Anticoncentration theorems for schemes showing a quantum speedup

D. Hangleiter, J. Bermejo-Vega, M. Schwarz, and J. Eisert

Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany
May 16, 2018

One of the main milestones in quantum information science is to realise quantum devices that exhibit an exponential computational advantage over classical ones without being universal quantum computers, a state of affairs dubbed quantum speedup, or sometimes "quantum computational supremacy". The known schemes heavily rely on mathematical assumptions that are plausible but unproven, prominently results on anticoncentration of random prescriptions. In this work, we aim at closing the gap by proving two anticoncentration theorems and accompanying hardness results, one for circuit-based schemes, the other for quantum quench-type schemes for quantum simulations. Compared to the few other known such results, these results give rise to a number of comparably simple, physically meaningful and resource-economical schemes showing a quantum speedup in one and two spatial dimensions. At the heart of the analysis are tools of unitary designs and random circuits that allow us to conclude that universal random circuits anticoncentrate as well as an embedding of known circuit-based schemes in a 2D translation-invariant architecture.

1   Introduction

Realising a quantum device that computationally outperforms state-of-the art classical supercomputers for a certain task that is provably intractable classically has become a key milestone in the field of quantum simulation and computing. This goal is often referred to as "quantum (computational) supremacy" [1] or quantum speedup. Such a quantum speedup is not merely meant in the sense of quantum dynamics being no longer tractable on classical supercomputers using the best known algorithms to date, for which there is evidence already today [2–4]. Instead, to make sure that the inefficient classical simulation is not victim of

D. Hangleiter: dominik.hangleiter@fu-berlin.de

J. Bermejo-Vega: jbermejovega@gmail.com

M. Schwarz: martin.schwarz@gmail.com

J. Eisert: jense@zedat.fu-berlin.de

a lack of imagination, such a quantum speedup is usually meant to refer to schemes for which the speedup can be related to a notion of computational complexity. For a quantum speedup scheme to be physically realisable *in principle* in the absence of quantum error correction, it is crucial that the hardness of the task is robust under physically realistic errors. To have any hope of realising such a scheme in the near term one would moreover wish for the resources required for an implementation of the architecture in the intractable regime to be achievable with present-day (or near-term) technology.

There are only very few quantum speedup architectures that are robust against physically realistic constant total-variation distance errors [5–12]. Even fewer of those are physically realistic when it comes to an implementation in present-day technology in that they require only nearest-neighbour interactions and are feasible in the available experimental platforms such as linear optics [5], superconducting qubits [7, 8], ion traps or cold atoms in optical lattices [10]. The computational task that is solved in all of these proposals is a sampling task, in particular, the task of sampling from the output distribution of a certain random time-evolution. That random time evolution may take the form of a Haar-random unitary applied to a bosonic state [5], a random circuit from a gate set [8], IQP circuits [7] applied to an all-zero state, or even a translation-invariant nearest-neighbour Ising Hamiltonian that is applied to a random product state [10]. In addition to this discussion, there is the question to what extent schemes showing a quantum speedup can be certified in their correctness [9, 10, 12–16].

The central ingredient of all existing quantum speedup proofs is Stockmeyer's algorithm [17] that implies a collapse of the Polynomial Hierarchy if sampling from the output distribution of the respective circuits is #P-hard on average. In order for this hardness argument to be valid, one crucially requires so-called anticoncentration bounds for the output probability distribution of the respective random circuits [18].

Indeed, it has been shown that one can efficiently classically sample from output distributions of certain circuit families, including IQP circuits, the output distribution of which concentrates on a polynomially small subset of the sample space [19]. This shows

that concentrated output distributions are in many relevant cases simulable, rendering anticoncentration a necessary condition for classical hardness for these cases.

Despite of their central role in the hardness argument of quantum speedup proposals, only few proofs of anticoncentration bounds are known so far [6, 7, 11]. In all other speedup architectures – boson sampling [5], universal random circuits [8], and translation-invariant Ising models [9, 10] – there exists none or merely numerical evidence for anticoncentration of the respective circuit families and the validity of the anticoncentration assumption needs to be conjectured. This still gives rise to plausible schemes, but in order to complete the program of realising quantum schemes showing a quantum speedup, these gaps must necessarily be closed. Rigorous anticoncentration results for classically intractable circuit families are therefore both of crucial importance to corroborate the validity of those existing speedup proposals, as well as to shine light on the conditions of them coming about which are highly debated in the literature.

In this work, we provide rigorous anticoncentration results for two types of such quantum architectures that are at the same time not classically simulable. First, we show that random circuits drawn from a unitary 2-design anticoncentrate. We then apply this result to both show that random circuits comprised of nearest-neighbour gates that are drawn from a universal gate set containing inverses anticoncentrate in linear depth, and propose two new schemes based on this insight. Second, we prove that the output distribution of a particular nearest-neighbour quantum quench architecture based on the time evolution of product states under certain translation-invariant Ising models anticoncentrate in constant depth. Thus, we consider two types of architectures tailored towards different kinds of experimental platforms in the following sense. In platforms in which achieving large numbers of qubits is expensive and local control feasible, circuit-based schemes such as universal random circuits can be reasonably implemented. A paradigmatic example of such a platform might be constituted of superconducting qubits. In contrast, there are physically most natural settings of quantum simulators in which local control on the level of individual gates is difficult to achieve, but for which extremely large numbers of local constituents can be reached. Cold atoms in optical lattices, in which $10^4 - 10^5$ atoms are readily reachable, provide the most prominent example of such an architecture. Our quench-type architecture is tailored toward such settings.

The application of our first result to universal random circuits complies with the intuition that due to the ballistic spread of correlations anticoncentration will generically arise in depth that scales linearly with the diameter of the system under consideration, and

hence, linearly in a one-dimensional architecture [8]. Still, to the best of our knowledge there is no rigorous proof for anticoncentration of universal random circuits. In contrast, in the light of this intuition the second result is quite surprising: It has even been argued [18] that it cannot be expected to reach anticoncentrating output distributions in constant depth, retaining its classical intractability. We conjecture the scaling of both results to be optimal in the settings considered (unstructured circuits in one dimension and highly structured circuits in two dimensions).

To prove the first result (Theorem 5) we make use of properties of approximate unitary 2-designs and apply the Paley-Zygmund inequality. In applying this result to show that several schemes anticoncentrate we use the fact that all these schemes form approximate 2-designs. This includes, universal random circuits as in [8], Clifford circuits acting on input product magic states [20], and certain models of diagonal unitaries [21]. What is more, we provide new complexity-theoretic evidence for the hardness of classically simulating these random-circuit families in the approximate sampling sense. In doing so, we focus on universal random circuits, which attain this property already in *linear* $O(n)$ depth [22, 23]. For this case, we derive a matching classical-hardness upper bound by proving that a broad class of 1D random quantum circuits are hard to simulate classically already in linear $O(n)$ circuit depth, in the strong simulation sense (Lemma 8). The novel aspects of our work are the generality of these results and the minimal linear-depth requirements for achieving both classical hardness and anticoncentration on a 1D nearest-neighbour architecture. This drastically improves over prior work [8] on universal random circuits, which gave only numerical evidence for anticoncentration (in depth $O(\sqrt{n})$, and in 2D), and no matching classical-hardness depth bound. Prior to us, analogous depth results were only available for nearest-neighbour IQP circuit models supplemented by SWAP gates: for these, Ref. [7] gave $O(\sqrt{n}\log n)$ depth bounds for anticoncentration and hardness in 2D architectures, and Ref. [10] proved linear $O(n)$ ones in 1D layouts.

For the second result (Theorem 9) and Corollaries 10–12 thereof, we use the facts that IQP circuits anticoncentrate [7] and can be implemented in a 1D architecture in linear depth [10]. Our technical contribution is to provide an embedding of such IQP circuits in the constant-time evolution of a product state under a translation-invariant Ising model on a two-dimensional lattice. Via this embedding we are thus able to show the first anticoncentration bound for translation-invariant constant-depth schemes that exhibit a quantum speedup [9, 10].

This work is structured as follows: First, in Sec. 2, we will introduce the formal statement of anticoncentration and show how it is used in the Stock-

meyer hardness proof. In Sec. 3 we will both state and prove the anticoncentration result for approximate unitary 2-designs and then apply this result to relevant examples, most importantly, universal random circuits in linear depth. In Sec. 4 we will then prove the anticoncentration result for the constant-time evolution of a random product state under a certain nearest-neighbour translation-invariant Ising Hamiltonian. Finally, we discuss the implications of our results in the context of the timely literature in Sec. 5, and conclude in Sec. 6.

