn-mirror configuration. DNP is mediated by cross-relaxation between spin packets in the inhomogeneously broadened ESR line of TEMPO [33]. (B) DNP enhancement of 38 for the same system obtained with the double-horn configuration.

spins in the inhomogeneously broadened ESR line [33]. The sample volume was $\sim 2 \text{ mm}^3$. The 40 mM TEMPO solution is prepared by dissolving 0.05 g of 4-amino-TEMPO with a solution of 4.38 ml glycerol and 2.92 ml water. Fig. 9 shows the DNP enhancement obtained with the two probes. It is seen that experimentally the two cavities yield almost the same enhancement. This enhancement corresponds to a proton polarization of about 1.7%. With the double-resonance probe we were able to transfer this enhanced proton polarization to the carbon spins using Hartmann–Hahn cross-polarization. Fig. 10 shows the carbon signal enhancement following cross-polarization from the protons at 100 K and 6 K without DNP, and with DNP at 6 K.

Fig. 11 shows the proton spectra of the frozen glycerol/water solution obtained using a one-pulse experiment and




Fig. 10. ^{13}C spectrum of the frozen glycerol/water solution, showing the signal enhancement following proton DNP and cross-polarization from the protons to the carbon.

an MREV-8 sequence [34]. It is seen that the single 70 kHz line seen in the one-pulse experiment is narrowed significantly and is about a kilohertz wide in the MREV-8 spectrum. The residual linewidth is most likely dominated by the proton CSA. Protons in ice are known to have an axially symmetric CSA ($\sigma_{\parallel} = 15 \pm 2 \text{ ppm}$, $\sigma_{\perp} = -19 \pm 2 \text{ ppm}$) [35]. The zero frequency peak in the MREV-8 spectrum corresponds to a residual spin-locking signal. In running the multiple-pulse experiments, the probe was initially tuned up at room temperature using a liquid sample, and the characteristic tuning curve of the probe measured. Following cool down to liquid helium temperatures, the




Fig. 11. One-pulse proton spectrum of the frozen glycerol/water solution, and a spectrum obtained with an MREV-8 sequence applied at a cryostat temperature of 4.2 K (sample temperature about 6 K). The inset shows the MREV-8 data expanded. The peak centered at 1.5 kHz corresponds to the water/glycerol signal, while the zero frequency peak is a residual spin-locking signal.

probe was re-tuned to try and match the measured room temperature characteristics.

We have shown that these probes can be used successfully to dynamically polarize the nuclear spins, and to implement both cross-polarization and multiple-pulse techniques. Additional modifications are possible to improve microwave performance, in particular by using overmoded waveguide to minimize losses during transmission to the cavity. Thermal losses could be minimized by using gold-plated stainless-steel waveguide. The sample in vacuum design ultimately limits the base sample temperatures achievable. We were able to reach a temperature of approximately 6 K. It is possible to lower this a little by improving the design, but as the thermal conductivity of sapphire drops sharply around 4 K, this design is unlikely to be useful at lower temperatures. At lower temperatures, a sample in helium design ensures that the sample reaches base temperature. If the RF electronics is maintained in liquid helium, arcing is prevented, and we are once again able to run high duty-cycle multiple-pulse experiments.

Acknowledgments

We thank Peter Allen for his help in assembling the probes outlined here. This work was supported in part by the National Security Agency (NSA) under Army Research Office (ARO) contract numbers DAAD19-03-1-0125 and W911NF-05-1-0469, and by DARPA DSO and the Air Force Office of Scientific Research.

References

- [1] D.G. Cory, R. Laflamme, E. Knill, L. Viola, T.F. Havel, N. Boulant, G. Boutis, E. Fortunato, S. Lloyd, R. Martinez, C. Negrevergne, M. Pravia, Y. Sharf, G. Teklemariam, Y.S. Weinstein, W.H. Zurek, NMR based quantum information processing: achievements and prospects, *Fortschr. Phys.* 48 (2000) 875–907.
- [2] G.P. Berman, G.W. Brown, M.E. Hawley, V.I. Tsifrinovich, Solid-state quantum computer based on scanning tunneling microscopy, *Phys. Rev. Lett.* 87 (2001) 097902.
- [3] J. Wrachtrup, S.Y. Kilin, A.P. Nizovtsev, Quantum computation using the C-13 nuclear spins near the single nv defect center in diamond, *Opt. Spectrosc. + 91* (2001) 429–437.
- [4] D. Suter, K. Lim, Scalable architecture for spin-based quantum computers with a single type of gate, *Phys. Rev. A* 65 (2002) 052309.
- [5] T.D. Ladd, J.R. Goldman, F. Yamaguchi, Y. Yamamoto, All-silicon quantum computer, *Phys. Rev. Lett.* 89 (2002) 017901.
- [6] E. Abe, K.M. Itoh, T.D. Ladd, J.R. Goldman, F. Yamaguchi, Y. Yamamoto, Solid-state silicon NMR quantum computer, *J. Supercond.* 16 (2003) 175–178.
- [7] W. Zhang, D.G. Cory, First experimental measurement of the spin diffusion rate in a homogeneous solid, *Phys. Rev. Lett.* 80 (1998) 1324–1327.
- [8] G.S. Boutis, D. Greenbaum, H. Cho, D.G. Cory, C. Ramanathan, Spin diffusion of correlated two-spin states in a dielectric solid, *Phys. Rev. Lett.* 92 (2004) 137201.
- [9] C. Ramanathan, H. Cho, P. Cappellaro, G.S. Boutis, D.G. Cory, Encoding multiple quantum coherences in non-commuting bases, *Chem. Phys. Lett.* 369 (2003) 311–317.
- [10] H. Cho, T.D. Ladd, J. Baugh, D.G. Cory, C. Ramanathan, Multi-spin dynamics of the solid state NMR free induction decay, *Phys. Rev. B* 72 (2005) 054427.
- [11] H. Cho, P. Cappellaro, D.G. Cory, C. Ramanathan, Decay of highly correlated spin states in a dipolar-coupled solid: NMR study of CaF₂, *Phys. Rev. B* 74 (2006) 224434.
- [12] G.M. Leskowitz, N. Ghaderi, R.A. Olsen, L.J. Mueller, Three-qubit nuclear magnetic resonance quantum information processing with a single-crystal solid, *J. Chem. Phys.* 119 (2003) 1643–1649.
- [13] C. Ramanathan, S. Sinha, J. Baugh, T.F. Havel, D.G. Cory, Selective coherence transfers in homogeneous dipolar coupled spin systems, *Phys. Rev. A* 71 (2005) 020303.
- [14] J. Baugh, O. Moussa, C.A. Ryan, A. Nayak, R. Laflamme, Experimental implementation of heat-bath algorithmic cooling using solid-state nuclear magnetic resonance, *Nature* 438 (2005) 470–473.
- [15] J. Baugh, O. Moussa, C.A. Ryan, R. Laflamme, C. Ramanathan, T.F. Havel, D.G. Cory, Solid-state NMR three-qubit homonuclear system for quantum information processing: control and characterization, *Phys. Rev. A* 73 (2006) 022305.
- [16] R.A. Wind, F.E. Anthonio, M.J. Duijvestijn, J. Smidt, J. Trommel, G.M.C. de Vette, Experimental setup for enhanced ¹³C NMR spectroscopy in solids using dynamic nuclear polarization, *J. Magn. Reson.* 52 (1983) 424–434.
- [17] D.J. Singel, H. Seidel, R.D. Kendrick, C.S. Yannoni, A spectrometer for EPR, DNP, and multinuclear high-resolution NMR, *J. Magn. Reson.* 81 (1989) 145–161.
- [18] R.A. Wind, R.A. Hall, A. Jurkiewicz, H. Lock, G.E. Maciel, Two novel DNP-NMR probes, *J. Magn. Reson.* 110 (1994) 33–37.
- [19] V. Weis, M. Bennati, M. Rosay, J.A. Bryant, R.G. Griffin, High-field DNP and ENDOR with a novel multiple-frequency resonance structure, *J. Magn. Reson.* 140 (1999) 293–299.
- [20] W. de Boer, T.O. Niinikoski, Dynamic proton polarization in propanediol below 0.5 K, *Nucl. Instrum. Methods* 114 (1974) 495–498.
- [21] M.S. Conradi, Low-temperature NMR techniques, *Concepts Magn. Reson.* 5 (1993) 243–262.
- [22] P.L. Kuhns, S.H. Lee, C. Coretsopoulos, P.C. Hammel, O. Gonen, J.S. Waugh, A low temperature NMR probe for use in a dilution refrigerator, *Rev. Sci. Instr.* 62 (1991) 2159–2163.
- [23] R.E. Norberg, W.L. Earl, Y.W. Kim, Cryogenic probe with low-loss transmission line for nuclear magnetic resonance, *J. Magn. Reson. A* 116 (1995) 139–144.
- [24] J.H. Walton, M.S. Conradi, Probe-tuning adjustments—need they be in the probe? *J. Magn. Reson.* 81 (1989) 623–627.
- [25] Y.W. Kim, W.L. Earl, R.E. Norberg, Cryogenic probe with low-loss transmission line for nuclear magnetic resonance, *J. Magn. Reson. A* 116 (1995) 139–144.
- [26] V.D. Kodibagkar, M.S. Conradi, Remote tuning of NMR probe circuits, *J. Magn. Reson.* 144 (2000) 53–57.
- [27] R.A. McKay, Double-tuned single coil probe for nuclear magnetic resonance spectrometer. US Patent 4,446,431, 01 May 1984.
- [28] E.N. Smith, R.C. Richardson, *Experimental Techniques in Condensed Matter Physics at Low Temperatures*, Addison Wesley Publishing Company, Boston, 1998.
- [29] A.Z. Damyanovich, J. Peternelj, M.M. Pintar, Design of a high-power NMR probe for low-temperature studies, *J. Magn. Reson.* 144 (2000) 1–5.
- [30] R. Prigl, U. Haebleren, The theoretical and practical limits of resolution in multiple-pulse high-resolution NMR of solids, *Adv. Magn. Opt. Reson.* 19 (1996) 1–58.
- [31] E.S. Gravlín, J.A. Cowen, Simple microwave resonance spectrometer, *Am. J. Phys.* 27 (1959) 566.
- [32] H. Nishiguchi, S. Sugito, N. Hirota, ESR spectrometers using electromagnetic horns, *J. Magn. Reson.* 68 (1986) 40–51.
- [33] C.T. Farrar, D.A. Hall, G.J. Gerfen, S.J. Inati, R.G. Griffin, Mechanism of dynamic nuclear polarization in high magnetic field, *J. Chem. Phys.* 114 (2001) 4922–4933.
- [34] U. Haebleren, *High Resolution NMR in Solids*, Academic Press, New York, 1976.
- [35] A. Pines, D.J. Ruben, S. Vega, M. Mehring, New approach to high-resolution proton NMR in solids: deuterium spin decoupling by multiple-quantum transitions, *Phys. Rev. Lett.* 36 (1976) 110–113.

