Classical Simulation of Yang-Baxter Gates

Gorjan Alagic[1], Aniruddha Bapat[1], and Stephen Jordan[2]

1   Institute for Quantum Information and Matter
     California Institute of Technology
     Pasadena, CA
2   National Institute of Standards and Technology
     Gaithersburg, MD

──── **Abstract** ────

A unitary operator that satisfies the constant Yang-Baxter equation immediately yields a unitary representation of the braid group B_n for every $n \geq 2$. If we view such an operator as a quantum-computational gate, then topological braiding corresponds to a quantum circuit. A basic question is when such a representation affords universal quantum computation. In this work, we show how to classically simulate these circuits when the gate in question belongs to certain families of solutions to the Yang-Baxter equation. These include all of the qubit (i.e., $d = 2$) solutions, and some simple families that include solutions for arbitrary $d \geq 2$. Our main tool is a probabilistic classical algorithm for efficient simulation of a more general class of quantum circuits. This algorithm may be of use outside the present setting.

1   Introduction

The Yang-Baxter equation, named after C. N. Yang and R. J. Baxter, appears in a number of areas of mathematics and physics. Yang encountered the equation while working on two-dimensional quantum field theory, while Baxter applied it to exactly solvable models in statistical mechanics [2]. An accessible review of some of the many applications of the Yang-Baxter equation can be found in [21]. In this work, we will consider what is typically called the constant quantum Yang-Baxter equation, and is defined as follows. Let V be a finite-dimensional complex Hilbert space and R a linear operator on $V \otimes V$. Then R satisfies the quantum Yang-Baxter equation (YBE) if

$$(R \otimes I)(I \otimes R)(R \otimes I) = (I \otimes R)(R \otimes I)(I \otimes R),$$

where I denotes the identity operator on V. In this case, we say that R is a Yang-Baxter operator. The YBE bears a close resemblance to the relation

$$\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$$

of the braid group B_n. Indeed, a Yang-Baxter operator naturally gives the space $V^{\otimes n}$ the structure of a representation $\rho_{(R,n)}$ of B_n. Turaev showed that if R also satisfies the so-called Markov property, then it corresponds to an invariant of links [24]. The invariant is given by the (appropriately scaled) trace of $\rho_{(R,n)}$, evaluated at any braid whose trace closure is equal to the link. More generally, one can derive a link invariant from the trace of any representation of B_n which satisfies the Markov property. This is the case for the famous

Jones representation and the corresponding Jones Polynomial invariant [15]. Freedman, Kitaev, Larsen and Wang [8, 9, 10] showed that the Jones representation has significant meaning in quantum computation. Informally speaking, the Jones representation provides a functionality-preserving "dictionary" between quantum circuits and braids. One consequence of these results is that additively approximating the Jones Polynomial is a universal problem for quantum computation. It also appears that this dictionary could correspond to a physically plausible implementation of quantum computers by means of exotic particles called non-abelian anyons [22]. One downside of the Jones representation in this context is that topological locality of braiding does not translate naturally into tensor-product locality of the corresponding quantum circuit. In particular, it is not the case that braiding two adjacent strands correponds to applying a Yang-Baxter operator on the space of two adjacent qubits. One might hope that the Jones representation could be made to look this way, e.g., by changing bases or manipulating the multiplicities of its irreducible summands. However, Rowell and Wang recently showed that this is impossible unless the Jones representation in question[1] is in fact *not* quantum-universal (see Corollary 4.2 in [23].)

Alternatively, one may ask if there exist other representations of the braid groups with the desired local structure and which exhibit computational universality. This amounts to finding unitary solutions to the YBE and determining if they are universal gates. In this work, we investigate low-dimensional solutions with this motivation in mind. All of the qubit (i.e., $d = \dim V = 2$) solutions to the YBE were found by Hietarinta [12]; the unitary ones among those were identified by Dye [5]. It was previously known that, when their eigenvalues are roots of unity, these solutions yield braid group representations with finite image [7, 6]. We show how to classically approximate the matrix entries of any quantum circuit constructed from a particular kind of two-qudit gate. Most of the qubit solutions to the YBE, as well as some solution families of arbitrary dimension, are special cases of this gate. For the remaining qubit solutions, we give a different result: how to classically simulate a quantum computation that begins in any product state, and ends with a measurement of an observable on logarithmically many qubits. This is typically considered sufficient to rule out quantum universality. However, some caution is called for: there are gate sets which are known to be classically simulable in this sense but become hard to simulate when one is allowed to measure all the output qubits in the computational basis [17, 3].

We remark that, as pointed out by Lomonaco and Kauffman [18], some qubit solutions to the YBE are entangling gates, and any entangling gate together with arbitrary single-qubit gates is universal [4]. However, in that case we are no longer computing with representations of the braid group. Indeed, a primary motivation for the topological approach to quantum computation is to rely on the topological stability of braiding for fault-tolerance. Applying single-qubit gates fault-tolerantly as part of this approach would require additional ideas. For this reason, we restrict ourselves to just one gate, which acts on two qubits and is a solution to the YBE. Some classes of entangling gates that have previously been shown to be classically simulatable are given in [16, 11].

2   Preliminaries

2.1   Gates, circuits, and universality

We briefly review basic notions about quantum gates, circuits, and computational universality. For more details, we refer the reader to the text of Nielsen and Chuang [20]. Given an integer

---

[1] Recall that, just like the Jones polynomial, the Jones representation has a parameter (in addition to n) which is typically a root of unity. Quantum universality holds for most but not all values of this parameter.

$d \geq 1$, let $[d] = \{0, 1, \ldots, d-1\}$. Let $V = \mathbb{C}[d]$ be a d-dimensional complex Hilbert space with distinguished orthonormal basis $\{|i\rangle : i \in [d]\}$. We refer to copies of $[d]$ as dits and copies of V as qudits. For any k and any $x \in [d]^k$, set $|x\rangle = |x_1\rangle \otimes |x_2\rangle \otimes \cdots \otimes |x_k\rangle$. The space $V^{\otimes k}$ has a preferred basis $\{|x\rangle : x \in [d]^k\}$, which we will call the computational basis. A unitary operator on $V^{\otimes k}$ is called a k-qudit gate.