2 Preliminaries: Anticoncentration and quantum speedups

Throughout this work, we consider quantum systems consisting of n qubits (with obvious generalisation to d-dimensional local constituents). To start with, let us make precise, what is meant by anticoncentration of the output distribution of a unitary U drawn from a certain measure μ. We call the distribution of probabilities $|\langle x|U|0\rangle|^2$ of obtaining $x \in \{0,1\}^n$ when applying a unitary $U \in U(N)$, $N = 2^n$, to an initial state vector $|0\rangle := |0\rangle^{\otimes n}$ and measuring in the computational basis, the output distribution. We say that this output distribution anticoncentrates if there exist universal constants $\alpha, \beta > 0$ such that for any $x \in \{0,1\}^n$, the probabilities $|\langle x|U|0\rangle|^2$ of this unitary anticoncentrate,

$$\Pr_{U \sim \mu}\left(|\langle x|U|0\rangle|^2 \geq \frac{\alpha}{N} \right) > \beta. \tag{1}$$

We can interpret this probability as the probability that an arbitrarily chosen entry x of the first column of a μ-randomly chosen U is larger than α/N. Throughout this work, we say that a quantity X is approximated by a quantity $\tilde{X}$ with *multiplicative error* c if $X/c \leq \tilde{X} \leq cX$, with *relative error* r if $(1-r)X \leq \tilde{X} \leq (1+r)X$, and with *additive error* a if $\|X - \tilde{X}\|_* \leq a$ for some norm $\|\cdot\|_*$.

In the commonly used proof technique for quantum speedups [24, 25] Stockmeyer's algorithm [17] is applied to show a collapse of the polynomial hierarchy if for an arbitrary such x the amplitude $|\langle x|U|0\rangle|^2$ is #P-hard to approximate multiplicatively. Anticoncentration comes into this proof when hardness is shown not up to multiplicative but up to an additive error in total-variation distance. More specifically, to prove a quantum speedup with constant total-variation distance errors for sampling from the output distribution of a circuit family $\mathcal{F}$ using the argument developed in Refs. [5, 6] one requires three ingredients: (i) The output distribution of $\mathcal{F}$ anticoncentrates in the sense of Eq. (3). (ii) post$\mathcal{F}$ = postBQP. By the result of Refs. [26, 27] the output probabilities are then #P-hard to approximate up to relative error $1/4$. (iii) The output probabilities of $\mathcal{F}$ are #P-hard to approximate up to

multiplicative errors in the *average case*. This needs to be conjectured for all quantum speedup schemes[1], preferably in terms of a universal quantity such as the imaginary-time partition function of Ising models [6], the permanent [5], or the Jones polynomial [29]. Together, (ii) and (iii) permit a reduction from hardness of strong simulation up to multiplicative error to hardness of weak simulation up to an additive error using Stockmeyer's algorithm [17] in the third level of the Polynomial Hierarchy. As a result, very often three conjectures need to be made when proving a quantum speedup using this technique:

C1 The Polynomial Hierarchy cannot collapse to its 3rd level [30–32].

C2 If it is #P-hard to approximate the output probability of a circuit drawn from $\mathcal{F}$ up to a constant relative error, then the same problem is #P-hard for a constant fraction of the instances[2].

C3 The output distribution of $\mathcal{F}$ anticoncentrates.

In the following, we will prove the anticoncentration conjecture C3 in the sense of equation (1) both for certain circuit-based schemes, in fact, those that form an approximate 2-design, and a quantum-quench architecture in the mindset of Ref. [10].

3 Anticoncentration of circuit-based schemes

In this section we will begin by introducing and proving our first result, namely, that approximate unitary 2-designs anticoncentrate, and then apply this result to three relevant examples of circuit-based schemes, most prominently, universal random circuits.

3.1 Anticoncentration of unitary 2-designs

Unitary k-designs approximate the uniform (Haar) measure on the unitary group (see App. A.1) in the sense that the first k moments of a unitary k-design and the Haar measure match (exactly or approximately). The definition of a k-design is motivated by the fact that in experiments samples from a unitary k-design are much easier to realise than samples from the full Haar measure. In order to define the notion of a k design, we need the notion of the k^{th}-moment operator that acts as a unitary twirl with respect to some measure μ on the unitary group maps on an operator.

---

[1]In the exact case, one can even prove average-case results for the permanent [5] and random-circuit based schemes [28].

[2]This conjecture may be regarded as a qubit analogue of the "permanent-of-Gaussians" conjecture of Ref. [5].

Definition 1 (k^{th}-moment operator). *Let M_μ^k be the k-th moment operator on $\mathcal{L}(\mathcal{H}^{\otimes k})$ with respect to a distribution μ on $U(N)$, $N = 2^n = \dim \mathcal{H}$ defined as*

$$X \mapsto M_\mu^k(X) := \mathbb{E}_\mu \left[U^{\otimes k} X (U^\dagger)^{\otimes k} \right]$$
$$= \int_{U(N)} U^{\otimes k} X (U^\dagger)^{\otimes k} \mu(U). \quad (2)$$

We can now define unitary k-designs [22, 33].

Definition 2 (Unitary k-design). *Let μ be a distribution on the unitary group $U(N)$. Then μ is an exact unitary k-design if*

$$M_\mu^k = M_{\mu_{\text{Haar}}}^k .$$

In all of what follows, we will need to relax this notion to the notion of an *approximate* unitary k-design. In such a definition we can allow for both relative and additive errors on the equality (2) [23, 34]:

Definition 3 (Approximate unitary k-designs). *Let μ be a distribution on the unitary group $U(N)$. Then μ is*

1. an additive ϵ-approximate unitary k-design if

$$\|M_\mu^k - M_{\mu_{\text{Haar}}}^k\|_\diamond \leq \epsilon ,$$

2. a relative ϵ-approximate unitary k-design if

$$(1 - \epsilon)M_{\mu_{\text{Haar}}}^k \leq M_\mu^k \leq (1 + \epsilon)M_{\mu_{\text{Haar}}}^k .$$

Since the former definition is much more common in the literature, let us remark that the two definitions are closely related via the following Lemma of Ref. [23] in which, however, a factor of the dimension enters.

Lemma 4 (Additive and relative approximate designs). *If μ is a relative ϵ-approximate unitary k-design then $\|M_\mu^k - M_{\mu_{\text{Haar}}}^k\|_\diamond \leq 2\epsilon$. Conversely, if $\|M_\mu^k - M_{\mu_{\text{Haar}}}^k\|_\diamond \leq \epsilon$, then μ is a relative ϵN^{2k}-approximate unitary k-design.*

We are now ready to state our first anticoncentration result on unitary 2-designs.

Theorem 5 (Anticoncentration of unitary 2-designs). *Let μ be a relative ϵ-approximate unitary 2-design on the group $U(N)$. Then the output probabilities $|\langle x|U|0\rangle|^2$ for $x \in \{0,1\}^n$ of a μ-random unitary $U \in U(N)$ anticoncentrate in the sense that for $0 \leq \alpha \leq 1$*

$$\mathbb{P}_{U \sim \mu} \left(|\langle x|U|0\rangle|^2 > \frac{\alpha(1-\epsilon)}{N} \right) \geq \frac{(1-\alpha)^2(1-\epsilon)^2}{2(1+\epsilon)} . \quad (3)$$

We point out that Theorem 5 also holds in exactly the same way for relative ϵ-approximate state 2-designs. This is a weaker condition than the unitary design condition since any (approximate) unitary 2-design generates an (approximate) state 2-design via

application to an arbitrary reference state. Also note that the fact that μ is a relative ϵ-approximate 1-design (cf. App. A.3) is crucial for the bound (3) to become non-trivial. If instead μ was an additive design the lower bound would asymptotically tend to zero as $1/N$ and hence not stay larger than a constant. However, the 1-design condition holds even exactly for many distributions μ, although for the higher moments it may only hold approximately.

Proof of Theorem 5. Our proof of the anticoncentration bound (3) has two steps and relies on two ingredients: In the first step, we prove anticoncentration of a single *but fixed* entry of Haar random unitaries. To this end we make use of the *Paley-Zygmund inequality* and an explicit expression of the distribution of matrix elements of Haar-random unitaries. In the second step, we extend this result to full anticoncentration of all output probabilities in the sense of equation (3).

The Paley-Zygmund inequality is a lower-bound analogue of Markov-type tail bounds and can be stated as follows. If $Z \geq 0$ is a random variable with finite variance, and if $0 \leq \alpha \leq 1$

$$\mathbb{P}(Z > \alpha \mathbb{E}[Z]) \geq (1-\alpha)^2 \frac{\mathbb{E}[Z]^2}{\mathbb{E}[Z^2]} . \quad (4)$$

That is, it lower bounds the probability that a positive random variable is small in terms of its mean and variance.