D.3 Experimental Implementation of a Quantum Walk

One of the first projects in my thesis was to demonstrate a discrete-time quantum walk on an NMR QIP. To achieve high-fidelity control we had to try several molecules and control techniques. This allowed us to compare implementing the same algorithm on two different systems. In a quantum walk the state starts on one of the nodes of the graph. A quantum coin is then flipped (rotated to a superposition state) and then conditioned on the state of the coin, the state is moved to another node of the graph. We implemented this using three qubits: two qubits denoted the state register and the other, the coin. The two qubit position register corresponds to a square and the state of the coin determines whether the state moves up/down or left/right on the square.

We implemented the algorithm on the three qubit molecule tri-chloro-ethylene which has one proton and two carbons. However, one of the carbons is coupled to quadrupolar chlorine nuclei which have a fast T_1 and induce dephasing on the carbon qubit. In addition, the carbon nuclei are strongly coupled and were difficult to control as we did not yet have optimal control theory tools. As a contrast we also used three carbon nuclei of the crotonic acid molecule where control is much easier and decoherence much less of a problem. The classical random walk will rapidly diffuse through the graph. The quantum version however will show interference and recurrences. We were able to demonstrate the periodic nature of the quantum walk in both molecules by applying quantum state tomography after each step and reconstructing the density matrices. In addition, we were able to show that when the coin is dephased, the quantum nature of the random walk is lost and it reverts to a classical random walk.

Experimental implementation of a discrete-time quantum random walk on an NMR quantum-information processor

C. A. Ryan,^{*} M. Laforest, J. C. Boileau, and R. Laflamme*Institute for Quantum Computing, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1*

(Received 27 July 2005; published 9 December 2005)

We present an experimental implementation of the coined discrete-time quantum walk on a square using a three-qubit liquid-state nuclear-magnetic-resonance (NMR) quantum-information processor (QIP). Contrary to its classical counterpart, we observe complete interference after certain steps and a periodicity in the evolution. Complete state tomography has been performed for each of the eight steps, making a full period. The results have extremely high fidelity with the expected states and show clearly the effects of quantum interference in the walk. We also show and discuss the importance of choosing a molecule with a natural Hamiltonian well suited to a NMR QIP by implementing the same algorithm on a second molecule. Finally, we show experimentally that decoherence after each step makes the statistics of the quantum walk tend to that of the classical random walk.

DOI: [10.1103/PhysRevA.72.062317](https://doi.org/10.1103/PhysRevA.72.062317)

PACS number(s): 03.67.Lx, 05.40.Fb

I. INTRODUCTION

The idea of exploiting the quantum-mechanical behavior of a device to gain power in simulating quantum systems was first introduced by Feynman [1]. The field of quantum computing has since grown enormously with the discovery of two algorithmic pillars: Shor's factoring algorithm [2] and Grover's search algorithm [3]. Both of these demonstrate a clear speedup over their classical counterparts. Following in this path, many other quantum algorithms have been developed that provide a speedup [4–6]. A more recent addition to the family of quantum algorithms which demonstrate an exponential speedup are those based on the quantum random walk—the quantum version of the successful classical random walk [7].

There is, however, a need to explore more than the simple computational properties of the algorithms. They must also be experimentally tested in real devices and their relative ease of implementation compared and considered. In particular, in quantum-information-processor (QIP) devices where we are controlling the natural Hamiltonian, it is important to choose a system where the Hamiltonian is amenable to automatic and systematic control. This can be explored by implementing the same algorithm in different molecules and contrasting the performance. Although many different implementation schemes have been proposed for the quantum-random-walk algorithm, using, for example, trapped ions [8], an optical lattice [9], cavity QED [10], or an optical cavity [11], these have not been tested. The only experimental test of a quantum walk is the continuous-time version of a quantum walk on a square using a two-qubit nuclear magnetic resonance (NMR) QIP [12]. This work showed the contrast between a classical and quantum random walk and showed the influence of entanglement on the probability distribution of the quantum walk. Here, we present an experimental proof of principal experiment of a discrete-time quantum walk on a

square. The effects of decoherence on the quantum random walk has been investigated by several authors and, indeed, it may offer some benefits [13,14]. Therefore, we also explored the quantum to classical transition of our walk under the addition of decoherence to the quantum register. Furthermore, we compared and contrasted two different control schemes and molecules by implementing the algorithm on two molecules.

II. QUANTUM RANDOM WALKS

In the development of deterministic classical randomized algorithms, the methods of Markov chains and random walks have played a fundamental role [15]. These algorithms can be divided into two categories: continuous-time random walks when the walker has a probability per unit time to make a move and discrete-time random walks where the walker moves at defined time steps. Since these processes are stochastic, it is not surprising that they have quantum counterparts. The quantum versions, however, show remarkable differences with their classical analogs. The continuous-time quantum walk (CTQW) [16] has been shown to provide an exponential speedup in propagation through a graph [7,17]. The discrete-time quantum walk (DTQW) [18] plays an important role in the speedup of a quantum algorithm design for spatial searching [19–21].

One step of a classical discrete-time random walk on a circle with n nodes, denoted by $\{0, \dots, n-1\}$, is performed by repetition of the following two steps: (1) the walker first flips a coin and then (2) moves either clockwise or counterclockwise depending on the outcome of the coin toss.

If we perform a quantum-mechanical treatment of the situation, we can label the nodes with a mutually orthonormal set of state vectors $\{|i\rangle\}_{i=0}^{n-1}$. The coined DTQW on the circle can be seen as “quantumly” flipping a coin degree of freedom using a unitary operation and then coherently moving the walker position degree of freedom clockwise, or counterclockwise, conditioned on the state of the coin [22].

^{*}Electronic address: c4ryan@iqc.ca




FIG. 1. Comparison of the dynamics of the classical (left) and quantum (right) random walk on a square for three steps. H or T represents the state of the coin and the number the position of the walker at one of the four nodes of the square. p is the probability of each classical state and A is the probability amplitude for the quantum state. Part (a) for each step is the coin flip and part (b) the movement around the square. In both cases the walker starts at node 0 with the coin in the heads state. After one step he has a 50% probability of being at either node 1 or 3. Then, in the second step he goes to either 0 or 2 with a 50% probability. In the third step, however, the two types of walk diverge. The classical walk continues to oscillate and the probability remains spread out. In the quantum walk, on the other hand, the probability amplitudes interfere and cancel out, leaving all the probability in one corner after three steps.

For a Hadamard walk, the coin flipping operation is simply the Hadamard gate described by the matrix

$$\hat{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \quad (1)$$

Now the conditional shift operator is defined as

$$\hat{S}|H\rangle|i\rangle = |H\rangle|i \oplus 1\rangle, \quad (2)$$

$$\hat{S}|T\rangle|i\rangle = |T\rangle|i \oplus 1\rangle, \quad (3)$$

where \oplus and \ominus are understood to be addition and subtraction modulo n and $|H\rangle$ and $|T\rangle$ describe the two basis states of the coin. Therefore, if the walker is in position $|i\rangle$, he will move clockwise to the position $|i \oplus 1\rangle$ if the coin is in the state $|H\rangle$ or counterclockwise to $|i \oplus 1\rangle$ if the coin in the state $|T\rangle$. We can write this operator as

$$\hat{S} = \sum_{i=0}^{n-1} (|H\rangle\langle H| \otimes |i \oplus 1\rangle\langle i| + |T\rangle\langle T| \otimes |i \oplus 1\rangle\langle i|). \quad (4)$$

Then one step of the DTQW is defined as applying the operator

$$\hat{W} = \hat{S}(\hat{H} \otimes 1). \quad (5)$$

On a circle, this type of algorithm shows destructive interference effects and a probability distribution that is periodic in time. The contrasting dynamics for the classical and quantum random walks are shown in Fig. 1. As opposed to the classical walk where the probability is always spread out, the quantum walk has steps where the probability amplitudes interfere such that all the probability comes back to one node. Furthermore, this walk is periodic in that after eight steps, the corresponding propagator is equal to the identity and the system comes back to its original state.

