A Polynomial Quantum Algorithm for Approximating the Jones Polynomial

Dorit Aharonov [*]   Vaughan Jones [†]   Zeph Landau [‡]

November 26, 2024

Abstract

The Jones polynomial, discovered in 1984 [18], is an important knot invariant in topology. Among its many connections to various mathematical and physical areas, it is known (due to Witten [32]) to be intimately connected to Topological Quantum Field Theory (TQFT). The works of Freedman, Kitaev, Larsen and Wang [13, 14] provide an efficient simulation of TQFT by a quantum computer, and vice versa. These results implicitly imply the existence of an efficient quantum algorithm that provides a certain additive approximation of the Jones polynomial at the fifth root of unity, $e^{2\pi i/5}$, and moreover, that this problem is BQP-complete. Unfortunately, this important algorithm was never explicitly formulated. Moreover, the results in [13, 14] are heavily based on TQFT, which makes the algorithm essentially inaccessible to computer scientists.

We provide an explicit and simple polynomial quantum algorithm to approximate the Jones polynomial of an n strands braid with m crossings at any primitive root of unity $e^{2\pi i/k}$, where the running time of the algorithm is polynomial in m, n and k. Our algorithm is based, rather than on TQFT, on well known mathematical results (specifically, the path model representation of the braid group and the uniqueness of the Markov trace for the Temperly Lieb algebra). By the results of [14], our algorithm solves a BQP complete problem.

The algorithm we provide exhibits a structure which we hope is generalizable to other quantum algorithmic problems. Candidates of particular interest are the approximations of other downwards self-reducible #P-hard problems, most notably, the important open problem of efficient approximation of the partition function of the Potts model, a model which is known to be tightly connected to the Jones polynomial [33].

1 Introduction

Since Shor's breakthrough discovery in 1994 [28], quantum algorithms with an exponential speedup over the best known classical algorithms have been shown for a number of problems (e.g., [10, 30, 15, 22]). All these problems and algorithms share some common features: the problems are group or number theoretic in nature and the key component of each algorithm is the quantum Fourier transform[1]. Arguably, the greatest challenge of quantum computation is the discovery of new algorithmic techniques.

In this paper we describe a polynomial time quantum algorithm that approximates the #P-hard problem of evaluating the Jones polynomial at certain roots of unity. The best classical algorithm for this problem is exponential. Our algorithm is significantly different from all previously known quantum algorithms that achieve an exponential speed up in the following three important ways: 1) it solves a problem which is

---

[*]School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel. doria@cs.huji.ac.il.

[†]Department of Mathematics, U.C.Berkeley

[‡]Department of Mathematics, The City College of New York, NY

[1]One interesting exception is the beautiful result of [8] who used random walks to achieve an exponential algorithmic speed-up for an oracle graph problem. This technique, however, has not yet found applications.

combinatorial rather than group or number theoretic in nature, 2) it does so not by using the Fourier transform, but instead, by exploiting a certain structure of the problem and encoding it into the nature of the unitary gates being used, and 3) it solves a problem that is BQP-hard [12], that is, a problem that captures all the power of the quantum model.

The connection between quantum computation and the Jones polynomial was first made through the series of papers [11, 12, 13, 14]. A model of quantum computation based on Topological Quantum Field Theory (TQFT) and Chern-Simons theory was defined in [11, 12], and Kitaev, Larsen, Freedman and Wang showed that this model is polynomially equivalent in computational power to the standard quantum computation model in [13, 14]. These results, combined with a deep connection between TQFT and the value of the Jones polynomial at particular roots of unity discovered by Witten 20 years ago [32] , implicitly implies an efficient quantum algorithm for the approximation of the Jones polynomial at the value $e^{2\pi i/5}$. This connection is also discussed, from the point of view of TQFT, in Preskill's notes [25]. Unfortunately, the important quantum algorithm implied by these intriguing results, though referred to in [7], was never explicitly formulated.

In this paper we use a different route to connect quantum computation and the Jones polynomial, one that does not involve TQFT. We present an explicit and simple to state algorithm for the above problem, which is based purely on algebraic results from more than 20 years ago. Moreover, our algorithm works for all roots of unity of the form $e^{2\pi i/k}$, going beyond the discussions in previous works involving only constant k's. We now describe the precise problem that we solve.