Let $\mathcal{R}$ be a set of gates which act on k or fewer qudits. Fix $n > 0$ and, for each l-qudit gate $R \in \mathcal{R}$, define $R_j \in \mathrm{U}(V^{\otimes n})$ to be the operator that applies R to qudits $j, \ldots, j+l$ and the identity operator I to the rest. Define $\mathcal{R}^{(n)}$ to be the set of all R_j, for every $R \in \mathcal{R}$ and every valid index j. An n-qudit quantum circuit over the gate set $\mathcal{R}$ (or $\mathcal{R}$-circuit for short) is a finite sequence

$$C = (U_1, U_2, \ldots, U_m)$$

where for each i, $U_i \in \mathcal{R}^{(n)}$ or $U_i^{-1} \in \mathcal{R}^{(n)}$. We will sometimes denote the number of gates in the circuit C by $|C| = m$. The circuit defines an operator

$$C = U_m \cdot U_{m-1} \cdots U_1 \in \mathrm{U}(V^{\otimes n}).$$

Note that we have overloaded notation so that C refers to both the sequence of gates and the operator implemented by their composition. Pictorially, an $\mathcal{R}$-circuit is represented by a diagram like the following, where each wire corresponds to one qudit.
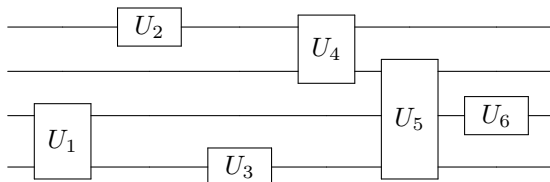
For pictorial convenience, the gates shown in the figure only act on nearest neighbors. While the nearest-neighbor condition is needed for certain other types of circuits to be classically simulatable (e.g. matchgates [16]), our results do not require it. We adopt here the common convention that circuits are applied from left to right (unfortunately, the opposite of the case for operators.) Of general interest are gate sets which allow for universal quantum computation.

▶ **Definition 1.** A gate set $\mathcal{R}$ is **universal** if there exists $N > 0$ such that N-qudit $\mathcal{R}$-circuits form a dense subset of $U(V^{\otimes N})$.

The Solovay-Kitaev theorem [20] tells us that, for universal $\mathcal{R}$, any unitary operator in $U(V^{\otimes N})$ can be approximated to precision ϵ with an N-qudit $\mathcal{R}$-circuit of length polylog($1/\epsilon$). Standard arguments also show that density can be extended from N to any $n \geq N$.

Quantum-computational power can also be defined in terms of complexity classes. The class that is typically associated with efficient quantum computation is called BQP, which stands for bounded-error quantum polynomial time. A drawback of BQP is the lack of known complete problems, i.e., problems which are both in BQP and at least as hard (under classical polynomial-time reduction) as any other problem in BQP. The classical analogue BPP (bounded-error probabilistic polynomial time) suffers from the same drawback. For this reason, we will work with promise versions of these two classes, i.e., PromiseBQP and PromiseBPP. We will not need the formal definitions of these classes (see, e.g., [14]). For us it will suffice to refer to the following.

▶ **Definition 2.** Given a set $\mathcal{R}$ of quantum gates, the problem $\mathcal{I}(\mathcal{R})$ is defined as follows. Given an n-qudit $\mathcal{R}$-circuit C and ditstrings x and y, as well as a promise that either $\langle x|C|y\rangle > 2/3$ or $\langle x|C|y\rangle < 1/3$, decide which is the case.

We may define PromiseBQP as the class of problems which reduce to $\mathcal{I}(\mathcal{R})$ for some universal set of quantum gates $\mathcal{R}$. Interestingly, there are gate sets $\mathcal{R}$ which are not universal in the density sense but for which $\mathcal{I}(\mathcal{R})$ is nonetheless PromiseBQP-hard; an example is $\mathcal{R} = \{\text{Hadamard}, \text{Toffoli}\}$. This gate set is dense over the special orthogonal group, but since the matrix entries are all real, it cannot be dense over the unitary group.

Later on, we will show that when $\mathcal{R}$ consists of a single gate which belongs to certain solution families of the Yang-Baxter equation, then $\mathcal{I}(\mathcal{R}) \in$ PromiseBPP. This means that $\mathcal{R}$ is not quantum universal under either of the above definitions, unless the widely believed conjecture that quantum computation is more powerful than classical computation is false.

2.2   Pauli group and Clifford group

Recall that the single-qubit Pauli operators are defined by

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Each Pauli operator is self-adjoint and unitary. In the n-qubit case, we set

$$X_j = I^{\otimes j-1} \otimes X \otimes I^{\otimes n-j}$$

and likewise for Y_j and Z_j. We define the n-qubit Pauli group $\mathcal{P}_n$ to be the group generated by $\{X_j, Y_j, Z_j : j = 1, \ldots, n\}$. An important property for us is that $\mathcal{P}_n$ spans the space of n-qubit Hermitian operators.

The Clifford group on n qubits is defined to be the normalizer of the Pauli group inside the unitary group, i.e.,

$$\mathcal{C}_n = \{U \in U(2^n) : UPU^\dagger \in \mathcal{P}_n \text{ for all } P \in \mathcal{P}_n\}.$$

By direct computation, it's easy to check that the following gates are elements of $\mathcal{C}_n$ for any $n \geq 2$:

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad P = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \qquad CNOT = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

It is a theorem (see [11]) that the above gates, when applied to arbitrary qubits or pairs of qubits, actually generate $\mathcal{C}_n$. We will thus call any circuit made up of these gates a Clifford circuit. Since $\mathcal{P}_n \subset \mathcal{C}_n$, we can also add the Pauli operators to this gate set for free. We remark that the conjugation action of a Clifford circuit on an element of $\mathcal{P}_n$ is easy to compute in a direct, gate-by-gate fashion. For details, see [11].

Due to the frequent appearance of $\mathcal{C}_n$ in various areas of quantum information, the computational power of Clifford circuits is well-studied. While $\mathcal{C}_n$ is finite and not universal, adding any gate outside $\mathcal{C}_n$ results in a universal set [19]. A thorough analysis of the computational power of Clifford circuits under various models is performed in [17].