Now let μ be a relative ϵ-approximate unitary 2-design. Then for $l = 2, 4$, it holds that

$$(1-\epsilon)\mathbb{E}_{U \sim \text{Haar}} \left[|\langle a|U|b\rangle|^l \right] \leq \mathbb{E}_{U \sim \mu} \left[|\langle a|U|b\rangle|^l \right]$$
$$\leq (1+\epsilon)\mathbb{E}_{U \sim \text{Haar}} \left[|\langle a|U|b\rangle|^l \right] . \quad (5)$$

This is due to the fact that for any unitary k-design μ the expectation value of an arbitrary polynomial P of degree 2 in the matrix elements of both U and $U^\dagger$ over μ equals the same expectation value but taken over the Haar measure up to a relative error $\epsilon > 0$ [35]. To see this, observe that averaging a monomial in the matrix elements of U over the k-design μ can be expressed as $\langle i_1, \ldots, i_k | M_\mu^k(|j_1, \ldots, j_k\rangle\langle j_1', \ldots, j_k'|)|i_1', \ldots, i_k'\rangle$. Hence, if $M_\mu^k = M_{\text{Haar}}^k$, "then any polynomial of degree k in the matrix elements of U will have the same expectation over both distributions" [35]. This gives rise to

$$\mathbb{P}_{U \sim \mu}(|\langle x|U|0\rangle|^2 > \alpha(1-\epsilon)\mathbb{E}_{U \sim \text{Haar}}[|\langle x|U|0\rangle|^2])$$
$$\geq \mathbb{P}_{U \sim \mu}(|\langle x|U|0\rangle|^2 > \alpha \mathbb{E}_{U \sim \mu}[|\langle x|U|0\rangle|^2])$$
$$\geq (1-\alpha)^2 \frac{\mathbb{E}_{U \sim \mu}[|\langle x|U|0\rangle|^2]^2}{\mathbb{E}_{U \sim \mu}[|\langle x|U|0\rangle|^4]} \quad (6)$$
$$\geq (1-\alpha)^2 \frac{(1-\epsilon)^2 \mathbb{E}_{U \sim \text{Haar}}[|\langle x|U|0\rangle|^2]^2}{(1+\epsilon)\mathbb{E}_{U \sim \text{Haar}}[|\langle x|U|0\rangle|^4]} .$$

Lemma 6 (Marginal output distribution). *The distribution of the marginal output probabilities $p = $*

$|\langle x|U|0\rangle|^2$ of Haar random unitaries U and arbitrary but fixed x is given by

$$P_{\text{Haar}}(p) = (N-1)(1-p)^{N-2} \xrightarrow{N\gg 1} N\exp(-Np). \tag{7}$$

In particular, P_{Haar}'s first and second moments are given by

$$\mathbb{E}_{\text{Haar}}[p] = \frac{1}{N}, \quad \mathbb{E}_{\text{Haar}}[p^2] = \frac{2}{N(N+1)}. \tag{8}$$

We prove this lemma in App. B. Inserting the expressions Eqs. (8) for $\mathbb{E}_{U\sim\text{Haar}}[|\langle x|U|0\rangle|^2]$ and $\mathbb{E}_{U\sim\text{Haar}}[|\langle x|U|0\rangle|^4]$, we find

$$P_{U\sim\mu}\left(|\langle x|U|0\rangle|^2 > \frac{\alpha(1-\epsilon)}{N}\right)$$
$$\geq (1-\alpha)^2 \frac{N(N+1)}{2N^2}\frac{(1-\epsilon)^2}{(1+\epsilon)} \geq (1-\alpha)^2\frac{(1-\epsilon)^2}{2(1+\epsilon)},$$

which completes the proof. $\qquad\square$

Note that the moments (8) can alternatively be obtained for both state and unitary 2-designs exploiting Schur-Weyl duality. This yields an explicit expression of the k^{th}-moment operators $M_\mu^k(X)$ as the projector onto the span of the symmetric group on k tensor copies of the Hilbert space $\mathcal{H}$. Moreover, the moments of the output probabilities of state 2-designs are also given by (8). This can be seen similarly using the fact that the expectation value over a state k-design is given by the projection onto the k-partite symmetric subspace of $\mathcal{H}^{\otimes k}$ [36]. We note that the output distribution (7) of a Haar random unitary asymptotically approaches the exponential (Porter-Thomas) distribution. This behaviour has already been observed numerically in many different contexts involving pseudo-random operators [22, 37], non-adaptive measurement-based quantum computation [38], and universal random circuits [8].

3.2 Applications: a "recipe" for quantum speedups

Our anticoncentration theorem 5 for approximate unitary (and state) 2-designs leads to a generic "recipe" for the identification of quantum circuit families and input states that are hard to simulate classically under plausible complexity-theoretic conjectures, building upon the approach of Refs. [5, 6]. The strategy goes in three steps parallel to ingredients (i-iii) of the proof based on Stockmeyer's algorithm introduced in Sec. 2. In the following, we apply this general strategy to a few examples of random circuits, most prominently, universal random circuits.

Universal random circuits. The first example that we also focus on are random quantum circuits constructed from single- and two-qubit
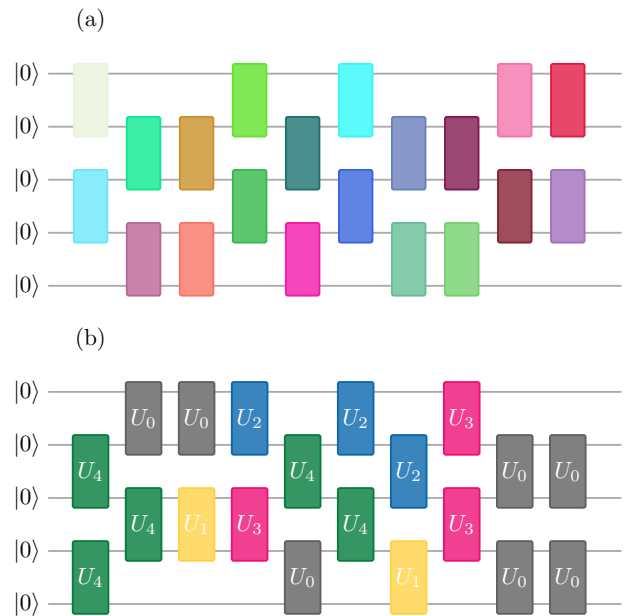

Figure 1: Layout of the parallel random circuit families. In each step either the even or odd configuration of parallel two-qubit unitaries is applied with probability $1/2$. Every two-qubit gate is chosen from the respective measure on $U(4)$ – (a) the Haar measure, (b) the uniform distribution on the gate set G. Here we depict a five-qubit random instance of depth 10 where in (a) the colour choice represents different gates, and in (b) the gate set consists of 5 two-qubit unitaries $G = \{U_0, U_2, \ldots, U_4\}$.

gates, most prominently, the gate set $\mathcal{G}_{\text{BIS}} = \{CZ, H, \sqrt{X}, \sqrt{Y}, T\}$ studied by Boixo *et al.* [8, 39]. In presenting this example we put particular emphasis on the *circuit depth* required to reach a scheme that shows a provable quantum speedup. As a first step (i) of the general strategy, the following Corollary establishes that the output distribution of a random circuit formed in a particular fashion from $\mathcal{G}_{\text{BIS}}$ anticoncentrates. This holds already in *linear depth*.

Corollary 7 (Universal random circuits anticoncentrate). *The output probabilities of universal random circuits in one dimension from the following two circuit families (illustrated in Fig. 1) anticoncentrate in a depth that scales as $O(n\log(1/\epsilon))$ in the sense of Eq. (3).*

- Parallel local random circuits*: In each step either the unitary $U_{1,2} \otimes U_{3,4} \otimes \cdots \otimes U_{n-1,n}$ or the unitary $U_{2,3} \otimes U_{4,5} \otimes \cdots \otimes U_{n-2,n-1}$ is applied (each with probability $1/2$), with $U_{j,j+1}$ independent unitaries drawn from the Haar measure on $U(4)$. (This assumes n is even.)*

- Universal gate sets*: Let $G \coloneqq \{g_i\}_{i=1}^m$ with each $g_i \in U(4)$ be a universal gate set containing inverses with elements composed of algebraic identities, i.e., a gate set G such that the group generated by G is dense in $U(4)$ and satisfying*

$g_i \in G \Rightarrow g_i^{-1} \in G$. *In each step either the unitary* $U_{1,2} \otimes U_{3,4} \otimes \cdots \otimes U_{n-1,n}$ *or the unitary* $U_{2,3} \otimes U_{4,5} \otimes \cdots \otimes U_{n-2,n-1}$ *is applied (each with probability* $1/2$*), with* $U_{j,j+1}$ *independent unitaries drawn uniformly from* G.

Proof of Corollary 7. The central ingredient of our proof of Corollary 7 is the result of Ref. [23]. There, the authors show that the two random circuit families are relative ϵ-approximate unitary k-designs on $U(2^n)$ in depth $\mathrm{poly}(k) \cdot O(n \log(1/\epsilon))$ (Corollary 6 and 7 in Ref. [23]).

Hence, in particular, these random circuits are relative ϵ-approximate unitary 2-designs in depth $O(n \log(1/\epsilon))$, i.e., linear in the number of qubits and logarithmic in $1/\epsilon$. Applying Theorem 5 to the output probabilities $|\langle x|\mathcal{C}|0\rangle|^2$ of a random circuit $\mathcal{C}$ applied to an initial all-zero state yields the claimed anticoncentration bound for the output probabilities of such circuits. $\qquad\square$

To prove a quantum speedup using the Stockmeyer technique the second required ingredient is #P-hardness of strong classical simulation of the output probabilities (ii). Indeed, since the gate set $\mathcal{G}_{\mathrm{BIS}}$ is universal, the post$\mathcal{F}$ =postBQP connection is immediate. Boixo *et al.* [8] moreover showed that the output probabilities can be expressed in terms of the imaginary-time partition function of a random Ising model, suggesting that the average-case conjecture for random circuits is a natural one (iii). It remains to be shown that random universal circuits are both not classically strongly simulable and anticoncentrate in *linear depth* in a one-dimensional setting. Lemma 8 (below) establishes this is indeed the case for a large class of finite gate sets with efficiently-computable matrix entries (so that they cannot artificially encode solutions to hard problems). It is an open question whether this can be improved to square-root depth in a two-dimensional setting such as that of Refs. [8, 39].

