

# Average-case complexity versus approximate simulation of commuting quantum computations

Michael J. Bremner,<sup>1</sup> Ashley Montanaro,<sup>2</sup> and Dan J. Shepherd<sup>3</sup>

<sup>1</sup>*Centre for Quantum Computation and Intelligent Systems,  
Faculty of Engineering and Information Technology, University of Technology Sydney, NSW 2007, Australia\**

<sup>2</sup>*Department of Computer Science, University of Bristol, UK*

<sup>3</sup>*CESG, Hubble Road, Cheltenham, GL51 0EX, UK*

We use the class of commuting quantum computations known as IQP (Instantaneous Quantum Polynomial time) to strengthen the conjecture that quantum computers are hard to simulate classically. We show that, if either of two plausible average-case hardness conjectures holds, then IQP computations are hard to simulate classically up to constant additive error. One conjecture relates to the hardness of estimating the complex-temperature partition function for random instances of the Ising model; the other concerns approximating the number of zeroes of random low-degree polynomials. We observe that both conjectures can be shown to be valid in the setting of worst-case complexity. We arrive at these conjectures by deriving spin-based generalisations of the Boson Sampling problem that avoid the so-called permanent anticoncentration conjecture.

Quantum computers are conjectured to outperform classical computers for a variety of important tasks ranging from integer factorisation [1] to the simulation of quantum mechanics [2]. However, to date there is relatively little rigorous evidence for this conjecture. It is well established that quantum computers can yield an exponential advantage in the query and communication complexity models. But in the more physically meaningful model of time complexity, there are no *proven* separations known between quantum and classical computation.

This can be seen as a consequence of the extreme difficulty of proving bounds on the power of classical computing models, such as the famous P vs. NP problem. Given this difficulty, the most we can reasonably hope for is to show that quantum computations cannot be simulated efficiently classically, assuming some widely believed complexity-theoretic conjecture. For example, any set of quantum circuits that can implement Shor's algorithm [1] provides a canonical example, with the unlikely consequence of efficient classical simulation of this class of quantum circuits being the existence of an efficient classical factoring algorithm. However, one could hope for the existence of other examples that have wider-reaching complexity-theoretic consequences.

With this in mind, in both [3] and [4] it was shown that the existence of an efficient classical sampler from a distribution that is close to the output distribution of an arbitrary quantum circuit, to within a small multiplicative error in each output probability, would imply that post-selected classical computation is equivalent to post-selected quantum computation. This consequence is considered very unlikely as it would collapse the infinite tower of complexity classes known as the Polynomial Hierarchy [5] to its third level. In both works this was proven even for non-universal quantum circuit families: commuting quantum circuits in the case of [3], and linear-optical networks in [4]. These non-universal families are of physical interest because they are simpler to implement, and easier to analyse because of the elegant mathematical structures on which they are based. However, in physically realistic scenar-

ios where the quantum computer and its corresponding classical simulator are allowed to be accurate up to a small additive error, these results have little meaning.

One important recent step to addressing this was proposed by Aaronson and Arkhipov [4], who gave a sophisticated argument based on counting complexity that approximately sampling from the output probability distribution of a network of noninteracting photons (a problem known as Boson Sampling) should be classically hard, even up to a reasonable additive error bound. This major breakthrough rests on two tantalising but as yet unproven conjectures: the so-called *permanent anticoncentration conjecture* and the *permanent-of-Gaussians conjecture*.

In this paper we propose a generalisation of the Boson Sampling argument of [4] that is native to spin systems, specifically to the class of commuting quantum circuits known as IQP (Instantaneous Quantum Polynomial time), which was introduced in [6] and [3]. In our opinion this leads to a mathematically simpler setting, while still apparently retaining the essential complexity-theoretic ingredients. This simplicity allows us to prove the IQP analogues of the *permanent anticoncentration conjecture*. The only remaining conjecture to prove is an IQP analogue of the *permanent-of-Gaussians* conjecture, of which we find two natural examples.

Informally (see [3] for the formal definition) an  $n$ -qubit IQP circuit  $\mathcal{C}$  is a quantum circuit which takes as input the state  $|0\rangle^{\otimes n}$ , whose gates are diagonal in the Pauli-X basis, and whose  $n$ -qubit output is measured in the computational basis. We say that a classical sampler of an IQP circuit is accurate up to error  $\epsilon$  in  $\ell_1$  norm if its output probability distribution has total variation distance at most  $\epsilon/2$  from that of  $\mathcal{C}$ . We can now state the main result of this paper:

**Theorem 1.** *Assume either Conjecture 2 or 3 below is true. If it is possible to classically sample from the output probability distribution of any IQP circuit  $\mathcal{C}$  in polynomial time, up to an error of  $1/192$  in  $\ell_1$  norm, then any problem in  $\#P$  can be solved in  $\text{FBPP}^{\text{NP}}$ . Hence the Polynomial Hierarchy collapses to its third level.*

Loosely speaking, the complexity class  $\#P$  appearing in this theorem is the class of problems reducible to counting the number of solutions of an NP problem [5];  $\text{FBPP}^{\text{NP}}$  is the

\* michael.bremner@uts.edu.au

class of problems that can be solved by randomised classical polynomial-time computation equipped with an oracle that can solve any problem in NP. Theorem 1 is based around two natural conjectures, one of which is native to computer science, and other common in condensed-matter physics. Each conjecture straightforwardly relates to a family of IQP circuits. Essentially the result states that if we assume either of our conjectures, then there is no way of efficiently simulating these families of quantum circuits without a major re-evaluation of the existing status-quo of complexity theory.

The first conjecture is based on the complexity of one of the most commonly studied models of statistical physics, the Ising model. Consider the partition function

$$Z(\omega) = \sum_{z \in \{\pm 1\}^n} \omega^{\sum_{i < j} w_{ij} z_i z_j + \sum_{k=1}^n v_k z_k}, \quad (1)$$

where the exponentiated sum is over the complete graph on  $n$  vertices,  $w_{ij}$  and  $v_k$  are real edge and vertex weights, and  $\omega \in \mathbb{C}$ . Then, for any  $\omega = e^{i\theta}$ ,  $Z(\omega)$  arises straightforwardly as an amplitude of some IQP circuit  $\mathcal{C}$ :  $\langle 0|^{\otimes n} \mathcal{C} |0\rangle^{\otimes n} = Z(\omega)/2^n$  (see Appendix A and [7–11]). Our conjecture is:

**Conjecture 2.** *Let  $Z(\omega)$  be the partition function of a random instance of the Ising model, picked by choosing uniformly random weights from the set  $\{0, \dots, 15\}$  on the vertices and edges of the complete graph on  $n$  vertices, evaluated at  $\omega$ . Then it is #P-hard to approximate  $|Z(e^{i\pi/8})|^2$  up to multiplicative error  $1/4 + o(1)$  for a  $1/24$  fraction of instances over the choice of vertex and edge weights.*

It is known that the family of partition functions  $Z(\omega)$  parametrised as above is #P-hard to compute in the worst case up to the above multiplicative error bound [10, 11]. Conjecture 2 thus states that this worst-case hardness result can be improved to an average-case hardness result.