1.1   Background on the Jones Polynomial

A central issue in low dimensional topology is that of *knot invariants*. A knot invariant is a function on knots (or links –i.e. circles embedded in R^3) which is invariant under isotopy of the link, i.e., it does not change under stretching, moving, etc., but no cutting. In 1984, Jones [18] discovered a new knot invariant, now called the Jones polynomial $V_L(t)$, which is a Laurent polynomial in $\sqrt{t}$ with integer coefficients, and which is an invariant of the link L. In addition to the important role it has played in low dimensional topology, the Jones polynomial has found applications in numerous fields, from DNA recombination [24], to statistical physics [33].

The Jones polynomial can also be defined as a function of braids. A braid of n strands and m crossings is described pictorially by n strands hanging alongside each other, with m crossings, each of two adjacent strands. A braid may be "closed" to form a link by tying its ends together. In this paper we will be interested in two ways to perform such closures, namely, the *trace closure* and the *plat closure* (to be defined later). We will be interested in the Jones polynomial of links that are trace or plat closures of braids.

The essential aspect of the Jones polynomial of a link L can be computed by the following process: project L to the plane keeping track crossings to get what is called a *link diagram*. Now replace every crossing by a particular linear combination (with coefficients being functions of the parameter t) of the pictures and . The result is a linear combination of diagrams containing only closed loops. Replace each of these diagrams with a particular function of the parameter t and the number of loops. The resulting function of t is a scaled version of the Jones polynomial $V_L(t)$.

From the moment of the discovery of the Jones polynomial, the question of how hard it is to compute was important. There is a very simple inductive algorithm (essentially due to Conway [9]) to compute it by changing crossings in a link diagram, but, naively applied, this takes exponential time in the number of crossings. It was shown [16] that the computation of $V_L(t)$ is #P-hard for all but a few values of t where $V_L(t)$ has an elementary interpretation. Thus a polynomial time algorithm for computing $V_L(t)$ for any value of t other than those elementary ones is unlikely. Of course, the #P-hardness of the problem does not rule

out the possibility of good approximations; see, e.g., [17]. Still, the best classical algorithms to approximate the Jones polynomial at all but trivial values are exponential.

1.2 Our Results

We show an efficient, explicit, and simple quantum algorithm to approximate the Jones polynomial at all points of the form $t = e^{2\pi i/k}$. It will in fact be easier to use the parameter $A = t^{-1/4}$. We prove the following for the trace closure B^{tr} and the plat closure B^{pl} of a braid B:

Theorem 1.1 *For a given braid B with n strands and m crossings, and a given integer k, there is a quantum algorithm which is polynomial in n, m, k which with all but exponentially small probability, outputs a complex number r with $|r - V_{B^{tr}}(e^{2\pi i/k})| < \varepsilon d^{n-1}$ where $d = -A^2 - A^{-2}$, and ε is inverse polynomial in n, k, m.*

Theorem 1.2 *For a given braid B with n strands and m crossings, and a given integer k, there is a quantum algorithm which is polynomial in n, m, k which with all but exponentially small probability, outputs a complex number r with $|r - V_{B^{pl}}(e^{2\pi i/k})| < \varepsilon d^{3n/2}/N$ where $d = -A^2 - A^{-2}$ and ε is inverse polynomial in n, k, m (N is an exponentially large factor to be defined later).*

We remark that the approximation we provide here is additive, namely the result lies in a given window, whose size is independent of the actual value we are trying to approximate. This of course is much weaker than a multiplicative approximation, which is what one might desire (see discussion in [7]). One might wonder if under such weak requirements, the problem remains meaningful at all. It turns out that, in fact, this additive approximation problem is hard for quantum computation:

Theorem 1.3 Adapted from Freedman, Larsen and Wang [14] *The problem of approximating the Jones polynomial of the plat closure of a braid at $e^{2\pi i/k}$ for constant k, to within the accuracy given in Theorem 1.2, is* BQP*-hard.*

This result was recently strengthened by Aharonov and Arad [1] to any k which is polynomial in the size of the input, namely, for all the plat closure cases for which our algorithm is polynomial in the size of the braid.

Curiously, the hardness results of [14, 1] are not known to hold (regardless of k) for the approximation of the trace closure for which we give an algorithm as well. We discuss the difference between the two problems in the open questions section.

1.3 Description of the Algorithm

The essence of the algorithm lies in the fact that for braids with n strands, the pictures that one gets by replacing each crossing by one of the two pictures { $\overset{\cup}{\cap}$, $\mid\mid$ }, can be assigned the structure of an algebra. This algebra is called the Temperley Lieb algebra, and denoted by TL_n. In fact, the map from the crossing to the appropriate linear combination of the above two pictures defines a *representation* of the group B_n of braids of n strands, inside the TL_n algebra. The Jones polynomial of the trace closure of a braid can be seen as a certain trace function (i.e., a linear function satisfying $tr(AB) = tr(BA)$) on the image of the braid in the TL_n algebra.