2.3   Yang-Baxter operators and representations of the braid group

Let $V = \mathbb{C}[d]$ and $R \in U(V \otimes V)$. Then R satisfies the quantum Yang-Baxter equation (YBE) if

$$(R \otimes I)(I \otimes R)(R \otimes I) = (I \otimes R)(R \otimes I)(I \otimes R), \tag{1}$$

where I denotes the identity operator on V. In this case, we say that R is a Yang-Baxter operator. Let $T : |a \otimes b\rangle \mapsto |b \otimes a\rangle$ denote the swap operator on $V \otimes V$. By comparing circuit diagrams, it's not hard to see that R is a solution to (1) if and only if $S = RT$ is a solution to

$$S_{12}S_{13}S_{23} = S_{23}S_{13}S_{12}\,, \tag{2}$$

where

$$S_{12} = S \otimes I\,, \qquad S_{13} = (I \otimes T)(S \otimes I)(I \otimes T)\,, \qquad S_{23} = I \otimes S\,.$$

Equation (2) is sometimes called the algebraic Yang-Baxter equation.

Recall that the braid group B_n is a finitely generated group with generators $\sigma_1, \sigma_2, \cdots, \sigma_{n-1}$ and relations

$$\begin{aligned} \sigma_i\sigma_j &= \sigma_j\sigma_i & \forall\,|i - j| \geq 2 \\ \sigma_i\sigma_{i+1}\sigma_i &= \sigma_{i+1}\sigma_i\sigma_{i+1} & \forall\,i. \end{aligned}$$

In 1925 Artin proved that the abstract group defined above precisely captures the topological equivalence of braided strings [1]. Pictorially, braids are represented with a diagram; an example diagram for $\sigma_3^{-1}\sigma_2^{-1}\sigma_3\sigma_1^{-1}$ is shown below. We read such diagrams left-to-right, keeping the same convention as with circuits. The second generating relation of B_n is known as the Yang-Baxter relation. A solution $R \in \mathrm{U}(V \otimes V)$ of the Yang-Baxter equation yields a unitary representation $\rho_{(R,n)}$ of B_n on the space $V^{\otimes n}$ for every n. It is defined by

$$\rho_{(R,n)}(\sigma_i) = I^{\otimes(i-1)} \otimes R \otimes I^{\otimes n-i-1}\,.$$

The images of braids under $\rho_{(R,n)}$ are precisely the R-circuits on n qudits, where $d = \dim V$. For example, the braid $\sigma_3^{-1}\sigma_2^{-1}\sigma_3\sigma_1^{-1}$ and the corresponding R-circuit are shown below.
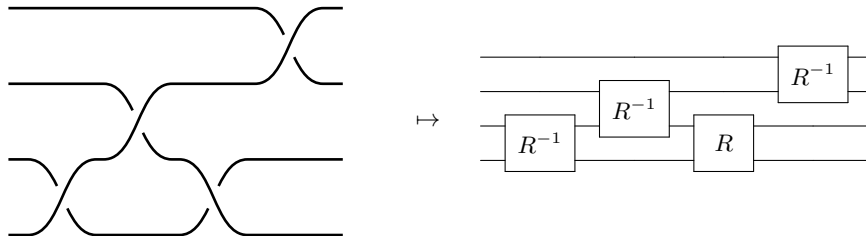
Under a plausible physical interpretation, a computation is performed by braiding particle-like excitations whose exchange statistics are described by R. If R is a universal gate, this model would result in universal topological computation. Such a model could provide a basis for a quantum computer architecture with inherent fault-tolerance [22].

3   Classical simulation of certain quantum circuits

In this section, we prove a general result about simulating certain quantum circuits with a classical probabilistic algorithm. We begin with two straightforward lemmas about classical sampling. (See A for proofs).

▶ **Lemma 3.** *Let $\{P_j\}_{j=1}^n$ be probability distributions on $[d]$ and let $P = \Pi_j P_j$ be the corresponding product distribution over $[d]^n$. Suppose that we can calculate $P_j(k)$ for every j and every k in total time* poly(n, d). *Then there's a classical probabilistic algorithm that runs in time* poly(n, d) *and samples from $[d]^n$ according to a probability distribution D such that $|P - D| \leq 1/2^{\mathrm{poly}(n)}$.*

We will also require the following Chernoff-Hoeffding bound for complex-valued random variables.

▶ **Lemma 4.** *Let $X_1, X_2, \ldots, X_n$ be independent complex-valued random variables with $\mathbb{E}[X_j] = \mu$ and $|X_j| \leq b$ for all j. Let $S = \sum_j X_j/n$. Then*

$$\Pr\left[|S - \mu| \geq \epsilon\right] \leq 4 \exp\left(-n\epsilon^2/8b^2\right) .$$

Let $\mathcal{S}_d$ denote the symmetric group, i.e., the group of permutations of d letters. We denote the action of $\pi \in \mathcal{S}_d$ on an integer $1 \leq j \leq d$ by πj.

▶ **Definition 5.** Let Q be an invertible $d \times d$ matrix over $\mathbb{C}$, and G a subgroup of $\mathcal{S}_d$. Define matrices A, B by setting $A_{ij} = |Q_{ij}|$ and $B_{ij} = |(Q^{-1})_{ij}|$. We say that Q satisfies property (G) if for every $\pi \in G$ and every k, l, we have $\sum_j A_{k,\pi j} B_{jl} \leq 1$.

If Q is unitary, then by Cauchy-Schwarz and the orthonormality of the rows of Q,

$$\sum_j A_{k,\pi j} B_{jl} \leq \left(\sum_j |A_{k,\pi j}|^2 \sum_i |B_{il}|^2\right)^{1/2} = 1.$$

It follows that unitary matrices satisfy property $(\mathcal{S}_d)$.

We are now ready to present the main classical simulation algorithm. When we refer to the matrix entries of operators in $\mathrm{GL}(\mathbb{C}[d]) \cong \mathrm{GL}_d(\mathbb{C})$, it will always be in the computational basis. We say that such an operator is computable if its entries can be computed exactly by a classical algorithm in $\mathrm{poly}(d)$ time. Recall that $T : a \otimes b \mapsto b \otimes a$ is the swap operator, and that for a subset S of a group G, $\langle S \rangle$ denotes the subgroup of G generated by S.

▶ **Theorem 6.** *Let $\mathcal{R} = \{R_1, R_2, \ldots, R_k\}$ be a set of unitary 2-qudit gates, each one a composition*

$$R_i = (Q \otimes Q)D_i P_i(C_i \otimes C_i)(Q \otimes Q)^{-1} \tag{3}$$

of computable, invertible operators. Suppose that for each i, D_i is a diagonal unitary, C_i is a $d \times d$ permutation matrix, and $P_i = I$ or $P_i = T$. Finally, let Q satisfy property (G) where $G = \langle \{C_i\}_{i=1}^k \rangle \leq \mathcal{S}_d$. Then there exists a classical probabilistic algorithm which, given an n-qudit $\mathcal{R}$-circuit U and strings $x, z \in [d]^n$ and $\epsilon > 0$, outputs a number r in time $\mathrm{poly}(n, |U|, 1/\epsilon)$ such that $|r - \langle x|U|z\rangle| < \epsilon$ except with probability exponentially small in n and $1/\epsilon$.

