Asymptotically Optimal Approximation of Single Qubit Unitaries by Clifford and T Circuits Using a Constant Number of Ancillary Qubits

Vadym Kliuchnikov, Dmitri Maslov, and Michele Mosca

Phys. Rev. Lett. **110**, 190502 — Published 8 May 2013

DOI: 10.1103/PhysRevLett.110.190502

Asymptotically optimal approximation of single qubit unitaries by Clifford and T circuits using a constant number of ancillary qubits

Vadym Kliuchnikov,[1,*] Dmitri Maslov,[2,†] and Michele Mosca[1,‡]

[1] *Institute for Quantum Computing, University of Waterloo, Waterloo, ON, Canada*
[2] *National Science Foundation, Arlington, VA, USA*

Decomposing unitaries into a sequence of elementary operations is at the core of quantum computing. Information theoretic arguments show that approximating a random unitary with precision ε requires $\Omega(\log(1/\varepsilon))$ gates. Prior to our work, the state of the art in approximating a single qubit unitary included the Solovay-Kitaev algorithm that requires $O(\log^{3+\delta}(1/\varepsilon))$ gates and does not use ancillae and the phase kickback approach that requires $O(\log^2(1/\varepsilon)\log\log(1/\varepsilon))$ gates, but uses $O(\log^2(1/\varepsilon))$ ancillae. Both algorithms feature upper bounds that are far from the information theoretic lower bound. In this letter, we report an algorithm that saturates the lower bound, and as such it guarantees asymptotic optimality. In particular, we present an algorithm for building a circuit that approximates single qubit unitaries with precision ε using $O(\log(1/\varepsilon))$ Clifford and T gates and employing up to two ancillary qubits. We connect the unitary approximation problem to the problem of constructing solutions corresponding to Lagrange's four-square theorem, and thereby develop an algorithm for computing an approximating circuit using an average of $O(\log^2(1/\varepsilon)\log\log(1/\varepsilon))$ operations with integers.

PACS numbers: 03.67.Ac, 03.67.Lx

Introduction. The circuit-based model of quantum computation requires the ability to accurately implement quantum operations, specified by unitary matrices. These unitary operations are implemented in practice via classical control protocols, that must be designed to yield the desired quantum mechanical evolution of the system and optimized to achieve efficient control. Since quantum errors and decoherence appear to be unavoidable [1], fault tolerance must be employed to give any hope of scaling quantum computational devices to a point where they can outperform classical devices. Fortunately, quantum fault tolerance protocols allow only moderate overhead on the amount of physical resources required to accomplish the desired scaling [1, 2]. However, there is a downside—the transformations that may be implemented in such fault-tolerant protocols are limited to circuits over very specific gate libraries. In particular, quantum circuits based on the Clifford and T gates naturally arose in this context.

The Clifford gates often allow efficient implementation on the physical level [2, 3], and the T gate is required to accomplish quantum computations beyond those simulable classically [3], and thus to use quantum mechanics to its full computational advantage. Fault tolerant implementations of the T gate have been well-studied in the relevant literature [2, 3]. As an important point, independently of the details of the quantum information processing proposal used or the control protocol, Clifford and T circuits arose as one of the most widely accepted solution dictated by the requirements of fault tolerance.

The efficient approximation of a unitary evolution using a discrete universal gate set is crucial for building a scalable quantum computing device. Understanding the minimum possible size of approximating circuits is both a fundamental question in quantum information theory, and also a critical question for harnessing the power of quantum information for computing in practice. The efficiency of constructive solutions will play a significant role in determining the point at which available quantum computing resources will outperform existing classical computers. We show that the fundamental lower bounds on gate complexity for approximating an arbitrary unitary operation on a quantum fault tolerant processor may be achieved with efficient, constructive algorithms.