Our second conjecture is based on the hardness of computing the *gap* of degree-3 polynomials over  $\mathbb{F}_2$ ,  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , which are expressible (up to an additive constant) as

$$f(x) = \sum_{i,j,k} \alpha_{ijk} x_i x_j x_k + \sum_{i,j} \beta_{ij} x_i x_j + \sum_i \gamma_i x_i \pmod{2},$$

where  $\alpha_{ijk}, \beta_{ij}, \gamma_i \in \{0, 1\}$ . The *gap* is defined by  $\text{gap}(f) := |\{x : f(x) = 0\}| - |\{x : f(x) = 1\}|$ . It can be shown that, for any degree-3 polynomial  $f$ ,  $\langle 0|^{\otimes n} \mathcal{C}_f |0\rangle^{\otimes n} = \text{gap}(f)/2^n$  for some IQP circuit  $\mathcal{C}_f$  (see Appendix B). We write  $\text{ngap}(f) = \text{gap}(f)/2^n$ . Then we have the following conjecture:

**Conjecture 3.** *Let  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  be a uniformly random degree-3 polynomial over  $\mathbb{F}_2$ . Then it is #P-hard to approximate  $\text{ngap}(f)^2$  up to a multiplicative error of  $1/4 + o(1)$  for a  $1/24$  fraction of polynomials  $f$ .*

It has been known for some time that  $\text{ngap}(f)$  is #P-hard to compute exactly in the worst case [12]. We show in Appendix D, using IQP techniques, that this worst-case hardness still holds for approximating  $\text{ngap}(f)^2$  up to multiplicative error less than  $1/2$ . Just as with Conjecture 2, what remains is to lift this worst-case hardness result to average-case hardness.

We remark that an additional piece of evidence provided in [4] for the average-case hardness of multiplicative approximations to Boson Sampling (the permanent-of-Gaussians conjecture) was a direct proof that exact simulation of Boson Sampling probability distributions is hard on average. This was based on average-case hardness results for computation of the permanent, for which we do not know IQP analogues. However, currently known techniques do not seem sufficient to extend these exact average-case hardness results for Boson Sampling to approximate hardness results [4, Section 9.2].

As with the case of Boson Sampling, the worst-case hardness of multiplicative approximations to both  $Z(\omega)$  and  $\text{ngap}(f)$  ([11] and Appendix D) implies via standard results on random-self-reducibility [13] that there exists *some* distribution over the choices of these functions that is #P-hard on average – but not necessarily the distributions that we require for Conjectures 2 and 3.

There are a number of technical ingredients of Theorem 1 which will be discussed in full below. The basic idea is that, for the class of problems underlying Conjectures 2 and 3, any classical sampler that is accurate up to a good additive error bound in the worst case, is forced to also be accurate to within a reasonable multiplicative error on average. This is shown by using a classic result of complexity theory, the so-called Stockmeyer counting algorithm [14]. Another essential ingredient is anticoncentration results for  $\text{ngap}(f)$  (for Conjecture 3) and the partition function of the random Ising model (for Conjecture 2). That such anticoncentration results can be proven is a consequence of the elegant mathematical structures upon which IQP circuits are based.

We first mention that recent independent work of Fefferman and Umans [15] has explored an alternative way to generalise the ideas of Aaronson and Arkhipov [4]. This work uses Quantum Fourier Sampling to construct states whose corresponding probability distributions are hard to sample from classically, under similar conjectures to [4]. An appealing aspect of the construction of [15] is that it shows that there are specific, and rather simple, quantum states which are hard to simulate classically, assuming an anticoncentration conjecture holds. However, constructing these states appears to require the full power of quantum computation, unlike the results described here and in [4].

We now describe our results. We assume basic knowledge of standard complexity classes [5, 16] such as P, #P and BPP. A crucial complexity-theoretic tool we will use is a result of Stockmeyer on approximate counting:

**Theorem 4** (Stockmeyer [14], see [4] for statement here). *There exists an FBPP<sup>NP</sup> machine which, for any boolean function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , can approximate*

$$p = \Pr_x[f(x) = 1] = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} f(x)$$

*to within a multiplicative factor of  $(1 + \epsilon)$  for all  $\epsilon = \Omega(1/\text{poly}(n))$ , given oracle access to  $f$ .*

We refer to [3, 6] for formal definitions of the IQP model. We commonly consider IQP circuits which consist of an  $n$ -fold tensor product of a Hadamard gate  $H^{\otimes n}$ , a component

whose gates are diagonal in the Pauli-Z basis, and a further application of  $H^{\otimes n}$ . For the main results of this paper we will consider IQP circuits whose diagonal-Z component consists of either Z, CZ (Controlled-Z, aka Controlled-Phase), and CCZ (Controlled-Controlled-Z) gates or Z,  $\sqrt{\text{CZ}}$ , and  $T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$  gates. Both of these sets of gates would be universal if we could also perform Hadamard (H) gates at any point in the circuit – which we cannot do in IQP because this gate does not commute with the Z gate. However, if we allow the unphysical resource of postselection, these Hadamard gates can effectively be implemented [3], allowing IQP circuit amplitudes  $\langle y|C|x\rangle$  to express any quantum circuit amplitude (up to a known constant). See Appendix C for a description of this construction. A consequence is that exact classical computation of such amplitudes is #P-hard, even up to constant multiplicative error [3] or exponentially small additive error [9].

**Approximation of general IQP circuits** – We first prove a key technical ingredient, which relates approximate sampling from the output distributions of IQP circuits to approximating individual output probabilities. This is essentially the same argument as used in [4] for the permanent, although we believe it becomes substantially simpler in the setting of IQP. The intuition behind this result is that adding random X gates to an IQP circuit randomly permutes the output probabilities. This allows the user of a sampler which is accurate for all circuits to obfuscate from the sampler which one of the output probabilities the user is interested in.

**Lemma 5.** *Let  $\mathcal{C}$  be an arbitrary IQP circuit on  $n$  qubits. Let  $\mathcal{C}_x$ , for  $x \in \{0, 1\}^n$ , be the circuit produced by appending an X gate to  $\mathcal{C}$  for each  $i$  such that  $x_i = 1$ . Assume there exists a classical polynomial-time algorithm  $\mathcal{A}$  which, for any IQP circuit  $\mathcal{C}'$ , can sample from a probability distribution which approximates the output probability distribution of  $\mathcal{C}'$  up to additive error  $\epsilon$  in  $\ell_1$  norm. Then, for any  $\delta$  such that  $0 < \delta < 1$ , there is a FBPP<sup>NP</sup> algorithm which, given access to  $\mathcal{A}$ , approximates  $|\langle 0|\mathcal{C}_x|0\rangle|^2$  up to additive error*

$$O((1 + o(1))\epsilon/(2^n \delta) + |\langle 0|\mathcal{C}_x|0\rangle|^2 / \text{poly}(n))$$

with probability at least  $1 - \delta$  (over the choice of  $x$ ).