Proof. Set $S_i = D_i P_i(C_i \otimes C_i)$. If we expand each R_i-gate to turn U into a circuit made from S_i-gates and Q-gates, then all of the Q-gates except the initial and final ones are cancelled, as in the example below. We are thus left with a circuit of the form $Q^{\otimes n} V (Q^{-1})^{\otimes n}$ where V is an $\{S_i\}$-circuit. We remark that, in this expanded form, the entire circuit is not necessarily a proper quantum circuit, since Q might not be unitary. The circuit V, on the other hand, is quantum since all of its gates are unitary.

Before we proceed, note that a non-nearest-neighbor gate can be written as a nearest-neighbor gate conjugated with a swap gate. We depict our gates as acting on nearest neighbors for

convenience only, but this condition is not needed for the result to hold. The action of an S_i-gate on the j-th and $(j+1)$-st qudits of a computational basis state is simple to compute. The values of the two qudits are both in $[d]$ initially, and remain in $[d]$ after the action of C_i. Second, these new values are either swapped or left unchanged by P_i. Third, the D_i-gate adds an overall phase factor to the state. By composing these easily-computable actions, the action of V on a computational basis state can be computed in time polynomial in n, d, and $|V|$. Up to phases, this action consists of permuting the n qudits by some $\pi \in S_n$, and applying some bijection $f_j : [d] \to [d]$ to the initial value of the $\pi(j)$-th qudit. Each f_j is a composition of C_i-gates, in the order specified by V. Explicitly, for a basis state $|y\rangle = |y_1 y_2 \dots y_n\rangle$, we write

$$V|y\rangle = e^{i\phi(y)}|f_1 y_{\pi 1} \otimes f_2 y_{\pi 2} \otimes \cdots \otimes f_n y_{\pi n}\rangle,$$

where $\phi(y)$ is the overall phase resulting from the D_i-gates. For simplicity of notation, we denoted the image of k under the permutation π as πk, and wrote $f_j y_{\pi j}$ in place of $f_j(y_{\pi j})$.

Next we consider the matrix element

$$\langle x|U|z\rangle = \langle x|(Q)^{\otimes n}V(Q^{-1})^{\otimes n}|z\rangle = \sum_{y\in[d]^n} \langle x|(Q)^{\otimes n}V|y\rangle\langle y|(Q^{-1})^{\otimes n}|z\rangle$$

$$= \sum_{y\in[d]^n} e^{i\phi(y)} \prod_{j=1}^{n} \langle x_j|Q|f_j y_{\pi j}\rangle\langle y_j|Q^{-1}|z_j\rangle.$$

We expand the matrix elements of Q and Q^{-1} in terms of magnitudes and phases:

$$\langle r|Q|s\rangle = A(r,s)e^{i\alpha(r,s)}$$
$$\langle r|Q^{-1}|s\rangle = B(r,s)e^{i\beta(r,s)}$$

where A, B, α, β are real-valued and $r, s \in [d]$. Then

$$\langle x|(Q)^{\otimes n}V(Q^{-1})^{\otimes n}|z\rangle = \sum_{y\in[d]^n} e^{i\theta(y)} \prod_{j=1}^{n} A(x_j, f_j y_{\pi j})B(y_j, z_j)$$

$$= \sum_{y\in[d]^n} e^{i\theta(y)} \prod_{j=1}^{n} A(x_j, f_{\sigma j} y_j)B(y_j, z_j),$$

where $\sigma = \pi^{-1}$ and

$$\theta(y) = \phi(y) + \sum_{j=1}^{n} \big(\alpha(x_j, f_{\sigma j} y_j) + \beta(y_j, z_j)\big).$$

Now we introduce the following normalization factor:

$$\rho = \sum_{y\in[d]^n} \prod_{j=1}^{n} A(x_j, f_{\sigma j} y_j)B(y_j, z_j) = \prod_{j=1}^{n} \sum_{k\in[d]} A(x_j, f_{\sigma j}k)B(k, z_j).$$

This allows us to define a natural probability distribution over $[d]^n$ by

$$P(y) = \frac{1}{\rho} \prod_{j=1}^{n} A(x_j, f_{\sigma j} y_j)B(y_j, z_j),$$

which factorizes as $P(y) = \prod_{j=1}^{n} P_j(y_j)$, where

$$P_j(l) = \frac{A(x_j, f_{\sigma j}l)B(l, z_j)}{\sum_{k \in [d]} A(x_j, f_{\sigma j}k)B(k, z_j)} \, .$$

Note that ρ and all of the $P_j(l)$ can be computed in time linear in n and d. By Lemma 3, we can efficiently sample from $[d]^n$ according to P, with error exponentially small in n.

In order to estimate $\langle x|U|z \rangle$, sample repeatedly from this distribution, obtaining outcomes $\xi(j) \in [d]^n$ for $j \in \{1, 2, \ldots\}$ and output the average of the random variables $X_j := \rho \exp(i\theta(\xi(j)))$. Observe that, for each j,

$$\mathbb{E}[X_j] = \sum_{z \in [d]^n} \rho e^{i\theta(z)} P(z) = \langle x|U|z \rangle \, .$$

To control the absolute value, recall that $f_{\sigma j}$ is a composition of the permutation matrices C_i, and is thus an element of $\langle \{C_i\}_{i=1}^k \rangle \leq \mathcal{S}_d$. Since Q satisfies property $(\langle \{C_i\}_{i=1}^k \rangle)$, we have

$$|X_j|^2 = |\rho|^2 = \prod_{j=1}^{n} \left| \sum_{k \in [d]} A(x_j, f_{\sigma j}k)B(k, z_j) \right|^2 \leq \prod_{j=1}^{n} 1^2 \leq 1.$$

by Cauchy-Schwarz, for each j. Now set $S(r) = \sum_{j=1}^{r} X_j/r$. By Lemma 4, for $r \geq 8n/\epsilon^3$ we have

$$\Pr\left[|S(r) - \langle x|U|z \rangle| \geq \epsilon\right] \leq 4\exp(-r\epsilon^2/8) \leq 4\exp(-n/\epsilon) \, .$$

◀

An immediate corollary is that, for $\mathcal{R}$ as in the theorem, $\mathcal{I}(\mathcal{R})$ is in PromiseBPP. We will also need the following simple result about simulating circuits constructed from conjugated Clifford gates.

