

Pseudo-randomness and Learning in Quantum Computation

Richard Andrew Low

A dissertation submitted to the University of Bristol in accordance with the requirements for award of the degree of Doctor of Philosophy in the Faculty of Engineering, Department of Computer Science, September 2009

40,000 Words

Abstract

This thesis discusses the young fields of quantum pseudo-randomness and quantum learning algorithms. We present techniques for derandomising algorithms to decrease randomness resource requirements and improve efficiency. One key object in doing this is a k -design, which is a distribution on the unitary group whose k^{th} moments match those of the unitarily invariant Haar measure. We show that for a natural model of a random quantum circuit, the distribution of random circuits quickly converges to a 2-design. We then present an efficient unitary k -design construction for any k , provided the number of qubits n satisfies $k = O(n/\log n)$. In doing this, we provide an efficient construction of a quantum tensor product expander, which is a generalisation of a quantum expander which in turn generalises classical expanders. We then discuss applications of k -designs. We show that they can be used to improve the efficiency of many existing algorithms and protocols and also find new applications to derandomising large deviation bounds. In particular, we show that many large deviation bound results for Haar random unitaries carry over to k -designs for $k = \text{poly}(n)$.

In the second part of the thesis, we present some learning and testing algorithms for the Clifford group. We find an optimal algorithm for identifying an unknown Clifford operation. We also give an algorithm to test if an unknown operation is close to a Clifford or far from every Clifford.

Acknowledgements

I am greatly indebted to my supervisor Aram Harrow for his teaching, support and inspiration throughout my PhD and for this I greatly thank him. I would also like to thank the whole quantum information group at Bristol and members of the Computer Science department for their support and help. In particular, I thank Mick Bremner, Raphaël Clifford, Toby Cubitt, Richard Jozsa, Will Matthews, Ashley Montanaro, Ben Sach, Dan Shepherd and Andreas Winter.

I also thank my family and friends for their encouragement and interest and members of the Hornstars, Dr Doctor, Cattle Market and Les Rosbifs for welcome distraction from this thesis.

I also acknowledge funding from the ARO grant ASTQIT and the EPSRC grant QIP-IRC.

Contents

1	Introduction	2
1.1	Brief Introduction to Quantum Mechanics	4
1.2	Preliminaries	6
1.2.1	Pauli Matrices	6
1.2.2	The Symmetric Group and Permutation Operators	6
1.2.3	Asymptotic Notation	8
1.2.4	Norms and Superoperator Norms	8
1.3	Previous Publications	10
I	Quantum Pseudo-randomness	12
2	Introduction to Quantum Pseudo-randomness	13
2.1	Random Unitaries	14
2.2	k -designs	16
2.2.1	k -wise Independence	16
2.2.2	Exact Designs	17
2.2.3	Approximate k -designs	21
3	Random Quantum Circuits	30
3.1	Introduction: Pseudo-random Quantum Circuits	30
3.1.1	Random Circuits	31
3.2	Preliminaries	33
3.2.1	Pauli expansion	33
3.2.2	Random Circuits as k -designs	33
3.3	Summary of Results	35
3.4	Analysis of the Moments	38
3.4.1	$k = 1$	42
3.4.2	$k = 2$	42
3.4.3	Moments for General Universal Random Circuits	44
3.5	Convergence	47
3.5.1	First Moments Convergence	50
3.5.2	Second Moments Convergence	53

3.5.3	Markov Chain Analysis	55
3.5.4	Convergence Proof	61
3.6	Tight Analysis for the $U(4)$ Case	66
3.6.1	Second Moments Convergence	66
3.6.2	Markov Chain of Coefficients	68
3.6.3	Proof of Theorem 3.6.4	72
3.6.4	Proof of Theorem 3.6.5	76
3.7	Main Result	82
3.8	Conclusions	84
3.9	Proofs	84
3.9.1	Zero chain mixing time proofs	84
3.9.2	Moment Generating Function Calculations	98
3.9.3	Mixing Times	99
4	Quantum Tensor Product Expanders and an Efficient Unitary Design Construction	101
4.1	Introduction	101
4.1.1	Quantum Expanders	102
4.1.2	Main Result	105
4.2	Proof of Theorem 4.1.5	106
4.2.1	Proof overview	106
4.2.2	Action of a Classical $2k$ -TPE	107
4.2.3	Fixed Points of a Quantum Expander	112
4.2.4	Fourier Transform in the Matrix Element Basis	113
4.2.5	Proof of Lemma 4.2.2	117
4.3	Conclusions	120
5	Applications of Designs	122
5.1	Review of Applications	122
5.1.1	Quantum Cryptography	122
5.1.2	Measurement	124
5.1.3	Average Gate Fidelity	125
5.1.4	Data Hiding	126
5.1.5	Decoupling and Evolution of Black Holes	127
5.1.6	Applications for Larger k	128
5.2	Derandomising Large Deviation Bounds	129
5.2.1	Introductory Problem: Entanglement of a 2-design	131
5.2.2	Main Results	132
5.2.3	Optimality of Results	134
5.3	Main Technique	135
5.3.1	Step 1: Concentration for uniform randomness	135
5.3.2	Step 2: A bound on the moments	136
5.3.3	Step 3: A concentration bound for a k -design	137
5.4	Application 1: Entropy of a k -design	139

5.5	Application 2: k -designs and Statistical Mechanics	142
5.6	Application 3: Using k -designs for Measurement-Based Quantum Computing	145
5.7	Conclusions	149
5.8	Proof of Theorem 5.2.3	150
II Quantum Learning		152
6	Learning and Testing Algorithms for the Clifford Group	153
6.1	Introduction	153
6.2	The Pauli and Clifford Groups and the Gottesman-Chuang Hierarchy	155
6.3	Learning Gottesman-Chuang Operations	156
6.3.1	Learning Pauli Operations	156
6.3.2	Learning Clifford Operations	156
6.3.3	Learning Gottesman-Chuang Operations	159
6.4	Learning Unitaries Close to \mathcal{C}_k Elements	160
6.5	Clifford Testing	162
6.6	Conclusions and Further Work	165
6.7	Proof of Lemma 6.4.4	165
6.8	Miscellaneous Lemmas	167
Bibliography		170

Chapter 1

Introduction

Landauer famously said that information is physical [Lan92]. A corollary of this is that computation is a physical process. It is simply the evolution of a physical state, governed by the laws of physics. A classical computer is therefore an information processor where the physical evolution is restricted to that of classical physics. A quantum computer is more general: quantum evolution is allowed. One might therefore reasonably expect that quantum computers are more powerful. It might be that the extra possibilities allowed by quantum evolution allow states to be processed more efficiently to speed up the computation. Determining which problems a quantum computer can solve faster than a classical computer is the central problem in the theory of quantum computation.

Significant progress has already been made in answering this question. Shor's algorithm [Sho94] shows that factoring of integers is possible in polynomial time on a quantum computer. In contrast, it is not known if factoring is possible in polynomial time on a classical computer. Also, Grover's unstructured search algorithm [Gro97] allows a marked item in an unsorted database to be found using only the square root of the time required on a classical computer. Finding other algorithms and provable separations between quantum and classical computation is an important area of current research.

This thesis makes some progress towards finding such new algorithms. In classical

computer science, randomness and pseudo-randomness have been key tools in the development of new and faster algorithms. In the first part of this thesis, we discuss applications and constructions of quantum analogues of these pseudo-random objects. Whilst we do not come up with new algorithms based on these, we hope that in the future the tools we build will find application in this area. In the second part of this thesis, we discuss problems in the theory of machine learning, which is an area in which quantum computers could outperform their classical counterparts.

A side theme in this thesis is the idea that computational complexity must be considered in physical models. The converse of our opening statement is also true: physical systems store and process information; they are computers. Therefore physical systems that could solve problems that are provably difficult do not exist in nature. This can rule out models that provide too much computational power.

In Part I, we discuss quantum pseudo-randomness. We introduce the subject in Chapter 2 and provide motivation from the classical computer science literature. The main idea is to use pseudo-randomness instead of full randomness to decrease the amount of randomness required. This is desirable in classical computing because random bits are expensive to produce. In quantum computing random bits can be obtained by measurement but uniformly random unitaries and states (formally defined in Section 2.1) cannot be produced efficiently so pseudo-randomness is necessary if efficiency is desired. In classical computing random bits are often saved by limiting dependence, for example by using k -wise independent random variables. These are variables where the distribution of any k variables is the same as for fully independent random variables but dependencies become apparent when observing more than k of the variables. We discuss a quantum analogue of this known as a k -design.

In Chapter 3 we show that short random quantum circuits (see Section 1.1 for background on quantum circuits) are 2-designs, giving an efficient method for producing a 2-design. Then in Chapter 4 we provide an efficient k -design construction for all k , giving the first construction for $k > 2$. In order to do this, we present an efficient construction of a quantum k -tensor product expander, which is a quantum

analogue of a classical tensor product expander which in turn is a generalisation of the standard expander used in classical computer science. We then summarise known applications of k -designs in Chapter 5 as well as providing our own to show that k -designs exhibit measure concentration which in some cases is almost as strong as for uniformly random unitaries.

In Part II we turn to problems in learning theory. In particular, we consider the problem of identifying a given black box unitary with as few queries as possible. We find an algorithm with optimal asymptotic query complexity to identify an unknown unitary from the Clifford group (defined in Chapter 6). We also show how this can be done if the unitary only approximately implements a Clifford and we also present a testing algorithm to determine if a given operation is close to a Clifford or far from every Clifford.

1.1 Brief Introduction to Quantum Mechanics

We now briefly mention some key concepts in quantum mechanics and define some notation. For a more complete introduction see [NC00].

The state of a d -dimensional quantum system is represented by a vector in the complex space \mathbb{C}^d . If $d = 2$, we call the system a *qubit* and often we will take $d = 2^n$ and say the system has n qubits.

We will normally use *Dirac notation* for quantum states. We write column vectors as $|\psi\rangle$, with the associated conjugate row vector as $\langle\psi|$. The inner product between two states is written as $\langle\psi_1|\psi_2\rangle$. We will write ψ for the projector $|\psi\rangle\langle\psi|$. States can also be probabilistic mixtures of pure states. If the state is $|\psi_i\rangle$ with probability p_i then it has *density matrix* $\sum_i p_i |\psi_i\rangle\langle\psi_i|$.

It is often convenient to break the space up into different components, for example the system and its environment. Mathematically, systems are combined by using the *tensor product*. The combined state of system A in state $|\psi_A\rangle$ and system B in state $|\psi_B\rangle$ is written $|\psi_A\rangle \otimes |\psi_B\rangle$. This leads to the phenomenon of *entanglement*, which is when the combined state cannot be written in this product form. For example, the

state $\frac{1}{\sqrt{2}}(|0_A\rangle \otimes |0_B\rangle + |1_A\rangle \otimes |1_B\rangle)$ cannot be written in the product form $|\psi_A\rangle \otimes |\psi_B\rangle$ and so is entangled. To find the state of a subsystem A from a density matrix ρ_{AB} , we take the *partial trace*. Write $\rho_{AB} = \sum_{ijkl} \rho_{ijkl} |i_A\rangle \langle j_A| \otimes |k_B\rangle \langle l_B|$. Then the reduced state is $\rho_A = \text{tr}_B \rho_{AB} = \sum_{ijk} \rho_{ijkk} |i_A\rangle \langle j_A|$.

Measurement of a quantum system can be written mathematically in terms of a POVM (positive operator valued measure). This is a set of positive semi-definite operators P_i such that $\sum_i P_i = I$. Then the measurement outcomes are the labels i and outcome i occurs with probability $\text{tr} P_i \rho$ if the state being measured is ρ .

The evolution of a closed quantum system is unitary. That is, the quantum state at a later time t is related by a unitary to the initial quantum state: $\rho_t = U_t \rho_0 U_t^\dagger$. If the system of interest is part of some larger system then the dynamics need not be unitary. The most general form of evolution can be written in the Kraus decomposition: $\rho_t = \sum_i A_i \rho_0 A_i^\dagger$ where A_i are any operators normalised so that $\sum_i A_i^\dagger A_i = I$.

Finally we mention that it is often convenient to think of unitary evolution as a quantum circuit, built up of smaller elementary unitary gates. This is in direct analogy to the use of circuits in classical computing. Classically, a NAND gate suffices to produce any other gate so all classical circuits can be made up of just NAND gates. Similarly, there exist sets of unitary gates from which any unitary can be built. An example is the following three gates:

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad R_{\pi/4} = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$$

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

We often seek to build a family of circuits that act on n qubits for all n , where the gates are chosen from some elementary set, such as that above. We say that the

circuits are efficient if the number of gates grows only polynomially with n .

1.2 Preliminaries

Here we define some notation and concepts that are used throughout this thesis.

1.2.1 Pauli Matrices

We will often use the Pauli matrices:

$$\begin{aligned} \sigma_0 = \sigma_I &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} & \sigma_1 = \sigma_x &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \\ \sigma_2 = \sigma_y &= \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} & \sigma_3 = \sigma_z &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \end{aligned} \quad (1.2.1)$$

We can extend these to matrices on n qubits by taking tensor products. Let $p \in \{0, 1, 2, 3\}^n$ and $\sigma_p = \sigma_{p_1} \otimes \sigma_{p_2} \otimes \cdots \otimes \sigma_{p_n}$ where p_i is the value at the i^{th} position in the string p . We will sometimes use the alternative notation of $p \in \{I, x, y, z\}^n$. We will refer to σ_p as Pauli matrices on n qubits. There are 4^n Pauli matrices and they are orthogonal i.e. $\text{tr } \sigma_p \sigma_q = 2^n \delta_{pq}$. Note also that $\sigma_p^2 = \sigma_0$, the identity. Also, Pauli matrices either commute or anticommute.

The Pauli matrices form an orthogonal basis for matrices in $\mathbb{C}^{2^n \times 2^n}$. Therefore any such matrix A can be written in the form $\sum_p \gamma(p) \sigma_p$, with $\gamma(p) = \frac{1}{2^n} \text{tr } \sigma_p A$. Sometimes we will choose a different normalisation for the Pauli coefficients $\gamma(p)$ but will make this clear from the context.

1.2.2 The Symmetric Group and Permutation Operators

The symmetric group is the group of all permutations. Let \mathcal{S}_N be the symmetric group on N objects. Then for $\pi \in \mathcal{S}_N$ define the corresponding permutation operator

$$B(\pi) := \sum_{i=1}^N |\pi(i)\rangle \langle i| \quad (1.2.2)$$

to be the matrix that permutes the basis states $|1\rangle, \dots, |N\rangle$ according to π .

On the other hand, if we have k N -dimensional systems then for $\pi \in \mathcal{S}_k$ define the subsystem permutation operator $S(\pi)$ by

$$S(\pi) := \sum_{n_1=1}^N \cdots \sum_{n_k=1}^N |n_{\pi^{-1}(1)}, \dots, n_{\pi^{-1}(k)}\rangle \langle n_1, \dots, n_k|. \quad (1.2.3)$$

Now we present two useful lemmas about subsystem permutation operators.

Lemma 1.2.1. *Let C be a cycle of length c in \mathcal{S}_c . Then*

$$\text{tr}(C(A_1 \otimes A_2 \otimes \dots \otimes A_c)) = \text{tr}(A_{C(1)} A_{C^{\circ 2}(1)} A_{C^{\circ 3}(1)} \dots A_1).$$

Proof. We have

$$\begin{aligned} \text{tr}(C(A_1 \otimes A_2 \otimes \dots \otimes A_c)) &= \sum_{i_1, i_2, \dots, i_c} \langle i_1 i_2 \dots i_c | C(A_1 \otimes A_2 \otimes \dots \otimes A_c) | i_1 i_2 \dots i_c \rangle \\ &= \sum_{i_1, i_2, \dots, i_c} \langle i_1 | A_{C(1)} | i_{C(1)} \rangle \langle i_2 | A_{C(2)} | i_{C(2)} \rangle \dots \langle i_c | A_{C(c)} | i_{C(c)} \rangle \\ &= \sum_{i_1, i_2, \dots, i_c} \langle i_1 | A_{C(1)} | i_{C(1)} \rangle \langle i_{C(1)} | A_{C^{\circ 2}(1)} | i_{C^{\circ 2}(1)} \rangle \dots \langle i_{C^{\circ c-1}(1)} | A_1 | i_1 \rangle \end{aligned}$$

since $C^{\circ c}(1) = 1$. Evaluate the sum using the resolution of the identity to get the result. \square

A simple example of this Lemma is that

$$\text{tr}(\mathcal{F}(A \otimes B)) = \text{tr} AB \quad (1.2.4)$$

where \mathcal{F} is the swap operator.

We also work out the Pauli expansion of the swap operator. To stress that this result does not depend on the choice of orthogonal basis we prove it in full generality.

Lemma 1.2.2. *The swap operator \mathcal{F} on two d -dimensional systems can be written as*

$$\frac{1}{d} \sum_p \sigma_p \otimes \sigma_p.$$

where $\{\sigma_p\}$ form a Hermitian orthogonal basis with $\text{tr} \sigma_p^2 = d$.

Proof. Expand \mathcal{F} in the basis and use Lemma 1.2.1:

$$\begin{aligned} \text{tr}((\sigma_p \otimes \sigma_q) \mathcal{F}) &= \text{tr} \sigma_p \sigma_q \\ &= \begin{cases} d & p = q \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

The given sum has the correct coefficients in the basis therefore $\frac{1}{d} \sum_p \sigma_p \otimes \sigma_p = \mathcal{F}$. \square

1.2.3 Asymptotic Notation

We will use the following standard asymptotic notation.

Definition 1.2.3. $f(n) = O(g(n))$ if there exists $c, n_0 > 0$ such that $0 \leq f(n) \leq cg(n)$ for all $n \geq n_0$.

Definition 1.2.4. $f(n) = \Omega(g(n))$ if there exists $c, n_0 > 0$ such that $f(n) \geq cg(n) \geq 0$ for all $n \geq n_0$.

Definition 1.2.5. $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$.

Definition 1.2.6. $f(n) = o(g(n))$ if $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$.

1.2.4 Norms and Superoperator Norms

Norms

We will make heavy use of Schatten p -norms:

Definition 1.2.7. For A a $d \times d$ matrix, the Schatten p -norm is given by

$$\|A\|_p = \left(\sum_{i=1}^d \sigma_i^p \right)^{1/p} \quad (1.2.5)$$

where σ_i are the singular values of A .

In particular, $\|A\|_1 = \sum_{i=1}^d \sigma_i = \text{tr} \sqrt{A^\dagger A}$, $\|A\|_2 = \sqrt{\sum_{i=1}^d \sigma_i^2} = \sqrt{\text{tr} A^\dagger A}$ and $\|A\|_\infty = \max_i \sigma_i$.

These norms satisfy the following simple relationships:

$$\|A\|_2 \leq \|A\|_1 \leq \sqrt{d} \|A\|_2 \quad (1.2.6)$$

$$\|A\|_\infty \leq \|A\|_1 \leq d \|A\|_\infty \quad (1.2.7)$$

$$\|A\|_\infty \leq \|A\|_2 \leq \sqrt{d} \|A\|_\infty. \quad (1.2.8)$$

Superoperator Norms

Just as state norms can be used to bound the distinguishability of states, superoperator norms bound how easy it is to tell different superoperators apart. We start with the 1-norm:

Definition 1.2.8. The 1-norm of a superoperator \mathcal{E} is given by

$$\|\mathcal{E}\|_{1 \rightarrow 1} = \sup_{X \neq 0} \frac{\|\mathcal{E}(X)\|_1}{\|X\|_1}.$$

The main problem with this definition is that the 1-norm is not stable under tensoring with the identity i.e. there exist channels with $\|\mathcal{E} \otimes \text{id}_d\|_{1 \rightarrow 1} > \|\mathcal{E}\|_{1 \rightarrow 1}$, where id_d is the identity channel on d dimensions. This means that some channels are easier to distinguish by inputting entangled states. If the norm is to measure the distinguishability of channels it should take this into account. To overcome this problem, the diamond norm is defined:

Definition 1.2.9 ([KSV02]). *The diamond norm of a superoperator \mathcal{E} is given by*

$$\|\mathcal{E}\|_{\diamond} = \sup_d \|\mathcal{E} \otimes \text{id}_d\|_{1 \rightarrow 1} = \sup_d \sup_{X \neq 0} \frac{\|(\mathcal{E} \otimes \text{id}_d)X\|_1}{\|X\|_1}.$$

It follows immediately that $\|\mathcal{E}\|_{1 \rightarrow 1} \leq \|\mathcal{E}\|_{\diamond}$. Also it is shown in [KSV02] that the diamond norm satisfies $\|\mathcal{E} \otimes \text{id}_d\|_{\diamond} = \|\mathcal{E}\|_{\diamond}$ for all channels \mathcal{E} and dimensions d and that the dimension d in the supremum can be taken to be the same as the dimension of the system \mathcal{E} acts on. Operationally, the diamond norm of the difference between two quantum operations tells us the largest possible probability of distinguishing the two operations if we are allowed to have them act on part of an arbitrary, possibly entangled, state.

We will also use the 2-norm:

Definition 1.2.10. *The 2-norm of a superoperator \mathcal{E} is given by*

$$\|\mathcal{E}\|_{2 \rightarrow 2} = \sup_{X \neq 0} \frac{\|\mathcal{E}(X)\|_2}{\|X\|_2}.$$

In [van02] Appendix C, the following relationships between the superoperator norms are proven:

$$\|\mathcal{E}\|_{2 \rightarrow 2} \leq \sqrt{d} \|\mathcal{E}\|_{1 \rightarrow 1} \tag{1.2.9}$$

$$\|\mathcal{E}\|_{1 \rightarrow 1} \leq \sqrt{d} \|\mathcal{E}\|_{2 \rightarrow 2} \tag{1.2.10}$$

$$\|\mathcal{E}\|_{\diamond} \leq d \|\mathcal{E}\|_{1 \rightarrow 1} \tag{1.2.11}$$

$$\|\mathcal{E}\|_{\diamond} \leq d \|\mathcal{E}\|_{2 \rightarrow 2}. \tag{1.2.12}$$

1.3 Previous Publications

The majority of this thesis has been published previously and some is work in collaboration.

Chapter 3 is joint work with Aram Harrow and is available as “Random Quantum Circuits are Approximate 2-designs”, Communications in Mathematical Physics, Vol-

ume 291, Number 1, Pages 257-302. It is also available as a pre-print: arXiv:0802.1919.

Chapter 4 is also joint work with Aram Harrow and is available as “Efficient Quantum Tensor Product Expanders and k -Designs”, Proceedings of RANDOM 2009, LNCS, Volume 5687, Pages 548-561. It is also available as a pre-print: arXiv:0811.2597.

Chapter 5 from Section 5.2 onwards is available as “Large deviation bounds for k -designs”, Proceedings of the Royal Society A, Volume 465, Number 2111, Pages 3289-3308. It is also available as a pre-print: arXiv:0903.5236.

Chapter 6 is available as “Learning and Testing Algorithms for the Clifford Group”, Physical Review A, Volume 80, Number 5, Page 052314. It is also available as a pre-print: arXiv:0907.2833.

Part I

Quantum Pseudo-randomness

Chapter 2

Introduction to Quantum Pseudo-randomness

Randomness is an important resource in both classical and quantum computing. It has applications in virtually all areas of computer science, including algorithms, cryptography and networking. Randomness can improve efficiency or, as in the case of cryptography, allow us to perform tasks that we would not be able to do with deterministic resources.

An example of an algorithm where a randomised algorithm is faster than any known deterministic algorithm is polynomial identity testing. Here, the task is to determine if two polynomials are identically equal. By evaluating the polynomials on random inputs, identity testing can be done in polynomial time whereas no polynomial time deterministic algorithm is known.

Also, a commonly used randomised algorithm is that of randomised quicksort. In quicksort, a pivot element is chosen and elements smaller than this are placed to the left and larger elements to the right. Then these two parts are sorted recursively. However, the choice of pivot element greatly affects the run-time of the algorithm. If chosen poorly (for example so that there is only one element smaller than the pivot), the algorithm runs in $O(n^2)$ time. If chosen well, the algorithm runs in $O(n \log n)$ time. Choosing the pivot element randomly will be a good choice on average, giving expected

run-time $O(n \log n)$ [MR95]. However, this run-time can be achieved deterministically using deterministic median finding [BFP⁺72] but in practice the randomised method is more efficient.

As another example, many primality testing algorithms are randomised because of their simplicity, even though a deterministic polynomial-time algorithm is now known. Also, in the field of communication complexity, separations between deterministic and randomised algorithms can be proven. The deterministic complexity of evaluating the equality function (to determine if Alice and Bob's strings are equal) is $\Theta(n)$, whereas the randomised complexity is $\Theta(\log n)$ [KN96]. As yet another example of randomness in classical computer science, in networking a random delay is often inserted after a collision so the nodes wait different times so are likely to avoid another collision.

In this part, we seek to extend some of these gains of using randomness to quantum computing. We wish to find applications of randomness to find new quantum algorithms and constructions.

Besides the computer science applications, there are also physical reasons for studying randomness in quantum mechanics. Some systems can be modelled as interacting randomly and it is interesting to ask what the limiting state (or distribution on states) is for such a system. Also of great interest is how quickly the system reaches this stationary state. If the time taken grows too quickly with the size of the system (for example, exponentially) then for any system apart from the most trivial, the stationary state will never be reached and will not be seen in physical systems. However, if the time is small (for example, a small polynomial), then the stationary state can be reached quickly and will be observed in real systems. It is in problems like this that physicists must consider the computer science aspects of their models. We study problems of this kind in Chapters 3 and 5.

2.1 Random Unitaries

In quantum computing, operations are unitary gates and randomness is often used in the form of random unitary operations. Random unitaries have algorithmic uses

(e.g. [Sen05]), cryptographic applications (e.g. [AS04, HLSW04]) and applications to fundamental quantum protocols (e.g. [BHL⁺05, HHL04]). For information-theoretic applications, it is often convenient to use unitary matrices drawn from the uniform distribution on the unitary group, also known as the Haar measure. This measure is the unique unitarily invariant measure i.e. the only measure dU on the unitary group $\mathcal{U}(d)$ where $\int_{\mathcal{U}(d)} f(U)dU = \int_{\mathcal{U}(d)} f(UV)dU$ for all functions f and unitaries V . For random states, we write the unitarily invariant measure on d -dimensional states as $d\psi$. This can be thought of as a Haar distributed unitary applied to any fixed pure state. It is also known as the Fubini-Study metric.

However, in both classical and quantum computing, obtaining random bits is often expensive, and so it is often desirable to minimise their use. For example, in classical computing, expanders (discussed in Chapter 4) and k -wise independent functions (see Section 2.2) have been developed for this purpose and have found wide application. We will spend a great deal of time exploring quantum analogues of these: quantum expanders and k -designs.

In addition to randomness being expensive, there is an even more pressing problem when using random unitaries and states. An n -qubit unitary is defined by 4^n real parameters, and so cannot even be approximated efficiently using a subexponential amount of time or randomness. So any application that requires a random unitary cannot be efficient. Instead, we will seek to construct efficient pseudo-random ensembles of unitaries which resemble the Haar measure for certain applications. For example, a k -design (often referred to as a t -design, or a (k, k) -design), as mentioned above, is a distribution on unitaries which matches the first k moments of the Haar distribution. k -designs have found many uses which are explored in Chapter 5.

In Section 2.2, we formally define k -designs and summarise known constructions. Then in Chapter 3 we show that, for a natural model of a random quantum circuit, the distribution quickly converges to that of a 2-design. This gives an efficient approximate 2-design construction and also has physical applications. In Chapter 4, we provide an efficient construction of a unitary k -design for any k (although there are restrictions

on the dimension, see later). Then in Chapter 5, we discuss applications of designs, including to derandomising constructions that use large deviation bounds.

Parts of this chapter have been published previously in [HL09b, HL09a, Low09a] and parts are joint work with Aram Harrow.

2.2 k -designs

A unitary k -design is a distribution of unitaries that gives the same expectations of polynomials of degree at most k as the Haar measure. This is just like Gaussian quadrature, where integrals of polynomials are calculated by sums. Gaussian quadrature says that there exist sample points $\{x_i\}$ and weights $\{w_i\}$ so that for all polynomials f of degree at most $2b - 1$,

$$\sum_{i=1}^b w_i f(x_i) = \int_p^q dx f(x) \quad (2.2.1)$$

for some fixed limits p and q . This allows the integrals to be calculated much more efficiently. A unitary k -design is the same, except the polynomial is on elements of unitary matrices from the unitary group rather than numbers on the real line. The k refers to the degree of the polynomial. We will also discuss state designs, where the function is on coefficients of states rather than unitaries.

2.2.1 k -wise Independence

k -designs can also be thought of as a quantum analogue of k -wise independence. A sequence of random variables X_1, \dots, X_n is k -wise independent if, for any subset of size $j \leq k$,

$$\mathbb{P}(X_{i_1} = x_{i_1}, \dots, X_{i_j} = x_{i_j}) = \mathbb{P}(X_{i_1} = x_{i_1}) \dots \mathbb{P}(X_{i_j} = x_{i_j}). \quad (2.2.2)$$

As a simple example of how this can save randomness, consider the set

$$\{000, 011, 101, 110\}. \tag{2.2.3}$$

If an element is chosen uniformly at random from this set, the probability distribution of the values of any two bits is the same as if all three bits were chosen independently. This is therefore a 2-wise independent set, and saves one bit of randomness. In general, if $k \ll n$, an exponential saving in randomness can be made in this way. Efficient constructions of exactly k -wise independent sets are known [ABI86] and more efficient approximate constructions are given in [NN90].

A related concept is that of k -wise independent permutations. These are sets of permutations with the property that, when a permutation is chosen randomly from this set and applied to n points, the distribution of the positions of any k points is the same as if a uniformly random permutation was applied. For example, a random cyclic shift is a 1-wise independent permutation. Again, an exponential saving of randomness is possible [KNR09].

We seek to construct quantum k -designs to achieve a similar saving of randomness for quantum algorithms. We now formally define k -designs.

2.2.2 Exact Designs

We will use the following notation to distinguish the measure we are using. Write \mathbb{E} for the expectation with $\mathbb{E}_{U \sim \nu}$ meaning the expectation when U is chosen from the measure ν . If the measure is the Haar measure in dimension d we will write $\mathbb{E}_{U \sim \mathcal{U}(d)}$. We use the same subscripts for probabilities so $\mathbb{P}_{U \sim \mathcal{U}(d)}$ denotes the probability when U is chosen from the Haar measure, etc.. When considering random states, we will write $\mathbb{E}_{|\psi\rangle \sim \mathcal{S}(d)}$, etc..

State designs

A k -design is an ensemble of states such that, when one state is chosen from the ensemble and copied k times, it is indistinguishable from a uniformly random state. The state k -design definition we use is due to Ambainis and Emerson [AE07]:

Definition 2.2.1 ([AE07], Definition 1). *An ensemble of quantum states $\nu = \{p_i, |\psi_i\rangle\}$ is a state k -design if*

$$\mathbb{E}_{|\psi\rangle\sim\nu} [(|\psi\rangle\langle\psi|)^{\otimes k}] = \mathbb{E}_{|\psi\rangle\sim\mathcal{S}(d)} [(|\psi\rangle\langle\psi|)^{\otimes k}] \quad (2.2.4)$$

We can evaluate the integral on the right hand side:

Lemma 2.2.2.

$$\int_{\psi} (|\psi\rangle\langle\psi|)^{\otimes k} d\psi = \frac{\Pi_{+k}}{\binom{k+d-1}{k}} \quad (2.2.5)$$

where Π_{+k} is the projector onto the symmetric subspace of k d -dimensional spaces.

Proof. The standard proof (see e.g. [GW98] or [BBD⁺97]) involves showing that $\int_{\psi} (|\psi\rangle\langle\psi|)^{\otimes k} d\psi$ commutes with all elements of an irreducible representation (irrep) of the unitary group that acts on the symmetric subspace so by Schur's lemma must be proportional to the projector onto the symmetric subspace. However, here we give an alternative proof that introduces a technique we will use later.

By the unitary invariance of the Haar measure, $\int_{\psi} (|\psi\rangle\langle\psi|)^{\otimes k} d\psi$ commutes with $U^{\otimes k}$ for all unitaries U . By Schur-Weyl duality (see e.g. [GW98]), this implies that the integral is a linear combination of subsystem permutation operators. Therefore we have

$$\int_{\psi} (|\psi\rangle\langle\psi|)^{\otimes k} d\psi = \sum_{\pi \in S_k} \alpha_{\pi} S(\pi). \quad (2.2.6)$$

However, the integral is invariant under permutations so α_{π} must be the same for all permutations π . Using $\Pi_{+k} = \frac{1}{k!} \sum_{\pi \in S_k} S(\pi)$ and finding the normalisation by taking the trace (the dimension of the symmetric subspace is $\binom{k+d-1}{k}$) proves the result. \square

We will also state equivalent definitions of designs in terms of polynomials of

matrix elements of the unitary or coefficients of the state. First we must define what we mean by the degree of a polynomial:

Definition 2.2.3. *A monomial in elements of a matrix U or state $|\psi\rangle$ is of degree (k_1, k_2) if it contains k_1 conjugated elements and k_2 unconjugated elements. We call it balanced if $k_1 = k_2$ and will simply say a balanced monomial has degree k if it is degree (k, k) . A balanced polynomial is of degree k if it is a sum of balanced monomials of degree at most k , with at least one monomial with degree equal to k .*

So that, in this definition, $U_{pq}U_{rs}^*$ is a balanced monomial of degree $(1, 1)$ and $U_{pq}U_{rs}$ is a monomial of degree $(2, 0)$ and is unbalanced. For the state $|\psi\rangle = \sum_i \alpha_i |i\rangle$, $\alpha_i \alpha_j^*$ is a balanced monomial of degree $(1, 1)$.

We can then define state k -designs in terms of monomials:

Definition 2.2.4 ([AE07], Definition 3). *An ensemble of quantum states ν is a state k -design if, for all balanced monomials M of degree at most k ,*

$$\mathbb{E}_{|\psi\rangle \sim \nu} M(|\psi\rangle) = \mathbb{E}_{|\psi\rangle \sim \mathcal{S}(d)} M(|\psi\rangle) \quad (2.2.7)$$

This is an equivalent definition to Definition 2.2.1:

Lemma 2.2.5 ([AE07], Theorem 5). *The state design definitions 2.2.1 and 2.2.4 are equivalent.*

Proof. Firstly, we only need to prove the result for M of degree exactly k , since by partial tracing this implies the result for any smaller k .

Each entry in the matrix $\mathbb{E}_{|\psi\rangle \sim \nu} \left[(|\psi\rangle\langle\psi|)^{\otimes k} \right]$ is the expectation of a monomial of degree k , with the state chosen from the design. Further, the corresponding entry in $\mathbb{E}_{|\psi\rangle \sim \mathcal{S}(d)} \left[(|\psi\rangle\langle\psi|)^{\otimes k} \right]$ is the expectation of the same monomial but with the state chosen from the Haar measure. If the ensemble of states satisfies Definition 2.2.4 then these are equal, so the ensemble also satisfies Definition 2.2.1.

On the other hand, for every balanced monomial of degree k , there is an entry in $\mathbb{E}_{|\psi\rangle \sim \nu} \left[(|\psi\rangle\langle\psi|)^{\otimes k} \right]$ equal to its expectation. Therefore, if the ensemble of states satisfies Definition 2.2.1 then it also satisfies Definition 2.2.4. \square

Unitary designs

Consider having k d -dimensional systems in any initial state. A unitary k -design is an ensemble of unitaries such that when a unitary is randomly selected from it and applied to each of the k systems, the overall state is indistinguishable from choosing a uniformly random unitary. This can be seen as a generalisation of state designs in that any column of a unitary k -design is a state k -design. Formally, we have:

Definition 2.2.6. *Let ν be an ensemble of unitary operators. Define*

$$\mathcal{G}_\nu(\rho) = \mathbb{E}_{U \sim \nu} \left[U^{\otimes k} \rho (U^\dagger)^{\otimes k} \right] \quad (2.2.8)$$

and

$$\mathcal{G}_H(\rho) = \mathbb{E}_{U \sim \mathcal{U}(d)} \left[U^{\otimes k} \rho (U^\dagger)^{\otimes k} \right] \quad (2.2.9)$$

Then the ensemble is a unitary k -design if $\mathcal{G}_\nu(\rho) = \mathcal{G}_H(\rho)$ for all $d^k \times d^k$ matrices ρ (not necessarily physical states).

For convenience we have defined this for all matrices ρ although it is equivalent to only require equality for physical states, since all matrices can be obtained from linear combinations of physical states.

Like state designs, unitary designs can also be defined in terms of polynomials:

Definition 2.2.7 ([DCEL06]). *ν is a unitary k -design if, for all balanced monomials M of degree k ,*

$$\mathbb{E}_{U \sim \nu} M(U) = \mathbb{E}_{U \sim \mathcal{U}(d)} M(U). \quad (2.2.10)$$

Again, these definitions are equivalent:

Lemma 2.2.8. *The unitary design definitions 2.2.6 and 2.2.7 are equivalent.*

Proof. The proof is very similar to the state design case. Again, we only consider monomials of degree k since by partial tracing this implies the result for smaller k .

Consider matrices ρ of the form $|i_1, i_2, \dots, i_k\rangle\langle j_1, j_2, \dots, j_k|$ in Definition 2.2.6. Then each element of $U^{\otimes k} \rho (U^\dagger)^{\otimes k}$ is a balanced monomial of degree k and, for some

choice of indices in $|i_1, i_2, \dots, i_k\rangle\langle j_1, j_2, \dots, j_k|$, each balanced monomial of degree k appears. \square

2.2.3 Approximate k -designs

While exact designs have desirable properties, it is often much easier to construct approximate designs which, for many applications, are sufficient. Also, approximate designs can have fewer unitaries than exact designs. For example, it was shown in [AMTd00] that 2^{2n} unitaries are necessary and sufficient for an exact unitary 1-design. However, an approximate 1-design can be implemented with only $2^{n+o(n)}$ unitaries which gives almost a factor of 2 saving in random bits.

Approximate state designs

Our approximate state design definition is as follows:

Definition 2.2.9. ν is an ϵ -approximate state k -design if

$$\left\| \mathbb{E}_{|\psi\rangle\sim\nu} \left[(|\psi\rangle\langle\psi|)^{\otimes k} \right] - \mathbb{E}_{|\psi\rangle\sim\mathcal{S}(d)} \left[(|\psi\rangle\langle\psi|)^{\otimes k} \right] \right\|_{\infty} \leq \frac{\epsilon}{\binom{k+d-1}{k}}. \quad (2.2.11)$$

$\binom{k+d-1}{k}$ appears because it is the dimension of the symmetric subspace. In [AE07], a similar definition was proposed but with the additional requirement that the ensemble also forms a 1-design (exactly), i.e.

$$\mathbb{E}_{|\psi\rangle\sim\nu} |\psi\rangle\langle\psi| = \mathbb{E}_{|\psi\rangle\sim\mathcal{S}(d)} |\psi\rangle\langle\psi|$$

This requirement was necessary there only so that a suitably normalised version of the ensemble would form a POVM. We will not use it.

By taking the partial trace one can show that a k -design is a k' -design for $k' \leq k$. Thus approximate k -designs are always at least approximate 1-designs.

Approximate unitary designs

We have many choices to make when defining an approximate design. Here we give four definitions which are convenient in different contexts. In Lemma 2.2.14 we show that they are all equivalent, up to polynomial dimension factors.

If the unitary design is considered a quantum channel that applies a random unitary from the distribution to the input, then a relevant measure is the diamond norm difference between the approximate design and an exact design. Because the diamond norm is related to the distinguishability of channels, having a low diamond norm distance means that it is difficult to detect that an approximate design was given rather than exact. One approximate design definition is therefore:

Definition 2.2.10 (DIAMOND¹, See Chapter 3). ν is an ϵ -approximate unitary k -design if

$$\|\mathcal{G}_\nu - \mathcal{G}_H\|_\diamond \leq \epsilon, \quad (2.2.12)$$

where \mathcal{G}_ν and \mathcal{G}_H are defined in Definition 2.2.6.

In [DCEL06], they consider approximate twirling, which is implemented using an approximate 2-design. They give an alternative definition of closeness which is more convenient for this application:

Definition 2.2.11 (TWIRL, [DCEL06]). ν is an ϵ -approximate twirl if

$$\max_{\Lambda} \left\| \left| \mathbb{E}_{U \sim \nu} U^\dagger \Lambda(U \rho U^\dagger) U - \mathbb{E}_{U \sim \mathcal{U}(d)} U^\dagger \Lambda(U \rho U^\dagger) U \right| \right\|_\diamond \leq \frac{\epsilon}{d^2}. \quad (2.2.13)$$

The maximisation is over channels Λ and d is the dimension.

In Chapter 4, unitary designs are constructed from quantum tensor product expanders. A quantum k -TPE is defined as an ensemble ν of unitaries such that

$$\left\| \mathbb{E}_{U \sim \nu} \left[U^{\otimes k, k} \right] - \mathbb{E}_{U \sim \mathcal{U}(d)} \left[U^{\otimes k, k} \right] \right\|_\infty \leq \lambda \quad (2.2.14)$$

¹We name the definitions to help distinguish them

for $\lambda < 1$ and $U^{\otimes k, k} = U^{\otimes k} \otimes (U^*)^{\otimes k}$ (the motivation for this definition is explained in Chapter 4). From this a natural k -design definition follows:

Definition 2.2.12 (TRACE, See Chapter 4). ν is an ϵ -approximate unitary k -design if

$$\left\| \mathbb{E}_{U \sim \nu} [U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(d)} [U^{\otimes k, k}] \right\|_1 \leq \epsilon. \quad (2.2.15)$$

In Theorem 4.1.3 we prove the simple result that a unitary design can be constructed by iterating the TPE.

We will also need a definition in terms of monomials:

Definition 2.2.13 (MONOMIAL, See Chapter 5). ν is an ϵ -approximate unitary k -design if, for all balanced monomials M of degree $\leq k$,

$$|\mathbb{E}_{U \sim \nu} M(U) - \mathbb{E}_{U \sim \mathcal{U}(d)} M(U)| \leq \frac{\epsilon}{d^k} \quad (2.2.16)$$

We would now like to show that all these definitions are equivalent. By equivalent, we mean that, if ν is an ϵ -approximate unitary design by one definition, then it is an ϵ' -approximate unitary design by any other definition, where $\epsilon' = \text{poly}(d^k)\epsilon$.

Lemma 2.2.14. *Definitions 2.2.10 (DIAMOND), 2.2.12 (TRACE) and 2.2.13 (MONOMIAL) are all equivalent. Also Definition 2.2.11 (TWIRL) is equivalent to the other definitions for an approximate 2-design only.*

Proof. To prove this, we will consider yet another possible definition (OPERATOR-2-NORM):

$$\|\mathcal{G}_\nu - \mathcal{G}_H\|_{2 \rightarrow 2} \leq \epsilon. \quad (2.2.17)$$

Note that this is equivalent to

$$\left\| \mathbb{E}_{U \sim \nu} [U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(d)} [U^{\otimes k, k}] \right\|_\infty \leq \epsilon \quad (2.2.18)$$

which is the same as Definition 2.2.12 (TRACE) except the norm is the ∞ -norm rather than the 1-norm. We shall prove that the other k -design definitions are equivalent

to this. We then show that Definition 2.2.11 (TWIRL) is equivalent to Definition 2.2.13 (MONOMIAL) for $k = 2$. We use notation $A \xrightarrow{s} B$ to mean that if ν is an ϵ -approximate unitary k -design according to definition A then it is a $s\epsilon$ -approximate unitary k -design according to definition B. If $s = 1$ we omit the superscript.

A diagram showing the different parts to the proof is given in Figure 2.1. We remark that direction 2 is unneeded but is included since it provides tighter bounds and has a simple proof.

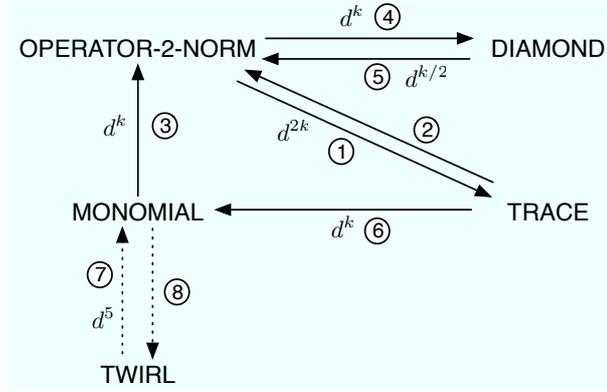


Figure 2.1: A diagram showing the different parts of the proof of Lemma 2.2.14. The dotted arrows show the correspondence is only for $k = 2$. The circled digits refer to the enumerated items below and the factors by the arrows indicate the precision lost in the approximation when converting between the definitions.