*Proof.* For each  $x, y \in \{0, 1\}^n$ , set

$$p_{xy} := \Pr[\mathcal{C}_x \text{ outputs } y], q_{xy} := \Pr[\mathcal{A} \text{ outputs } y \text{ on input } \mathcal{C}_x],$$

where in both cases the probability is taken over the algorithm's internal randomness. For any  $y$  of our choice, we can apply Theorem 4 to  $\mathcal{A}$  to estimate  $q_{0y}$  (where we write 0 for  $0^n$ ). This gives us an FBPP<sup>NP</sup> algorithm which produces an estimate  $\tilde{q}_y$  such that

$$|\tilde{q}_y - q_{0y}| \leq q_{0y} / \text{poly}(n).$$

Then

$$\begin{aligned} |\tilde{q}_y - p_{0y}| &\leq |\tilde{q}_y - q_{0y}| + |q_{0y} - p_{0y}| \\ &\leq q_{0y} / \text{poly}(n) + |q_{0y} - p_{0y}| \\ &\leq (p_{0y} + |q_{0y} - p_{0y}|) / \text{poly}(n) + |q_{0y} - p_{0y}| \\ &= p_{0y} / \text{poly}(n) + |q_{0y} - p_{0y}|(1 + 1 / \text{poly}(n)). \end{aligned}$$

As  $\mathcal{A}$  approximates the output probability distribution of  $\mathcal{C}_0$  up to  $\ell_1$  error  $\epsilon$ , it follows from Markov's inequality that

$$\Pr_y[|p_{0y} - q_{0y}| \geq \epsilon/(2^n \delta)] \leq \delta$$

for any  $0 < \delta < 1$ , where  $y$  is picked uniformly at random. Thus

$$|\tilde{q}_y - p_{0y}| \leq p_{0y} / \text{poly}(n) + \frac{\epsilon(1 + 1 / \text{poly}(n))}{2^n \delta}$$

with probability at least  $1 - \delta$ , over the choice of  $y$ . But

$$p_{0y} = \Pr[\mathcal{C}_0 \text{ outputs } y] = |\langle y|\mathcal{C}_0|0\rangle|^2 = |\langle 0|\mathcal{C}_y|0\rangle|^2.$$

This completes the proof.  $\square$

If  $|\langle 0|\mathcal{C}_x|0\rangle|^2 = \Omega(2^{-n})$ , the algorithm of Lemma 5 gives a good approximation – i.e. a multiplicative approximation up to a  $1 + 1 / \text{poly}(n)$  factor. We encapsulate this formally as the following corollary.

**Corollary 6.** *Let  $\mathcal{F}$  be a family of IQP circuits on  $n$  qubits. Pick a random circuit  $\mathcal{C}$  by choosing a circuit from  $\mathcal{F}$  uniformly at random, then appending X gates on a uniformly random subset of the qubits. Assume that there exist universal constants  $\alpha, p > 0$  such that  $\Pr[|\langle 0|\mathcal{C}|0\rangle|^2 \geq \alpha \cdot 2^{-n}] \geq p$ . Further assume there exists a classical polynomial-time algorithm  $\mathcal{A}$  which, for any IQP circuit  $\mathcal{C}'$ , can sample from a probability distribution which approximates the output probability distribution of  $\mathcal{C}'$  up to additive error  $\epsilon = \alpha p / 8$  in  $\ell_1$  norm. Then there is a FBPP<sup>NP</sup> algorithm which, given access to  $\mathcal{A}$ , approximates  $|\langle 0|\mathcal{C}|0\rangle|^2$  up to multiplicative error  $1/4 + o(1)$  on at least a  $p/2$  fraction of circuits  $\mathcal{C}$ .*

*Proof.* With probability at least  $p$  over the choice of  $\mathcal{C}$ ,  $|\langle 0|\mathcal{C}|0\rangle|^2 \geq \alpha \cdot 2^{-n}$ . Setting  $\epsilon = \alpha p / 8$ ,  $\delta = p/2$  in Lemma 5, there is a FBPP<sup>NP</sup> algorithm which, given access to  $\mathcal{A}$ , approximates  $|\langle 0|\mathcal{C}|0\rangle|^2$  up to additive error  $O((1 + o(1))\alpha \cdot 2^{-n-2} + |\langle 0|\mathcal{C}|0\rangle|^2 / \text{poly}(n))$  with probability at least  $1 - p/2$  (over the choice of  $\mathcal{C}$ ). Therefore, with probability at least  $p/2$ , the algorithm approximates  $|\langle 0|\mathcal{C}|0\rangle|^2$  up to multiplicative error  $1/4 + o(1)$ .  $\square$

We next show that the first precondition of Corollary 6 is indeed satisfied for two interesting families of IQP circuits.

**Anticoncentration bounds** – Fix a family  $\mathcal{F}$  of IQP circuits. We would like to show that  $|\langle 0|\mathcal{C}|0\rangle|^2$  is likely to be high for a circuit  $\mathcal{C}$  formed by picking a random circuit  $\mathcal{D}$  from  $\mathcal{F}$ , then appending X gates on a uniformly random subset  $S$  of the qubits. We will use the following fact:

**Fact 7** (Paley-Zygmund inequality). *If  $R$  is a non-negative random variable with finite variance, then for any  $0 < \alpha < 1$ ,  $\Pr[R \geq \alpha \mathbb{E}[R]] \geq (1 - \alpha)^2 \mathbb{E}[R]^2 / \mathbb{E}[R^2]$ .*

We will apply Fact 7 to the random variable  $R = |\langle 0|\mathcal{C}|0\rangle|^2$ , first observing that

$$\begin{aligned} \mathbb{E}_{\mathcal{C}}[|\langle 0|\mathcal{C}|0\rangle|^2] &= \mathbb{E}_{\mathcal{D}, x}[|\langle x|\mathcal{D}|0\rangle|^2] \\ &= \frac{1}{2^n} \mathbb{E}_{\mathcal{D}} \sum_{x \in \{0, 1\}^n} |\langle x|\mathcal{D}|0\rangle|^2 = \frac{1}{2^n}, \end{aligned}$$

where in the second expectation  $x$  is picked uniformly at random from  $\{0, 1\}^n$ . This deals with the numerator; to handle the denominator, we need to upper-bound  $\mathbb{E}[|0|C|0\rangle|^4]$ .

The first family of circuits we consider corresponds to polynomials over  $\mathbb{F}_2$ . We prove in Appendix E that for uniformly random degree-3 polynomials  $f$ ,  $\mathbb{E}_f[\text{ngap}(f)^4] \leq 3 \cdot 2^{-2n}$ . Based on this, and the tight connection between IQP circuits over the gate set  $\{Z, CZ, CCZ\}$  and degree-3 polynomials, we have the following result:

**Theorem 8.** *Assume there exists a classical polynomial-time algorithm  $\mathcal{A}$  which, for any IQP circuit  $\mathcal{C}$ , can sample from a probability distribution which approximates the output probability distribution of  $\mathcal{C}$  up to additive error  $1/192$  in  $\ell_1$  norm. Then there is an FBPP<sup>NP</sup> algorithm which, given access to  $\mathcal{A}$ , approximates  $\text{ngap}(f)^2$  up to multiplicative error  $1/4 + o(1)$  on at least a  $1/24$  fraction of degree-3 polynomials  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ .*

*Proof.* Combining Fact 7 and the bound on  $\mathbb{E}_f[\text{ngap}(f)^4]$ , we have  $\Pr_f[\text{ngap}(f)^2 \geq \alpha/2^n] \geq (1 - \alpha)^2/3$  for any  $0 < \alpha < 1$ . Fixing  $\alpha = 1/2$ , we get  $\Pr_f[\text{ngap}(f)^2 \geq 2^{-n-1}] \geq 1/12$ . The claim then follows by inserting these parameters into Corollary 6.  $\square$

We next consider the Ising model, where we are interested in evaluating the partition function  $Z(\omega)$  of a randomly weighted graph at  $\omega = e^{i\pi/8}$  (see (1)). Each edge of the complete graph has a weight  $w_{ij}$ , and each vertex has a weight  $v_k$ , each picked uniformly at random from the set  $\{0, \dots, 15\}$ .