Barenco *et al.* [4] showed that any unitary may be implemented by a circuit with CNOT and single qubit gates, effectively reducing the problem to that of the single qubit unitary synthesis/approximation. In this letter, we report a constructive algorithm to saturate the information-theoretic lower bound on the number of gates required to approximate an arbitrary single qubit unitary to precision ε, using an additional resource in the form of two ancillae initialized to a simple state $|0\rangle$. The significance of the improvement provided by our approach is best seen when, for a fixed precision ε, all of the approximating circuit parameters such as depth, the number of gates, and total number of qubits used are combined into one aggregate figure, such as, e.g., the product of the three of these parameters. To further illustrate, $O(\text{Depth}\times\text{GateCount}\times\text{QubitsUsed})$ is $O(\log^{7.94}(1/\varepsilon))$ for the standard version of the Solovay-Kitaev algorithm [5], and $O(k\log^2(1/\varepsilon)\log(\log(1/\varepsilon)) + \log^4(1/\varepsilon)\log^3(\log(1/\varepsilon)))$ for implementing k single-qubit gates by phase kickback algorithm [6], whereas it is only $O(\log^2(1/\varepsilon))$ for our alrogithm reported in this paper. We next discuss the existing approaches to the solution of the single qubit synthesis problem as well as how our

approach improves the state of the art.

Background. Technically, the problem of the single qubit circuit synthesis is formulated as follows: given a discrete universal gate set or "library", find a sequence of gates in it that approximates a given unitary to precision ε. Parameter ε determines complexity of the resulting approximation.

Computing an approximation using the standard version of the Solovay-Kitaev algorithm [5] takes $O(\log^{2.71}(1/\varepsilon))$ steps on a classical computer and the number of gates in the resulting quantum circuit is $O(\log^{3.97}(1/\varepsilon))$. The best known upper bound on the circuit size resulting from the application of the Solovay-Kitaev algorithm is $O(\log^{3+\delta}(1/\varepsilon))$, where δ can be chosen arbitrary small [6]. Our gate count is $O(\log(1/\varepsilon))$, however, our circuits employ two ancillae.

From the other side, Harrow *et al.* [7] show an $\Omega(\log(1/\varepsilon))$ lower bound on the number of gates in the approximating circuit. A certain library of quantum gates that allows approximating a single qubit unitary to precision ε with a circuit containing at most $O(\log(1/\varepsilon))$ gates is also reported in [7]. However, authors did not provide an efficient algorithm to construct a circuit meeting the lower bound in the number of gates. Also, the gate set used, $\frac{I+2i\{X,Y,Z\}}{\sqrt{5}}$, is not considered to be well-suited for a fault-tolerant implementation, in contrast to the Clifford and T library. To the best of our knowledge, constructive saturation of the logarithmic lower bound in the Clifford and T library has not been shown yet, however, numerical evidence supports the theory that this is the case [8] (based on an exponential-time breadth first search algorithm).

Allowing additional resources helps to achieve interesting improvements over the Solovay-Kitaev algorithm. For example, using a special resource state $|\gamma\rangle$ on $O(\log(1/\varepsilon))$ qubits allows to achieve the desired accuracy of approximation by a depth $O(\log(\log(1/\varepsilon)))$ circuit containing $O(\log(1/\varepsilon)$ gates [6], also known as phase kickback algorithm. However, the resource state preparation requires $O(\log^2(1/\varepsilon))$ ancillary qubits and a circuit of depth $O(\log^2(\log(1/\varepsilon)))$ containing $O(\log^2(1/\varepsilon)\log\log(1/\varepsilon))$ gates. Furthermore, exact preparation of the resource state $|\gamma\rangle$ is not possible using gates from the Clifford and T library and qubits initialized to the state $|0\rangle$ [9, 10]. In comparison, in our work, we employ only two ancillae prepared in the simple state $|0\rangle$, which results in achieving the approximating accuracy of ε using a circuit with $O(\log(1/\varepsilon))$ gates. Also, our circuit is asymptotically optimal.

