Efficient Quantum Pseudorandomness

Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki

Phys. Rev. Lett. **116**, 170502 — Published 29 April 2016

DOI: 10.1103/PhysRevLett.116.170502

Efficient Quantum Pseudorandomness

Fernando G.S.L. Brandão,[1,2] Aram W. Harrow,[3] and Michał Horodecki[4]
[1]*Microsoft Research Redmond*
[2]*Computer Science Department, University College London*
[3]*Center for Theoretical Physics, Massachusetts Institute of Technology*
[4]*National Quantum Information Center of Gdansk*

Randomness is both a useful way to model natural systems and a useful tool for engineered systems, e.g. in computation, communication and control. Fully random transformations require exponential time for either classical or quantum systems, but in many case pseudorandom operations can emulate certain properties of truly random ones. Indeed in the classical realm there is by now a well-developed theory of such pseudorandom operations. However the construction of such objects turns out to be much harder in the quantum case. Here we show that random quantum unitary time evolutions ("circuits") are a powerful source of quantum pseudorandomness. This gives the for the first time a polynomial-time construction of quantum unitary designs, which can replace fully random operations in most applications, and shows that generic quantum dynamics cannot be distinguished from truly random processes. We discuss applications of our result to quantum information science, cryptography and to understanding self-equilibration of closed quantum dynamics.

Random processes are ubiquitous both in natural and engineered systems. They are both an effective way to model many systems, and also a vital tool in algorithms, communication, control, cryptography, and elsewhere. However a random function on n bits is known to require $\exp(\Omega(n))$ elementary operations[43] to implement and a similar number of random bits to even specify [40], meaning that in fact such random functions are neither found in nature, nor can they can be designed on a computer. Instead, we now know many methods for engineering *pseudorandom* functions using far less randomness. These pseudorandom functions can be proven to be indistinguishable from truly random functions either by any test that examines their first k moments [2] (in which case they are called k-designs), or by any computationally limited test [18] (for which case the term *pseudorandom function* is usually reserved). This Letter will focus on k-designs.

While these constructions mean that a carefully designed algorithm can simulate a random function in many circumstances, they do not speak to the question of whether we should expect natural processes to also resemble random functions. However, it was proved by [9, 27] that even reasonably short random reversible circuits yield approximate k-designs, meaning that they approximate well the first k moments of a truly random function. These circuits are defined to be sequences of basic reversible operations, each involving three bits, which is the simplest type of reversible circuit that is computationally universal. As such, they form plausible toy models for the dynamics of actual systems, and provide some rigorous justification for the intuition that generic dynamics cannot easily be distinguished from fully random functions.

In recent decades, quantum mechanics has been found to dramatically change the nature of information and information-processing [36], implying among other things potential new applications such as quantum cryptography and computation. The above story needs then to be modified to account for quantum mechanics. The problem of quantum pseudorandomness was posed in [16] where it was asked to what extent random quantum circuits can mimic random quantum transformations. The simplest quantum systems are 2-level systems, called *qubits*, and any larger quantum system can be decomposed into some number n of qubits (e.g. the state of m fermions in n modes can be expressed using $\lceil \binom{n}{m} \rceil$ qubits). A quantum circuit is a sequence of gates, each acting on a constant number of qubits. Short quantum circuits are roughly equivalent in power to time evolution by local Hamiltonians for short times. The uniform distribution over unitary matrices is called the *Haar measure* and, as in the classical case, has been extensively studied as a model of natural generic dynamics [20] with applications to black holes [26, 41], quantum information processing [1, 24, 39] and elsewhere. In a further parallel to the classical case, Haar-random unitary matrices on n qubits cannot be implemented, even approximately, without $\Omega(4^n)$ elementary operations and $\Omega(4^n)$ random bits [29].