We show in Appendix E that  $\langle 0|C|0\rangle = Z(\omega)/2^n$  for an IQP circuit  $\mathcal{C}$  consisting of gates from the set  $\{\text{diag}(1, 1, 1, i), \text{diag}(1, e^{i\pi/4})\}$ ; and further that we can consider a random circuit of this form as being chosen by picking a random circuit using this gate set, then following it by a random choice of X gates. In addition,  $\mathbb{E}_{w,v}[|Z(\omega)|^4] \leq 3 \cdot 2^{2n}$ . This implies the following result, whose proof is essentially the same as that of Theorem 8:

**Theorem 9.** *Let  $Z(\omega)$  be the partition function of a random instance of the Ising model, picked by choosing random weights from the set  $\{0, \dots, 15\}$  on the vertices and edges of the complete graph on  $n$  vertices, evaluated at  $\omega$ . Assume there exists a classical polynomial-time algorithm  $\mathcal{A}$  which, for any IQP circuit  $\mathcal{C}$ , can sample from a probability distribution which approximates the output probability distribution of  $\mathcal{C}$  up to additive error  $1/192$  in  $\ell_1$  norm. Then there is a FBPP<sup>NP</sup> algorithm which, given access to  $\mathcal{A}$ , approximates  $|Z(e^{i\pi/8})|^2$  up to multiplicative error  $1/4 + o(1)$  with probability at least  $1/24$  (over the choice of weights).*

Combining Theorems 8 and 9 gives our main result, Theorem 1.

**Outlook** – In our opinion, the most compelling question is how best to physically implement the classes of IQP circuits

discussed in this paper, as they are arguably among the simplest classes of quantum computations that are unlikely to be classically simulable. For both the Ising model and degree-3 polynomial case the most obvious experimental implementations require non-local gates, which is quite challenging. In terms of the Ising model case 2-qubit interactions might need to be generated between any two qubits in a system. Such interactions do not arise naturally in lattice geometries and as such will need to be engineered. From an experimental perspective there has already been a large amount of progress in these directions. The dynamics of the Ising model with local interactions have been digitally simulated in ion traps [17, 18] and very recently non-local interactions have been utilised in the digital simulation of fermionic systems with superconductors [19]. As technologies such as cavity buses for superconducting systems [20] become more reliable, we expect that an increasing number of systems will be able to implement IQP circuits in a regime that is likely not to be classically simulable.

Theoretically there are a number of natural questions that remain to be answered, the most obvious of which is whether or not Conjectures 2 and 3 are true. To the best of our knowledge there are no techniques that are capable of lifting the known worst-case hardness results for the Ising model and gaps of degree-3 polynomials to average-case results. However, recent breakthroughs [21–23] in categorising the complexity of statistical mechanical systems via the underlying interaction graph properties give some hope that these conjectures can be resolved.

As well as the connections used here between IQP, the Ising model and low-degree polynomials, it is known that IQP circuits are closely related to Tutte polynomials and weight enumerator polynomials of binary linear codes [7]. It is a compelling question as to whether the ideas developed here can be extended to these models. It would also be very interesting to explore the connection between nonadaptive measurement-based quantum computing and IQP [24, 25] in order to find subclasses of IQP that are physically straightforward to implement, and for which one can prove conditional hardness results similar to those shown here.

## ACKNOWLEDGMENTS

We would like to thank Aram Harrow, Richard Jozsa, Gavin Brennen, and Peter Rohde for helpful discussions and comments on this manuscript. AM was supported by the UK EPSRC under Early Career Fellowship EP/L021005/1. MJB has received financial support from the Australian Research Council via the Future Fellowship scheme (grant FT110101044), and Lockheed Martin Corporation.

[1] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer,” SIAM J.

- [2] I. Georgescu, S. Ashhab, and F. Nori, “Quantum simulation,” *Rev. Mod. Phys.* **86**, 153 (2014), arXiv:1308.6253.
- [3] M. Bremner, R. Jozsa, and D. Shepherd, “Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy,” *Proc. R. Soc. A* **467**, 459–472 (2011), arXiv:1005.1407.
- [4] S. Aaronson and A. Arkhipov, “The computational complexity of linear optics,” *Theory of Computing* **9**, 143–252 (2013), arXiv:1011.3245.
- [5] C. Papadimitriou, *Computational Complexity* (Addison-Wesley, 1994).
- [6] D. Shepherd and M. J. Bremner, “Temporally unstructured quantum computation,” *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **465**, 1413–1439 (2009), arXiv:0809.0847.
- [7] D. Shepherd, “Binary Matroids and Quantum Probability Distributions,” (2010), arXiv:1005.1744.
- [8] S. Iblisdir, M. Cirio, O. Boada, and G. K. Brennen, “Low depth quantum circuits for Ising models,” *Annals of Physics* **340**, 205–251 (2014), arXiv:1208.3918.
- [9] X. Ni and M. Van den Nest, “Commuting quantum circuits: efficient classical simulations versus hardness results,” *Quantum Information and Computation* **13**, 0054–0072 (2013), arXiv:1204.4570.
- [10] K. Fujii and T. Morimae, “Quantum commuting circuits and complexity of Ising partition functions,” (2013), arXiv:1311.2128.
- [11] L. A. Goldberg and H. Guo, “The complexity of approximating complex-valued Ising and Tutte partition functions,” (2014), arXiv:1409.5627.
- [12] A. Ehrenfeucht and M. Karpinski, “The computational complexity of (XOR, AND)-counting problems,” (1990), technical Report 8543-CS.
- [13] J. Feigenbaum and L. Fortnow, “Random-Self-Reducibility of Complete Sets,” *SIAM Journal on Computing* **22**, 994–1005 (1993).
- [14] L. Stockmeyer, “On approximation algorithms for #P,” *SIAM J. Comput.* **14**, 849–861 (1985).
- [15] B. Fefferman and C. Umans, “On the power of Quantum Fourier Sampling,” (2015), poster presentation, QIP 2015.
- [16] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- [17] B. P. Lanyon, C. Hempel, D. Nigg, M. Miller, R. Gerritsma, F. Zhringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos, “Universal Digital Quantum Simulation with Trapped Ions,” *Science* **334**, 57–61 (2011), arXiv:1109.1512.
- [18] J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. Joseph Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, “Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins,” *Nature* **484**, 489–492 (2012), arXiv:1204.5789.
- [19] R. Barends, L. Lamata, J. Kelly, L. Garca-Ivarex, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, E. Solano, and John M. Martinis, “Digital quantum simulation of fermionic models with a superconducting circuit,” (2015), arXiv:1501.07703.
- [20] J. Majer, J. M. Chow, J. M. Gambetta, Jens Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, A. Wallraff, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, “Coupling superconducting qubits via a cavity bus,” *Nature* **449**, 443–447 (2007), arXiv:0709.2135.
- [21] A. Sly and N. Sun, “The Computational Hardness of Counting in Two-Spin Models on d-Regular Graphs,” in *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS)* (2012) pp. 361–369, arXiv:1203.2602.
- [22] A. Sinclair, P. Srivastava, and M. Thurley, “Approximation Algorithms for Two-State Anti-Ferromagnetic Spin Systems on Bounded Degree Graphs,” *Journal of Statistical Physics* **155**, 666–686 (2014), arXiv:1107.2368.
- [23] A. Galanis, D. Stefankovic, and E. Vigoda, “Inapproximability for Antiferromagnetic Spin Systems in the Tree Non-uniqueness Region,” in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC ’14* (ACM, New York, NY, USA, 2014) pp. 823–831, arXiv:1305.2902.
- [24] D. Browne, E. Kashefi, and S. Perdrix, “Computational depth complexity of measurement-based quantum computation,” in *Proc. 6th Conference on the Theory of Quantum Computation, Communication, and Cryptography (TQC’11)* (2011) pp. 35–46, arXiv:0909.4673.
- [25] M. J. Hoban, J. J. Wallman, H. Anwar, N. Usher, R. Raussendorf, and D. E. Browne, “Measurement-Based Classical Computation,” *Physical Review Letters* **112**, 140505 (2014), arXiv:1304.2667.
- [26] C. Dawson, H. Haselgrove, A. Hines, D. Mortimer, M. Nielsen, and T. Osborne, “Quantum computing and polynomial equations over the finite field  $Z_2$ ,” *Quantum Inf. Comput.* **5**, 102–112 (2005), quant-ph/0408129.
- [27] T. Rudolph, “A simple encoding of a quantum circuit amplitude as a matrix permanent,” *Phys. Rev. A* **80**, 054302 (2009), arXiv:0909.3005.
- [28] S. Fenner, L. Fortnow, and S. Kurtz, “Gap-definable counting classes,” *J. Comput. Syst. Sci.* **48**, 116–148 (1994).
- [29] S. Aaronson, “A linear-optical proof that the permanent is #P-hard,” *Proc. Roy. Soc. Ser. A* **467**, 3393–3405 (2011), arXiv:1109.1674.