1. OPERATOR-2-NORM $\xrightarrow{d^{2k}}$ TRACE:

Use the equivalence between Equations 2.2.17 and 2.2.18 and that

$$\left\| \mathbb{E}_{U \sim \nu}[U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(N)}[U^{\otimes k, k}] \right\|_1 \leq d^{2k} \left\| \mathbb{E}_{U \sim \nu}[U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(d)}[U^{\otimes k, k}] \right\|_\infty.$$

2. TRACE \rightarrow OPERATOR-2-NORM:

Use

$$\left\| \mathbb{E}_{U \sim \nu}[U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(d)}[U^{\otimes k, k}] \right\|_\infty \leq \left\| \mathbb{E}_{U \sim \nu}[U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(d)}[U^{\otimes k, k}] \right\|_1$$

and the equivalence between Equations 2.2.17 and 2.2.18.

3. MONOMIAL $\xrightarrow{d^k}$ OPERATOR-2-NORM:

Choose any $\rho \in \mathbb{C}^{d^k \times d^k}$ and write it as $\rho = \sum_{ij} \rho_{ij} |i\rangle\langle j|$. Then

$$\begin{aligned} \|\mathcal{G}_\nu(\rho) - \mathcal{G}_H(\rho)\|_2 &\leq \sum_{ij} |\rho_{ij}| \|\mathcal{G}_\nu(|i\rangle\langle j|) - \mathcal{G}_H(|i\rangle\langle j|)\|_2 \\ &= \sum_{ij} |\rho_{ij}| \sqrt{\sum_{kl} |(\mathcal{G}_\nu(|i\rangle\langle j|) - \mathcal{G}_H(|i\rangle\langle j|))_{kl}|^2} \end{aligned}$$

using the fact that the 2-norm squared is the sum of the squares of the matrix elements. Now, we have a bound on the matrix elements of $\mathcal{G}_\nu(|i\rangle\langle j|) - \mathcal{G}_H(|i\rangle\langle j|)$ from Definition 2.2.13 (MONOMIAL):

$$|(\mathcal{G}_\nu(|i\rangle\langle j|) - \mathcal{G}_H(|i\rangle\langle j|))_{kl}| \leq \frac{\epsilon}{d^k}$$

so

$$\begin{aligned} \|\mathcal{G}_\nu(\rho) - \mathcal{G}_H(\rho)\|_2 &\leq \epsilon \sum_{ij} |\rho_{ij}| \\ &\leq d^k \epsilon \|\rho\|_2. \end{aligned}$$

4. OPERATOR-2-NORM $\xrightarrow{d^k}$ DIAMOND:

This follows from the superoperator norm relationship given in Eqn. 1.2.12.

5. DIAMOND $\xrightarrow{d^{k/2}}$ OPERATOR-2-NORM:

This uses the operator norm inequalities $\|\phi\|_{1 \rightarrow 1} \leq \|\phi\|_\diamond$ and Eqn. 1.2.9.

6. TRACE $\xrightarrow{d^k}$ MONOMIAL:

Let M be a balanced monomial of degree k and write it as

$$M = U_{p_1 q_1} \dots U_{p_k q_k} U_{r_1 s_1}^* \dots U_{r_k s_k}^*.$$

Then let $\hat{M} = |p_1\rangle\langle q_1| \otimes \dots \otimes |p_k\rangle\langle q_k| \otimes |r_1\rangle\langle s_1| \otimes \dots \otimes |r_k\rangle\langle s_k|$. Then $M(U) = \text{tr} \hat{M} U^{\otimes k, k}$ and $\|\hat{M}\|_\infty = 1$. Now we use the fact that for any operator A

$$\|A\|_1 = \max_B \{ \text{tr} AB : \|B\|_\infty \leq 1 \} \quad (2.2.19)$$

to rewrite the TRACE definition:

$$\begin{aligned} \|\mathbb{E}_{U \sim \nu}[U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(d)}[U^{\otimes k, k}]\|_1 &= \\ \max_B \left\{ \text{tr} \left(\mathbb{E}_{U \sim \nu}[U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(d)}[U^{\otimes k, k}] \right) B : \|B\|_\infty \leq 1 \right\} & \\ \geq \left| \text{tr} \left(\mathbb{E}_{U \sim \nu}[U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(d)}[U^{\otimes k, k}] \right) \hat{M} \right| & \\ = \left| \mathbb{E}_{U \sim \nu} M(U) - \mathbb{E}_{U \sim \mathcal{U}(d)} M(U) \right|. & \end{aligned}$$

7. MONOMIAL $\xrightarrow{d^5}$ TWIRL (for $k = 2$):

Write $\Lambda(\rho)$ in the Kraus decomposition as

$$\Lambda(\rho) = \sum_k A_k \rho A_k^\dagger \quad (2.2.20)$$

with

$$\sum_k A_k^\dagger A_k = I. \quad (2.2.21)$$

Let $\Lambda_U(\rho) = \sum_k U^\dagger A_k U \rho U^\dagger A_k^\dagger U$. Then the p, q matrix element of $\Lambda_U(\rho)$ is

$$\sum_{krstuij} \rho_{ij} U_{si} U_{uq} U_{rp}^* U_{tj}^* A_{krs} A_{kut}^*. \quad (2.2.22)$$

From Definition 2.2.13 (MONOMIAL) we have that

$$\left| (\mathbb{E}_{U \sim \nu} - \mathbb{E}_{U \sim \mathcal{U}(d)}) U_{si} U_{uq} U_{rp}^* U_{tj}^* \right| \leq \epsilon/d^2$$

(treating expectation as an operator). This implies that

$$\left| \sum_{krstuij} \rho_{ij} A_{krs} A_{kut}^* (\mathbb{E}_{U \sim \nu} - \mathbb{E}_{U \sim \mathcal{U}(d)}) U_{si} U_{uq} U_{rp}^* U_{tj}^* \right| \leq \sum_{krstuij} |\rho_{ij} A_{krs} A_{kut}^*| \epsilon / d^2. \quad (2.2.23)$$

Now, $\sum_{rs} |A_{krs}| \leq d \|A_k\|_2$ and $\sum_{ij} |\rho_{ij}| \leq d \|\rho\|_2$ and, taking the trace of the normalisation condition Eqn. 2.2.21 we find

$$d = \sum_k \text{tr} A_k^\dagger A_k = \sum_k \|A_k\|_2^2.$$

So we find Eqn. 2.2.23 is upper bounded by

$$\frac{\epsilon}{d^2} d \|\rho\|_2 \sum_k d^2 \|A_k\|_2^2 \leq \epsilon d^2 \|\rho\|_2.$$

Using the fact that the 2-norm squared is the sum of the squares of the matrix elements we find that

$$\|\mathbb{E}_{U \sim \nu} \Lambda_U(\rho) - \mathbb{E}_{U \sim \mathcal{U}(d)} \Lambda_U(\rho)\|_2 \leq \epsilon d^4 \|\rho\|_2.$$

Using $\|\cdot\|_\diamond \leq d \|\cdot\|_2$ (Eqn. 1.2.12) we prove the result.

8. TWIRL \rightarrow MONOMIAL for $k = 2$:

Let $A_\sigma = |p\rangle\langle q| + \sigma|r\rangle\langle s|$ where $\sigma \in \{+1, -1, +i, -i\}$. Let $B = I - |q\rangle\langle q| - |s\rangle\langle s|$. Then A_σ and B are the Kraus operators of a valid channel, provided $p \neq r$, which we assume for now. Further, let

$$\Lambda_{U,\sigma}(\rho) = U^\dagger \Lambda_\sigma(U \rho U^\dagger) U \quad (2.2.24)$$

where Λ_σ is the channel with Kraus operators A_σ and B . Now let

$$\Lambda_{U,s}(\rho) = \Lambda_{U,+1}(\rho) - \Lambda_{U,-1}(\rho) + i\Lambda_{U,+i}(\rho) - i\Lambda_{U,-i}(\rho). \quad (2.2.25)$$

We see that

$$\Lambda_{U,s}(\rho) = 4U^\dagger|p\rangle\langle q|U\rho U^\dagger|s\rangle\langle r|U. \quad (2.2.26)$$

Now, from Definition 2.2.11 (TWIRL) and the triangle inequality (using $\|\cdot\|_2 \leq \|\cdot\|_1$), we have

$$\|\mathbb{E}_{U \sim \nu} \Lambda_{U,s}(\rho) - \mathbb{E}_{U \sim \mathcal{U}(d)} \Lambda_{U,s}(\rho)\|_2 \leq \frac{4\epsilon\|\rho\|_1}{d^2}. \quad (2.2.27)$$

This implies that each matrix element is small i.e.

$$|(\mathbb{E}_{U \sim \nu} - \mathbb{E}_{U \sim \mathcal{U}(d)})\langle c|U^\dagger|p\rangle\langle q|U\rho U^\dagger|s\rangle\langle r|U|d\rangle| \leq \frac{\epsilon\|\rho\|_1}{d^2}. \quad (2.2.28)$$

Now let $\rho = |e\rangle\langle f|$. We do not have to choose a physical state since the diamond-norm bound is true for all matrices. This gives us

$$|(\mathbb{E}_{U \sim \nu} - \mathbb{E}_{U \sim \mathcal{U}(d)})U_{pc}^*U_{qe}U_{sf}^*U_{rd}| \leq \frac{\epsilon}{d^2} \quad (2.2.29)$$

as required.

For $p = r$, we also assume that $s = q$ since if not, just take $p \neq r$ and $s = q$ and swap the labels. Here take $A_\pm = \pm|p\rangle\langle q|$ and $B = I - |q\rangle\langle q|$ and consider $\Lambda_{U,+}(\rho) - \Lambda_{U,-}(\rho) = 2U^\dagger|p\rangle\langle q|U\rho U^\dagger|q\rangle\langle p|U$. \square

We remark that other types of approximate definitions are possible. For cryptographic uses, a computationally secure approximate design may be sufficient, rather than the information theoretic security discussed above. A computationally secure approximate design would be nearly indistinguishable from an exact design in polynomial time. Applications and constructions of such objects remain open problems.

Constructions

Here we summarise the known constructions of unitary and state designs. We will say that a k -design construction is efficient if the effort required to sample a state or

unitary from the design is polynomial in n and k . Note that we do not require the number of states or unitaries to be polynomial because, even for approximate designs, an exponential number is required. Rather, the number of random bits needed to specify an element of the design should be $\text{poly}(n, k)$.

We start with state design constructions since these have been studied far more than unitary designs. Firstly, exact efficient state 1-designs are trivial: simply choose a random state from any basis. Numerous examples of exact efficient state 2-design constructions are known (e.g. [Bar02]). Hayashi et al. [HHH06] give an inefficient construction of state k -designs for any n and k but general exact constructions are not efficient in n and k . However, Ambainis and Emerson provide an efficient approximate construction for any k with $d \geq 2k$. Aaronson [Aar09] also gives an efficient approximate construction.

Less is known about efficient constructions for unitary designs. It is straightforward to prove that the Pauli matrices form an exact 1-design and in [DLT02, Dan05] it is shown that the Clifford group (see Chapter 6 for a definition) forms an exact 2-design although no efficient exact sampling method is known. However, an approximate sampling method is given in [DLT02] and a more efficient approximate 2-design construction is given in [DCEL06]. The structure of unitary 2-designs is considered in [GAE07], providing lower bounds on the number of unitaries in the design.

In Chapter 4 we give the first efficient approximate unitary k -design construction for $k > 2$. The construction works in $O(kn + \log 1/\epsilon)$ time for $k = O(n/\log n)$. Through Lemma 2.2.14, the construction is efficient for all the equivalent definitions above. We also conjecture in Chapter 3 that random quantum circuits of length $\text{poly}(n, k)$ are approximate unitary k -designs although we only prove this for $k = 2$.

Chapter 3

Random Quantum Circuits

3.1 Introduction: Pseudo-random Quantum Circuits

Random circuits are a natural object to consider when looking at the complexity of random operations. They are circuits where the gates and their positions are chosen randomly from some given distribution. If the gate set that the random circuit chooses from is universal then, as we show below, the random circuit will converge to the uniform Haar measure. The advantage of considering a random circuit rather than a random unitary on the whole system is it is naturally efficient to implement, for polynomial length circuits. Random circuits of some fixed length are also a new measure on the unitary group which, as we show later, reproduces some of the properties of the Haar measure for polynomial length. As well as the computer science aspects, this has applications in physics since randomly interacting systems could be modelled as a random circuit. These systems will only reach their equilibrium if the random circuit converges quickly. Thus proving convergence of the random circuit shows that some physical systems will have some properties of Haar random systems after evolving for a short amount of time.

We consider a general class of random circuits where a series of two-qubit gates are chosen from a universal gate set. We give a framework for analysing the k^{th} moments of these circuits. Our conjecture, based on an analogous classical result [BH08], is

that a random circuit on n qubits of length $\text{poly}(n, k)$ is an approximate k -design. While we do not prove this, we instead give a tight analysis of the $k = 2$ case. We find that in a broad class of natural random circuit models (described in Section 3.1.1), a circuit of length $O(n(n + \log 1/\epsilon))$ yields an ϵ -approximate 2-design. The approximate design definition used in this section is the diamond-norm definition given in Definition 2.2.10 and, through Lemma 2.2.14, applies to the alternative definitions given above. Moreover, our results also apply to random stabiliser circuits, meaning that a random stabiliser circuit of length $O(n(n + \log 1/\epsilon))$ will be an ϵ -approximate 2-design. This both simplifies the construction and tightens the efficiency of the approach of [DLT02], which constructed ϵ -approximate 2-designs in time $O(n^6(n^2 + \log 1/\epsilon))$ using $O(n^3)$ elementary quantum gates.

3.1.1 Random Circuits

The random circuit we will use is the following. Choose a 2-qubit gate set that is universal on $U(4)$ (or on the stabiliser subgroup of $U(4)$). One example of this is the set of all one qubit gates together with the controlled-NOT gate. Another is simply the set of all of $U(4)$. Then, at each step, choose a random pair of qubits and apply a gate from the universal set chosen uniformly at random. For the $U(4)$ case, the distribution will be the Haar measure on $U(4)$. One such circuit is shown in Fig. 3.1 for $n = 4$ qubits. This is based on the approach used in [ODP07, DOP07] but our analysis is both simpler and more general.

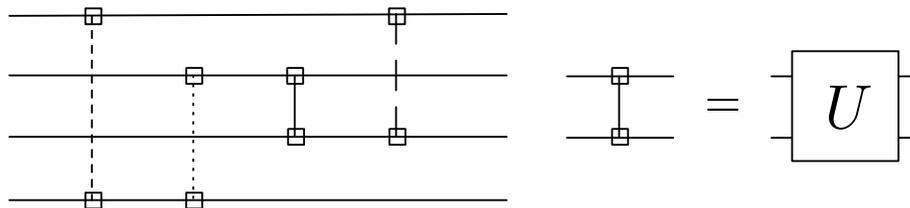


Figure 3.1: An example of a random circuit. Different lines indicate a different gate is applied at each step.

Since the universal set can generate the whole of $U(2^n)$ in this way, such random

circuits can produce any unitary. Further, since this process converges to a unitarily invariant distribution and the Haar distribution is unique, the resulting unitary must be uniformly distributed amongst all unitaries [ELL05]. Therefore this process will eventually converge to a Haar distributed unitary from $U(2^n)$. This is proven rigorously in Lemma 3.4.7. However, since a Haar unitary cannot be produced in polynomial time, this process will not converge in polynomial time. We address this problem by considering only the lower-order moments of the distribution and showing these are nearly the same for random circuits as for Haar-distributed unitaries. This claim is formally described in Theorem 3.3.1.

This chapter is organised as follows. In Section 3.2 we explain how a random circuit could be used to construct a k -design. We then summarise the results of this chapter in Section 3.3. In Section 3.4 we work out how the state evolves after a single step of the random circuit. We then extend this to multiple steps in Section 3.5 and prove our general convergence results. A key simplification will be (following [ODP07]) to map the evolution of the second moments of the quantum circuit onto a classical Markov chain. We then prove a tight convergence result for the case where the gates are chosen from $U(4)$ in Section 3.6. This section contains most of the technical content of the chapter. Using our bounds on mixing time we put together the proof that random circuits yield approximate unitary 2-designs in Section 3.7. Section 3.8 concludes with some discussion of applications.

The majority of this chapter, with the exception of Section 3.6.4, has been published previously as [HL09b] and is joint work with Aram Harrow.

3.2 Preliminaries

3.2.1 Pauli expansion

Much of the following will be done in the Pauli basis. In this chapter, we choose the normalisation so that ρ is written in the Pauli basis as

$$\rho = 2^{-n/2} \sum_p \gamma(p) \sigma_p. \quad (3.2.1)$$

With this normalisation,

$$\sum_p \gamma(p)^2 = \text{tr } \rho^2 \quad (3.2.2)$$

which is 1 for pure ρ . In general,

$$\sum_p \gamma^2(p) \leq 1$$

with equality if and only if ρ is pure. Note also that $\text{tr } \rho = 1$ is equivalent to $\gamma(0) = 2^{-n/2}$.

This notation is extended to states on nk qubits by treating γ as a function of k strings from $\{0, 1, 2, 3\}^n$. Thus a state ρ on nk qubits is written as

$$\rho = 2^{-nk/2} \sum_{p_1, \dots, p_k} \gamma_0(p_1, \dots, p_k) \sigma_{p_1} \otimes \dots \otimes \sigma_{p_k}. \quad (3.2.3)$$

3.2.2 Random Circuits as k -designs

If a random circuit is to be an approximate k -design then Eqn. 2.2.12 must be satisfied where the unitaries in \mathcal{G}_ν are the different possible random circuits. We can think of this as applying the random circuit not once but k times to k different systems.

Suppose that applying t random gates yields the random circuit W . If $W^{\otimes k}$ acts on an nk -qubit state ρ , then the resulting state is

$$\rho_W := W^{\otimes k} \rho (W^\dagger)^{\otimes k} = 2^{-nk/2} \sum_{p_1, \dots, p_k} \gamma_0(p_1, \dots, p_k) W \sigma_{p_1} W^\dagger \otimes \dots \otimes W \sigma_{p_k} W^\dagger. \quad (3.2.4)$$

For this to be a k -design, the expectation over all choices of random circuit should match the expectation over Haar-distributed $W \in U(2^n)$.

We are now ready to state our main results. Our results apply to a large class of gate sets which we define below:

Definition 3.2.1. Let $\mathcal{E} = \{p_i, U_i\}$ be a discrete ensemble of elements from $U(d)$. Define an operator $G_{\mathcal{E}}$ by

$$G_{\mathcal{E}} := \sum_i p_i U_i^{\otimes k, k} \quad (3.2.5)$$

where $U^{\otimes k, k} = U^{\otimes k} \otimes (U^*)^{\otimes k}$. More generally, we can consider continuous distributions. If μ is a probability measure on $U(d)$ then we can define G_{μ} by analogy as

$$G_{\mu} := \int_{U(d)} d\mu(U) U^{\otimes k, k} \quad (3.2.6)$$

Then \mathcal{E} (or μ) is k -copy gapped if $G_{\mathcal{E}}$ (or G_{μ}) has only $k!$ eigenvalues with absolute value equal to 1.

For any discrete ensemble $\mathcal{E} = \{p_i, U_i\}$, we can define a measure $\mu = \sum_i p_i \delta_{U_i}$. Thus, it suffices to state our theorems in terms of μ and G_{μ} . We also remark that the k -copy gapped property is the same as the k -tensor product expander property for any non-zero gap as defined in Chapter 4.

The condition on G_{μ} in the above definition may seem somewhat strange. We will see in Section 3.4 that when $d \geq k$ there is a $k!$ -dimensional subspace of $(\mathbb{C}^d)^{\otimes 2k}$ that is acted upon trivially by any G_{μ} . Additionally, when μ is the Haar measure on $U(d)$ then G_{μ} is the projector onto this space. Thus, the k -copy gapped condition implies that vectors orthogonal to this space are shrunk by G_{μ} .

We will see that G_{μ} is k -copy gapped in a number of important cases. First, we give a definition of universality that can apply not only to discrete gates sets, but to arbitrary measures on $U(4)$.

Definition 3.2.2. Let μ be a distribution on $U(4)$. Suppose that for any open ball $S \subset U(4)$ there exists a positive integer ℓ such that $\mu^{\star \ell}(S) > 0$. Then we say μ is

universal [for $U(4)$].

Here $\mu^{\star\ell}$ is the ℓ -fold convolution of μ with itself; i.e.

$$\mu^{\star\ell} = \int \delta_{U_1 \dots U_\ell} d\mu(U_1) \cdots d\mu(U_\ell).$$

When μ is a discrete distribution over a set $\{U_i\}$, Definition 3.2.2 is equivalent to the usual definition of universality for a finite set of unitary gates.

Theorem 3.2.3. *The following distributions on $U(4)$ are k -copy gapped:*

- (i) *Any universal gate set. Examples are $U(4)$ itself, any entangling gate together with all single qubit gates, or the gate set considered in [ODP07].*
- (ii) *Any approximate (or exact) unitary k -design on 2 qubits, such as the uniform distribution over the 2-qubit Clifford group, which is an exact 2-design.*

Proof.

(i) This is proven in Lemma 3.4.7.

(ii) This follows straight from Definition 2.2.6. □

3.3 Summary of Results

Theorem 3.3.1. *Let μ be a 2-copy gapped distribution and W be a random circuit on n qubits obtained by drawing t random unitaries according to μ and applying each of them to a random pair of qubits. Then there exists C (depending only on μ) such that for any $\epsilon > 0$ and any $t \geq C(n(n + \log 1/\epsilon))$, \mathcal{G}_W is an ϵ -approximate unitary 2-design according to either Definition 2.2.10 (DIAMOND) or Definition 2.2.11 (TWIRL).*

To prove Theorem 3.3.1, we show that the second moments of the random circuits converge quickly to those of a uniform Haar distributed unitary. For W a circuit as in Theorem 3.3.1, write $\gamma_W(p_1, p_2)$ for the Pauli coefficients of $\rho_W = W^{\otimes 2} \rho (W^\dagger)^{\otimes 2}$. Then write $\gamma_t(p_1, p_2) = \mathbb{E}_W \gamma_W(p_1, p_2)$ where W is a circuit of length t . Then we have

Lemma 3.3.2. *Let μ and W be as in Theorem 3.3.1. Let the initial state be ρ with $\gamma_0(p, p) \geq 0$ and $\sum_p \gamma_0(p, p) = 1$ (for example the state $|\psi\rangle\langle\psi| \otimes |\psi\rangle\langle\psi|$ for any pure state $|\psi\rangle$). Then there exists a constant C (possibly depending on μ) such that for any $\epsilon > 0$*

$$(i) \quad \sum_{\substack{p_1, p_2 \\ p_1 p_2 \neq 00}} \left(\gamma_t(p_1, p_2) - \delta_{p_1 p_2} \frac{1}{2^n(2^n + 1)} \right)^2 \leq \epsilon \quad (3.3.1)$$

for $t \geq Cn \log 1/\epsilon$.

$$(ii) \quad \sum_{\substack{p_1, p_2 \\ p_1 p_2 \neq 00}} \left| \gamma_t(p_1, p_2) - \delta_{p_1 p_2} \frac{1}{2^n(2^n + 1)} \right| \leq \epsilon \quad (3.3.2)$$

for $t \geq Cn(n + \log 1/\epsilon)$ or, when μ is the uniform distribution on $U(4)$ or its stabiliser subgroup, $t \geq Cn \log \frac{n}{\epsilon}$.

We can then extend this to all states by a simple corollary:

Corollary 3.3.3. *Let μ , W and γ_W be as in Lemma 3.3.2. Then, for any initial state $\rho = \frac{1}{2^n} \sum_{p_1, p_2} \gamma_0(p_1, p_2) \sigma_{p_1} \otimes \sigma_{p_2}$, there exists a constant C (possibly depending on μ) such that for any $\epsilon > 0$*

$$(i) \quad \sum_{\substack{p_1, p_2 \\ p_1 p_2 \neq 00}} \left(\gamma_t(p_1, p_2) - \delta_{p_1 p_2} \frac{\sum_{p \neq 0} \gamma_0(p, p)}{4^n - 1} \right)^2 \leq \epsilon \quad (3.3.3)$$

for $t \geq Cn(n + \log 1/\epsilon)$.

$$(ii) \quad \sum_{\substack{p_1, p_2 \\ p_1 p_2 \neq 00}} \left| \gamma_t(p_1, p_2) - \delta_{p_1 p_2} \frac{\sum_{p \neq 0} \gamma_0(p, p)}{4^n - 1} \right| \leq \epsilon \quad (3.3.4)$$

for $t \geq Cn(n + \log 1/\epsilon)$.

By the diamond-norm definition of an approximate design (Definition 2.2.10), we only need convergence in the 2-norm (Eqn. 3.3.3), which is implied by 1-norm

convergence (Eqn. 3.3.4) but weaker. However, Definition 2.2.11 (TWIRL), which requires the map to be close to the twirling operation, requires 1-norm convergence (i.e. Eqn. 3.3.4). Thus, Theorem 3.3.1 for Definition 2.2.10 (DIAMOND) follows from Corollary 3.3.3(i) and Theorem 3.3.1 for Definition 2.2.11 (TWIRL) follows from Corollary 3.3.3(ii). Theorem 3.3.1 is proved in Section 3.7 and Corollary 3.3.3 in Section 3.5.

We note that we do not need to separately prove the result for Definition 2.2.11 (TWIRL) since the result follows from the equivalence of the k -design definitions (Lemma 2.2.14). However, we include the proof since, if our bounds were improved to show convergence in $O(n \log \frac{n}{\epsilon})$ time, if we simply applied Lemma 2.2.14, this would only imply that $O(n(n + \log 1/\epsilon))$ time was needed for Definition 2.2.11 (TWIRL).

We also emphasise that, in the course of proving Lemma 3.3.2, we prove that the eigenvalue gap (defined in Section 3.5.3) of the Markov chain that gives the evolution of the $\gamma(p, p)$ terms is $O(1/n)$. It is easy to show that this bound is tight for some gate sets.

Related work: Here we compare our work with other related results and efficient constructions of approximate unitary 2-designs.

- The uniform distribution over the Clifford group on n qubits is an exact 2-design [DLT02]. Moreover, [DLT02] described how to sample from the Clifford group using $O(n^8)$ classical gates and $O(n^3)$ quantum gates. Our results show that applying $O(n(n + \log 1/\epsilon))$ random two-qubit Clifford gates also achieve an ϵ -approximate 2-design (although not necessarily a distribution that is within ϵ of uniform on the Clifford group).
- Dankert et al. [DCEL06] gave a specific circuit construction of an approximate 2-design. To achieve small error in the sense of Definition 2.2.10 (DIAMOND), their circuits require the same $O(n(n + \log 1/\epsilon))$ gates that our random circuits do. However, when we use Definition 2.2.11 (TWIRL), the circuits from [DCEL06] only need $O(n \log 1/\epsilon)$ gates while we only show that random circuits of length $O(n(n + \log 1/\epsilon))$ suffice.

- The closest results to our own are in the papers by Oliveira et al. [ODP07, DOP07], which considered a specific gate set (random single qubit gates and a controlled-NOT) and proved that the second moments converge in time $O(n^2(n + \log 1/\epsilon))$. Our strategy of analysing random quantum circuits in terms of classical Markov chains is also adapted from [ODP07, DOP07]. In Section 3.4, we generalise this approach to analyse the k^{th} moments for arbitrary k .

Our main results extend the results of [ODP07, DOP07] to a larger class of gate sets and improve their convergence bounds. Some of these improvements have been conjectured by [Zni07], where the author presented numerical evidence in support of them.

- An algorithmic application of random circuits was given in [HH08], where they were used to construct a new class of superpolynomial quantum speedups. In that paper, random circuits of length $O(n^3)$ were used in order to guarantee that they were so-called “dispersing” circuits. Our results immediately imply that circuits of length $O(n^2)$ would instead suffice. We believe that this could be further improved with a specialised argument, since [HH08] assumed that the input to the random circuit was always a computational basis state.

3.4 Analysis of the Moments

In order to prove our results, we need to understand how the state evolves after each step of the random circuit. In this section we consider just one step and a fixed pair of qubits. Later on we will extend this to prove convergence results for multiple steps with random pairs of qubits drawn at every step. We consider first the Haar distribution over the full unitary group and then will discuss the more general case of any 2-copy gapped distribution.

In this section, we work in general dimension d and with a general Hermitian orthogonal basis $\sigma_0, \dots, \sigma_{d^2-1}$. Later we will take d to be either 4 or 2^n and the σ_i to be Pauli matrices. However, in this section we keep the discussion general to

emphasise the potentially broader applications.

Fix an orthonormal basis for $d \times d$ Hermitian matrices: $\sigma_0, \dots, \sigma_{d^2-1}$, normalised so that $\text{tr} \sigma_p \sigma_q = d \delta_{p,q}$. Let σ_0 be the identity. We need to evaluate the quantity

$$\mathbb{E}_U \left(U^{\otimes k} \sigma_{p_1} \otimes \dots \otimes \sigma_{p_k} (U^\dagger)^{\otimes k} \right) =: T(\mathbf{p}) \quad (3.4.1)$$

where the expectation is over Haar distributed $U \in U(d)$. We will need this quantity in two cases. Firstly, for $d = 2^n$, these are the moments obtained after applying a uniformly distributed unitary so we know what the random circuit must converge to. Secondly, for $d = 4$, this tells us how a random $U(4)$ gate acts on any chosen pair.

Call the quantity in Eqn. 3.4.1 $T(\mathbf{p})$ (we use **bold** to indicate a k -tuple of coefficients; take $\mathbf{p} = (p_1, \dots, p_k)$) and write it in the σ_p basis as

$$T(\mathbf{p}) = \sum_{\mathbf{q}} \hat{G}(\mathbf{q}; \mathbf{p}) \sigma_{q_1} \otimes \dots \otimes \sigma_{q_k}. \quad (3.4.2)$$

Here, $\hat{G}(\mathbf{q}; \mathbf{p})$ is the coefficient in the Pauli expansion of $T(\mathbf{p})$ and we define \hat{G} as the matrix with entries equal to $\hat{G}(\mathbf{q}; \mathbf{p})$. We have left off the usual normalisation factor because, as we shall see, with this normalisation \hat{G} is a projector. Inverting this, we have

$$\begin{aligned} \hat{G}(\mathbf{q}; \mathbf{p}) &= d^{-k} \text{tr} (\sigma_{q_1} \otimes \dots \otimes \sigma_{q_k} T(\mathbf{p})) \\ &= d^{-k} \mathbb{E}_U \text{tr} \left((\sigma_{q_1} \otimes \dots \otimes \sigma_{q_k}) U^{\otimes k} (\sigma_{p_1} \otimes \dots \otimes \sigma_{p_k}) (U^\dagger)^{\otimes k} \right) \end{aligned} \quad (3.4.3)$$

Note that \hat{G} is real since T and the basis are Hermitian.

We can gain all the information we need about the Haar integral in Eqn. 3.4.1 with the following observations:

Lemma 3.4.1. $T(\mathbf{p})$ commutes with $U^{\otimes k}$ for any unitary U .

Proof. Follows from the invariance of the Haar measure on the unitary group. \square

Corollary 3.4.2. $T(\mathbf{p})$ is a linear combination of permutations from the symmetric group S_k .

Proof. This follows from Schur-Weyl duality (see e.g. [GW98]). \square

From this, we can prove that \hat{G} is a projector and find its eigenvectors.

Theorem 3.4.3. \hat{G} is symmetric, i.e. $\hat{G}(\mathbf{q}; \mathbf{p}) = \hat{G}(\mathbf{p}; \mathbf{q})$.

Proof. Follows from the invariance of the trace under cyclic permutations. \square

Theorem 3.4.4. $S(\pi)$ is an eigenvector of \hat{G} with eigenvalue 1 for any subsystem permutation operator $S(\pi)$ i.e.

$$\sum_{\mathbf{q}} \hat{G}(\mathbf{p}; \mathbf{q}) \operatorname{tr}(\sigma_{q_1} \otimes \dots \otimes \sigma_{q_k} S(\pi)) = \operatorname{tr}(\sigma_{p_1} \otimes \dots \otimes \sigma_{p_k} S(\pi)).$$

Further, any vector orthogonal to this set has eigenvalue 0.

Proof. For the first part,

$$\begin{aligned} & \sum_{\mathbf{q}} \hat{G}(\mathbf{p}; \mathbf{q}) \operatorname{tr}(\sigma_{q_1} \otimes \dots \otimes \sigma_{q_k} S(\pi)) \\ &= d^{-k} \sum_{\mathbf{q}} \mathbb{E}_U \operatorname{tr}(\sigma_{q_1} U \sigma_{p_1} U^\dagger) \dots \operatorname{tr}(\sigma_{q_k} U \sigma_{p_k} U^\dagger) \operatorname{tr}(\sigma_{q_1} \otimes \dots \otimes \sigma_{q_k} S(\pi)) \\ &= d^{-k} \operatorname{tr} \left(S(\pi) \mathbb{E}_U \sum_{q_1} \operatorname{tr}(\sigma_{q_1} U \sigma_{p_1} U^\dagger) \sigma_{q_1} \otimes \dots \otimes \sum_{q_k} \operatorname{tr}(\sigma_{q_k} U \sigma_{p_k} U^\dagger) \sigma_{q_k} \right) \end{aligned} \tag{3.4.4}$$

Writing $U^\dagger \sigma_p U$ in the σ_p basis, we find

$$\frac{1}{d} \sum_q \operatorname{tr}(\sigma_q U \sigma_p U^\dagger) \sigma_q = U \sigma_p U^\dagger.$$

Therefore Eqn. 3.4.4 becomes

$$\operatorname{tr} \left(S(\pi) \mathbb{E}_U U^\dagger \sigma_{p_1} U \otimes \dots \otimes U^\dagger \sigma_{p_k} U \right) = \operatorname{tr}(\sigma_{p_1} \otimes \dots \otimes \sigma_{p_k} S(\pi)).$$

For the second part, consider any vector v which is orthogonal to the permutation operators (we can neglect the complex conjugate because $S(\pi)$ is real in this basis), i.e.

$$\sum_{\mathbf{q}} \text{tr}(\sigma_{q_1} \otimes \dots \otimes \sigma_{q_k} S(\pi)) v(\mathbf{q}) = 0 \quad (3.4.5)$$

for any permutation π . Then

$$\sum_{\mathbf{q}} \hat{G}(\mathbf{p}; \mathbf{q}) v(\mathbf{q}) = d^{-k} \sum_{\mathbf{q}} \text{tr}(\sigma_{q_1} \otimes \dots \otimes \sigma_{q_k} T(\mathbf{p})) v(\mathbf{q})$$

which is zero since $T(\mathbf{p})$ is a linear combination of permutations and v is orthogonal to this by Eqn. 3.4.5. \square

Theorem 3.4.5. $\hat{G}^2 = \hat{G}$, i.e. $\sum_{\mathbf{q}'} \hat{G}(\mathbf{p}; \mathbf{q}') \hat{G}(\mathbf{q}'; \mathbf{q}) = \hat{G}(\mathbf{p}; \mathbf{q})$.

Proof. Using Eqn. 3.4.3,

$$\sum_{\mathbf{q}'} \hat{G}(\mathbf{p}; \mathbf{q}') \hat{G}(\mathbf{q}'; \mathbf{q}) = \sum_{\mathbf{q}'} \hat{G}(\mathbf{p}; \mathbf{q}') d^{-k} \text{tr}(\sigma_{q'_1} \otimes \dots \otimes \sigma_{q'_k} T(\mathbf{q})).$$

From Corollary 3.4.2, $T(\mathbf{q})$ is a linear combination of permutations. This implies, using Theorem 3.4.4 that

$$\begin{aligned} \sum_{\mathbf{q}'} \hat{G}(\mathbf{p}; \mathbf{q}') d^{-k} \text{tr}(\sigma_{q'_1} \otimes \dots \otimes \sigma_{q'_k} T(\mathbf{q})) &= d^{-k} \text{tr}(\sigma_{p_1} \otimes \dots \otimes \sigma_{p_k} T(\mathbf{q})) \\ &= \hat{G}(\mathbf{p}; \mathbf{q}) \end{aligned}$$

as required. \square

Corollary 3.4.6. \hat{G} is a projector so has eigenvalues 0 and 1.

We now evaluate \hat{G} and T for the cases of $k = 1$ and $k = 2$ since these are the cases we are interested in for the remainder of the chapter.

3.4.1 $k = 1$

The $k = 1$ case is clear: the random unitary completely randomises the state. Therefore all terms in the expansion are set to zero apart from the identity i.e.

$$T(p) = \begin{cases} \sigma_0 & p = 0 \\ 0 & p \neq 0. \end{cases} \quad (3.4.6)$$

3.4.2 $k = 2$

For $k = 2$, there are just two permutation operators, identity I and swap \mathcal{F} . Therefore there are just two eigenvectors with non-zero eigenvalue ($n > 1$). In normalised form, taking them to be orthogonal, their components are

$$\begin{aligned} f_1(q_1, q_2) &= \delta_{q_1 0} \delta_{q_2 0} \\ f_2(q_1, q_2) &= \frac{1}{d^2 - 1} \delta_{q_1 q_2} (1 - \delta_{q_1 0}) \end{aligned}$$

We will now prove three properties of \hat{G} that we need:

1. $\hat{G}(p_1, p_2; q_1, q_2) = 0$ if $p_1 \neq p_2$ or $q_1 \neq q_2$.

Proof. Consider the function $f(q_1, q_2) = \delta_{q_1 a} \delta_{q_2 b}$ with $a \neq b$. This function has zero overlap with the eigenvectors f_1 and f_2 so it goes to zero when acted on by \hat{G} . Therefore $\hat{G}(p_1, p_2; a, b) = 0$. The claim follows from the symmetry property (Theorem 3.4.3). \square

With this we will write $\hat{G}(p; q) \equiv \hat{G}(p_1, p_2; q_1, q_2)$.

2. $\hat{G}(p; 0) = \delta_{p 0}$.

Proof. Let \hat{G} act on eigenvector f_1 . \square

3. $\hat{G}(p; a) = \frac{1}{d^2 - 1}$ for $a, p \neq 0$.

Proof. Let \hat{G} act on the input δ_{qa} . This has zero overlap with f_1 and overlap $\frac{1}{d^2-1}$ with f_2 . \square

Therefore we have

$$\hat{G}(p_1, p_2; q_1, q_2) = \begin{cases} 0 & p_1 \neq p_2 \text{ or } q_1 \neq q_2 \\ 1 & p_1 = p_2 = q_1 = q_2 = 0 \\ \frac{1}{d^2-1} & p_1 = p_2 \neq 0, q_1 = q_2 \neq 0 \end{cases} \quad (3.4.7)$$

Since $T(p_1, p_2) = \sum_{q_1, q_2} \hat{G}(p_1, p_2; q_1, q_2) \sigma_{q_1} \otimes \sigma_{q_2}$, we have

$$T(p_1, p_2) = \begin{cases} 0 & p_1 \neq p_2 \\ \sigma_0 \otimes \sigma_0 & p_1 = p_2 = 0 \\ \frac{1}{d^2-1} \sum_{p' \neq 0} \sigma_{p'} \otimes \sigma_{p'} & p_1 = p_2 \neq 0. \end{cases} \quad (3.4.8)$$

Therefore the terms $\sigma_{p_1} \otimes \sigma_{p_2}$ with $p_1 \neq p_2$ are set to zero. Further, the sum of the diagonal coefficients $\gamma(p, p)$ is conserved. This allows us to identify this with a probability distribution (after renormalising) and use Markov chain analysis. To see this, write again the starting state

$$\rho = \frac{1}{d} \sum_{q_1, q_2} \gamma_0(q_1, q_2) \sigma_{q_1} \otimes \sigma_{q_2}$$

with state after application of any unitary W

$$\rho_W = \frac{1}{d} \sum_{q_1, q_2} \gamma_W(q_1, q_2) \sigma_{q_1} \otimes \sigma_{q_2} = 2^{-n} \sum_{q_1, q_2} \gamma(q_1, q_2) \left(W \sigma_{q_1} W^\dagger \right) \otimes \left(W \sigma_{q_2} W^\dagger \right).$$

Then

$$\begin{aligned}
\sum_q \gamma_W(q, q) &= \frac{1}{d} \sum_q \text{tr}(\sigma_q \otimes \sigma_q \rho_W) \\
&= \text{tr}(\mathcal{F} \rho_W) \\
&= \frac{1}{d} \sum_{q_1, q_2} \gamma(q_1, q_2) \text{tr}\left(\mathcal{F}\left(W \sigma_{q_1} W^\dagger\right) \otimes \left(W \sigma_{q_2} W^\dagger\right)\right) \\
&= \frac{1}{d} \sum_{q_1, q_2} \gamma(q_1, q_2) \text{tr}(\sigma_{q_1} \sigma_{q_2}) \\
&= \sum_q \gamma(q, q)
\end{aligned}$$

as required, where \mathcal{F} is the swap operator and we have used Lemmas 1.2.2 and 1.2.1.

3.4.3 Moments for General Universal Random Circuits

We now consider universal distributions μ that in general may be different from the uniform (Haar) measure on $U(d)$. Our main result in this section will be to show that a universal distribution on $U(4)$ is also 2-copy gapped. In fact, we will phrase this result in slightly more general terms and show that a universal distribution on $U(d)$ is also k -copy gapped for any k . Universality (Definition 3.2.2) generalises in the obvious way to $U(d)$, whereas when we say that μ is k -copy gapped, we mean that

$$\|G_\mu - G_{U(d)}\|_\infty < 1, \quad (3.4.9)$$

where $G_\mu = \mathbb{E}_U U^{\otimes k, k}$, with the expectation taken over μ for G_μ or over the Haar measure for $G_{U(d)}$.

The reason Eqn. 3.4.9 represents our condition for μ to be k -copy gapped is as follows: Observe that \hat{G} and G are unitarily related, so the definition of k -copy gapped could equivalently be given in terms of \hat{G} . We have shown above that $\hat{G}_{U(d)}$ (and thus $G_{U(d)}$) has all eigenvalues equal to 0 or 1 i.e. it is a projector. By contrast, G_μ may not even be Hermitian. However, we will prove below that all eigenvectors of $G_{U(d)}$ with eigenvalue 1 are also eigenvectors of G_μ with eigenvalue 1. Thus, Eqn. 3.4.9 will

imply that $\lim_{t \rightarrow \infty} (\hat{G}_\mu)^t = \hat{G}_{U(d)}$, just as we would expect for a gapped random walk.

We would like to show that Eqn. 3.4.9 holds whenever μ is universal. This result was proved in [AK62] (and was probably known even earlier) when μ had the form $(\delta_{U_1} + \delta_{U_2})/2$. Here we show how to extend the argument to any universal μ .

Lemma 3.4.7. *Let μ be a distribution on $U(d)$. Then all eigenvectors of $G_{U(d)}$ with eigenvalue 1 are eigenvectors of G_μ with eigenvalue 1. Additionally, if μ is universal then μ is k -copy gapped for any positive integer k (cf. Eqn. 3.4.9).*

In particular, if $k = 2$ this Lemma implies that μ is 2-copy gapped (cf. Theorem 3.2.3).

Proof. Let $V \cong \mathbb{C}^d$ be the fundamental representation of $U(d)$, where the action of $U \in U(d)$ is simply U itself. Let V^* be its dual representation, where U acts as U^* . The operators G_μ and $G_{U(d)}$ act on the space $V^{\otimes k} \otimes (V^*)^{\otimes k}$. We will see that $G_{U(d)}$ is completely determined by the decomposition of $V^{\otimes k} \otimes (V^*)^{\otimes k}$ into irreducible representations (irreps). Suppose that the multiplicity of (r_λ, V_λ) in $V^{\otimes k} \otimes (V^*)^{\otimes k}$ is m_λ , where the V_λ 's are the irrep spaces and $r_\lambda(U)$ the corresponding representation matrices. In other words

$$V^{\otimes k} \otimes (V^*)^{\otimes k} \cong \bigoplus_{\lambda} V_{\lambda} \otimes \mathbb{C}^{m_{\lambda}} \quad (3.4.10)$$

$$U^{\otimes k} \otimes (U^*)^{\otimes k} \sim \sum_{\lambda} |\lambda\rangle\langle\lambda| \otimes r_{\lambda}(U) \otimes I_{m_{\lambda}} \quad (3.4.11)$$

Here \sim indicates that the two sides are related by conjugation by a fixed (U independent) unitary.

