

Efficient Discrete Approximations of Quantum Gates

Aram W. Harrow
aram@mit.edu
MIT Media Laboratory
20 Ames Street
Cambridge, MA 02139

Benjamin Recht
brecht@media.mit.edu
MIT Media Laboratory
20 Ames Street
Cambridge, MA 02139

Isaac L. Chuang
ichuang@media.mit.edu
MIT Media Laboratory
20 Ames Street
Cambridge, MA 02139

July 11, 2002

Abstract

Quantum compiling addresses the problem of approximating an arbitrary quantum gate with a string of gates drawn from a particular finite set. It has been shown that this is possible for almost all choices of base sets and furthermore that the number of gates required for precision ϵ is only polynomial in $\log 1/\epsilon$. Here we prove that using certain sets of base gates quantum compiling requires a string length that is linear in $\log 1/\epsilon$, a result which matches the lower bound from counting volume up to constant factor.

1 Introduction

Quantum computation generalizes computer science to utilize novel quantum physical resources as elementary building blocks for information processing [1, 2, 3, 4]. Quantum algorithms, like their classical analogues, can be written in a number of nearly equivalent ways. While a classical program is typically composed of a series of simple boolean functions, such as NAND and FANOUT, a quantum algorithm is typically written as a product of unitary gates, such as the Hadamard transform H , the controlled-NOT (CNOT), and the $\pi/8$ -gate T [5]. For classical computers, a common problem is that of compiling a program, in which one typically wishes to express the program in as few elementary operations as possible. By analogy, we can raise the principal questions of *quantum compiling*: which sets of gates can be composed to form what sorts of quantum algorithm, how many of them are necessary, and what efficient algorithms can be devised to express quantum programs in terms of a particular set of base gates?

Mathematically, a gate on n quantum bits (qubits) is represented by a unitary transformation on a 2^n -dimensional vector space. We will denote the set of all determinant-one unitary transformations of a d -dimensional vector space by $SU(d)$. This space is a manifold and is hence parameterized by a *continuum* of real parameters; for example, the 2×2 unitary transforms

$$\begin{pmatrix} e^{i\alpha} \cos \theta & e^{i\beta} \sin \theta \\ -e^{-i\beta} \sin \theta & e^{-i\alpha} \cos \theta \end{pmatrix} \quad (1)$$

parameterized by α, β, θ represent the group $SU(2)$ of valid single qubit gates.

In contrast, digital quantum algorithms compute with only a *finite set* of base gates (such as those mentioned previously: H , T , and CNOT). This is a reasonable restriction in real circuit implementations, since the presence of noise reduces the number of reliably distinguishable gates to a finite subset of the continuous set. Finite gate sets are also intrinsic to fault-tolerant quantum computation, the art of constructing arbitrarily reliable circuits from unreliable parts.[6, 7, 8, 9] Thus, in general we do not desire perfect computational universality, but only the ability to approximate any quantum algorithm, preferably without using too many more gates than originally required.

A set of base gates $\mathcal{A} \subset SU(d)$ is *computationally universal* if given any gate U , we can find a string consisting of gates from \mathcal{A} and their inverses, such that the product of the gates in the string approximates U to arbitrary precision. Equivalently, \mathcal{A} must generate a dense subgroup of $SU(d)$.

Which sets of base gates are computationally universal? It turns out that probabilistically speaking, almost all of them are [10, 11]. If base gates are chosen at random, then all but a set of measure zero are computationally universal. The idea is that if the eigenvalues of the base gates have phases that are irrationally related to π (which occurs with probability one), then taking powers of them allows each base gate to approximate a one-parameter subgroup to arbitrary precision, just as integer multiples of a random vector modulo a lattice will almost always fill space. Furthermore, the base gates will almost always lie on different one-parameter subgroups, which will generate all of $SU(d)$ with probability one.