## Appendix A: The Ising model and IQP circuits

Consider the general Ising model:

$$H_I = \sum_{i < j} w_{ij} z_i z_j + \sum_i v_i z_i$$

where  $z \in \{-1, +1\}^n$ ,  $i, j$  label vertices in the complete graph,  $w_{ij}$  is a weight assigned to the edge  $(i, j)$  and  $v_i$  is a weight assigned to the vertex  $i$ . We want to show that the partition function of  $H_I$  evaluated at  $e^{i\theta}$ ,  $Z = \text{Tr}[e^{i\theta H_I}]$ , is proportional to an amplitude of an  $n$ -qubit IQP circuit.

Imagine we have an IQP circuit constructed from the gates  $e^{i\theta X \otimes X}$  and  $e^{i\theta X}$ , and consider the circuit amplitude

$$\langle 0 |^{\otimes n} e^{i\theta(\sum_{i < j} w_{ij} X_i X_j + \sum_i v_i X_i)} | 0 \rangle^{\otimes n}$$

where  $X$  is a Pauli- $X$  operator. Then

$$\begin{aligned}
\langle 0|^{\otimes n} e^{i\theta(\sum_{i<j} w_{ij} X_i X_j + \sum_i v_i X_i)} |0\rangle^{\otimes n} &= \langle 0|^{\otimes n} H^{\otimes n} e^{i\theta(\sum_{i<j} w_{ij} Z_i Z_j + \sum_i v_i Z_i)} H^{\otimes n} |0\rangle^{\otimes n} \\
&= \frac{1}{2^n} \sum_{x,y \in \{0,1\}^n} \langle y | e^{i\theta(\sum_{i<j} w_{ij} Z_i Z_j + \sum_i v_i Z_i)} |x\rangle \\
&= \frac{1}{2^n} \sum_{x \in \{0,1\}^n} e^{i\theta(\sum_{i<j} w_{ij} (-1)^{x_i x_j} + \sum_i v_i (-1)^{x_i})} \\
&= \frac{1}{2^n} \sum_{z \in \{-1,1\}^n} e^{i\theta(\sum_{i<j} w_{ij} z_i z_j + \sum_i v_i z_i)} \\
&= \frac{1}{2^n} \text{Tr}[e^{i\theta H_I}].
\end{aligned}$$

If we consider IQP circuits constructed from the gates  $e^{i\frac{\pi}{t} X \otimes X}$  and  $e^{i\frac{\pi}{t} X}$ , then it is clear that amplitudes of such circuits give rise to partition functions of the form  $Z(\omega) = \text{Tr}[\omega^{H_I}]$  where  $\omega = e^{i\frac{\pi}{t}}$  and  $w_{ij}, v_i \in \{0, \dots, 2t-1\}$ .

See Appendix E below for an alternative connection between IQP and the Ising model using a different gate set.

### Appendix B: Gaps of degree-3 polynomials

We can also use IQP circuit amplitudes to express the gap of any degree-3 polynomial over  $\mathbb{F}_2$  (with no constant term),  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ :

$$f(x) = \sum_{i,j,k} \alpha_{ijk} x_i x_j x_k + \sum_{i,j} \beta_{ij} x_i x_j + \sum_i \gamma_i x_i,$$

where  $\alpha_{ijk}, \beta_{ij}, \gamma_i \in \{0, 1\}$ .

This can be seen by noting the action of the Z, CZ, and CCZ gates on the computational basis states. Given a single qubit in a computational basis state,  $|x_1\rangle$ , the Pauli-Z gate has action  $Z|x_1\rangle = (-1)^{x_1}|x_1\rangle$ . Similarly we see  $CZ_{12}|x_1 x_2\rangle = (-1)^{x_1 x_2}|x_1 x_2\rangle$  and  $CCZ_{123}|x_1 x_2 x_3\rangle = (-1)^{x_1 x_2 x_3}|x_1 x_2 x_3\rangle$ . Then, in order to generate the phase  $(-1)^{f(x)}|x\rangle$ , we simply apply for each non-zero  $\alpha_{ijk}, \beta_{ij}, \gamma_i$  in  $f$  a corresponding CCZ, CZ, or Z gate to the  $n$ -qubit computational basis state,  $|x\rangle$ . Let  $\tilde{\mathcal{C}}_f$  denote this circuit and  $\mathcal{C}_f = H^{\otimes n} \tilde{\mathcal{C}}_f H^{\otimes n}$  denote its Hadamard transform, which is an IQP circuit.

We want to show that  $\langle 0|^{\otimes n} \mathcal{C}_f |0\rangle^{\otimes n}$  is proportional to the gap of  $f$ . This is straightforward:

$$\langle 0|^{\otimes n} H^{\otimes n} \tilde{\mathcal{C}}_f H^{\otimes n} |0\rangle^{\otimes n} = \frac{1}{2^n} \sum_{x,y \in \{0,1\}^n} \langle y | \tilde{\mathcal{C}}_f |x\rangle = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} = \frac{1}{2^n} \text{gap}(f),$$

recalling that  $\text{gap}(f)$  is defined as  $\text{gap}(f) := |\{x : f(x) = 0\}| - |\{x : f(x) = 1\}|$ . Note that similar ideas were previously used in [26, 27] for different classes of circuits.