▶ **Theorem 7.** *Let $S \in \mathcal{C}_2$, and $R = (Q \otimes Q)S(Q \otimes Q)^\dagger$ where Q is a single-qubit gate. Let U be a $\{R\}$-circuit on n qubits, M a Hermitian operator on $O(\log(n))$ qubits, and $|\psi\rangle, |\phi\rangle$ arbitrary n-qubit product states. Then $\langle \psi|U^\dagger(M \otimes I)U|\phi \rangle$ can be computed exactly in $O(poly(n))$ classical time.*

Proof. We first apply the procedure from (4) as before, and write

$$U = Q^{\otimes n} V (Q^\dagger)^{\otimes n}$$

where V is described by a circuit consisting only of S gates. The unitary operator implemented by V is an element of $\mathcal{C}_n$. Now let M be a Hermitian operator on $m = c\log(n)$ qubits, and suppose for simplicity that it acts only on the first m qubits. Let I denote the identity operator on the $(m+1)$st through nth qubits. We write

$$\langle \psi|U^\dagger(M \otimes I)U|\phi \rangle = \langle \psi|Q^{\otimes n}V^\dagger Q^{\dagger \otimes n}(M \otimes I)Q^{\otimes n}VQ^{\dagger \otimes n}|\phi \rangle$$
$$= \langle \psi|Q^{\otimes n}V^\dagger(M' \otimes I)VQ^{\dagger \otimes n}|\phi \rangle \, ,$$

where $M' = Q^{\otimes m}MQ^{\dagger \otimes m}$.

As discussed earlier, a basis for the space of Hermitian operators on m qubits is the m-qubit Pauli group $\mathcal{P}_m$, which has size $O(poly(n))$. The expansion of M' in that basis can be computed in polynomial time by basic linear algebra. Embedding the first m qubits into

all n qubits gives the obvious embedding of $\mathcal{P}_m$ into $\mathcal{P}_n$, and this also gives (the same, still polynomial-size) expansion of M' into n-qubit Paulis. We write

$$M' = \sum_{\sigma \in \mathcal{P}_n \cap \mathcal{P}_m} \alpha_\sigma \sigma \,.$$

We emphasize that this is a sum over polynomially many terms, and that each coefficient can be calculated from knowledge of M and Q in polynomial time. Moreover, since V is a Clifford circuit, its conjugation action $\sigma \mapsto \sigma^V := V^\dagger \sigma V$ on a Pauli group element $\sigma \in \mathcal{P}_n$ is easily computed by direct gate-by-gate matrix multiplication (see, e.g., [11]).

We now return to the main calculation, to see that

$$\langle \psi | U^\dagger (M \otimes I) U | \phi \rangle = \langle \psi | Q^{\otimes n} V^\dagger (M' \otimes I) V Q^{\dagger \otimes n} | \phi \rangle$$

$$= \sum_{\sigma \in \mathcal{P}_n \cap \mathcal{P}_m} \alpha_\sigma \langle \psi | Q^{\otimes n} V^\dagger \sigma V Q^{\dagger \otimes n} | \phi \rangle$$

$$= \sum_{\sigma \in \mathcal{P}_n \cap \mathcal{P}_m} \alpha_\sigma \langle \psi | Q^{\otimes n} \sigma^V Q^{\dagger \otimes n} | \phi \rangle$$

$$= \sum_{\sigma \in \mathcal{P}_n \cap \mathcal{P}_m} \alpha_\sigma \prod_{j=1}^n \langle \psi_j | Q \sigma_j^V Q^\dagger | \phi_j \rangle \,.$$

The sum and product in the final expression are both of polynomial size, and each term in the product can be computed in constant time. ◀

4 Qubit solutions to Yang-Baxter

4.1 The four solution families

Hietarinta classified all solutions to the Yang-Baxter equation in the qubit (i.e., 4×4) case [12]. The qubit solutions which are also unitary operators were identified by Dye [5]. All of these are of the form

$$R = k(Q \otimes Q)ST(Q \otimes Q)^{-1} \tag{5}$$

where k is a unit-norm scalar, T is the swap gate, and

$$Q = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

is an invertible matrix. The trivial solution is $S = T$ which implies $R = kI$. There are four nontrivial solution families, depending on the possible values taken by S, which are listed below, along with the required conditions on the matrix entries.

$$S_1 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & p & 0 & 0 \\ 0 & 0 & q & 0 \\ 0 & 0 & 0 & r \end{pmatrix} \qquad\qquad 1 = |p| = |q| = |r| \ ; \ c = -a\bar{b}/\bar{d}$$

$$S_2 = \begin{pmatrix} 0 & 0 & 0 & p \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ q & 0 & 0 & 0 \end{pmatrix} \qquad\qquad p = \frac{(b\bar{b} + d\bar{d})(\bar{a}b + \bar{c}d)}{(a\bar{a} + c\bar{c})(a\bar{b} + c\bar{d})} \ ; \ q = 1/p \ ; \ c \neq -a\bar{b}/\bar{d}$$

$$S_3 = \begin{pmatrix} 0 & 0 & 0 & p \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ q & 0 & 0 & 0 \end{pmatrix} \qquad p\bar{p} = \frac{(d\bar{d})^2}{(a\bar{a})^2} \; ; \; q\bar{q} = \frac{(a\bar{a})^2}{(d\bar{d})^2} \; ; \; |pq| = 1 \; ; \; c = -a\bar{b}/\bar{d}$$

$$S_4 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & -1 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix} \qquad |a| = |d| \; ; \; c = -a\bar{b}/\bar{d}.$$

For $j = 1, 2, 3, 4$, let R_j be the Yang-baxter operator (5) resulting from choosing $S = S_j$.

4.2   Families one, two and three are unlikely to be universal

We will show that Theorem 6 applies to the single-element gate sets $\{R_1\}$, $\{R_2\}$, and $\{R_3\}$. We assume that all of the above matrix entries are exactly computable in constant time via a classical algorithm.

The gate R_1 has the form (3) where $C_i = I$, $P_i = T$, and $D_i = kS_1$. It remains to check that Q satisfies property (G) where G is the trivial group consisting only of the identity; this is confirmed by Lemma 8 below.