Let $\lambda = 0$ denote the trivial irrep: i.e. $V_0 = \mathbb{C}$ and $r_0(U) = 1$ for all U . We claim that $\mathbb{E}_U r_\lambda(U) = 0$ whenever $\lambda \neq 0$ and the expectation is taken over the Haar measure. To show this, note that $\mathbb{E}_U r_\lambda(U)$ commutes with $r_\lambda(V)$ for all $V \in U(d)$ and thus, by Schur's Lemma, we must have $\mathbb{E}_U r_\lambda(U) = cI$ for some $c \in \mathbb{C}$. However, by the translation-invariance of the Haar measure we have $cI = \mathbb{E}_U r_\lambda(U) = \mathbb{E}_U r_\lambda(UV) =$

$c r_\lambda(V)$ for all $V \in U(d)$. Since $\lambda \neq 0$, we cannot have $r_\lambda(V) = I$ for all V and so it must be that $c = 0$.

Thus, if we write $G_{U(d)}$ and G_μ using the basis on the RHS of Eqn. 3.4.11, we have

$$G_{U(d)} = |0\rangle\langle 0| \otimes I_{m_0} \quad (3.4.12)$$

where $|0\rangle\langle 0|$ is a projector onto the trivial irrep. On the other hand,

$$G_\mu = |0\rangle\langle 0| \otimes I_{m_0} + \sum_{\lambda \neq 0} |\lambda\rangle\langle \lambda| \otimes \left(\int r_\lambda(U) d\mu(U) \right) \otimes I_{m_\lambda} \quad (3.4.13)$$

Thus, every eigenvector of $G_{U(d)}$ with eigenvalue one is also fixed by G_μ . For the remainder of the space, the direct sum structure means that

$$\|G_{U(d)} - G_\mu\|_\infty = \max_{\substack{\lambda \neq 0 \\ m_\lambda \neq 0}} \left\| \int r_\lambda(U) d\mu(U) \right\|_\infty. \quad (3.4.14)$$

Note that this maximisation only includes λ with $\dim V_\lambda > 1$. This is because non-trivial one-dimensional irreps of $U(d)$ have the form $\det U^m$ for some non-zero integer m . Under the map $U \mapsto e^{i\phi}U$, such irreps pick up a phase of $e^{im\phi}$. However, $U^{\otimes k} \otimes (U^*)^{\otimes k}$ is invariant under $U \mapsto e^{i\phi}U$. Thus $V^{\otimes k} \otimes (V^*)^{\otimes k}$ cannot contain any non-trivial one-dimensional irreps.

Now suppose by contradiction that there exists $\lambda \neq 0$ with $m_\lambda \neq 0$ and

$$\left\| \int r_\lambda(U) d\mu(U) \right\|_\infty = 1.$$

(We do not need to consider the case $\| \int r_\lambda(U) d\mu(U) \|_\infty > 1$, since $\|r_\lambda(U)\|_\infty = 1$ for all U and $\|\cdot\|_\infty$ obeys the triangle inequality.) Indeed, the triangle inequality further implies that there exists a unit vector $|v\rangle \in V_\lambda$ such that

$$\int d\mu(U) r_\lambda(U) |v\rangle = \omega |v\rangle,$$

for some $\omega \in \mathbb{C}$ with $|\omega| = 1$.

By the above argument we can assume that $\dim V_\lambda > 1$. Since V_λ is irreducible, it cannot contain a one-dimensional invariant subspace, implying that there exists $U_0 \in U(d)$ such that

$$|\langle v|r_\lambda(U_0)|v\rangle| = 1 - \delta,$$

for some $\delta > 0$. Since $U \mapsto |\langle v|r_\lambda(U)|v\rangle|$ is continuous, there exists an open ball S around U_0 such that $|\langle v|r_\lambda(U)|v\rangle| \leq 1 - \delta/2$ for all $U \in S$. Define $\bar{S} := U(d) \setminus S$.

Now we use the fact that μ is universal to find an ℓ such that $\mu^{*\ell}(S) > 0$. Next, observe that $\int d\mu^{*\ell}(U) \langle v|r_\lambda(U)|v\rangle = \omega^\ell$. Taking the absolute value of both sides yields

$$\begin{aligned} 1 &= \left| \int_{U(d)} d\mu^{*\ell}(U) \langle v|r_\lambda(U)|v\rangle \right| \\ &\leq \int_{U(d)} d\mu^{*\ell}(U) |\langle v|r_\lambda(U)|v\rangle| \\ &= \int_S d\mu^{*\ell}(U) |\langle v|r_\lambda(U)|v\rangle| + \int_{\bar{S}} d\mu^{*\ell}(U) |\langle v|r_\lambda(U)|v\rangle| \\ &\leq \mu^{*\ell}(S) \left(1 - \frac{\delta}{2}\right) + \left(1 - \mu^{*\ell}(S)\right) \\ &< 1, \end{aligned}$$

a contradiction. We conclude that $\|G_{U(d)} - G_\mu\|_\infty < 1$. □

3.5 Convergence

In the previous section we saw that iterating any universal gate set on $U(d)$ eventually converges to the uniform distribution on $U(d)$. Since the set of all two-qubit unitaries is universal on $U(2^n)$, this implies that random circuits eventually converge to the Haar measure. In this section, we turn to proving upper bounds on this convergence rate, focusing on the first two moments.

Let $\hat{G}^{(ij)}$ be the matrix with \hat{G} (with $d = 4$) acting on qubits i and j and the identity on the others. Then, if the pair (i, j) is chosen at step t , we can find the

expected coefficients at step $t + 1$ by multiplying by $\hat{G}^{(ij)}$. In general, a random pair is chosen at each step. So

$$\gamma_{t+1}(\mathbf{p}) = \sum_{\mathbf{q}} \frac{1}{n(n-1)} \sum_{i \neq j} \hat{G}^{(ij)}(\mathbf{p}; \mathbf{q}) \gamma_t(\mathbf{q}) \quad (3.5.1)$$

where γ_{t+1} are the expected coefficients at step t . We can think of this evolution as repeated application of the matrix

$$P = \frac{1}{n(n-1)} \sum_{i \neq j} \hat{G}^{(ij)}. \quad (3.5.2)$$

For $k = 2$, the key idea of Oliveira et al. [ODP07] was to map the evolution of the $\gamma(p, p)$ coefficients to a Markov chain. The $\gamma(p_1, p_2)$ coefficients with $p_1 \neq p_2$ just decay as each qubit is chosen and can be analysed directly.

However, we can only map the $\gamma(p, p)$ coefficients to a probability distribution when they are non-negative, which is not the case for general states. Most of the rest of the chapter is dedicated to proving Lemma 3.3.2, which only applies to states with $\gamma(p, p) \geq 0$ and normalised so their sum is 1. Corollary 3.3.3 then extends this to all states:

Proof of Corollary 3.3.3. Lemma 3.3.2 still applies to the $\gamma(p_1, p_2)$ terms with $p_1 \neq p_2$. Therefore we just need to show how to apply Lemma 3.3.2 to states that initially have some negative $\gamma(p, p)$ terms.

For the $\gamma(p, p)$ terms, Lemma 3.3.2 says that the random walk starting with any initial probability distribution converges to uniform in some bounded time t . Let $g_t(p, p; q, q)$ be the coefficients after t steps of the walk starting at a particular point q (i.e. $g_0(p, p; q, q) = \delta_{p,q}$). Now, for any starting state ρ , let the initial coefficients be $\gamma_0(p, p)$. Then, by linearity, we can write the expected coefficients after t steps $\gamma_t(p, p) := \mathbb{E}\gamma_W(p, p)$ as

$$\gamma_t(p, p) = \sum_{q \neq 0} \gamma_0(q, q) g_t(p, p; q, q) \quad (3.5.3)$$

for $p \neq 0$.

We can now prove convergence rates for the expected coefficients $\gamma_t(p, p)$:

(i) For the 2-norm, we have from Lemma 3.3.2 that for $t \geq Cn \log 1/\epsilon$

$$\sum_{p \neq 0} \left(g_t(p, p; q, q) - \frac{1}{4^n - 1} \right)^2 \leq \epsilon \quad (3.5.4)$$

for any q . Note that the normalisation for the $\gamma(p, p)$ terms with $p \neq 0$ has changed from Lemma 3.3.2 since we are neglecting the $\gamma(0, 0)$ term here. Now

$$\begin{aligned} & \sum_{p \neq 0} \left(\gamma_t(p, p) - \frac{\sum_{q \neq 0} \gamma_0(q, q)}{4^n - 1} \right)^2 \\ &= \sum_{p \neq 0} \left(\sum_{q \neq 0} \gamma_0(q, q) \left(g_t(p, p; q, q) - \frac{1}{4^n - 1} \right) \right)^2 \\ &\leq \sum_{q \neq 0} \gamma_0(q, q)^2 \sum_{q' \neq 0} \sum_{p \neq 0} \left(g_t(p, p; q', q') - \frac{1}{4^n - 1} \right)^2 \\ &\leq (4^n - 1) \epsilon \sum_{q \neq 0} \gamma_0(q, q)^2 \\ &\leq 4^n \epsilon \sum_{q_1, q_2} \gamma_0(q_1, q_2)^2 \\ &= 4^n \epsilon \operatorname{tr} \rho^2 \\ &\leq 4^n \epsilon \end{aligned}$$

where the first inequality is the Cauchy-Schwarz inequality. Therefore for $t \geq Cn(n + \log 4^n/\epsilon)$, the 2-norm distance from stationarity for the $\gamma(p, p)$ terms is at most ϵ . Choose C' such that $C'n(n + \log 1/\epsilon) \geq Cn(n + \log 4^n/\epsilon)$ to obtain the result.

(ii) For the 1-norm, Lemma 3.3.2 says that for $t \geq Cn(n + \log 1/\epsilon)$

$$\sum_{p \neq 0} \left| g_t(q; p, p) - \frac{1}{4^n - 1} \right| \leq \epsilon. \quad (3.5.5)$$

We can then proceed much as for the 2-norm case:

$$\begin{aligned}
& \sum_{p \neq 0} \left| \gamma_t(p, p) - \frac{\sum_{q \neq 0} \gamma_0(q, q)}{4^n - 1} \right| \\
&= \sum_{p \neq 0} \left| \sum_{q \neq 0} \gamma_0(q, q) \left(g_t(p, p; q, q) - \frac{1}{4^n - 1} \right) \right| \\
&\leq \sum_{q \neq 0} |\gamma_0(q, q)| \sum_{p \neq 0} \left| g_t(p, p; q, q) - \frac{1}{4^n - 1} \right| \\
&\leq \epsilon \sum_{q \neq 0} |\gamma_0(q, q)| \\
&\leq 2^n \epsilon \sum_{q \neq 0} \gamma_0^2(q, q) \\
&\leq 2^n \epsilon.
\end{aligned}$$

Therefore for $t \geq Cn(n + \log 2^n / \epsilon)$, the 1-norm distance from stationarity for the $\gamma(p, p)$ terms is at most ϵ . \square

We now proceed to prove Lemma 3.3.2. Firstly, we will consider the simple case of $k = 1$ to prove this process forms a 1-design as this will help us to understand the more complicated case of $k = 2$.

3.5.1 First Moments Convergence

Recall that $\rho = 2^{-n/2} \sum_p \gamma(p) \sigma_p$ and we wish to evaluate the moments of the coefficients. So for the first moments to converge, we want to know $\mathbb{E}\gamma(p)$.

For $k = 1$, the $U(4)$ random circuit uniformly randomises each pair that is chosen. More precisely, a pair of sites i, j are chosen at random and all the coefficients with $p_i \neq 0$ or $p_j \neq 0$ are set to zero. Thus we get an exact 1-design when all sites have been hit. For other gate sets, the terms do not decay to zero but decay by a factor depending on the gap of \hat{G} . Call the gap Δ ; for $U(4)$ $\Delta = 1$ and for others $0 < \Delta \leq 1$ and Δ is independent of n . Therefore once each site has been hit m times the terms have decayed by a factor $(1 - \Delta)^m$.

For a bound like the mixing time (see Section 3.5.3 for definition), we want to bound the quantity $\sum_{p \neq 0} |\mathbb{E}_W \gamma_W(p)|$ where $\gamma_W(p)$ is the Pauli coefficient after applying the random circuit W . We also want 2-norm bounds, so we bound $\sum_{p \neq 0} (\mathbb{E}_W \gamma_W(p))^2$ too. We will in fact find bounds on

$$\sum_{p \neq 0} \mathbb{E}_W |\gamma_W(p)|$$

and

$$\sum_{p \neq 0} (\mathbb{E}_W |\gamma_W(p)|)^2,$$

which are stronger.

A standard problem in the theory of randomised algorithms is the *coupon collector problem*. If a magazine comes with a free coupon, which is chosen uniformly randomly from n different types, how many magazines should you buy to have a high probability of getting all n coupons? It is not hard to show that $n \ln \frac{n}{\epsilon}$ samples (magazines) have at least a $1 - \epsilon$ probability of including all n coupons. Using this, we expect all sites to be hit with probability at least $1 - \epsilon$ after $\Theta(n \log \frac{n}{\epsilon})$ steps. This argument can be made precise in this context by bounding the non-identity coefficients. We find, as expected, that the sum is small after $O(n \log n)$ steps:

Lemma 3.5.1. *After $O(n \log 1/\epsilon)$ steps*

$$\sum_{p \neq 0} (\mathbb{E}_W |\gamma_W(p)|)^2 \leq \epsilon$$

and after $O(n \log \frac{n}{\epsilon})$ steps,

$$\sum_{p \neq 0} \mathbb{E}_W |\gamma_W(p)| \leq \epsilon. \tag{3.5.6}$$

Proof. At each step, a pair of sites is chosen at random and any terms with non-identity coefficients for this pair decay by a factor $(1 - \Delta)$. For example, the term $\sigma_1 \otimes \sigma_0^{\otimes(n-1)}$ decays whenever the first site is chosen. Thus the probability of each term decaying depends on the number of zeroes. We start with the 1-norm bound.

Suppose the circuit applied after t steps is W_t . Consider $\mathbb{E}_{W_t} |\gamma_{W_t}(p)|$ for any p with d non-zeroes. Since the state ρ is physical, $\text{tr } \rho^2 \leq 1$ so $\sum_p \gamma_0^2(p) \leq 1$. Now, in each step, if any site is chosen where p is non-zero, this term decays by a factor $(1 - \Delta)$. This occurs with probability $1 - \frac{(d-n)(d-n-1)}{n(n-1)} \geq d/n$, the probability of choosing a pair where at least one site is non-zero. Therefore

$$\mathbb{E} |\gamma_{W_t}(p)| \leq ((1 - \Delta)d/n + (1 - d/n)) |\gamma_{W_{t-1}}(p)|$$

where the expectation is over the circuit applied at step t . If we iterate this t times we find

$$\mathbb{E}_W |\gamma_W(p)| \leq \exp(-\Delta t d/n) |\gamma_0(p)|$$

where the expectation here is over all random circuits for the t steps. We now sum over all p :

$$\sum_{p \neq 0} \mathbb{E}_W |\gamma_W(p)| \leq \sum_{d=1}^n \exp(-\Delta t d/n) \sum_{d(p)=d} |\gamma_0(p)|$$

where $d(p)$ is the number of non-zeroes in p . For the 1-norm bound, we can simply bound $|\gamma_0(p)| \leq 1$ to give $\sum_{d(p)=d} |\gamma_0(p)| \leq \binom{n}{d} 3^d$ so

$$\sum_{p \neq 0} \mathbb{E}_W |\gamma_W(p)| \leq (1 + 3 \exp(-\Delta t/n))^n - 1$$

where we have used the binomial theorem. Now let $t = \frac{n}{\Delta} \ln \frac{3n}{\epsilon}$. This gives

$$\sum_{p \neq 0} \mathbb{E}_W |\gamma_W(p)| \leq (1 + \epsilon/n)^n - 1 = O(\epsilon).$$

For the 2-norm bound,

$$\begin{aligned}
\sum_{p \neq 0} (\mathbb{E}_W |\gamma_W(p)|)^2 &\leq \sum_{p \neq 0} \exp(-2\Delta t d/n) \gamma_0^2(p) \\
&= \sum_{d=1}^n \exp(-2\Delta t d/n) \sum_{d(p)=d} \gamma_0^2(p) \\
&\leq \sum_{d=1}^n \exp(-2\Delta t d/n) \\
&\leq \frac{\exp(-2\Delta t/n)}{1 - \exp(-2\Delta t/n)}
\end{aligned}$$

where we have used $\sum_p \gamma_0^2(p) \leq 1$. We find after $\frac{n}{2\Delta} \ln 1/\epsilon$ steps that

$$\sum_{p \neq 0} (\mathbb{E}_W |\gamma_W(p)|)^2 \leq \frac{\epsilon}{1 - \epsilon} \quad \square$$

3.5.2 Second Moments Convergence

Firstly, the $\sigma_{p_1} \otimes \sigma_{p_2}$ terms for $p_1 \neq p_2$ decay in a similar way to the non-identity terms in the 1-design analysis. In fact, the proof of Lemma 3.5.1 carries over almost identically to this case to give

Lemma 3.5.2. *After $O(n \log 1/\epsilon)$ steps*

$$\sum_{p_1 \neq p_2} (\mathbb{E}_W |\gamma_W(p_1, p_2)|)^2 \leq \epsilon$$

and after $O(n(n + \log 1/\epsilon))$ steps

$$\sum_{p_1 \neq p_2} \mathbb{E}_W |\gamma_W(p_1, p_2)| \leq \epsilon.$$

Proof. Instead of the number of zeroes governing the decay rate, we need to count the number of places where p_1 and p_2 differ. This gives

$$\mathbb{E} |\gamma_{W_t}(p_1, p_2)| \leq ((1 - \Delta)d/n + (1 - d/n)) |\gamma_{W_{t-1}}(p_1, p_2)|$$

where now d is the number of differing sites. There are $\binom{n}{d}12^d4^{n-d}$ states that differ in d places so we find

$$\sum_{p_1 \neq p_2} \mathbb{E}_W |\gamma_W(p_1, p_2)| \leq 4^n [(1 + 3 \exp(-\Delta t/n))^n - 1].$$

Set $t = \frac{n}{\Delta}(n \ln 4 + \ln 1/\epsilon)$ to make this $O(\epsilon)$. The 2-norm bound follows in the same way as for Lemma 3.5.1. \square

We now need to prove the $\gamma(p, p)$ terms converge quickly. We have seen above that the sum of the terms $\gamma(p, p)$ is conserved and, for the purposes of proving Lemma 3.3.2, we assume the sum is 1 and $\gamma(p, p) \geq 0$ for all p .

To illustrate the evolution, consider the simplest case when the gates are chosen from $U(4)$. We have evaluated \hat{G} in Section 3.4.2 for $k = 2$ for this case. Translated into coefficients this yields the following update rule, where we have written it for the case when qubits 1 and 2 are chosen:

$$\begin{aligned} & \gamma_{t+1}(r_1, r_2, r_3, \dots, r_n, s_1, s_2, s_3, \dots, s_n) \\ &= \begin{cases} 0 & (r_1, r_2) \neq (s_1, s_2) \\ \gamma_t(0, 0, r_3, \dots, r_n, 0, 0, s_3, \dots, s_n) & (r_1, r_2) = (s_1, s_2) = (0, 0) \\ \frac{1}{15} \sum_{\substack{r'_1, r'_2 \\ r'_1 r'_2 \neq 0}} \gamma_t(r'_1, r'_2, r_3, \dots, r_n, r'_1, r'_2, s_3, \dots, s_n) & (r_1, r_2) = (s_1, s_2) \neq (0, 0). \end{cases} \end{aligned} \tag{3.5.7}$$

The key idea of Oliveira et al. [ODP07] was to map the evolution of the $\gamma(p, p)$ coefficients to a Markov chain. We can apply this here to get, on state space $\{0, 1, 2, 3\}^n$, the evolution:

1. Choose a pair of sites uniformly at random.
2. If the state is 00 it remains 00.
3. Otherwise, choose the state uniformly at random from $\{0, 1, 2, 3\}^2 \setminus \{00\}$.

This is the correct evolution since, if the initial state is distributed according to $\gamma_t(q, q)$, the final state is distributed according to $\gamma_{t+1}(p, p)$.

The evolution for other gate sets will be similar, but the states will not be chosen uniformly randomly in the third step. However, the state 00 will remain 00 and the stationary distribution on the other 15 states is the same. We will find the convergence times for general gate sets and then consider the $U(4)$ gate set since we can perform a tight analysis for this case.

3.5.3 Markov Chain Analysis

Before finding the convergence rate for our problem, we will briefly introduce the basics of Markov chain mixing time analysis. All of these standard results can be found in [MT06] and references therein.

A process is Markov if the evolution only depends on the current state rather than the full state history. Therefore the evolution of the state can be thought of as a matrix, the *transition matrix*, acting on a vector which represents the current distribution. We will only be interested in discrete time processes so the state after t steps is given by the t^{th} power of the transition matrix acting on the initial distribution.

We say a Markov chain is *irreducible* if it is possible to get from one state to any other state in some number of steps. Further, a chain is *aperiodic* if it does not return to a state at regular intervals. If a chain is both irreducible and aperiodic then it is said to be *ergodic*. A well known result of Markov chain theory is that all ergodic chains converge to a unique stationary distribution. In matrix language this says that the transition matrix P has eigenvalue 1 with no multiplicity and all other eigenvalues have absolute value strictly less than 1. We will also need the notion of *reversibility*. A Markov chain is reversible if the time reversed chain has the same transition matrix, with respect to some distribution. This condition is also known as *detailed balance*:

$$\pi(x)P(x, y) = \pi(y)P(y, x). \tag{3.5.8}$$

It can be shown that a reversible ergodic Markov chain is only reversible with respect to the stationary distribution. So above $\pi(x)$ is the stationary distribution of P . An immediate consequence of this is that for a chain with uniform stationary distribution, it is reversible if and only if it is symmetric (i.e. $P(x, y) = P(y, x)$). Note also that reversible chains have real eigenvalues, since they are similar to the symmetric matrix $\sqrt{\frac{\pi(x)}{\pi(y)}}P(x, y)$ (using the similarity transform $\delta_{xy}\sqrt{\pi(x)}$).

With these definitions and concepts, we can now ask how quickly the Markov chain converges to the stationary distribution. This is normally defined in terms of the 1-norm mixing time. We use (half the) 1-norm distance to measure distances between distributions:

$$\|s - t\| = \frac{1}{2} \|s - t\|_1 = \frac{1}{2} \sum_i |s_i - t_i|. \quad (3.5.9)$$

We assume all distributions are normalised so then $0 \leq \|s - t\| \leq 1$. We can now define the mixing time:

Definition 3.5.3. *Let π be the stationary distribution of P . Then if P is ergodic the mixing time τ is*

$$\tau(\epsilon) = \max_s \min_t \{t \geq 0 : \|P^t s - \pi\| \leq \epsilon\}. \quad (3.5.10)$$

We will also use the (weaker) 2-norm mixing time (note this is not the same as τ_2 in [MT06]):

Definition 3.5.4. *Let π be the stationary distribution of P . Then if P is ergodic the 2-norm mixing time τ_2 is*

$$\tau_2(\epsilon) = \max_s \min_t \{t \geq 0 : \|P^t s - \pi\|_2 \leq \epsilon\}. \quad (3.5.11)$$

Unless otherwise stated, when we say mixing time we are referring to the 1-norm mixing time.

There are many techniques for bounding the mixing time, including finding the second largest eigenvalue of P . This gives a good measure of the mixing time because

components parallel to the second largest eigenvector decay the slowest. We have (for reversible ergodic chains)

Theorem 3.5.5 (see [MT06], Corollary 1.15).

$$\tau(\epsilon) \leq \frac{1}{\Delta} \ln \frac{1}{\pi_* \epsilon}$$

where $\pi_* = \min \pi(x)$ and $\Delta = \min(1 - \lambda_2, 1 + \lambda_{min})$ where λ_2 is the second largest eigenvalue and λ_{min} is the smallest. Δ is known as the gap.

If the chain is irreversible, it may not even have real eigenvalues. However, we can bound the mixing time in terms of the eigenvalues of the reversible matrix PP^* where $P^*(x, y) = \frac{\pi(y)}{\pi(x)}P(y, x)$. In this case we have ([MT06], Corollary 1.14)

$$\tau(\epsilon) \leq \frac{2}{\Delta_{PP^*}} \ln \frac{1}{\pi_* \epsilon} \tag{3.5.12}$$

where now Δ_{PP^*} is the gap of the chain PP^* . Note that for a reversible chain $P = P^*$ and $\Delta_{PP^*} \approx 2\Delta$ so the bounds are approximately the same.

This can also be converted into a 2-norm mixing time bound:

$$\tau_2(\epsilon) \leq \frac{2}{\Delta_{PP^*}} \ln 1/\epsilon. \tag{3.5.13}$$

To bound the gap, we will use the comparison theorem in Theorem 3.5.6 below. In this Theorem, we are thinking of the Markov chain as a directed graph where the vertices are the states and there are edges for allowed transitions (i.e. transitions with non-zero probability). For irreducible chains, it is possible to make a path from any vertex to any other; we call the path length the number of transitions in such a path (which will in general depend on the choice of path).

Theorem 3.5.6 (see [MT06], Theorem 2.14). *Let P and \hat{P} be two Markov chains on the same state space Ω with the same stationary distribution π . Then, for every $x \neq y \in \Omega$ with $\hat{P}(x, y) > 0$ define a directed path γ_{xy} from x to y along edges in P*

and let its length be $|\gamma_{xy}|$. Let Γ be the set of all such paths. Then

$$\Delta \geq \hat{\Delta}/A$$

for the gaps Δ and $\hat{\Delta}$ where

$$A = A(\Gamma) = \max_{a \neq b, P(a,b) \neq 0} \frac{1}{\pi(a)P(a,b)} \sum_{x \neq y: (a,b) \in \gamma_{xy}} \pi(x) \hat{P}(x,y) |\gamma_{xy}|.$$

For example, when comparing 1-dimensional random walks there is no choice in the paths; they must pass through every point between x and y . Further, the walk can only progress one step at a time so (without loss of generality, for reversible chains) let $b = a + 1$ to give

$$\begin{aligned} A &= \max_a \frac{1}{\pi(a)P(a, a+1)} \sum_{x \leq a} \sum_{y \geq a+1} \pi(x) \hat{P}(x,y) (y-x) \\ &= \max_a \frac{\hat{P}(a, a+1)}{P(a, a+1)}. \end{aligned} \tag{3.5.14}$$

A generalisation of the comparison theorem involves constructing flows, which are weighted sets of paths between states. This can give a tighter bound since bottlenecks are averaged over. This gives a modified comparison theorem:

Theorem 3.5.7 ([DS93], Theorem 2.3). *Let P and \hat{P} be two Markov chains on the same state space Ω with the same stationary distribution π . Then, for every $x \neq y \in \Omega$ with $\hat{P}(x,y) > 0$, construct a set of directed paths \mathcal{P}_{xy} from x to y along edges in P . We define the flow function f which maps each path $\gamma_{xy} \in \mathcal{P}_{xy}$ to a real number in the interval $[0, 1]$ such that*

$$\sum_{\gamma_{xy} \in \mathcal{P}_{xy}} f(\gamma_{xy}) = \hat{P}(x,y).$$

Again, let the length of each path be $|\gamma_{xy}|$. Then

$$\Delta \geq \hat{\Delta}/A$$

for the gaps Δ and $\hat{\Delta}$ where

$$A = A(f) = \max_{a \neq b, P(a,b) \neq 0} \frac{1}{\pi(a)P(a,b)} \sum_{x \neq y, \gamma_{xy} \in \mathcal{P}_{xy}: (a,b) \in \gamma_{xy}} \pi(x) f(\gamma_{xy}) |\gamma_{xy}|. \quad (3.5.15)$$

Note that we recover the comparison theorem when there is just one path between each x and y .

Yet another generalisation is to allow general length functions instead of simply counting the edges. This only appears in the literature as a comparison to the chain $\hat{P}(x,y) = \pi(y)$ although it can easily be generalised to allow comparison with any chain.

Theorem 3.5.8 ([Kah96], Proposition 1). *Let P be a Markov chain on the state space Ω with stationary distribution π . Then, for every $x \neq y \in \Omega$ define a directed path γ_{xy} from x to y along edges in P and let its length be*

$$|\gamma_{xy}|_l = \sum_{(a,b) \in \gamma_{xy}} l(a,b) \quad (3.5.16)$$

for any positive length function $l(a,b)$, defined on the edges of the path. Let Γ be the set of all such paths. Then

$$\Delta \geq 1/A$$

where

$$A = A(\Gamma) = \max_{a \neq b, P(a,b) \neq 0} \frac{1}{l(a,b)\pi(a)P(a,b)} \sum_{x \neq y: (a,b) \in \gamma_{xy}} \pi(x) \hat{P}(x,y) |\gamma_{xy}|_l.$$

Decomposition

For some Markov chains, it is easier to consider different parts of the chain separately to prove convergence results. This allows, for example, different convergence techniques to be used on different parts of the chain. The separate parts are combined using the decomposition theorem:

Theorem 3.5.9 ([MR00], Theorem 4.2). *Let $P(x, y)$ be the transition matrix for a reversible Markov chain with state space Ω and stationary distribution $\pi(x)$. Then let Ω_i be disjoint subsets of Ω such that $\cup_i \Omega_i = \Omega$. Let*

$$P_i(x, y) = \begin{cases} P(x, y) & x, y \in \Omega_i, x \neq y \\ 1 - \sum_{y' \in \Omega_i, y' \neq x} P(x, y') & x = y \in \Omega_i \\ 0 & \text{otherwise.} \end{cases} \quad (3.5.17)$$

Further let $w_i = \sum_{x \in \Omega_i} \pi(x)$ and

$$\bar{P}(i, j) = \frac{1}{w_i} \sum_{x \in \Omega_i, y \in \Omega_j} \pi(x) P(x, y). \quad (3.5.18)$$

Then

$$\Delta \geq \frac{1}{2} \bar{\Delta} \min_i \Delta_i \quad (3.5.19)$$

where Δ is the gap of P , $\bar{\Delta}$ for \bar{P} and Δ_i for P_i .

log-Sobolev Constant

We will need tighter, but more complicated, mixing time results to prove the tight result for the $U(4)$ case. We use the log-Sobolev constant:

Definition 3.5.10. *The log-Sobolev constant ρ of a chain with transition matrix P and stationary distribution π is*

$$\rho = \min_f \frac{\sum_{x \neq y} (f(x) - f(y))^2 P(x, y) \pi(y)}{\sum_x \pi(x) f(x)^2 \log \frac{f(x)^2}{\sum_y \pi(y) f(y)^2}}.$$

The mixing time result is:

Lemma 3.5.11 (see [DS96], Theorem 3.7'). *The mixing time of a finite, reversible, irreducible Markov chain is*

$$\tau(\epsilon) = O\left(\frac{1}{\rho} \log \log \frac{1}{\pi_*} + \frac{1}{\Delta} \log \frac{d}{\epsilon}\right) \quad (3.5.20)$$

where ρ is the Sobolev constant, π_* is the smallest value of the stationary distribution, Δ is the gap and d is the size of the state space.

Further, the comparison theorem (Theorem 3.5.6) works just the same to give

$$\rho \geq \hat{\rho}/A.$$

We will need one more result, due to Diaconis and Saloff-Coste:

Lemma 3.5.12 ([DS96], Lemma 3.2). *Let P_i , $i = 1, \dots, d$, be Markov chains with gaps Δ_i and Sobolev constants ρ_i . Now construct the product chain P . This chain has state space equal to the product of the spaces for the chains P_i and at each step one of the chains is chosen at random and run for one step. Then P has spectral gap given by:*

$$\Delta = \frac{1}{d} \min_i \Delta_i$$

and Sobolev constant:

$$\rho = \frac{1}{d} \min_i \rho_i.$$

3.5.4 Convergence Proof

We now prove the Markov chain convergence results to show that the $\gamma(p, p)$ terms converge quickly. We have already shown that the $\gamma(p_1, p_2)$ terms with $p_1 \neq p_2$ converge quickly and that there is no mixing between these terms and the $\gamma(p, p)$ terms. Therefore, in this section, we remove such terms from \hat{G} .

We want to prove the Markov chain with transition matrix (Eqn. 3.5.2)

$$P = \frac{1}{n(n-1)} \sum_{i \neq j} \hat{G}^{(ij)}$$

converges quickly. Firstly, we know from Section 3.4.3 that P has two eigenvectors with eigenvalue 1. The first is the identity state ($\sigma_0 \otimes \sigma_0$) and the second is the uniform sum of all non-identity terms ($\frac{1}{4^n-1} \sum_{p \neq 0} \sigma_p \otimes \sigma_p$). From now on, we remove the identity state. This makes the chain irreducible. Since we know it converges, it must be aperiodic also so the chain is ergodic and all other eigenvalues are strictly between 1 and -1 .

We show here that the gap of this chain, up to constants, does not depend on the choice of 2-copy gapped gate set. In the second half of the chapter we find a tight bound on the gap for the $U(4)$ case which consequently gives a tight bound on the gap for all universal sets.

Since the stationary distribution is uniform, the chain is reversible if and only if P is a symmetric matrix. A sufficient condition for P to be symmetric is for $\hat{G}^{(ij)}$ to be symmetric. We saw in Theorem 3.4.3 that for the $U(4)$ gate set case $\hat{G}^{(ij)}$ is symmetric. In fact, the proof works identically to show that $\hat{G}^{(ij)}$ is symmetric for any gate set, provided the set is invariant under Hermitian conjugation. However, 2-copy gapped gate sets do not necessarily have this property so the Markov chain is not necessarily reversible. We will find equal bounds (up to constants) for the gaps of both P (if \hat{G} is symmetric) and PP^* (if \hat{G} is not symmetric) below:

Theorem 3.5.13. *Let μ be any 2-copy gapped distribution of gates. If μ is invariant under Hermitian conjugation then let Δ_P be the eigenvalue gap of the resulting Markov chain matrix P . Then*

$$\Delta_P = \Omega(\Delta_{U(4)}) \tag{3.5.21}$$

where $\Delta_{U(4)}$ is the eigenvalue gap of the $U(4)$ chain. If μ is not invariant under Hermitian conjugation then let Δ_{PP^} be the eigenvalue gap of the resulting Markov*

chain matrix PP^* . Then

$$\Delta_{PP^*} = \Omega(\Delta_{U(4)}). \quad (3.5.22)$$

Proof. We will use the comparison method with flows (Theorem 3.5.7). Firstly consider the case where μ is closed under Hermitian conjugation i.e. \hat{G} is symmetric.

We will compare P to the $U(4)$ chain, which we call $P_{U(4)}$. Recall that this chain chooses a pair at random and does nothing if the pair is 00 and chooses a random state from $\{0, 1, 2, 3\}^2 \setminus \{00\}$ otherwise.

To apply Theorem 3.5.7, we need to construct the flows between transitions in $P_{U(4)}$. We will choose paths such that only one pair is modified throughout. For example (with $n = 4$), the transition $1000 \rightarrow 2000$ is allowed in $P_{U(4)}$. To construct a path in P , we need to find allowed transitions between these two paths in P . \hat{G} may not include the transition $10 \rightarrow 20$ directly, however, \hat{G} is irreducible on this subspace of just two pairs. This means that a path exists and can be of maximum length 14 if it has to cycle through all intermediate states (in fact, since \hat{G} is symmetric the maximum path length is 8; all that is important here is that it is constant). For example, the transitions $10 \rightarrow 11 \rightarrow 20$ might be allowed. Then we could choose the full path to be $1000 \rightarrow 1100 \rightarrow 2000$. In this case we have chosen the path to involve transitions pairing sites 1 and 2. However, we could equally well have chosen any pairing; we could pair the first site with any of the others. We can choose 3 paths in this way. For this example, the flow we want to choose will be all 3 of these paths equally weighted. We now use this idea to construct flows between all transitions in $P_{U(4)}$ to prove the result.

Let $x \neq y \in \Omega$ and let $d(x, y)$ be the Hamming distance between the states ($d(x, y)$ gives the number of places at which x and y differ). There are two cases where $P_{U(4)}(x, y) \neq 0$:

1. $d(x, y) = 2$. Here we must choose a unique pairing, specified by the two sites that differ. Make all transitions in P using this pair giving just one path.
2. $d(x, y) = 1$. For this case, choose all possible pairings of the changing site

that give allowed transitions in $P_{U(4)}$. For each pairing, construct a path in P modifying only this pair. If the differing site is initially non-zero then there are $n - 1$ such pairings; if the differing site is initially zero then there are $n - z(x)$ pairings where $z(x)$ is the number of zeroes in the state x .

All the above paths are of constant length since we have to (at most) cycle through all states of a pair. We must now choose the weighting $f(\gamma_{xy})$ for each path such that

$$\sum_{\mathcal{P}_{xy}} f(\gamma_{xy}) = P_{U(4)}(x, y) \quad (3.5.23)$$

where \mathcal{P}_{xy} is the set of all paths from x to y constructed above. We choose the weighting of each path to be uniform. We just need to calculate the number of paths in \mathcal{P}_{xy} to find f :

1. $d(x, y) = 2$. There is just one path so $f(\gamma_{xy}) = P_{U(4)}(x, y) = \Theta(1/n^2)$.
2. $d(x, y) = 1$. If the differing site is initially non-zero then $P_{U(4)}(x, y) = \Theta(1/n)$ and there are $n - 1$ paths so $f(\gamma_{xy}) = \frac{P_{U(4)}(x, y)}{n-1} = \Theta(1/n^2)$. If the differing site is initially zero then $P_{U(4)}(x, y) = \Theta\left(\frac{n-z(x)}{n^2}\right)$ and there are $n - z(x)$ paths so $f(\gamma_{xy}) = \frac{P_{U(4)}(x, y)}{n-z(x)} = \Theta(1/n^2)$.

So for all paths, $f = \Theta(1/n^2)$. We now just need to know how many times each edge (a, b) in P is used to calculate A :

$$A = \max_{a \neq b, P(a, b) \neq 0} A(a, b) \quad (3.5.24)$$

where

$$A(a, b) = \frac{1}{P(a, b)} \sum_{x \neq y, \gamma_{xy} \in \mathcal{P}_{xy}: (a, b) \in \gamma_{xy}} f(\gamma_{xy}). \quad (3.5.25)$$

We have cancelled the factors of $\pi(x)$ because the stationary distribution is uniform.

We have also ignored the lengths of the paths since they are all constant.

To evaluate $A(a, b)$, we need to know how many paths pass through each edge (a, b) . We again consider the two possibilities separately:

1. $d(a, b) = 2$. Suppose a and b differ at sites i and j . Firstly, we need to count how many transitions from x to y in $P_{U(4)}$ could use this edge, and then how many paths for each transition actually use the edge.

To find which x and y could use the edge, note that x and y must differ at sites i, j or both. Furthermore, the values at the sites other than i and j must be the same as for a (and therefore b). There is a constant number of x, y pairs that satisfy this condition. Now, for each x, y pair satisfying this, paths that use this edge must use the pairing i, j for all transitions. Since in the paths we have chosen above there is a unique path from x to y for each pairing, there is at most one path for each x, y pair that uses edge a, b .

For $d(a, b) = 2$, $P(a, b) = \Theta(1/n^2)$ so $A(a, b)$ is a constant for this case.

2. $d(a, b) = 1$. Let there be r pairings that give allowed transitions in P between a and b . As above, each pairing gives a constant number of paths. So the numerator is $\Theta(r/n^2)$. Further, $P(a, b) = \Theta(r/n^2)$. So again $A(a, b)$ is constant.

Combining, A is a constant so the result is proven for the case \hat{G} is symmetric.

We now turn to the irreversible case. We now need to bound the gap of $PP^* = PP^T$. This chain selects two (possibly overlapping) pairs at random and applies \hat{G} to one of them and \hat{G}^T to the other. We can use the above exactly by choosing \hat{G} to perform the transitions above and \hat{G}^T to just loop the states back to themselves. By aperiodicity (the greatest common divisor of loop lengths is 1), we can always find constant length paths that do this. \square

Now we need to know the gap of the $U(4)$ chain. We can, by a simple application of the comparison theorem, show it is $\Omega(1/n^2)$. However, in the second half of this chapter we show it is $\Theta(1/n)$. This gives us (using Theorem 3.5.5):

Corollary 3.5.14. *The Markov chain P has mixing time $O(n(n + \log 1/\epsilon))$ and 2-norm mixing time $O(n \log 1/\epsilon)$.*

We conjecture that the mixing time (as well as Lemma 3.5.2) can be tightened to $\Theta(n \log \frac{n}{\epsilon})$, which is asymptotically the same as for the $U(4)$ case:

Conjecture 3.5.15. *The second moments for the case of general 2-copy gapped distributions have 1-norm mixing time $\Theta(n \log \frac{n}{\epsilon})$.*

It seems likely that an extension of our techniques in Section 3.6 could be used to prove this.

Combining the convergence results we have proved our general result Lemma 3.3.2:

Proof of Lemma 3.3.2. Combining Corollary 3.5.14 (for the $\gamma(p, p)$ terms) and Lemma 3.5.2 (for the $\gamma(p_1, p_2)$, $p_1 \neq p_2$ terms) proves the result. \square

We have now shown that the first and second moments of random circuits converge quickly. For the remainder of the chapter we prove the tight bound for the gap and mixing time of the $U(4)$ case and show how mixing time bounds relate to the closeness of the 2-design to an exact design. Only for the $U(4)$ case is the matrix \hat{G} a projector so in this sense the $U(4)$ random circuit is the most fundamental. While we expect the above mixing time bound is not tight, we can prove a tight mixing time result for the $U(4)$ case. However, using our definition of an approximate k -design, the gap rather than the mixing time governs the degree of approximation.

3.6 Tight Analysis for the $U(4)$ Case

We have already found tight bounds for the first moments in Lemma 3.5.1: just set $\Delta = 1$.

3.6.1 Second Moments Convergence

We need to prove a result analogous to Lemma 3.5.2 for the terms $\sigma_{p_1} \otimes \sigma_{p_2}$ where $p_1 \neq p_2$. We already have a tight bound for the 2-norm decay, by setting $\Delta = 1$ into Lemma 3.5.2. We tighten the 1-norm bound:

Lemma 3.6.1. *After $O(n \log \frac{n}{\epsilon})$ steps*

$$\sum_{p_1 \neq p_2} \mathbb{E}_W |\gamma_W(p_1, p_2)| \leq \epsilon \tag{3.6.1}$$

Proof. We will split the random circuits up into classes depending on how many qubits have been hit. Let H be the random variable giving the number of different qubits that have been hit. We can work out the distribution of H and bound the sum of $|\gamma_W(p_1, p_2)|$ for each outcome.

Firstly we have, after t steps,

$$\mathbb{P}(H \leq h) \leq \binom{n}{h} \left(\frac{h(h-1)}{n(n-1)} \right)^t \leq \binom{n}{h} (h/n)^t.$$

Now, for each qubit hit, each coefficient which has p_1 and p_2 differing in this place is set to zero. So after h have been hit, there are only (at most) $16^{(n-h)}$ terms in the sum in Eqn. 3.6.1. As before, the state is a physical state, $\text{tr } \rho^2 \leq 1$ so $\sum_{p_1 p_2} \gamma^2(p_1, p_2) \leq 1$ so $\sum_{p_1 p_2} |\gamma(p_1, p_2)| \leq \sqrt{N}$ if there are at most N non-zero terms in the sum. Therefore we have, after t steps,

$$\begin{aligned} \sum_{p_1 \neq p_2} \mathbb{E}_W |\gamma_W(p_1, p_2)| &\leq \sum_{h=1}^{n-1} \mathbb{P}(H = h) 16^{(n-h)/2} \\ &\leq \sum_{h=1}^{n-1} \mathbb{P}(H \leq h) 4^{(n-h)} \\ &\leq \sum_{h=1}^{n-1} \binom{n}{h} (h/n)^t 4^{(n-h)} \\ &= \sum_{h=1}^{n-1} \binom{n}{h} (1 - h/n)^t 4^h \quad h \rightarrow n - h \\ &\leq \sum_{h=1}^{n-1} \binom{n}{h} \exp(-ht/n) 4^h. \end{aligned}$$

Now, let $t = n \ln \frac{n}{\epsilon}$:

$$\begin{aligned} \sum_{p_1 \neq p_2} \mathbb{E}_W |\gamma_W(p_1, p_2)| &\leq \sum_{h=1}^{n-1} \binom{n}{h} \left(\frac{4\epsilon}{n} \right)^h \\ &= \left(1 + \frac{4\epsilon}{n} \right)^n - 1 - \left(\frac{4\epsilon}{n} \right)^n = O(\epsilon) \end{aligned}$$

where the last line follows from the binomial theorem. \square

This, combined with the mixing time result we prove below, completes the proof that the second moments of the random circuit converge in time $O(n \log \frac{n}{\epsilon})$.