Our goal is then to design an algorithm that approximates this trace. To this end we use an important fact about this trace: it satisfies an additional property called the *Markov property*. Moreover, this property makes it unique; any trace function on the TL_n algebra (or a representation of it) that satisfies this property is equal to the above trace! This leads us to the key idea of the algorithm: suppose we can define a representation of the TL_n algebra by matrices operating on qubits, and we can identify and estimate the trace that satisfies the Markov property on this representation. Then by the uniqueness of this trace we can estimate the Jones polynomial.

3

But what is the representation that should be used? If our intent is to design a quantum algorithm, it is best if the representation induced on the braid group be unitary, so that we can hope to approximate its trace by a quantum computer. Fortunately, it is in fact possible to give representations of the Temperley Lieb algebra which induce *unitary* representations of the braid group. Such representations were constructed in [20, 18] and are called the *path model* representations. If we want to evaluate the Jones polynomial $V_L(t)$ for L a closure of a braid in B_n, and $t = e^{2\pi i/k}$, we use the kth path model representation of B_n. It is fairly straight forward to adapt these representations to work on the space of n qubits, and moreover, to show that the resulting unitary matrices (namely, the images of the generators of the braid group) can be applied efficiently by a quantum computer. We find that the image (by the path model representation) of an entire braid B can be applied efficiently by a quantum computer. Let us call the unitary matrix corresponding to a braid $Q(B)$.

To approximate the Jones polynomial of a trace closure of the braid, it suffices to approximate the Markov trace of $Q(B)$. This is done using standard quantum and classical algorithmic techniques, including the well known Hadamard test. The algorithm for the plat closure builds on similar ideas, though it is not directly stated in terms of traces. Thus we obtain a polynomial quantum algorithm for the BQP-complete problem of approximating the Jones polynomial of a plat closure of a braid.

We remark that after we completed this work, we learned about a previous independent attempt to prove similar results [26]. Unfortunately, the work of [26] is greatly flawed, and in particular claims to provide an exact solution to the #P-hard problem.

1.4 Conclusions and Further Directions

We have provided a simple algorithm for a BQP-complete problem, which is different in its methods than previous quantum algorithms. In essence, what it does is to isolate a certain local structure of the problem, and assign gates which somehow exhibit the same local structure. Our hope is that this more combinatorial direction in quantum algorithms will lead to further progress in the area.

In particular, one very interesting related question is an important problem from mathematical physics: that of approximating the partition function of the Potts model [31], which is known to be tightly connected to the Jones polynomial [33], and its exact evaluation is once again #P-hard [31]. We hope that the results of this paper will lead to progress in this question, or in other questions related to approximating #P-complete problems.

We briefly discuss the relation between the plat and the trace closures problems. It is known that any plat closure of a braid can be transformed efficiently into a trace closure of some other braid [29]. The reader might therefore find it curious that one of these problems is BQP-complete, while the other one is not known to be so. The explanation lies in the fact that the quality of the approximation in both algorithms depends *exponentially* on the number of strands in the braid. The transformation from plat to trace closures requires, in the worst case, a significant increase in the number of strands. This, unfortunately, degrades the quality of the approximation exponentially. The computational complexity of the trace closure problem is left open.

Finally, we believe that this paper helps to clarify and demystify (at least one direction of) the intriguing equivalence between quantum computation and the problem of approximating the Jones polynomial. We hope this connection leads to a deeper understanding of quantum computation complexity.

Organization of paper: Section 2 provides the background and the necessary definitions, starting from quantum computation, the Hadamard test, the braid group, algebras and representations, the Jones polynomial, the Temperley Lieb algebras and the path model representation. Using these notions we describe the algorithms in Section 3 and prove their correctness.

2 Background

2.1 Quantum Computation

For background on quantum computation, see [23]. We merely mention here that, strictly speaking, we use here a quantum-classical hybrid model of computation, in which a classical probabilistic Turing machine performs calls to a quantum computer, and uses its outcomes to perform some classical computations. It is standard that this model can be simulated efficiently by the standard quantum computation model.