Given two $O(1)$-local gate sets A and B, we say that A *exactly synthesises* B if every gate $V \in B$ can be exactly implemented via a polynomial-time computable constant-size circuit of gates in A.

Lemma 8 (Hardness of strong classical simulation)**.** *Let* G *be any finite universal gate set with algebraic efficiently computable matrix entries that can exactly synthesise either the* $\{e^{i\frac{\pi}{8}X}, e^{i\frac{\pi}{4}X \otimes X}, \mathrm{SWAP}\}$ *or* $\{e^{i\frac{\pi}{8}X \otimes X}, \mathrm{SWAP}\}$. *Then, approximating the output probabilities of* $O(n)$*-depth circuits of* G *nearest-neighbour gates in one dimension up to relative error* $1/4$ *is #P-hard.*

Let us highlight that Lemma 8 applies to many well-studied universal gate sets, including $\mathcal{G}_{\mathrm{BIS}}$, the ubiquitous Clifford+T [40], Hadamard+controlled-$\sqrt{Z}$ [41], Hadamard+Toffoli [42, 43] and others [44–46]. Interestingly, Lemma 8 holds also for non-universal gate sets, though the latter may not always anticoncentrate. We now prove Lemma 8.

Proof of Lemma 8. We begin by showing that both given target gate sets can exactly implement subgroups of the 2-qubit dense IQP circuits of Ref. [6]. Specifically, the first gate set gives us the group $\mathcal{G}_1$ generated by $\exp\left(i\frac{\pi}{8}X_i\right)$, $\exp\left(i\frac{\pi}{4}X_iX_j\right)$ gates acting on a complete graph, while the second gives the group generated by arbitrary long range $\exp\left(i\frac{\pi}{8}X_iX_j\right)$ gates. In both cases, long range interactions are obtained via the available SWAPs.

Next, we show that, like the circuits in Ref. [6], both $\mathcal{G}_1$ and $\mathcal{G}_2$ are universal under post-selection. Indeed, both can adaptively implement a single-qubit Hadamard via gate teleportation [47] (see also [24, 48]), and non-adaptively, if we can post-select. The claim follows from the universality of known gate sets [40, 41].

Last, due to Refs. [27, 49], the output probabilities of post-selected universal quantum circuits are #P-hard to approximate up to multiplicative error $\sqrt{2}$ (relative error $1/4$). The previous fact implies that this holds for the dense IQP circuits in $\mathcal{G}_1$ and $\mathcal{G}_2$. Furthermore, n-qubit dense IQP circuit can be exactly implemented in $O(n)$ depth on a 1D nearest-neighbour architecture using SWAP gates [10, Lemma 6]. It follows that the output probabilities of linear-depth circuits in $\mathcal{G}_1$ or $\mathcal{G}_2$ are #P-hard to approximate. This readily extends to any circuit family that can exactly synthesise either of the former, since this process only introduces a constant depth overhead. $\qquad\square$

We do not know whether Lemma 8 extends to arbitrary gate sets since applying some Solovay-Kitaev type gate synthesis algorithms [41, 50, 51] should introduce a polynomial overhead factor in depth. This is because due to Chernoff-Hoeffding's bound #P-hard-to-approximate quantum probabilities need to be (at least) super-polynomially small, for otherwise they could be inferred in quantum polynomial time by mere sampling, which is not believed possible [52, 53]. To approximate such small probabilities via the Solovay-Kitaev algorithm requires $\Omega(n^\alpha)$ overhead for some $\alpha > 0$ assuming the counting exponential time hypothesis [54]. These issues are closely related to the open question of whether or not the power of post-selected quantum circuits is gate set independent given some $\tilde{O}(n^\alpha)$ depth bound [26].

Commuting circuits. As a second example, we consider circuits of diagonal unitaries composed of controlled-phase type one- and two-qubit gates of the form $\mathrm{diag}(1, 1, 1, e^{i\phi})$, and an input state $|+\rangle^{\otimes n}$. By the result of Ref. [21] this gate set yields a state 2-design if the phases are picked from discrete sets ($\{0, \pi\}$ for the two-qubit gates, and $\{0, 2\pi/3, 4\pi/3\}$ for the single qubit gates), and thus satisfies anticoncentration in the sense of Eq. 3 (i). Adding the S-gate to the gate set and measuring all qubits in the X-basis we obtain post$\mathcal{F}$ = postBQP by Refs. [6, 27]

Circuit family $\mathcal{F}$	Input state $|\psi_0\rangle$	(State) 2-design property	Worst-case hardness (post$\mathcal{F}$ = postBQP)	Average-case conjecture in terms of universal quantity
$\mathcal{G}_{\mathrm{BIS}}$	$|0\rangle^{\otimes n}$	[23]	[8, 40]	Ising part. func./Jones polyn.
Diagonal unitaries	$|+\rangle^{\otimes n}$	[21]	[26, 27]	Ising partition function
Clifford circuits	$(T|0\rangle)^{\otimes n}$	[33]	[20]	Ising partition function

Table 1: Examples of random circuit families that exhibit a provable quantum speedup up to total-variation distance errors.

as IQP circuits are an instance of diagonal unitaries (ii). Here, we have used the fact that adding the S-gate and post-selection gives us access to the universal gate set Clifford + $\pi/12$ [55, 56]. Again, the average-case conjecture can be phrased in terms of an Ising partition function (iii). Last, the circuits can be implemented in linear depth if either long-range interactions or nearest-neighbour SWAPs are allowed [10].

Clifford circuits with product-state inputs. A similar argument can be applied to Clifford circuits which are known to be an exact 2-design [33, 57] applied to magic input states. By the result of Ref. [58] an arbitrary element of the Clifford group in 2^n dimensions can be decomposed into $O(n^3)$ elementary Clifford gates. The result of Ref. [57] even achieves an exact 2-design using only quasi-linearly many one- and two-qubit Clifford gates. The post$\mathcal{F}$ =postBQP for this case is due to Ref. [20]. We summarise these examples in Table 1.

4 Anticoncentration of quenched many-body dynamics

In this section, we investigate anticoncentration of a particular type of quantum simulation scheme exhibiting a quantum speedup based on the architecture recently introduced in Ref. [10] (specifically, architectures I-II therein). These implement quenched (constant-time) dynamical evolutions [59, 60] under many-body Ising Hamiltonians on the square lattice whose graph we denote by $\mathcal{L} = (V, E)$. More specifically, we consider a particular variant of such a setting and prove both anticoncentration for its output distribution and the hardness of strongly classically simulating it. In virtue of the Stockmeyer-type argument presented in Sec. 2 this gives rise to a new quantum speedup result for this architecture.

4.1 A new quantum quench architecture

Let us start by reviewing the idea of the quench architectures introduced in Ref. [10]. There, the computation in the circuit model amounts to, first, preparing a product state vector $|\psi_\beta\rangle = \bigotimes_{i \in V}(|0\rangle + \mathrm{e}^{\mathrm{i}\beta_i}|1\rangle)/\sqrt{2}$ with β_i chosen randomly from a finite set of angles; second, implementing a constant-time evolution

$U := \mathrm{e}^{-\mathrm{i}H}$ under a nearest-neighbour translation-invariant Ising Hamiltonian

$$H := \sum_{(i,j) \in E} J_{i,j} Z_i Z_j - \sum_{i \in V} h_i Z_i, \qquad (9)$$

and, third, measuring all qubits in the X basis. Ref. [10] proved that quantum simulations of this form cannot be efficiently classically sampled from up to constant total-variation distance assuming variants of the average-case and anticoncentration conjecture. As supporting evidence for the anticoncentration conjecture, Ref. [10] built a link to the anticoncentration of certain families of universal random circuits and provided numerical data.

We note that these architectures are closely related to that of Gao *et al.* [9]. The similarities and differences are spelled out in Ref. [10].

We now introduce a new variant of such a quantum quench architecture, named $\mathcal{Q}_{\mathrm{ac}}$ and illustrated in Fig. 2 that produces provably hard-to-approximate anticoncentrated distributions which cannot be classically sampled from if an average-case conjecture holds and the Polynomial Hierarchy does not collapse. The architecture picks a uniformly-random input product state from a finite family $\mathcal{S}_{\mathrm{ac}} = \{|\psi_\beta\rangle\}_\beta$ and lets it evolve under a nearest-neighbour translation-invariant Hamiltonian H_{ac} (defined below).

Specifically, the architecture uses $n(m) := m(2m + 1)$ qubits, arranged in an m-row $(2m + 1)$-column square lattice. Boundary qubits on even-rows are initialised on $|0\rangle$ or $|1\rangle$ uniformly at random. The remaining ones are divided in two groups, named "blue" and "yellow", using a lattice 2-colouring that places no blue qubit on the top-left and top-right columns. Blue qubits are initialised on $(|0\rangle + |1\rangle)/\sqrt{2}$; yellow ones on $\mathrm{e}^{-\mathrm{i}k_i\pi Z_i/8}|+\rangle$ with uniformly-random $k_i \in \{0, 1, 2, 3\}$.