3.6.2 Markov Chain of Coefficients

The Markov chain acting on the coefficients is reducible because the state $\{0\}^n$ is isolated. However, if we remove it then the chain becomes irreducible. The presence of self loops implies aperiodicity therefore the chain is ergodic. We have already seen that the chain converges to the Haar uniform distribution (in Section 3.1.1) therefore the stationary state is the uniform state $\pi(x) = 1/(4^n - 1)$. Further, since the chain is symmetric and has uniform stationary distribution, the chain satisfies detailed balance (Eqn. 3.5.8) so is reversible. We now turn to obtaining bounds on the mixing time of this chain.

We want to show that the full chain converges to stationarity in time $\Theta(n \log \frac{n}{\epsilon})$. To prove this, we will construct another chain called the zero chain. This is the chain that counts the number of zeroes in the state. Since it is the zeroes that slow down the mixing, this chain will accurately describe the mixing time of the full chain.

Lemma 3.6.2. *The zero chain has transition matrix P on state space (we count non-zero positions) $\Omega = \{1, 2, \dots, n\}$.*

$$P(x, y) = \begin{cases} 1 - \frac{2x(3n-2x-1)}{5n(n-1)} & y = x \\ \frac{2x(x-1)}{5n(n-1)} & y = x - 1 \\ \frac{6x(n-x)}{5n(n-1)} & y = x + 1 \\ 0 & \text{otherwise} \end{cases} \quad (3.6.2)$$

for $1 \leq x, y \leq n$.

Proof. Suppose there are $n - x$ zeroes (so there are x non-zeroes). Then the only way the number of zeroes can decrease (i.e. for x to increase) is if a non-zero item is paired

with a zero item and one of the 9 (out of 15) new states is chosen with no zeroes. The probability of choosing such a pair is $\frac{2x(n-x)}{n(n-1)}$ so the overall probability is $\frac{9}{15} \frac{2x(n-x)}{n(n-1)}$.

The number of zeroes can increase only if a pair of non-zero items is chosen and one of the 6 states is chosen with one zero. The probability of this occurring is $\frac{6}{15} \frac{x(x-1)}{n(n-1)}$.

The probability of the number of zeroes remaining unchanged is simply calculated by requiring the probabilities to sum to 1. \square

We see that the zero chain is a one-dimensional random walk on the line. It is a lazy random walk because the probability of moving at each step is < 1 . However, as the number of zeroes decreases, the probability of moving increases monotonically:

$$1 - P(x, x) = \frac{2x(3n - 2x - 1)}{5n(n - 1)} \geq 2x/5n. \quad (3.6.3)$$

Lemma 3.6.3. *The stationary distribution of the zero chain is*

$$\pi_0(x) = \frac{3^x \binom{n}{x}}{4^n - 1}. \quad (3.6.4)$$

Proof. This can be proven by multiplying the transition matrix in Lemma 3.6.2 by the state Eqn. 3.6.4. Alternatively, it can be proven by counting the number of states with $n - x$ zeroes. There are $\binom{n}{x}$ ways of choosing which sites to make non-zero and each non-zero site can be one of three possibilities: 1, 2 or 3. The total number of states is $4^n - 1$, which gives the result. \square

Below we will prove the following theorem:

Theorem 3.6.4. *The zero chain mixes in time $\Theta(n \log \frac{n}{\epsilon})$.*

We prove this using direct arguments about the convergence of the random walk. However, we also include a less complex method that only bounds the gap:

Theorem 3.6.5. *The zero chain has gap $\Omega(1/n)$.*

This only implies the mixing time is $O(n(n + \log 1/\epsilon))$ which is weaker than Theorem 3.6.4, although still sufficient to prove our main result Theorem 3.3.1,

using a modification of Corollary 3.6.7 to show that the full chain mixing time is $O(n(n + \log 1/\epsilon))$.

Knowing the gap allows us to easily work out the 2-norm mixing time:

Theorem 3.6.6. *The zero chain has 2-norm mixing time $O(n \log 1/\epsilon)$.*

Proof. Use the bound on the gap in Theorem 3.6.5 and Eqn. 3.5.13. \square

Before proving Theorem 3.6.4, we will show how the mixing time of the full chain follows from this.

Corollary 3.6.7. *The full chain mixes in time $\Theta(n \log \frac{n}{\epsilon})$.*

Proof. Once the zero chain has approximately mixed, the distribution of zeroes is almost correct. We need to prove that the distribution of non-zeroes is correct after $O(n \log \frac{n}{\epsilon})$ steps too.

Once each site of the full chain has been hit, meaning it is chosen and paired with another site so not both equal zero, the chain has mixed. This is because, after each site has been hit, the probability distribution over the states is uniform. When the zero chain has approximately mixed, a constant fraction of sites are zero so the probability of hitting a site at each step is $\Theta(1/n)$. By the coupon collector argument, each site will have been hit with probability at least $1 - \epsilon$ in time $O(n \log \frac{n}{\epsilon})$. Once the zero chain has mixed to ϵ' , we can run the full chain this extra number of steps to ensure each site has been hit with high probability. Since the mixing of the zero chain only increases with time, the distance to stationarity of the full chain is now $1 - \epsilon - \epsilon'$. We make this formal below.

After $t_0 = O(n \log \frac{n}{\epsilon'})$ steps, the number of zeroes is ϵ' -close to the stationary distribution π_0 by Theorem 3.6.4 and only gets closer with more steps since the distance to stationarity decreases monotonically. The stationary distribution Eqn. 3.6.4 is approximately a Gaussian peaked at $3n/4$ with $O(n)$ variance. This means that, with high probability, the number of non-zeroes is close to $3n/4$. We will in fact only need that there is at least a constant fraction of non-zeroes; with probability at least $1 - \epsilon' - \exp(-\Omega(n))$ there will be at least $n/2$.

To prove the mixing time, we run the chain for time t_0 so the zero chain mixes to ϵ' . Then run for t_1 additional steps. Let $H_{i,t}$ be the event that site i is hit at step t . Let $H_i = \cup_{t=t_0+1}^{t_0+t_1} H_{i,t}$ and $H = \cap_{i=1}^n H_i$. We want to show $\mathbb{P}(H)$ is close to 1, or, in other words, that all sites are hit with high probability. Further let X_t be the random variable giving the number of non-zeroes at step t .

If at step $t-1$ site i is non-zero then the event $H_{i,t}$ occurs if the qubit is chosen, which occurs with probability $2/n$. If, however, it was zero then it must be paired with a non-zero thing for $H_{i,t}$ to hold. Conditioned on any history with $X_{t-1} \geq n/2$, this probability is $\geq 1/n$. In particular, we can condition on not having previously hit i and the bound does not change. Combining we have

$$\mathbb{P} \left(H_{i,t}^c \mid [X_{t-1} \geq n/2] \cap \left(\bigcap_{t'=t_0+1}^{t-1} H_{i,t'}^c \right) \right) \leq 1 - 1/n.$$

Then, after t_1 extra steps,

$$\mathbb{P} \left(H_i^c \mid \bigcap_{t=t_0}^{t_0+t_1-1} [X_t \geq n/2] \right) \leq (1 - 1/n)^{t_1}$$

which, using the union bound, gives

$$\mathbb{P} \left(H^c \mid \bigcap_{t=t_0}^{t_0+t_1-1} [X_t \geq n/2] \right) \leq n(1 - 1/n)^{t_1}.$$

Now, since the zero chain has mixed to ϵ' ,

$$\mathbb{P} \left(\bigcap_{t=t_0}^{t_0+t_1-1} [X_t \geq n/2] \right) \leq t_1 \left(\sum_{x=n/2}^{n-1} \pi_0(x) + \epsilon' \right) \leq t_1 (\exp(-O(n)) + \epsilon')$$

so

$$\mathbb{P}(H^c) \leq n(1 - 1/n)^{t_1} + t_1 (\exp(-O(n)) + \epsilon').$$

Now, choose $t_1 = n \ln \frac{2n}{\epsilon}$ so that $\mathbb{P}(H^c) \leq \delta$ where $\delta = \epsilon + t_1 (\exp(-O(n)) + \epsilon')$. Choose $\epsilon = 1/n$ and $\epsilon' = 1/n^3$ so that δ is $1/\text{poly}(n)$. Now, using the bound on

$\mathbb{P}(H^c)$, we can write the state v after $t_1 = O(n \log n)$ steps as

$$v = (1 - \delta)\pi + \delta\pi'$$

where π is the stationary distribution and π' is any other distribution. Using this,

$$\|v - \pi\| \leq \delta.$$

We now apply Lemma 3.9.13 to show that after $O(n \log \frac{n}{\epsilon})$ steps the distance to stationarity of the full chain is ϵ . \square

3.6.3 Proof of Theorem 3.6.4

We will now proceed to prove Theorem 3.6.4. We present an outline of the proof here; the details are in Section 3.9.1.

Firstly, note that by the coupon collector argument, the lower bound on the time is $\Omega(n \log n)$. We need to prove an upper bound equal to this. Intuition says that the mixing time should take time $O(n \log n)$ because the walk has to move a distance $\Theta(n)$ and the waiting time at each step is proportional to $n, n/2, n/3, \dots$ which sums to $O(n \log n)$, provided each site is not hit too often. We will show that this intuition is correct using Chernoff bound and log-Sobolev (see later) arguments.

We will first work out concentration results of the position after some number of *accelerated* steps. The zero chain has some probability of staying still at each step. The accelerated chain is the zero chain conditioned on moving at each step. We define the accelerated chain by its transition matrix:

Definition 3.6.8. *The transition matrix for the accelerated chain is*

$$P_a(x, y) = \begin{cases} 0 & y = x \\ \frac{x-1}{3n-2x-1} & y = x - 1 \\ \frac{3(n-x)}{3n-2x-1} & y = x + 1 \\ 0 & \text{otherwise.} \end{cases} \quad (3.6.5)$$

We use the accelerated chain in the proof to firstly prove the accelerated chain mixes quickly, then to bound the waiting time at each step to obtain a mixing time bound for the zero chain.

To prove the mixing time bound, we will split the walk up into three phases. We will split the state space into three (slightly overlapping) parts and the phase can begin at any point within that space. So each phase has a state space $\Omega_i \subset [1, n]$, an entry space $E_i \subset \Omega_i$ and an exit condition T_i . We say that a phase completes successfully if the exit condition is satisfied in time $O(n \log n)$ for an initial state within the entry space. When the exit condition is satisfied, the walk moves onto the next phase.

The phases are:

1. $\Omega_1 = [1, n^\delta]$ for some constant δ with $0 < \delta < 1/2$. $E_1 = \Omega_1$ (i.e. it can start anywhere) and T_1 is satisfied when the walk reaches n^δ . For this part, the probability of moving backwards (gaining zeroes) is $O(n^{\delta-1})$ so the walk progresses forwards at each step with high probability. This is proven in Lemma 3.9.6. We show that the waiting time is $O(n \log n)$ in Lemma 3.9.7.
2. $\Omega_2 = [n^\delta/2, \theta n]$ for some constant θ with $0 < \theta < 3/4$. $E_2 = [n^\delta, \theta n]$ and T_2 is satisfied when the walk reaches θn . Here the walk can move both ways with constant probability but there is a $\Omega(1)$ forward bias. Here we use a

monotonicity argument: the probability of moving forward at each step is

$$\begin{aligned} p(x) &= \frac{3(n-x)}{3n-2x-1} \\ &\geq \frac{3(n-x)}{3n-2x} \\ &\geq \frac{3(1-\theta)}{3-2\theta}. \end{aligned}$$

If we model this random walk as a walk with constant bias equal to $\frac{3(1-\theta)}{3-2\theta}$ we will find an upper bound on the mixing time since mixing time increases monotonically with decreasing bias. Further, the waiting time at $x = a$ stochastically dominates the waiting time at $x = b$ for $b \geq a$. The true bias decreases with position so the walk with constant bias spends more time at the early steps. Thus the position of this simplified walk is stochastically dominated by the position of the real walk while the waiting time stochastically dominates the waiting time of the real walk.

3. $\Omega_3 = [\frac{\theta}{2}n, n]$ and $E_3 = [\theta n, n]$. T_3 is satisfied when this restricted part of the chain has mixed to distance ϵ . Here the bias decreases to zero as the walk approaches $3n/4$ but the moving probability is a constant. We show that this walk mixes quickly by bounding the log-Sobolev constant of the chain.

Showing these three phases complete successfully will give a mixing time bound for the whole chain.

We now prove in Section 3.9 that the phases complete successfully with probability at least $1 - 1/\text{poly}(n)$:

Lemma 3.6.9.

$$\mathbb{P}(\text{Phase 1 completes successfully}) \geq 1 - n^{2\delta-1} - 2n^{-\delta}$$

Lemma 3.6.10.

$\mathbb{P}(\text{Phase 2 completes successfully}) \geq$

$$1 - \exp\left(-\frac{2}{3}\mu\theta n\right) - \left(\frac{4}{\theta n}\right)^{\frac{3}{2\mu}} - \frac{2 \exp\left(\frac{-\mu n^\delta}{4}\right)}{1 - \exp(-\mu/2)} - (q/p)^{n^\delta/2}$$

where $\mu = \frac{6(1-\theta)}{3-2\theta} - 1$.

Lemma 3.6.11.

$$\mathbb{P}(\text{Phase 3 completes successfully}) \geq 1 - \left(\frac{\theta}{3(2-\theta)}\right)^{\theta n/2}$$

We can now finally combine to prove our result:

Proof of Theorem 3.6.4. The stationary distribution has exponentially small weight in the tail with lots of zeroes. We show that, provided the number of zeroes is within phase 3, the walk mixes in time $O(n \log \frac{n}{\epsilon})$. We also show that if the number of zeroes is initially within phase 1 or 2, after $O(n \log n)$ steps the walk is in phase 3 with high probability. We can work out the distance to the stationary distribution as follows.

Let p_f be the probability of failure. This is the sum of the error probabilities in Lemmas 3.6.9, 3.6.10 and 3.6.11. The key point is that $p_f = 1/\text{poly}(n)$. Then after $O(n \log \frac{n}{\epsilon})$ steps (the sum of the number of steps in the 3 phases), the state is equal to $(1 - p_f)v_3 + p_f v'$ where v_3 is the state in the phase 3 space and v' is any other distribution, which occurs if any one of the phases fails. Since the distance to stationarity in phase 3 is ϵ , $\|v_3 - \pi_3\| \leq \epsilon$, where π_3 is the stationary distribution on the state space of phase 3. In Lemma 3.9.11 we show that $\pi_3(x) = \pi(x)/(1-w)$ where $w = \sum_{x=1}^{\theta n/2-1} \pi(x)$. Since $\pi(x)$ is exponentially small in this range, w is exponentially small in n . Now use the triangle inequality to find

$$\|v_3 - \pi\| \leq \|v_3 - \pi_3\| + \|\pi_3 - \pi\|. \tag{3.6.6}$$

Since the chain in phase 3 has mixed to ϵ , the first term is $\leq \epsilon$. We can evaluate

$\|\pi_3 - \pi\|$:

$$\begin{aligned}
\|\pi_3 - \pi\| &= \frac{1}{2} \sum_{x=1}^n \|\pi_3(x) - \pi(x)\| \\
&= \frac{1}{2} \left(\sum_{x=1}^{\theta n/2-1} \pi(x) + \sum_{x=\theta n/2}^n (\pi(x)/(1-w) - \pi(x)) \right) \\
&= \frac{1}{2} (w + 1 - (1-w)) = w.
\end{aligned}$$

So now,

$$\begin{aligned}
\|(1-p_f)v_3 + p_f v' - \pi\| &= \|(1-p_f)(v_3 - \pi) + p_f(v' - \pi)\| \\
&\leq (1-p_f)\|v_3 - \pi\| + p_f\|v' - \pi\| \\
&\leq (1-p_f)(\epsilon + w) + p_f \\
&\leq \delta
\end{aligned}$$

where $\delta = \epsilon + w + p_f$. We are free to choose ϵ : choose it to be $1/n$ so that δ is $1/\text{poly}(n)$. So now the running time to get a distance δ is $t = O(n \log n)$. We then apply Lemma 3.9.13 to obtain the result.

This concludes the proof of Theorem 3.6.4 so Corollary 3.6.7 is proved. \square

We have now proven Lemma 3.3.2 and consequently Corollary 3.3.3. We are now ready to show how Theorem 3.3.1 follows, but first give the alternative proof that the zero chain gap is $\Omega(1/n)$. The remainder of the proof of Theorem 3.3.1 is in Section 3.7.

3.6.4 Proof of Theorem 3.6.5

Here we prove that the gap of the zero chain is $\Omega(1/n)$. While this can be deduced from Theorem 3.6.4 and provides weaker mixing time bounds, this bound on the gap is sufficient to prove our main result so we present it as a simpler alternative proof.

We use the method of decomposition (Theorem 3.5.9), whereby the Markov chain

is split up into disjoint state spaces. This works well here because, for the first part of the walk with many zeroes, the walker remains stationary most of the time whereas when there is a constant fraction of zeroes, the walker moves on most steps. Using the decomposition method allows us to use different techniques in these different regimes.

We therefore divide the walk up into two parts, P_1 and P_2 , which are shown in Figure 3.2. The chain \bar{P} is the chain that links the two parts, according to the decomposition theorem, Theorem 3.5.9.

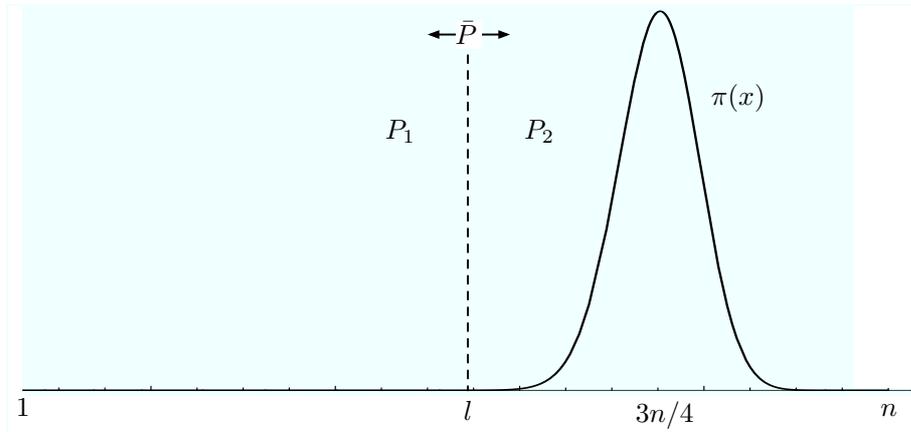


Figure 3.2: The decomposition of the zero chain into P_1 and P_2 . The graph plotted is the zero chain stationary distribution $\pi(x)$.

- P_1 : Let P_1 have state space $\Omega_1 = \{1, \dots, m\}$. This chain has transition matrix

$$P_1(x, y) = \begin{cases} 0 & x > m \text{ or } y > m \\ 1 - P(m, m-1) & x = y = m \\ P(x, y) & \text{otherwise} \end{cases} \quad (3.6.7)$$

and stationary distribution

$$\pi_1(x) = \frac{\pi(x)}{b_m} \quad (3.6.8)$$

where

$$b_m = \sum_{x=1}^m \pi(x). \quad (3.6.9)$$

- P_2 : Let P_2 be on state space $\Omega_2 = \{m + 1, \dots, n\}$. This chain has transition matrix

$$P_2(x, y) = \begin{cases} 0 & x \leq m \text{ or } y \leq m \\ 1 - P(m + 1, m + 2) & x = y = m + 1 \\ P(x, y) & \text{otherwise} \end{cases} \quad (3.6.10)$$

and stationary distribution

$$\pi_2(x) = \frac{\pi(x)}{c_m} \quad (3.6.11)$$

where

$$c_m = \sum_{x=m+1}^n \pi(x) = 1 - b_m. \quad (3.6.12)$$

We will take $m = \theta n$ where $0 < \theta < 0.49$ (we could in principle just have $\theta < 3/4$ but this restriction makes the calculations simpler; see Lemma 3.6.12 for the origin of the upper bound on θ). Note that P_2 is the same as phase 3 used in the direct mixing time proof (up to relabelling $m + 1$ to m). Therefore we already have, from the proof of Lemma 3.9.11, that the gap Δ_2 of P_2 is $\Omega(1/n)$. To find the gap of the whole zero chain we need to find the gaps Δ_1 and $\bar{\Delta}$. An ingredient to proving this is an exponential bound on the tail of the stationary distribution:

Lemma 3.6.12.

$$\pi(\theta n) \leq 2 \left(\frac{1}{4} \left(\frac{3e}{\theta} \right)^\theta \right)^n \quad (3.6.13)$$

and for $\theta < \theta_0 \approx 0.49$, $\pi(\theta n) = e^{-\Omega(n)}$.

Proof. Use $\binom{n}{k} \leq \left(\frac{ne}{k}\right)^k$ and $4^n - 1 \geq \frac{4^n}{2}$ to prove the bound. When $\frac{1}{4} \left(\frac{3e}{\theta}\right)^\theta < 1$, $\pi(\theta n)$ is exponentially small. θ_0 is the solution to $\frac{1}{4} \left(\frac{3e}{\theta}\right)^\theta = 1$. \square

From this we can bound the gap of \bar{P} :

Lemma 3.6.13. *The gap of \bar{P} is $\Omega(1/n)$.*

Proof. We first need to work out the transition matrix for \bar{P} . From the definition of \bar{P} in the decomposition theorem,

$$\begin{aligned}\bar{P}(1, 2) &= \frac{1}{b_m} \pi(m) P(m, m+1) \\ \bar{P}(2, 1) &= \frac{1}{c_m} \pi(m+1) P(m+1, m).\end{aligned}$$

We can find the two diagonal elements using the fact that the transition matrix is stochastic. Because the zero chain is reversible, \bar{P} is also reversible and by direct calculation of the eigenvalues has gap

$$\bar{\Delta} = 1 - |1 - (\bar{P}(1, 2) + \bar{P}(2, 1))|. \quad (3.6.14)$$

However, we can remove the modulus signs since, for n large enough, $\bar{P}(1, 2) + \bar{P}(2, 1) \leq 1$. This is because, using reversibility and $\pi(m) \leq b_m$

$$\bar{P}(1, 2) + \bar{P}(2, 1) = \frac{\pi(m)P(m, m+1)}{c_m b_m} \leq \frac{P(m, m+1)}{c_m}. \quad (3.6.15)$$

Using Lemma 3.6.12 we find that $c_m \geq 1 - e^{-\Omega(n)}$ for $m = \theta n$ and $\theta < \theta_0$. Using $P(m, m+1) \leq 3/5$, we find that for n large enough, $\bar{P}(1, 2) + \bar{P}(2, 1) \leq 1$.

Now we need to show that $\bar{P}(1, 2) + \bar{P}(2, 1) = \Omega(1)$. Again using Lemma 3.6.12,

we find $\bar{P}(2, 1) = e^{-\Omega(n)}$. We just need a bound on $\bar{P}(1, 2)$. First we bound $\pi(m)/b_m$:

$$\begin{aligned}
\frac{\pi(m)}{b_m} &= \frac{\pi(m)}{\sum_{x=1}^m \pi(x)} \\
&= \frac{1}{\sum_{x=1}^m \frac{\pi(x)}{\pi(m)}} \\
&= \frac{1}{\sum_{x=1}^m 3^{x-m} \frac{\binom{n}{x}}{\binom{n}{m}}} \\
&\geq \frac{1}{\sum_{x=1}^m 3^{x-m}} \\
&= \frac{1}{\sum_{x=0}^{m-1} 3^{-x}} \\
&\geq \frac{1}{\sum_{x=0}^{\infty} 3^{-x}} \\
&= \frac{2}{3}
\end{aligned}$$

using $\binom{n}{x} \leq \binom{n}{m}$ for $x \leq m \leq n/2$. For $m = \theta n$, we have $P(m, m+1) = \Omega(1)$ so overall $\bar{P}(1, 2) = \Omega(1)$, proving the bound on the gap. \square

Lemma 3.6.14. *The gap of P_1 is $\Omega(1/n)$.*

Proof. We use the comparison method with length functions as stated in Theorem 3.5.8. The length function we choose is, for $x \leq y$, $l(x, y) = r^x$ for some constant r satisfying $0 < r < 1$.

Let

$$A_z = \frac{1}{l(z, z+1)\pi_1(z)P_1(z, z+1)} \sum_{x=1}^z \sum_{y=z+1}^m \pi_1(x)\pi_1(y) \sum_{s=x}^{y-1} l(s, s+1). \quad (3.6.16)$$

Then, according to Theorem 3.5.8, $\Delta_1 \geq 1/A$ where $A = \max_z A_z$. We need to find an upper bound for A :

$$\begin{aligned}
A_z &= \frac{1}{r^z \pi(z) P(z, z+1) b_m} \sum_{x=1}^z \sum_{y=z+1}^m \pi(x) \pi(y) \frac{r^x - r^y}{1-r} \\
&= \frac{1}{r^z (1-r) \pi(z) P(z, z+1) b_m} \left((b_m - b_z) \sum_{x=1}^z r^x \pi(x) - b_z \sum_{x=z+1}^m r^x \pi(x) \right) \\
&= \frac{1}{r^z (1-r) \pi(z) P(z, z+1) b_m} \left((1 - b_z - (1 - b_m)) \sum_{x=1}^z r^x \pi(x) - (1 - (1 - b_z)) \sum_{x=z+1}^m r^x \pi(x) \right) \\
&= \frac{1}{r^z (1-r) \pi(z) P(z, z+1) b_m} \left((1 - b_z) \sum_{x=1}^m r^x \pi(x) - (1 - b_m) \sum_{x=1}^z r^x \pi(x) - \sum_{x=z+1}^m r^x \pi(x) \right) \\
&\leq \frac{1}{r^z (1-r) \pi(z) P(z, z+1) b_m} \left(\sum_{x=1}^m r^x \pi(x) - (1 - b_m) \sum_{x=1}^z r^x \pi(x) - \sum_{x=z+1}^m r^x \pi(x) \right) \\
&= \frac{\sum_{x=1}^z r^x \pi(x)}{r^z (1-r) \pi(z) P(z, z+1)}
\end{aligned}$$

where the inequality comes from $1 - b_z \leq 1$. Now let $h_z(r) = \frac{1}{r^z \pi(z)} \sum_{x=1}^z r^x \pi(x)$ then, plugging in the value of $P(z, z+1)$, we find

$$A_z \leq \frac{5n(n-1)}{6z(n-z)(1-r)} h_z(r).$$

Now, $\max_z \frac{5n(n-1)}{6z(n-z)(1-r)} = O(n)$ so showing $\max_z h_z(r) = O(1)$ is sufficient to prove the bound we require. We evaluate h_z recursively:

Firstly,

$$h_z(r) = 1 + \frac{1}{r^z \pi(z)} \sum_{x=1}^{z-1} r^x \pi(x).$$

Then evaluate the sum:

$$\begin{aligned}
\sum_{x=1}^{z-1} r^x \pi(x) &= \sum_{x=1}^{z-1} \frac{r^x \pi(x)}{r^{x+1} \pi(x+1)} r^{x+1} \pi(x+1) \\
&= \frac{1}{3r} \sum_{x=1}^{z-1} \frac{x+1}{n-x} r^{x+1} \pi(x+1) \\
&= \frac{1}{3r} \sum_{x=2}^z \frac{x}{n-x+1} r^x \pi(x) \\
&< \frac{1}{3r} \sum_{x=1}^z \frac{x}{n-x+1} r^x \pi(x) \\
&< \frac{1}{3r} \frac{z}{n-z+1} \sum_{x=1}^z r^x \pi(x).
\end{aligned}$$

Combining,

$$h_z(r) < 1 + \frac{z}{3r(n-z+1)} h_z(r)$$

or

$$h_z(r) < \frac{1}{1 - \frac{z}{3r(n-z+1)}}.$$

Since $1 \leq z \leq \theta n$, $h_z(r)$ is constant in this range, proving the result. \square

We can now combine the results to prove the bound on the zero chain gap:

Proof of Theorem 3.6.5. Using Lemmas 3.6.14, 3.9.11 and 3.6.13 together with Theorem 3.5.9 proves the result. \square

3.7 Main Result

We will now show how the mixing time results imply that we have an approximate 2-design.

Proof of Theorem 3.3.1: We will go via the 2-norm since this gives a tight bound when working with the Pauli operators. We write ρ in the Pauli basis as usual (as Eqn. 3.2.3) and note that ρ is not necessarily a physical state so the coefficients may

not be real.

$$\begin{aligned}
\|\mathcal{G}_W - \mathcal{G}_H\|_\diamond^2 &= \sup_{\rho} \frac{1}{\|\rho\|_1^2} \|(\mathcal{G}_W \otimes \text{id}_{2^{2n}})(\rho) - (\mathcal{G}_H \otimes \text{id}_{2^{2n}})(\rho)\|_1^2 \\
&\leq 2^{4n} \sup_{\rho} \frac{1}{\|\rho\|_1^2} \|(\mathcal{G}_W \otimes \text{id}_{2^{2n}})(\rho) - (\mathcal{G}_H \otimes \text{id}_{2^{2n}})(\rho)\|_2^2 \\
&= \sup_{\rho} \frac{1}{\|\rho\|_1^2} \left\| \sum_{\substack{p_1, p_2, p_3, p_4 \\ p_1 p_2 \neq 00}} \gamma_0(p_1, p_2, p_3, p_4) (\mathcal{G}_W(\sigma_{p_1} \otimes \sigma_{p_2}) \otimes \sigma_{p_3} \otimes \sigma_{p_4}) \right. \\
&\quad \left. - \mathcal{G}_H(\sigma_{p_1} \otimes \sigma_{p_2}) \otimes \sigma_{p_3} \otimes \sigma_{p_4} \right\|_2^2
\end{aligned}$$

Now, write (for $p_1 p_2 \neq 00$) $\mathcal{G}_W(\frac{1}{2^n} \sigma_{p_1} \otimes \sigma_{p_2}) = \frac{1}{2^n} \sum_{\substack{q_1, q_2 \\ q_1 q_2 \neq 00}} g_t(q_1, q_2; p_1, p_2) \sigma_{q_1} \otimes \sigma_{q_2}$.

We get

$$\begin{aligned}
&\sup_{\rho} \frac{1}{\|\rho\|_1^2} \left\| \sum_{\substack{p_1, p_2, p_3, p_4, q_1, q_2 \\ p_1 p_2 \neq 00, q_1 q_2 \neq 00}} \gamma_0(p_1, p_2, p_3, p_4) \left(g_t(q_1, q_2; p_1, p_2) - \frac{\delta_{q_1 q_2} \delta_{p_1 p_2}}{2^n (2^n + 1)} \right) \right. \\
&\quad \left. \sigma_{q_1} \otimes \sigma_{q_2} \otimes \sigma_{p_3} \otimes \sigma_{p_4} \right\|_2^2 \\
&= 2^{4n} \sup_{\rho} \frac{1}{\|\rho\|_1^2} \sum_{\substack{p_1, p_2, p_3, p_4, q_1, q_2 \\ p_1 p_2 \neq 00, q_1 q_2 \neq 00}} |\gamma_0(p_1, p_2, p_3, p_4)|^2 \left(g_t(q_1, q_2; p_1, p_2) - \frac{\delta_{q_1 q_2} \delta_{p_1 p_2}}{2^n (2^n + 1)} \right)^2 \\
&\leq 2^{4n} \epsilon^2 \sup_{\rho} \frac{\sum_{\substack{p_1, p_2, p_3, p_4 \\ p_1 p_2 \neq 00}} |\gamma_0(p_1, p_2, p_3, p_4)|^2}{\|\rho\|_1^2} \\
&\leq 2^{4n} \epsilon^2 \sup_{\rho} \frac{\|\rho\|_2^2}{\|\rho\|_1^2} \\
&= 2^{4n} \epsilon^2
\end{aligned}$$

where the first equality comes from the orthogonality of the Pauli operators under the Hilbert-Schmidt inner product. This proves the result for the diamond norm, Definition 2.2.10. For the distance measure defined in Definition 2.2.11 (TWIRL), the argument in [DCEL06] can be used together with the 1-norm bound to prove the result. \square

3.8 Conclusions

We have proved tight convergence results for the first two moments of a random circuit. We have used this to show that random circuits are efficient approximate 1- and 2-unitary designs. Our framework readily generalises to k -designs for any k and the next step in this research is to prove that random circuits give approximate k -designs for all k .

We have shown that, provided the random circuit uses gates from a universal gate set that is also universal on $U(4)$, the circuit is still an efficient 2-design. We also see that the random circuit with gates chosen uniformly from $U(4)$ is the most natural model. We note that the gates from $U(4)$ can be replaced by gates from any approximate 2-design on two qubits without any change to the asymptotic convergence properties.

Finally, random circuits are interesting physical models in their own right. The original purpose of [ODP07] was to answer the physical question of how quickly entanglement grows in a system with random two party interactions. Lemma 3.3.2(i) shows that $O(n(n + \log 1/\epsilon))$ steps suffice (in contrast to $O(n^2(n + \log 1/\epsilon))$ which they prove) to give almost maximal entanglement in such a system.

3.9 Proofs

3.9.1 Zero chain mixing time proofs

Asymmetric Simple Random Walk

We will use some facts about asymmetric simple random walks i.e. a random walk on a 1D line with probability p of moving right at each step and probability $q = 1 - p$ of moving left.

The position of the walk after k steps is tightly concentrated around $k(p - q)$:

Lemma 3.9.1. *Let X_k be the random variable giving the position of a random walk after k steps starting at the origin with probability p of moving right and probability*

$q = 1 - p$ of moving left. Let $\mu = p - q$. Then for any $\eta > 0$,

$$\mathbb{P}(X_k \geq \mu k + \eta) \leq \exp\left(-\frac{\eta^2}{2k}\right)$$

and

$$\mathbb{P}(X_k \leq \mu k - \eta) \leq \exp\left(-\frac{\eta^2}{2k}\right).$$

Proof. The standard Chernoff bound for 0/1 variables \tilde{Y}_i gives, with \tilde{Y}_i equal to 1 with probability p and for $Y_k = \sum_{i=1}^k \tilde{Y}_i$,

$$\begin{aligned} \mathbb{P}(Y_k \geq kp + \eta) &\leq \exp\left(-\frac{2\eta^2}{k}\right) \\ \mathbb{P}(Y_k \leq kp - \eta) &\leq \exp\left(-\frac{2\eta^2}{k}\right). \end{aligned}$$

For our case, set $\tilde{Y}_i = 2\tilde{X}_i - 1$ to give the desired result. \square

This result is for a walk with constant bias. We will need a result for a walk with varying (but bounded from below) bias:

Lemma 3.9.2. *Let X_k be the random variable giving the position of a random walk after k steps starting at the origin with probability $p_i \geq p$ of moving right and probability $q_i \leq p$ of moving left at step i . Let $\mu = p - (1 - p)$. Then for any $\eta > 0$,*

$$\mathbb{P}(X_k \geq \mu k + \eta) \leq \exp\left(-\frac{\eta^2}{2k}\right)$$

and

$$\mathbb{P}(X_k \leq \mu k - \eta) \leq \exp\left(-\frac{\eta^2}{2k}\right).$$

Proof. Let \tilde{Y}_i be a random variable equal to 1 with probability p and 0 with probability $1 - p$. Then let \tilde{Z}_i be a random variable equal to 1 with probability p_i and 0 with probability $1 - p_i$. Let $Y_k = \sum_{i=1}^k \tilde{Y}_i$ and $Z_k = \sum_{i=1}^k \tilde{Z}_i$. Then following the standard

Chernoff bound derivation (for $\lambda > 0$),

$$\begin{aligned} \mathbb{P}(Z_k \geq kp + \eta) &= \mathbb{P}\left(e^{\lambda Z_k} \geq e^{\lambda(kp + \eta)}\right) \\ &\leq \frac{e^{\lambda(kp + \eta)}}{\mathbb{E}e^{\lambda Z_k}} \\ &\leq \frac{e^{\lambda(kp + \eta)}}{\mathbb{E}e^{\lambda Y_k}} \\ &\leq \exp\left(-\frac{2\eta^2}{k}\right). \end{aligned}$$

We can then, as above, set $\tilde{Z}_i = 2\tilde{X}_i - 1$. The calculation is similar for the bound on $\mathbb{P}(X_k \leq \mu k - \eta)$. \square

From Lemma 3.9.1 we can prove a result about how often each site is visited. If the walk runs for t steps the walk is at position $t\mu$ with high probability so we might expect from symmetry that each site will have been visited about $1/\mu$ times. Below is a weaker concentration result of this form but is strong enough for our purposes. It says that the amount of time spent $\leq x$ is about x/μ .

Lemma 3.9.3. *For $\gamma > 2$ and integer $x > 0$,*

$$\mathbb{P}\left(\sum_{k=1}^{\infty} \mathbb{I}(X_k \leq x) \geq \gamma x/\mu\right) \leq 2 \exp\left(-\frac{\mu x(\gamma - 2)}{2}\right),$$

where \mathbb{I} is the indicator function.

Proof. Let $Y_k = \mathbb{I}(X_k \leq x)$. From Lemma 3.9.1,

$$\mathbb{P}(Y_k = 0) \leq \exp\left(-\frac{(k\mu - x)^2}{2k}\right)$$

for $k \leq x/\mu$ and

$$\mathbb{P}(Y_k = 1) \leq \exp\left(-\frac{(k\mu - x)^2}{2k}\right)$$

for $k \geq x/\mu$.

Then the quantity to evaluate is

$$\mathbb{P} \left(\sum_{k=1}^{\infty} Y_k \geq \gamma x / \mu \right).$$

We use a standard trick to split this into two mutually exclusive possibilities and then bound the probabilities separately. Write

$$\begin{aligned} \mathbb{P} \left(\sum_{k=1}^{\infty} Y_k \geq \gamma x / \mu \right) &= \mathbb{P} \left(\left(\sum_{k=1}^{\infty} Y_k \geq \gamma x / \mu \right) \cap \left(\bigcap_{j=1}^{\lceil \gamma x / \mu \rceil} [Y_j = 1] \right) \right) + \\ &\quad \mathbb{P} \left(\left(\sum_{k=1}^{\infty} Y_k \geq \gamma x / \mu \right) \cap \left(\bigcup_{j=1}^{\lceil \gamma x / \mu \rceil} [Y_j = 0] \right) \right). \end{aligned} \quad (3.9.1)$$

We can bound the first term:

$$\begin{aligned} \mathbb{P} \left(\left(\sum_{k=1}^{\infty} Y_k \geq \gamma x / \mu \right) \cap \left(\bigcap_{j=1}^{\lceil \gamma x / \mu \rceil} [Y_j = 1] \right) \right) &= \mathbb{P} \left(\bigcap_{k=1}^{\lceil \gamma x / \mu \rceil} Y_k = 1 \right) \\ &\leq \mathbb{P} (Y_{\lceil \gamma x / \mu \rceil} = 1) \\ &\leq \exp \left(-\frac{\mu x (\gamma - 1)^2}{2\gamma} \right) \\ &\leq \exp \left(-\frac{\mu x (\gamma - 2)}{2} \right) \end{aligned}$$

The second term similarly:

$$\begin{aligned} \mathbb{P} \left(\left(\sum_{k=1}^{\infty} Y_k \geq \gamma x / \mu \right) \cap \left(\bigcup_{j=1}^{\lceil \gamma x / \mu \rceil} [Y_j = 0] \right) \right) &\leq \mathbb{P} \left(\bigcup_{k=\lceil \frac{\gamma x}{\mu} + 1}^{\infty} [Y_k = 1] \right) \\ &\leq \sum_{k=\lceil \frac{\gamma x}{\mu} + 1}^{\infty} \mathbb{P} (Y_k = 1) \\ &\leq \sum_{k=\lceil \frac{\gamma x}{\mu} + 1}^{\infty} \exp \left(-\frac{(k\mu - x)^2}{2k} \right) \\ &\leq \exp \left(-\frac{\mu x (\gamma - 2)}{2} \right). \quad \square \end{aligned}$$

The last fact we need about asymmetric simple random walks is a bound on the probability of going backwards. If $p > q$ then we expect the walk to go right in the majority of steps. The probability of going left a distance a is exponentially small in a . This is a well known result, often stated as part of the gambler's ruin problem:

Lemma 3.9.4 (See e.g. [GW86]). *Consider an asymmetric simple random walk that starts at $a > 0$ and has an absorbing barrier at the origin. The probability that the walk eventually absorbs at the origin is 1 if $p \leq q$ and $(q/p)^a$ otherwise.*

This result is for infinitely many steps. If we only consider finitely many steps, the probability of absorption must be at most this.

Waiting Time

From above we saw that the probability of moving is at least $2x/5n$ when at position x . The length of time spent waiting at each step is therefore stochastically dominated by a geometric distribution with parameter $2x/5n$. The following concentration result will be used to bound the waiting time (in our case $\beta = 2/5$):

Lemma 3.9.5. *Let the waiting time at each site be $W(x) \sim \text{Geo}(\beta x/n)$, the total waiting time $W = \sum_{x=1}^t W(x)$ and $t' = \frac{n \ln t}{\beta}$. Then*

$$\mathbb{P}(W \geq Ct') \leq 2t^{(1-C)/2}.$$

Proof. By Markov's inequality for $\lambda > 0$,

$$\mathbb{P}(W \geq Ct') \leq \frac{\mathbb{E}e^{\lambda W}}{e^{\lambda Ct'}}.$$

The $W(x)$ are independent so

$$\mathbb{E}e^{\lambda W} = \prod_{x=1}^t \mathbb{E}e^{\lambda W(x)}.$$

Summing the geometric series we find

$$\mathbb{E}e^{\lambda W(x)} = \frac{\frac{\beta x}{n}}{e^{-\lambda} - 1 + \frac{\beta x}{n}}$$

provided $e^\lambda < \frac{1}{1 - \frac{\beta x}{n}}$ for all $1 \leq x \leq t$. Therefore e^λ is of the form $\frac{1}{1 - \frac{\alpha \beta}{n}}$ where $0 < \alpha < 1$. With this,

$$\mathbb{E}e^{\lambda W(x)} = \frac{x}{x - \alpha}$$

and

$$\mathbb{E}e^{\lambda W} = \frac{t! \Gamma(1 - \alpha)}{\Gamma(t + 1 - \alpha)}.$$

We are free to choose α within its range to optimise the bound. However, for simplicity, we will choose $\alpha = 1/2$. From Lemma 3.9.12,

$$\mathbb{E}e^{\lambda W} \leq 2\sqrt{t}.$$

The result follows, using the inequality $1 - x \leq e^{-x}$. □

Phase 1

Here we prove that phase 1 completes successfully with high probability. The bias here is large so the walk moves right every time with high probability:

Lemma 3.9.6. *The probability that the accelerated chain moves right at each step, starting from $x = 1$ for t steps, is at least*

$$1 - t^2/n.$$

Proof. The probability of moving right at each step is

$$\begin{aligned} \prod_{x=1}^t \frac{3(n-x)}{3n-2x-1} &= \frac{(n-2)(n-3)\dots(n-t)}{(n-5/3)(n-7/3)\dots(n-(2t+1)/3)} \\ &\geq (1-2/n)(1-3/n)\dots(1-t/n) \\ &\geq (1-t/n)^t \geq 1-t^2/n \quad \square \end{aligned}$$

Let $t = n^\delta$. Provided $\delta < 1/2$ this probability is close to one. Therefore, with high probability, the walk moves to n^δ in n^δ steps. Using Lemma 3.9.5 the waiting time can be bounded:

Lemma 3.9.7. *Let $W^{(1)}$ be the waiting time during phase 1. Let H be the event that the walk moves right at each step. Then*

$$\mathbb{P}\left(W^{(1)} \geq Ct' | H\right) \leq 2n^{\delta(1-C)/2} \quad (3.9.2)$$

where $t' = \frac{5\delta n \ln n}{2}$.