### Appendix C: The Hadamard gadget

Arbitrary Hadamard operations can be implemented in IQP using post-selection [3] – something that cannot be done efficiently in physical systems but nonetheless turns out to be useful as a mathematical tool. We perform Hadamard gates by introducing *Hadamard gadgets* to our IQP circuits. Imagine that we have an  $n$ -qubit state  $|\psi\rangle$  upon which we wish to apply a Hadamard gate to qubit  $j$ , and denote this by  $H_j|\psi\rangle$ . We now add an extra qubit  $e$  initialised to the state  $H_e|0\rangle_e$ , recalling that all lines of an IQP circuit in the Z basis begin and end with a H. Then we perform some gates to produce the state  $H_j CZ_{je} H_e |\psi\rangle |0\rangle_e$  and measure qubit  $j$  in the computational basis, post-selecting on the 0 outcome. We note that the gate  $H_j$  must be present as this is an IQP circuit. It is easily checked that the effect of this protocol is to teleport the state of qubit  $j$  to qubit  $e$  while performing a Hadamard operation on this qubit. If we had desired to perform more gates on qubit  $j$  we now perform them on qubit  $e$ .

There are a few important things to note about the Hadamard gadget. Firstly for each ‘‘intermediate’’ Hadamard to be performed there is a cost of 1 qubit, 1 CZ gate, and 1 qubit of post-selection. Secondly, if we did not perform post-selection there would be a probability of 1/2 for this procedure to implement a Hadamard gate, so if a circuit had  $m$  intermediate Hadamard gates the success probability would be  $1/2^m$ , making this procedure impractical.

Finally, it is clear that the Hadamard gadget *preserves multiplicative approximations of circuit amplitudes*. Imagine we had an  $n$ -qubit quantum circuit  $\mathcal{U}$  which is expressed in terms of gates that are diagonal in the Z basis and  $m$  intermediate  $H$  gates.

Then using the Hadamard gadget  $m$  times we get the IQP circuit  $\mathcal{C}$  such that  $\langle 0 |^{\otimes n} \mathcal{U} | 0 \rangle^{\otimes n} = 2^{m/2} \langle 0 |^{\otimes (n+m)} \mathcal{C} | 0 \rangle^{\otimes (n+m)}$ . Furthermore, if we had a multiplicative approximation,  $A$ , to this latter quantity within multiplicative error  $\gamma$  then  $|A - \langle 0 |^{\otimes (n+m)} \mathcal{C} | 0 \rangle^{\otimes (n+m)}| \leq \gamma \langle 0 |^{\otimes (n+m)} \mathcal{C} | 0 \rangle^{\otimes (n+m)}$ . Then it is clear that  $A' = 2^{m/2} A$  also provides a multiplicative approximation to  $\langle 0 |^{\otimes n} \mathcal{U} | 0 \rangle^{\otimes n}$  within  $\gamma$ .

#### Appendix D: Hardness of multiplicative approximations

We now show that approximating the squared gap of a degree-3 polynomial over  $\mathbb{F}_2$  is GapP-complete, even up to constant multiplicative error. GapP is a complexity class containing #P [28]. The majority of the proof is very similar to arguments used in other settings, such as for the permanent [29] or the Ising model [11]. However, we believe that the connection to IQP again somewhat simplifies the argument. A similar argument could be used to prove hardness of multiplicative approximation for the Ising model; we omit this as it was shown in [11].

**Proposition 10.** *Let  $f$  be a degree-3 polynomial over  $\mathbb{F}_2$ . The problem of producing an estimate  $\tilde{z}$  such that  $||\text{ngap}(f)| - \tilde{z}| \leq \epsilon |\text{ngap}(f)|$ , for any  $\epsilon < 1/2$ , is GapP-complete.*

*Proof.* Let  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  be a function described by a polynomial-size classical circuit  $C$ . By definition, computing  $\text{ngap}(f)$  for arbitrary functions  $f$  is GapP-complete [28]. We now show that this problem can be reduced to approximating  $|\text{ngap}(f)|$  for degree-3 polynomials  $f$ , up to sufficiently high precision.

The first step is to show a connection between a ‘‘shifted’’ version of the general problem and IQP. For any  $m = O(n)$  and for any  $c \in [-1, 1]$  such that  $c$  is an integer multiple of  $2^{1-m}$ , there is a polynomial-size classical circuit  $D_c$  on  $m$  input bits computing a function  $g$  such that  $\text{ngap}(g) = -c$ . By adding a control bit which determines whether to execute  $C$  or  $D_c$ , we obtain a polynomial-size classical circuit on  $n + m + 1$  bits computing a function  $f_c$  such that  $\text{ngap}(f_c) = \frac{1}{2}(\text{ngap}(f) - c)$ .

By universality of the Toffoli gate for reversible classical computation, we can write down a quantum circuit  $Q_c$  consisting of Toffoli, CNOT and X gates such that  $Q_c$  performs the map

$$|0\rangle^{\otimes a} |x\rangle |0\rangle \mapsto |0\rangle^{\otimes a} |x\rangle |f_c(x)\rangle,$$

where the first register contains  $a$  ancilla qubits, with  $a = \text{poly}(n)$ . Let the circuit  $Q'_c$  be defined as follows: start in the state  $|0\rangle^{\otimes a} |0\rangle^{\otimes (n+m+1)} |0\rangle$ ; apply an X gate to the third register; apply Hadamard gates to the second and third registers; execute  $Q_c$ ; apply Hadamard and X gates to the third register; and finally apply Hadamard gates to the second register. The state near the end of this circuit, just before the last Hadamard gates are applied, is

$$|0\rangle^{\otimes a} \left( \frac{1}{\sqrt{2^{n+m+1}}} \sum_{x \in \{0,1\}^{n+m+1}} (-1)^{f_c(x)} |x\rangle \right) |0\rangle.$$

Therefore,  $\langle 0 | Q'_c | 0 \rangle = \text{ngap}(f_c) = \frac{1}{2}(\text{ngap}(f) - c)$ , where we henceforth abbreviate  $|0\rangle^{\otimes n} = |0\rangle$  for any integer  $n$ . As discussed in Appendix C, there exists an IQP circuit  $Q''_c$  such that  $\langle 0 | Q''_c | 0 \rangle = \alpha \langle 0 | Q'_c | 0 \rangle$  for some easily computed real number  $\alpha$ . An approximation of  $\langle 0 | Q''_c | 0 \rangle$  up to multiplicative error  $\epsilon$  therefore implies an approximation of  $\langle 0 | Q'_c | 0 \rangle$  up to the same error.

Assume we have a deterministic algorithm  $\mathcal{A}$  such that, on taking as input an arbitrary IQP circuit  $C$ ,  $\mathcal{A}$  outputs  $\tilde{z}$  such that  $||\langle 0 | C | 0 \rangle - \tilde{z}| \leq \epsilon |\langle 0 | C | 0 \rangle|$ , for some  $\epsilon < 1$ . We will use  $\mathcal{A}$  to output estimates  $\tilde{d}$  such that  $|\tilde{d} - |\text{ngap}(f) - c|| \leq \epsilon |\text{ngap}(f) - c|$ . Observe that  $\mathcal{A}$  can be used to distinguish with certainty between the two cases that  $c = \text{ngap}(f)$ , and  $c \neq \text{ngap}(f)$ . So if we can guess  $c = \text{ngap}(f)$ ,  $\mathcal{A}$  certifies that our guess is correct.