We thus have the same need for quantum pseudorandomness and an analogous notion of unitary k-designs. Again unitary k-designs can be used in place of Haar-random unitaries in most applications (e.g. for encoding quantum information to protect from errors [1], or realizing more efficient quantum process tomography [16]), but it has been much harder to prove that efficient unitary k-designs exist. It was long conjectured that random quantum circuits yield approximate k-designs, but for a long time, this was known only for $k = 1$; in other words, only the first moments of a quantum state were known to rapidly equilibrate under random dynamics. However, 1-designs can be realized even without entangling operations (e.g. a random product of Pauli matrices will suffice), while Haar-uniform unitaries create states with large amounts of entanglement [25], so 1-designs do not give a qualitatively good fit for the Haar measure. A better, but still imperfect, goal is to achieve a 2-design. It has been known how to efficiently engineer a 2-design on a quantum computer by selecting a random element of the so-called *Clifford group*—a restricted class of quantum gates—which involves creating significant entanglement while still performing operations far simpler than those resulting from the Haar measure [14].

Initial numerical work suggested that indeed random quantum circuits were approximate 2-designs [3, 16]. Later work was able to establish that random circuits matched the entanglement properties of 2-designs [13, 37], and finally that they in fact were approximate 2-designs [3, 15, 22, 28]. Since even the Clifford group, which is not universal for quantum computation, yields a 2-design, this too is a limited proxy for the Haar measure. Later work achieved 3-designs [7, 22] (see also [11]). In this work we settle the question and achieve k-designs for any k via circuits of length $\mathrm{poly}(n, k)$. Full details are given in [6], where it is also shown that this work cannot be substantially improved. We follow part of the framework of [10], which conjectured our result and gave a mean-field argument supporting it.

Definitions: Let us give a more precise definition of approximate unitary designs. First we say a probability measure μ on $\mathbb{U}(d)$ (the group of $d \times d$ unitary matrices) is a unitary k-design if for every monomial $q(U) = U_{i_1 j_1} \dots U_{i_k j_k} U^*_{m_1 n_1} \dots U^*_{m_k n_k}$ of degree at most k, in the entries of the unitary U_{ij} and their complex conjugate U^*_{nm}, the average of $q(U)$ over the Haar measure is the same as the average over μ.

In turn, we say a measure μ on $\mathbb{U}(d)$ forms an ε-approximate k-design if

$$|\mathbb{E}_{\text{Haar}} q(U) - \mathbb{E}_\mu q(U)| \leq \varepsilon \qquad (1)$$

with $\mathbb{E}_{\text{Haar}}, \mathbb{E}_\mu$ the expectations over the Haar measure and μ, respectively. There are other definitions of ε-approximate k-designs [6], but they turn out to be equivalent to the one above, up to a rescale of the approximation factor [32].

We model random quantum circuits as random walks on $\mathbb{U}(2^n)$ following Refs. [7, 26]. In each step of the walk an index i is chosen uniformly at random from $\{1, \dots, n-1\}$ and a unitary $U_{i,i+1}$ drawn from the Haar measure on $\mathbb{U}(4)$ is applied to the two neighboring qubits i and $i+1$. This is illustrated in Fig. 1. Other choices of random circuits are possible (e.g. considering non-nearest-neighbor gates or gates from different universal sets as in Ref. [17]) and variants of our results will apply to them as well. However here for concreteness we focus on the model above. A related model of random walk on the unitary (or orthogonal) group is Kac's random walk, extensively studied in connection to statistical mechanics (see e.g. [38]).

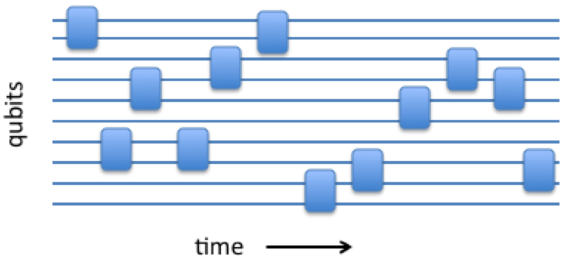

FIG. 1: In our model of a random quantum circuit, there are n qubits on a line, here arranged vertically. Time proceeds from left to right. In each time step, two adjacent qubits are chosen at random and a random two-qubit unitary operator is applied to them.