Proof. This follows directly from Lemma 3.9.5, since each site is hit exactly once. \square

We now combine these two lemmas to prove that phase 1 completes successfully with high probability:

Proof of Lemma 3.6.9. In Lemma 3.9.6, we show that in n^δ accelerated steps, the walk moves right at each step with probability $\geq 1 - n^{2\delta-1}$. Call this event H . Then $\mathbb{P}(H) \geq 1 - n^{2\delta-1}$. Lemma 3.9.7 shows that the waiting time $W^{(1)}$ is bounded with high probability (choosing $C = 3$):

$$\mathbb{P}(W^{(1)} \leq 15n\delta \ln n/2 | H) \geq 1 - 2n^{-\delta}.$$

Then we can bound the probability of phase 1 completing successfully:

$$\begin{aligned}
\mathbb{P}(\text{Phase 1 completes successfully}) &\geq \mathbb{P}(H \cap W^{(1)} \leq 15n\delta \ln n/2) \\
&= \mathbb{P}(H)\mathbb{P}(W^{(1)} \leq 15n\delta \ln n/2|H) \\
&\geq (1 - n^{2\delta-1})(1 - 2n^{-\delta}) \\
&\geq 1 - n^{2\delta-1} - 2n^{-\delta}. \quad \square
\end{aligned}$$

Phase 2

Phase 2 starts at $n^\delta/2$ and finishes when the walk has reached θn for some constant $0 < \theta < 3/4$. We show that, with high probability, this also takes time $O(n \log n)$. The probability of moving right during this phase is at least $p = \frac{3(1-\theta)}{3-2\theta}$. We first define some constants that we will derive bounds in terms of. Let γ be a constant > 2 . Let $\mu = p - (1-p)$ and $\tilde{\mu} = \mu/\gamma$. Finally let $s = \tilde{\mu}t$ for some t (which will be the number of accelerated steps). Then, with high probability, the walk will have passed s after t steps:

Lemma 3.9.8. *Let X_t be the position of the walk at accelerated step t , where $X_0 = n^\delta$. Then*

$$\mathbb{P}(X_t \leq s) \leq \exp(-\mu^2 t(1 - 1/\gamma)^2/2).$$

Proof. Let $X'_t = X_t - n^\delta$. Then from Lemma 3.9.2,

$$\mathbb{P}(X'_t \leq \mu t - \eta) \leq \exp\left(-\frac{\eta^2}{2t}\right).$$

Now let $\eta = \mu t - s$ and use

$$\begin{aligned}
\mathbb{P}(X_t \leq s) &= \mathbb{P}(X'_t \leq s - n^\delta) \\
&\leq \mathbb{P}(X'_t \leq s)
\end{aligned}$$

to complete the proof. □

We now prove a bound on the waiting time:

Lemma 3.9.9. *Let $W^{(2)}$ be the waiting time in phase 2. Then, assuming the walk does not go back beyond $n^\delta/2$,*

$$\mathbb{P}\left(W^{(2)} \geq \frac{15n \ln s}{\mu}\right) \leq (4/s)^{3/2\mu} + \frac{2 \exp\left(\frac{-\mu n^\delta}{4}\right)}{1 - \exp\left(\frac{-\mu}{2}\right)}. \quad (3.9.3)$$

Proof. Let $W_k \sim \text{Geo}\left(\frac{2X_k}{5n}\right)$ where X_k is the position of the walk at accelerated step k ($X_0 = n^\delta$). We want to bound (w.h.p.) the waiting time $W^{(2)} = \sum_{k=1}^t W_k$ of t steps of the accelerated walk.

Define the event H to be

$$H = \left\{ \bigcap_{x \geq n^\delta/2} \left[\sum_{k=1}^{\infty} \mathbb{I}(X_k \leq x) \leq x/\tilde{\mu} \right] \right\}. \quad (3.9.4)$$

If H occurs, no sites have been hit too often and the walk has not gone back further than $n^\delta/2$. It is important that we also use the restriction that $X_k \geq n^\delta/2$ because the waiting time grows the longer the walk moves back. However, it is very unlikely that the walk will go backwards (even to $n^\delta/2$).

We now define some more notation to bound the waiting time. Let $\mathbf{X} = (X_1, X_2, \dots, X_t)$ be a tuple of positions and let $N_x(\mathbf{X})$ be the number of times that x appears in \mathbf{X} and let $\mathbf{N}(\mathbf{X}) = (N_1(\mathbf{X}), N_2(\mathbf{X}), \dots, N_n(\mathbf{X}))$. Then we have $\sum_x N_x(\mathbf{X}) = t$.

As we said above, the waiting time at $x = a$ stochastically dominates the waiting time at $x = b$ for $b \geq a$. In other words,

$$W_k \supseteq W_{k'} \text{ if } X_k \leq X_{k'} \quad (3.9.5)$$

where $X \supseteq Y$ means that X stochastically dominates Y . Now write the waiting time

for all steps

$$\begin{aligned} W^{(2)}(\mathbf{X}) &= \sum_{k=1}^t W_k \\ &= \sum_x \sum_{h=1}^{N_x(\mathbf{X})} W_h(x) \end{aligned} \quad (3.9.6)$$

where $W_h(x) \sim \text{Geo}\left(\frac{2x}{5n}\right)$.

If event H occurs, we can put some bounds on N_x . We find that, for all $x \geq n^\delta/2$,

$$\sum_{y=n^\delta/2}^x N_y(\mathbf{X}) \leq x/\tilde{\mu} \quad (3.9.7)$$

and $N_x(\mathbf{X}) = 0$ for $x < n^\delta/2$. Now let \mathbf{X}_m be such that $N_{n^\delta/2}(\mathbf{X}_m) = \frac{n^\delta}{2\tilde{\mu}}$ and $N_x(\mathbf{X}_m) = 1/\tilde{\mu}$ for $x > n^\delta/2$. Then

$$\sum_{y=n^\delta/2}^x N_y(\mathbf{X}_m) = x/\tilde{\mu}. \quad (3.9.8)$$

Now we introduce the relation \preceq :

Definition 3.9.10. *Let \mathbf{x} and \mathbf{y} be n -tuples. Then $x \preceq y$ if*

$$\sum_{i=1}^k x_i \leq \sum_{i=1}^k y_i \quad (3.9.9)$$

for all $1 \leq k \leq n$ with equality for $k = n$.

Note that this is like majorisation, except the elements of the tuples are not sorted. Using this, we find that $\mathbf{N}(\mathbf{X}) \preceq \mathbf{N}(\mathbf{X}_m)$ (Using $\sum_y N_y(\mathbf{X}) = \sum_y N_y(\mathbf{X}') = t$ for all \mathbf{X}, \mathbf{X}' .)

If we combine Equations 3.9.5 and 3.9.6 we find that $W^{(2)}(\mathbf{X}) \geq W^{(2)}(\mathbf{X}')$ if $\mathbf{N}(\mathbf{X}) \succeq \mathbf{N}(\mathbf{X}')$. Roughly speaking, this is simply saying that the waiting time is larger if the earlier sites are hit more often. But since for all \mathbf{X} that satisfy H , $\mathbf{X} \preceq \mathbf{X}_m$, we have $W^{(2)}(\mathbf{X}) \leq W^{(2)}(\mathbf{X}_m)$ provided H occurs. We will simplify further

by noting that $\mathbf{X}_m \preceq \mathbf{X}_0$ where $N_x(\mathbf{X}_0) = 1/\tilde{\mu}$ for $1 \leq x \leq \tilde{\mu}t = s$ and zero elsewhere. Therefore

$$\mathbb{P}\left(W^{(2)}(\mathbf{X}) \geq \frac{5Cn \ln s}{2\tilde{\mu}} \middle| H\right) \leq \mathbb{P}\left(W^{(2)}(\mathbf{X}_0) \geq \frac{5Cn \ln s}{2\tilde{\mu}}\right).$$

We can bound this by applying Lemma 3.9.5. Let $W_h = \sum_{x=1}^s W_h(x)$. From Lemma 3.9.5,

$$\mathbb{P}(W_h \geq Ct') \leq 2s^{\frac{1-C}{2}} \quad (3.9.10)$$

where $t' = \frac{5n \ln s}{2}$. However, we want a bound on $\mathbb{P}\left(\sum_{h=1}^{1/\tilde{\mu}} W_h \geq Ct'/\tilde{\mu}\right)$. The same reasoning as in Lemma 3.9.5 bounds this as

$$\mathbb{P}\left(\sum_{h=1}^{1/\tilde{\mu}} W_h \geq Ct'/\tilde{\mu}\right) \leq \left(2s^{\frac{1-C}{2}}\right)^{1/\tilde{\mu}}. \quad (3.9.11)$$

Therefore

$$\mathbb{P}\left(W^{(2)}(\mathbf{X}_0) \geq \frac{5Cn \ln s}{2\tilde{\mu}}\right) \leq 2^{1/\tilde{\mu}} s^{\frac{(1-C)/2}{\tilde{\mu}}}. \quad (3.9.12)$$

To complete the proof, we just need to find $\mathbb{P}(H^c)$. We can bound it using the union bound and Lemma 3.9.3:

$$\begin{aligned}
\mathbb{P}(H^c) &= \mathbb{P} \left(\bigcup_{x=n^\delta/2}^n \left[\sum_{k=1}^{\infty} \mathbb{I}(X_k \leq x) > x/\tilde{\mu} \right] \right) \\
&\leq \sum_{x=n^\delta/2}^n \mathbb{P} \left(\sum_{k=1}^{\infty} \mathbb{I}(X_k \leq x) \geq x/\tilde{\mu} \right) \\
&\leq \sum_{x=n^\delta/2}^n 2 \exp \left(\frac{-\mu x(\gamma-2)}{2} \right) \\
&\leq \sum_{x=n^\delta/2}^{\infty} 2 \exp \left(\frac{-\mu x(\gamma-2)}{2} \right) \\
&= \frac{2 \exp \left(\frac{-\mu n^\delta(\gamma-2)}{4} \right)}{1 - \exp \left(\frac{-\mu(\gamma-2)}{2} \right)}
\end{aligned}$$

Now, for any events A and B

$$\begin{aligned}
\mathbb{P}(A) &= \mathbb{P}(A \cap B) + \mathbb{P}(A \cap B^c) \\
&= \mathbb{P}(A|B)\mathbb{P}(B) + \mathbb{P}(A \cap B^c) \\
&\leq \mathbb{P}(A|B) + \mathbb{P}(B^c)
\end{aligned}$$

and set $C = 2$ and $\gamma = 3$ to obtain the result. □

We now combine these two lemmas to prove that phase 2 completes successfully with high probability:

Proof of Lemma 3.6.10. Phase 2 can fail if:

- The walk does not reach θn . The probability of this is bounded by Lemma 3.9.8:

$$\mathbb{P}(X_t \leq \theta n) \leq \exp \left(-\frac{2}{3} \mu \theta n \right).$$

This follows from setting $t = \frac{3\theta n}{\mu}$ and $\gamma = 3$.

- The waiting time is too long. This probability is bounded by Lemma 3.9.9:

$$\mathbb{P}\left(W^{(2)} \geq \frac{15n \ln(\theta n)}{\mu}\right) \leq \left(\frac{4}{\theta n}\right)^{\frac{3}{2\mu}} + \frac{2 \exp\left(\frac{-\mu n^\delta}{4}\right)}{1 - \exp(-\mu/2)} + (q/p)^{n^\delta/2}.$$

- The walk gets back to $n^\delta/2$. This is bounded by Lemma 3.9.4:

$$\mathbb{P}\left(\text{Walk gets to } n^\delta/2\right) \leq (q/p)^{n^\delta/2}.$$

So, using the union bound we can bound the overall probability of failure:

$$\mathbb{P}(\text{Phase 2 fails}) \leq \exp\left(-\frac{2}{3}\mu\theta n\right) + \left(\frac{4}{\theta n}\right)^{\frac{3}{2\mu}} + \frac{2 \exp\left(\frac{-\mu n^\delta}{4}\right)}{1 - \exp(-\mu/2)} + (q/p)^{n^\delta/2}. \quad \square$$

Phase 3

This phase starts at θn . We show that this mixes quickly using log-Sobolev arguments.

Lemma 3.9.11. *The zero chain on the restricted state space $x \in [m, n]$ where $m = \theta n$ for $\theta > 0$ has mixing time $O\left(n \log \frac{n}{\epsilon}\right)$.*

Proof. We restrict the Markov chain to only run from m by adjusting the holding probability at m , $P(m, m)$. Construct the chain P' with transition matrix

$$P'(x, y) = \begin{cases} 0 & x < m \text{ or } y < m \\ 1 - P(m, m+1) & x = y = m \\ P(x, y) & \text{otherwise} \end{cases} \quad (3.9.13)$$

where P is the transition matrix of the full zero chain. This chain then has stationary distribution

$$\pi'(x) = \begin{cases} \pi(x)/(1-w) & m \leq x \leq n \\ 0 & \text{otherwise} \end{cases} \quad (3.9.14)$$

where $w = \sum_{x=1}^{m-1} \pi(x)$. To see this, first note that the distribution is normalised. We

want to show that

$$\sum_{x=m}^n P'(x, y) \pi'(x) = \pi'(y). \quad (3.9.15)$$

When $y = m$ we are required to prove that $P'(m, m)\pi'(m) + P'(m + 1, m)\pi'(m + 1) = \pi'(m)$. This follows from the reversibility of the unrestricted zero chain, using $P'(m, m) = 1 - P(m, m + 1)$. For $y > m$, Eqn. 3.9.15 is satisfied simply because $\pi(x)$ is the stationary distribution of P and related by a constant factor to $\pi'(x)$.

We can now prove this final mixing time result, making use of Lemma 3.5.12. Let Q_i be the chain that uniformly mixes site i . This converges in one step and has a log-Sobolev constant independent of n ; call it ρ_1 . Let Q be the chain that chooses a site at random and then uniformly mixes that site. This is the product chain of the Q_i so, by Lemma 3.5.12, has gap $1/n$ and Sobolev constant $\rho_Q = \rho_1/n$. We can construct the zero chain for this and find its Sobolev constant.

The Sobolev constant is defined (Definition 3.5.10) in terms of a minimisation over functions on the state space. For the chain Q we can write

$$\rho_Q = \inf_{\phi} f(\phi).$$

If we restrict the infimum to be over functions ϕ with $\phi(x) = \phi(y)$ for x and y containing the same number of zeroes then we obtain the Sobolev constant for the zero- Q chain, ρ_{Q_0} , which is the chain which counts the number of zeroes in the full chain Q . Since taking the infimum over less functions cannot give a smaller value,

$$\rho_{Q_0} \geq \rho_Q \geq \rho_1/n.$$

We can now compare this chain to the zero- P chain. The stationary distributions are

the same. The transition matrix for the zero- Q chain is

$$Q_0(x, y) = \begin{cases} \frac{n+2x}{4n} & y = x \\ \frac{x}{4n} & y = x - 1 \\ \frac{3(n-x)}{4n} & y = x + 1 \\ 0 & \text{otherwise} \end{cases}$$

Then construct Q'_0 by restricting the space to only run from m in exactly the same way as P' is constructed from P . Q'_0 has the same stationary distribution as P' . Now we can perform the comparison. From Eqn. 3.5.14:

$$\begin{aligned} A &= \max_{a \geq m} \frac{Q'_0(a, a+1)}{P'(a, a+1)} \\ &= \max_{a \geq m} \frac{5(n-1)}{8a} \leq \frac{5}{8\theta}. \end{aligned}$$

Therefore $\rho_{P'} \geq \frac{8\theta\rho_1}{5n}$. Exactly the same argument applies to show the gap is $\Omega(1/n)$ so the mixing time is (from Eqn. 3.5.20) $O(n \log \frac{n}{\epsilon})$. \square

Now we can prove that phase 3 completes successfully with high probability:

Proof of Lemma 3.6.11. In Lemma 3.9.11, we show that after $O(n \log \frac{n}{\epsilon})$ steps the chain mixes to distance ϵ . We just need to show that the walk goes back to $\theta n/2$ with small probability. This follows from Lemma 3.9.4. \square

3.9.2 Moment Generating Function Calculations

The following lemma is needed in the moment generating function calculations.

Lemma 3.9.12. *For Integer $s > 0$,*

$$\frac{\Gamma(s+1)\Gamma(1/2)}{\Gamma(s+1/2)} \leq 2\sqrt{s} \tag{3.9.16}$$

Proof. From expanding the Γ functions, Eqn. 3.9.16 becomes

$$\begin{aligned} \frac{s!2^s}{(2s-1)!!} &= \frac{2 \times 4 \times 6 \times \dots \times 2(s-1) \times 2s}{1 \times 3 \times 5 \times \dots \times (2s-3) \times (2s-1)} \\ &= \prod_{x=1}^s \frac{2x}{2x-1} \end{aligned}$$

We then proceed by induction. $\prod_{x=1}^1 \frac{2x}{2x-1} = 2$ and by the inductive hypothesis

$$\prod_{x=1}^{s+1} \frac{2x}{2x-1} \leq \frac{2(s+1)}{2(s+1)-1} 2\sqrt{s}.$$

It is easy to show that $\frac{2(s+1)}{2(s+1)-1} \leq \sqrt{\frac{s+1}{s}}$ and the result follows. \square

3.9.3 Mixing Times

We find bounds for the mixing time above that are valid with high probability. Below we turn these into full mixing time bounds.

Lemma 3.9.13. *If after $O(n \log n)$ steps the state v of a random walk satisfies*

$$\|v - \pi\| \leq \delta$$

where π is the stationary distribution and δ is $1/\text{poly}(n)$ then the number of steps required to be at most a distance ϵ from stationarity is

$$O\left(n \log \frac{n}{\epsilon}\right).$$

Proof. Let s be the slowest mixing initial state. Then, after $t = O(n \log n)$ steps we have at worst the state

$$(1 - \delta)\pi + \delta s$$

and if we repeat kt times δ becomes δ^k . So to get a distance ϵ , $k = \left\lceil \frac{\log \epsilon}{\log \delta} \right\rceil$.

Now we evaluate the mixing time:

$$\begin{aligned} kt &= O(n \log n) \left\lceil \frac{\log \epsilon}{\log \delta} \right\rceil = O(n \log n) \left\lceil \frac{\log 1/\epsilon}{\log 1/\delta} \right\rceil \\ &= O(n \max(\log n, \log 1/\epsilon)) \\ &= O\left(n \log \frac{n}{\epsilon}\right). \end{aligned} \quad \square$$

Chapter 4

Quantum Tensor Product Expanders and an Efficient Unitary Design Construction

4.1 Introduction

In this chapter, we give an efficient construction of a unitary k -design on n qubits for any k up to $O(n/\log(n))$. We will do this by first finding an efficient construction of a quantum k -copy tensor product expander (k-TPE), which can then be iterated to produce a k -design. We will therefore need to understand some of the theory of expanders before presenting our construction.

Classical expander graphs have the property that a marker executing a random walk on the graph will have a distribution close to the stationary distribution after a small number of steps. We consider a generalisation of this, known as a k -tensor product expander (TPE) and due to [HH09], to graphs that randomise k different markers carrying out correlated random walks on the same graph. This is a stronger requirement than for a normal ($k = 1$) expander because the correlations between walkers (unless they start at the same position) must be broken. We then generalise

quantum expanders in the same way, so that the unitaries act on k copies of the system. We give an efficient construction of a quantum k -TPE which uses an efficient classical k -TPE as its main ingredient. We then give as a key application the first efficient construction of a unitary k -design for any k .

While randomised constructions yield k -designs (by a modification of Theorem 5 of [ABW09]) and k -TPEs (when the dimension is polynomially larger than k [HH09]) with near-optimal parameters, these approaches are not efficient. Previous efficient constructions of k -designs were known only for $k = 1, 2$, and no efficient constant-degree, constant-gap quantum k -TPEs were previously known, except for the $k = 1$ case corresponding to quantum expanders [BT07, AS04, Har08, GE08].

In Section 4.1.1, we will define quantum expanders and other key terms. Then in Section 4.1.2 we will describe our main result which will be proved in Section 4.2. In this chapter, we will use N to denote the dimension rather than d to be consistent with the rest of the quantum expander literature.

This chapter has been published previously as [HL09a] and is joint work with Aram Harrow.

4.1.1 Quantum Expanders

We will only consider D -regular expander graphs here. We can think of a random walk on such a graph as selecting one of D permutations of the vertices randomly at each step. We construct the permutations as follows. Label the vertices from 1 to N . Then label each edge from 1 to D so that each edge label appears exactly once on the incoming and outgoing edges of each vertex. This gives a set of D permutations. Choosing one of these permutations at random (for some fixed probability distribution) then defines a random walk on the graph.

We now define a classical k -TPE:

Definition 4.1.1 ([HH09]). *Let ν be a probability distribution on \mathcal{S}_N with support on*

$\leq D$ permutations. Then ν is a classical (N, D, λ, k) -TPE if

$$\left\| \mathbb{E}_{\pi \sim \nu} \left[B(\pi)^{\otimes k} \right] - \mathbb{E}_{\pi \sim \mathcal{S}_N} \left[B(\pi)^{\otimes k} \right] \right\|_{\infty} = \left\| \sum_{\pi \in \mathcal{S}_N} \left(\nu(\pi) - \frac{1}{N!} \right) B(\pi)^{\otimes k} \right\|_{\infty} \leq \lambda. \quad (4.1.1)$$

with $\lambda < 1$. Here $\mathbb{E}_{\pi \sim \nu}$ means the expectation over π drawn according to ν and $\mathbb{E}_{\pi \sim \mathcal{S}_N}$ means the expectation over π drawn uniformly from \mathcal{S}_N .

Setting $k = 1$ recovers the usual spectral definition of an expander. Note that a (N, D, λ, k) -TPE is also a (N, D, λ, k') -TPE for any $k' \leq k$. The largest meaningful value of k is $k = N$, corresponding to the case when ν describes a Cayley graph expander on \mathcal{S}_N .

The degree of the map is $D = |\text{supp } \nu|$ and the gap is $1 - \lambda$. Ideally, the degree should be small and gap large. To be useful, these should normally be independent of N and possibly k . We say that a TPE construction is efficient if it can be implemented in $\text{poly}(\log N)$ steps. There are known constructions of efficient classical TPEs. The construction of Hoory and Brodsky [BH08] provides an expander with $D = \text{poly}(\log N)$ and $\lambda = 1 - 1/\text{poly}(k, \log N)$ with efficient running time. An efficient TPE construction is also known, due to Kassabov [Kas05], which has constant degree and gap (independent of N and k).

Similarly, we define a quantum k -TPE:

Definition 4.1.2 ([HH09]). *Let ν be a distribution on $\mathcal{U}(N)$, the group of $N \times N$ unitary matrices, with $D = |\text{supp } \nu|$. Then ν is a quantum (N, D, λ, k) -TPE if*

$$\left\| \mathbb{E}_{U \sim \nu} \left[U^{\otimes k, k} \right] - \mathbb{E}_{U \sim \mathcal{U}(N)} \left[U^{\otimes k, k} \right] \right\|_{\infty} \leq \lambda \quad (4.1.2)$$

with $\lambda < 1$. Here $\mathbb{E}_{U \sim \mathcal{U}(N)}$ means the expectation over U drawn from the Haar measure.

Again, normally we want D and λ to be constants and setting $k = 1$ recovers the usual definition of a quantum expander. Note that an equivalent statement of the

above definition is that, for all ρ ,

$$\left\| \mathbb{E}_{U \sim \nu} \left[U^{\otimes k} \rho(U^\dagger)^{\otimes k} \right] - \mathbb{E}_{U \sim \mathcal{U}(N)} \left[U^{\otimes k} \rho(U^\dagger)^{\otimes k} \right] \right\|_2 \leq \lambda \|\rho\|_2 \quad (4.1.3)$$

A natural application of this is to make an efficient unitary k -design. The definition we use here is the same as for a k -TPE, except with closeness in the 1-norm rather than the ∞ -norm. This is given in Definition 2.2.12 (TRACE).

We can make an ϵ -approximate unitary k -design from a quantum k -TPE with $O(k \log N)$ overhead:

Theorem 4.1.3. *If \mathcal{U} is a quantum (N, D, λ, k) -TPE then iterating the map $m = \frac{1}{\log 1/\lambda} \log \frac{N^{2k}}{\epsilon}$ times gives an ϵ -approximate unitary k -design according to Definition 2.2.12 (TRACE) with D^m unitaries.*

Proof. Iterating the TPE m times gives

$$\left\| \mathbb{E}_{U \sim \nu} [U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(N)} [U^{\otimes k, k}] \right\|_\infty \leq \lambda^m$$

This implies that

$$\left\| \mathbb{E}_{U \sim \nu} [U^{\otimes k, k}] - \mathbb{E}_{U \sim \mathcal{U}(N)} [U^{\otimes k, k}] \right\|_1 \leq N^{2k} \lambda^m$$

We take m such that $N^{2k} \lambda^m = \epsilon$ to give the result. \square

Corollary 4.1.4. *A construction of an efficient quantum (N, D, λ, k) -TPE yields an efficient approximate unitary k -design, provided $\lambda = 1 - 1/\text{poly}(\log N)$. Further, if D and λ are constants, the number of unitaries in the design is $N^{O(k)}$.*

Our approach to construct an efficient quantum k -TPE will be to take an efficient classical $2k$ -TPE and mix it with a quantum Fourier transform. The degree is thus only larger than the degree of the classical expander by one. Since the quantum Fourier transform on \mathbb{C}^N requires $\text{poly}(\log N)$ time, it follows that if the classical expander is efficient then the quantum expander is as well. The main technical difficulty is to

show for suitable values of k that the gap of the quantum TPE is not too much worse than the gap of the classical TPE.

A similar approach to ours was first used in [HH09] to construct a quantum expander (i.e. a 1-TPE) by mixing a classical 2-TPE with a phase. However, regardless of the set of phases chosen, this approach will not yield quantum k -TPEs from classical $2k$ -TPEs for any $k \geq 2$.

4.1.2 Main Result

Let $\omega = e^{2\pi i/N}$ and define the N -dimensional Fourier transform to be

$$\mathcal{F} = \frac{1}{\sqrt{N}} \sum_{m=1}^N \sum_{n=1}^N \omega^{mn} |m\rangle \langle n|. \quad (4.1.4)$$

Define $\delta_{\mathcal{F}}$ to be the distribution on $\mathcal{U}(N)$ consisting of a point mass on \mathcal{F} . Our main result in this chapter is that mixing $\delta_{\mathcal{F}}$ with a classical $2k$ -TPE yields a quantum k -TPE for appropriately chosen k and N .

Theorem 4.1.5. *Let ν_C be a classical $(N, D, 1 - \epsilon_C, 2k)$ -TPE, and for $0 < p < 1$, define $\nu_Q = p\nu_C + (1 - p)\delta_{\mathcal{F}}$. Suppose that*

$$\epsilon_A := 1 - 2(2k)^{4k} / \sqrt{N} > 0. \quad (4.1.5)$$

Then ν_Q is a quantum $(N, D + 1, 1 - \epsilon_Q, k)$ -TPE where

$$\epsilon_Q \geq \frac{\epsilon_A}{12} \min(p\epsilon_C, 1 - p) > 0 \quad (4.1.6)$$

The bound in Eqn. 4.1.6 is optimised when $p = 1/(1 + \epsilon_C)$, in which case we have

$$\epsilon_Q \geq \frac{\epsilon_A \epsilon_C}{24}. \quad (4.1.7)$$

This means that any constant-degree, constant-gap classical $2k$ -TPE gives a quantum k -TPE with constant degree and gap. If the classical TPE is efficient then

the quantum TPE is as well. Using Corollary 4.1.4, we obtain approximate unitary k -designs with polynomial-size circuits.

Unfortunately the construction does not work for all dimensions; we require that $N = \Omega((2k)^{8k})$, so that ϵ_A is lower-bounded by a positive constant. However, in applications normally k is fixed. An interesting open problem is to find a construction that works for all dimensions, in particular a $k = \infty$ expander. (Most work on $k = \infty$ TPEs so far has focused on the $N = 2$ case [BG06].) We suspect our construction may work for k as large as cN for a small constant c . On the other hand, if $2k > N$ then the gap in our construction drops to zero.

4.2 Proof of Theorem 4.1.5

4.2.1 Proof overview

First, we introduce some notation. Define $\mathcal{E}_{S_N}^{2k} = \mathbb{E}_{\pi \sim S_N}[B(\pi)^{\otimes 2k}]$ and $\mathcal{E}_{U(N)}^k = \mathbb{E}_{U \sim \mathcal{U}(N)}[U^{\otimes k, k}]$. These are both projectors onto spaces which we label V_{S_N} and $V_{U(N)}$ respectively. Since $V_{U(N)} \subset V_{S_N}$, it follows that $\mathcal{E}_{S_N}^{2k} - \mathcal{E}_{U(N)}^k$ is a projector onto the space $V_0 := V_{S_N} \cap V_{U(N)}^\perp$. We also define $\mathcal{E}_{\nu_C}^{2k} = \mathbb{E}_{\pi \sim \nu_C}[B(\pi)^{\otimes 2k}]$ and $\mathcal{E}_{\nu_Q}^k = \mathbb{E}_{U \sim \nu_Q}[U^{\otimes k, k}]$.

The idea of our proof is to consider $\mathcal{E}_{\nu_C}^{2k}$ a proxy for $\mathcal{E}_{S_N}^{2k}$; if λ_C is small enough then this is a reasonable approximation. Then we can restrict our attention to vectors in V_0 , which we would like to show all shrink substantially under the action of our expander. This in turn can be reduced to showing that $\mathcal{F}^{\otimes k, k}$ maps any vector in V_0 to a vector that has $\Omega(1)$ amplitude in $V_{S_N}^\perp$. This last step is the most technically involved step of the chapter, and involves careful examination of the different vectors making up V_{S_N} .

Thus, our proof reduces to two key Lemmas. The first allows us to substitute $\mathcal{E}_{\nu_C}^{2k}$ for $\mathcal{E}_{S_N}^{2k}$ while keeping the gap constant.

Lemma 4.2.1 ([HH09] Lemma 1). *Let Π be a projector and let X and Y be operators such that $\|X\|_\infty \leq 1$, $\|Y\|_\infty \leq 1$, $\Pi X = X \Pi = \Pi$, $\|(I - \Pi)X(I - \Pi)\|_\infty \leq 1 - \epsilon_C$ and*

$\|\Pi Y \Pi\|_\infty \leq 1 - \epsilon_A$. Assume $0 < \epsilon_C, \epsilon_A < 1$. Then for any $0 < p < 1$, $\|pX + (1 - p)Y\|_\infty < 1$. Specifically,

$$\|pX + (1 - p)Y\|_\infty \leq 1 - \frac{\epsilon_A}{12} \min(p\epsilon_C, 1 - p). \quad (4.2.1)$$

We will restrict to $V_{\mathcal{U}(N)}^\perp$, or equivalently, subtract the projector $\mathcal{E}_{\mathcal{U}(N)}^k$ from each operator. Thus we have $X = \mathcal{E}_{\nu_C}^{2k} - \mathcal{E}_{\mathcal{U}(N)}^k$, $\Pi = \mathcal{E}_{\mathcal{S}_N}^{2k} - \mathcal{E}_{\mathcal{U}(N)}^k$ and $Y = \mathcal{F}^{\otimes k, k} - \mathcal{E}_{\mathcal{U}(N)}^k$. According to Definition 4.1.1, we have the bound

$$\|(I - \Pi)X(I - \Pi)\|_\infty = \|\mathcal{E}_{\nu_C}^{2k} - \mathcal{E}_{\mathcal{S}_N}^{2k}\|_\infty \leq 1 - \epsilon_C. \quad (4.2.2)$$

It will remain only to bound $\lambda_A := 1 - \epsilon_A = \left\| \left(\mathcal{E}_{\mathcal{S}_N}^{2k} - \mathcal{E}_{\mathcal{U}(N)}^k \right) \mathcal{F}^{\otimes k, k} \left(\mathcal{E}_{\mathcal{S}_N}^{2k} - \mathcal{E}_{\mathcal{U}(N)}^k \right) \right\|_\infty$.

Lemma 4.2.2. For $N \geq (2k)^2$,

$$\lambda_A = \left\| \left(\mathcal{E}_{\mathcal{S}_N}^{2k} - \mathcal{E}_{\mathcal{U}(N)}^k \right) \mathcal{F}^{\otimes k, k} \left(\mathcal{E}_{\mathcal{S}_N}^{2k} - \mathcal{E}_{\mathcal{U}(N)}^k \right) \right\|_\infty \leq 2(2k)^{4k} / \sqrt{N}. \quad (4.2.3)$$

Combining Eqn. 4.2.2, Lemma 4.2.2 and Lemma 4.2.1 now completes the proof of Theorem 4.1.5.

4.2.2 Action of a Classical $2k$ -TPE

We start by analysing the action of a classical $2k$ -TPE. (We consider $2k$ -TPEs rather than general k -TPEs since our quantum expander construction only uses these.) The fixed points are states which are unchanged when acted on by $2k$ copies of any permutation matrix. Since the same permutation is applied to all copies, any equal indices will remain equal and any unequal indices will remain unequal. This allows us to identify the fixed points of the classical expander: they are the sums over all states with the same equality and difference constraints. For example, for $k = 1$ (corresponding to a 2-TPE), the fixed points are $\sum_{n_1} |n_1, n_1\rangle$ and $\sum_{n_1 \neq n_2} |n_1, n_2\rangle$ (all off-diagonal entries equal to 1). In general, there is a fixed point for each partition of the set $\{1, 2, \dots, 2k\}$ into at most N non-empty parts. If $N \geq 2k$, which is the only

case we consider, the $2k^{\text{th}}$ Bell number β_{2k} gives the number of such partitions (see e.g. [Sta86]).

We now write down some more notation to further analyse this. If Π is a partition of $\{1, \dots, 2k\}$, then we write $\Pi \vdash 2k$. We will see that $\mathcal{E}_{S_N}^{2k}$ projects onto a space spanned by vectors labelled by partitions. For a partition Π , say that $(i, j) \in \Pi$ if and only if elements i and j are in the same block. Now we can write down the fixed points of the classical expander. Let

$$I_\Pi = \{(n_1, \dots, n_{2k}) : n_i = n_j \text{ iff } (i, j) \in \Pi\}. \quad (4.2.4)$$

This is a set of tuples where indices in the same block of Π are equal and indices in different blocks are not equal. The corresponding state is

$$|I_\Pi\rangle = \frac{1}{\sqrt{|I_\Pi|}} \sum_{\mathbf{n} \in I_\Pi} |\mathbf{n}\rangle \quad (4.2.5)$$

where $\mathbf{n} = (n_1, \dots, n_{2k})$. Note that the $\{|I_\Pi\rangle\}_{\Pi \vdash 2k}$ form a partition $\{1, \dots, N\}^{2k}$ and thus the $\{|I_\Pi\rangle\}_{\Pi \vdash 2k}$ form an orthonormal basis for V_{S_N} . This is because, when applying the same permutation to all indices, indices that are the same remain the same and indices that differ remain different. This implies that

$$\mathcal{E}_{S_N}^{2k} = \sum_{\Pi \vdash 2k} |I_\Pi\rangle \langle I_\Pi|. \quad (4.2.6)$$

To evaluate the normalisation, use $|I_\Pi| = (N)_{|\Pi|}$ where $(N)_n$ is the falling factorial $N(N-1)\dots(N-n+1)$ and $|\Pi|$ is the number of blocks in Π . We will later find it useful to bound $(N)_n$ with

$$\left(1 - \frac{n^2}{2N}\right) N^n \leq (N)_n \leq N^n. \quad (4.2.7)$$

We will also make use of the refinement partial order:

Definition 4.2.3. *The refinement partial order \leq on partitions $\Pi, \Pi' \in \text{Par}(2k, N)$*

is given by

$$\Pi \leq \Pi' \text{ iff } (i, j) \in \Pi \Rightarrow (i, j) \in \Pi'. \quad (4.2.8)$$

For example, $\{\{1, 2\}, \{3\}, \{4\}\} \leq \{\{1, 2, 4\}, \{3\}\}$. Note that $\Pi \leq \Pi'$ implies that $|\Pi| \geq |\Pi'|$.

Turning Inequality Constraints into Equality Constraints.

In the analysis, it will be easier to consider just equality constraints rather than both inequality and equality constraints as in I_Π . Therefore we make analogous definitions:

$$E_\Pi = \{(n_1, \dots, n_{2k}) : n_i = n_j \forall (i, j) \in \Pi\} \quad (4.2.9)$$

and

$$|E_\Pi\rangle = \frac{1}{\sqrt{|E_\Pi|}} \sum_{\mathbf{n} \in E_\Pi} |\mathbf{n}\rangle. \quad (4.2.10)$$

Then $|E_\Pi| = N^{|\Pi|}$. For E_Π , indices in the same block are equal, as with I_Π , but indices in different blocks need not be different.

We will need relationships between I_Π and E_Π . First, observe that E_Π can be written as the union of some $I_{\Pi'}$ sets:

$$E_\Pi = \bigcup_{\Pi' \geq \Pi} I_{\Pi'}. \quad (4.2.11)$$

To see this, note that for $\mathbf{n} \in E_\Pi$, we have $n_i = n_j \forall (i, j) \in \Pi$, but we may also have an arbitrary number of additional equalities between n_i 's in different blocks. The (unique) partition Π' corresponding to these equalities has the property that Π is a refinement of Π' ; that is, $\Pi' \geq \Pi$. Thus for any $\mathbf{n} \in E_\Pi$ there exists a unique $\Pi' \geq \Pi$ such that $\mathbf{n} \in I_{\Pi'}$. Conversely, whenever $\Pi' \geq \Pi$, we also have $I_{\Pi'} \subseteq E_{\Pi'} \subseteq E_\Pi$ because each inclusion is achieved only by relaxing constraints.

Using Eqn. 4.2.11, we can obtain a useful identity involving sums over partitions:

$$N^{|\Pi|} = |E_\Pi| = \sum_{\Pi' \geq \Pi} |I_{\Pi'}| = \sum_{\Pi' \geq \Pi} N_{(|\Pi'|)}. \quad (4.2.12)$$

Additionally, since both sides in Eqn. 4.2.12 are degree $|\Pi|$ polynomials and are equal on $\geq |\Pi| + 1$ points (we can choose any N in Eqn. 4.2.12 with $N \geq 2k$), it implies that $x^{|\Pi|} = \sum_{\Pi' \geq \Pi} x_{(|\Pi'|)}$ as an identity on formal polynomials in x .

The analogue of Eqn. 4.2.11 for the states $|E_\Pi\rangle$ and $|I_\Pi\rangle$ is similar but has to account for normalisation factors. Thus we have

$$\sqrt{|E_\Pi|} |E_\Pi\rangle = \sum_{\Pi' \geq \Pi} \sqrt{|I_{\Pi'}|} |I_{\Pi'}\rangle. \quad (4.2.13)$$

We would also like to invert this relation, and write $|I_\Pi\rangle$ as a sum over various $|E_{\Pi'}\rangle$. Doing so will require introducing some more notation. Define $\zeta(\Pi, \Pi')$ to be 1 if $\Pi \leq \Pi'$ and 0 if $\Pi \not\leq \Pi'$. This can be thought of as a matrix that, with respect to the refinement ordering, has ones on the diagonal and is upper-triangular. Thus it is also invertible. Define $\mu(\Pi, \Pi')$ to be the matrix inverse of ζ , meaning that for all Π_1, Π_2 , we have

$$\sum_{\Pi' \geq 2k} \zeta(\Pi_1, \Pi') \mu(\Pi', \Pi_2) = \sum_{\Pi' \geq 2k} \mu(\Pi_1, \Pi') \zeta(\Pi', \Pi_2) = \delta_{\Pi_1, \Pi_2},$$

where $\delta_{\Pi_1, \Pi_2} = 1$ if $\Pi_1 = \Pi_2$ and $= 0$ otherwise. Thus, if we rewrite Eqn. 4.2.13 as

$$\sqrt{|E_\Pi|} |E_\Pi\rangle = \sum_{\Pi' \geq 2k} \zeta(\Pi, \Pi') \sqrt{|I_{\Pi'}|} |I_{\Pi'}\rangle, \quad (4.2.14)$$

then we can use μ to express $|I_\Pi\rangle$ in terms of the $|E_\Pi\rangle$ as

$$\sqrt{|I_\Pi|} |I_\Pi\rangle = \sum_{\Pi' \geq 2k} \mu(\Pi, \Pi') \sqrt{|E_{\Pi'}|} |E_{\Pi'}\rangle. \quad (4.2.15)$$

This approach is a generalisation of inclusion-exclusion known as Möbius inversion,

and the function μ is called the Möbius function (see Chapter 3 of [Sta86] for more background). For the case of the refinement partial order, the Möbius function is known:

Lemma 4.2.4 ([Rot64], Section 7).

$$\mu(\Pi, \Pi') = (-1)^{|\Pi| - |\Pi'|} \prod_{i=1}^{|\Pi'|} (b_i - 1)!$$

where b_i is the number of blocks of Π in the i^{th} block of Π' .

We can use this to evaluate sums involving the Möbius function for the refinement order.

Lemma 4.2.5.

$$\sum_{\Pi' \geq \Pi} |\mu(\Pi, \Pi')| x^{|\Pi'|} = x^{(|\Pi|)} \quad (4.2.16)$$

where x is arbitrary and $x^{(n)}$ is the rising factorial $x(x+1)\cdots(x+n-1)$.

Proof. Start with $|\mu(\Pi, \Pi')| = (-1)^{|\Pi| - |\Pi'|} \mu(\Pi, \Pi')$ to obtain

$$\begin{aligned} \sum_{\Pi' \geq \Pi} |\mu(\Pi, \Pi')| x^{|\Pi'|} &= (-1)^{|\Pi|} \sum_{\Pi' \geq \Pi} \mu(\Pi, \Pi') (-x)^{|\Pi'|} \\ &= (-1)^{|\Pi|} \sum_{\Pi' \geq \Pi} \mu(\Pi, \Pi') \sum_{\Pi'' \geq \Pi'} \zeta(\Pi', \Pi'') (-x)_{(|\Pi''|)} \end{aligned}$$

using Eqn. 4.2.12. Then use Möbius inversion and $(-x)_{(n)} = (-1)^n x^{(n)}$ to prove the result. \square

We will mostly be interested in the special case $x = 1$:

Corollary 4.2.6.

$$\sum_{\Pi' \geq \Pi} |\mu(\Pi, \Pi')| = |\Pi|! \quad (4.2.17)$$

Using $|\mu(\Pi, \Pi')| \geq 1$ and the fact that $\Pi \geq \{\{1\}, \dots, \{n\}\}$ for all $\Pi \vdash n$, we obtain a bound on the total number of partitions:

Corollary 4.2.7. *The Bell numbers β_n satisfy $\beta_n \leq n!$.*

4.2.3 Fixed Points of a Quantum Expander

We now turn to $V_{\mathcal{U}(N)}$, the space fixed by the quantum expander. As in Chapter 2, the only operators on $(\mathbb{C}^N)^{\otimes k}$ to commute with $U^{\otimes k}$ for all U are linear combinations of subsystem permutations. The equivalent statement for $V_{\mathcal{U}(N)}$ is that the only states invariant under all $U^{\otimes k,k}$ are of the form

$$\frac{1}{\sqrt{N^k}} \sum_{n_1, \dots, n_k \in [N]} |n_1, \dots, n_k, n_{\pi(1)}, \dots, n_{\pi(k)}\rangle, \quad (4.2.18)$$

for some permutation $\pi \in \mathcal{S}_k$. Since $\mathcal{E}_{\mathcal{U}(N)}^k = \mathbb{E}[U^{\otimes k,k}]$ projects onto the set of states that is invariant under all $U^{\otimes k,k}$, it follows that $V_{\mathcal{U}(N)}$ is equal to the span of the states in Eqn. 4.2.18.

Now we relate these states to our previous notation.

Definition 4.2.8. *For $\pi \in \mathcal{S}_k$, define the partition corresponding to π by*

$$P(\pi) = \{\{1, k + \pi(1)\}, \{2, k + \pi(2)\}, \dots, \{k, k + \pi(k)\}\}.$$

Then the state in Eqn. 4.2.18 is simply $|E_{P(\pi)}\rangle$, and so

$$V_{\mathcal{U}(N)} = \text{span}\{|E_{P(\pi)}\rangle : \pi \in \mathcal{S}_k\}. \quad (4.2.19)$$

Note that the classical expander has many more fixed points than just the desired $|E_{P(\pi)}\rangle$. The main task in constructing a quantum expander from a classical one is to modify the classical expander to decay the fixed points that should not be fixed by the quantum expander.