Given that compiling is generically possible, it is vital to determine *how short a string of base gates is typically required to approximate a given gate to a specified precision*; this is the question we consider in this paper. The construction described by Lloyd [10] requires using a number of base gates exponential in $\log 1/\epsilon$ to achieve a precision of ϵ . This is an unreasonable cost for many applications. However, Solovay [12] and Kitaev [13] have independently described an efficient (meaning its running time is polynomial in $\log 1/\epsilon$) algorithm for quantum compiling that produces strings of length only $O(\log^c(1/\epsilon))$, where c is a constant between 3 and 4. [14] The algorithm works by constructing successively finer ϵ -nets; finite sets of gates that can approximate any element of $SU(d)$ to an accuracy of ϵ .

On the other hand, as we will later discuss, since a ball of radius ϵ in $SU(d)$ has volume proportional to ϵ^{d^2-1} , it takes $O((1/\epsilon)^{d^2-1})$ different strings of gates to approximate every element of $SU(d)$ to a precision of ϵ . Therefore, no algorithm will ever be able to reduce c below 1. Furthermore, it is unlikely that the successive approximation method used by the Solovay-Kitaev theorem will be able to do better than $c = 2$. [14] This still leaves open the question of whether some other technique could establish an upper bound asymptotically smaller than the one achieved by the Solovay-Kitaev theorem.

Here, our main result is that for at least some universal sets of base gates only $O(\log 1/\epsilon)$ gates are sufficient to approximate any gate to a precision ϵ (i.e. $c = 1$). This is within a constant factor of the lower bound obtained from counting arguments. We say that these base gates are not only computationally universal, but also *efficiently universal*, since using them for quantum compiling requires a string length that is optimal up to a constant multiplicative factor.

We present this result as follows. The set of strings from a fixed computationally universal set of base gates cover $SU(d)$ increasingly densely and uniformly, as the string length grows. [15] First, in Section 2, we quantify how quickly this occurs by introducing a framework for comparing the distribution of strings with the uniform distribution. We use this formalism in Section 3 to identify a condition on base sets that implies their efficient universality. In Section 4 we then combine this condition with results from the literature to show that efficiently universal gate sets for Hilbert spaces of any finite dimension. Section 5 discusses lower bounds for compilation and demonstrates the optimality of the result; we conclude with open questions and further directions.

2 Preliminaries

We begin by developing a metric of how well strings drawn from a finite set of gates approximate arbitrary elements of $SU(d)$.

Let dg be the Haar measure on $SU(d)$ normalized so that $\int dg = 1$. Consider the Hilbert space $L^2(SU(d))$ with norm defined by the usual inner product $\langle \psi, \varphi \rangle \equiv \int \psi(g)^* \varphi(g) dg$. The norm of a linear transformation on $L^2(SU(d))$ is given by

$$|M| \equiv \sup \{ \|Mf\| \mid f \in L^2(SU(d)), \|f\| = 1 \} . \quad (2)$$

When M is bounded and hermitian, the norm is simply the supremum of its spectrum and as a result, $|M^n| = |M|^n$.

Define a representation $U \mapsto \tilde{U}$ of $SU(d)$ on $L^2(SU(d))$ by

$$\tilde{U}f(x) = f(U^{-1}x) . \quad (3)$$

Using the right invariance of the Haar measure, we see that \tilde{U} is unitary. For any finite set $\mathcal{A} \subset SU(d)$, define the mixing operator $T(\mathcal{A})$ by

$$T(\mathcal{A}) = \frac{1}{2|\mathcal{A}|} \sum_{A \in \mathcal{A}} \tilde{A} + \tilde{A}^{-1} . \quad (4)$$

All such T are hermitian and have norm one. We will often simply write T instead of $T(\mathcal{A})$. These represent averaging the action of the elements of \mathcal{A} and their inverses on a function; when the function is a probability distribution on $SU(d)$ we can think of T as multiplying by a random element of \mathcal{A} .