One other recent approach uses resource states [11] and probabilistic circuits with classical feedback. The circuit itself, excluding state preparation, requires on average a constant number of operations and a constant number of ancilla qubits. The method requires precomputed ancillae in the states $R_Z(2^n\phi)H|0\rangle$ to implement $R_Z(2^m\phi)$. The other recently developed method to approximate $R_Z(\phi)$ that also relies on special resource states, measurements and classical feedback presented in [12]. Our algorithm does not rely on the measurements and classical feedback, and our circuit is deterministic. More importantly, our algorithm does not employ sophisticated ancilla states that, in turn, may require approximation, as they may not be possible to prepare exactly in the Clifford and T library [9, 10].

In our previous work [9], we showed that any single qubit unitary with entries u_{ij} in the ring $\mathbb{Z}\left[i, 1/\sqrt{2}\right]$ can be synthesized exactly using single qubit Clifford and T gates. We presented an asymptotically optimal algorithm for finding a circuit with the minimal number of Hadamard and T gates and asymptotically minimal total number of gates. More precisely, if the square of the norm of an element of the single qubit unitary matrix, $|u_{ij}|^2$, can be represented as $(a+\sqrt{2}b)/2^n$, where a and b are integers such that $GCD(a,b)$ is odd, the total number of gates required to synthesize the unitary is in $\Theta(n)$. This work opened the door for bypassing the Solovay-Kitaev algorithm for fast circuit approximation of single qubit unitaries by efficiently approximating arbitrary unitaries with unitaries over the ring $\mathbb{Z}\left[i, 1/\sqrt{2}\right]$. However, as of the time of this (original) writing, no efficient ring round-off procedure was reported, and it remains an important open problem.

Giles and Selinger [10] recently found an elegant way to prove the conjecture formulated in [9] stating that multiple qubit unitaries over the ring $\mathbb{Z}\left[i, 1/\sqrt{2}\right]$ may be synthesized exactly using Clifford and T library. In this letter, we employ some of their results to show that, by adding at most two ancilla qubits, we can achieve asymptotically optimal approximation of the single qubit unitaries in the Clifford and T library.

Main result. We focus on the approximation of the following operator:

$$\Lambda(e^{i\phi}) : \alpha|0\rangle + \beta|1\rangle \mapsto \alpha|0\rangle + \beta e^{i\phi}|1\rangle .$$

We note that any single qubit unitary can be decomposed in terms of a constant number of Hadamard gates and $\Lambda(e^{i\phi})$ (see solution to Problem 8.1 in [6]). Therefore, the ability to approximate $\Lambda(e^{i\phi})$ implies the ability to approximate any single qubit unitary.

There are two main steps in our algorithm:

1. Find a circuit C consisting of Clifford and T gates such that the result of applying C to $|00\rangle$ is close to $e^{i\phi}|00\rangle$.

2. Apply circuit C controlled on the first qubit to perform a transformation close to:

$$\alpha|000\rangle + \beta|100\rangle \mapsto \alpha|000\rangle + \beta e^{i\phi}|100\rangle .$$

It can be observed that the net effect of such transformation may be described as the application of $\Lambda(e^{i\phi})$ to

the first qubit. To accomplish the first step we approximate $e^{i\phi}|00\rangle$ with a four dimensional vector $|v\rangle$ with entries in the ring $\mathbb{Z}\left[i, 1/\sqrt{2}\right]$. We then employ an algorithm for multiple qubit exact synthesis to find a circuit C that prepares $|v\rangle$ starting from $|00\rangle$ using at most one ancilla qubit. It was shown in [13] that any circuit using Clifford and T gates can be transformed into its exact (meaning no further approximation is required) controlled version with only a linear overhead in the number of gates, and using at most one ancilla qubit in the state $|0\rangle$ that is returned unchanged. Our analysis shows that, however, on this step we do not need to use this additional ancilla. The resulting total number of ancillae is thus at most two.