In our experimental setup we have three qubits available, which allows one qubit to describe the coin state and two for the position state. Thus, we have $n=4$, and we are performing a discrete quantum walk on a square. The shift operator

defined in Eq. (4) would require a complicated quantum circuit involving a Toffoli gate. We can simplify the circuit required by using a shifting operator that moves the walker along a direction vector—i.e., horizontally or vertically (this also is analogous to the random walk on the hypercube [23]). Therefore, if we label the corners of the square as shown in Fig. 2, the shift operator on the three qubit register becomes

$$\begin{aligned} \hat{S} &= \hat{P}_H \hat{X}^2 + \hat{P}_T \hat{X}^3 = (\hat{P}_H \hat{X}^2 + \hat{P}_T) (\hat{P}_T \hat{X}^3 + \hat{P}_H) \\ &= (\hat{X}^1 \text{CNOT}^{1,2} \hat{X}^1) \text{CNOT}^{1,3}, \end{aligned} \quad (6)$$

where CNOT denotes controlled-NOT, X denotes the standard σ_x Pauli matrix, $P_{H/T}$ are the projectors on the two coin states, and the superscript indicates on which of the qubits the action is performed. Here, it is understood that the first qubit represents the coin and the second and third the position register. The resulting probabilities for each step are shown in Table I.

III. LIQUID-STATE NMR QUANTUM-INFORMATION PROCESSING

A. Basic principles

A liquid-state NMR QIP consists of an ensemble of roughly 10^{20} identical molecules dissolved in a liquid sol-




FIG. 2. Logical labeling of the nodes on which we implemented the DTQW. With this labeling, flipping the first qubit corresponds to a horizontal move and flipping the second qubit a vertical move.

TABLE I. Probability to be in each of the corner states as denoted in Fig. 2. While in the classical random walk the probability always remains spread out between two corners, in the quantum random walk all the probability returns to one corner at certain time steps.

Corner	Classical				Quantum			
	0	1	2	3	0	1	2	3
Step 0	1				1			
Step 1		0.5		0.5		0.5		0.5
Step 2	0.5		0.5		0.5		0.5	
Step 3		0.5		0.5				1
Step 4	0.5		0.5				1	
Step 5		0.5		0.5		0.5		0.5
Step 6	0.5		0.5		0.5		0.5	
Step 7		0.5		0.5		1		
Step 8	0.5		0.5		1			

vent. Due to the fast tumbling motion of the molecules, they are essentially decoupled from each other; ideally, all the molecules have the same evolution. We can think of the quantum register made of qubits that correspond to the spin- $\frac{1}{2}$ nuclei within each molecule. The sample is placed in a strong homogeneous magnetic field which provides the quantization axis and causes the spins to precess around the axis of the field. It is possible to implement single-qubit gates using radio-frequency (rf) pulses resonant with the precession frequency, which can effect a rotation about any axis orthogonal to the axis of the external field. Two-qubit gates are effected through the coupling from the natural Hamiltonian, which produces a controlled phase gate [24].

If the molecule used contains n distinguishable nuclei and the magnetic field is aligned along the z axis, then the system Hamiltonian is approximated by

$$\hat{H} = \pi \sum_{i=1}^n \nu_i \hat{Z}_i + \frac{\pi}{2} \sum_{i>j} J_{ij} \hat{Z}_i \otimes \hat{Z}_j, \quad (7)$$

where ν_i is the Larmor frequency of spin i in Hz, J_{ij} is the coupling strength between spin i and j in Hz and Z is the conventional Pauli operator σ_z . The interaction part of the Hamiltonian can be approximated to the above Ising form (weak-coupling regime or secular approximation) only if the difference between any two nuclei Larmor frequencies is much greater than the coupling between the nuclei. We can also turn off the coupling between any two spins as needed by applying refocusing rf pulses.

B. Implementing dephasing in NMR

We can apply a controllable amount of decoherence to selected spins using gradient techniques in NMR. Consider only one nucleus with state ρ , and suppose we work in the rotating frame of that spin. On a NMR spectrometer, it is possible to apply a gradient to the external magnetic field. During the time that the gradient is applied, the spins will precess at different frequencies depending on their position in the sample. The state of the ensemble will then be given by an average over the observable sample,

$$\rho' = \frac{1}{2a} \int_{-a}^a e^{-i(\alpha' \gamma t z/2) \hat{Z}} \rho e^{i(\alpha' \gamma t z/2) \hat{Z}} dz, \quad (8)$$

where $2a$ is the length of the sample, t is the interval of time the gradient is being applied, and $\alpha' = \alpha/\hbar$ and $\gamma = \nu/B_z$, the gyromagnetic ratio of the nucleus. If we compute the integral, it can be shown that

$$\rho' = (1-p)\rho + p \hat{Z} \rho \hat{Z}, \quad (9)$$

$$p = \frac{1}{2} \left(1 - \frac{1}{\alpha' \gamma t a} \sin(\alpha' \gamma t a) \right),$$

which is the exact form of a z -dephasing decoherence. The amount of dephasing can be controlled by the strength and time of the gradient pulse. Particular spins can be protected from the applied decoherence by applying a 180° rotation and applying a second gradient of the same strength and time. This second gradient will reverse the dephasing of the rotated spins and double it on the spins that were not rotated.

IV. EXPERIMENT

We implemented the quantum walk algorithm on two molecules: *trans*-crotonic acid and trichloroethylene (TCE). This allowed us to compare the quality of two different methods of control and the merits of the two molecules.

A. Implementation on crotonic acid

The seven-qubit molecule *trans*-crotonic acid (four carbons, two hydrogens, and one methyl group) has been used in experimental demonstrations of quantum algorithms, such as quantum error correction [25,26] and quantum simulations [27]. In this experiment, we used the carbon backbone of labeled *trans*-crotonic acid in a solution of deuterated acetone. The hydrogen nuclei were decoupled using standard heteronuclear decoupling techniques [28]. We used C_3 as the coin and C_2 and C_4 as the position register (see Fig. 3). C_1 was used as a labeling spin to ease the creation of the initial state. On a Bruker DRX Avance 600 NMR spectrometer, the molecule has the Hamiltonian parameters shown in Fig. 3.




FIG. 3. (Color online) Molecular structure of *trans*-crotonic acid and its Hamiltonian parameters. The chemical shifts are given as the diagonal elements and the coupling strength (Hz) by the off-diagonal elements. Note that since the darkly shaded unlabeled nuclei are oxygen whose natural abundance of ^{16}O with 0 spin is close to 100%, the two oxygen nuclei do not couple with the rest of the molecule and can be ignored. Lightly shaded unlabeled nuclei are hydrogen which were decoupled during the experiment.

Since our system is homonuclear, the control of individual qubits is achieved through soft Gaussian-like rf pulses at the Larmor frequency of the target nucleus. The length of the soft pulses is of the order of the inverse of the smallest chemical shift difference with the other nuclei. In our experiment the length of the selective pulses on C_1 and C_2 , C_3 , C_4 were 192 μs and 704 μs , respectively.

1. Initial-state preparation

The experiment required the initial state

$$\rho_{\text{in}} = |000\rangle\langle 000| = \frac{1}{8} (1 + Z)(1 + Z)(1 + Z). \quad (10)$$

We created the labeled pseudopure state Z000 (using the notation $C_1C_2C_3C_4$) following the spatial averaging technique elaborated in [29].

2. Pulse sequence implementation

The unitary of one step of the DTQW from Eq. (6) was translated to a sequence of pulses and coupling gates as shown in Fig. 4. Although many pulse sequences are possible through the use of commutation rules, this particular one was designed to be the most efficient due to the cancellations possible during multiple-step sequences. Moreover, the ZZ gates are achieved simultaneously, which shortens the overall




FIG. 4. NMR pulse sequence representing one step of DTQW. The notation R_i^θ means a rotation of an angle θ around the axis i . Refocusing pulses are not shown. Since each nucleus is tracked in its own rotating frame, rotations about the z axis are implemented instantaneously through a change of reference frame.

pulse sequence, thus reducing decoherence effects; the slight mismatch in coupling strengths was taken care of with a refocusing pulse on C_4 near the end of the coupling period [30]. Commutation rules were also used to cancel pulses between the final step and the readout pulses.

The ideal pulse sequence of rotations and couplings was then input into a pulse sequence compiler which numerically optimized the timing and phases of the pulses. During the selective excitations, the first-order deviations from the ideal pulse are simply phase (the Bloch-Siegert shift) rotations on the other qubits and ZZ couplings that occur during the finite-length pulse. The compiler presimulates the selective rf pulses using an efficient pairwise simulation and then decomposes the simulated unitary into phase and coupling errors sandwiching the ideal selective pulse. These errors can then be taken into account by the refocusing scheme and phase of the pulses, so that the overall unitary is as close to the desired one as possible.