Next, the prepared state evolves under a translation-invariant Ising Hamiltonian H_{ac}. Letting $[i, j]$ denote the qubit on the i-th row and j-th column lattice (in left-to-right top-to-bottom order), the latter reads

$$H_{\mathrm{ac}} = \sum_{\substack{i \leq k, \, j \leq l \\ ([i,j],[k,l]) \in E}} \tfrac{\pi}{4} \delta_{i,j}^{k,l} Z_{[i,j]} Z_{[k,l]} - \sum_{v \in V} \tfrac{\pi}{4} \deg_{\mathcal{I}}(v) Z_v,$$

(10)

$$\delta_{i,j}^{k,l} := \begin{cases} 0 & \text{if } (i \neq k) \wedge (j = 0 \bmod 4) \wedge ([i,j] \text{ is blue}), \\ 0 & \text{if } (i \neq k) \wedge (j = 2 \bmod 4) \wedge ([i,j] \text{ is yellow}), \\ 1 & \text{otherwise.} \end{cases}$$
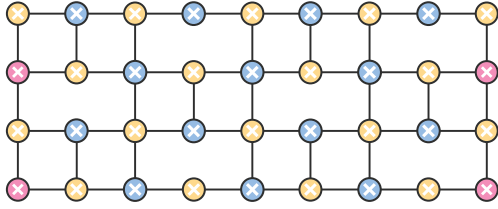
Figure 2: Quantum quench architecture. Circles denote qubits, lines denote interactions. Blue sites are initialised on $|+\rangle$; pink ones on $|0\rangle$ or $|1\rangle$ at random; yellow ones on $e^{-ik_i\pi Z/8}|+\rangle$ with random $k_i \in \{0,1,2,3\}$. The Hamiltonian evolution (10) implements CZ gates on connected qubit pairs. Qubits are measured in the X basis.

Above, $\delta_{i,j}^{k,l}$ is the indicator function of the edge set $E_{\mathcal{I}}$ of a (4,2)-periodic interaction sub-lattice $\mathcal{I} = (V, E_{\mathcal{I}}) \subset \mathcal{L}$ (Fig. 2) and $\deg_{\mathcal{I}}(v)$ is the degree of $v \in V$ in $\mathcal{I}$. It is easily seen that $\mathcal{I}$ is a brick-work pattern of 2-square-cells with closed boundaries (Fig. 2). The net effect of the dynamics is to implement a controlled-Z gate on every pair of neighbouring qubits in $\mathcal{I}$ before all qubits are measured in the X basis.

In what follows the central quantity will be the probability distribution defined by the probabilities

$$q_{ac}(x, \beta) := \left| \langle x | H^{\otimes n} e^{-iH_{ac}} | \psi_\beta \rangle \right|^2 / |\mathcal{S}_{ac}| \quad (11)$$

of measuring the outcomes x after picking $|\psi_\beta\rangle$ and evolving it under H_{ac} for unit time. We also let x_R be x's sub-string of rightmost column's outcomes in the square lattice, and $x_L := x - x_R$ be its set-theoretic complement.

4.2 Anticoncentration and classical hardness of $\mathcal{Q}_{ac}$

Our second main result from which anticoncentration and classical hardness follow, shows that this quantum quench architecture is closely related to the "dense" IQP circuit family of Ref. [6] consisting of circuits of $e^{i\theta_i \pi X_i}$, $e^{i\theta_{i,j} \pi X_i X_j}$ gates acting at arbitrary pairs of qubits, with $\theta_i, \theta_{i,j}$ chosen fully and uniformly at random from $\{k\pi/8 : k \in \{0,\dots,7\}\}$. Dense IQP circuits form a commutative finite group under multiplication $\mathcal{G}_{IQP}$ with Haar measure μ_{IQP}. The implementation of a dense IQP circuit proposed in [6] requires a fully-connected architecture, and the enactment of $\Theta(m^2)$ long-range gates for m-qubits in average. Here, we show that our constant-depth nearest-neighbour architecture $\mathcal{Q}_{ac}$ implements *exact* sampling over dense IQP circuits with a linear $(2m+1)$ overhead-factor in qubit number.

Theorem 9 (Quantum quench architecture). *For $n(m) = m(2m+1)$ qubits, the output probability distribution q_{ac} of architecture $\mathcal{Q}_{ac}$ fulfils*

$$q_{ac}(x_L|\beta) = \frac{1}{2^{n-m}}, \quad q_{ac}(x_R|x_L, \beta) = |\langle x_R | V_{x_L, \beta} | 0 \rangle|^2$$

for some x_L, β-dependent m-qubit dense IQP circuit $V_{x_L,\beta} \in \mathcal{G}_{IQP}$ such that $\Pr_{(x_L,\beta)\sim q_{ac}}(V_{x_L,\beta}) = \mu_{IQP}(V_{x_L,\beta})$.

Before we turn to proving this theorem, let us state the two Corollaries important for us, namely, that both the anticoncentration result and the simulability result proven for dense IQP circuits in Ref. [6] carry over to the quench architecture $\mathcal{Q}_{ac}$.

Corollary 10 (Anticoncentration from quenched dynamics). *The distribution q_{ac} (11) of the quench architecture $\mathcal{Q}_{ac}$ described below satisfies $q_{ac}(\beta) = 1/|\mathcal{S}_{ac}|$ and*

$$\Pr_{\beta \sim q_{ac}} \left(q_{ac}(x|\beta) \geq \frac{1}{2N} \right) \geq \frac{1}{12}. \quad (12)$$

Corollary 11 (Hardness of approximation). *Approximating either $q_{ac}(x, \beta)$ or $q_{ac}(x|\beta)$ up to relative error $1/4$ is #P-hard.*

Corollary 10 proves anticoncentration for $\mathcal{Q}_{ac}$'s output probabilities, while Corollary 11 shows that the latter are #P-hard to approximate. As an application of these two technical statements, a quantum-speedup result follows from a Stockmeyer argument. This result is based on the average case conjecture C1*:

C1* Let $H_v := H + \sum_{i \in V} v_i Z_i$ be the random Ising model derived from (9) by adding uniformly random on-site fields $v_i Z_i$ with random angles $v_i = (\beta_i + x_i)/2$, where β_i are the phases of the input state $|\psi_\beta\rangle$ and x_i are the measurement outcomes. The conjecture states that, if its #P-hard to approximate the imaginary temperature partition function $\mathrm{tr}\,[\exp(iH_v)]$ up to a constant relative error, then the same problem is #P-hard for a constant fraction of the instances—intuitively, because random Ising models have no visible structure making this problem easier in average (see also Refs. [6–8]).

Corollary 12 (Intractability of classical sampling). *If conjectures C1-C1* hold, then a classical computer cannot sample from the outcome distribution of architecture $\mathcal{Q}_{ac}$ up to ℓ_1-error $1/192$ in time $O(\mathrm{poly}(n))$.*

The significance of Corollary 12 is that, as shown below, architecture $\mathcal{Q}_{ac}$ defines a resource-wise plausible, certifiable experiment for demonstrating a quantum speedup with experimental demands competitive to those in Ref. [10]. However, unlike the latter, the speedup of $\mathcal{Q}_{ac}$ relies only on a natural average-case hardness conjecture about Ising models and a Polynomial Hierarchy collapse. Hence, we believe this corollary should help in the analysis of quantum speedups in near-term quantum devices.

4.3 Proof of Theorem 9 and its Corollaries

Proof of Theorem 9. We make use of two circuit gadgets, named "odd" and "even", illustrated next,

Odd gadget   **Even gadget**
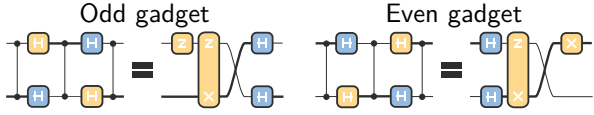
which represent quantum circuit identities modulo terminal Pauli operator corrections. Vertical links represent CZ gates. Crossing qubit lines perform SWAP gates. Blue "H" blocks implement Hadamard gates preceded by uniformly-random $Z_i^x, x \in \{0,1\}$, single-qubit gates. Yellow "H" blocks, $H_i e^{-i\pi k Z_i/8} Z_i^{x'}$ gates with uniformly-random $0 \le k \le 3$, $x' \in \{0,1\}$, where $e^{-ik\pi Z_i/8} Z_i^{x'}$ is a uniformly-random power of $e^{-i\pi Z_i/8}$ up to a global phase since the latter has order 8 and $Z_i \propto e^{-i4\pi Z_i/8}$. Analogously, yellow "Z" (resp. "ZX" blocks) perform uniformly-random powers of $e^{-i\frac{\pi}{8} Z_i}$ (resp. $e^{-i\frac{\pi}{8} Z_i X_{i+1}}$). The correctness of the identities is easily verified using the stabiliser formalism [51]. Pauli corrections correspond to "by-product" Zs in blue blocks, which we can propagate to the end of the circuit by flipping some of the $e^{-ik\pi Z_i/8}$ gates' angles in yellow blocks, which leaves them invariant.