Approximating $e^{i\phi}|00\rangle$. The key is the reduction of the approximation problem to expressing an integer number as a sum of four squares. In particular, we are looking for an approximation of:

$$e^{i\phi}|00\rangle = (\cos(\phi) + i\sin(\phi), 0, 0, 0)$$

by a unit vector:

$$|v\rangle := \frac{1}{2^k}\left(\left\lfloor 2^k\cos(\phi)\right\rfloor + i\left\lfloor 2^k\sin(\phi)\right\rfloor, 0, a+ib, c+id\right),$$

where $k \in \mathbb{N}; a, b, c, d \in \mathbb{Z}$. Without loss of generality we can assume that $0 \leq \phi \leq \frac{\pi}{4}$. The power k of the denominator determines precision of our approximation and complexity of the resulting circuit. As $|v\rangle$ must be a unit vector, the remaining four parameters (a, b, c, and d) should satisfy the integer equation:

$$a^2 + b^2 + c^2 + d^2 = 4^k - \left\lfloor 2^k\cos(\phi)\right\rfloor^2 - \left\lfloor 2^k\sin(\phi)\right\rfloor^2.$$

Lagrange's four square theorem states that this equation always has a solution. Furthermore, there exists an efficient probabilistic algorithm for finding a solution to the Diophantine equation. For the right hand side M it requires on average $O(\log^2(M)\log\log M)$ operations with integers smaller than M. It is described in Theorem 2.2 in [14]. We get such a simple round off procedure and reduction to such a simple Diophantine equation at the expense of using two qubits instead of one.

Furthermore, in estimating the classical complexity of the algorithm for finding the approximating circuit, we will rely on an observation that

$$4^k - \left\lfloor 2^k\cos(\phi)\right\rfloor^2 - \left\lfloor 2^k\sin(\phi)\right\rfloor^2 \leq 4\times 2^k + Const \in O(2^k).$$

The exact synthesis method for finding a circuit that prepares $|v\rangle$ given $|0\rangle$ is based on the connection between the form of the elements of vector $|v\rangle$ and the complexity of the corresponding circuit. More precisely, square of the absolute value of each element of $|v\rangle$ can be written as $(a+\sqrt{2}b)/\sqrt{2}^n$, where n is minimized across all equivalent representations. The maximum of such n over all elements of $|v\rangle$ defines the complexity of the state preparation. In particular, it was shown in [10] that it is always possible to reduce maximal n or the number of elements of $|v\rangle$ with maximal n using finitely many two-level unitaries. Furthermore, each of those two-level unitaries can be implemented exactly using finitely many Clifford and T gates. In summary, one can always find a sequence of Clifford and T gates reducing maximal n to 0. This sequence defines the circuit synthesizing the desired state $|v\rangle$ given $|0\rangle$.

Precision and complexity analysis. Let us introduce $\gamma = \left(\left\lfloor 2^k\cos(\phi)\right\rfloor + i\left\lfloor 2^k\sin(\phi)\right\rfloor\right)/2^k$ and express $|v\rangle$ as $|v\rangle = \gamma|00\rangle + |1\rangle\otimes|g\rangle$. The application of the circuit C controlled on the first qubit will transform $(\alpha|0\rangle + \beta|1\rangle)\otimes|00\rangle$ into $\alpha|000\rangle + \beta\gamma|100\rangle + \beta|01\rangle\otimes|g\rangle$. The distance of the result to the desired state $\alpha|000\rangle + \beta e^{i\phi}|100\rangle$ is:

$$\sqrt{|\beta(e^{i\phi}-\gamma)|^2 + |\beta|^2\,\||g\rangle\|^2}.$$

By the choice of γ we have $\left|\gamma - e^{i\phi}\right| \leq \frac{\sqrt{2}}{2^k}$, therefore the first term in the sum above is in $O(1/2^{2k})$. The norm squared of $|g\rangle$ equals $1-|\gamma|^2$. The complex number γ approximates $e^{i\phi}$, and the distance of its absolute value to identity can be estimated using the triangle inequality, $\left||\gamma| - \left|e^{i\phi}\right|\right| \leq \left|\gamma - e^{i\phi}\right|$. Therefore, $1-|\gamma|^2$ is in $O(1/2^k)$. In summary, the distance to the approximation is in $O(1/2^{0.5k})$.