4.2.4 Fourier Transform in the Matrix Element Basis

Since we make use of the Fourier transform, we will need to know how it acts on a matrix element. We find

$$\mathcal{F}^{\otimes k,k}|\mathbf{m}\rangle = \frac{1}{N^k} \sum_{\mathbf{n}} \omega^{\mathbf{m}\cdot\mathbf{n}}|\mathbf{n}\rangle$$

where

$$\mathbf{m}\cdot\mathbf{n} = m_1n_1 + \dots + m_kn_k - m_{k+1}n_{k+1} - \dots - m_{2k}n_{2k} \quad (4.2.20)$$

We will also find it convenient to estimate the matrix elements $\langle E_{\Pi_1}|\mathcal{F}^{\otimes k,k}|E_{\Pi_2}\rangle$. The properties we require are proven in the following lemmas.

Lemma 4.2.9. *Choose any $\Pi_1, \Pi_2 \vdash 2k$. Let $\mathbf{m} \in \Pi_1$ and $\mathbf{n} \in \Pi_2$. Call the free indices of \mathbf{m} \tilde{m}_i for $1 \leq i \leq |\Pi_1|$. Then let $\mathbf{m}\cdot\mathbf{n} = \sum_{i=1}^{|\Pi_1|} \sum_{j=1}^{2k} \tilde{m}_i A_{i,j} n_j$ where $A_{i,j}$ is a $|\Pi_1| \times 2k$ matrix with entries in $\{0, 1, -1\}$ which depends on Π_1 (but not Π_2). Then*

$$\langle E_{\Pi_1}|\mathcal{F}^{\otimes k,k}|E_{\Pi_2}\rangle = N^{-k + \frac{|\Pi_1| - |\Pi_2|}{2}} \sum_{\mathbf{n} \in E_{\Pi_2}} \mathbb{I} \left(\sum_j A_{i,j} n_j \equiv 0 \pmod{N} \forall i \right) \quad (4.2.21)$$

where \mathbb{I} is the indicator function.

Proof. Simply perform the \mathbf{m} sum in

$$\langle E_{\Pi_1}|\mathcal{F}^{\otimes k,k}|E_{\Pi_2}\rangle = N^{-\left(k + \frac{|\Pi_1| + |\Pi_2|}{2}\right)} \sum_{\mathbf{m} \in E_{\Pi_1}} \sum_{\mathbf{n} \in E_{\Pi_2}} \omega^{\mathbf{m}\cdot\mathbf{n}} \quad (4.2.22)$$

□

Lemma 4.2.10. $\langle E_{\Pi_1}|\mathcal{F}^{\otimes k,k}|E_{\Pi_2}\rangle$ is real and positive.

Proof. Since all entries in the sum in Eqn. 4.2.21 are nonnegative and at least one ($\mathbf{n} = 0$) is strictly positive, Lemma 4.2.9 implies the result. □

Lemma 4.2.11. If $\Pi'_1 \leq \Pi_1$ and $\Pi'_2 \leq \Pi_2$ then

$$\sqrt{|E_{\Pi_1}| \cdot |E_{\Pi_2}|} \langle E_{\Pi_1}|\mathcal{F}^{\otimes k,k}|E_{\Pi_2}\rangle \leq \sqrt{|E_{\Pi'_1}| \cdot |E_{\Pi'_2}|} \langle E_{\Pi'_1}|\mathcal{F}^{\otimes k,k}|E_{\Pi'_2}\rangle \quad (4.2.23)$$

Proof. We prove first the special case when $\Pi'_1 = \Pi_1$, but $\Pi'_2 \leq \Pi_2$ is arbitrary. Recall that $\Pi'_2 \leq \Pi_2$ implies that $E_{\Pi_2} \subseteq E_{\Pi'_2}$. Now the LHS of Eqn. 4.2.23 equals

$$\begin{aligned}
& N^{-k} \sum_{\mathbf{m} \in E_{\Pi_1}, \mathbf{n} \in E_{\Pi_2}} \exp\left(\frac{2\pi i}{N} \mathbf{m} \cdot \mathbf{n}\right) \\
&= N^{|\Pi_1| - k} \sum_{\mathbf{n} \in E_{\Pi_2}} \mathbb{I}\left(\sum_j A_{i,j} n_j \equiv 0 \pmod{N} \forall i\right) \\
&= N^{|\Pi_1| - k} \sum_{\mathbf{n} \in E_{\Pi'_2}} \mathbb{I}(\mathbf{n} \in E_{\Pi_2}) \mathbb{I}\left(\sum_j A_{i,j} n_j \equiv 0 \pmod{N} \forall i\right) \\
&\leq N^{|\Pi_1| - k} \sum_{\mathbf{n} \in E_{\Pi'_2}} \mathbb{I}\left(\sum_j A_{i,j} n_j \equiv 0 \pmod{N} \forall i\right) \\
&= \sqrt{|E_{\Pi_1}| |E_{\Pi'_2}|} \langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi'_2} \rangle,
\end{aligned}$$

as desired. To prove Eqn. 4.2.23 we repeat this argument, interchanging the roles of Π_1 and Π_2 and use the fact that $\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle$ is symmetric in Π_1 and Π_2 . \square

Lemma 4.2.12.

$$\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle \leq N^{-\frac{1}{2}|2k - (|\Pi_1| + |\Pi_2|)|} \quad (4.2.24)$$

Proof. Here, there are two cases to consider. The simpler case is when $|\Pi_1| + |\Pi_2| \leq 2k$. Here we simply apply the inequality

$$\sum_{\mathbf{m} \in E_{\Pi_1}, \mathbf{n} \in E_{\Pi_2}} \exp\left(\frac{2\pi i}{N} \mathbf{m} \cdot \mathbf{n}\right) \leq |E_{\Pi_1}| |E_{\Pi_2}| = N^{|\Pi_1| + |\Pi_2|}$$

to Eqn. 4.2.22, and conclude that $\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle \leq N^{\frac{|\Pi_1| + |\Pi_2|}{2} - k}$.

Next, we would like to prove that

$$\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle \leq N^{k - \frac{|\Pi_1| + |\Pi_2|}{2}}. \quad (4.2.25)$$

Here we use Lemma 4.2.11 with $\Pi'_1 = \Pi_1$ and $\Pi'_2 = \{\{1\}, \{2\}, \dots, \{2k\}\}$, the maximally refined partition. Note that $|E_{\Pi'_2}| = N^{2k}$ and $\mathcal{F}^{\otimes k, k}|E_{\Pi'_2}\rangle = |0\rangle$. Thus

$$\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle \leq N^{k - \frac{|\Pi_2|}{2}} \langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi'_2} \rangle = N^{k - \frac{|\Pi_2|}{2}} \langle E_{\Pi_1} | 0 \rangle = N^{k - \frac{|\Pi_1| + |\Pi_2|}{2}},$$

establishing Eqn. 4.2.25. □

Lemma 4.2.13. *If $\Pi_1 = \Pi_2 = P(\pi)$ then $\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle = 1$. If, for any Π_1, Π_2 with $|\Pi_1| + |\Pi_2| = 2k$, either condition isn't met (i.e. either $\Pi_1 \neq \Pi_2$ or there does not exist $\pi \in \mathcal{S}_k$ such that $P(\pi) = \Pi_1 = \Pi_2$) then*

$$\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle \leq \frac{2k}{N} \tag{4.2.26}$$

for $N > k$.

Proof. In Lemma 4.2.14, we introduce the $\Pi_1 \times \Pi_2$ matrix \tilde{A} with the property that

$$\mathbf{m} \cdot \mathbf{n} = \sum_{i=1}^{|\Pi_1|} \sum_{j=1}^{|\Pi_2|} \tilde{m}_i \tilde{A}_{i,j} \tilde{n}_j \tag{4.2.27}$$

for all $\mathbf{m} \in \Pi_1$ and $\mathbf{n} \in \Pi_2$ where \tilde{m}_i and \tilde{n}_j are the free indices of \mathbf{m} and \mathbf{n} . This is similar to the matrix A introduced in Lemma 4.2.9 except only the free indices of \mathbf{n} are considered.

For $\Pi_1 = \Pi_2 = P(\pi)$, Lemma 4.2.14 implies that $\tilde{A} = 0$, or equivalently $\mathbf{m} \cdot \mathbf{n} = 0$ for all $\mathbf{m}, \mathbf{n} \in P(\pi)$. Using $|\Pi_1| + |\Pi_2| = 2k$, $\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle = 1$.

Otherwise we have $(\Pi_1, \Pi_2) \notin \{(P(\pi), P(\pi)) : \pi \in \mathcal{S}_k\}$ with $|\Pi_1| + |\Pi_2| = 2k$. For all these, Lemma 4.2.14 implies that \tilde{A} is nonzero (for $N > k$, no entries in \tilde{A} can be $> N$ or $< -N$ so $\tilde{A} \equiv 0 \pmod{N}$ is equivalent to $\tilde{A} = 0$). Fix an i for which the i^{th} row of \tilde{A} is nonzero. We wish to count the number of $(\tilde{n}_1, \dots, \tilde{n}_{|\Pi_2|})$ such that $\sum_j \tilde{A}_{i,j} \tilde{n}_j \equiv 0 \pmod{N}$. Assume that each $\tilde{A}_{i,j}$ divides N and is nonnegative; if not, we can replace $\tilde{A}_{i,j}$ with $\text{GCD}(|\tilde{A}_{i,j}|, N)$ by a suitable change of variable for \tilde{n}_j .

Now choose an arbitrary j such that $\tilde{A}_{i,j} \neq 0$. For any values of $\tilde{n}_1, \dots, \tilde{n}_{j-1}$,

$\tilde{n}_{j+1}, \dots, \tilde{n}_{|\Pi_2|}$, there are $|\tilde{A}_{i,j}| \leq 2k$ choices of \tilde{n}_j such that $\sum_j \tilde{A}_{i,j} \tilde{n}_j \equiv 0 \pmod{N}$. Thus, there are $\leq 2kN^{|\Pi_2|-1}$ choices of \tilde{n} such that $\sum_j \tilde{A}_{i,j} \tilde{n}_j \equiv 0 \pmod{N}$. Substituting this into Eqn. 4.2.21 (which we can trivially modify to apply for \tilde{A} rather than just A), we find that

$$\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle \leq \frac{2k}{N} N^{-k + \frac{|\Pi_1| + |\Pi_2|}{2}} = \frac{2k}{N},$$

thus establishing Eqn. 4.2.26. \square

Lemma 4.2.14. *Let \tilde{A} be the matrix such that $\mathbf{m} \cdot \mathbf{n} = \sum_{i=1}^{|\Pi_1|} \sum_{j=1}^{|\Pi_2|} \tilde{m}_i \tilde{A}_{i,j} \tilde{n}_j$ for all $\mathbf{m} \in \Pi_1$ and $\mathbf{n} \in \Pi_2$ where \tilde{m}_j and \tilde{n}_j are the free indices of \mathbf{m} and \mathbf{n} . Then $\tilde{A} = 0$ if and only if $\Pi_1, \Pi_2 \geq P(\pi)$ for some $\pi \in \mathcal{S}_k$.*

Proof. We first consider $\Pi_1 = \Pi_2 = P(\pi)$ for the “if” direction. Note that for any $\mathbf{m}, \mathbf{n} \in E_{P(\pi)}$, we have

$$\mathbf{m} \cdot \mathbf{n} = \sum_{j=1}^k m_j n_j - \sum_{j=1}^k m_{\pi(j)} n_{\pi(j)} = 0. \quad (4.2.28)$$

This implies that $\tilde{A} = 0$. Now, choose any $\Pi_1 \geq P(\pi)$ and $\Pi_2 \geq P(\pi)$. Then for any $\mathbf{m} \in \Pi_1$ and $\mathbf{n} \in \Pi_2$, $\mathbf{m}, \mathbf{n} \in P(\pi)$. This means Eqn. 4.2.28 holds for this case so $\tilde{A} = 0$ also.

On the other hand, suppose that $\tilde{A} = 0$. We will argue that this implies the existence of a permutation π such that $\Pi_1, \Pi_2 \geq P(\pi)$, thus establishing the “only if” direction.

Let $\Pi_{1,j}$ (resp. $\Pi_{2,j}$) denote the j^{th} block of Π_1 (resp. Π_2). Then

$$\tilde{A}_{i,j} = \sum_{\substack{i' \in \Pi_{1,i} \\ j' \in \Pi_{2,j}}} \Lambda_{i',j'},$$

where $\Lambda_{i',j'}$ is defined to be

$$\Lambda_{i',j'} = \begin{cases} 1 & \text{if } i' = j' \in \{1, \dots, k\} \\ -1 & \text{if } i' = j' \in \{k+1, \dots, 2k\} \\ 0 & \text{if } i' \neq j' \end{cases}$$

If $\tilde{A} = 0$ then for each i, j we have

$$|\Pi_{1,i} \cap \Pi_{2,j} \cap \{1, \dots, k\}| = |\Pi_{1,i} \cap \Pi_{2,j} \cap \{k+1, \dots, 2k\}|. \quad (4.2.29)$$

Denote the *meet* of Π_1 and Π_2 , $\Pi_1 \wedge \Pi_2$ to be the greatest lower bound of Π_1 and Π_2 , or equivalently the unique partition with the fewest blocks that satisfies $\Pi_1 \wedge \Pi_2 \leq \Pi_1$ and $\Pi_1 \wedge \Pi_2 \leq \Pi_2$. The blocks of $\Pi_1 \wedge \Pi_2$ are simply all of the nonempty sets $\Pi_{1,i} \cap \Pi_{2,j}$, for $i = 1, \dots, |\Pi_1|$ and $j = 1, \dots, |\Pi_2|$. Thus, Eqn. 4.2.29 implies that each block of $\Pi_1 \wedge \Pi_2$ contains an equal number of indices from $\{1, \dots, k\}$ as it does from $\{k+1, \dots, 2k\}$. This implies the existence of a permutation $\pi \in \mathcal{S}_k$ such that $\{i, k+\pi(i)\}$ is contained in a single block of $\Pi_1 \wedge \Pi_2$ for each $i = 1, \dots, k$. Equivalently $\Pi_1 \wedge \Pi_2 \geq P(\pi)$, implying that $\Pi_1 \geq P(\pi)$ and $\Pi_2 \geq P(\pi)$. \square

4.2.5 Proof of Lemma 4.2.2

Proof. We would like to show that, for any unit vector $|\psi\rangle \in V_0$, $|\langle \psi | \mathcal{F}^{\otimes k, k} | \psi \rangle|^2 \leq 2(2k)^{4k} / \sqrt{N}$. Our strategy will be to calculate the matrix elements of $\mathcal{F}^{\otimes k, k}$ in the $|I_\Pi\rangle$ and $|E_\pi\rangle$ bases. While the $|I_\Pi\rangle$ states are orthonormal, we will see that the $\langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle$ matrix elements are easier to calculate. We then use Möbius functions to express $|I_\Pi\rangle$ in terms of $|E_\Pi\rangle$.

Consider the matrix $\mathcal{E}_{\mathcal{S}_N}^{2k} \mathcal{F}^{\otimes k, k} \mathcal{E}_{\mathcal{S}_N}^{2k}$. It has $k!$ unit eigenvalues, corresponding to the $k!$ -dimensional space $V_{\mathcal{U}(N)}$. Call the $k! + 1^{\text{st}}$ largest eigenvalue λ_A . We bound λ_A

with

$$\begin{aligned}
k! + \lambda_A^2 &\leq \text{tr} \left(\mathcal{E}_{\mathcal{S}_N}^{2k} \mathcal{F}^{\otimes k, k} \mathcal{E}_{\mathcal{S}_N}^{2k} \right)^2 \\
&= \sum_{\Pi_1, \Pi_2 \vdash 2k} \left| \langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle \right|^2.
\end{aligned} \tag{4.2.30}$$

We divide the terms in Eqn. 4.2.30 into four types.

1. The leading-order contribution comes from the $k!$ terms of the form $\Pi_1 = \Pi_2 = P(\pi)$ for $\pi \in \mathcal{S}_k$. We bound them with the trivial upper bound

$$\left| \langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle \right|^2 \leq 1 \tag{4.2.31a}$$

(which turns out to be nearly tight). We will then show that the remaining terms are all $k^{O(k)}/N$.

2. If $|\Pi_1| + |\Pi_2| < 2k$ then

$$\begin{aligned}
\left| \langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle \right|^2 &= \frac{1}{|I_{\Pi_1}| \cdot |I_{\Pi_2}| N^{2k}} \left| \sum_{\substack{\mathbf{m} \in \Pi_1 \\ \mathbf{n} \in \Pi_2}} e^{\frac{2\pi i \mathbf{m} \cdot \mathbf{n}}{N}} \right|^2 \\
&\leq \frac{|I_{\Pi_1}| \cdot |I_{\Pi_2}|}{N^{2k}} \\
&\leq N^{|\Pi_1| + |\Pi_2| - 2k} \leq \frac{1}{N},
\end{aligned} \tag{4.2.31b}$$

where in the last line we have used the fact that $|I_{\Pi}| \leq |E_{\Pi}| = N^{|\Pi|}$.

3. If $|\Pi_1| + |\Pi_2| > 2k$ then we will show that

$$\left| \langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle \right|^2 \leq \frac{4 \cdot (2k!)^2}{N} \tag{4.2.31c}$$

4. If $|\Pi_1| + |\Pi_2| = 2k$ but either $\Pi_1 \neq \Pi_2$ or there is no $\pi \in \mathcal{S}_k$ satisfying $P(\pi) =$

$\Pi_1 = \Pi_2$, then we will show that

$$\left| \langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle \right|^2 \leq \frac{((2k)! + 2k)^2}{N^2} \leq \frac{4 \cdot (2k!)^2}{N} \quad (4.2.31d)$$

To establish these last two claims, we will find it useful to express $|I_{\Pi}\rangle$ in terms of the various $|E_{\Pi}\rangle$ states.

Lemmas 4.2.12 and 4.2.13 can now be used together with the Möbius function to bound $|\langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle|^2$. First, suppose $|\Pi_1| + |\Pi_2| > 2k$. Then

$$\begin{aligned} \left| \langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle \right| &= \left| \sum_{\substack{\Pi'_1 \geq \Pi_1 \\ \Pi'_2 \geq \Pi_2}} \sqrt{\frac{|E_{\Pi'_1}| |E_{\Pi'_2}|}{|I_{\Pi_1}| |I_{\Pi_2}|}} \mu(\Pi_1, \Pi'_1) \mu(\Pi_2, \Pi'_2) \langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle \right| \\ &\leq \sum_{\substack{\Pi'_1 \geq \Pi_1 \\ \Pi'_2 \geq \Pi_2}} \sqrt{\frac{|E_{\Pi'_1}| |E_{\Pi'_2}|}{|I_{\Pi_1}| |I_{\Pi_2}|}} |\mu(\Pi_1, \Pi'_1) \mu(\Pi_2, \Pi'_2)| \langle E_{\Pi'_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi'_2} \rangle \\ &\leq \frac{N^k}{\sqrt{|I_{\Pi_1}| |I_{\Pi_2}|}} \sum_{\substack{\Pi'_1 \geq \Pi_1 \\ \Pi'_2 \geq \Pi_2}} |\mu(\Pi'_1, \Pi_1) \mu(\Pi'_2, \Pi_2)| \end{aligned}$$

by Lemma 4.2.12. Then using by Corollary 4.2.6 we find

$$\begin{aligned} \left| \langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle \right| &= \frac{N^k |\Pi_1|! |\Pi_2|!}{\sqrt{(N)_{|\Pi_1|} (N)_{|\Pi_2|}}} \\ &\leq \frac{2 \cdot (2k)!}{\sqrt{N}} \end{aligned} \quad (4.2.32)$$

In the last step, we have assumed that $4k^2 < N$, so that $(N)_\ell \geq N^\ell/2$ for any $\ell \leq 2k$. We have also made use of the fact that (still assuming $4k^2 < N$) Eqn. 4.2.32 is maximised when $|\Pi_1| + |\Pi_2| = 2k + 1$, and in particular, when one of $|\Pi_1|, |\Pi_2|$ is equal to $2k$ and the other is equal to 1.

A similar analysis applies to the pairs Π_1, Π_2 with $|\Pi_1| + |\Pi_2| = 2k$, but with

$(\Pi_1, \Pi_2) \notin \{(P(\pi), P(\pi)) : \pi \in \mathcal{S}_k\}$. In this case,

$$\begin{aligned} \langle I_{\Pi_1} | \mathcal{F}^{\otimes k, k} | I_{\Pi_2} \rangle &= \sqrt{\frac{|E_{\Pi_1}| |E_{\Pi_2}|}{|I_{\Pi_1}| |I_{\Pi_2}|}} \langle E_{\Pi_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi_2} \rangle + \\ &\quad \sum_{\substack{\Pi'_1 \geq \Pi_1, \Pi'_2 \geq \Pi_2 \\ (\Pi'_1, \Pi'_2) \neq (\Pi_1, \Pi_2)}} \sqrt{\frac{|E_{\Pi'_1}| |E_{\Pi'_2}|}{|I_{\Pi_1}| |I_{\Pi_2}|}} \mu(\Pi_1, \Pi'_1) \mu(\Pi_2, \Pi'_2) \langle E_{\Pi'_1} | \mathcal{F}^{\otimes k, k} | E_{\Pi'_2} \rangle \end{aligned} \quad (4.2.33)$$

We now use Lemmas 4.2.13 and 4.2.12 to bound each of the two terms. For the first term, we use Eqn. 4.2.26 to upper bound it with $2k/N$. For each choice of Π'_1 and Π'_2 in the second sum, we have $|\Pi'_1| + |\Pi'_2| \leq 2k - 1$. Thus we can upper bound the absolute value of the second term in Eqn. 4.2.33 with

$$\begin{aligned} \frac{1}{\sqrt{|I_{\Pi_1}| |I_{\Pi_2}|}} \sum_{\substack{\Pi'_1 \geq \Pi_1, \Pi'_2 \geq \Pi_2 \\ (\Pi'_1, \Pi'_2) \neq (\Pi_1, \Pi_2)}} |\mu(\Pi_1, \Pi'_1) \mu(\Pi_2, \Pi'_2)| N^{|\Pi'_1| + |\Pi'_2| - k} &\leq \frac{2 \cdot |\Pi_1|! \cdot |\Pi_2|!}{N} \\ &\leq \frac{(2k)!}{N}. \end{aligned}$$

We combine the two terms and square to establish Eqn. 4.2.31d.

We now put together the components from Eqn. 4.2.31 to upper bound Eqn. 4.2.30, and find that

$$k! + \lambda_A^2 \leq k! + \beta_{2k}^2 \frac{4 \cdot (2k!)^2}{N},$$

implying that $\lambda_A \leq 2\beta_{2k}(2k!)/\sqrt{N} \leq 2(2k)^{4k}/\sqrt{N}$. This concludes the proof of Lemma 4.2.2. \square

4.3 Conclusions

We have shown how efficient quantum tensor product expanders can be constructed from efficient classical tensor product expanders. This immediately yields an efficient construction of unitary k -designs for any k . Unfortunately our results do not work for all dimensions; we require the dimension N to be $\Omega((2k)^{8k})$. While tighter analysis of

our construction could likely improve this, our construction does not work for $N < 2k$.
Constructions of expanders for all dimensions remains an open problem.

Chapter 5

Applications of Designs

In this chapter we first survey known applications of designs from a wide variety of areas. Then we present new results applying designs to derandomise some large deviation bounds.

5.1 Review of Applications

As we have already discussed, random unitaries and random states have many applications. For some of these applications a design is sufficient since only the first few moments of the distribution are required to be equal to those of the Haar measure.

5.1.1 Quantum Cryptography

The first applications we discuss are to quantum cryptography. In classical cryptography, the one-time pad is the most basic operation that perfectly encrypts a message using a key that is the same length as the message. In quantum cryptography the analogue is a quantum operation \mathcal{E} such that for all input states ρ , $\mathcal{E}(\rho) = \rho_0$ with the requirement that given a secret key Bob can decode Alice's message perfectly. Since all states are encoded to ρ_0 Eve cannot learn anything about the message without knowing the key. In this section all logs will be taken to base 2.

If ρ_0 is the identity, then the map \mathcal{E} is a unitary 1-design. Therefore using $2n$

bits of key to label the Pauli operators provides a quantum one-time pad. In fact, in [AMTd00] it is shown that $2n$ bits of key are also necessary for a quantum one-time pad. Therefore this unitary 1-design is of optimal size.

It is interesting to note that $2n$ bits of key are required rather than just n for the classical one-time pad. This is related to the fact that quantum states allow superdense coding [BW92], which allows two classical bits to be sent per qubit. Although it is not possible to use a shorter key for exact encryption, it would be desirable to shorten the key if we can tolerate Eve learning a small amount of information about the message. In [AS04], they consider closeness in the 1-norm, and set $\rho_0 = I$. They define a map \mathcal{E} as an ϵ -approximate quantum encryption scheme if for all ρ

$$\|\mathcal{E}(\rho) - I/d\|_1 \leq \epsilon. \quad (5.1.1)$$

Thus we see an ϵ -approximate unitary 1-design according to (for example) Definition 2.2.10 (DIAMOND) suffices. In [AS04], they present an efficient construction that satisfies Eqn. 5.1.1 with $n + 2 \log n + 2 \log(\frac{1}{\epsilon}) + O(1)$ bits of key. While this does not immediately provide an ϵ -approximate 1-design according to any of our definitions with only $(1 + o(1))n$ bits of key, Eqn. 5.1.1 is a valid definition of an ϵ -approximate unitary 1-design. This key length was further improved by Dickinson and Nayak [DN06] to $n + 2 \log(\frac{1}{\epsilon}) + O(1)$ and their construction is efficient.

A stronger definition for ϵ -approximate encryption was given in [HLSW04]. They define a map \mathcal{E} to be an ϵ -approximate quantum encryption scheme if for all ρ

$$\|\mathcal{E}(\rho) - I/d\|_\infty \leq \epsilon/d. \quad (5.1.2)$$

This implies the 1-norm bound in Eqn. 5.1.1 but a dimension factor is lost when converting the other way. This could be used as yet another approximate 1-design definition. However, there are no known efficient constructions of such ∞ -norm randomising maps. In [HLSW04] they provide an inefficient randomised construction with key length $n + \log n + 2 \log(\frac{1}{\epsilon}) + O(1)$. Their method is to show that with

non-zero probability random unitaries suffice.

This result was improved by Aubrun in [Aub09] to reduce the key length to $n + 2 \log(\frac{1}{\epsilon}) + O(1)$. The method is the same as [HLSW04] except the analysis is tighter. Aubrun also makes a step towards finding an efficient construction by showing that the unitaries can be Pauli matrices, which can be implemented efficiently, although the sampling is still inefficient.

Besides cryptographic applications, it is shown in [HLSW04] that ∞ -norm randomising maps can be used to hide correlations from local operations and classical communication (LOCC) and have applications to data hiding (see later) and locking of classical correlations [DHL⁺04], whereby classical correlations can be hidden but unlocked by a very short key.

The last cryptography example we give is that of non-malleable encryption given in [ABW09]. Here the authors not only consider hiding information from Eve but they also require that she cannot change the message. Of course, Eve could always replace the message with some fixed state or do nothing, so according to [ABW09], an encryption scheme is non-malleable if these (or a convex combination) are the only operations Eve can perform on the encoded data. The main result of this paper is that a unitary 2-design is necessary and sufficient. They then show, as do Gross et al. [GAE07], that a 2-design requires at least $(d^2 - 1)^2 + 1$ unitaries i.e. the key must be at least $4n - o(1)$ bits long. Even for approximate encryption (which can be seen as an approximate 2-design) the key length is essentially the same.

5.1.2 Measurement

In some cases a random measurement is a good choice but cannot be performed efficiently. One example of such a result is:

Theorem 5.1.1 (Sen, [Sen05]). *Let ρ_1 and ρ_2 be any mixed states with $r(\rho_1) + r(\rho_2) \leq \sqrt{d}/C$ for a sufficiently large constant C . Here, $r(\rho)$ is the rank of the state ρ . Then*

$$\mathbb{E}_M \|M(\rho_1) - M(\rho_2)\|_1 = \Omega(\|\rho_1 - \rho_2\|_2) \quad (5.1.3)$$

where M is an orthonormal basis picked from the Haar measure. Here, $M(\rho)$ is the probability distribution of outcomes according to the POVM M .

Since a large 1-norm distance between probability distributions means the distributions are easily distinguishable, this result places a lower bound on the distinguishability of the states ρ_1 and ρ_2 in terms of their 2-norm distance.

In [AE07], Ambainis and Emerson show that a POVM made from a state 4-design achieves the bound in Eqn. 5.1.3. In fact, an ϵ -approximate state 4-design suffices, provided that $\epsilon = O(\|\rho_1 - \rho_2\|_2^4)$. To ensure the POVM is suitably normalised, we insist here that the approximate 4-design is also an exact 1-design rather than an ϵ -approximate 1-design, which is all that Definition 2.2.9 ensures.

In [IR06], Iblisdir and Roland consider a slightly different measurement problem for which a random measurement achieves the best outcome. The setting is that Alice chooses a random pure state (the authors only consider the case that Alice's system is 2-dimensional i.e. a single qubit) from the Haar measure and makes k copies of it. Bob then has to find a state with high overlap with the given state. The POVM that achieves the optimum is [MP95]

$$(k + 1)|\psi\rangle\langle\psi|^{\otimes k} d\psi. \tag{5.1.4}$$

From Lemma 2.2.2, the average of this is the projector onto the symmetric subspace of k qubits. While this is not the identity, no other outcomes are possible because the input state is symmetric. The states in the POVM can be replaced by a state k -design and in [IR06] the authors present a construction of a state k -design for all k , although only for one qubit.

5.1.3 Average Gate Fidelity

When implementing a quantum operation, we would like to know how far the actual operation is from the desired. One way of measuring this is the average gate fidelity

[Nie02]:

$$\bar{F}(\mathcal{E}, U) = \int d\psi \langle \psi | U^\dagger \mathcal{E}(|\psi\rangle\langle\psi|) U | \psi \rangle \quad (5.1.5)$$

where \mathcal{E} is the operation implemented and U is the desired unitary. We see immediately, following [DCEL06], that the integrand is a balanced polynomial of degree 2 so the Haar measure on states can be replaced by a state 2-design. We can even use an approximate design if the average fidelity only needs to be known approximately. By repeatedly sampling from the design we can obtain an estimate of the average to $1/\text{poly}(\log d)$ accuracy efficiently whereas naively sampling random states will not be efficient.

5.1.4 Data Hiding

Data hiding was introduced by Terhal, DiVincenzo and Leung [TDL01, DLT02] as a fundamentally quantum concept. The setting is that Alice and Bob share a quantum state which contains secret bits. However, the state is chosen so that if they can only communicate using LOCC then they cannot learn this secret bit. To encode one secret bit, the “hider” constructs one of two orthogonal mixed states $\rho_0^{(m)}$ and $\rho_1^{(m)}$ and hands half to Alice and the other half to Bob. $\rho_0^{(m)}$ is the state with m random Bell pairs chosen subject to the constraint that the number of singlets is even. $\rho_1^{(m)}$ is the same state except with an odd number of singlets. The parameter m controls the degree of security.

The way that designs help here is in the construction of these states using minimal resources. The authors show that ρ_0 can be obtained from twirling any initial pure state of the form $|\psi\rangle\langle\psi| \otimes |\psi\rangle\langle\psi|$ ¹:

$$\rho_0^{(m)} = \int_{\mathcal{U}(2^m)} dU (U \otimes U) |\psi\rangle\langle\psi| \otimes |\psi\rangle\langle\psi| (U \otimes U)^\dagger. \quad (5.1.6)$$

$\rho_1^{(m)}$ can be created from $\rho_0^{(m-1)}$.

The authors consider replacing the Haar integral with a sum over a unitary 2-

¹By unitary invariance of the Haar measure, the choice of $|\psi\rangle$ does not affect the resultant state.

design. If errors can be tolerated then an approximate 2-design can be used and the state can be prepared efficiently.

5.1.5 Decoupling and Evolution of Black Holes

For various tasks in quantum Shannon theory, it is desirable to decouple a system from the environment. In [HHYW07] and [ADHW06], it is shown that for most random unitaries applied to the system the resulting overall state is close to a product state.

The setting is that there is a system S with two parts S_1 and S_2 . The environment is E . Let the initial state be ψ_{SE} and let

$$\sigma_{S_2E}(U) = \text{tr}_{S_1} \left[(U \otimes I_E) \psi_{SE} (U^\dagger \otimes I_E) \right]. \quad (5.1.7)$$

Then we have

Theorem 5.1.2 ([ADHW06], Theorem 4.2).

$$\int_{\mathcal{U}(S)} \|\sigma_{S_2E}(U) - \sigma_{S_2}(U) \otimes \sigma_R(U)\|_1^2 dU \leq \frac{d_S d_E}{d_{S_1}^2} \left(\text{tr} \left[(\psi_{SE})^2 \right] + \text{tr} \left[(\psi_S)^2 \right] \text{tr} \left[(\psi_E)^2 \right] \right) \quad (5.1.8)$$

where $\sigma_{S_2}(U) = \text{tr}_{S_1E} \sigma_{SE}(U)$, etc..

The proof uses the 2-norm squared, which is a polynomial of degree 2 in the matrix elements of the random unitary. Therefore the same result holds when U is selected from a unitary 2-design instead and, as above, an approximate design can be used to allow an efficient implementation. This allows the encoding circuits in [ADHW06] to be made efficient although unfortunately the decoding circuits are still inefficient.

Decoupling has also been used in the study of the evolution of black holes. While many aspects of quantum gravity are not understood, some attempts have been made to understand how black holes leak information. Two examples are by Hayden and Preskill [HP07] and Sekino and Susskind [SS08]. We concentrate on the approach in [HP07] here. The idea is that Alice wishes to destroy some quantum information by

throwing it into a black hole. However, Bob has been watching it and storing the Hawking radiation emitted. The question they ask is how long does Bob have to wait before he can recover Alice's information.

Imagine that Alice's information is maximally entangled with a system N held by Charlie. Should Bob acquire a state from the emitted radiation that is maximally entangled with N then we say he has successfully recovered Alice's information. Decoupling is used because, if what remains of the black hole after some evaporation is uncorrelated with N , then the emitted radiation must be maximally entangled with N and Bob has succeeded. We therefore require that the evolution of the black hole produces a decoupling unitary. If the evolution is random then, using Theorem 5.1.2, this will likely happen, provided enough radiation has been emitted. In fact, if Bob holds a system that is maximally entangled with the black hole's internal state before Alice throws in her message, then he can recover her state with fidelity $1 - 2^{-c}$ by reading in only the $k + c$ qubits emitted after Alice deposits her information, where k is the number of qubits in Alice's message.

This model is not physically realistic because most unitaries cannot be implemented efficiently so the black hole would take far too long to apply the decoupling unitary. However, as we said above, only a 2-design is required. In fact, in [HP07] they consider the case that the evolution of a black hole is a local random quantum circuit. This is similar to the random circuits discussed in Chapter 3 except they assume that the unitaries are only applied to nearest-neighbour qubits. Should the random circuit converge to a 2-design quick enough (as Hayden and Preskill conjecture) then the evolution will be sufficiently fast for Bob to find Alice's state. While our results do not prove this they could readily be extended to cover the local case considered here.

5.1.6 Applications for Larger k

So far we have only used k -designs for $k \leq 4$. However, the higher k is the more similar a k -design is to a random unitary. In the next section we consider replacing

random unitaries with k -designs in large deviation bounds, thus finding applications for larger k .

5.2 Derandomising Large Deviation Bounds

The remainder of this chapter has been published previously as [Low09a].

There are many results in quantum information theory that show generic properties of states or unitaries (e.g. [HLW06, HLSW04]). Often, these results say that, with high probability, a random state or unitary has some property, for example high entropy. However, as we have seen above, neither random unitaries nor random states can be implemented efficiently. This limits the usefulness of such results since no physical systems will behave truly randomly. To make such results more physically relevant, it would be desirable to show that these properties are generic properties of unitaries from some natural distribution that can be implemented efficiently. Only then could we conclude that we would expect to see such properties in natural systems.

In many cases, the generic properties of unitaries are desirable but randomised constructions given by the large deviation bounds are inefficient. We would like to come up with distributions which can be implemented efficiently that have similar generic properties. One example where the best known construction is an inefficient randomised one is the ∞ -norm randomising map (see Section 5.1.1). Another example is locking of classical correlations [DHL⁺04, HLSW04], which is a quantum phenomenon whereby a small amount of communication can greatly enhance the classical correlation between two parties. To prove the randomised constructions, the authors show that, with some non-zero probability, random unitaries have the required property. However, there are no known efficient constructions of unitaries with these properties. If, on the other hand, we could show that unitaries drawn randomly from a set that can be implemented efficiently have the property with non-zero probability, we could move an important step closer to finding efficient constructions. (It would not actually provide an efficient construction unless we could find an efficient sampling method.) In fact, for the case of ∞ -norm randomisation, this was done by Aubrun in

[Aub09].

In this section we continue the theme of replacing the Haar measure with a k -design. The reason for using k -designs is two-fold. Firstly, because the first k moments are the same we would expect similar (although weaker) measure concentration results. Secondly, for $k = \text{poly}(n)$ (when the design is on n qubits), we might expect to be able to implement the k -design efficiently (i.e. in $\text{poly}(n)$ time). Indeed, for $k = O(n/\log n)$, we can use the construction from Chapter 4, provided we allow for approximate designs. However, in the applications we consider here we can always make the approximation good enough to make the error negligible.

Not only can k -designs be constructed efficiently, they may even be the product of generic dynamics. In Chapter 3, we show that random quantum circuits quickly converge to a 2-design for a quite general model of such circuits. We also conjecture in Chapter 3 that random circuits give k -designs for $k > 2$ and $k = \text{poly}(n)$ in polynomial time. If a physical system can be accurately modelled by a random circuit then, assuming this conjecture, the naturally occurring states will be k -designs rather than fully random states.

We now summarise some related results in this area. Smith and Leung [SL06] and Dahlsten and Plenio [DP06] found large deviation bounds for stabiliser states. They showed that, in certain regimes, stabiliser states are very likely to have large entanglement. Stabiliser states are state 2-designs so our results can be seen as a generalisation of this to k -designs for $k > 2$ and to other problems. There are also some recent classical results related to the present work. Alon and Nussboim [AN08] consider replacing full randomness with k -wise independence, a classical analogue of k -designs, in random graph theory. They show that k -wise independent random graphs with $k = \text{poly}(\log N)$ (N is the number of vertices) have similar generic properties to fully random graphs.

In the remainder of this chapter, unless otherwise stated, we will use the definition of an ϵ -approximate unitary design given in terms of monomials, as in Definition 2.2.13. Using the tensor product expander construction of Chapter 4 together with

Lemma 2.2.14 gives an efficient construction for $k = O(\log d / \log \log d)$ for this definition.

5.2.1 Introductory Problem: Entanglement of a 2-design

We now illustrate our main idea by showing a large deviation bound for the entanglement of a 2-design, but in a different way to [SL06, DP06].

It has been known for a long time that random states are highly entangled across any bipartition [Pag93, FK94, San95]. Further, in [HLW06], it is shown that random unitaries generate almost maximally entangled states with high probability. However, generating random states is inefficient so it is an interesting question to ask if random efficiently obtainable states are highly entangled.

Let the system be $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_E$, where we label the two systems S and E . Let the dimensions be d_S and d_E and $d = d_S d_E$. Let the overall initial state be any fixed pure state ρ_0 . Then consider applying a random unitary U to SE to get the state $\psi = U \rho_0 U^\dagger$. Then the von Neumann entropy $S(\psi_S) = -\text{tr} \psi_S \log \psi_S$ of the reduced state $\psi_S = \text{tr}_E \psi$ is close to $\log_2 d_S$ (the maximal) with high probability:

Theorem 5.2.1 ([HLW06] Theorem 3.3). *Let $d_E \geq d_S \geq 3$. Then for unitaries chosen from the Haar measure*

$$\mathbb{P}(S(\psi_S) \leq \log_2 d_S - \alpha - \beta) \leq \exp\left(-\frac{(d-1)C\alpha^2}{(\log_2 d_S)^2}\right) \quad (5.2.1)$$

where $C = \frac{1}{8\pi^2}$ and $\beta = \frac{1}{\ln 2} \frac{d_S}{d_E}$.

Now, consider choosing the unitary from a 2-design instead. Later on (Lemma 5.4.1), we show that $\mathbb{E} \text{tr} \psi_S^2 = \frac{d_S + d_E}{d + 1} =: \mu$. Since purity is a polynomial of degree 2, it does not matter if we take the expectation over the Haar measure or the 2-design.

We now apply Markov's inequality:

$$\begin{aligned} \mathbb{P}(\operatorname{tr} \psi_S^2 \geq \mu\gamma) &\leq \frac{\mathbb{E} \operatorname{tr} \psi_S^2}{\mu\gamma} \\ &= \frac{1}{\gamma}. \end{aligned}$$

Using the bound $S(\psi_S) \geq -\log_2 \operatorname{tr} \psi_S^2$ and some manipulations (the details are in Section 5.4), this can be written as

$$\mathbb{P}(S(\psi_S) \leq \log_2 d_S - \alpha - \beta) \leq 2^{-\alpha} \quad (5.2.2)$$

where β is as in Theorem 5.2.1. This bound is much weaker than the bound in Theorem 5.2.1 and, in particular, does not show stronger concentration as d increases. Later in the chapter, we will show that choosing unitaries from a k -design with larger k will give a much stronger bound that does give sharp concentration results for large d .

5.2.2 Main Results

We will now state our main results.

Our most general result is:

Theorem 5.2.2. *Let f be a polynomial of degree K . Let $f(U) = \sum_i \alpha_i M_i(U)$ where $M_i(U)$ are monomials and let $\alpha(f) = \sum_i |\alpha_i|$. Suppose that f has probability concentration*

$$\mathbb{P}_{U \sim \mathcal{U}(d)}(|f - \mu| \geq \delta) \leq C e^{-a\delta^2} \quad (5.2.3)$$

and let ν be an ϵ -approximate unitary k -design. Then

$$\mathbb{P}_{U \sim \nu}(|f - \mu| \geq \delta) \leq \frac{1}{\delta^{2m}} \left(C \left(\frac{m}{a} \right)^m + \frac{\epsilon}{d^k} (\alpha + |\mu|)^{2m} \right) \quad (5.2.4)$$

for integer m with $2mK \leq k$.

We therefore take a bound for Haar random unitaries of the form Eqn. 5.2.3 and

turn it into a bound for k -designs. Often, we will use Levy's Lemma (Lemma 5.3.2) to give the initial concentration bound in Eqn. 5.2.3. In this case, $a = \Theta(d)$ (provided the Lipschitz constant (see later) is constant).

We then apply this to entropy, as a generalisation of Section 5.2.1. We go via the 2-norm since the entropy function is not a polynomial. We find

Theorem 5.2.3. *Let ν be a 4^{-n^2} -approximate unitary $\frac{n}{10 \log_2 n}$ -design on dimension 2^n with $n \geq 19$. Let $d_S d_E = 2^n$ and $2 \leq d_S \leq 2^{n/10}$ and $\alpha \geq 2$. Then*

$$\mathbb{P}_{U \sim \nu}(S(\psi_S) \leq \log_2 d_S - \alpha - \beta) \leq 8 \exp_2 \left(-\frac{n}{80 \log_2 n} \left(\frac{n}{5} + \alpha \right) \right) \quad (5.2.5)$$

where $\beta = \frac{1}{\ln 2} \frac{d_S}{d_E}$ and \exp_2 is the exponential function base 2.

We choose a k -design for $k = \frac{n}{10 \log_2 n}$ since this is (up to constants) the largest k for which we have an efficient unitary k -design construction (using the construction of Chapter 4).