For the gate R_2, we set

$$M = \begin{pmatrix} 0 & \sqrt{p} \\ 1/\sqrt{p} & 0 \end{pmatrix}$$

and check that $M \otimes M = S_2$. It follows that $R_2 = kT(QMQ^{-1} \otimes QMQ^{-1})$ is not an entangling gate. Since R_2 is unitary, so is QMQ^{-1}. By the spectral theorem, there exist diagonal V and unitary U such that $UVU^{-1} = QMQ^{-1}$. Observe that $R_2 = (U \otimes U)k(V \otimes V)T(U \otimes U)^{-1}$ satisfies the conditions of Theorem 6.

For the gate R_3, we first rewrite the matrices as follows. Set

$$N = \begin{pmatrix} p^{-1/4} & 0 \\ 0 & p^{1/4} \end{pmatrix}$$

and $Q' = QN^{-1}$ and $S_3' = (N \otimes N)S_3(N \otimes N)^{-1}$. It's not hard to check that

$$R_3 = k(Q \otimes Q)S_3 T(Q \otimes Q)^{-1} = k(Q' \otimes Q')S_3' T(Q' \otimes Q')^{-1} \, ,$$

and that Q' and S_3' satisfy the conditions of the third YBE solution family, with the additional property that $p = 1$ and $|q| = 1$. Note further that $S_3'(X \otimes X)$ is a diagonal unitary operator, where

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \, .$$

We now see that R_3 is of the form (3) from Theorem 6, where $C_i = X$, $D_i = kS_3'(X \otimes X)$, and $P = T$. It remains to check that Q' satisfies property $(\langle X \rangle)$, which is done in Lemma 9 below.

▶ **Lemma 8.** *Let Q be an invertible 2×2 matrix defined by*

$$Q = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \, ,$$

such that $c = -a\bar{b}/\bar{d}$. Then Q satisfies property (I).

Proof. Define the relevant matrices

$$A = \begin{pmatrix} |a| & |b| \\ |c| & |d| \end{pmatrix} \qquad \text{and} \qquad B = \frac{1}{|ad - bc|} \begin{pmatrix} |d| & |b| \\ |c| & |a| \end{pmatrix}.$$

Note that $a = 0$ implies $c = 0$, which would make Q non-invertible.

We compute each case separately. First let $k = l = 1$.

$$A_{11}B_{11} + A_{12}B_{21} = \frac{|a||d| + |b||c|}{|ad - bc|} = \frac{|a||d| + |b||a\bar{b}/\bar{d}|}{|ad + ba\bar{b}/\bar{d}|}$$

$$= \frac{|\bar{d}|(|a||d|^2 + |a||b|^2)}{|\bar{d}|(|ad\bar{d} + ab\bar{b}|)} = \frac{|a|(|d|^2 + |b|^2)}{|a||d\bar{d} + b\bar{b}|} = 1.$$

Next, let $k = l = 2$, and we again get

$$A_{21}B_{12} + A_{22}B_{22} = \frac{|c||b| + |d||a|}{|ad - bc|} = 1.$$

Now suppose $k = 1$ and $l = 2$.

$$A_{11}B_{12} + A_{12}B_{22} = \frac{|a||b| + |b||a|}{|ad - bc|} = \frac{2|a||b|}{|ad + ab\bar{b}/\bar{d}|}$$

$$= \frac{2|a||b||d|}{|ad\bar{d} + ab\bar{b}|} = \frac{2|b||d|}{|d|^2 + |b|^2}.$$

It remains to note that

$$|b|^2 + |d|^2 - 2|b||d| = (|b| - |d|)^2 > 0.$$

Finally, we choose $k = 2$ and $l = 1$.

$$A_{21}B_{11} + A_{22}B_{21} = \frac{|c||d| + |d||c|}{|ad - bc|} = \frac{2|a||b|}{|ad - bc|} \leq 1,$$

by two applications of $c = -a\bar{b}/\bar{d}$ and the previous case. ◀

▶ **Lemma 9.** *Let Q be an invertible 2×2 matrix defined by*

$$Q = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

such that $c = -a\bar{b}/\bar{d}$ and $|a|^2 = |d|^2$. Then Q satisfies property $(\mathcal{S}_2)$.

Proof. Define the matrices A and B as in Lemma 8. The case of π equal to the trivial permutation is handled by Lemma 8. We compute the remaining cases. Set $\pi = (12)$ and $k = l = 1$. Then

$$A_{12}B_{11} + A_{11}B_{21} = \frac{|a||c| + |b||d|}{|ad - bc|} = \frac{|aa\bar{b}/\bar{d}| + |bd|}{|ad - ab\bar{b}/\bar{d}|}$$

$$= \frac{|aa\bar{b}| + |bd\bar{d}|}{|ad\bar{d} + ab\bar{b}|} = \frac{|aa\bar{b}| + |ba\bar{a}|}{|aa\bar{a} + ab\bar{b}|} = \frac{|a\bar{b}| + |b\bar{a}|}{|a|^2 + |b|^2},$$

where we have applied the facts $c = -a\bar{b}/\bar{d}$ and $a\bar{a} = d\bar{d}$ and $a \neq 0$. Now note that

$$|a|^2 + |b|^2 - (|a\bar{b}| + |b\bar{a}|) = |a|^2 + |b|^2 - 2|a||b| = (|a| - |b|)^2 \geq 0.$$

Hence $(|a\bar{b}| + |b\bar{a}|)/(|a|^2 + |b|^2) \leq 1$. For $k = l = 2$, we again get

$$A_{22}B_{12} + A_{21}B_{22} = \frac{|a||c| + |b||d|}{|ad - bc|} \leq 1 \,.$$

Now set $k = 1$ and $l = 2$. Then

$$\begin{aligned}
A_{12}B_{12} + A_{11}B_{22} &= \frac{|a|^2 + |b|^2}{|ad - bc|} = \frac{|a|^2 + |b|^2}{|ad + ab\bar{b}/\bar{d}|} \\
&= \frac{|\bar{d}|(|a|^2 + |b|^2)}{|ad\bar{d} + ab\bar{b}|} = \frac{|d|(|a|^2 + |b|^2)}{|a|(|d|^2 + |b|^2)} = 1 \,.
\end{aligned}$$

Finally, for $k = 2$ and $l = 1$, write $b = -\bar{c}d/\bar{a}$ and calculate

$$\begin{aligned}
A_{22}B_{11} + A_{21}B_{21} &= \frac{|c|^2 + |d|^2}{|ad - bc|} = \frac{|c|^2 + |d|^2}{|ad + dc\bar{c}/\bar{a}|} \\
&= \frac{|\bar{a}|(|c|^2 + |d|^2)}{|da\bar{a} + dc\bar{c}|} = \frac{|a|(|c|^2 + |d|^2)}{|d|(|c|^2 + |a|^2)} = 1 \,.
\end{aligned}$$

◀

To conclude, we have shown the following.