The same estimate is true if we consider the circuit C as a part of a larger system. In this case we should start with the state $(\alpha|\phi_0\rangle\otimes|0\rangle + \beta|\phi_1\rangle\otimes|1\rangle)\otimes|00\rangle$. Similar analysis shows that the distance to approximation remains $O(1/2^{0.5k})$.

As shown in [10], it is possible to find a circuit that prepares $|v\rangle$ using $O(k)$ Clifford and T gates ([10], Lemma 20 (Column lemma)). The classical complexity of constructing a quantum circuit implementing $|v\rangle$ is in $O(k)$. In the controlled version of this circuit the number of gates remains $O(k)$ ([13], Theorem 1). In summary, we need $O(\log(1/\varepsilon))$ gates to achieve precision ε. The complexity of the classical algorithm for constructing the entire approximating circuit is thus dominated by the complexity of finding a solution to the Diophantine equation, which is in $O(\log^2(1/\varepsilon)\log\log(1/\varepsilon))$, counting operations over integers of size $O(\log(1/\varepsilon)$.

How many ancillae are needed? A straightforward calculation shows that the number of ancillae used is three. However, we can get around using only two ancillae. To understand how, we need to go into the details of the proof of Lemma 20 (Column lemma) from [10]. It shows how to find a sequence of two-level unitaries of type iX, $T^{-m}(iH)T^m$, and W [10] and length $O(k)$ that allows to prepare a state with the denominator 2^k. A controlled version of the two level unitary is again a two level unitary. In [10], Lemma 24, it was also shown that any such unitary required can be implemented using no extra ancillae. Therefore, the controlled version of the

circuit C will not use any additional ancilla and we need only two of them in total.

Lower bound on the number of gates when ancillae are allowed. We use a volume argument to show the lower bound. Suppose we can approximate any element of the group of N by N unitary matrices $\mathbb{U}(N)$ with precision ε by a circuit over gate library G that uses at most k gates. This implies that we can cover $\mathbb{U}(N)$ with $|G|^k$ sets, where each set contains such unitaries from $\mathbb{U}(N)$ that can be approximated by some particular circuit of length k. By showing that the Haar measure of each of the mentioned sets is in $O(\varepsilon^{N^2})$, and using the notion that the measure of the set union is smaller then the sum of the measures of the individual sets, we obtain the required bound on k.

The idea is similar to the derivation of the lower bound for the case when no ancillae allowed, originally found in [7]. The difference is that we have to deal with the circuits acting on $n+m$ qubits and consider a more complicated notion of approximation, in contrast to the usual distance between two unitaries. The precise statement of the lower bound is achieved by the following lemma:

Lemma 1 *Let G be a universal library, and let M_V be a set of unitaries, that simulate a unitary V acting on n qubits, using m ancillary qubits:*

$$M_V := \left\{U \in \mathbb{U}\left(2^{m+n}\right) | U\left(|0\rangle \otimes |\phi\rangle\right) = |0\rangle \otimes \left(V|\phi\rangle\right)\right\}.$$

Then, for any ε there always exists a unitary $V(\varepsilon)$ such that the number of gates from G needed to construct a unitary within the distance ε to $M_{V(\varepsilon)}$ is in $\Omega(\log(1/\varepsilon))$.

The proof of this lemma may be found in the Supplemental Material.

Conclusions and future work. Our work answers a fundamental and important question for both theoretical and practical quantum information science: up to constant factors, the fundamental limits for approximating single qubit unitaries to a given precision may be attained by efficient algorithms.

Our work also opens up several other interesting and important questions (in no specific order): what are the constants hidden behind the big-O notation in our approach, and can they be optimized (while further optimizations are only possible up to a multiplicative factor they are, nevertheless, important for practical purposes)? What are the possible trade-offs between adding/reducing ancillae and the gate count? Is it possible to use other efficiently solvable Diophantine equations to discover approximations of other types of gates? Lastly, does there exist an efficient algorithm to round off single-qubit unitaries to those single-qubit unitaries over the ring $\mathbb{Z}\left[i, 1/\sqrt{2}\right]$ and avoid the need for ancillary qubits altogether?