Our main result is the following:

Theorem 1. *Random circuits with $O(nk^9(nk + \log(1/\epsilon)))$ gates form ϵ-approximate k-designs.*

We comment briefly on these parameters. The n^2 dependence cannot be improved without changing the graph or using non-random circuits, since it corresponds to $O(n)$ layers of $O(n)$ gates each and with fewer layers the circuit would not be able to destroy correlations between the endpoints of the chain. The $\log(1/\epsilon)$ scaling also cannot be improved, because of the possibility of an unlucky choice of the random gates. As for the k dependence, Proposition 8 of [6] shows that the circuit size should grow at least linearly with k. There is evidence for this scaling from a mean-field argument [10] and numerical studies [12]. However, no previous work could rule out the possibility that the scaling with k would be exponential or even worse. Our result is the first proof that the circuit size scales polynomially with k. While the degree of the polynomial is too large to use in practical applications, we expect that a more careful analysis would reduce its degree perhaps even to being linear in k.

Applications: As k-wise independent distributions (or classical k-designs) have widespread applications, so too do approximate unitary k-designs [1, 16, 31].

Here we briefly outline one application of our result to the problem of understanding equilibration in closed quantum dynamics. Consider the unitary time evolution of a system, initially in a fixed state, say all spins up $|\uparrow\rangle^{\otimes n}$. The total state at any particular time is pure and hence does not appear to equilibrate in any sense. However a long sequence of investigations, starting with von Neumann in 1929 [35], has elucidated that the state does equilibrate if one imposes constraints on the kind of observations possible [4, 19, 30, 42]. For instance suppose that one only has access to measurements on a few of the particles. Then it turns out that the building up of entanglement in the quantum state leads to local equilibration of every small subset of particles, for almost all times [30]. The limits of equilibration in closed quantum dynamics is an interesting problem. What is the largest class of observables for which equilibration holds? Our result on unitary designs allows us advance this question significantly.

Define the circuit complexity of a measurement as the minimum size of any circuit of two-qubit gates that implements the measurement. Physical measurements (e.g. measurement of magnetization or heat capacity or even topological invariants) generally have low complexity. We show that in generic quantum dynamics given by random circuits (which model the case of generic evolutions under time-dependent Hamiltonians), the system equilibrates with respect to all measurements of low complexity. A general quantum measurement with two outcomes can be represented by operators M and $I - M$, with $0 \leq M \leq I$. We then have:

Corollary 2. *For every $k \geq 1$, for sufficiently large n and almost all random circuits U of size $O(n^{11k+9})$ on n qubits,*

$$\left| \langle \uparrow^{\otimes n} | U^\dagger M U | \uparrow^{\otimes n} \rangle - \frac{\mathrm{tr}(M)}{2^n} \right| \leq 2^{-n/4}. \qquad (2)$$

for every measurement $\{M, I - M\}$ of circuit complexity less than n^k.

Proof sketch: If M is fixed and U is Haar uniform then (2) holds with high probability; large deviations are suppressed with probability exponential in the dimension, meaning $\exp(-2^{\Omega(n)})$. If instead U is drawn from a t-design then this probability becomes $\exp(-t)$. We can approximate any low-complexity M with a measurement drawn from a set of size roughly $\exp(n^k)$. Thus (2) holds (approximately) for *all* low-complexity M with probability $\leq \exp(n^k - t)$. For this to be $\ll 1$ we need $t \gg n^k$, which according to Theorem 1 can be achieved by a random circuit of length $O(n^{11k+9})$. The full details of this proof are in [6].

One interpretation of Corollary 2 is that long random quantum circuits appear Haar-random when tested by significantly shorter quantum circuits, even when the tester has a full description of the random circuit. Thus it is crucial that the tester be weaker than the circuit U being tested; otherwise the tester could apply the random circuit, then apply $U^\dagger$ and verify that the resulting dynamics are trivial (say by applying them to half of a maximally entangled state).