Applying T^n represents averaging over the action of words of length n . Denote the set of words of length n made up of elements of \mathcal{A} and their inverses by $W_n(\mathcal{A})$, or when the set \mathcal{A} is understood, simply W_n . This set comprises $(2|\mathcal{A}|)^n$ words, though as matrices there are generally some duplicates since substrings such as $AA^{-1} = \mathbf{1}$ for all $A \in \mathcal{A}$. For any positive integer n , expanding T^n gives

$$T^n = \sum_{w \in W_n} \frac{\tilde{w}}{(2|\mathcal{A}|)^n}. \quad (5)$$

We want to compare T^n to the integral operator P .

$$Pf(h) = \int f(gh)dg = \int f(g)dg. \quad (6)$$

Note that P is the projection operator onto the set of constant functions on $SU(d)$, and hence $P = P^\dagger$ and $P^2 = P$. It is not hard to show that $TP = P = PT$ and consequently

$$(T - P)^n = T^n - P. \quad (7)$$

The metric for comparing $T(\mathcal{A})$ to P is given by

$$\Lambda(\mathcal{A}) \equiv |T(\mathcal{A}) - P|. \quad (8)$$

From Eq. 7 and the hermiticity of T and P , it follows that

$$\Lambda(\mathcal{A})^n = |T^n(\mathcal{A}) - P|. \quad (9)$$

If one thinks of T^n as a Riemann sum then Λ serves as to quantify how quickly T^n converges to the integral. It has been shown [15] that if \mathcal{A} is a computationally universal set that all the eigenvalues of $T - P$ have absolute value strictly less than one. However, this only implies that $\Lambda(\mathcal{A}) \leq 1$, since $T - P$ has an infinite number of eigenvalues.

The proof of main result of our paper—that efficiently universal sets of gates exist—is divided in the next two sections. In Section 3 we show that $\Lambda(\mathcal{A}) < 1$ implies that \mathcal{A} is efficiently universal and in Section 4 we demonstrate that for any d an efficiently universal set of gates can be found in $SU(d)$.

3 A condition for efficient universality

Theorem 1 *For any $\mathcal{A} \subset SU(d)$ such that $\Lambda(\mathcal{A}) < 1$, \mathcal{A} is efficiently universal. Specifically, there exists a constant C such that for all $U \in SU(d)$, $\epsilon > 0$, and $n > C \log 1/\epsilon$, there is a $w \in W_n(\mathcal{A})$ such that $|w - U| < \epsilon$.*

Before proving the theorem, we will need to note a fact about the geometry of $SU(d)$. For any d and r_0 , if $V(r)$ is the Haar measure of a ball of radius r in $SU(d)$, then there exist constants k_1 and k_2 such that

$$k_1 r^{d^2-1} < V(r) < k_2 r^{d^2-1}. \quad (10)$$

for all $r \in (0, r_0)$. This is true because $SU(d)$ is a $d^2 - 1$ -dimensional manifold and because $V(r)$ does not depend on the center of the ball under the Haar measure.

Now we can proceed with the proof of Theorem 1:

Proof Define $\chi \in L^2(SU(d))$ by

$$\chi(g) = \begin{cases} 1 & \text{for } |g - I| < \epsilon/2 \\ 0 & \text{otherwise} \end{cases}. \quad (11)$$

Let $V = \|P\chi\| = \|\chi\|^2$ be the measure of the ball around the identity of radius $\epsilon/2$. We will not perform this integration, but recall from Eq. 10 that $V > k_1(\epsilon/2)^{d^2-1}$.

Let $T = T(\mathcal{A})$ and $\Lambda = \Lambda(\mathcal{A})$.