Further development of the ideas reported in this letter has already led to some interesting results. An efficient algorithm for approximating a unitary by Clifford and T circuits without using ancillae and leading to shorter sequences may be found in [15]. [16] allows to find even shorter approximating circuits at the expense of a more intensive (classical) computation. Finally, [17] shows how to use similar ideas to efficiently approximate unitaries over the gate set $\frac{I+2i\{X,Y,Z\}}{\sqrt{5}}$.

Authors supported in part by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center Contract number DllPC20l66. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC or the U.S. Government.

This material is based upon work partially supported by the National Science Foundation (NSF), during D. Maslov's assignment at the Foundation. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Michele Mosca is also supported by Canada's NSERC, MPrime, CIFAR, and CFI. IQC and Perimeter Institute are supported in part by the Government of Canada and the Province of Ontario.

We wish to thank Alex Bocharov, Troy W. Borneman, Martin Roetteler, and Peter Selinger for their comments and helpful discussions.

* v.kliuchnikov@gmail.com; David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada

† dmitri.maslov@gmail.com; Department of Physics & Astronomy, University of Waterloo, Waterloo, ON, Canada

‡ mmosca@iqc.ca; Deptartment of Combinatorics & Optimization, University of Waterloo, Waterloo, ON, Canada; Perimeter Institute for Theoretical Physics, Waterloo, ON, Canada

[1] J. Preskill, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **454**, 385 (Jan. 1998), ISSN 1364-5021.

[2] A. G. Fowler, A. M. Stephens, and P. Groszkowski, Physical Review A **80**, 052312 (Nov. 2009), ISSN 1050-2947, arXiv:arXiv:0803.0272v4.

[3] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*, 10th ed. (Cambridge University Press, New York, NY, USA, 2011) ISBN 9781107002173.

[4] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Physical Review A **52**, 3457 (Nov. 1995), ISSN 1050-2947, arXiv:quant-ph/9503016.

[5] C. M. Dawson and M. A. Nielsen, Quantum Information & Computation **6**, 81 (May 2005), arXiv:quant-ph/0505030.
[6] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation*, Graduate studies in mathematics, v. 47 (American Mathematical Society, Boston, MA, USA, 2002) ISBN 9780821832295.
[7] A. W. Harrow, B. Recht, and I. L. Chuang, Journal of Mathematical Physics **43**, 4445 (Nov. 2002), ISSN 00222488, arXiv:quant-ph/0111031.
[8] A. G. Fowler, Quantum Information & Computation **11**, 8 (Nov. 2011), arXiv:quant-ph/0411206.
[9] V. Kliuchnikov, D. Maslov, and M. Mosca, Quantum Information & Computation **13**, 0567 (Jul. 2013), arXiv:1206.5236.
[10] B. Giles and P. Selinger, Physical Review A **87**, 032332 (Mar 2013), arXiv:1212.0506.
[11] N. C. Jones, J. D. Whitfield, P. L. McMahon, M.-H. Yung, R. V. Meter, A. Aspuru-Guzik, and Y. Yamamoto, New Journal of Physics **14**, 115023 (2012), arXiv:1204.0567.
[12] G. Duclos-Cianci and K. M. Svore(Oct. 2012), arXiv:1210.1980.
[13] M. Amy, D. Maslov, M. Mosca, and M. Roetteler(Jun. 2012), arXiv:1206.0758.
[14] M. O. Rabin and J. O. Shallit, Communications on Pure and Applied Mathematics **39**, S239 (1986), ISSN 00103640.
[15] P. Selinger(Dec. 2012), arXiv:1212.6253.
[16] V. Kliuchnikov, D. Maslov, and M. Mosca(Dec. 2012), arXiv:1212.6964.
[17] A. Bocharov, Y. Gurevich, and K. M. Svore(Mar. 2013), arXiv:1303.1411.