Since we are concerned with the final state of only three qubits in this experiment, complete state tomography is still feasible. On a three-qubit system in NMR, only seven different readout pulses are required to rotate each term of the density matrix into observable simple single coherences.¹ And since we were operating on a homonuclear system, observing the signal from all spins in one experiment was possible, with some post-processing to adjust for the correct phase of each individual rotating frame. The coupling between the labeling spin C_1 and the other three qubits is resolvable, and so the presence of the labeling spin does not interfere with the tomography of C_2 , C_3 , and C_4 .

3. Experimental results

For the state tomography each of the peaks in the spectra were fitted using absorption and dispersion Lorentzian peaks. The full density matrix was then reconstructed by appropriately summing up the corresponding Pauli terms. Where two experiments gave values for the same density matrix terms, the values were simply averaged. As we observed only C_2 , C_3 , and C_4 , the term ZIII could not be determined. A suitable amount of that term was subsequently added to the density matrix so as to make the initial state as close to Z000 as possible. This amount was then kept constant for the density matrix reconstruction in subsequent experiments.

To quantify the success of our experiments, we computed the fidelity of the experimental density matrix to both the ideal and simulated results. In NMR, all states are nearly completely mixed and the fidelity measure introduced in [31] is appropriate. We can compare one density matrix to another using the formula

$$F^{A,B} = \frac{\text{Tr}(\rho^A \rho^B)}{\sqrt{\text{Tr}[(\rho^A)^2]} \sqrt{\text{Tr}[(\rho^B)^2]}}. \quad (11)$$

We made two comparisons. First, we compare the experimentally determined density matrix to the theoretically expected result. The theoretical result is achieved by multiplying the ideal initial state by the ideal propagator. To

¹Readout pulses yII,Ily,Ilx,yyI,Ixx,yyy,xxx are sufficient.

TABLE II. Fidelities (in percent) of experimental and simulated results. The first column gives the fidelity of the experimental density matrices determined from the tomography, with respect to the theoretical expected density matrix. The second column gives the fidelity of the simulation results. Errors are estimated from the fitting procedure. Note that since computer simulation of the spatial averaging that occurs during the pseudopure preparation is difficult and inaccurate, the initial state for the simulation was the experimental pseudopure state determined from the tomography. The fluctuations observed in the fidelity come from uncertainties in the fit and instabilities in the spectrometer over the course of the experiment.

	Experimental	Simulated
Step 0	98±5	—
Step 1	97±5	98
Step 2	98±5	98
Step 3	92±5	98
Step 4	99±5	98
Step 5	94±5	97
Step 6	96±5	97
Step 7	96±5	97
Step 8	87±4	97

investigate how well we understand our control of the system, we also compare the fidelity of the results from a simulation of the experiment to the theoretical result.

The fidelities of simulated and experimental results are compared in Table II and Fig. 5 give a sample of the fully reconstructed density matrices for C_2 , C_3 , and C_4 . The loss of fidelity in our experiment, over and above that of the simulated control errors, is explained from three sources which are not taken into account by the simulation. We have losses from T_2 relaxation. Although our pulse sequence is short compared with the T_2 relaxation times, during the quantum-walk algorithm, the state is often in high coherences, which decay much faster than the simple T_2 time. Inhomogeneities in the strong magnetic field also cause extra relaxation and dephasing. Further losses come from inhomogeneities of the rf field used to implement rotations and pulse-angle miscalibration.

4. Addition of decoherence on the coin

In a subsequent experiment, we added dephasing decoherence to the entire qubit register using the technique described in Sec. III B. We expect that the behavior of the quantum walk should converge to the classical walk as the decoherence becomes complete after each step. To demonstrate this claim experimentally, we implemented the quantum random walk for four steps, adding decoherence of a certain strength between each step of the walk. The differences between the quantum and classical walks are manifested in the different probabilities of being in each of the corners after each step. The results are shown in Table III for gradient strengths corresponding to no, partial, and full decoherence.

The divergence between the classical and quantum walks shows most clearly in steps 3 and 4. Whether the walk is

classical or quantum, steps 1 and 2 yield the same measurement probabilities for the position (however, the quantum version with decoherence will have coherent superposition states). Analyzing the data from steps 3 and 4, one can see that the quantum interference, present so clearly in the quantum walk with no decoherence, is less obvious as the amount of decoherence increases. Instead of the probability all collecting in one corner, it remains spread out between two opposite corners—the same as in the classical walk. This can also be seen in part (b) of Fig. 5 by the reduction of the off-diagonal terms of the density matrix as the decoherence becomes stronger and the appearance of diagonal terms which were previously canceled by the quantum interference.

The probabilities even with zero gradient strength do not correspond perfectly to the ideal quantum walk. We believe these errors come from two sources. Because the gradient does not commute with any pulses, we were not able to use commutation rules to reduce the number of pulses during multiple-step experiments. Furthermore, gradient methods are hampered by diffusion and multiple gradients may lead to a return of signal that was “erased” by a previous gradient.

B. Comparison with the TCE molecule

For comparison purposes and to show the importance of choosing a molecule with good characteristics in liquid-state NMR quantum-information processing, we show our results from our initial attempt to implement the quantum walk on the molecule trichloroethylene—a molecule with which we have much less control due to the presence of strong coupling. The molecule has been used for some initial demonstrations of quantum algorithms [32]. A diagram of the molecule and the parameters of its Hamiltonian are shown in Fig. 6.

1. Pseudopure-state preparation

Since the TCE molecule contains only three qubits, we are unable to create the labeled pseudopure state that we used in the crotonic acid experiments. Instead, we chose to use temporal averaging and add three separate experiments to achieve the initial state $|000\rangle$. The three different initial states we used are

$$\begin{aligned}\rho_1 &= Z \otimes (1 + Z) \otimes (1 + Z), \\ \rho_2 &= 1 \otimes Z \otimes (1 + Z), \\ \rho_3 &= 1 \otimes 1 \otimes (1 + Z).\end{aligned}\quad (12)$$

If we add the results of these three experiments, it is equivalent to having performed the algorithm on the initial state:

$$\rho_{\text{in}} = \rho_1 + \rho_2 + \rho_3 = (1 + Z) \otimes (1 + Z) \otimes (1 + Z) = |000\rangle\langle 000|. \quad (13)$$

Since there is only one hydrogen nucleus in the molecule, we can use broadband hard pulses to control it. One useful property of the TCE molecule in a 600-Mhz spectrometer is that the J coupling between the two carbons is almost exactly 10.5 times smaller than the difference in chemical shift ($\Delta\nu$).

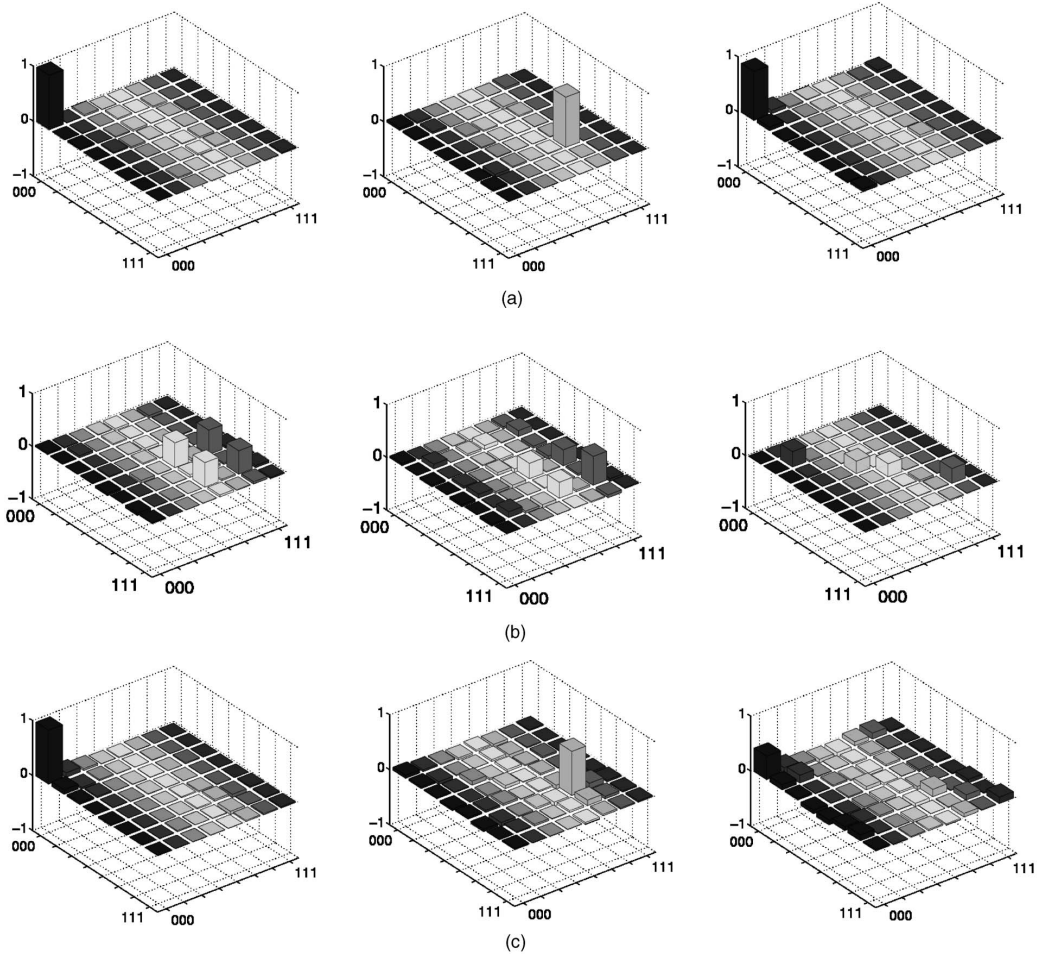


FIG. 5. Examples of the real part of reconstructed density matrices after (a) steps 0, 4, and 8 (left to right) of the quantum walk implemented on crotonic acid; (b) step 3 with no, partial, and full decoherence applied after each step; and (c) steps 0, 4, and 8 of the quantum walk implemented on TCE. The effects of the quantum interference returning all the probability to one corner is clearly evident in steps 4 and 8; however, the fidelity in the TCE case is much worse. The density matrices from the decoherence experiments show how destruction of the off-diagonal coherences prevents all the probability returning to one corner after three steps.