Another interpretation of Corollary 2 is in the context of quantum cryptography. It gives a procedure for so called *quantum data hiding* against a computationally-bounded adversary, meaning that information is present in a state but cannot be measured without using a long quantum computation (cf. [14]). Indeed, Corollary 2 gives that all but a $2^{-\Omega(n)}$-fraction of states generated by circuits of size $O(n^{11k+9})$ cannot be distinguished from the maximally mixed state with bias larger than $n^{-\Omega(1)}$ by any circuit of size n^k. So whether one has the particular pure state or the maximally mixed state is hidden from any adversary that is constrained to run in time n^k.

One situation where the assumption is satisfied is in the somewhat unrealistic setting when the adversary has less computational power than the honest parties. A more realistic situation is when the time it takes the adversary to decode the message is longer than the time is takes to send the message from one honest party to the other. In this case the honest parties can abort the protocol if the message is not delivered in time.

Proof overview: The remainder of the paper will give a high-level description of the proof of Theorem 1. A full proof is in [6]. The proof is based on an interplay of techniques from quantum many-body theory [33], representation theory [23], and the theory of Markov chains [38], and we believe similar ideas might find further applications elsewhere.

Expressing the problem in terms of spectral properties of matrices: Classical case. As a warmup to understanding the properties of our random circuits, consider the classical analogue. If C is a t-gate reversible classical circuit acting on n bits, then we can think of it as a permutation matrix of size 2^n. Since mixing over the set of all permutations requires exponentially long circuits, we instead examine the behavior of the moments of the circuit. To represent the circuit's k'th moments, we can examine its action on sets of k inputs, each of which are n-bit strings. Using the tensor product, this action can be also be described as a matrix: $C^{\otimes k}$, which maps $|i_1\rangle \otimes \ldots \otimes |i_k\rangle$ to $C|i_1\rangle \otimes \cdots \otimes C|i_k\rangle$. The advantage of this representation is that the average over t-step circuits of $C^{\otimes k}$ (call it $A_{t,k}$) is simply the t^{th} power of the average over one-step circuits $A_{1,k}$; i.e. $A_{t,k} = A_{1,k}^t$. Moreover, if the gate set is universal then $A_{t,k}$ will approach the average over all permutations as $t \to \infty$.

Determining the rate of convergence now reduces to an eigenvalue problem. Since $A_{\infty,k} = A_{1,k}^\infty$, it must have only eigenvalues 0 or 1. The 1-eigenspace corresponds to the degrees of freedom that are preserved when the same circuit is applied to each of $i_1, \ldots, i_k$; e.g. information about whether $i_1 = i_2$ or $i_1 \neq i_2$. When the set of gates is universal, the matrix $A_{1,k}$ has the same eigenspace with eigenvalue one and has all other eigenvalues smaller than one (i.e. there are no additional "constants of motion"). Thus everything orthogonal to this subspace will decay to 0 as $t \to \infty$ at a rate controlled by the eigenvalues of $A_{1,k}$.

Our distance after t steps to the average over random permutations can be quantified by $\|A_{t,k} - A_{\infty,k}\|$. Due to the above arguments, this is given just by $(1-\delta)^t$, where $1-\delta$ is the second-largest eigenvalue of $A_{1,k}$ (disregarding multiplicity). This is the source of the exponential convergence typically exhibited by Markov chains on discrete state spaces. As a result, error ϵ is achieved by taking circuit length $t \geq \delta^{-1} \log \frac{1}{\epsilon}$. By proving [9, 27] that $\delta \geq 1/\operatorname{poly}(k, n)$, it follows that n-bit circuits of length $\operatorname{poly}(k, n)$ have k-th moments that approximate those of random permutations.