Supplemental Material

Proof of Lemma 1:

Let $N = 2^n$, ρ be the distance induced by Frobenius norm and μ be the Haar measure on $\mathbb{U}(N)$. For a unitary U from $\mathbb{U}(2^{m+n})$ we define a set of unitaries from $\mathbb{U}(N)$ that can be approximated by U with precision ε as:

$$b\left(U,\varepsilon\right) := \left\{V \in \mathbb{U}(N) | \rho\left(M_V, U\right) \leq \varepsilon\right\}.$$

Let G^k be the set of all unitaries that can be constructed using k gates from the library G. Suppose that for any unitary V we can find a unitary U from G^k within the distance ε from M_V. In other words, we can cover $\mathbb{U}(N)$ with sets $b\left(U,\varepsilon\right)$ when U goes through all unitaries that can be implemented using circuits of length k. This implies:

$$\mu\left(\mathbb{U}(N)\right) \leq \sum_{U\in G^k} \mu\left(b\left(U,\varepsilon\right)\right) \leq |G|^k \max_{U\in G^k} \mu\left(b\left(U,\varepsilon\right)\right).$$

We will show that $\mu\left(b\left(U,\varepsilon\right)\right)$ is upper bounded by $C_0\varepsilon^{N^2}$, for some constant C_0, and, therefore:

$$k \geq \frac{1}{\log|G|} \log\left(\frac{\mu\left(\mathbb{U}\left(N\right)\right)}{C_0\varepsilon^{N^2}}\right). \tag{1}$$

We next show that we can always find a unitary V_U from $\mathbb{U}(N)$ such that the set $b\left(U,\varepsilon\right)$ is contained in the ball $b(V_U, 2\varepsilon)$. This will give us the required bound on $\mu\left(b\left(U,\varepsilon\right)\right)$. Indeed, Haar measure of the ball $b(V_U, 2\varepsilon)$ does not depend on V_U and equals to $b(I, 2\varepsilon)$. As $\mathbb{U}(N)$ is a smooth manifold of dimension N^2, the quantity $\mu(b(I, 2\varepsilon))$ is upper bounded by $C_0\varepsilon^{N^2}$ for some positive constant C_0. We proceed to the construction of V_U.

It suffices to consider only the cases when the set $b\left(U,\varepsilon\right)$ is non-empty. Let V_U be any element of $b\left(U,\varepsilon\right)$. We first show that $b\left(U,\varepsilon\right)$ is contained in the ball $b(U_0,\varepsilon)$, where U_0 is an $N\times N$ complex matrix, but not necessary unitary. We second show that $\rho(V_U, U_0) \leq \varepsilon$. Let U_0 be a submatrix of U defined as follows:

$$U_0 := \left\{\langle e_i| \otimes \langle 0|\right) U \left(|0\rangle \otimes |e_j\rangle\right\},$$

where $\{|e_i\rangle\}$ is the standard (computational) basis in $\mathbb{C}(N)$. Taking into account that the distance ρ is induced by Frobenius norm, we write $\rho\left(U, M_V\right) \geq \rho\left(U_0, V\right)$. This implies:

$$b(U,\varepsilon) \subseteq b(U_0,\varepsilon) = \left\{V \in \mathbb{U}(N) | \rho\left(U_0, V\right) \leq \varepsilon\right\}.$$

As V_U is an element of $b\left(U,\varepsilon\right)$ we also have $\rho\left(U_0, V_U\right) \leq \varepsilon$.

Estimate (1) on k shows that we need circuits of the size at least $\Omega(\log(1/\varepsilon))$ to cover the full group $\mathbb{U}\left(N\right)$. If k is chosen in such a way that the inequality (1) does not hold, due to the volume argument, there exists a unitary $V\left(\varepsilon\right)$ such that it is not possible to approximate any unitary from $M_{V(\varepsilon)}$ with precision ε using at most k gates.