We then move on to apply our results to ideas in statistical mechanics from Popescu et al. [PSW06]. In this paper, the authors show that, for almost all pure states of the universe, any subsystem is very close to the canonical state, which is the state obtained by assuming a uniform distribution over all allowed states of the universe (defined in Eqn. 5.5.2). This could be achieved if the dynamics of the universe produced a random unitary, but this would take exponential time in the size of the universe. We show that the random unitary can be replaced by a k -design, showing that the canonical state can be reached in polynomial time:

Theorem 5.2.4. *Let Ω_S be the canonical state of the system (defined in Eqn. 5.5.2) and ρ_S be the state after choosing a unitary from an ϵ -approximate k -design. Let d_R be the dimension of the universe's Hilbert space subject to the arbitrary constraint R (normally this will be a total energy constraint). Then for $\epsilon \leq \frac{3}{2} \left(\frac{4d_S^3}{d_R} \right)^{k/8}$, $k \leq \frac{4d_S^2}{9\pi^3}$*

$$\mathbb{P}_{U \sim \nu}(\|\rho_S - \Omega_S\|_1 \geq \delta) \leq 6 \left(\frac{4d_S^3}{d_R \delta^2} \right)^{k/8}. \quad (5.2.6)$$

Finally, we use results from [GFE09] to show that most states in an $O(1)$ -approximate state n^2 -design on n qubits are useless for measurement-based quantum computing, in the sense that any computation using such states could be simulated efficiently on a classical computer. We do this, following [GFE09], by showing that the states are so entangled that the measurement outcomes are essentially random.

5.2.3 Optimality of Results

An important question is how close our results are to optimal, in terms of their scaling with dimension d . In Theorem 5.2.2, we will normally have $a = \Theta(d)$ so for m constant, we obtain polynomial bounds, rather than the exponential bounds for full randomness. This is to be expected:

Theorem 5.2.5. *Let ν be an ϵ -approximate unitary k -design. Suppose also that it is discrete i.e. contains a finite number S of unitaries. Let $f(U)$ be any function on matrix elements of U and μ be any constant. Then either $f(U) = \mu$ for all U in ν or for some $\delta > 0$*

$$\mathbb{P}_{U \sim \nu}(|f - \mu| \geq \delta) \geq p_{\min} \quad (5.2.7)$$

where p_{\min} is the probability of choosing the least probable unitary from ν . If the probability is uniform, $p_{\min} = 1/S$.

Proof. There exists at least one U such that $|f(U) - \mu| \geq \delta$ for some $\delta > 0$; the probability of selecting one such U is at least p_{\min} . \square

Corollary 5.2.6. *Our results are polynomially related to the optimal (i.e. the optimal bounds can be obtained by raising ours to a constant power).*

Proof. Our results apply for any design, so must obey the bound in Theorem 5.2.5 for all designs. The unitary design construction we use (from Chapter 4 using Lemma 2.2.14) has $p_{\min} = d^{-O(k)}$ hence the bounds cannot scale better than this. \square

We can also almost recover the tail bound for full randomness in Theorem 5.2.2.

Suppose for simplicity that we have an exact design (i.e. $\epsilon = 0$), so that

$$\mathbb{P}_{U \sim \nu}(|f - \mu| \geq \delta) \leq C \left(\frac{m}{a\delta^2} \right)^m.$$

The optimal m is $a\delta^2/e$, which gives

$$\mathbb{P}_{U \sim \nu}(|f - \mu| \geq \delta) \leq C e^{-a\delta^2/e}.$$

So our result allows us to interpolate from Markov's inequality, which gives weak bounds, all the way to full Haar randomness and is within a polynomial correction of optimal for the full range.

The remainder of the chapter is organised as follows. In Section 5.3 we present our main technique for finding large deviation bounds for k -designs. We then apply this to entropy in Section 5.4, to ideas in statistical mechanics in Section 5.5 and to using k -designs for measurement-based quantum computing in Section 5.6. We then conclude in Section 5.7.

5.3 Main Technique

The main idea in this chapter can be summarised in three steps. Let $f : \mathcal{U}(d) \rightarrow \mathbb{C}$ be a balanced polynomial of degree K in the matrix elements of a unitary U . Then to get a concentration bound on f when U is chosen from a k -design:

1. Find some measure concentration result for $|f(U) - \mu|$ when the unitaries are chosen uniformly at random from the Haar measure. Normally μ will be the expectation of f .
2. Use this to bound the moments $\mathbb{E}|f(U) - \mu|^{2m}$ for some integer $m \leq \frac{k}{2K}$.
3. Then use Markov's inequality and the fact that for a (approximate) k -design the moments are (almost) the same as for uniform randomness. We then optimise the bound for m , which will often involve setting m close to the maximum,

$$\lfloor \frac{k}{2K} \rfloor.$$

We will now work through each of these steps and finish with a proof of Theorem 5.2.2.

5.3.1 Step 1: Concentration for uniform randomness

For the first step, we will often start with Levy's Lemma. This states, roughly speaking, that slowly varying functions in high dimensions are approximately constant. We quantify 'slowly varying' by the Lipschitz constant:

Definition 5.3.1. *The Lipschitz constant η (with respect to the Euclidean norm) for a function f is*

$$\eta = \sup_{U_1, U_2} \frac{|f(U_1) - f(U_2)|}{\|U_1 - U_2\|_2}. \quad (5.3.1)$$

Then we have Levy's lemma:

Lemma 5.3.2 (Levy, see e.g. [Led01]). *Let f be an η -Lipschitz function on $U(d)$ with mean $\mathbb{E}f$. Then*

$$\mathbb{P}(|f - \mathbb{E}f| \geq \delta) \leq 4 \exp\left(-\frac{C_1 d \delta^2}{\eta^2}\right) \quad (5.3.2)$$

where C_1 can be taken to be $\frac{2}{9\pi^3}$.

5.3.2 Step 2: A bound on the moments

Levy's Lemma says that f is close to its mean. This means that $\mathbb{E}|f - \mathbb{E}f|^m$ should be small. We will bound the moments for slightly more general concentration results:

Lemma 5.3.3. *Let X be any random variable with probability concentration*

$$\mathbb{P}(|X - \mu| \geq \delta) \leq C e^{-a\delta^2}. \quad (5.3.3)$$

(Normally μ will be the expectation of X , although the bound does not assume this.)

Then

$$\mathbb{E}|X - \mu|^m \leq C\Gamma(m/2 + 1)a^{-m/2} \leq C\left(\frac{m}{2a}\right)^{m/2} \quad (5.3.4)$$

for any $m > 0$.

Proof. This proof is based on the proof of an analogous result by Bellare and Rompel [BR94], Lemma A.1.

Note that, for any random variable $Y \geq 0$,

$$\mathbb{E}Y = \int_0^\infty \mathbb{P}(Y \geq y)dy. \quad (5.3.5)$$

Therefore

$$\begin{aligned} \mathbb{E}|X - \mu|^m &= \int_0^\infty \mathbb{P}(|X - \mu|^m \geq x)dx \\ &= \int_0^\infty \mathbb{P}(|X - \mu| \geq x^{1/m})dx \\ &\leq C \int_0^\infty \exp(-ax^{2/m})dx \end{aligned}$$

where in the last line we used the assumed large deviation bound Eqn. 5.3.3. To evaluate this integral, use the change of variables $y = ax^{2/m}$ to get

$$\begin{aligned} \mathbb{E}|X - \mu|^m &\leq \frac{Cm}{2}a^{-m/2} \int_0^\infty e^{-y}y^{m/2-1}dy \\ &= Ca^{-m/2}\Gamma(m/2 + 1) \\ &\leq C \left(\frac{m}{2a}\right)^{m/2}. \quad \square \end{aligned}$$

5.3.3 Step 3: A concentration bound for a k -design

Now we show how to obtain a measure concentration result for polynomials when the unitaries are selected from an approximate k -design. We first show that the moments of $|f - \mu|$ for f a polynomial are close to the Haar measure moments:

Lemma 5.3.4. *Let f be a balanced polynomial of degree K and μ be any constant. Let $f = \sum_{i=1}^t \alpha_i M_i$ where each M_i is a monomial. Let $\alpha(f) = \sum_i |\alpha_i|$. Then for m*

an integer with $2mK \leq k$ and ν an ϵ -approximate k -design,

$$\mathbb{E}_{U \sim \nu} |f - \mu|^{2m} \leq \mathbb{E}_{U \sim \mathcal{U}(d)} |f - \mu|^{2m} + \frac{\epsilon}{d^k} (\alpha + |\mu|)^{2m}. \quad (5.3.6)$$

Proof. For simplicity, we assume that f and μ are real. Our proof easily generalises to the complex case.

Firstly we calculate $|\mathbb{E}_{U \sim \nu} f^i - \mathbb{E}_{U \sim \mathcal{U}(d)} f^i|$ using the multinomial theorem:

$$\begin{aligned} & |\mathbb{E}_{U \sim \nu} f^i - \mathbb{E}_{U \sim \mathcal{U}(d)} f^i| \\ &= \left| \sum_{k_1 + \dots + k_t = i} \binom{i}{k_1, \dots, k_t} \alpha_1^{k_1} \dots \alpha_t^{k_t} (\mathbb{E}_{U \sim \nu} - \mathbb{E}_{U \sim \mathcal{U}(d)}) M_1^{k_1} \dots M_t^{k_t} \right| \\ &\leq \sum_{k_1 + \dots + k_t = i} \binom{i}{k_1, \dots, k_t} |\alpha_1|^{k_1} \dots |\alpha_t|^{k_t} |(\mathbb{E}_{U \sim \nu} - \mathbb{E}_{U \sim \mathcal{U}(d)}) M_1^{k_1} \dots M_t^{k_t}| \\ &\leq \frac{\epsilon}{d^k} \sum_{k_1 + \dots + k_t = i} \binom{i}{k_1, \dots, k_t} |\alpha_1|^{k_1} \dots |\alpha_t|^{k_t} \\ &= \frac{\epsilon}{d^k} \alpha^i. \end{aligned}$$

We now calculate $\mathbb{E}_{U \sim \nu} |f - \mu|^{2m}$:

$$\begin{aligned} |\mathbb{E}_{U \sim \nu} |f - \mu|^{2m} - \mathbb{E}_{U \sim \mathcal{U}(d)} |f - \mu|^{2m}| &= |\mathbb{E}_{U \sim \nu} (f - \mu)^{2m} - \mathbb{E}_{U \sim \mathcal{U}(d)} (f - \mu)^{2m}| \\ &= \left| \sum_{i=0}^{2m} \binom{2m}{i} (\mathbb{E}_{U \sim \nu} f^i - \mathbb{E}_{U \sim \mathcal{U}(d)} f^i) (-\mu)^{2m-i} \right| \\ &\leq \sum_{i=0}^{2m} \binom{2m}{i} |\mathbb{E}_{U \sim \nu} f^i - \mathbb{E}_{U \sim \mathcal{U}(d)} f^i| |\mu|^{2m-i} \\ &\leq \frac{\epsilon}{d^k} \sum_{i=0}^{2m} \binom{2m}{i} \alpha^i |\mu|^{2m-i} \\ &= \frac{\epsilon}{d^k} (\alpha + |\mu|)^{2m}. \quad \square \end{aligned}$$

Now we can simply apply Markov's inequality to prove Theorem 5.2.2.

Proof of Theorem 5.2.2. Apply Markov's inequality and Lemmas 5.3.3 and 5.3.4:

$$\begin{aligned}
\mathbb{P}_{U \sim \nu}(|f - \mu| \geq \delta) &= \mathbb{P}_{U \sim \nu}(|f - \mu|^{2m} \geq \delta^{2m}) \\
&\leq \frac{\mathbb{E}_{U \sim \nu} |f - \mu|^{2m}}{\delta^{2m}} \\
&\leq \frac{1}{\delta^{2m}} \left(C \left(\frac{m}{a} \right)^m + \frac{\epsilon}{d^k} (\alpha + |\mu|)^{2m} \right). \quad \square
\end{aligned}$$

We finish this section with two remarks. Firstly, provided $\alpha(f)$ (the sum of the absolute value of all the coefficients) is at most polynomially large in d , we can choose ϵ to be polynomially small to cancel this at no change to the asymptotic efficiency. Secondly, when applying the theorem we will optimise the choice of m (and normally choose $k = 2mK$). Often $a = \Theta(d)$ and the optimal choice of m is often $\Theta(d)$ as well. However, we will not take m so large because we can only implement an efficient k -design for $k = O(\log d / \log \log d)$.

5.4 Application 1: Entropy of a k -design

We now apply the above to show that most unitaries in a k -design generate large amounts of entropy across any bipartition, provided the dimensions are sufficiently far apart. This means that, for any initial state, for most choices of a unitary from a k -design applied to the state, the resulting state will be highly entangled. We go via the purity of the reduced density matrix, since the entropy function is not a polynomial.

We will call the two systems S (the 'system') and E (the 'environment') and calculate the purity of the reduced state. That the purity, $\text{tr} \left[\left(\text{tr}_E U \rho U^\dagger \right)^2 \right]$, is a balanced polynomial of degree 2 is easily seen by noting that the trace is linear and the reduced state is squared. However, we should check that there are not too many terms or terms with large coefficients. To do this, we should calculate α to apply Theorem 5.2.2.

There is a general method for calculating $\alpha(f)$ which we will use. Write $f(U) = \sum_i \alpha_i M_i(U)$ for monomials M_i . To evaluate $\alpha(f) = \sum_i |\alpha_i|$, calculate $f(A)$ where A

is the matrix with all entries equal to 1 (so that $M_i(A) = 1$) and replace α_i with $|\alpha_i|$. Using this here we find

$$\begin{aligned}\alpha &= d^2 \left(\sum_{ij} |\rho_{ij}| \right)^2 \\ &\leq d^4 \sum_{ij} |\rho_{ij}|^2 \\ &= d^4 \|\rho\|_2^2 \\ &\leq d^4.\end{aligned}$$

We now calculate the expected purity:

Lemma 5.4.1. *The expected purity of the reduced state is $\frac{d_S + d_E}{d+1}$, where d_S is the dimension of subsystem S and $d_E = d/d_S$ is the dimension of subsystem E .*

Proof. We have

$$\mathbb{E}_{U \sim \mathcal{U}(d)} \|\psi_S\|_2^2 = \mathbb{E}_{U \sim \mathcal{U}(d)} \left[\text{tr } \mathcal{F}_{S_1 S_2} (\text{tr}_E U \rho U^\dagger \otimes \text{tr}_E U \rho U^\dagger) \right] \quad (5.4.1)$$

where $\mathcal{F}_{S_1 S_2}$ is swap acting between systems S_1 and S_2 . By linearity of the trace, we can commute the $\mathbb{E}_{U \sim \mathcal{U}(d)}$ through and use $\mathbb{E}_{U \sim \mathcal{U}(d)} [U \rho U^\dagger \otimes U \rho U^\dagger] = \frac{I_{12} + \mathcal{F}_{12}}{d(d+1)}$ to find

$$\begin{aligned}\mathbb{E}_{U \sim \mathcal{U}(d)} \|\psi_S\|_2^2 &= \text{tr} \left[\frac{\mathcal{F}_{S_1 S_2}}{d(d+1)} (d_E^2 I_{S_1 S_2} + d_E \mathcal{F}_{S_1 S_2}) \right] \\ &= \frac{1}{d(d+1)} (d_E^2 d_S + d_E d_S^2) \\ &= \frac{d_S + d_E}{d+1} \quad \square\end{aligned}$$

Working out the higher moments in this way is difficult (although has been done in [Gir07]) so we use Levy's Lemma and Lemma 5.3.3. To use Levy's Lemma, all we have to do is find the Lipschitz constant for the purity:

Lemma 5.4.2. *The Lipschitz constant for purity is ≤ 2 .*

Proof.

$$\begin{aligned}\eta &= \sup_{\psi, \phi} \frac{|||\psi_S||_2^2 - ||\phi_S||_2^2|}{\|\psi - \phi\|_2} \\ &= \sup_{\psi, \phi} \frac{|||\psi_S||_2 - ||\phi_S||_2| (||\psi_S||_2 + ||\phi_S||_2)}{\|\psi - \phi\|_2}\end{aligned}$$

Now we use $|||S||_2 - |||T||_2| \leq \|S - T\|_2$ to find

$$\eta \leq \sup_{\psi, \phi} (||\psi_S||_2 + ||\phi_S||_2) \leq 2$$

using the fact that the purity is upper bounded by 1. \square

Lemma 5.4.3. For $\mu = \frac{d_S + d_E}{d+1}$ and m an integer with $m \leq k/4$ and ν an ϵ -approximate k -design,

$$\mathbb{P}_{U \sim \nu}(S(\psi_S) \leq -\log_2 \mu - \alpha) \leq \frac{1}{(\mu(2^\alpha - 1))^{2m}} \left(4 \left(\frac{4m}{C_1 d} \right)^m + \frac{\epsilon}{d^k} (d^4 + \mu)^{2m} \right). \quad (5.4.2)$$

Proof. We use the fact that von Neumann entropy is lower bounded by the Renyi 2-entropy i.e. $-\log_2 ||\psi_S||_2^2$:

$$S(\psi_S) \geq S_2(\psi_S) = -\log_2 ||\psi_S||_2^2. \quad (5.4.3)$$

Then

$$\begin{aligned}\mathbb{P}_{U \sim \nu}(S(\psi_S) \leq -\log_2(1 + \delta)\mu) &\leq \mathbb{P}_{U \sim \nu}(S_2(\psi_S) \leq -\log_2(1 + \delta)\mu) \\ &= \mathbb{P}_{U \sim \nu}(||\psi_S||_2^2 \geq (1 + \delta)\mu) \\ &\leq \mathbb{P}_{U \sim \nu}(||\psi_S||_2^2 - \mu| \geq \delta\mu) \\ &\leq \frac{1}{(\mu\delta)^{2m}} \left(4 \left(\frac{4m}{C_1 d} \right)^m + \frac{\epsilon}{d^k} (d^4 + \mu)^{2m} \right)\end{aligned}$$

using Theorem 5.2.2 in the last line. \square

We have written this in a more convenient form in Theorem 5.2.3 which is proved

in Section 5.8. This is to be compared with the Haar random version Theorem 5.2.1. As expected, we have $n = \log_2 d$ appearing in the exponent rather than d . Note also that our bound does not work well for $d_S \approx d_E$. In fact, in this case, we do not get a bound that improves with dimension. In order to achieve such a bound in this regime a different technique will be necessary.

5.5 Application 2: k -designs and Statistical Mechanics

We can also apply these ideas to partially derandomise some of the arguments on the foundations of statistical mechanics in [PSW06]. In this paper, the authors develop the idea that the uncertainty in statistical mechanics comes from entanglement rather than the traditional assumption of the principle of equal a priori probabilities. They consider the universe being in a pure quantum state and that the uncertainty in the state of a subsystem comes from the entanglement between this system and the rest of the universe.

The setting is that there is an arbitrary global linear constraint R . Often this will be a total energy constraint although this is not assumed. Let the Hilbert space of states satisfying R be \mathcal{H}_R . Then let the system and environment Hilbert spaces be \mathcal{H}_S and \mathcal{H}_E respectively. Then

$$\mathcal{H}_R \subseteq \mathcal{H}_S \otimes \mathcal{H}_E. \quad (5.5.1)$$

Let the dimensions be d_R , d_S and d_E and let $\mathcal{E}_R = \frac{I_R}{d_R}$. Note that $d_R \leq d_S d_E$, unlike in the above where we took $d = d_S d_E$. Normally we will have $d_S \ll d_R$. The principle of equal a priori probabilities says that the state of the universe is \mathcal{E}_R which implies the subsystem state is the canonical state, given by

$$\Omega_S = \text{tr}_E(\mathcal{E}_R). \quad (5.5.2)$$

The main result of [PSW06] (the ‘principle of *apparently* equal a priori probabilities’)

is that, for almost all pure states of the universe, the subsystem state is almost exactly the canonical state.

Theorem 5.5.1 (Theorem 1 of [PSW06]). *For a randomly chosen state $|\phi\rangle \in \mathcal{H}_R \subseteq \mathcal{H}_S \otimes \mathcal{H}_E$ and arbitrary $\epsilon > 0$, the distance between the reduced density matrix of the system $\rho_S = \text{tr}_E(|\phi\rangle\langle\phi|)$ and the canonical state Ω_S (Eqn. 5.5.2) is given probabilistically by*

$$\mathbb{P}_{U \sim \mathcal{U}(d)} \left(\|\rho_S - \Omega_S\|_1 \geq \epsilon + \sqrt{\frac{d_S}{d_E^{\text{eff}}}} \right) \leq 2 \exp(-C_2 d_R \epsilon^2) \quad (5.5.3)$$

where $C_2 = 1/(18\pi^3)$ and $d_E^{\text{eff}} = \frac{1}{\text{tr} \Omega_E^2} \geq \frac{d_R}{d_S}$.

This result gives compelling evidence to replace the principle of equal a priori probabilities with the principle of apparently equal a priori probabilities, but it does not address the problem of how the system reaches this state. It will take an extremely (exponentially) long time for the universe to reach a random pure state, in contrast to the observed fact that thermalisation occurs quickly. Here, we show that for almost all unitaries in a k -design applied to the universe, the subsystem state is close to the canonical state. Since these unitaries can be implemented and sampled efficiently, this means that equilibrium could be reached quickly to match observations.

We are now ready to show that a k -design gives a small $\|\rho_S - \Omega_S\|_1$. First, we have to modify Lemma 5.3.3 slightly:

Lemma 5.5.2. *Let X be any non-negative random variable with probability concentration*

$$\mathbb{P}(X \geq \delta + \eta) \leq C e^{-a\delta^2} \quad (5.5.4)$$

where $\eta \geq 0$. Then

$$\mathbb{E}X^m \leq C \left(\frac{2m}{a} \right)^{m/2} + (2\eta)^m \quad (5.5.5)$$

for any $m > 0$.

The proof is very similar to the proof of Lemma 5.3.3.

Now we state and prove the main result in this section:

Theorem 5.5.3. *Let ν be an ϵ -approximate unitary k -design. Then*

$$\mathbb{P}_{U \sim \nu}(\|\rho_S - \Omega_S\|_1 \geq \delta) \leq \left(\frac{d_S}{\delta^2}\right)^{k/8} \left(2 \left(\frac{k}{2C_2 d_R}\right)^{k/8} + \left(\frac{4d_S^2}{d_R}\right)^{k/8} + \frac{\epsilon}{d_R^k} (d_R^2 + 1)^{k/2}\right). \quad (5.5.6)$$

In particular, with $\epsilon = \frac{3}{2} \left(\frac{4d_S^3}{d_R}\right)^{k/8}$, $k \leq 8C_2 d_S^2$,

$$\mathbb{P}_{U \sim \nu}(\|\rho_S - \Omega_S\|_1 \geq \delta) \leq 6 \left(\frac{4d_S^3}{d_R \delta^2}\right)^{k/8}. \quad (5.5.7)$$

Again, we need d_S to be polynomially smaller than d_R to obtain non-trivial bounds.

Proof. We go via the 2-norm and use Lemmas 5.5.2 and 5.3.4.

We have from Theorem 5.5.1 that

$$\mathbb{P}_{U \sim \mathcal{U}(d)}(\|\rho_S - \Omega_S\|_1 \geq \delta + \eta) \leq 2e^{-C_2 d_R \delta^2} \quad (5.5.8)$$

where $\eta = \sqrt{\frac{d_S}{d_E^{\text{eff}}}} \leq \frac{d_S}{\sqrt{d_R}}$. Since $\|\rho_S - \Omega_S\|_2 \leq \|\rho_S - \Omega_S\|_1$,

$$\mathbb{P}_{U \sim \mathcal{U}(d)}(\|\rho_S - \Omega_S\|_2 \geq \delta + \eta) \leq 2e^{-C_2 d_R \delta^2}. \quad (5.5.9)$$

We now apply Lemma 5.5.2 to get

$$\mathbb{E}_{U \sim \mathcal{U}(d)} \|\rho_S - \Omega_S\|_2^{2m} \leq 2 \left(\frac{4m}{C_2 d_R}\right)^m + (2\eta)^{2m}. \quad (5.5.10)$$

So for $m \leq k/4$, using Markov's inequality and Lemma 5.3.4 (with $\mu = 0$) on the polynomial $\|\rho_S - \Omega_S\|_2^2$:

$$\mathbb{P}_{U \sim \nu}(\|\rho_S - \Omega_S\|_2 \geq \delta) \leq \frac{1}{\delta^{2m}} \left(2 \left(\frac{4m}{C_2 d_R}\right)^m + (2\eta)^{2m} + \frac{\epsilon}{d_R^k} (d_R^2 + 1)^{4m}\right). \quad (5.5.11)$$

Here, we used an estimate of α , the sum of the moduli of the coefficients:

$$\alpha \leq (d_R^2 + 1)^2 \quad (5.5.12)$$

which we obtain via a similar calculation to that in Section 5.4.

Now we go back to the 1-norm, using $\|\rho_S - \Omega_S\|_1 \leq \sqrt{d_S} \|\rho_S - \Omega_S\|_2$ to get

$$\mathbb{P}_{U \sim \nu}(\|\rho_S - \Omega_S\|_1 \geq \delta) \leq \mathbb{P}_{U \sim \nu}(\|\rho_S - \Omega_S\|_2 \geq \delta/\sqrt{d_S}) \quad (5.5.13)$$

$$\leq \left(\frac{d_S}{\delta^2}\right)^m \left(2 \left(\frac{4m}{C_2 d_R}\right)^m + (2\eta)^{2m} + \frac{\epsilon}{d_R^k} (d_R^2 + 1)^{4m}\right). \quad (5.5.14)$$

To obtain the result in Eqn. 5.5.6, we just use $\eta \leq \frac{d_S}{\sqrt{d_R}}$ and set $m = k/8$.

To prove the simplified version, first use, as in Section 5.4, that $(d_R^2 + 1)^{4m} \leq 2d_R^{8m}$ for $m \leq d_R^2/8$. This is implied by $k \leq 8C_2 d_S^2$. We then set $m = k/8$ to find

$$\mathbb{P}_{U \sim \nu}(\|\rho_S - \Omega_S\|_1 \geq \delta) \leq 2 \left(\frac{k d_S}{2C_2 d_R \delta^2}\right)^{k/8} + \left(\frac{4d_S^3}{d_R \delta^2}\right)^{k/8} + 2\frac{\epsilon}{\delta^{k/4}}. \quad (5.5.15)$$

Then, using $k \leq 8C_2 d_S^2$, with $\epsilon \leq \frac{3}{2} \left(\frac{4d_S^3}{d_R}\right)^{k/8}$, we obtain the simplified result Eqn. 5.5.7. \square

5.6 Application 3: Using k -designs for Measurement-Based Quantum Computing

Here we apply our ideas to partially derandomise some results of Gross, Flammia and Eisert in [GFE09] and Bremner, Mora and Winter in [BMW09]. The main result in these two papers is that most states do not offer any advantage over classical computation when used in the measurement-based quantum computing (MBQC) model. In MBQC, a classical computer is given access to a large quantum state on which it can do single qubit measurements. Some states allow for universal quantum computation

whereas others do not add any extra power to the classical computer. These results are concerned with the question of characterising which states do and do not work. Showing that random states do not give any speed up shows that useful states for MBQC are not generic and so must be carefully constructed.

While the results in these two papers are similar, we will concentrate on the methods from [GFE09] since their methods are simpler to apply here. They prove their result by showing that most states are very entangled in the geometric measure (see Definition 5.6.1). They then use this to show that the measurement outcomes of even the best possible measurement scheme are almost completely random. In fact, the state could be thrown away and the measurement outcomes replaced with random numbers to solve the computational problem just as efficiently. This shows that you can classically simulate any quantum computation that uses these highly entangled states. The measure of entanglement they use is the geometric measure:

Definition 5.6.1. *The geometric measure of entanglement of a state $|\Psi\rangle$ is [Shi95, BL01]*

$$E_g(|\Psi\rangle) = -\log_2 \sup_{\alpha \in \mathcal{P}} |\langle \alpha | \Psi \rangle|^2. \quad (5.6.1)$$

where \mathcal{P} is the set of all product states.

They show that any MBQC using a state $|\Psi\rangle$ with $E_g(|\Psi\rangle) = n - O(\log_2 n)$ can be efficiently simulated classically. They then show that

Theorem 5.6.2 ([GFE09], Theorem 2). *For $n \geq 11$,*

$$\mathbb{P}_{|\psi\rangle \sim \mathcal{S}(d)}(E_g(|\Psi\rangle) \leq n - 2 \log_2 n - 3) \leq e^{-n^2}. \quad (5.6.2)$$

This shows that most states are useless. We partially derandomise this result to show that most states in an ϵ -approximate (ϵ can be taken as a constant) state n^2 -design have high geometric measure of entanglement and thus are useless in the same way.

We could apply our technique and use Theorem 5.2.2 but in this case, it is simpler

to directly bound the probability using Markov's inequality.

Lemma 5.6.3.

$$\mathbb{P}_{|\Psi\rangle\sim\nu}(|\langle\Phi|\Psi\rangle|^2 \geq \delta) \leq (1+\epsilon)\frac{m!}{(d\delta)^m} \leq (1+\epsilon)\left(\frac{m}{d\delta}\right)^m \quad (5.6.3)$$

where $|\Psi\rangle$ is chosen from an ϵ -approximate state k -design ν , $m \leq k$ and a positive integer and $|\Phi\rangle$ is any fixed state.

Proof. We prove this bound directly using Markov's inequality:

$$\begin{aligned} \mathbb{P}_{|\Psi\rangle\sim\nu}(|\langle\Phi|\Psi\rangle|^2 \geq \delta) &= \mathbb{P}_{|\Psi\rangle\sim\nu}(|\langle\Phi|\Psi\rangle|^{2m} \geq \delta^m) \\ &\leq \frac{\mathbb{E}_{|\Psi\rangle\sim\nu}|\langle\Phi|\Psi\rangle|^{2m}}{\delta^m} \\ &= \frac{\mathbb{E}_{|\Psi\rangle\sim\nu}\langle\Phi|^{\otimes m}|\Psi\rangle^{\otimes m}\langle\Psi|^{\otimes m}|\Phi\rangle^{\otimes m}}{\delta^m} \\ &= \frac{\langle\Phi|^{\otimes m}\mathbb{E}_{|\Psi\rangle\sim\nu}[|\Psi\rangle^{\otimes m}\langle\Psi|^{\otimes m}]|\Phi\rangle^{\otimes m}}{\delta^m} \\ &\leq \frac{\langle\Phi|^{\otimes m}(1+\epsilon)\frac{\Pi_m^{\text{sym}}}{\binom{m+d-1}{d-1}}|\Phi\rangle^{\otimes m}}{\delta^m} \\ &= \frac{1+\epsilon}{\binom{m+d-1}{d-1}\delta^m} \\ &\leq \frac{(1+\epsilon)m!}{(d\delta)^m} \leq (1+\epsilon)\left(\frac{m}{d\delta}\right)^m. \quad \square \end{aligned}$$

We now prove the main result in this section:

Theorem 5.6.4. For $|\Psi\rangle$ randomly drawn from an ϵ -approximate state k -design with $d = 2^n$

$$\mathbb{P}_{|\Psi\rangle\sim\nu}(E_g(|\Psi\rangle) \leq n - \delta) \leq (1+\epsilon)\exp_2(k \log_2 2k + 4n \log_2 10n - k\delta + 4n(n - \delta)). \quad (5.6.4)$$

In particular, for $k = n^2$, $\delta = 3 \log_2 n + 5$ and $\epsilon = 1$,

$$\mathbb{P}_{|\Psi\rangle\sim\nu}(E_g(|\Psi\rangle) \leq n - 3 \log_2 n - 5) \leq 2 \cdot n^{-n^2}. \quad (5.6.5)$$

We note that this bound is almost the same as in Theorem 5.6.2. It only works for slightly larger deviations from n , which is why we obtain a slightly better probability bound. Note also that we can obtain an exponential bound in n (not $d = 2^n$) because the design is exponentially large in n .

Proof. This proof closely mirrors the proof of Theorem 2 in [GFE09]. We use the idea of a γ -net. $\mathcal{N}_{\gamma,n}$ is a γ -net on product states if

$$\sup_{|\alpha\rangle \in \mathcal{P}} \inf_{|\tilde{\alpha}\rangle \in \mathcal{N}_{\gamma,n}} \left\| |\alpha\rangle - |\tilde{\alpha}\rangle \right\|_2 \leq \gamma/2. \quad (5.6.6)$$

In [GFE09], it is shown that such a net exists with $|\mathcal{N}_{\gamma,n}| \leq (5n/\gamma)^{4n}$. We then proceed by showing that most states in the state design have small overlap with every state in the net using the union bound and Lemma 5.6.3. Finally, since every state is close to one in the net, we can show that most states in the design have small overlap with every product state.

We now formalise the above. Using Lemma 5.6.3 and the union bound,

$$\mathbb{P}_{|\Psi\rangle \sim \nu} \left(\sup_{|\tilde{\alpha}\rangle \in \mathcal{N}_{\gamma,n}} |\langle \tilde{\alpha} | \Psi \rangle|^2 \geq \delta'/2 \right) \leq |\mathcal{N}_{\gamma,n}| (1+\epsilon) \left(\frac{2k}{d\delta'} \right)^k \leq \left(\frac{5n}{\gamma} \right)^{4n} (1+\epsilon) \left(\frac{2k}{2^n \delta'} \right)^k. \quad (5.6.7)$$

Now, we need to bound

$$\begin{aligned} \mathbb{P}_{|\Psi\rangle \sim \nu} (E_g(|\Psi\rangle) \leq n - \delta) &= \mathbb{P}_{|\Psi\rangle \sim \nu} \left(-\log_2 \sup_{|\alpha\rangle \in \mathcal{P}} |\langle \alpha | \Psi \rangle|^2 \leq n - \delta \right) \\ &= \mathbb{P}_{|\Psi\rangle \sim \nu} \left(\sup_{|\alpha\rangle \in \mathcal{P}} |\langle \alpha | \Psi \rangle|^2 \geq 2^{-(n-\delta)} \right). \end{aligned}$$

We now claim that

$$\sup_{|\alpha\rangle \in \mathcal{P}} |\langle \alpha | \Psi \rangle|^2 \geq \delta' \Rightarrow \sup_{|\tilde{\alpha}\rangle \in \mathcal{N}_{\delta'/2,n}} |\langle \tilde{\alpha} | \Psi \rangle|^2 \geq \delta'/2. \quad (5.6.8)$$

To prove this claim, let $|\alpha\rangle$ be the state that achieves the supremum on the left hand side, and let $|\tilde{\alpha}\rangle$ be the state closest to it in the $\delta'/2$ -net. It is shown in [GFE09] that

this implies for any $|\Psi\rangle$

$$||\langle\alpha|\Psi\rangle|^2 - |\langle\tilde{\alpha}|\Psi\rangle|^2| \leq \delta'/2. \quad (5.6.9)$$

Therefore

$$\begin{aligned} |\langle\tilde{\alpha}|\Psi\rangle|^2 &\geq |\langle\alpha|\Psi\rangle|^2 - \delta'/2 \\ &\geq \delta'/2. \end{aligned}$$

This implies that the supremum over all states in the net must be at least $\delta'/2$ to prove the claim.

We can now finish the proof. Set $\delta' = 2^{-(n-\delta)}$ in Eqn. 5.6.8 and use Eqn. 5.6.7 with $\gamma = \delta'/2$ to find

$$\begin{aligned} &\mathbb{P}_{|\Psi\rangle\sim\nu} \left(\sup_{|\alpha\rangle\in\mathcal{P}} |\langle\alpha|\Psi\rangle|^2 \geq 2^{-(n-\delta)} \right) \\ &\leq \mathbb{P}_{|\Psi\rangle\sim\nu} \left(\sup_{|\tilde{\alpha}\rangle\in\mathcal{N}_{2^{-(n-\delta)-1},n}} |\langle\tilde{\alpha}|\Psi\rangle|^2 \geq 2^{-(n-\delta)-1} \right) \\ &\leq (1 + \epsilon) \exp_2(k \log_2 2k + 4n \log_2 10n - k\delta + 4n(n - \delta)). \quad \square \end{aligned}$$

Combining this with the arguments of [GFE09] shows that most states in a state n^2 -design on n qubits are useless for MBQC. This shows that even many efficiently preparable states are useless.

5.7 Conclusions

We have seen how to turn large deviation bounds for Haar-random unitaries into bounds for k -designs. The main technique was applied to show that unitaries from k -designs generate large amounts of entanglement. Then we showed that, if the dynamics of the universe produced a k -design, the entanglement generated would be sufficient to reproduce the principle of equal a priori probabilities. Finally we showed that most states in sufficiently large state designs are useless for measurement-based

quantum computing, in the sense that computation using them can be efficiently simulated classically.

However, there are other bounds for which our technique does not work. Since we cannot obtain exponential bounds for polynomially sized designs, our technique cannot directly derandomise some bounds. Some results, for example showing that the ∞ -norm of the reduced state of a random pure state is close to $1/d_S$ [HHL04], are proven by using an ϵ -net of states and the union bound. Since the ϵ -net is exponentially large, exponentially small bounds are required. We do not know how to apply our idea to results of this kind and still have $k = \text{poly}(\log d)$. (Note that we could cope with the ϵ -net in Section 5.6 since it was just a net on product states which is considerably smaller.)

It is also possible that our ideas could be used to completely derandomise some constructions (e.g. locking [HLSW04, DHL⁺04]). If we could show that unitaries drawn from a k -design work with non-zero probability, and come up with an efficient sampling method, then we could obtain efficient randomised constructions.

5.8 Proof of Theorem 5.2.3

Here we prove the more convenient form of Lemma 5.4.3 stated as Theorem 5.2.3.

Proof of Theorem 5.2.3. Firstly, we will write the left hand side of Eqn. 5.4.2 in a more useful way. Using $\ln(1+x) \leq x$, we find

$$-\log_2 \mu \geq \log_2 d_S - \beta$$

where $\beta = \frac{1}{\ln 2} \frac{d_S}{d_E}$, following the notation in [HLW06]. This means

$$\begin{aligned} \mathbb{P}_{U \sim \nu}(S(\psi_S) \leq \log_2 d_S - \alpha - \beta) &\leq \mathbb{P}_{U \sim \nu}(S(\psi_S) \leq -\log_2 \mu - \alpha) \\ &\leq \frac{1}{(\mu(2^\alpha - 1))^{2m}} \left(4 \left(\frac{4m}{C_1 d} \right)^m + \frac{\epsilon}{d^k} (d^4 + \mu)^{2m} \right). \end{aligned}$$

We now simplify the right hand side. Let $\delta = 2^\alpha - 1$. For $d_S \geq 2$, we have $\mu \geq 1/d_S$.

We shall also assume that $m = k/8$. This gives us (using $\mu \leq 1$)

$$\mathbb{P}_{U \sim \nu}(S(\psi_S) \leq \log_2 d_S - \alpha - \beta) \leq \left(\frac{d_S}{\delta}\right)^{k/4} \left(4 \left(\frac{k}{2C_1 d}\right)^{k/8} + \epsilon \left(1 + \frac{1}{d^4}\right)^{k/4}\right). \quad (5.8.1)$$

Now, one can easily show (e.g. by induction on n) that

$$(1 + \delta)^n \leq 2 \quad (5.8.2)$$

for $2n\delta \leq 1$. We use this for $n = k/4$ and $\delta = 1/d^4$. The condition is then $k \leq 2d^4$, which we shall assume (we will set $k = \log d / \log \log d$ later). We now obtain

$$\mathbb{P}_{U \sim \nu}(S(\psi_S) \leq \log_2 d_S - \alpha - \beta) \leq \left(\frac{d_S}{\delta}\right)^{k/4} \left(4 \left(\frac{k}{2C_1 d}\right)^{k/8} + 2\epsilon\right). \quad (5.8.3)$$

We will now take $\epsilon = 2 \left(\frac{k}{2C_1 d}\right)^{k/8}$, so that the two terms are the same. $\log 1/\epsilon$ is $\text{poly}(\log d)$ so this remains efficient. Now

$$\mathbb{P}_{U \sim \nu}(S(\psi_S) \leq \log_2 d_S - \alpha - \beta) \leq 8 \left(\frac{d_S^2 k}{2C_1 d \delta^2}\right)^{k/8}. \quad (5.8.4)$$

Assuming that $\delta^2 > \frac{k d_S^2}{2C_1 d}$, we should take k as large as possible up to $\frac{2C_1 \delta^2 d}{e d_S^2}$, when the right hand side is maximised. We then find the result after further simplification. \square

Part II

Quantum Learning

Chapter 6

Learning and Testing Algorithms for the Clifford Group

6.1 Introduction

A central problem in quantum computing is to determine an unknown quantum state from measurements of multiple copies of the state. This process is known as quantum state tomography (see [NC00] and references therein). By making enough measurements, the probability distributions of the outcomes can be estimated from which the state can be inferred. A related problem is that of quantum process tomography, where an unknown quantum evolution is determined by applying it to certain known input states. There are several methods for doing this, including what are known as Standard Quantum Process Tomography [CN97, PCZ97] and Ancilla Assisted Process Tomography [DLP01, Leu03]. These methods work by using state tomography on the output states for certain input states.

However, all these procedures share one important downside: the number of measurements required increases exponentially with the number of qubits. This already presents problems even with systems achievable with today's technology, for which complete tomographical measurements can take hours (e.g. [HHR⁺05]) making tomography of larger systems unfeasible. Unfortunately this exponential cost is necessary to

determine a completely unknown state or process, since there are exponentially many parameters to measure. To make tomography feasible for larger systems, we need to find a restriction that requires fewer measurements, ideally polynomially many.

One way to improve the measurement, or query, complexity is to assume some prior knowledge of the process. For example, suppose the process was known to be one of a small number of unitaries, then the task is just to decide which. This is the approach we take here. As a simple example, consider being given a black box implementing an unknown Pauli matrix. By applying this to half a maximally entangled state, the Pauli can be identified with one query. This is essentially superdense coding [BW92] and is explained in Section 6.3.1. Indeed, if the black box performed a tensor product of arbitrary Paulis on n qubits then it too can be identified with just one query.

We extend this to work for elements of the Clifford group (the normaliser of the Pauli group; see Definition 6.2.1) and show that any member of the Clifford group can be learnt with $O(n)$ queries, which we show is optimal. The Clifford group is an important subgroup of the unitary group that has found uses in quantum error correction and fault tolerance [CRSS97, Sho96, Got98].

Then generalising further, we show that elements of the Gottesman-Chuang hierarchy [GC99] (see Definition 6.2.2), also known as the \mathcal{C}_k hierarchy, can also be learnt efficiently. As the level k increases, the set \mathcal{C}_k includes more and more unitaries so this implies ever larger sets can be learnt, although the number of queries scales exponentially with k . Our methods also work if the unitary is known to be close to a Clifford (or any element of \mathcal{C}_k for some known k) rather than exactly a Clifford.

We also give a Clifford testing algorithm, which determines whether an unknown unitary is close to a Clifford or far from every Clifford. This is an extension of the Pauli testing algorithm given in [MO08]. Indeed, our results are closely related to results in [MO08] and we use some of the algorithms presented there as ingredients. Our results can also be compared with [Aar07], which contains methods to approximately learn quantum states. Another related result is that of Aaronson and Gottesman [AG09], which provides a method of learning stabiliser states with linearly many copies.

We only consider query complexity although, at least for the Clifford group results, our methods are computationally efficient too.

The rest of the chapter is organised as follows. In Section 6.2, we define the Pauli and Clifford groups and the Gottesman-Chuang hierarchy. In Section 6.3 we present our algorithm for exact learning of Clifford and \mathcal{C}_k elements. In Section 6.4 we show how to find the closest element of \mathcal{C}_k to an unknown unitary. In Section 6.5 we present our Clifford testing algorithm and then conclude in Section 6.6.

This chapter has been published previously as [Low09b].

6.2 The Pauli and Clifford Groups and the Gottesman-Chuang Hierarchy

Firstly, we define the Pauli group. Call the set of all Pauli matrices on n qubits $\hat{\mathcal{P}}$. We then have $|\hat{\mathcal{P}}| = 4^n$. We write matrices in the Pauli basis using the normalisation $\rho = \sum_p \gamma(p)\sigma_p$. To make $\hat{\mathcal{P}}$ into a group, the Pauli group \mathcal{P} , we must include each matrix in $\hat{\mathcal{P}}$ with phases $\{\pm 1, \pm i\}$.

We can now define the Clifford group:

Definition 6.2.1 (The Clifford group). *The Clifford group is the normaliser of the Pauli group i.e.*

$$\mathcal{C} = \{U \in \mathcal{U}(2^n) : U\mathcal{P}U^\dagger \subseteq \mathcal{P}\}.$$

Then the Gottesman-Chuang hierarchy is a generalisation:

Definition 6.2.2 (The Gottesman-Chuang hierarchy [GC99]). *Let \mathcal{C}_1 be the Pauli group \mathcal{P} . Then level \mathcal{C}_k of the hierarchy is defined recursively:*

$$\mathcal{C}_k = \{U \in \mathcal{U}(2^n) : U\mathcal{P}U^\dagger \subseteq \mathcal{C}_{k-1}\}.$$

By definition, \mathcal{C}_2 is the Clifford group \mathcal{C} . For $k > 2$, \mathcal{C}_k is no longer a group but contains a universal gate set, whereas \mathcal{C}_1 and \mathcal{C}_2 are not universal.