Therefore, during the time for a $(\pi/2)\hat{Z}\hat{Z}$ coupling gate between the two carbons ($\Delta t = 1/2J_{C_1C_2}$), the relative chemical shift evolution of C_2 with respect to C_1 will be $\theta = 2\pi\Delta\nu\Delta t = \pi\Delta\nu/J_{C_1C_2} = -10.5\pi = -\pi/2 \pmod{2\pi}$. Therefore, in the reference frame rotating at the Larmor frequency of C_1 , every time there is a $\pi/2$ coupling between the carbons, an extra $R_z^{-\pi/2}$ is naturally performed on C_2 .

The chemical shift difference between the two carbons is small and the coupling between them large, so selective pulses were impossible to achieve using the same technique of Gaussian-shaped pulses used in the crotonic acid experiments. These pulses would be very long (roughly 5 ms) and the large coupling errors that would occur during the pulse would be difficult to refocus. Instead, it was possible to perform selective pulses using hard pulses and the chemical

shift evolution. To illustrate the technique, we demonstrate how to perform a selective $\pi/2$ rotation of C_2 . If we use a reference frame rotating at the Larmor frequency of C_1 , then, during a time $\tau = 1/4\Delta\nu$, the spin C_1 will not precess while C_2 will undergo a rotation of $-\pi/2$ around the z axis. Since $1/4\Delta\nu$ is much less than the coupling time $1/2J_{C_1C_2}$, we can ignore the coupling between the two carbons and refocus only the hydrogen. Using this selective z rotation combined with hard pulses that rotate the two carbons together, we can perform a $\pi/2$ rotation with phase ϕ on only C_2 as follows:

$$\begin{aligned} & (R_{\phi-\pi/2}^{\pi/2} \otimes R_{\phi-\pi/2}^{\pi/2})(1 \otimes R_z^{-\pi/2})(R_{\phi+\pi/2}^{\pi/2} \otimes R_{\phi+\pi/2}^{\pi/2}) \\ &= 1 \otimes R_{\phi-\pi/2}^{\pi/2} R_z^{-\pi/2} R_{\phi+\pi/2}^{\pi/2} = 1 \otimes R_{\phi}^{\pi/2}. \end{aligned} \quad (14)$$

Similar pulse sequences can be derived to perform a π rotation on C_2 and selective pulses on C_1 . Because of the

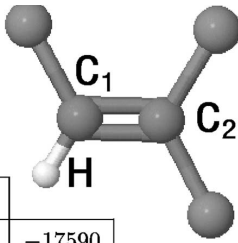
TABLE III. Estimate of quantum probability to be in each corner of the square for one through four steps (cf. Table I). The results were obtained for gradient strengths corresponding to no, partial, and full decoherence. The probabilities were obtained by reconstructing the density matrix using the same fitting software used before and then applying the position measurement projectors to the reconstructed density matrix.

Corner	Quantum walk with decoherence											
	None				Partial				Full			
	0	1	2	3	0	1	2	3	0	1	2	3
Step 0	100±8	0±1	4±1	-2±1	100±8	0±1	4±1	-2±1	100±8	0±1	4±1	-2±1
Step 1	2±1	57±4	-1±1	43±3	2±1	58±4	-2±1	44±3	0±1	59±4	-2±1	45±4
Step 2	57±4	1±1	44±3	-1±1	50±4	7±1	40±3	5±1	51±4	4±1	46±40	1±1
Step 3	7±1	14±1	3±1	78±6	2±1	14±1	1±1	84±6	-1±1	53±4	-3±1	53±4
Step 4	15±1	-1±1	84±6	3±1	19±2	1±1	78±6	4±1	50±4	0±1	53±4	-1±1

different form of selective pulses used, the pulse sequences were written and optimized by hand. This required a different pulse-sequence implementation of the quantum-walk unitary, which avoided as much as possible selective pulses and z rotations where possible. The one z rotation used is a natural outcome of the C_1 - C_2 coupling gate as described above. This alternative pulse sequence is shown in Fig. 7.

2. Experimental results

Fidelity results, similar to those calculated for the crotonic acid experiments, are shown in Table IV, and a sample of reconstructed density matrices can be seen in Fig. 5(c). Clearly this experiment was not as successful as the implementation on the crotonic acid molecule. There are two main reasons for this loss of fidelity. First, the chemical shift difference between the two carbons is very small. Because of this, the secular approximation no longer holds and thus the coupling between the two carbon spins can no longer be approximated by the Ising form $Z_{C_1} \otimes Z_{C_2}$. Indeed, it has to take all the strong-coupling terms into account—i.e., $\vec{S}_{C_1} \cdot \vec{S}_{C_2} = X_{C_1} \otimes X_{C_2} + Y_{C_1} \otimes Y_{C_2} + Z_{C_1} \otimes Z_{C_2}$.



H	-3891		
C ₁	201	-17590	
C ₂	8.9	103	-18675
	H	C ₁	C ₂

FIG. 6. Diagram of ^{13}C labeled TCE. The chemical shifts and couplings are given in the table. Note that since the chlorine nuclei (unlabeled) have a spin of $\frac{3}{2}$, they have an electric-quadrupole moment which causes them to decohere quickly and they have a very small coupling to the rest of the molecule which we can ignore in the natural Hamiltonian of the molecule.

Unfortunately, this strong coupling renders our ideal ZZ gates much less precise. Every coupling gate performed added XX and YY error terms which we could not refocus. This coupling also caused problems during our selective carbon rotations. Although the coupling is small, there is an unrefocusable coupling of $\pi J_{C_1 C_2} / 4\Delta\nu \approx 4.27^\circ$. Our only way to minimize these errors was to optimize the delay times analytically and from numerical simulations. However, these did not correspond well to the experimentally determined optimal values. This point also clearly demonstrates the second reason for the less satisfactory results on TCE. We were unable to use the numerical optimization of the pulse-sequence compiler used for crotonic acid. The compiler provides a systematic and reliable way to produce pulse sequences that implement unitaries with high fidelity and is clearly superior to writing and optimizing pulse sequences by hand. These experiments also showed the limits of our simulator. For the crotonic acid experiments, where only soft pulses were used, the rf power applied changed slowly and the simulator was faithful to what rf power the spins were experiencing. In TCE, where control was achieved only through short hard pulses, other effects such as phase transients enter and the spins might experience an rf field much different from the ideal square pulse simulated. To fully understand the issues surrounding hard-pulse control a much more detailed study of the probe response must be undertaken. This underlines a key point: control of a more complex and strongly coupled system could be obtained through sophisticated control techniques such as strongly modulating




FIG. 7. Circuit used to implement one step of the DTQW on the TCE. The z rotation on C_2 occurs naturally during the coupling with C_1 . Note also that the refocusing pulses are not shown in this pulse sequence.

TABLE IV. Experimental and simulated fidelities (in percent) for the implementation of eight steps of the DTQW on the molecule TCE. Again, simulation of the pseudopure state was not performed.

	Experimental	Simulated
Step 0	98±6	—
Step 1	85±5	96
Step 2	82±4	94
Step 3	70±4	93
Step 4	80±4	90
Step 5	76±4	89
Step 6	65±4	86
Step 7	53±4	84
Step 8	43±4	83

pulses [31]; however, it seems prudent to invest the effort in a wise choice of molecule.

V. CONCLUSION

We have presented the first experimental implementation of a coined discrete-time quantum walk. It showed a clear difference with the classical coined quantum walk, since the DTQW possesses destructive interference and periodicity in its evolution. A proof of principle like this lays down the path to more elaborate experiments using discrete quantum walks, such as the database searches, walks on a hypercube or N -node circle, or a more profound study of the effect of decoherence on the walk. This paper also demonstrates the importance of choosing a natural Hamiltonian well suited to automated control in the context of quantum-information processing.

ACKNOWLEDGMENTS

All of the authors would like to thank M. J. Ditty for technical support with the work on the spectrometer and C. Negrevergne for helpful discussions. This work has been supported by NSERC, ARDA, and CFI.