First we use the Cauchy-Schwartz inequality to give

$$\begin{aligned} \left| \langle \chi, (T^n - P)\tilde{U}\chi \rangle \right| &\leq \|\chi\| \|(T^n - P)\tilde{U}\chi\| \\ &\leq \|\chi\|^2 |(T^n - P)\tilde{U}| < \Lambda^n V. \end{aligned} \quad (12)$$

Another way to compute the same inner product is

$$\langle \chi, (T^n - P)\tilde{U}\chi \rangle = \langle \chi, T^n \tilde{U}\chi \rangle - V^2. \quad (13)$$

Combining Eq. 12 and Eq. 13 gives that $|\langle \chi, T^n \tilde{U}\chi \rangle - V^2| < \Lambda^n V$. This means that there exists C which depends only on \mathcal{A} such that if $n > C \log 1/\epsilon$ then $\Lambda^n < V$ and $\langle \chi, T^n \tilde{U}\chi \rangle > 0$. Specifically, it suffices to choose

$$n > \frac{d^2 - 1}{\log(1/\Lambda)} \log(1/\epsilon) + \frac{\log(2^{d^2-1}/k_1)}{\log(1/\Lambda)}. \quad (14)$$

When this occurs it means that

$$\int \chi(g) \sum_{w \in W_n} \frac{\chi(wU^{-1}g)}{(2|\mathcal{A}|)^n} dg > 0, \quad (15)$$

which implies that $\exists g \in SU(d)$ and $w \in W_n$ such that $\chi(g) \neq 0$ and $\chi(wU^{-1}g) \neq 0$. Thus $|g - I| < \epsilon/2$ and $|wU^{-1}g - I| < \epsilon/2$, implying that $|w - g^{-1}U| < \epsilon/2$. Combining these and using the triangle inequality gives $|w - U| < \epsilon$. ■

4 A class of efficiently universal gate sets

In this section we show that for each d there exists a set of gates \mathcal{G}_d in $SU(d)$ such that $\Lambda(\mathcal{G}_d) < 1$ (and thus \mathcal{G}_d is efficiently universal). We begin with a result demonstrating this for $SU(2)$ and then extend it to $SU(d)$.

Lemma 2 (Lubotsky, Phillips and Sarnak) *Let*

$$\begin{aligned} V_1 &= \frac{1}{\sqrt{5}} \begin{pmatrix} 1 & 2i \\ 2i & 1 \end{pmatrix}, V_2 = \frac{1}{\sqrt{5}} \begin{pmatrix} 1 & 2 \\ -2 & 1 \end{pmatrix} \\ \text{and } V_3 &= \frac{1}{\sqrt{5}} \begin{pmatrix} 1+2i & 0 \\ 0 & 1-2i \end{pmatrix}. \end{aligned} \quad (16)$$

Then $\lambda = \Lambda(\{V_1, V_2, V_3\}) = \frac{\sqrt{5}}{3} < 1$. Furthermore, for any $U_1, U_2, U_3 \in SU(2)$, $\Lambda(\{U_1, U_2, U_3\}) \geq \lambda$.

The proof of this Lemma is presented in [16, 17]. Let $\mathcal{G}_2 = \{V_1, V_2, V_3\}$, as it is a family of quantum gates from $SU(2)$ for which Λ is strictly less than one. The optimality of Λ for this set is an interesting aside, but has little bearing on what follows.

Extending the result to $SU(d)$ will require slightly more effort. To this end, if I_k denotes the $k \times k$ identity matrix, then, for any $U \in SU(2)$ and $2 \leq j \leq d$, define $\beta_j^{(d)}(U)$ to be

$$\beta_j^{(d)}(U) = \begin{pmatrix} I_{j-2} & 0 & 0 \\ 0 & U & 0 \\ 0 & 0 & I_{d-j} \end{pmatrix} \in SU(d). \quad (17)$$

We will typically omit the (d) where it is understood.