Expressing the problem in terms of spectral properties of matrices: Quantum case. In the quantum case, the picture is similar [10] if we replace the action on n-bit strings with the action on 2^n-dimensional density matrices. The k^{th} moments of this action can be expressed [10, 21] in terms of the matrix

$$G_\mu = \int_{\mathbb{U}(2^n)} U^{\otimes k} \otimes (U^*)^{\otimes k} \mu(\mathrm{d}U). \tag{3}$$

where μ is a distribution over unitary transformations of n qubits. This matrix can also be thought of as the matrix form of the map that sends ρ to $\int U^{\otimes k} \rho (U^\dagger)^{\otimes k} \mu(\mathrm{d}U)$. If μ is taken to be Haar measure, we obtain an analogue of $A_{\infty,k}$. To obtain an analogue of $A_{1,k}$ one sets $\mu = \nu \equiv \nu_n$ with ν_n representing an average over $n-1$ choices of pairs of neighboring qubits and over a random choice of two-qubit gate applied to those qubits. As in the classical case, G_{Haar} is the projector onto the 1-eigenspace of G_ν, which we will argue below has dimension $k!$. Let $1-\delta$ denote the next largest eigenvalue of G_ν. Then we again have that

$$\|(G_\nu)^t - G_{\text{Haar}}\| = (1-\delta)^t, \tag{4}$$

so the length of the circuit ensuring $\|(G_\nu)^t - G_{\text{Haar}}\| \leq \epsilon$ is given by

$$t = \delta(n,k)^{-1} \log \frac{1}{\epsilon}, \tag{5}$$

where we have made explicit the dependence of δ on n, k. Now, our main result is the following estimate

$$\delta(n,k) \geq \Omega\left(\frac{1}{nk^{8.1}\log^2(k)}\right), \tag{6}$$

which implies that random circuit of length $t = O(nk^{8.1}\log^2(k)(nk\log(d) + \log(1/\epsilon)))$ approximates up to ϵ the k^{th} moment of random unitary, thus proving Theorem 1.

Connection to many-body theory. The matrix G_ν can be expressed in terms of a *local Hamiltonian*, bringing the problem within the scope of quantum many-body theory. The quantity $\delta(n,k)$ that determines circuit length in our case was shown in [7, 10, 28] to be directly related to the spectral gap of some Hamiltonian consisting of nearest-neighbor interactions between n $D = 4^k$-dimensional systems on a line. This Hamiltonian does not correspond to a physical system, but tools from many-body theory still apply and can help evaluate the gap.

To construct the Hamiltonian, note that our random circuit ν consists of picking with probability $\frac{1}{n-1}$ a random gate on two adjacent qubits. Thus $G_\nu = \frac{1}{n-1}\sum_{i=1}^{n-1} P_{i,i+1}$ where $P_{i,i+1} := \int_{\mathbb{U}(4)} (U_{i,i+1})^{\otimes k,k} \mu_{\text{Haar}}(\mathrm{d}U)$ and $U_{i,i+1}$ acts on the i-th and $(i+1)$-th qubit. Define

$$H_{n,k} := (n-1)(I - G_\nu) = \sum_{i=1}^{n-1} h_{i,i+1} \tag{7}$$

with local terms $h_{i,i+1} := I - P_{i,i+1}$ where I is the identity operator. It can then be shown that the "energy" of a ground state of the Hamiltonian is zero (corresponding to the 1-eigenspace of G_ν), while that of the first excited level is $(n-1)\delta(n,k)$. In other words, the spectral gap of $H_{n,k}$ (denoted $\Delta(H_{n,k})$) is directly related to the difference of the largest and second largest eigenvalue of G_ν according to $\delta(n,k) = \frac{\Delta(H_{n,k})}{n-1}$. Thus to determine the length of the random circuit, that approximates k-design, it sufficies to lower bound the spectral gap of $H_{n,k}$.