6.3 Learning Gottesman-Chuang Operations

Before we give our algorithm for learning Gottesman-Chuang operations, we present a simple method for learning Pauli operations, which we use as the main ingredient.

6.3.1 Learning Pauli Operations

This is due to [MO08] and is in fact identical to the superdense-coding protocol [BW92].

Theorem 6.3.1 ([MO08], Proposition 20). *Pauli operations can be identified with one query and in time $O(n)$.*

Proof. Apply the operator σ_p to half of the maximally entangled state

$$|\psi\rangle = 2^{-n/2} \sum_i |ii\rangle.$$

For different choices of σ_p , the resulting states are orthogonal so can be perfectly distinguished:

$$\begin{aligned} \langle\psi| (\sigma_p \otimes I) (\sigma_q \otimes I) |\psi\rangle &= 2^{-n} \sum_{ij} \langle ii | \sigma_p \sigma_q \otimes I | jj \rangle \\ &= 2^{-n} \sum_{ij} \langle i | \sigma_p \sigma_q | j \rangle \langle i | j \rangle \\ &= 2^{-n} \sum_i \langle i | \sigma_p \sigma_q | i \rangle \\ &= 2^{-n} \text{tr } \sigma_p \sigma_q \\ &= \delta_{pq}. \end{aligned}$$

The time complexity $O(n)$ comes from the preparation and measurement operations. □

6.3.2 Learning Clifford Operations

We can now present our algorithm for learning Clifford operations to illustrate our main idea for learning unitaries in the Gottesman-Chuang hierarchy. We will use the fact that knowing how a unitary acts by conjugation on all elements of $\hat{\mathcal{P}}$ identifies it uniquely (up to phase):

Lemma 6.3.2. *Knowing $U\sigma_p U^\dagger$ for all $\sigma_p \in \hat{\mathcal{P}}$ uniquely determines U , up to global phase.*

Proof. The Pauli matrices form a basis for all $2^n \times 2^n$ matrices so knowing the action of U on the Paulis is enough to determine the action of U on any matrix up to phase. The phase cannot be determined because action by conjugation does not reveal the phase. \square

Now let $G = \{\sigma_{x_i}, \sigma_{z_i}\}_{i=1}^n$ where σ_{x_i} (σ_{z_i}) is the matrix with σ_x (σ_z) acting on qubit i and trivially elsewhere. We think of this as a set of generators for $\hat{\mathcal{P}}$ since each element of $\hat{\mathcal{P}}$ can be written as a product of elements of G , up to phase. Using this, knowledge of how U acts on elements of G is sufficient to determine the action on all of $\hat{\mathcal{P}}$:

Lemma 6.3.3. *$U\sigma_p U^\dagger$ for any $\sigma_p \in \hat{\mathcal{P}}$ can be calculated from knowledge of $U\sigma_g U^\dagger$ for each $\sigma_g \in G$.*

Proof. Let $\sigma_p = \alpha\sigma_{g_1} \dots \sigma_{g_m}$ for $\sigma_{g_i} \in G$ where α is a phase. Then

$$U\sigma_p U^\dagger = \alpha U\sigma_{g_1} \dots \sigma_{g_m} U^\dagger = \alpha U\sigma_{g_1} U^\dagger \dots U\sigma_{g_m} U^\dagger. \quad \square$$

With these definitions and observations, we can now present the Clifford learning algorithm.

Theorem 6.3.4. *Given oracle access to an unknown Clifford operation C and its conjugate C^\dagger , C can be determined exactly (up to global phase) with $2n + 1$ queries to C and $2n$ to C^\dagger . The algorithm runs in time $O(n^2)$.*

Proof. From the definition of the Clifford group, $C\sigma_p C^\dagger \in \mathcal{P}$ for all $\sigma_p \in \hat{\mathcal{P}}$. Note that $C\sigma_p C^\dagger$ is not necessarily a Pauli operator in $\hat{\mathcal{P}}$ because there is a phase of ± 1 (complex phases are not allowed because $C\sigma_p C^\dagger$ is Hermitian). Determining which Pauli operator and phase for every σ_p would be sufficient to learn C using Lemma 6.3.2. But from Lemma 6.3.3, we only need to know $C\sigma_g C^\dagger$ for each $\sigma_g \in G$.

Let $C\sigma_{x_i} C^\dagger = \alpha_i \sigma_{a_i}$ and $C\sigma_{z_i} C^\dagger = \beta_i \sigma_{b_i}$, where $\alpha_i, \beta_i = \pm 1$. Knowing just σ_{a_i} and σ_{b_i} is enough to specify C up to a Pauli correction factor σ_q which gives the phases α_i and β_i . Choosing σ_q that anticommutes with σ_{x_i} flips the sign of α_i and similarly for σ_{z_i} . We now present the algorithm:

1. Apply $C\sigma_{x_i} C^\dagger$ and $C\sigma_{z_i} C^\dagger$ for each i and use Theorem 6.3.1 to determine σ_{a_i} and σ_{b_i} . This uses $2n$ queries to both C and C^\dagger .
2. Let C' be such that $C'\sigma_{x_i} C'^\dagger = \sigma_{a_i}$ and $C'\sigma_{z_i} C'^\dagger = \sigma_{b_i}$ i.e. the phases are all $+1$. Then, choosing a phase for C' , we can write $C = C'\sigma$ where

$$\sigma = \prod_{i:\alpha_i=-1} \sigma_{z_i} \prod_{i:\beta_i=-1} \sigma_{x_i}. \quad (6.3.1)$$

Then implement $C'^\dagger C$ to determine σ using Theorem 6.3.1. This uses one query to C . We can now calculate the phases α_i and β_i .

To work out the time complexity, note that in step 1 the $O(n)$ time Pauli learning algorithm is called $2n$ times. Then for step 2, the Clifford C' can be implemented in $O(n^2)$ time using for example Theorem 10.6 of [NC00]. \square

We now show that this algorithm is optimal, in terms of number of queries, up to constant factors:

Lemma 6.3.5. *Any method of learning a Clifford gate requires at least n queries.*

Proof. Each application of the gate C can give at most $2n$ bits of mutual information about C . This follows from the optimality of superdense coding [BW92]. The Clifford

group (modulo global phase) is of size [CRSS98] $2^{n^2+2n+3} \prod_{j=1}^n (4^j - 1) \geq 2^{2n^2+n+3}$. To identify an element with m queries, we therefore need

$$2^{2nm} \geq 2^{2n^2+n+3} \tag{6.3.2}$$

which implies $m \geq n$. □

It is unfortunate that access to C^\dagger is also required, but we do not know a method with optimal query complexity that works without C^\dagger . There are however methods that use $O(n^2)$ queries that do not use C^\dagger . The result of [HW06] can be used to show that $O(n^2)$ queries to C are sufficient, by distinguishing the states $C \otimes I|\psi\rangle$ for different Cliffords C and where $|\psi\rangle$ is the maximally entangled state. We can use Lemma 6.4.4 to show that these states are far apart in the distance measure used in [HW06], allowing us to apply their result.

6.3.3 Learning Gottesman-Chuang Operations

Theorem 6.3.4 can easily be generalised to learning any operation from the \mathcal{C}_k hierarchy:

Theorem 6.3.6. *Given oracle access to an unknown operation $C \in \mathcal{C}_k$ and its conjugate C^\dagger , C can be determined exactly (up to phase) with $\frac{(2n)^k - 1}{2n - 1}$ queries to C and $(2n)^{k-1}$ to C^\dagger .*

Proof. The proof is by induction. The base case is for the Paulis and is proven in Theorem 6.3.1. Then, to learn $C \in \mathcal{C}_{k+1}$, we assume we have a learning algorithm for members of \mathcal{C}_k . Apply $C\sigma_g C^\dagger$ for each $\sigma_g \in G$. These operations are elements of \mathcal{C}_k so use the learning algorithm for \mathcal{C}_k to determine these up to phase. Then use the last step of Theorem 6.3.4 to determine the phases.

We now determine the number of queries to C and C^\dagger . Let $T(k)$ be the number

of queries to C and $T'(k)$ the number of queries to C^\dagger . We have the recurrences

$$\begin{aligned} T(k+1) &= 2nT(k) + 1 \\ T(1) &= 1 \end{aligned} \tag{6.3.3}$$

and

$$\begin{aligned} T'(k+1) &= 2nT'(k) \\ T'(2) &= 2n \end{aligned} \tag{6.3.4}$$

which have solutions $T(k) = \frac{(2n)^k - 1}{2n - 1}$ and $T'(k) = (2n)^{k-1}$ (with $T'(1) = 0$). \square

6.4 Learning Unitaries Close to \mathcal{C}_k Elements

Here we suppose that we are given a unitary that is known to be close to an element of \mathcal{C}_k for some given k . We present a method for finding this element. But first we must define our distance measure.

We would like our distance measure to not distinguish between unitaries that differ by just an unobservable global phase. We define a ‘distance’ D below with this property. However, firstly define the distance D^+ to be a normalised 2-norm distance:

Definition 6.4.1. For U_1 and U_2 $d \times d$ matrices,

$$D^+(U_1, U_2) := \frac{1}{\sqrt{2d}} \|U_1 - U_2\|_2.$$

where $\|A\|_2 = \sqrt{\text{tr } A^\dagger A}$.

We have chosen the normalisation so that $0 \leq D^+(U_1, U_2) \leq 1$. We now define our phase invariant ‘distance’:

Definition 6.4.2. For U_1 and U_2 $d \times d$ matrices,

$$D(U_1, U_2) := \frac{1}{\sqrt{2d^2}} \|U_1 \otimes U_1^* - U_2 \otimes U_2^*\|_2$$

This is not a true distance since $D(U_1, U_2) = 0$ does not imply $U_1 = U_2$, but that U_1 and U_2 are the same up to a phase so the difference is unobservable. From the 2-norm definition, we can show:

Lemma 6.4.3.

$$D^+(U_1, U_2) = \sqrt{1 - \operatorname{Re} \frac{\operatorname{tr} U_1 U_2^\dagger}{d}} \quad (6.4.1)$$

and

$$D(U_1, U_2) = \sqrt{1 - \left| \frac{\operatorname{tr} U_1 U_2^\dagger}{d} \right|^2}. \quad (6.4.2)$$

From this we can easily see that $0 \leq D(U_1, U_2) \leq 1$ with equality if and only if U_1 and U_2 are orthogonal. Further note that by the unitary invariance of the 2-norm, both D and D^+ are unitarily invariant and from the triangle inequality for the 2-norm they both obey the triangle inequality.

Our approximate learning method will find the unique closest element of \mathcal{C}_k to U . In order to guarantee uniqueness, the distance must be upper bounded:

Lemma 6.4.4. *If $D(U, C) < \frac{1}{2^{k-1/2}}$ for some $C \in \mathcal{C}_k$ then C is unique up to phase.*

The proof is in Section 6.7.

Theorem 6.4.5. *Given oracle access to U and U^\dagger and k such that $D(U, C) \leq \epsilon$ for some $C \in \mathcal{C}_k$ with*

$$\epsilon' := \sqrt{2(1 - (2^{k-1}\epsilon)^2)} - 1 > 0 \quad (6.4.3)$$

then C can be determined with probability at least $1 - \delta$ with

$$O\left(\frac{1}{\epsilon'^2} (2n)^{k-1} \log \frac{(2n+1)^{k-1}}{\delta}\right)$$

queries.

Proof. By Lemma 6.4.4, C is unique up to phase. We now prove the Theorem by induction.

For $k = 1$, use Proposition 21 of [MO08] to learn the closest Pauli operator. This works by repeating the Pauli learning method Theorem 6.3.1 and taking the majority vote. This uses $O\left(\frac{1}{\epsilon'^2} \log \frac{1}{\delta}\right)$ queries to succeed with probability at least $1 - \delta$.

Now for the inductive step. Assume we have a learning algorithm for level k . Then for $C \in \mathcal{C}_{k+1}$, let $C\sigma_{g_i}C^\dagger = C_{g_i}$ for $\sigma_{g_i} \in G$. By Lemma 6.8.1, we have $D(U\sigma_{g_i}U^\dagger, C_{g_i}) \leq 2\epsilon$. Use the learning algorithm for level k to determine C_{g_i} up to phase for all i . Then to find the phases we use the same method as before: implement any C' with $C'\sigma_{g_i}C'^\dagger = \pm C_{g_i}$ for any (known) choice of phase. Then $C' = C\sigma_q$ for some Pauli operator σ_q . We can determine σ_q by implementing $C'^\dagger U$ and using the $k = 1$ learning algorithm since

$$\begin{aligned} D(C'^\dagger U, \sigma_q) &= D(U, C'\sigma_q) \\ &= D(U, C) \leq \epsilon. \end{aligned} \tag{6.4.4}$$

Now we calculate the success probabilities and number of queries. There are $2n + 1$ calls to the algorithm at lower levels, which all succeed with probability at least $1 - \delta$. So at this level the success probability is at least $1 - (2n + 1)\delta$. So to succeed with probability at least $1 - \delta$ we must replace δ with $\delta/(2n + 1)$. Then the overall number of queries is

$$2n \cdot O\left(\frac{1}{\epsilon'^2} (2n)^{k-1} \log \frac{(2n+1)^k}{\delta}\right) + 1 = O\left(\frac{1}{\epsilon'^2} (2n)^k \log \frac{(2n+1)^k}{\delta}\right). \tag{6.4.5}$$

□

We remark that there is only $O(k \log n)$ overhead (for constant ϵ' and δ) over the exact learning algorithm of Theorem 6.3.6.

6.5 Clifford Testing

Here we present an efficient algorithm to determine whether an unknown unitary operation is close to a Clifford or far from every Clifford. Whereas the previous results allow us to find the Clifford operator close to the given black box unitary,

in this section we are concerned with determining how far the given unitary is from any Clifford. We do not measure this directly, but provide an algorithm of low query complexity that decides if the given unitary is close to a Clifford or far from all. This type of algorithm is known in computer science as a property testing algorithm and has many applications, including the theory of probabilistically checkable proofs [ALM⁺98]. The result in this section could be extended to work for any level of the Gottesman-Chuang hierarchy although for simplicity we only present the version for Cliffords.

The key ingredient to our method will be a way of estimating the Pauli coefficients:

Lemma 6.5.1 (Lemma 23 of [MO08]). *For any $p \in \{I, x, y, z\}^n$ and unitary U , the Pauli coefficients $|\gamma(p)| = \frac{1}{2^n} |\text{tr } U \sigma_p|$ can be estimated to within $\pm\eta$ with probability $1 - \delta$ using $O\left(\frac{1}{\eta^2} \log \frac{1}{\delta}\right)$ queries.*

This is a generalisation of Theorem 6.3.1 and the method is similar. Instead of there being only one possible outcome, now the probability of obtaining the outcome corresponding to σ_p is estimated. This probability is equal to $|\gamma(p)|^2$.

Theorem 6.5.2. *Given oracle access to U and U^\dagger with the promise that for $0 < \epsilon < 1$ either*

a) *CLOSE: there exists $C \in \mathcal{C}$ such that $D(U, C) \leq \frac{\epsilon}{\sqrt{32n}}$ or*

b) *FAR: for all $C \in \mathcal{C}$, $D(U, C) > \epsilon$ and there exists $C \in \mathcal{C}$ such that $D(U, C) \leq 1/3$*

holds then there is a $O\left(\frac{n^3}{\epsilon^2} \log \frac{n}{\delta}\right)$ algorithm that determines which with probability at least $1 - \delta$.

Proof. In both cases, we have that $D(U, C) < 1/3$ for some C , which ensures that C is unique (using Lemma 6.4.4, since $\frac{1}{3} < \frac{1}{2\sqrt{2}}$) and can be found using Theorem 6.4.5 with $O\left(n \log \frac{n}{\delta}\right)$ queries. Then the algorithm is:

1. For each $\sigma_g \in G$, measure the Pauli coefficient of $C\sigma_g C^\dagger$ in $U\sigma_g U^\dagger$ (i.e. measure $|\text{tr } U\sigma_g U^\dagger C\sigma_g C^\dagger|/2^n$) to precision $\frac{\epsilon^2}{16n^2}$ using Lemma 6.5.1.

2. If all the coefficients are found to have modulus at least $1 - \frac{3\epsilon^2}{16n^2}$ then output *CLOSE* else output *FAR*.

This works because, for the two possibilities *CLOSE* and *FAR*:

- a) Using Lemma 6.8.1, $D(U, C) \leq \frac{\epsilon}{\sqrt{32n}}$ implies that for all $\sigma_p \in \hat{\mathcal{P}}$,

$$D(U\sigma_p U^\dagger, C\sigma_p C^\dagger) \leq \frac{2\epsilon}{\sqrt{32n}}. \quad (6.5.1)$$

Since we will only apply $U\sigma_g U^\dagger$ for $\sigma_g \in G$ we restrict this to only the generators to find that for all $\sigma_g \in G$,

$$D(U\sigma_g U^\dagger, C\sigma_g C^\dagger) \leq \frac{2\epsilon}{\sqrt{32n}} \quad (6.5.2)$$

giving

$$\left| \frac{\text{tr } U\sigma_g U^\dagger C\sigma_g C^\dagger}{2^n} \right|^2 \geq 1 - \frac{\epsilon^2}{8n^2} \quad (6.5.3)$$

for every generator σ_g . We need a bound on the non-squared coefficients, which follows directly:

$$\left| \frac{\text{tr } U\sigma_g U^\dagger C\sigma_g C^\dagger}{2^n} \right| \geq 1 - \frac{\epsilon^2}{8n^2}. \quad (6.5.4)$$

Therefore when measuring the coefficients to precision $\frac{\epsilon^2}{16n^2}$, all results will give at least $1 - \frac{3\epsilon^2}{16n^2}$.

- b) Using the contrapositive of Lemma 6.8.2, $D(U, C) > \epsilon$ implies that there exists $\sigma_p \in \hat{\mathcal{P}}$ such that

$$D^+(U\sigma_p U^\dagger, C\sigma_p C^\dagger) > \epsilon. \quad (6.5.5)$$

Using the contrapositive of Lemma 6.8.3 this in turn implies there exists $\sigma_g \in G$ such that

$$D^+(U\sigma_g U^\dagger, C\sigma_g C^\dagger) > \frac{\epsilon}{2n}, \quad (6.5.6)$$

which means that for at least one $\sigma_g \in G$, $U\sigma_g U^\dagger$ will have a small overlap with

$C\sigma_g C^\dagger$ i.e. there exists $\sigma_g \in G$ such that

$$\left| \frac{\text{tr } U\sigma_g U^\dagger C\sigma_g C^\dagger}{2^n} \right| < 1 - \frac{\epsilon^2}{4n^2}. \quad (6.5.7)$$

The C returned by the application of Theorem 6.4.5 is such that $\text{tr } U\sigma_g U^\dagger C\sigma_g C^\dagger$ is positive, which justifies inserting the absolute value signs above when using D^+ rather than D . This implies that at least one coefficient will be found to be less than $1 - \frac{3\epsilon^2}{16n^2}$ when measuring to precision $\frac{\epsilon^2}{16n^2}$. \square

6.6 Conclusions and Further Work

We have shown how to exactly identify an unknown Clifford operator in $O(n)$ queries, which we show is optimal. This is then extended to cover elements of the \mathcal{C}_k hierarchy and for unitaries that are only known to be close to \mathcal{C}_k operations. The key to the Clifford learning algorithm is to apply $C\sigma_p C^\dagger$ and then find the resulting Pauli operator.

A way of extending this idea could be to learn unitaries from larger sets. Suppose \mathcal{V} is a set of unitaries with the property that for every $V \in \mathcal{V}$, $V\sigma_p V^\dagger$ is a linear combination of a constant number of Paulis. Then V can be learnt in the same way as above, using the quantum Goldreich-Levin algorithm of [MO08], which can efficiently find which Paulis have large overlap with an input unitary. However, we have not been able to find interesting sets \mathcal{V} other than the Clifford group with this property.

We also presented a Clifford testing algorithm, which determines whether a given black-box unitary is close to a Clifford or far from every Clifford. This can be seen as a quantum generalisation of quadratic testing, just as Pauli testing can be seen as a quantum generalisation of linearity testing. Property testing of this form is used to prove the PCP theorem [ALM⁺98] so these quantum testing results could potentially be useful in proving a quantum PCP theorem. It would also be interesting to strengthen the testing method in Theorem 6.5.2 to remove the $O(1/n)$ difference

between the close and far conditions.

Finally, it would be interesting to see if it is possible to remove the requirement to have access to U^\dagger . However, using both U and U^\dagger is the key to our method so we do not know if a method without U^\dagger is possible with low query complexity.

6.7 Proof of Lemma 6.4.4

Proof of Lemma 6.4.4. The proof is by induction. The base case is for $k = 1$ when we have the Pauli group. Without loss of generality, assume C is a Pauli operator with no phase. Let $C = \sigma_p$.

Expand U in the Pauli basis:

$$U = \sum_q \gamma(q) \sigma_q. \quad (6.7.1)$$

Since U is unitary, we have $\sum_q |\gamma(q)|^2 = 1$. By Lemma 6.4.3,

$$D(U, \sigma_p)^2 = 1 - \left| \frac{\text{tr } \sigma_p U}{2^n} \right|^2 \quad (6.7.2)$$

which implies

$$|\gamma(p)|^2 \geq 1 - \epsilon^2. \quad (6.7.3)$$

Now, suppose for contradiction that there exists $\sigma_{p_1} \neq \sigma_{p_2}$ with $D(U, \sigma_{p_1}) \leq \epsilon$ and $D(U, \sigma_{p_2}) \leq \epsilon$. Then by the above, $|\gamma(p_1)|^2, |\gamma(p_2)|^2 \geq 1 - \epsilon^2$. But there is also the constraint $|\gamma(p_1)|^2 + |\gamma(p_2)|^2 \leq 1$ which combined give

$$\epsilon \geq \frac{1}{\sqrt{2}} \quad (6.7.4)$$

which is false by assumption. This implies $\sigma_{p_1} = \sigma_{p_2}$, which proves the base case.

To prove the inductive step, again assume for contradiction that there exist $C_1, C_2 \in \mathcal{C}_{k+1}$ with $C_1 \neq C_2$ and $D(U, C_1) \leq \epsilon$ and $D(U, C_2) \leq \epsilon$. Then there

exists $\sigma_g \in G$ with

$$C_1\sigma_g C_1^\dagger =: C_{1g} \neq C_{2g} := C_2\sigma_g C_2^\dagger. \quad (6.7.5)$$

Here, $C_{1g}, C_{2g} \in \mathcal{C}_k$.

Using Lemma 6.8.1, $D(U\sigma_g U^\dagger, C_{1g}) \leq 2\epsilon$ and $D(U\sigma_g U^\dagger, C_{2g}) \leq 2\epsilon$.

Now there are two cases. Firstly, suppose we can choose σ_g such that $C_1\sigma_g C_1^\dagger \neq \pm C_2\sigma_g C_2^\dagger$. Then C_{1g} and C_{2g} are not equivalent up to phase so, using the inductive hypothesis, we must have

$$2\epsilon \geq \frac{1}{2^{k-1/2}} \quad (6.7.6)$$

or

$$\epsilon \geq \frac{1}{2^{(k+1)-1/2}} \quad (6.7.7)$$

which is again false by assumption.

For the other case, $C_1\sigma_g C_1^\dagger = \pm C_2\sigma_g C_2^\dagger$ for all $\sigma_g \in G$. This implies that $C_2 = C_1\sigma_q$ for some Pauli $\sigma_q \neq I$. Then we have

$$\begin{aligned} D(U, C_1) &\leq \epsilon \\ D(U, C_1\sigma_q) &\leq \epsilon \end{aligned} \quad (6.7.8)$$

which by unitary invariance gives

$$\begin{aligned} D(C_1^\dagger U, I) &\leq \epsilon \\ D(C_1^\dagger U, \sigma_q) &\leq \epsilon. \end{aligned} \quad (6.7.9)$$

But we proved that this is impossible in this range of ϵ in the $k = 1$ proof above. \square

6.8 Miscellaneous Lemmas

Here we prove some miscellaneous lemmas used earlier in the chapter.

The first lemma says that for two close operators U_1 and U_2 , $U_1\sigma_p U_1^\dagger$ is close to $U_2\sigma_p U_2^\dagger$ for all Paulis σ_p :

Lemma 6.8.1. *If $D(U_1, U_2) \leq \delta$ then for all $\sigma_p \in \hat{\mathcal{P}}$,*

$$D(U_1\sigma_p U_1^\dagger, U_2\sigma_p U_2^\dagger) \leq 2\delta.$$

Proof. Let $U_1 = VU_2$ and $U_{2p} = U_2\sigma_p U_2^\dagger$. Then we simply apply the triangle inequality for D and unitary invariance:

$$\begin{aligned} D(U_1\sigma_p U_1^\dagger, U_2\sigma_p U_2^\dagger) &= D(VU_2\sigma_p U_2^\dagger V^\dagger, U_{2p}) \\ &= D(VU_{2p}, U_{2p}V) \\ &\leq D(VU_{2p}, U_{2p}) + D(U_{2p}, U_{2p}V) \\ &= D(V, I) + D(I, V) \\ &= 2D(U_1, U_2). \end{aligned} \quad \square$$

The next lemma is a converse to this:

Lemma 6.8.2. *If for all $\sigma_p \in \hat{\mathcal{P}}$*

$$D^+(U_1\sigma_p U_1^\dagger, U_2\sigma_p U_2^\dagger) \leq \delta \tag{6.8.1}$$

then

$$D(U_1, U_2) \leq \delta. \tag{6.8.2}$$

Proof. If $D^+(U_1\sigma_p U_1^\dagger, U_2\sigma_p U_2^\dagger) \leq \delta$ then $\frac{1}{2^n} \text{Re tr } U_1\sigma_p U_1^\dagger U_2\sigma_p U_2^\dagger \geq 1 - \delta^2$. Since this is true for all σ_p , we can take the average of this over the whole of $\hat{\mathcal{P}}$ and use the fact that for any $d \times d$ matrix A $\frac{1}{4^n} \sum_{\sigma_p \in \hat{\mathcal{P}}} \sigma_p A \sigma_p = \frac{I}{2^n} \text{tr } A$ (the Paulis are a *1-design*) to find

$$\frac{1}{2^n} \text{Re tr } U_1 \left(\frac{I}{2^n} \text{tr } U_1^\dagger U_2 \right) U_2^\dagger \geq 1 - \delta^2 \tag{6.8.3}$$

which simplified gives

$$\left| \frac{\text{tr } U_1 U_2^\dagger}{2^n} \right|^2 \geq 1 - \delta^2 \tag{6.8.4}$$

giving the desired result. \square

Now we show how to go from distances for just the generators G to distances for the whole of $\hat{\mathcal{P}}$:

Lemma 6.8.3. *If for all $\sigma_g \in G$*

$$D^+(U_1\sigma_gU_1^\dagger, U_2\sigma_gU_2^\dagger) \leq \delta \tag{6.8.5}$$

then for all $\sigma_p \in \hat{\mathcal{P}}$

$$D^+(U_1\sigma_pU_1^\dagger, U_2\sigma_pU_2^\dagger) \leq 2n\delta \tag{6.8.6}$$

Proof. The proof is by induction on the number of generators required to make σ_p , using the triangle inequality for D^+ . □

Bibliography

- [Aar07] S. Aaronson. The learnability of quantum states. *Proc. R. Soc. A*, 463:3089–3114, 2007. arXiv:quant-ph/0608142.
- [Aar09] S. Aaronson. Quantum Copy-Protection and Quantum Money. *IEEE Conference on Computational Complexity 2009*, 2009.
- [ABI86] N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(4):567–583, 1986.
- [ABW09] A. Ambainis, J. Bouda, and A. Winter. Nonmalleable encryption of quantum information. *J. Math. Phys.*, 50(4):042106, 2009. arXiv:0808.0353.
- [ADHW06] A. Abeyesinghe, I. Devetak, P. Hayden, and A. Winter. The mother of all protocols: Restructuring quantum information’s family tree, 2006. arXiv:quant-ph/0606225.
- [AE07] A. Ambainis and J. Emerson. Quantum t -designs: t -wise independence in the quantum world. *IEEE Conference on Computational Complexity 2007*, pages 129–140, 2007. arXiv:quant-ph/0701126v2.
- [AG09] S. Aaronson and D. Gottesman, 2009. Unpublished.
- [AK62] V. I. Arnold and A. L. Krylov. Uniform distribution of points on a sphere and some ergodic properties of solutions of linear ordinary differential equations in a complex domain. *Soviet Math. Dokl.*, 4(1), 1962.

- [ALM⁺98] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [AMTd00] A. Ambainis, M. Mosca, A. Tapp, and R. de Wolf. Private Quantum Channels. *41st Annual IEEE Symposium on Foundations of Computer Science*, pages 547–553, 2000.
- [AN08] N. Alon and A. Nussboim. k -Wise Independent Random Graphs. *49th Annual IEEE Symposium on Foundations of Computer Science*, 0:813–822, 2008. arXiv:0804.1268.
- [AS04] A. Ambainis and A. Smith. Small Pseudo-Random Families of Matrices: Derandomizing Approximate Quantum Encryption. *Proceedings of RANDOM 2004, LNCS*, 3122:249–260, 2004. arXiv:quant-ph/0404075.
- [Aub09] G. Aubrun. On Almost Randomizing Channels with a Short Kraus Decomposition. *Comm. Math. Phys.*, 288:1103–1116, 2009. arXiv:0805.2900.
- [Bar02] H. Barnum. Information-disturbance tradeoff in quantum measurement on the uniform ensemble and on the mutually unbiased bases, 2002. arXiv:quant-ph/0205155.
- [BBD⁺97] A. Barenco, A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello. Stabilization of Quantum Computations by Symmetrization. *SIAM J. Comput.*, 26(5):1541–1557, 1997. arXiv:quant-ph/9604028.
- [BFP⁺72] M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, and R. E. Tarjan. Linear time bounds for median computations. *4th Annual ACM Symposium on Theory of Computing*, pages 119–124, 1972.
- [BG06] J. Bourgain and A. Gamburd. New results on expanders. *C. R. Acad. Sci. Paris, Ser. I*, 342:717–721, 2006.

- [BH08] A. Brodsky and S. Hoory. Simple Permutations Mix Even Better. *Random Struct. Algorithms*, 32(3):274–289, 2008. arXiv:math/0411098.
- [BHL⁺05] C.H. Bennett, P. Hayden, D.W. Leung, P.W. Shor, and A. Winter. Remote preparation of quantum states. *IEEE Trans. Inform. Theory*, 51(1):56–74, 2005.
- [BL01] H. Barnum and N. Linden. Monotones and invariants for multi-particle quantum states. *J. Phys. A*, 34(35):6787–6805, 2001. arXiv:quant-ph/0103155.
- [BMW09] M. J. Bremner, C. Mora, and A. Winter. Are Random Pure States Useful for Quantum Computation? *Phys. Rev. Lett.*, 102(19):190502, 2009. arXiv:0812.3001.
- [BR94] M. Bellare and J. Rompel. Randomness-efficient oblivious sampling. *35th Annual IEEE Symposium on Foundations of Computer Science*, pages 276–287, 1994.
- [BT07] A. Ben-Aroya and A. Ta-Shma. Quantum expanders and the quantum entropy difference problem, 2007. arXiv:quant-ph/0702129.
- [BW92] C. H. Bennett and S. J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Phys. Rev. Lett.*, 69(20):2881–2884, 1992.
- [CN97] I. L. Chuang and M. A. Nielsen. Prescription for experimental determination of the dynamics of a quantum black box. *Journal of Modern Optics*, 44:2455–2467, 1997. arXiv:quant-ph/9610001.
- [CRSS97] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane. Quantum Error Correction and Orthogonal Geometry. *Phys. Rev. Lett.*, 78(3):405–408, 1997.

- [CRSS98] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane. Quantum Error Correction Via Codes Over $GF(4)$. *IEEE Trans. Inform. Theory*, 44:1369–1387, 1998.
- [Dan05] C. Dankert. Efficient Simulation of Random Quantum States and Operators. *MMath Thesis, University of Waterloo*, 2005. arXiv:quant-ph/0512217.
- [DCEL06] C. Dankert, R. Cleve, J. Emerson, and E. Livine. Exact and Approximate Unitary 2-Designs: Constructions and Applications, 2006. arXiv:quant-ph/0606161.
- [DHL⁺04] D. P. DiVincenzo, M. Horodecki, D. W. Leung, J. A. Smolin, and B. M. Terhal. Locking Classical Correlations in Quantum States. *Phys. Rev. Lett.*, 92(6):067902, 2004. arXiv:quant-ph/0303088.
- [DLP01] G. M. D’Ariano and P. Lo Presti. Quantum Tomography for Measuring Experimentally the Matrix Elements of an Arbitrary Quantum Operation. *Phys. Rev. Lett.*, 86(19):4195–4198, 2001. arXiv:quant-ph/0012071.
- [DLT02] D. DiVincenzo, D. Leung, and B. Terhal. Quantum Data Hiding. *IEEE Trans. Inform. Theory*, 48(3):580–598, 2002. arXiv:quant-ph/0103098.
- [DN06] P. A. Dickinson and A. Nayak. Approximate Randomization of Quantum States With Fewer Bits of Key. *Quantum Computing: Back Action 2006*, 864:18–36, 2006.
- [DOP07] O. C. O. Dahlsten, R. Oliveira, and M. B. Plenio. The emergence of typical entanglement in two-party random processes. *Journal of Physics A Mathematical General*, 40:8081–8108, 2007. arXiv:quant-ph/0701125.
- [DP06] O. Dahlsten and M. Plenio. Entanglement probability distribution of bipartite randomised stabilizer states. *Q. Info. Comp.*, 6(6):527–538, 2006. arXiv:quant-ph/0511119.

- [DS93] P. Diaconis and L. Saloff-Coste. Comparison Theorems for Reversible Markov Chains. *Ann. Appl. Probab.*, 3(3):696–730, 1993.
- [DS96] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *Ann. Appl. Probab.*, 6(3):695–750, 1996.
- [ELL05] J. Emerson, E. Livine, and S. Lloyd. Convergence conditions for random quantum circuits. *Phys. Rev. A*, 72(060302), 2005. arXiv:quant-ph/0503210.
- [FK94] S. K. Foong and S. Kanno. Proof of Page’s conjecture on the average entropy of a subsystem. *Phys. Rev. Lett.*, 72:1148–1151, 1994.
- [GAE07] D. Gross, K. Audenaert, and J. Eisert. Evenly distributed unitaries: On the structure of unitary designs. *J. Math. Phys.*, 48(052104), 2007. arXiv:quant-ph/0611002.
- [GC99] D. Gottesman and I. L. Chuang. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, 402:390–393, 1999. arXiv:quant-ph/9908010.
- [GE08] D. Gross and J. Eisert. Quantum Margulis Expanders. *Q. Info. Comp.*, 8(8/9):722–733, 2008. arXiv:0710.0651.
- [GFE09] D. Gross, S. T. Flammia, and J. Eisert. Most Quantum States Are Too Entangled To Be Useful As Computational Resources. *Phys. Rev. Lett.*, 102(19):190501, 2009. arXiv:0810.4331.
- [Gir07] O. Giraud. Distribution of bipartite entanglement for random pure states. *J. Phys. A*, 40:2793–2801, 2007. arXiv:quant-ph/0611285.
- [Got98] D. Gottesman. Theory of fault-tolerant quantum computation. *Phys. Rev. A*, 57(1):127–137, 1998. arXiv:quant-ph/9702029.
- [Gro97] L. Grover. Quantum Mechanics Helps in Searching for a Needle in a Haystack. *Phys. Rev. Lett.*, 79(2):325–328, 1997.

- [GW86] G. Grimmett and D. Welsh. *Probability: An Introduction*. Oxford University Press, Oxford, UK, 1986.
- [GW98] R. Goodman and N. Wallach. *Representations and Invariants of the Classical Groups*. Cambridge University Press, Cambridge, UK, 1998.
- [Har08] A. W. Harrow. Quantum expanders from any classical Cayley graph expander. *Q. Info. Comp.*, 8(8/9):715–721, 2008. arXiv:0709.1142.
- [HH08] S. Hallgren and A. W. Harrow. Superpolynomial Speedups Based on Almost Any Quantum Circuit. In *ICALP '08: Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*, pages 782–795, Berlin, Heidelberg, 2008. Springer-Verlag. arXiv:0805.0007.
- [HH09] M. B. Hastings and A. W. Harrow. Classical and Quantum Tensor Product Expanders. *Q. Info. Comp.*, 9:336, 2009. arXiv:0804.0011.
- [HHH06] A. Hayashi, T. Hashimoto, and M. Horibe. Reexamination of optimal quantum state estimation of pure states. *Phys. Rev. A*, 72(032325), 2006. arXiv:quant-ph/0410207.
- [HHL04] A. Harrow, P. Hayden, and D. Leung. Superdense Coding of Quantum States. *Phys. Rev. Lett.*, 92(18):187901, 2004. arXiv:quant-ph/0307221.
- [HHR⁺05] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-Al-Kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt. Scalable multiparticle entanglement of trapped ions. *Nature*, 438(7068):643–646, 2005. arXiv:quant-ph/0603217.
- [HHYW07] P. Hayden, M. Horodecki, J. Yard, and A. Winter. A decoupling approach to the quantum capacity. *Open Syst. Inf. Dyn.*, 15:7–19, 2007. arXiv:quant-ph/0702005.

- [HL09a] A. W. Harrow and R. A. Low. Efficient Quantum Tensor Product Expanders and k -Designs. *Proceedings of RANDOM 2009, LNCS*, 5687:548–561, 2009. arXiv:0811.2597.
- [HL09b] A. W. Harrow and R. A. Low. Random Quantum Circuits are Approximate 2-designs. *Comm. Math. Phys.*, 291(1):257–302, 2009. arXiv:0802.1919.
- [HLSW04] P. Hayden, D. Leung, P. W. Shor, and A. Winter. Randomizing Quantum States: Constructions and Applications. *Comm. Math. Phys.*, 250:371–391, 2004. arXiv:quant-ph/0307104.
- [HLW06] P. Hayden, D. W. Leung, and A. Winter. Aspects of Generic Entanglement. *Comm. Math. Phys.*, 265:95–117, 2006. arXiv:quant-ph/0407049.
- [HP07] P. Hayden and J. Preskill. Black holes as mirrors: quantum information in random subsystems. *Journal of High Energy Physics*, 09(120), 2007. arXiv:0708.4025.
- [HW06] A. W. Harrow and A. Winter. How many copies are needed for state discrimination?, 2006. arXiv:quant-ph/0606131.
- [IR06] S. Iblisdir and J. Roland. Optimal finite measurements and Gauss quadratures. *Phys. Lett. A*, 358:368–372, 2006. arXiv:quant-ph/0410237.
- [Kah96] N. Kahale. A Semidefinite Bound for Mixing Rates of Markov Chains. *Proceedings of the 5th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pages 190–203, 1996.
- [Kas05] M. Kassabov. Symmetric Groups and Expanders, 2005. arXiv:math/0503204.
- [KN96] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, Cambridge, UK, 1996.

- [KNR09] E. Kaplan, M. Naor, and O. Reingold. Derandomized Constructions of k -Wise (Almost) Independent Permutations. *Algorithmica*, 55(1):113–133, 2009.
- [KSV02] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, Boston, MA, USA, 2002.
- [Lan92] R. Landauer. Information is Physical. *1992 Workshop on Physics and Computation*, pages 1–4, 1992.
- [Led01] M. Ledoux. *The Concentration of Measure Phenomenon*. American Mathematical Society, Providence, RI, USA, 2001.
- [Leu03] D. W. Leung. Choi’s proof as a recipe for quantum process tomography. *J. Math. Phys.*, 44(2):528–533, 2003. arXiv:quant-ph/0201119.
- [Low09a] R. A. Low. Large Deviation Bounds for k -designs. *Proc. R. Soc. A*, 465(2111):3289–3308, 2009. arXiv:0903.5236.
- [Low09b] R. A. Low. Learning and Testing Algorithms for the Clifford Group. *Phys. Rev. A*, 80(5):052314, 2009. arXiv:0907.2833.
- [MO08] A. Montanaro and T. J. Osborne. Quantum boolean functions, 2008. arXiv:0810.2435.
- [MP95] S. Massar and S. Popescu. Optimal Extraction of Information from Finite Quantum Ensembles. *Phys. Rev. Lett.*, 74(8):1259–1263, 1995.
- [MR95] R. Motwani and P. Raghavan. *Randomized algorithms*. Cambridge University Press, Cambridge, UK, 1995.
- [MR00] R.A. Martin and D. Randall. Sampling adsorbing staircase walks using a new Markov chain decomposition method. *41st Annual IEEE Symposium on Foundations of Computer Science*, pages 492–502, 2000.

- [MT06] R. Montenegro and P. Tetali. Mathematical aspects of mixing times in Markov chains. *Found. Trends Theor. Comput. Sci.*, 1(3):237–354, 2006.
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, UK, 2000.
- [Nie02] M. A. Nielsen. A simple formula for the average gate fidelity of a quantum dynamical operation. *Phys. Lett. A*, 303:249–252, 2002. arXiv:quant-ph/0205035.
- [NN90] J. Naor and M. Naor. Small-bias probability spaces: efficient constructions and applications. *22nd Annual ACM Symposium on Theory of Computing*, pages 213–223, 1990.
- [ODP07] R. Oliveira, O. C. O. Dahlsten, and M. B. Plenio. Efficient Generation of Generic Entanglement. *Phys. Rev. Lett.*, 98(130502), 2007. arXiv:quant-ph/0605126.
- [Pag93] D. N. Page. Average entropy of a subsystem. *Phys. Rev. Lett.*, 71:1291, 1993.
- [PCZ97] J. F. Poyatos, J. I. Cirac, and P. Zoller. Complete Characterization of a Quantum Process: The Two-Bit Quantum Gate. *Phys. Rev. Lett.*, 78(2):390–393, 1997. arXiv:quant-ph/9611013.
- [PSW06] S. Popescu, A. J. Short, and A. Winter. Entanglement and the Foundations of Statistical Mechanics. *Nature Physics*, 2:754–758, 2006. arXiv:quant-ph/0511225.
- [Rot64] G.-C. Rota. On the foundations of combinatorial theory I. Theory of Möbius Functions. *Probability Theory and Related Fields*, 2(4):340–368, 1964.
- [San95] J. Sanchez-Ruiz. Simple proof of Page’s conjecture on the average entropy of a subsystem. *Phys. Rev. E*, 52:5653, 1995.

- [Sen05] P. Sen. Random measurement bases, quantum state distinction and applications to the hidden subgroup problem. *IEEE Conference on Computational Complexity 2006*, pages 274–287, 2005. arXiv:quant-ph/0512085.
- [Shi95] A. Shimony. Degree of Entanglement. *Ann. N.Y. Acad. Sci.*, 755:675, 1995.
- [Sho94] P. W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. *35th Annual IEEE Symposium on Foundations of Computer Science*, pages 124–134, 1994.
- [Sho96] P. W. Shor. Fault-tolerant quantum computation. *37th Annual IEEE Symposium on Foundations of Computer Science*, pages 56–65, 1996. arXiv:quant-ph/9605011.
- [SL06] G. Smith and D. Leung. Typical entanglement of stabilizer states. *Phys. Rev. A*, 74(6):062314, 2006. arXiv:quant-ph/0510232.
- [SS08] Y. Sekino and L. Susskind. Fast scramblers. *Journal of High Energy Physics*, 10:65, 2008. arXiv:0808.2096.
- [Sta86] R. Stanley. *Enumerative Combinatorics*. Cambridge University Press, Cambridge, UK, 1986.
- [TDL01] B. M. Terhal, D. P. Divincenzo, and D. W. Leung. Hiding Bits in Bell States. *Phys. Rev. Lett.*, 86:5807–5810, 2001. arXiv:quant-ph/0011042.
- [van02] W. K. van Dam. On Quantum Computation Theory. *PhD Thesis, University of Amsterdam*, 2002.
- [Zni07] M. Znidaric. Optimal two-qubit gate for generation of random bipartite entanglement. *Phys. Rev. A*, 76(012318), 2007. arXiv:quant-ph/0702240.