We now perform the following sequence of steps to compute  $\text{ngap}(f)$ . At each step, we have a guess  $c$  for  $\text{ngap}(f)$ , starting with  $c = 0$ . We use  $\mathcal{A}$  to output an estimate  $\tilde{d}$  of  $|\text{ngap}(f) - c|$ . We then consider the two possibilities  $c \pm \tilde{d}$ . We apply  $\mathcal{A}$  to estimate  $|\text{ngap}(f) - (c \pm \tilde{d})|$ , obtaining outputs  $\tilde{d}^+$ ,  $\tilde{d}^-$ , and choose as our new value of  $c$  either  $c + \tilde{d}$  (if  $\tilde{d}^+ \leq \tilde{d}^-$ ) or otherwise  $c - \tilde{d}$ . We then repeat until  $c = \text{ngap}(f)$ .

Write  $\tilde{d} = (1 + \gamma)|c - \text{ngap}(f)|$ , for some  $\gamma$  such that  $|\gamma| \leq \epsilon$ . If  $c < \text{ngap}(f)$ , we have

$$|c + \tilde{d} - \text{ngap}(f)| = |c + (1 + \gamma)(\text{ngap}(f) - c) - \text{ngap}(f)| = |\gamma(\text{ngap}(f) - c)| \leq \epsilon |c - \text{ngap}(f)|,$$

so if we choose  $c + \tilde{d}$  for the next step, the distance between  $c$  and  $\text{ngap}(f)$  is reduced by multiplying it by a factor smaller than  $\epsilon$ . A similar calculation holds in the case  $c > \text{ngap}(f)$  if we choose  $c - \tilde{d}$  for the next step. As  $\text{ngap}(f)$  is an integer multiple of  $2^{-n}$ , if we pick the correct choice at each step, we need only  $O(n)$  steps to obtain  $c = \text{ngap}(f)$ . We now show that, for sufficiently small  $\epsilon$ , the correct choice is indeed picked.

Assuming  $c < \text{ngap}(f)$ , we would like  $\tilde{d}^+ < \tilde{d}^-$ . This will hold if

$$\begin{aligned} (1 + \epsilon)|c + \tilde{d} - \text{ngap}(f)| &< (1 - \epsilon)|c - \tilde{d} - \text{ngap}(f)| \\ \Leftrightarrow (1 + \epsilon)|c + (1 + \gamma)|c - \text{ngap}(f)| - \text{ngap}(f)| &< (1 - \epsilon)|c - (1 + \gamma)|c - \text{ngap}(f)| - \text{ngap}(f)| \\ \Leftrightarrow (1 + \epsilon)|\gamma(\text{ngap}(f) - c)| &< (1 - \epsilon)|(2 + \gamma)(\text{ngap}(f) - c)| \\ \Leftrightarrow (1 + \epsilon)|\gamma| &< (1 - \epsilon)|2 + \gamma|. \end{aligned}$$

This inequality is easily seen to hold for all  $\gamma$  such that  $|\gamma| \leq \epsilon < 1/2$ . The case  $c > \text{ngap}(f)$  is similar.

One omission in this argument is that we can only approximate  $|\text{ngap}(f) - c|$  for values  $c$  which are integer multiples of  $2^{1-m}$ , but depending on the estimates output by  $\mathcal{A}$ , we may want to use values of  $c$  which are not integer multiples. In this case, we simply truncate  $c$  to the nearest integer multiple of  $2^{1-m}$ . For sufficiently large  $m = O(n)$ , the additional error introduced by this truncation is negligibly small.  $\square$

This immediately implies hardness of multiplicatively approximating  $\text{ngap}(f)^2$  to the same accuracy, because

$$|\text{ngap}(f)^2 - \tilde{z}^2| \leq \epsilon \text{ngap}(f)^2 \Rightarrow \left| |\text{ngap}(f)| - \tilde{z} \right| (|\text{ngap}(f)| + \tilde{z}) \leq \epsilon \text{ngap}(f)^2 \Rightarrow \left| |\text{ngap}(f)| - \tilde{z} \right| \leq \epsilon |\text{ngap}(f)|.$$

### Appendix E: Proofs of anticoncentration bounds

Here we prove the lemmas required for the anticoncentration bounds claimed in the main body of the paper. Our main technical tool will be the following result:

**Lemma 11.** *Let  $\omega = e^{2\pi i/r}$ ,  $\eta = e^{2\pi i/s}$  for some integers  $r$  and  $s$  such that  $r, s \geq 2$ , and if  $s = 2$  then  $r = 2$ . Fix integer  $n$ . For pairs of integers  $i < j$  between 1 and  $n$ , let  $\alpha_{ij}$  be picked uniformly at random from  $\{0, \dots, r-1\}$ . For  $k \in \{1, \dots, n\}$ , let  $\beta_k$  be picked uniformly at random from  $\{0, \dots, s-1\}$ . Then*

$$\sum_{w, x, y, z \in \{0, 1\}^n} \left| \mathbb{E}_\alpha \left[ \omega^{\sum_{i < j} \alpha_{ij} (w_i w_j + x_i x_j - y_i y_j - z_i z_j)} \right] \mathbb{E}_\beta \left[ \eta^{\sum_k \beta_k (w_k + x_k - y_k - z_k)} \right] \right| \leq 3 \cdot 2^{2n}.$$

*Proof.* Each term in the sum is clearly upper-bounded by 1. Our goal will be to show that many terms in the sum are actually zero, using the identities  $\mathbb{E}_\alpha[\omega^{\alpha t}] = 0$  for any integer  $t \not\equiv 0 \pmod r$ , and similarly  $\mathbb{E}_\beta[\eta^{\beta t}] = 0$  for  $t \not\equiv 0 \pmod s$ . First,

$$\mathbb{E}_\beta \left[ \eta^{\sum_k \beta_k (w_k + x_k - y_k - z_k)} \right] = \prod_{k=1}^n \mathbb{E}_{\beta_k} \left[ \eta^{\beta_k (w_k + x_k - y_k - z_k)} \right],$$

which equals 0 unless  $z \equiv w + x - y \pmod s$ . If this is the case, it also holds that  $z \equiv w + x - y \pmod r$ . To see this, observe that for  $s \geq 3$ , as  $w + x - y \in \{-1, 0, 1, 2\}^n$ ,  $z \equiv w + x - y \pmod s$  implies that in fact  $z = w + x - y$ . If  $s = 2$ , then  $r = s$  by assumption and the claim is immediate. Hence

$$\begin{aligned} \mathbb{E}_\alpha \left[ \omega^{\sum_{i < j} \alpha_{ij} (w_i w_j + x_i x_j - y_i y_j - z_i z_j)} \right] &= \mathbb{E}_\alpha \left[ \omega^{\sum_{i < j} \alpha_{ij} (w_i w_j + x_i x_j - y_i y_j - (w + x - y)_i (w + x - y)_j)} \right] \\ &= \prod_{i < j} \mathbb{E}_{\alpha_{ij}} \left[ \omega^{\alpha_{ij} (w_i (y_j - x_j) + x_i (y_j - w_j) + y_i (w_j + x_j) - 2y_i y_j)} \right]. \end{aligned}$$