Bounding the Spectral Gap. Despite of decades of research on many-body systems, the evaluation of spectral gaps of local Hamiltonians is often a formidable task. Fortunately, our Hamiltonian has a nice feature of being *non-frustrated*: ground states of the total Hamiltonian are at the same time a ground states of its local constituents. For such Hamiltonians, Nachtergaele [33] provided a sufficient condition for existence of a constant gap (in the number of qubits n), together with an estimate for the gap. Nachtergaele's criterion is given in terms of ground subspaces of Hamiltonians consisting of various numbers of systems $m \leq n$. One finds that the ground space of our Hamiltonian is spanned by the $k!$ product vectors $|\psi_\sigma\rangle^{\otimes n}$, labeled by k-element permutations σ. This originates from the fact that the only operators which commute with $U^{\otimes k}$, where U is arbitrary unitary transformation, are linear combinations of operators that permute systems (the vectors are actually isomorphic to those operators). Were the $|\psi_\sigma\rangle^{\otimes n}$ strictly orthogonal, the Nachtergaele criterion would apply immediately, but this does not hold here. Yet, by use of group representation theory, we obtain that the vectors $|\psi_\sigma\rangle^{\otimes n}$ have sufficiently small overlaps (see the Supplementary Materials), to enable us to apply the criterion and obtain a tight gap.

As a result, we obtain that the spectral gap of Hamiltonian over n qubits for any n can be bounded by the gap of Hamiltonian of a *fixed* number of qubits, depending only on the moment k:

Lemma 3. *For every integers n, k with $n \geq \lceil 2.5\log(4k)\rceil$,*

$$\Delta(H_{n,k}) \geq \frac{\Delta(H_{\lceil 2.5\log(4k)\rceil,k})}{4\lceil 2.5\log(4k)\rceil}. \tag{8}$$

where $\lceil x \rceil$ denotes the least integer no smaller than x.

It remains only to control the gap of the Hamiltonian for systems with $O(\log(k))$ qubits. Here, since we are concerned with a relatively small number of qubits, it suffices to establish that random circuits on m qubits mix after a number of steps that is exponential in m. We prove this in [6] using the path-coupling method of [38]. If $m = O(\log(k))$, then this means a number of steps that is polynomial in k. Translating these mixing bounds back into a statement about the gap of H, we find that $\Delta(H_{O(\log k),k}) \geq 1/\operatorname{poly}(k)$. This completes the proof of Theorem 1.

Discussion. Our main result shows that $O(n^2 k^{10})$ random nearest-neighbor unitary interactions yield a distribution over unitaries that approximately matches the first k moments of the Haar measure. The dependence on n is approximately optimal for 1-d; indeed, we can think of it (for fixed k) as $O(n)$ rounds of parallel nearest-neighbor interactions. Since it takes at least this much time for information to propagate along a line of n qubits, this n-dependence cannot be improved. On the other hand, we did not try hard to reduce the degree of the polynomial in k and it may even be that the number of gates required is independent of k, as in [5]. Another open question is whether the n-dependence could be improved for better connected geometries; e.g. nearest-neighbor gates in higher-dimensional lattices or gates with arbitrary connectivity.

Theorem 1 shows that k-designs can be implemented efficiently on a quantum computer and also describes a plausible natural process that can give rise to them. Other work has studied open systems with nonunitary dynamics [8] and closed systems with Hamiltonian dynamics [34]. Our model differs from both of these, but could plausibly arise from a randomly time-varying Hamiltonian. Still, a major question left open is to understand the conditions under which realistic physical systems thermalize.

Acknowledgements. We would like to thank Dorit Aharonov, Itai Arad, Winton Brown, Daniel Jonathan, Bruno Nachtergaele, Alex Russell, Tomasz Szarek, and Andreas Winter. FGSLB acknowledges support from the Swiss National Science Foundation, via the National Centre of Competence in Research QSIT, and the Ministry of Education and the National Research Foundation, Singapore. AWH was funded by NSF grants CCF-1111382 and CCF-1452616 and ARO contract W911NF-12-1-0486. MH is supported by EU grant QESSENCE, by Polish Ministry of Science and Higher Education grant N N202 231937. MH also acknowledges grant QUASAR of the National Centre for Research and Development of Poland for preparing the final version of the paper. Part of this work was done in the National Quantum Information Center of Gdansk. We thank the hospitality of Institute Mittag Leffler within the program "Quantum Information Science," where (another) part of this work was done.