-
- [1] R. P. Feynman, *Int. J. Man-Mach. Stud.* **21**, 467 (1982).
 [2] P. W. Shor, in *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, edited by S. Goldwasser (IEEE Computer Society, Los Alamitos, CA, 1994), pp. 124–134.
 [3] L. K. Grover, *Phys. Rev. Lett.* **79**, 4709 (1997).
 [4] A. Kitaev, e-print quant-ph/9511026.
 [5] D. Simon, *SIAM J. Comput.* **26**, 1474 (1997).
 [6] A. Ekert and R. Jozsa, *Philos. Trans. R. Soc. London, Ser. A* **356**, 1769 (1998).
 [7] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, in *Proceedings of the 35th ACM Symposium on Theory of Computing* (ACM Press, New York, 2002), pp. 59–68.
 [8] B. C. Travaglione and G. J. Milburn, *Phys. Rev. A* **65**, 032310 (2002).
 [9] W. Dür, R. Raussendorf, V. M. Kendon, and H. J. Briegel, *Phys. Rev. A* **66**, 052319 (2002).
 [10] B. C. Sanders, S. D. Bartlett, B. Tregenna, and P. L. Knight, *Phys. Rev. A* **67**, 042305 (2003).
 [11] P. L. Knight, E. Roldán, and J. E. Sipe, *Opt. Commun.* **227**, 147 (2003).
 [12] J. Du, H. Li, X. Xu, M. Shi, J. Wu, X. Zhou, and R. Han, *Phys. Rev. A* **67**, 042316 (2003).
 [13] V. Kendon and B. Tregenna, *Phys. Rev. A* **67**, 042315 (2002).
 [14] T. A. Brun, H. A. Carteret, and A. Ambainis, *Phys. Rev. Lett.* **91**, 130602 (2003).
 [15] R. Motwani and P. Raghavan, *Randomized Algorithms* (Cambridge University Press, Cambridge, England, 1995).
 [16] E. Farhi and S. Gutmann, *Phys. Rev. A* **58**, 915 (1998).
 [17] A. Childs, E. Farhi, and S. Gutmann, *Quantum Inf. Process.* **1**, 35 (2002).
 [18] Y. Aharonov, L. Davidovich, and N. Zagury, *Phys. Rev. A* **48**, 1687 (1993).
 [19] N. Shenvi, J. Kempe, and K. Birgitta Whaley, *Phys. Rev. A* **67**, 052307 (2003).
 [20] S. Aaronson and A. Ambainis *Theory of Computing Volume 1*, 47 (2005).
 [21] A. Ambainis, J. Kempe, and A. Rivosh, in *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, British Columbia, Canada, 2005*, (SIAM, 2005), pp. 1099–1108.
 [22] D. Aharonov, A. Ambainis, J. Kempe, and U. Vazirani, in *Proceedings of the 33th ACM Symposium on Theory of Computing* (ACM Press, New York, 2001), pp. 50–59.
 [23] C. Moore and A. Russell, in *Proceedings of RANDOM 2002*, edited by J. D. P. Rolim and S. Vadhan (Springer, Cambridge, MA, 2002), pp. 164–178.
 [24] R. Laflamme *et al.*, e-print quant-ph/0207172.
 [25] E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne, *Phys. Rev. Lett.* **86**, 5811 (2001).
 [26] N. Boulant, L. Viola, E. M. Fortunato, and D. G. Cory, *Phys. Rev. Lett.* **94**, 130501 (2005).
 [27] C. Negrevergne, R. Somma, G. Ortiz, E. Knill, and R. Laflamme, *Phys. Rev. A* **71**, 032344 (2005).
 [28] A. J. Shaka, J. Keepler, and R. Freeman, *J. Magn. Reson.* (1969-1992) **53**, 313 (1983).
 [29] E. Knill, R. Laflamme, R. Martinez, and C.-H. Tseng, *Nature (London)* **404**, 368 (2000).
 [30] M. D. Bowdrey, J. A. Jones, E. Knill, and R. Laflamme, *Phys. Rev. A* **72**, 032315 (2005).
 [31] E. M. Fortunato, M. A. Pravia, N. Boulant, G. Teklemariam, T. F. Havel, and D. G. Cory, *J. Chem. Phys.* **116**, 7599 (2002).
 [32] D. G. Cory, M. D. Price, W. Maas, E. Knill, R. Laflamme, W. H. Zurek, T. F. Havel, and S. S. Somaroo, *Phys. Rev. Lett.* **81**, 2152 (1998).

Appendix E

Publications

- **C.A. Ryan**, M. Laforest and R. Laflamme. Randomized benchmarking of single and multi-qubit control in liquid-state NMR quantum information processing. *arXiv*: 0808.3973v1 (2008).
- **C.A. Ryan**, C. Negrevergne, M. Laforest, E. Knill and R. Laflamme. Liquid-state nuclear magnetic resonance as a testbed for developing quantum control methods. *Physical Review A* **78**, 012328 (2008).
- **C.A. Ryan**, O. Moussa, J. Baugh and R. Laflamme. A spin based heat engine: multiple rounds of algorithmic cooling. *Physical Review Letters* **100**: 149905 (2008).
- Michael K. Henry, Chandrasekhar Ramanathan, Jonathan S. Hodges, **Colm A. Ryan**, Michael J. Ditty, Raymond Laflamme and David G. Cory. Fidelity enhancement by logical qubit encoding. *Physical Review Letters* **99**: 220501 (2007).
- Joseph Emerson, Marcus Silva, Osama Moussa, **Colm Ryan**, Martin Laforest, Jonathan Baugh, David G. Cory and Raymond Laflamme. Symmetrised Characterisation of Noisy Quantum Processes. *Science* **317**: 1893 (2007).
- HyungJoon Cho, Jonathan Baugh, **Colm A. Ryan**, David G. Cory and Chandrasekhar Ramanathan. Low temperature probe for dynamic nuclear polarization and multiple-pulse solid-state NMR. *Journal of Magnetic Resonance* **187**: 242-250 (2007).
- C. Negrevergne, T. Mahesh, **C.A. Ryan**, M. Ditty, F. Cyr-Racine, W. Power, N. Boulant, T. Havel, D. Cory, and R. Laflamme. Benchmarking quantum control methods on a 12 qubit system. *Physical Review Letters* **96**: 170501 (2006).

- J. Baugh, O. Moussa, **C.A. Ryan**, R. Laflamme, C. Ramanathan, T.F. Havel, D.G. Cory. Solid-state NMR three-qubit homonuclear system for quantum information processing: control and characterization. *Physical Review A* **73**: 022305 (2006).
- **C.A. Ryan**, M. Laforest, J.C. Boileau and R. Laflamme. Experimental implementation of a discrete-time quantum random walk on an NMR quantum-information processor. *Physical Review A* **72**: 062317 (2005).
- **C.A. Ryan**, J. Emerson, D. Poulin, C. Negrevergne and R. Laflamme. Characterization of complex quantum dynamics with a scalable NMR information processor. *Physical Review Letters* **95**: 250502 (2005).
- J. Baugh, O. Moussa, **C. A. Ryan**, A. Nayak and R. Laflamme. Experimental implementation of heat-bath algorithmic cooling using solid-state nuclear magnetic resonance. *Nature* **438**: 470-473 (2005)

Bibliography

- [1] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. 1
- [2] J. Preskill. Reliable quantum computers. *Proceedings of the Royal Society London A*, 454(1969):385–410, 1998. 1, 57
- [3] Leonard J. Schulman and Umesh V. Vazirani. Molecular scale heat engines and scalable quantum computation. In *STOC '99: Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 322–329, New York, NY, USA, 1999. ACM Press. 2
- [4] J. A. Jones and E. Knill. Efficient refocusing of one-spin and two-spin interactions for nmr quantum computation. *Journal of Magnetic Resonance*, 141(2):322–325, 1999. 4, 26, 28
- [5] C. A. Ryan, C. Negrevergne, M. Laforest, E. Knill, and R. Laflamme. Liquid-state nuclear magnetic resonance as a testbed for developing quantum control methods. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 78(1):012328, 2008. 5
- [6] Cleve Moler and Charles Van Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review*, 45(1):3–49, 2003. 5
- [7] T. Schulte-Herbrüggen, A. Spörl, and S. J. Glaser. Quantum cisc compilation by optimal control and scalable assembly of complex instruction sets beyond two-qubit gates. *arXiv*, page arXiv:0712.3227v1, 2007. 5, 33
- [8] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52(5):3457–3467, Nov 1995. 6
- [9] Viswanath Ramakrishna, Murti V. Salapaka, Mohammed Dahleh, Herschel Rabitz, and Anthony Peirce. Controllability of molecular systems. *Phys. Rev. A*, 51(2):960–966, Feb 1995. 6