Lemma 3 (Diaconis and Shahshahani) Let $\{G_j^i\}, 1 \leq i < j \leq d$ be a series of $\binom{d}{2}$ independent random matrices in $SU(2)$ that are chosen uniformly according to a Haar measure. Then

$$\prod_{i=1}^{d-1} \prod_{j=i+1}^d \beta_j(G_j^i) \quad (18)$$

is uniformly distributed in $SU(d)$.

This Lemma is proved in [18]. It means that if we had access to random elements of $SU(2)$ that were completely uniformly distributed, then we could generate uniformly distributed elements of $SU(d)$. When the elements of $SU(2)$ are only approximately uniform, we can bound the distance to uniformity of the words they form by using what is known as a hybrid argument: [19]

Lemma 4 (Bernstein and Vazirani) If $U_1, \dots, U_m, V_1, \dots, V_m$ are linear operators such that $|U_i| \leq 1, |V_i| \leq 1$ and $|U_i - V_i| < \delta$, then $|U_m \cdots U_2 U_1 - V_m \cdots V_2 V_1| < m\delta$.

Proof If we replace a single U_i in the product $U_m \cdots U_1$ with the corresponding V_i , then the entire product will still change by less than δ since $|AB| \leq |A| \cdot |B|$ for any operators A, B . Thus we can construct a series of $m+1$ ‘‘hybrid’’ operators, which start with $U_1 \cdots U_m$, end with $V_1 \cdots V_m$ and are each separated by less than δ . The proof follows from the triangle inequality. ■

We now combine all of the other results in this section to demonstrate a set of gates in $SU(d)$ for which Λ is strictly less than one.

Proposition 5 For any $d > 2$, define \mathcal{G}_d by

$$\mathcal{G}_d = \{\beta_j(V) \mid 1 \leq j \leq (d-1), V \in \mathcal{G}_2\}. \quad (19)$$

Then $\Lambda(\mathcal{G}_d) < 1$.

Proof

The approach of our proof will be to approximate the uniform distribution in Lemma 3, and then we show that this forces Λ to be less than one. To this end, let $R_m \subset W_{m \binom{d}{2}}(\mathcal{G}_d)$ be the set of all products of the form

$$\prod_{i=1}^{d-1} \prod_{j=i+1}^d \beta_j(G_j^i) \quad (20)$$

such that the G_j^i are selected from $W_m(\mathcal{G}_2)$.

From Lemma 2 we have that $\forall m, |T(V_1, V_2, V_3) - P|^m = \lambda^m$ for some $\lambda < 1$. There are $\binom{d}{2}$ terms in Eq. 18, each of which is approximated to within an accuracy of λ^m by the appropriate length m substring of R_m . Thus, using the hybrid argument and Lemma 3 gives that

$$\Lambda(R_m) = \left| \sum_{w \in R_m} \frac{\tilde{w} - P}{|R_m|} \right| \leq \binom{d}{2} \lambda^m. \quad (21)$$

Now, if we let R'_m denote $W_{m \binom{d}{2}} - R_m$ then

$$\begin{aligned} \Lambda(W_{m \binom{d}{2}}) &\leq \frac{|R'_m|}{|W|} \Lambda(R'_m) + \frac{|R_m|}{|W|} \Lambda(R_m) \\ &\leq \left(1 - \frac{|R_m|}{|W|}\right) + \frac{|R_m|}{|W|} \Lambda(R_m) \\ &= 1 - \frac{|R_m|}{|W|} (1 - \Lambda(R_m)). \end{aligned} \quad (22)$$

If we choose m large enough so that $\binom{d}{2} \lambda^m < 1$, then this last expression will be less than one, and $\Lambda(\mathcal{G}_d) < 1$. ■

Thus, efficient quantum compiling is possible for d -dimensional systems, given the appropriate choice of base gate set.