[1] A. Abeyesinghe, I. Devetak, P. Hayden, and A. Winter. The mother of all protocols: Restructuring quantum information's family tree. *Proc. Roc. Soc. A*, 465(2108):2537–2563, 2009, `arXiv:quant-ph/0606225`.

[2] N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(7):567, 1986.

[3] L. Arnaud and D. Braun. Efficiency of producing random unitary matrices with quantum circuits. *Phys. Rev. A*, 78:062329, Dec 2008, `arXiv:0807.0775`.

[4] M. C. Bañuls, J. I. Cirac, and M. B. Hastings. Strong and weak thermalization of infinite nonintegrable quantum systems. *Phys. Rev. Lett.*, 106:050405, Feb 2011, `arXiv:1007.3957`.

[5] J. Bourgain and A. Gamburd. A spectral gap theorem in $SU(d)$, 2011, `arXiv:1108.6264`.

[6] F. G. S. L. Brandão, A. W. Harrow, and M. Horodecki. Local random quantum circuits are approximate polynomial-designs, 2012, `arXiv:1208.0692`.

[7] F. G. Brandao and M. Horodecki. Exponential quantum speed-ups are generic. *Q. Inf. Comp.*, 13:0901, 2013, `arXiv:1010.3654`.

[8] H. Breuer and F. Petruccione. *The Theory of Open Quantum Systems*. Oxford University Press, 2002.

[9] A. Brodsky and S. Hoory. Simple permutations mix even better. *Random Struct. Algorithms*, 32:274–289, May 2008. arXiv:math/0411098.

[10] W. G. Brown and L. Viola. Convergence rates for arbitrary statistical moments of random quantum circuits. *Phys. Rev. Lett.*, 104:250501, Jun 2010, `arXiv:0910.0913`.

[11] P. Cwiklinski, M. Horodecki, M. Mozrzymas, L. Pankowski, and M. Studzinski. Efficient algorithm for multiqudit twirling for ensemble quantum computation. *J. Phys. A: Math. Theor.*, 46:305301, 2013. arXiv:1212.2556.

[12] P. Ćwikliński, M. Horodecki, M. Mozrzymas, Ł. Pankowski, and M. Studziński. Local random quantum circuits are approximate polynomial-designs: numerical results. *Journal of Physics A: Mathematical and Theoretical*, 46(30):305301, 2013, `arXiv:1212.2556`.

[13] O. C. O. Dahlsten, R. Oliveira, and M. B. Plenio. The emergence of typical entanglement in two-party random processes. *J. Phys. A: Math. Theor.*, 40(28):8081–8108, 2007, `arXiv:quant-ph/0701125`.

[14] B. M. T. David P. DiVincenzo, Debbie W. Leung. Quantum data hiding. *IEEE Trans. Inf. Theory*, 48(3):580, 2002.

[15] I. Diniz and D. Jonathan. Comment on "random quantum circuits are approximate 2-designs". *c*, 304:281–293, 2011, `arXiv:1006.4202`.

[16] J. Emerson, Y. S. Weinstein, M. Saraceno, S. Lloyd, and D. G. Cory. Pseudo-random unitary operators for quantum information processing. *Science*, 302(5653):2098–2100, 2003.

[17] J. Emerson, Y. S. Weinstein, M. Saraceno, S. Lloyd, and D. G. Cory. Pseudo-random unitary operators for quantum information processing. *Science*, 302(5653):2098–2100, 2003.

[18] O. Goldreich, S. Goldwasser, and S. Micali. How to con-

struct random functions. *Journal of the ACM*, 33(4):792, 1986.

[19] S. Goldstein, J. L. Lebowitz, R. Tumulka, and N. Zanghì. Canonical typicality. *Phys. Rev. Lett.*, 96:050403, Feb 2006, `arXiv:cond-mat/0511091`.