We next show that the computation carried out by $\mathcal{Q}_{ac}$ is equivalent to a 1D circuit of our odd and even gadgets composed in a brickwork layout. We begin by reminding the reader of the properties of X-teleportation circuits [47], namely, that given an $(r+1)$-qubit state vector $|\psi\rangle|+\rangle$, the effect of measuring the i-th qubit of $|\psi\rangle$ in the $D^\dagger X D$ basis after entangling it with $|+\rangle$ via a CZ gate is, first, to produce a uniformly-random bit x; second, teleport the value of the former qubit onto the latter; and, third, implement a single-qubit gate $H_{r+1} Z_{r+1}^x D_{r+1}$ on site $r+1$. Next, note that pink sites in Fig. 2 can be eliminated from the lattice by introducing uniformly-random simultaneous Z_i rotations on their neighbouring qubits. Combining these three facts and using induction, we obtain that $\mathcal{Q}_{ac}$ can be simulated exactly by an algorithm that first generates a uniformly-random classical bit-string $x_L \in \{0,1\}^{n-m}$ and then draws x_R from the output of the following network of random 1D nearest-neighbour quantum gates,
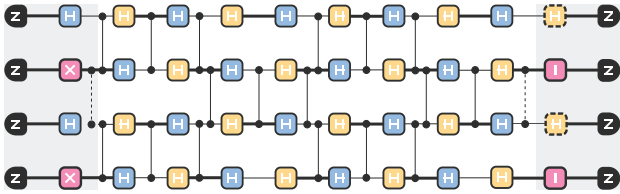
which we draw for $m = 4$ and explicate next. The "bulk" of this network (white area) contains an m-layered brickwork layout of odd and even gadgets with boundaries connected by pairs of blue and yellow blocks. Blue/yellow blocks act as before. $n-m$ out of these are placed in the bulk; their associated random $Z_i^{x_{L_i}}$ gates originally correspond to the by-product rotations introduced via X-teleportation, and are activated by the algorithm depending distinct bits val-

ues of x_L. Qubits are initialised on $|0\rangle$, followed by a "blue" Hadamard (resp. a "pink" uniformly random $\{I_i, X_i\}$) gate on odd (resp. even) rows. Even qubit lines are measured on the Z basis (preceded by "pink" identity gates in the figure); and odd ones on the X basis preceded by a $e^{i\frac{\pi k}{8} Z_i}$ gate. Straight-line random blocks are mutually uncorrelated (terminal "dashed" ones are not). Dashed CZs are "gauge gates" that can be included or removed from by inserting CNOT gates at predetermined input/output locations and reinterpreting the measurement outcomes. As before, we assume H block's by-product Pauli operators are w.l.o.g. conjugated to the end.

Next, we apply our odd/even gadgets to the bulk of our network to rewrite the full quantum circuit in an m-layered brickwork normal form
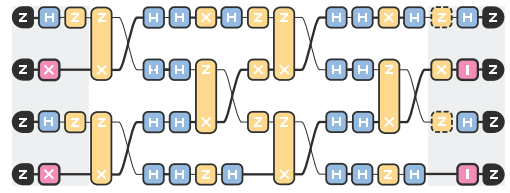
where odd layers execute random gates of the form

$$\prod_{\text{odd } i} \left[H_i H_{i+1} \text{SWAP}_{i,i+1} e^{-i\frac{\pi a_i}{8} Z_i X_{i+1}} e^{-i\frac{\pi b_i}{8} Z_i} \right],$$

with $a_i, b_i \in \mathbb{Z}_8$, followed by random-gate even layers of the form

$$\prod_{\text{even } i} \left[e^{-i\frac{\pi d_i}{8} X_i} \text{SWAP}_{i,i+1} e^{-i\frac{\pi c_i}{8} Z_i X_{i+1}} H_{i-1} H_i \right],$$

where $c_i, d_i \in \mathbb{Z}_8$ and we define $H_j = Z_j = X_j = \text{SWAP}_{k,k+1} = 1$ for $j, k < 1$ and $j, k+1 > n$. Trailing Hadamard gates in odd layers cancel out with their counterparts in even-layers. By a parity-counting argument, it follows that SWAP gates move qubits initially on odd (resp. even) rows travel down (resp. up) the circuit; the latter first undergo Z-type (resp. X-type) interactions, meet an odd number of H gates when they reach the bottom (resp. top) qubit line, and then undergo the opposite process. By propagating all Hadamards in the full circuit to the measurement step, we are left only with a bulk of n brickwork layers of uniformly-random $e^{-i\frac{\pi a_i}{8} X_i X_{i+1}}$, $e^{-i\frac{\pi b_i}{8} X_i}$ and SWAPs, and some additional IQP gates and random Pauli by-products in the preparation/measurement steps. It was shown in Ref. [10] that all pairs of qubits in a bulk circuit of the given form meet exactly once, hence, the network implements exact sampling over dense IQP circuits (crucially, due to their lack of temporal structure). Furthermore the remaining gates are either also dense IQP gates, which leave the Haar measure μ_{IQP} invariant, or terminal Pauli Z gates, which do not affect the final measurement statistics. $\square$

We now exploit the mapping in Lemma 9 between $\mathcal{Q}_{\mathrm{ac}}$'s and IQP circuits' output statistics to prove Corollaries 10 and 12.

Proof of Corollary 10. Recall that m-qubit dense IQP circuits fulfil

$$\Pr_{V \sim \mu_{\mathrm{IQP}}} \left[|\langle x|V|0\rangle|^2 \geq \tfrac{1}{2^{m+1}} \right] \geq \tfrac{1}{12}, \forall x \in \{0,1\}^m. \tag{13}$$

Since $V_{x_L,\beta}$ is drawn according to μ_{IQP} in Lemma 9, we get

$$\Pr_{(x_L,\beta) \sim q_{\mathrm{ac}}} \left[q_{\mathrm{ac}}(x_R|x_L,\beta) \geq \tfrac{1}{2^{m+1}} \right] \geq \tfrac{1}{12}. \tag{14}$$

Since $q_{\mathrm{ac}}(x_L|\beta) = 1/2^{n-m}$, we derive (12). Last, $q_{\mathrm{ac}}(\beta) = 1/|\mathcal{S}_{\mathrm{ac}}|$ by definition. $\square$

Proof of Corollary 12. The proof of Corollary 12 is analogous to those of Ref. [10, Theorem 1] and Ref. [6, Theorem 7], noting that X-measurements on qubits prepared in states $|0\rangle$ or $|1\rangle$ in $\mathcal{Q}_{\mathrm{ac}}$ are equivalent to the Z-measurements on qubits prepared in the $|+\rangle$-state of architecture III in Ref. [10]. Then, the same argument as in Ref. [10] shows that the output probabilities $q_{\mathrm{ac}}(x_L, x_R|\beta)$ are proportional to an Ising partition function as in conjecture C1*.

The only remaining difference with the proof of Theorem 1 in Ref. [10] is that we employ a different anticoncentration bound. Here, we use Eq. (12) of Theorem 10, which is the same bound used in Ref. [6, Theorem 7]. As a result, we obtain a bound of 1/192 for the allowed sampling error identical to that of Theorem 7 in [6]. $\square$

5 Implications and discussion

We now discuss the implications of our two main anticoncentration results and discuss possible improvements in both settings.

First, we conjecture that the linear-circuit-depth scaling in our anticoncentration result for universal random circuits in one dimension is optimal. Indeed, on the one hand, this result is in agreement with the intuition that anticoncentration arises as soon as correlations have spread across the entire system, a process that occurs ballistically and thus scales with the diameter of the system. On the other hand, for one-dimensional random universal circuits to be intractable classically, the depth needs to be polynomial in the number of qubits. Hence, our result only leaves room for a sub-linear improvement, since for circuits of poly-logarithmic depth there is a quasi-polynomial time classical simulation based on matrix-product states. However, as is argued in Refs. [7, 18], it would seem counter-intuitive that one can achieve sub-linear depth. Indeed, standard tensor network contraction techniques would allow any output probabilities of a circuit of depth t in one dimension to be computed in a time scaling as $O(2^t)$ [61]. Hence, if

the depth t as a function of n required for the classical hardness of generic circuits could be brought down to sub-linear, this would violate the counting exponential time hypothesis [62] and is therefore considered highly unlikely.

Second, we highlight that the anticoncentration result for the two-dimensional quenched-dynamics setting provably achieves the optimal asymptotic scaling of depth, namely, constant in the number of qubits. This is due to the highly specific structure of the dynamical evolution and not believed to hold in an approach that relies on sampling random gates such as Refs. [7, 8, 63]. Indeed, in such settings a scaling as $\Theta(\sqrt{n})$ is expected to be necessary and sufficient for an average-case hardness result and hence for anticoncentration. Again, this is due to the ballistic spreading of correlations in the system. Last, the discussed connections between 2D quenches and one-dimensional random circuits lead us to conjecture that the required lattice width in our result, $m \times (2m + 1) \in O(m^2)$, is also asymptotically optimal.

6 Conclusion

In summary, we have presented two anticoncentration theorems for quantum speedup schemes that are based on simple nearest-neighbour interactions and hence realisable with plausible physical architectures, filling a significant gap in the literature. We contrast the anticoncentration property of random circuits in one dimension that are sampled from a universal gate set with anticoncentration of the output distribution of quenched constant-time evolution of product states under translation-invariant nearest-neighbour Ising models. In the former setting the depth required to achieve classical hardness and at the same time anticoncentration of the output distribution scales with the diameter of the system size. In the latter setting a similar hardness and anticoncentration result is achieved after evolution for constant time. We argue that both results are optimal for the respective setting. We hope that this kind of endeavour significantly contributes to the quest of realising quantum devices that outperform classical supercomputers, equipped with strong complexity-theoretic claims.