5 Lower Bounds

This proves that sets of base gates exist which can achieve a precision of ϵ in $O(\log 1/\epsilon)$ gates, but can we do any better? An ϵ -ball in $SU(d)$ has measure of order ϵ^{d^2-1} , so if we expect to cover all of $SU(d)$ with strings of length n , then we will require $(2|\mathcal{A}|)^n k_2 \epsilon^{d^2-1} > 1$, or equivalently,

$$n \geq \frac{d^2 - 1}{\log 2|\mathcal{A}|} \log 1/\epsilon - \frac{\log k_2}{\log 2|\mathcal{A}|}. \quad (23)$$

Thus the result is optimal up to a constant factor. This fact is quite general, since it follows from simple counting arguments. However, if the assumptions of the problem are relaxed to allow many gates to act in parallel, then using ancilla qubits it is possible to approximate single-qubit gates with a circuit of *size* $\text{poly}(\log 1/\epsilon)$ but *depth* of only $\text{poly}(\log \log 1/\epsilon)$. [?] This construction, like the one in this paper, relies on having access to a specific set of base gates; to date, only the Solovay-Kitaev theorem applies to any computationally universal set.

In our original problem, though, eliminating the constant linear factor turns out to be impossible. Consider any set \mathcal{A} of l base gates that is *not* computationally universal. Let $B(\mathcal{A}, \delta)$ be the set of gates obtained by perturbing each gate in \mathcal{A} by no more than δ . Then $B(\mathcal{A}, \delta)$ has non-zero measure (in $SU(d)^l$), almost all of its elements are computationally universal and from the hybrid argument, any string of length n drawn from gates in $B(\mathcal{A}, \delta)$ will be within $n\delta$ of something in the (non-dense) group generated by \mathcal{A} . Since we can make δ arbitrarily small, any fixed prefactor in front of $\log 1/\epsilon$ will fail on a computationally universal set of non-zero measure for some values of ϵ .

Note that unlike most results about quantum compiling, this argument also holds if the base gates are parameterized; say, A_1, \dots, A_l are elements of the algebra $su(d)$ and a single operation now has the form $e^{\pm A_i t}$, for any $t > 0$. The above proof demonstrates that there exist sets with non-zero measure which require arbitrarily many steps, even if the steps are continuous. If we measure cost not in terms of number of steps, but by the total time taken, then we have to modify the argument slightly. For small values of t , $|e^{A_i t} - I|$ is on the order of $t\delta$, but for large t the difference never gets any higher than δ . This means that no matter how many steps we take, in time t , we will stay within $t\delta$ of some non-dense subgroup and the same result holds.

These results can be obtained more simply by considering the (non-zero measure) set of gates which are very close to the identity. If every gate does very little, then we will need a large number them in order to accomplish anything. The reason why universal sets that are very close to non-universal sets are interesting is because of their frequent appearance in actual physical systems, such as NMR under the weak coupling approximation.[20]

6 Conclusions

We have found a condition that implies the efficient universality of a set of gates and demonstrated a family of gate sets in $SU(d)$ that satisfy this condition. This means that given access to such a gate set, arbitrary quantum gates can be approximated to accuracy ϵ using only $O(\log 1/\epsilon)$ gates. Such knowledge will likely be invaluable in crafting future physical implementations of quantum information processing systems.

Many open questions remain, however. For example, determining or bounding Λ (even numerically) for a given set of base gates seems to be very difficult, though it is likely an important step in determining the prefactor C , which measures how effective a set of gates would be for compiling. The method used by [16, 17] involves specialized arguments from number theory that do not generalize easily to other sets of gates or to $SU(d)$ for $d > 2$. Our proof (like the Solovay-Kitaev algorithm) also requires the ability to perform the inverse of each gate in the base set. This restriction feels unnecessary, yet very little is known in the case where inverses are unavailable.

More broadly, it is also generally unknown which gate sets are efficiently universal and when $\Lambda < 1$. Note that $\Lambda(\mathcal{A}) < 1$ implies that \mathcal{A} is efficiently universal, but the converse is not known to be true. Thus it is possible that the questions of efficient universality and Λ being less than one will be settled separately.