We claim that if  $w, x$  and  $y$  are all distinct, then there exists a pair  $(i, j)$  (also distinct) such that

$$w_i (y_j - x_j) + x_i (y_j - w_j) + y_i (w_j + x_j) \not\equiv 2y_i y_j \pmod r, \tag{E1}$$

which implies that the product evaluates to 0. Since  $w \neq x$ , let  $i$  be such that  $w_i \neq x_i$ . We may assume without loss of generality that  $w_i = 0$  and  $x_i = 1$ , since the other case just relabels  $w$  for  $x$  and vice versa. Therefore it remains to find a  $j > i$  such that

$$y_j - w_j + y_i (w_j + x_j) \not\equiv 2y_i y_j \pmod r. \tag{E2}$$

In fact, finding a  $j < i$  is just as good, because the expression (E1) is also symmetric under exchange of  $i$  and  $j$ . If  $y_i = 0$  then since  $y \neq w$ , there must exist some  $j$  such that  $y_j \neq w_j$ . Moreover,  $j \neq i$  because  $y_j$  cannot differ from both  $w_i$  and  $x_i$  (which we have already seen are different). So  $y_j = 0$ ,  $w_j = 1$  (or vice versa), they cannot be equivalent modulo  $r$ , and (E2) holds in this case. On the other hand, if  $y_i = 1$  then since  $y \neq x$ , there must be some  $j$  such that  $y_j \neq x_j$ , and  $j \neq i$  because we already chose  $x_i = 1$ . So  $y_j = 0$ ,  $x_j = 1$  (or vice versa), they are not equivalent modulo  $r$ , and (E2) holds in this case too.

We have therefore shown that each term in the sum in the statement of the lemma evaluates to 0 unless both  $z \equiv w + x - y \pmod s$ , and also  $w = x$ ,  $w = y$  or  $x = y$ . The number of ways of choosing strings  $w, x, y, z$  to comply with this constraint is at most  $3 \cdot 2^{2n}$ . This completes the proof.  $\square$

We first apply Lemma 11 to random degree-3 polynomials. A degree-3 polynomial  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  can be picked uniformly at random in a straightforward manner. Assuming that  $f(0^n) = 0$  without loss of generality and writing

$$f(x) = \sum_{i,j,k} \alpha_{ijk} x_i x_j x_k + \sum_{i,j} \beta_{ij} x_i x_j + \sum_i \gamma_i x_i,$$

where  $\alpha_{ijk}, \beta_{ij}, \gamma_i \in \{0, 1\}$ , we simply pick each of these coefficients uniformly at random. Observe that the random  $\gamma_i$  coefficients correspond to applying X gates to a random subset of the qubits. We now show that, for each possible choice of the  $\alpha_{ijk}$  coefficients, the expected value of  $\text{ngap}(f)^4$  (taken over the  $\beta, \gamma$  coefficients) is low.

**Lemma 12.** *Let  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  be a polynomial over  $\mathbb{F}_2$  whose degree  $\leq 2$  part is uniformly random. Then  $\mathbb{E}_f[\text{ngap}(f)^4] \leq 3 \cdot 2^{-2n}$ .*

*Proof.* We have

$$\begin{aligned} \mathbb{E}_f [\text{ngap}(f)^4] &= \mathbb{E}_f \left[ \left( \mathbb{E}_{x \in \{0,1\}^n} (-1)^{f(x)} \right)^4 \right] \\ &= \mathbb{E}_f \left[ \mathbb{E}_{w,x,y,z \in \{0,1\}^n} (-1)^{f(w)+f(x)+f(y)+f(z)} \right] \\ &= \mathbb{E}_{w,x,y,z \in \{0,1\}^n} \mathbb{E}_f \left[ (-1)^{f(w)+f(x)+f(y)+f(z)} \right]. \end{aligned}$$

Let  $f_{\leq 2}$  and  $f_{>2}$  be the parts of  $f$  of degree  $\leq 2$  and degree  $> 2$ , respectively. Then

$$\mathbb{E}_f \left[ (-1)^{f(w)+f(x)+f(y)+f(z)} \right] = (-1)^{f_{>2}(w)+f_{>2}(x)+f_{>2}(y)+f_{>2}(z)} \mathbb{E}_{f_{\leq 2}} \left[ (-1)^{f_{\leq 2}(w)+f_{\leq 2}(x)+f_{\leq 2}(y)+f_{\leq 2}(z)} \right].$$

Expanding  $f_{\leq 2}(x) = \sum_{i,j} \beta_{ij} x_i x_j + \sum_k \gamma_k x_k$ , the expectation in this expression can be written as

$$\mathbb{E}_{\beta,\gamma} \left[ (-1)^{\sum_{i,j} \beta_{ij} (w_i w_j + x_i x_j + y_i y_j + z_i z_j) + \sum_k \gamma_k (w_k + x_k + y_k + z_k)} \right].$$

So

$$\begin{aligned} \mathbb{E}_f [\text{ngap}(f)^4] &\leq \frac{1}{2^{4n}} \sum_{w,x,y,z \in \{0,1\}^n} \left| \mathbb{E}_{\beta} \left[ (-1)^{\sum_{i,j} \beta_{ij} (w_i w_j + x_i x_j + y_i y_j + z_i z_j)} \right] \mathbb{E}_{\gamma} \left[ (-1)^{\sum_k \gamma_k (w_k + x_k + y_k + z_k)} \right] \right| \\ &\leq 3 \cdot 2^{-2n} \end{aligned}$$

by Lemma 11. □

Next we consider the partition function (1) of the Ising model for random weights  $w_{ij}, v_k \in \{0, \dots, 15\}$ , evaluated at  $\omega = e^{i\pi/8}$ . To make contact with the previous analysis, and with IQP circuits, we rewrite the expression (1) for  $Z(\omega)$  as

$$\begin{aligned} Z(\omega) &= \sum_{x \in \{0,1\}^n} \omega^{\sum_{i<j} w_{ij} (1-2x_i)(1-2x_j) + \sum_{k=1}^n v_k (1-2x_k)} \\ &= \omega^{\sum_{i<j} w_{ij} + \sum_{k=1}^n v_k} \sum_{x \in \{0,1\}^n} \omega^{4 \sum_{i<j} w_{ij} x_i x_j - 2 \sum_{k=1}^n v'_k x_k}, \end{aligned}$$

where  $v'_k = v_k + \sum_{j \neq k} w_{jk}$ . Observe that  $v'_k$  is uniformly random mod 16. Up to an easily computed phase, we can write

$$Z(\omega) = \sum_{x \in \{0,1\}^n} i^{\sum_{i<j} w_{ij} x_i x_j} e^{-(\pi i/4)(\sum_{k=1}^n v'_k x_k)}.$$

Then it is clear that  $Z(\omega)/2^n = \langle 0 | \otimes^n \mathcal{C} | 0 \rangle \otimes^n$  for an IQP circuit  $\mathcal{C}$  consisting of gates from the set  $\{\text{diag}(1, 1, 1, i), \text{diag}(1, e^{i\pi/4})\}$ . Further, because of the uniformly random choice of  $w$  and  $v'$ , we can consider such a circuit as being formed by picking a random circuit using this gate set, then following it by a random choice of X gates. It follows immediately from Lemma 11 that  $\mathbb{E}_{w,v'} [|Z(\omega)|^4] \leq 3 \cdot 2^{2n}$ .