▶ **Theorem 10.** *Let $R \in \{R_1, R_2, R_3\}$ be a unitary solution to the Yang-Baxter equation on qubits. Then $\mathcal{I}(\{R\})$ is in PromiseBPP.*

In particular, if one could perform (perhaps encoded) universal quantum computation with these circuits then PromiseBQP = PromiseBPP. We can also formulate the lack of universality for these solutions in the following terms.

▶ **Theorem 11.** *Let $R \in \{R_1, R_2, R_3\}$ be a unitary solution to the Yang-Baxter equation on qubits, and let $\rho_n : B_n \to SU(2^n)$ be the corresponding unitary representation of the braid group. Then the image of ρ_n is not dense in $SU(2^n)$ for any $n \geq 2$, unless PromiseBQP = PromiseBPP.*

Proof. (Sketch) For a contradiction, suppose there exists an $n \geq 2$ such that the image of ρ_n is dense. Let C be an arbitrary m-qubit quantum circuit. We can assume without loss of generality that C only consists of 2-qubit gates acting on adjacent qubits, and that n is even. For each of the m qubits, assign $n/2$ qubits from the space of ρ_n. By the density of the image of ρ_n, we can then simulate C inside $\rho_{mn/2}$ gate-by-gate via the Solovay-Kitaev theorem. Then we can use the classical algorithm from Theorem 6 to approximate the relevant matrix entry of the resulting R-circuit, thus solving the PromiseBQP-hard problem of approximating the corresponding entry of C. ◀

4.3 Family four is unlikely to be universal

Recall that the fourth solution family is of the form $R_4 = k(Q \otimes Q)S_4T(Q \otimes Q)^{-1}$. We begin by demonstrating a Clifford circuit which is equal to the gate S_4T.

$$S_4T = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & -1 & 1 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix} =$$

We also note that, in this solution family, Q is a scaled unitary operator. To see this, note that

$$Q^\dagger Q = \begin{pmatrix} |a|^2 + |c|^2 & \bar{a}b + \bar{c}d \\ a\bar{b} + c\bar{d} & |b|^2 + |d|^2 \end{pmatrix} = \begin{pmatrix} |a|^2 + |c|^2 & 0 \\ 0 & |b|^2 + |a|^2 \end{pmatrix} = (|a|^2 + |b|^2) \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

where we first applied the condition $c = -a\bar{b}/\bar{d}$ to the off-diagonal elements and the condition $|a|^2 = |d|^2$ to the diagonal ones; the last equality follows from combining these two conditions to get $|c|^2 = |b|^2$. Now set $\alpha = (|a|^2 + |b|^2)^{1/2}$ and $Q_1 = \alpha^{-1}Q$. Using the above, one easily checks that Q_1 is unitary and that $Q_1^\dagger = \alpha Q^{-1}$. It follows that

$$(Q \otimes Q)A(Q \otimes Q)^{-1} = (\alpha Q_1 \otimes \alpha Q_1)A(\alpha^{-1}Q_1^\dagger \otimes \alpha^{-1}Q_1^\dagger) = (Q_1 \otimes Q_1)A(Q_1 \otimes Q_1)^\dagger$$

for any A. For us it will thus suffice to assume that Q is in fact unitary. This allows us to apply Theorem 7 and get the following result.

▶ **Theorem 12.** *Let U be a $\{R_4\}$-circuit on n qubits, M a Hermitian operator on $O(\log(n))$ qubits, and $|\psi\rangle, |\phi\rangle$ arbitrary n-qubit product states. Then $\langle\psi|U^\dagger(M \otimes I)U|\phi\rangle$ can be computed exactly in $O(poly(n))$ classical time.*

5 Some simple high-dimensional solutions

Finally, we list some simple unitary solution families to the Yang-Baxter equation that exist in every dimension, and to which Theorem 6 applies. We begin by observing that, whenever a 2-qudit gate S is a solution, then by (4) so is $(Q \otimes Q)S(Q \otimes Q)^{-1}$ for any 1-qudit gate Q.

For $A, B \in \mathrm{U}(V)$, the operator $T(A \otimes B)$ is a solution to the Yang-Baxter equation if and only if A and B commute. This is easily seen by following the wires in the circuits below.
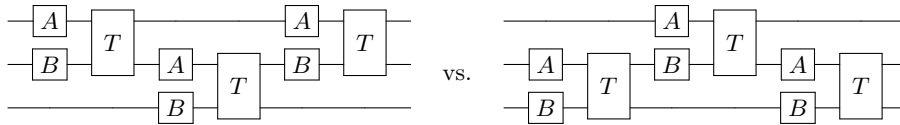
If A and B do commute, then there's a unitary change of basis Q on V such that $Q^{-1}AQ$ and $Q^{-1}BQ$ are both diagonal. Therefore, Theorem 6 applies to $T(A \otimes B)$, so any circuits using this gate are classically simulable. Of course, this is not surprising, as they do not even entangle the qudits.

More generally, suppose $S \in \mathrm{U}(V \otimes V)$ is diagonal in the computational basis, and set

$$\lambda_{ij} = \langle ij|S|ij\rangle \qquad \text{for} \qquad i, j \in [d],$$

where $d = \dim V$. Note that

$$S_{12} = S \otimes I = \bigoplus_{k \in [d]} P_k, \qquad S_{23} = I \otimes S = \bigoplus_{k \in [d]} S, \qquad I \otimes T = \bigoplus_{k \in [d]} T.$$

where $P_k = \oplus_{l \in [d]} \lambda_{kl} I$. We also have

$$S_{13} = (I \otimes T)(S \otimes I)(I \otimes T) = \bigoplus_{k \in [d]} TP_kT.$$

Substituting the above into the two sides of the algebraic Yang-Baxter equation (2), we get

$$\bigoplus_{k \in [d]} P_kTP_kTS \qquad \text{and} \qquad \bigoplus_{k \in [d]} STP_kTP_k$$

Clearly, P_k and S are symmetric. Since $\langle ab|T|cd\rangle = \delta_{ad}\delta_{bc} = \langle cd|T|ab\rangle$, so is T. By applying the transpose to one of the two sides above, we see that S satisfies algebraic Yang-Baxter. Thus ST is a solution to the YBE, one to which Theorem 6 clearly applies.