7 Acknowledgements

References

[1] J. Preskill, Bull. Am. Phys. Soc. **58** (2013), arXiv:1203.5813 .

[2] S. Trotzky, Y.-A. Chen, A. Flesch, I. P. McCulloch, U. Schollwöck, J. Eisert, and I. Bloch, Nature Phys. **8**, 325 (2012), arXiv:1101.2659 .

[3] J.-Y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, Science **352**, 1547 (2016), arXiv:1604.04178 .

[4] S. Braun, M. Friesdorf, S. S. Hodgman, M. Schreiber, J. P. Ronzheimer, A. Riera, M. del Rey, I. Bloch, J. Eisert, and U. Schneider, Proc. Natl. Ac. Sc. **112**, 3641 (2015), arXiv:1403.7199 .

[5] S. Aaronson and A. Arkhipov, Th. Comp. **9**, 143 (2013), arXiv:1011.3245 .

[6] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Phys. Rev. Lett. **117**, 080501 (2016), arXiv:1504.07999 .

[7] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Quantum **1**, 8 (2017).

[8] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Nature Physics , 1 (2018), arXiv:1608.00263 .

[9] X. Gao, S.-T. Wang, and L.-M. Duan, Phys. Rev. Lett. **118**, 040502 (2017), arXiv:1607.04947 .

[10] J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, and J. Eisert, Phys. Rev. X **8**, 021010 (2018), arXiv:1703.00466 .

[11] T. Morimae, Phys. Rev. A **96**, 040302 (2017), arXiv:1704.03640 .

[12] J. Miller, S. Sanders, and A. Miyake, Phys. Rev. A **96**, 062320 (2017), arXiv:1703.11002 .

[13] C. Gogolin, M. Kliesch, L. Aolita, and J. Eisert, "Boson sampling in the light of sample complexity," arXiv:1306.3995 .

[14] S. Aaronson and A. Arkhipov, "BosonSampling is far from uniform," arXiv:1309.7460 .

[15] D. Hangleiter, M. Kliesch, M. Schwarz, and J. Eisert, Quantum Sci. Technol. **2**, 015004 (2017), arXiv:1602.00703 .

[16] T. Kapourniotis and A. Datta, (2017), arXiv:1703.09568 .

[17] L. Stockmeyer, Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing, STOC '83, 118 (1983).

[18] A. P. Lund, M. J. Bremner, and T. C. Ralph, npj Quant. Inf. **3**, 15 (2017), arXiv:1702.03061 .

[19] M. Schwarz and M. V. den Nest, (2013), arXiv:1310.6749 .

[20] R. Jozsa and M. V. d. Nest, Quant. Inf. Comp **14**, 0633–0648 (2014), arXiv:1305.6190 .

[21] Y. Nakata, M. Koashi, and M. Murao, New J. Phys. **16**, 053043 (2014), arXiv:1311.1128 .

[22] D. Gross, K. Audenaert, and J. Eisert, J. Math. Phys. **48**, 052104 (2007), arXiv:quant-ph/0611002 .

[23] F. G. S. L. Brandão, A. W. Harrow, and M. Horodecki, Commun. Math. Phys. **346**, 397 (2016).

[24] M. J. Bremner, R. Jozsa, and D. J. Shepherd, Proc. Roy. Soc. **467**, 2126 (2010), arXiv:1005.1407 .

[25] B. M. Terhal and D. P. DiVincenzo, Quant. Inf. Comp. **4**, 134 (2004), arXiv:quant-ph/0205133 .

[26] G. Kuperberg, Theory of Computing **11**, 183 (2015).

[27] K. Fujii and T. Morimae, New J. Phys. **19**, 033003 (2017), arXiv:1311.2128 .

[28] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, (2018), arXiv:1803.04402 .

[29] R. L. Mann and M. J. Bremner, arXiv:1711.00686 .

[30] S. Aaronson, "P≠NP?" in Open problems in mathematics (Springer, 2016).

[31] L. Fortnow, in Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing, STOC '05 (ACM, 2005).

[32] R. M. Karp and R. J. Lipton, in Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing, STOC '80 (1980).

[33] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Phys. Rev. A **80**, 012304 (2009), arXiv:quant-ph/0606161 .

[34] E. Onorati, O. Buerschaper, M. Kliesch, W. Brown, A. H. Werner, and J. Eisert, Commun. Math. Phys. **355**, 905 (2017), arXiv:1606.01914 .

[35] A. W. Harrow and R. A. Low, Commun. Math. Phys. **291**, 257 (2009), arXiv:0802.1919 .

[36] H. Zhu, R. Kueng, M. Grassl, and D. Gross, (2016), arXiv:1609.08172 .

[37] J. Emerson, Y. S. Weinstein, M. Saraceno, S. Lloyd, and D. G. Cory, Science **302**, 2098 (2003).

[38] W. G. Brown, Y. S. Weinstein, and L. Viola, Phys. Rev. A **77**, 040303 (2008), arXiv:0802.2675 .

[39] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, R. Barends, B. Burkett, Y. Chen, and Z. Chen, (2017), arXiv:1709.06678 .

[40] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, Inf. Proc. Lett. **75**, 101 (2000), arXiv:quant-ph/9906054 .

[41] A. Kitaev, A. Shen, and M. Vyalyi, Classical and Quantum Computation, Graduate studies in

mathematics (American Mathematical Society, 2002).

[42] Y. Shi, Quant. Inf. Comp. **3**, 84 (2003), arXiv:quant-ph/0205115 .

[43] A. Paetznick and B. W. Reichardt, Phys. Rev. Lett. **111**, 090505 (2013), arXiv:1304.3709 .

[44] P. W. Shor, in *Proc. of 37th Conf. Found. Comp. Sci.* (1996) pp. 56–65, arXiv:quant-ph/9605011 .

[45] E. Knill, R. Laflamme, and W. Zurek, (1996), arXiv:quant-ph/9610011 .

[46] E. Knill, R. Laflamme, and W. H. Zurek, Proc. Roy. Soc. A **454**, 365 (1998).

[47] D. Gottesman and I. L. Chuang, Nature **402**, 390 (1999).

[48] R. Raussendorf, D. E. Browne, and H. J. Briegel, Phys. Rev. A **68**, 022312 (2003), arXiv:quant-ph/0301052 .

[49] L. A. Goldberg and H. Guo, (2014), arXiv:1409.5627 .

[50] A. Y. Kitaev, Russ. Math. Surv. **52**, 1191 (1997).

[51] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*, Cambridge Series on Information and the Natural Sciences (Cambridge University Press, 2000).

[52] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, SIAM J. Comp. **26**, 1510 (1997), arXiv:quant-ph/9701001 .

[53] S. Aaronson, Proc. Roy. Soc. A **461**, 2063 (2005), arXiv:quant-ph/0412187 .

[54] H. Dell, T. Husfeldt, D. Marx, N. Taslaman, and M. Wahlén, ACM Trans. Algorithms **10**, 21:1 (2014).

[55] S. X. Cui and Z. Wang, J. Math. Phys. **56**, 032202 (2015).

[56] A. Bocharov, M. Roetteler, and K. M. Svore, Phys. Rev. A **91**, 052317 (2015), arXiv:1409.3552 .

[57] R. Cleve, D. Leung, L. Liu, and C. Wang, Quant. Inf. Comp. **16**, 0721 (2016), arXiv:1501.04592 .

[58] R. Koenig and J. A. Smolin, J. Math. Phys. **55**, 122202 (2014), arXiv:1406.2170 .

[59] J. Eisert, M. Friesdorf, and C. Gogolin, Nature Phys **11**, 124 (2015), arXiv:1408.5148 .

[60] A. Polkovnikov, K. Sengupta, A. Silva, and M. Vengalattore, Rev. Mod. Phys. **83**, 863 (2011), arXiv:1007.5331 .

[61] R. Jozsa, (2006), arXiv:quant-ph/0603163 .

[62] R. Impagliazzo and R. Paturi, in *Proc. XIV IEEE Conf. Comp. Compl.* (1999) pp. 237–240.

[63] S. Aaronson and L. Chen, (2016), arXiv:1612.05903 .

[64] M. Ozols, *How to generate a random unitary matrix* (Mar, 2009).

[65] F. Mezzadri, (2006), arXiv:math-ph/0609050 .

[66] Y. S. Weinstein and C. S. Hellberg, Phys. Rev. A **72**, 022331 (2005).

[67] K. Zyczkowski and M. Kus, J. Phys. A **27**, 4235 (1994).

[68] M. Pozniak, K. Zyczkowski, and M. Kus, J. Phys. A **31**, 1059 (1998), arXiv:chao-dyn/9707006 .

[69] F. Haake, *Quantum signatures of chaos*, Springer Series in Synergetics, Vol. 54 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2010).

A Some facts on random matrix theory, the Haar measure, and unitary designs

A.1 The Haar measure

In this appendix, we give a precise definition of the Haar measure on the unitary group. To do so, let us first define a Radon measure.