However, if $\Lambda(\mathcal{A})$ were to be a continuous function of \mathcal{A} (for fixed $|\mathcal{A}|$), then the situation would simplify considerably. In this case, it is not hard to show that $\Lambda(\mathcal{A}) < 1$ if and only if \mathcal{A} is computationally universal, so that computational universality, efficient universality and $\Lambda < 1$ would all become equivalent conditions. We suspect that this is the case, but have been unable to prove it.

Finally, the techniques used in our results do not suggest any efficient (i.e. running time polynomial in $\log 1/\epsilon$) algorithms for quantum compiling. The most important, and possibly most difficult, open problem remaining is to find

a polynomial time algorithm to approximate any unitary gate by a fixed efficiently universal set of base gates with a string whose length saturates the $O(\log 1/\epsilon)$ bound.

We thank Persi Diaconis, Michael Freedman, Neil Gershenfeld, David Jerison, Seth Lloyd and Michael Nielsen for useful conversations and assistance, and are particularly grateful to Alexei Kitaev for pointing out a flaw in an earlier version of this paper and sharing with us a draft of [?]. AWH acknowledges support from the Army Research Office under the SUSPENSE program. ILC and BHR are supported by the Things That Think consortium. This work was also supported by the DARPA QuIST project on Quantum Architectures.

References

- [1] C. H. Bennett and D. P. DiVincenzo, *Nature* **404**, 247 (2000).
- [2] C. H. Bennett and P. W. Shor, *IEEE Transactions on Information Theory* **44**, 2724 (1998).
- [3] E. Rieffel and W. Polak, An introduction to quantum computing for non-physicists, <http://xxx.lanl.gov/abs/quant-ph/9809016>, 1998.
- [4] D. Aharonov, in *Annual Reviews of Computational Physics VI*, edited by D. Stauffer (World Scientific, Singapore, 1999).
- [5] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2000).
- [6] J. Preskill, *Proc. R. Soc. London A* **454**, 385 (1998).
- [7] D. Gottesman, *Phys. Rev. A* **57**, 127 (1998), arXiv e-print quant-ph/9702029.
- [8] P. O. Boykin *et al.*, 40th Annual Symposium on Foundations of Computer Science 486 (1999), arXiv e-print quant-ph/9906054.
- [9] D. Aharonov and M. Ben-Or, *SIAM J. Comp.* to appear (1999), arXiv e-print quant-ph/9906129.
- [10] S. Lloyd, *Phys. Rev. Lett.* **75**, 346 (1995).
- [11] D. Deutsch, A. Barenco, and A. Ekert, *Proc. R. Soc. London A* **449**, 669 (1995).
- [12] R. Solovay, *Lie Groups and Quantum Circuits*, MSRI, <http://www.msri.org/publications/ln/msri/2000/qcomputing/solovay/1/>, 2000.
- [13] A. Y. Kitaev, *Russ. Math. Surv.* **52**, 1191 (1997).
- [14] A. W. Harrow, Quantum compiling, undergraduate thesis, MIT Physics, 2001.
- [15] V. I. Arnold and A. L. Krylov, *Soviet Math. Dokl.* **4**, 1 (1962).
- [16] A. Lubotsky, R. Phillips, and P. Sarnak., *I. Comm. Pure Appl. Math.* **39**, S149 (1986).
- [17] A. Lubotsky, R. Phillips, and P. Sarnak., *I. Comm. Pure Appl. Math.* **40**, 401 (1987).
- [18] P. Diaconis and L. Saloff-Coste, *Comm. Math. Phys* **209**, 729 (2000).
- [19] E. Bernstein and U. Vazirani, *SIAM J. Comput.* **26**, 1411 (1997), arXiv e-print quant-ph/9701001.
- [20] R. R. Ernst, G. Bodenhausen, and A. Wokaun, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions* (Oxford University Press, Oxford, UK, 1997).