- [10] S. G. Schirmer, H. Fu, and A. I. Solomon. Complete controllability of quantum systems. *Phys. Rev. A*, 63(6):063410, May 2001. 6
- [11] Jennifer L. Dodd, Michael A. Nielsen, Michael J. Bremner, and Robert T. Thew. Universal quantum computation and simulation using any entangling hamiltonian and local unitaries. *Phys. Rev. A*, 65(4):040301, Apr 2002. 6
- [12] Richard R. Ernst. *Principles of Nuclear Magnetic Resonance in One and Two Dimensions*. International Series of Monographs on Chemistry. Oxford University Press, 1990. 7, 25
- [13] Joseph Fitzsimons and Jason Twamley. Globally controlled quantum wires for perfect qubit transport, mirroring, and computing. *Physical Review Letters*, 97(9):090502, 2006. 7
- [14] Malcolm H. Levitt. *Spin Dynamics*. Wiley, 2001. 8, 39
- [15] Ray Freeman. Shaped radiofrequency pulses in high resolution nmr. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 32(1):59–106, 1998. 9, 19, 22, 36
- [16] Jintong Mao, T. H. Mareci, K. N. Scott, and E. R. Andrew. Selective inversion radiofrequency pulses by optimal control. *Journal of Magnetic Resonance (1969)*, 70(2):310–318, 1986. 9
- [17] Alberto Carlini, Akio Hosoya, Tatsuhiko Koike, and Yosuke Okudaira. Time-optimal unitary operations. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 75(4):042308, 2007. 9
- [18] M. A. Nielsen. A simple formula for the average gate fidelity of a quantum dynamical operation. *Physics Letters A*, 303(4):249–252, October 2002. 11
- [19] Evan M. Fortunato, Marco A. Pravia, Nicolas Boulant, Grum Teklemariam, Timothy F. Havel, and David G. Cory. Design of strongly modulating pulses to implement precise effective hamiltonians for quantum information processing. *The Journal of Chemical Physics*, 116(17):7599–7606, 2002. 11, 101, 114
- [20] Raj Chakrabarti and Herschel Rabitz. Quantum control landscapes. *International Reviews in Physical Chemistry*, 26(4):671 – 735, 2007. 11
- [21] Tobias Gradl, A. Spörl, Thomas Huckle, Steffen J. Glaser, and Thomas Schulte-Herbrüggen. Parallelising matrix operations on clusters for an optimal control-based quantum compiler. In *Euro-Par*, pages 751–762, 2006. 11, 37

- [22] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbruggen, and Steffen J. Glaser. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2):296–305, 2005. 12, 72, 101, 102
- [23] Kyryl Kobzar, Thomas E. Skinner, Navin Khaneja, Steffen J. Glaser, and Burkhard Luy. Exploring the limits of broadband excitation and inversion pulses. *Journal of Magnetic Resonance*, 170(2):236–243, 2004. 13, 60
- [24] Ilya Kuprov, Nicola Wagner-Rundell, and P. J. Hore. Polynomially scaling spin dynamics simulation algorithm based on adaptive state-space restriction. *Journal of Magnetic Resonance*, 189(2):241–250, 2007. 13
- [25] J Werschnik and E K U Gross. Tailoring laser pulses with spectral and fluence constraints using optimal control theory. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S300–S312, 2005. 17
- [26] Thomas M. Barbara, Joel F. Martin, and Jon G. Wurl. Phase transients in nmr probe circuits. *Journal of Magnetic Resonance*, 93:497–508, 1991. 17, 60
- [27] Jonathan S. Hodges. Bandwidth limited grape. *unpublished*, 2007. 17
- [28] E Knill, R. Laflamme, R. Martinez, and C.-H Tseng. An algorithmic benchmark for quantum information processing. *Nature*, 404:368–370, March 2000. 20, 22, 61, 114
- [29] Lyndon Emsley and Geoffrey Bodenhausen. Phase shifts induced by transient Bloch-Siegert effects in NMR. *Chemical Physics Letters*, 168(3-4):297–303, May 1990. 19
- [30] F. Bloch and A. Siegert. Magnetic resonance for nonrotating fields. *Phys. Rev.*, 57(6):522–527, Mar 1940. 19
- [31] U. Haeberlen. *Advances in magnetic resonance*, chapter High resolution NMR in solids. Academic Press, 1976. 19, 51
- [32] L. M. K. Vandersypen. *Experimental Quantum Computation with Nuclear Spins in Liquid Solution*. PhD thesis, Stanford University, 2001. 22
- [33] W. M. Witzel, Rogerio de Sousa, and S. Das Sarma. Quantum theory of spectral-diffusion-induced electron spin decoherence. *Physical Review B (Condensed Matter and Materials Physics)*, 72(16):161306, 2005. 25
- [34] Debbie W. Leung, Isaac L. Chuang, Fumiko Yamaguchi, and Yoshihisa Yamamoto. Efficient implementation of coupled logic gates for quantum computation. *Phys. Rev. A*, 61(4):042310, Mar 2000. 28

- [35] M. D. Bowdrey, J. A. Jones, E. Knill, and R. Laflamme. Compiling gate networks on an ising quantum computer. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 72(3):032315, 2005. 28, 32
- [36] Yaakov S. Weinstein, Timothy F. Havel, Joseph Emerson, Nicolas Boulant, Marcos Saraceno, Seth Lloyd, and David G. Cory. Quantum process tomography of the quantum fourier transform. *The Journal of Chemical Physics*, 121(13):6117–6133, 2004. 33, 36
- [37] Garrett M. Leskowitz and Leonard J. Mueller. State interrogation in nuclear magnetic resonance quantum-information processing. *Phys. Rev. A*, 69(5):052302, May 2004. 40
- [38] Isaac L. Chuang and M. A. Nielsen. Prescription for experimental determination of the dynamics of a quantum black box. *Journal of Modern Optics*, 44(11-12):2455 – 2467, 1997. 40
- [39] Timothy F. Havel. Robust procedures for converting among lindblad, kraus and matrix representations of quantum dynamical semigroups. *Journal of Mathematical Physics*, 44(2):534–557, 2003. 40
- [40] N. Boulant, T. F. Havel, M. A. Pravia, and D. G. Cory. Robust method for estimating the lindblad operators of a dissipative quantum process from measurements of the density operator at multiple time points. *Phys. Rev. A*, 67(4):042322, Apr 2003. 40
- [41] M. Mohseni and D. A. Lidar. Direct characterization of quantum dynamics. *Physical Review Letters*, 97(17):170501, 2006. 40
- [42] Asher Peres. Stability of quantum motion in chaotic and regular systems. *Physical Review A*, 30(4):1610–1615, October 1984. 41
- [43] Ph. Jacquod, P.G. Silvestrov, and C.W.J. Beenakker. Golden rule decay versus lyapunov decay of the quantum loschmidt echo. *Physical Review E*, 64:055203, October 2001. 41
- [44] Giuliano Benenti and Giulio Casati. Quantum-classical correspondence in perturbed chaotic systems. *Physical Review E*, 65:066205, June 2002. 41
- [45] Joseph Emerson, Yaakov S. Weinstein, Seth Lloyd, and D.G. Cory. Fidelity decay as an efficient indicator of quantum chaos. *Physical Review Letters*, 89(28):284102, December 2002. 41, 47
- [46] Joseph Emerson, Robert Alicki, and Karol Zyczkowski. Scalable noise estimation with random unitary operators. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S347–S352, 2005. 41, 48, 49, 58, 68

- [47] David Poulin, Robin Blume-Kohout, Raymond Laflamme, and Harold Olivier. Exponential speedup with a single bit of quantum information: Measuring the average fidelity decay. *Physical Review Letters*, 92(17):177906, April 2004. 41, 42, 48
- [48] E. Knill and R. Laflamme. Power of one bit of quantum information. *Physical Review Letters*, 81(25):5672–5675, December 1998. 41
- [49] Cesar Miquel, Juan Pablo Paz, Marcos Saraceno, Emanuel Knill, Raymond Laflamme, and Camille Negrevergne. Interpretation of tomography and spectroscopy as dual forms of quantum computation. *Nature*, 418:59–62, July 2002. 42
- [50] Fritz Haake. *Quantum signatures of chaos*. Springer, New York, 2001. 43
- [51] Joseph Emerson, Yaakov S. Weinstein, Marcos Saraceno, Seth Lloyd, and D.G. Cory. Pseudo-random unitary operators for quantum information processing. *Science*, 302:2098–2100, December 2003. 43, 48, 49
- [52] Joseph Emerson, Etera Livine, and Seth Lloyd. Convergence conditions for random quantum circuits. *Physical Review A*, 72:060302, 2005. 43
- [53] T. R. Beals and C.E. Young. Spin-1/2 selection. internal technical document available on SpinWiki, May 2003. 45, 98
- [54] A. O. Caldeira and A. J. Leggett. Influence of damping on quantum interference: An exactly soluble model. *Phys. Rev. A*, 31(2):1059–1066, Feb 1985. 48
- [55] J.-P. Paz and W. H. Zurek. Environment-induced decoherence and the transition from quantum to classical. In R. Kaiser, C. Westbrook, and F. David, editors, *Coherent Atomic Matter Waves*, Session LXXII; july 27 - august 27, 1999, pages 533–614. Les Houches summer school, Springer, Berlin, 2001. 48
- [56] Robin Blume-Kohout and Wojciech H. Zurek. Decoherence from a chaotic environment: An upside-down "oscillator" as a model. *Phys. Rev. A*, 68(3):032104, Sep 2003. 48
- [57] C. A. Ryan, J. Emerson, D. Poulin, C. Negrevergne, and R. Laflamme. Characterization of complex quantum dynamics with a scalable nmr information processor. *Physical Review Letters*, 95(25):250502, 2005. 48
- [58] D.P. DiVincenzo, D.W. Leung, and B.M. Terhal. Quantum data hiding. *Information Theory, IEEE Transactions on*, 48(3):580–598, Mar 2002. 49