Definition 13 (Radon measure)**.** *Let $(X, \mathcal{T})$ be a topological space and $\mathcal{B}$ its Borel algebra. A Radon measure on X is a measure $\mu : \mathcal{B} \to [0, +\infty]$ such that*

  i for any compact set $K \subset X$, $\mu(K) < \infty$

  ii for any $B \in \mathcal{B}$, $\mu(B) = \inf\{\mu(V) : B \subset V \text{ and } V \text{ open}\}$

  iii for any open set $V \subset X$, $\mu(V) = \sup\{\mu(K) : K \subset V \text{ and } K \text{ compact}\}$.

Definition 14 (Haar measure on the unitary group)**.** *The Haar measure is the unique (up to a strictly positive scalar factor) Radon measure which is non-zero on non-empty open sets and is left- and right-translation invariant, i.e.*

$$\mu_{\text{Haar}}(U) > 0 \quad \text{for any non-empty open set } U \subset \mathcal{U} \tag{15}$$

and

$$\mu_{\text{Haar}}(B) = \mu_{\text{Haar}}(uB) = \mu_{\text{Haar}}(Bu) \tag{16}$$

for any $u \in \mathcal{U}$ and Borel set B of $\mathcal{U}$, where the left- and right-translate of B with respect to u is given by

$$uB = \{u\,b : b \in B\} \quad \text{and} \quad Bu = \{b\,u : b \in B\}. \tag{17}$$

A.2 Random matrix ensembles

For the calculation of the distribution matrix elements of Haar-random unitaries it is instructive to introduce a few important ensembles of random matrices. In this appendix we do so from a rather hands-on perspective.

- G(N) (Ginibre Ensemble): The set of matrices Z with complex Gaussian entries.

  G(N) is characterised by the measure $\mathrm{d}\mu_G(Z) \coloneqq \pi^{-N^2} \exp(-\operatorname{tr}(Z^\dagger Z)\mathrm{d}Z$, i.e., each individual entry $z_{i,j}$ is distributed as $\exp(-|z_{i,j}|^2)/\pi$.

- GUE(N) (Gaussian Unitary Ensemble): The set of $N \times N$ Hermitian matrices with complex Gaussian entries, i.e., $H \in$ GUE $\Leftrightarrow H = D + R + R^\dagger$, where D is a diagonal matrix with real Gaussian entries and R is an upper triangular matrix with complex Gaussian entries.

  GUE(N) is characterised by the measure $d\mu_{GUE} = Z_{\mathrm{GUE}(N)}^{-1} \exp(-N \mathrm{tr}(H^2)/2) dH$ on the space of Hermitian matrices.

- CUE(N) (Circular Unitary Ensemble): The set of Haar-random $N \times N$ unitary matrices.

  CUE(N) is characterised by the Haar measure $d\mu_{\mathrm{Haar}}$.

All of $d\mu_{GUE}, d\mu_G$, and $d\mu_{\mathrm{Haar}}$ are left- and right invariant under the action of $U(N)$. There are two ways of constructing Haar-random matrices.

1. Draw a Gaussian matrix $Z \in \mathrm{G}(N)$, and perform the unique QR decomposition such that $Z = QR$, with an orthogonal matrix Q and R is required to have positive diagonal entries. Setting $U = Q$, yields a Haar-random unitary [64, 65].

2. Draw a GUE matrix $Z \in$ GUE. Since Z is Hermitian, the eigenvectors $v_i, i = 1, \ldots, N$ of Z are orthonormal. Multiplying each eigenvector v_i by a random phase e^{ϕ_i} we can construct a Haar-random unitary matrix $U = (e^{\phi_1} v_1 \, e^{\phi_2} v_2 \cdots e^{\phi_n} v_N)$ writing those eigenvectors into the columns of U [66].

A.3 Unitary designs

It is a simple exercise to show that if μ is a unitary k-design, all up to the k^{th} moments of μ equal the moments of the Haar measure.

Lemma 15 ($k-1$ designs from k designs)**.** *Let μ be a distribution on the unitary group $U(N)$ that is an exact unitary k-design. Then μ is also a $(k-1)$-design.*

Proof. Let μ be a unitary k design. That means that it holds

$$\mathbb{E}_\mu \left[U^{\otimes k} X (U^\dagger)^{\otimes k} \right] = \mathbb{E}_{\mathrm{Haar}} \left[U^{\otimes k} X (U^\dagger)^{\otimes k} \right] \quad (18)$$

for all operators X acting on $\mathcal{L}(\mathcal{H}^{\otimes k})$. Choose $X = Y \otimes \mathrm{id}$ with Y being an arbitrary operator on $\mathcal{L}(\mathcal{H}^{\otimes k-1})$. Then

$$\mathbb{E}_\mu \left[U^{\otimes k-1} Y (U^\dagger)^{\otimes k-1} \right] = \mathbb{E}_\mu \left[U^{\otimes k} X (U^\dagger)^{\otimes k} \right] \quad (19)$$

i.e., μ is a unitary $(k-1)$-design. $\qquad \square$

Corollary 16 (Approximate $k-1$ designs from approximate k designs)**.** *Let μ be an (additive or relative) approximate unitary k-design. Then μ is also an approximate unitary $(k-1)$-design, i.e.,*

$$\|M_\mu^k - M_{\mathrm{Haar}}^k\|_\diamond \leq \epsilon \Rightarrow \|M_\mu^{k-1} - M_{\mathrm{Haar}}^{k-1}\|_\diamond \leq \epsilon \quad (20)$$

and likewise for relative errors.

B Matrix elements of Haar-random unitaries

Let us now derive the distribution of the amplitudes $|\langle a|U|b\rangle|^2$ of the matrix elements a Haar-random unitary U [66–69]. To this end we apply knowledge about the distribution of entries of eigenvectors of GUE matrices and their relation to Haar-random unitaries (see App. A.2). We follow Ref. [69], Chapter 4.9.

The eigenvectors v_i of a given operator $H \in$ GUE(N) have N complex components c_k and unit norm $\|v_i\|_2 = 1$. Since every eigenvector can be unitarily transformed into an arbitrary vector of unit norm, the only invariant characteristic of those eigenvectors is the norm itself. Thus, the joint probability for its components $\{c_k\}$ must read

$$P_{\mathrm{GUE}}(\{c_k\}) = \mathrm{const} \cdot \delta \left(1 - \sum_{k=1}^N |c_k|^2 \right), \quad (21)$$

where the constant is fixed by normalisation.

Assuming real entries for now (we can always go to complex ones by doubling N) we can calculate that normalisation by evaluating the integral on the N-dimensional unit sphere

$$\mathrm{const} = \int_{-\infty}^\infty \left(\prod_{i=1}^N dc_i \right) \delta \left(1 - \sum_{k=1}^N |c_k|^2 \right) \quad (22)$$

$$= \int d\omega^{N-1} \int_0^\infty dR \, R^{N-1} \delta(1 - R^2) \quad (23)$$

$$= \int d\omega^{N-1} \int_0^\infty dR \, R^{N-1} \frac{1}{2R} [\delta(1-R) + \delta(1+R)] \quad (24)$$

$$= \pi^{N/2} / \Gamma(N/2) \,. \quad (25)$$

Similarly, we can calculate the marginal distribution

$$P^{(N,l)}(c_1, \ldots, c_l) = \pi^{-N/2} \Gamma(N/2) \int_{-\infty}^{\infty} \left(\prod_{i=l+1}^{N} \mathrm{d}c_i \right) \delta \left(1 - \sum_{k=1}^{N} |c_k|^2 \right) \tag{26}$$

$$= \int \mathrm{d}\omega^{N-l-1} \int_0^{\infty} \mathrm{d}R \, R^{N-l-1} \delta(1 - R^2 - \sum_{k=1}^{l} |c_k|^2) \tag{27}$$

$$= \pi^{-l/2} \frac{\Gamma(N/2)}{\Gamma((N-l)/2)} \left(1 - \sum_{k=1}^{l} |c_k|^2 \right)^{(N-l-2)/2}. \tag{28}$$

For the GUE we then obtain the probability density for the amplitude $y = x_1^2 + x_2^2$ of a single complex entry $x_1 + \mathrm{i}x_2$ of an eigenvector to be the twofold integral over real and imaginary part

$$P_{\mathrm{GUE}}(y) = \int dx_1 dx_2 P^{(2N,2)}(x_1, x_2) \delta(y - x_1^2 - x_2^2)$$
$$= (N-1)(1-y)^{N-2}. \tag{29}$$

Since the eigenvectors of a GUE matrix are identically distributed (up to a global phase) as the columns of a CUE matrix, we obtain the same distribution as (29) for the amplitudes of the matrix elements of a CUE matrix [66]. Notably, as N becomes much larger than 1, we obtain

$$P_{\mathrm{Haar}}(p) = (N-1)(1-p)^{N-2} \xrightarrow{N \gg 1} N \exp(-Np) . \tag{30}$$

The first and second moments of P_{CUE} are then given by

$$\mathbb{E}_{\mathrm{Haar}}[p] = \frac{1}{N}, \tag{31}$$

$$\mathbb{E}_{\mathrm{Haar}}[p^2] = \frac{2}{N(N+1)}. \tag{32}$$