References

1 Emil Artin. Theorie der Zopfe. *Abh. Math. Sem. Univ. Hamburg*, 4:47–72, 1925.

2 John C. Baez. Braids and quantization (online lecture notes), May 1992.

3 M. Bremner, R. Jozsa, and D. J. Sheperd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proc. R. Soc. A*, 467:459, 2011.

4 J. L. Brylinski and R. Brylinski. Universal quantum gates. In *Mathematics of Quantum Computation*. Chapman & Hall, 2002.

5 H. A. Dye. Unitary solutions to the Yang-Baxter equation in dimension four. *Quantum Information Processing*, 2(1/2):117–152, April 2003.

6 Jennifer M. Franko. Braid group representations arising from the Yang-Baxter equation. *Journal of Knot Theory and Its Ramifications*, 19(04):525–538, 2010.

7 Jennifer M. Franko, Eric C. Rowell, and Zhenghan Wang. Extraspecial 2-groups and images of braid group representations. *Journal of Knot Theory and Its Ramifications*, 15(04):413–427, 2006.

8 Michael H. Freedman, Alexei Kitaev, and Zhenghan Wang. Simulation of topological field theories by quantum computers. *Communications in Mathematical Physics*, 227:587–603, 2002.

9 Michael H. Freedman, Michael J. Larsen, and Zhenghan Wang. A modular functor which is universal for quantum computation. *Communications in Mathematical Physics*, 227:605–622, 2002.

10 Michael H. Freedman, Michael J. Larsen, and Zhenghan Wang. The two-eigenvalue problem and density of Jones representation of braid groups. *Communications in Mathematical Physics*, 228:177–199, 2002.

11 Daniel Gottesman. The Heisenberg representation of quantum computers. *Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, 22:32–43, 1999.

12 Jarmo Hietarinta. Solving the two-dimensional constant quantum Yang-Baxter equation. *Journal of Mathematical Physics*, 34(5):1725–1756, 1993.

13 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.

14 Dominik Janzing and Pawel Wocjan. A simple PromiseBQP-complete matrix problem. *Theory of Computing*, 3(4):61–79, 2007.

15 Vaughan F. R. Jones. A polynomial invariant for knots via von Neumann algebras. *Bull. Amer. Math. Soc. (N.S.)*, 12(1):103–111, 1985.

16 R. Jozsa and A. Miyake. Matchgates and classical simulation of quantum circuits. *Royal Society of London Proceedings Series A*, 464:3089–3106, December 2008.

17 Richard Jozsa and Maarten Van den Nest. Classical simulation complexity of extended clifford circuits. *arXiv*, 2013.

18 Louis H. Kauffman and Samuel J. Lomonaco Jr. Braiding operators are universal quantum gates. *New Journal of Physics*, 6(1):134, 2004.

19 Gabriele Nebe, E. M. Rains, and N. J. A. Sloane. The invariants of the Clifford groups. *Designs, Codes and Cryptography*, 24(1):99–122, 2001.

20 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information.* Cambridge University Press, 2000.

21 J. H. H. Perk and H. Au-Yang. Yang-Baxter equation. In *Encyclopedia of Mathematical Physics*, pages 465–473. Oxford: Elsevier, 2006.

22 John Preskill. Topological quantum computation (online lecture notes), 2004.

23 Eric C. Rowell and Zhenghan Wang. Localization of unitary braid group representations. *Communications in Mathematical Physics*, 311(3):595–615, 2012.

24 V. G. Turaev. The Yang-Baxter equation and invariants of links. *Invent. Math.*, 92:527–553, 1988.

A   Appendix

We will now prove Lemmas 3 and 4.

Proof of Lemma 3. To sample from P_j, flip m unbiased coins to get an integer $0 \leq l \leq 2^m$. Subdivide 2^m into intervals according to

$$2^m = P_j(0)2^m + P_j(1)2^m + \cdots + P_j(d-1)2^m$$

and output k if l falls into the kth interval. Then the probability $D_j(k)$ with which you output k satisfies $|D_j(k) - P_j(k)| \leq 1/2^m$. Now do this for two indices, say 1 and 2 and note that

$$\begin{aligned}
|P_1(k)P_2(l) - D_1(k)D_2(l)| &= |P_1(k)P_2(l) - D_1(k)D_2(l) + D_1(k)P_2(l) - D_1(k)P_2(l)| \\
&\leq |P_2(l)(P_1(k) - D_1(k))| + |D_1(k)(P_2(l) - D_2(l))| \\
&\leq 2/2^m
\end{aligned}$$

Extending this to the case of multiplying all n distributions together, we get $|P(y) - D(y)| \leq n/2^m$ for all $y \in [d]^n$. The total variation distance then satisfies

$$|P - D| = \frac{1}{2} \sum_{x \in [d]^n} |P(x) - D(x)| \leq \frac{nd^n}{2^m} < 2^{-n}$$

so long as $m \geq 3n \log d$.   ◄

Proof of Lemma 4. We expand the X_j into real and imaginary parts and apply the standard bound. Set $S_r = \sum_j \mathrm{Re}[X_j]/n$ and $S_i = \sum_j \mathrm{Im}[X_j]/n$ and $\mu_r = \mathbb{E}[\mathrm{Re}[X_j]]$ and $\mu_i = \mathbb{E}[\mathrm{Im}[X_j]]$. Note that $|\mathrm{Re}[X_j]| \leq b$ and $|\mathrm{Im}[X_j]| \leq b$. By the Chernoff-Hoeffding bound for real-valued random variables [13], we have

$$\Pr\left[|S_r - \mu_r| \geq \epsilon/2\right] \leq 2\exp\left(-n\epsilon^2/8b^2\right),$$

and likewise for the imaginary part. Taking the union bound, we have that

$$|S - \mu| = |S_r - \mu_r + i(S_i - \mu_i)| \leq |S_r - \mu_r| + |S_i - \mu_i| \leq \epsilon/2 + \epsilon/2 = \epsilon$$

except with probability $4\exp\left(-n\epsilon^2/8b^2\right)$.   ◄