- [59] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. 49
- [60] Joseph Emerson, Marcus Silva, Osama Moussa, Colm Ryan, Martin Laforest, Jonathan Baugh, David G. Cory, and Raymond Laflamme. Symmetrized characterization of noisy quantum processes. *Science*, 317:1893–1896, 2007. 49, 56, 70, 72
- [61] Daniel Gottesman. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, Caltech, 1997. 49, 59
- [62] D. P. Burum. Magnus expansion generator. *Phys. Rev. B*, 24(7):3684–3692, Oct 1981. 51
- [63] R. W. Vaughan, D. D. Elleman, L. M. Stacey, W-K. Rhim, and J. W. Lee. A simple, low power, multiple pulse nmr spectrometer. *Review of Scientific Instruments*, 43(9):1356–1364, 1972. 51, 85
- [64] D. G. Cory, J. B. Miller, and A. N. Garroway. Time-suspension multiple-pulse sequences: applications to solid-state imaging. *Journal of Magnetic Resonance*, 90:205–213, 1990. 51, 85
- [65] Ralf Prigl and Ulrich Haeberlen. Limits of resolution in nmr of solids. *Advances in Magnetic and Optical Resonance*, 19:1–58, 1996. 53, 90
- [66] B. M. Fung, A. K. Khitrin, and Konstantin Ermolaev. An improved broadband decoupling sequence for liquid crystals and solids. *Journal of Magnetic Resonance*, 142(1):97–101, 2000. 56, 92
- [67] E. Knill. Quantum computing with realistically noisy devices. *Nature*, 434(7029):39–44, March 2005. 57, 59
- [68] BÈnÈdicte Elena, GaÛl de PaÛpe, and Lyndon Emsley. Direct spectral optimisation of proton-proton homonuclear dipolar decoupling in solid-state nmr. *Chemical Physics Letters*, 398(4-6):532–538, 2004. 57
- [69] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 77(1):012307, 2008. 58, 59
- [70] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs: Constructions and applications. *ArXiv e-prints*, pages quant-ph/0606161v1, June 2006. 58

- [71] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 71(2):022316, 2005. 59
- [72] E. Knill. Scalable quantum computing in the presence of large detected-error rates. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 71(4):042322, 2005. 59
- [73] S. Wimperis. Broadband, narrowband, and passband composite pulses for use in advanced nmr experiments. *Journal of Magnetic Resonance A*, 109:221, 1994. 60
- [74] Holly K. Cummins, Gavin Llewellyn, and Jonathan A. Jones. Tackling systematic errors in quantum logic gates with composite rotations. *Phys. Rev. A*, 67(4):042308, Apr 2003. 60
- [75] David Poulin. Macroscopic observables. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 71(2):022102, 2005. 62
- [76] J. Emerson. Open quantum systems course notes, 2007. 66
- [77] Aram Harrow and Richard Low. Random quantum circuits are approximate 2-designs. *ArXiv e-prints*, page arXiv:0802.1919v1, February 2008. 69
- [78] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70(5):052328, Nov 2004. 69, 73
- [79] Benjamin Lévi, Cecilia C. López, Joseph Emerson, and D. G. Cory. Efficient error characterization in quantum information processing. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 75(2):022314, 2007. 70
- [80] Lorenza Viola and Emanuel Knill. Random decoupling schemes for quantum dynamical control and error suppression. *Physical Review Letters*, 94(6):060502, 2005. 72
- [81] Michael K. Henry, Alexey V. Gorshkov, Yaakov S. Weinstein, Paola Cappellaro, Joseph Emerson, Nicolas Boulant, Jonathan S. Hodges, Chandrasekhar Ramanathan, Timothy F. Havel, Rudy Martinez, and David G. Cory. Signatures of incoherence in a quantum information processor. *Quantum Information Processing*, 6(6):431–444, December 2007. 72, 100
- [82] Thomas E. Skinner, Timo O. Reiss, Burkhard Luy, Navin Khaneja, and Stefan J. Glaser. Reducing the duration of broadband excitation pulses using optimal control with limited rf amplitude. *Journal of Magnetic Resonance*, 167(1):68–74, 2004. 73

- [83] C. A. Ryan, M. Laforest, and R. Laflamme. Randomized benchmarking of single and multi-qubit control in liquid-state nmr quantum information processing. *ArXiv e-prints*, page 0808.3973v1, 2008. 74
- [84] O Sorensen. Polarization transfer experiments in high-resolution nmr spectroscopy. *Progress in NMR Spectroscopy*, 21:503–569, 1989. 77
- [85] Leonard J. Schulman, Tal Mor, and Yossi Weinstein. Physical limits of heat-bath algorithmic cooling. *Physical Review Letters*, 94(12):120501, 2005. 78, 79
- [86] O. Moussa. On heat-bath algorithmic cooling and its implementation in solid-state nmr. Master’s thesis, University of Waterloo, 2005. 78, 91
- [87] Phillip Kaye. Cooling algorithms based on the 3-bit majority. *arXiv*, pages arXiv:quant-ph/0703194v3, March 2007. 79
- [88] Gilles Brassard, Yuval Elias, Jose M. Fernandez, Haggai Gilboa, Jonathan A. Jones, Tal Mor, Yossi Weinstein, and Li Xiao. Experimental heat-bath cooling of spins. *arXiv*, pages arXiv:quant-ph/0511156v1, 2005. 80
- [89] M. H. Levitt, D. Suter, and R. R. Ernst. Spin dynamics and thermodynamics in solid-state nmr cross polarization. *The Journal of Chemical Physics*, 84(8):4243–4255, 1986. 83
- [90] Luciano Müller, Anil Kumar, Thomas Baumann, and Richard R. Ernst. Transient oscillations in nmr cross-polarization experiments in solids. *Phys. Rev. Lett.*, 32(25):1402–1406, Jun 1974. 83, 84
- [91] D. P. Weitekamp, J. R. Garbow, and A. Pines. Determination of dipole coupling constants using heteronuclear multiple quantum nmr. *The Journal of Chemical Physics*, 77(6):2870–2883, 1982. 85
- [92] P. Caravatti, L. Braunschweiler, and R. R. Ernst. Heteronuclear correlation spectroscopy in rotating solids. *Chemical Physics Letters*, 100(4):305–310, 1983. 85
- [93] Sergey V. Dvinskikh, Kazutoshi Yamamoto, and Ayyalusamy Ramamoorthy. Heteronuclear isotropic mixing separated local field nmr spectroscopy. *The Journal of Chemical Physics*, 125(3):034507, 2006. 85
- [94] N. R. Jagannathan, S. S. Rajan, and E. Subramanian. Refinement of the crystal structure of malonic acid, c3h4o4. *Journal of Chemical Crystallography*, 24(1):75–78, 1994. 86

- [95] J. Baugh, O. Moussa, C.A. Ryan, A. Nayak, and R. Laflamme. Experimental implementation of heat-bath algorithmic cooling using solid-state nuclear magnetic resonance. *Nature*, 438:470–473, November 2005. 86
- [96] R.C. McCalley and Alvin L. Kwiram. Endor studies at 4.2 k of the radicals in malonic acid single crystals. *Journal of Physical Chemistry*, 97:2888–2903, 1993. 87
- [97] S. F. Sagnowski, S. Aravamudhan, and U. Haeberlen. Wide-line and high-resolution proton magnetic resonance in single crystals of malonic acid. *Journal of Magnetic Resonance*, 28(2):271–288, 1977. 87
- [98] Carmen Müller, S. Idziak, N. Pislewski, and U. Haeberlen. Deuteron nmr on single crystals. efgs, signs of quadrupole coupling constants, and assignment of $[\sigma]$ tensors in malonic acid. *Journal of Magnetic Resonance (1969)*, 47(2):227–239, 1982. 87
- [99] J. Krzystek, Andrew B. Kwiram, and Alvin L. Kwiram. Ednmr and endor study of the beta-gamma phase transition in malonic acid crystals. *Journal of Physical Chemistry*, 99(1):402–409, 1995. 87
- [100] R. G. Delaplane, W. I. F. David, R. M. Ibberson, and C. C. Wilson. The ab initio crystal structure determination of $[\alpha]$ -malonic acid from neutron powder diffraction data. *Chemical Physics Letters*, 201(1-4):75–78, 1993. 87
- [101] Hiroshi Takeuchi, Kaori Inoue, Yuta Ando, and Shigehiro Konaka. Efficient method to analyze nmr spectra of solutes in liquid crystals: The use of genetic algorithm and integral curves. *Chemistry Letters*, 29(11):1300, 2000. 88
- [102] Jonathan Baugh, Osama Moussa, Colm A. Ryan, Raymond Laflamme, Chandrasekhar Ramanathan, Timothy F. Havel, and David G. Cory. Solid-state nmr three-qubit homonuclear system for quantum-information processing: Control and characterization. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 73(2):022305, 2006. 90
- [103] Eiichi Fukushima and Stephen Roeder. *Experimental Pulse NMR: A Nuts and Bolts Approach*. Westview Press, 1993. 94
- [104] J. Harrison, M. J. Sellars, and N. B. Manson. Measurement of the optically induced spin polarisation of n-v centres in diamond. *Diamond and Related Materials*, 15:586–588, 2006. 94
- [105] J. S. Hodges, J. C. Yang, C. Ramanathan, and D. G. Cory. Universal control of nuclear spins via anisotropic hyperfine interactions. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 78(1):010303, 2008. 96

- [106] J. S. Hodges. *Engineering Coherent Control of Quantum Information in Spin Systems*. PhD thesis, MIT, August 2007. 96
- [107] Jonathan Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. Available online at <http://www.cs.cmu.edu/quake-papers/painless-conjugate-gradient.pdf>, August 1994. 108