[20] F. Haake. *Quantum signatures of chaos*. Springer: Berlin, 1991.

[21] A. Harrow and R. Low. Efficient quantum tensor product expanders and k-designs. In *Proc. of APPROX-RANDOM*, volume 5687 of *LNCS*, pages 548–561. Springer, 2009, `arXiv:0811.2597`.

[22] A. Harrow and R. Low. Random quantum circuits are approximate 2-designs. *Commun. Math. Phys.*, 291:257–302, 2009, `arXiv:0802.1919`.

[23] A. W. Harrow. Approximate orthogonality of permutations, 2013. unpublished manuscript.

[24] P. Hayden, P. W. S. D. W. Leung, and A. J. Winter. Randomizing quantum states: Constructions and applications. *Commun. Math. Phys.*, 250:371–391, 2004, `arXiv:quant-ph/0307104`.

[25] P. Hayden, D. W. Leung, and A. Winter. Aspects of generic entanglement. *Commun. Math. Phys.*, 265:95, 2006, `arXiv:quant-ph/0407049`.

[26] P. Hayden and J. Preskill. Black holes as mirrors: quantum information in random subsystems. *JHEP*, 0709:120, 2007, `arXiv:0708.4025`.

[27] S. Hoory and A. Brodsky. Simple Permutations Mix Even Better, 2004, `arXiv:math.CO/0411098`.

[28] M. Žnidarič. Exact convergence times for generation of random bipartite entanglement. *Phys. Rev. A*, 78:032324, Sep 2008, `arXiv:0809.0554`.

[29] E. Knill. Approximation by quantum circuits, 1995, `arXiv:quant-ph/9508006`.

[30] N. Linden, S. Popescu, A. J. Short, and A. Winter. Quantum mechanical evolution towards thermal equilibrium. *Phys. Rev. E*, 79:061103, Jun 2009, `arXiv:0812.2385`.

[31] R. A. Low. Large deviation bounds for k-designs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 465(2111):3289–3308, 2009, `arXiv:0903.5236`.

[32] R. A. Low. *Pseudo-randomness and Learning in Quantum Computation*. PhD thesis, University of Bristol, 2010, `arXiv:1006.5227`.

[33] B. Nachtergaele. The spectral gap for some spin chains with discrete symmetry breaking. *Communications in Mathematical Physics*, 175:565–606, 1996, `arXiv:cond-mat/9410110`.

[34] R. Nandkishore and D. A. Huse. Many-body localization and thermalization in quantum statistical mechanics. *Annual Review of Condensed Matter Physics*, 6(1):15–38, 2015, `arXiv:1404.0686`.

[35] J. Neumann. Beweis des ergodensatzes und des h-theorems in der neuen mechanik. *Zeitschrift für Physik*, 57(1-2):30–70, 1929, `arXiv:1003.2133`.

[36] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press,Cambridge, 2000.

[37] 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`.

[38] R. I. Oliveira. On the convergence to equilibrium of Kac's random walk on matrices. *Ann. Appl. Probab.*, 19(3):1200–1231, 2009, `arXiv:0705.2253`.

[39] P. Sen. Random measurement bases, quantum state distinction and applications to the hidden subgroup problem. In *IEEE Conference on Computational Complexity*, pages 274–287, 2006, `arXiv:quant-ph/0512085`.

[40] C. E. Shannon. The synthesis of two-terminal switching circuits. *BSTJ*, 28(1):459–466, 1949.

[41] L. Susskind. Computational complexity and black hole horizons. *arXiv:1402.5674*, 2014.

[42] S. Trotzky, Y.-A. Chen, A. Flesch, I. P. McCulloch, U. Schollwöck, J. Eisert, and I. Bloch. Probing the relaxation towards equilibrium in an isolated strongly correlated one-dimensional bose gas. *Nature Physics*, 8(4):325–330, 2012, `arXiv:1101.2659`.

[43] $\Omega(f(n))$ refers to a function that is $\geq cf(n)$ for some $c > 0$ and for sufficiently large n.