2.2 The Hadamard Test

The following fact is standard in quantum computation. If a state $|\alpha\rangle$ can be generated efficiently, and a unitary Q can be applied efficiently, then there exists an efficient quantum circuit whose output is a random variable $\in \{-1, 1\}$, and whose expectation is $\mathcal{R}e\langle\alpha|Q|\alpha\rangle$. Start with the two-register state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |\alpha\rangle$, apply Q conditioned on the first qubit to get the state $\frac{1}{\sqrt{2}}(|0\rangle \otimes |\alpha\rangle + |1\rangle \otimes Q|\alpha\rangle)$, apply a Hadamard gate on the first qubit, and measure. Output 1 if the measurement result is $|0\rangle$, -1 if the measurement result is $|1\rangle$. The expectation of the output is exactly $\mathcal{R}e\langle\alpha|Q|\alpha\rangle$. To get a random variable whose expectation is the imaginary part, start with the state $\frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle) \otimes |\alpha\rangle$ instead.

2.3 Algebra Background

An algebra is a vector space with a multiplication. The multiplication must be associative and distributive. A representation of a group G inside an algebra is a group homomorphism ρ from G to the group of invertible elements in the algebra, namely, we require $\rho(g_1)\rho(g_2) = \rho(g_1 g_2)$ for any $g_1, g_2 \in G$.

We say a representation is reducible if there exists a proper subspace of vectors which is invariant under the group action. If there is no such subspace, we say the representation is irreducible.

We shall sometimes refer to a representation of G without specifying the algebra; in these cases we shall mean a representation inside the algebra of $n \times n$ matrices. We shall be interested in algebra representations as well:

Definition 2.1 *An r dimensional representation Φ of an algebra is a linear mapping from the algebra into the set of $r \times r$ complex matrices M_r, such that for any two elements X, Y in the algebra, $\Phi(X)\Phi(Y) = \Phi(XY)$*
.

If a group is represented inside an algebra then any representation of the algebra gives a representation of the group by composition.

Often, an algebra or a group is defined using a set of generators and relations between them. In this case, a representation may be defined by specifying the images of the generators, provided the same relations hold for the images as for the generators.
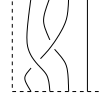
2.4 The Braid Group

Consider two horizontal bars, one on top of the other, with n pegs on each. By an *n strand braid* we shall mean a set of n strands such that: (1) Each strand is tied to one peg on the top bar and one peg on the bottom bar, (2) Every peg has exactly one end attached to it, (3) The strands may pass over and under each other, (4) The tangent vector of every strand at any point along the path from top to bottom always has a non-zero component in the downward direction. Here is an example of a 4-strand braid:
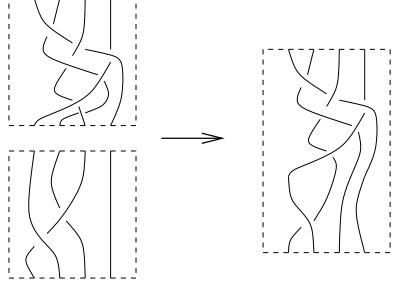

.
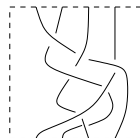
The set of n-strand braids, B_n, has a group structure with multiplication as follows. Given two n-strand braids b_1, b_2, place braid b_1 above b_2, remove the bottom b_1 bar and the top b_2 bar and fuse the bottom of the b_1 strands to the top of the b_2 strands.

The product of the above 4-strand braid with the 4-strand braid is:
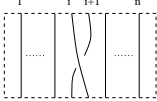

An algebraic presentation of the braid group due to Artin is as follows [3]: Let B_n be the group with generators $\{1, \sigma_1, \ldots \sigma_{n-1}\}$ and relations

1. $\sigma_i \sigma_j = \sigma_j \sigma_i$ for $|i - j| \geq 2$,

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

This algebraic description corresponds to the pictorial picture of braids: σ_i corresponds to the pictorial braid , and concatenating such pictures gives a general braid in B_n.

2.5 The Temperley-Lieb Algebras

Definition 2.2 *Given n an integer and d a complex number we define the Temperley-Lieb algebra $TL_n(d)$ to be the algebra generated by $\{1, E_1, \ldots, E_{n-1}\}$ with relations*

1. $E_i E_j = E_j E_i$, $|i - j| \geq 2$,

2. $E_i E_{i\pm1} E_i = E_i$,

3. $E_i^2 = d E_i$.

There is a well known geometric description of $TL_n(d)$ due to Kauffman [21]. It uses the notion of Kauffman n-diagrams, which is best explained by an example, e.g., a Kauffman 4-diagram:
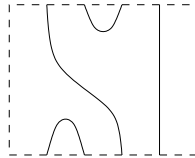
In general, a Kauffman n-diagram is a diagram as above, with n top pegs and n bottom pegs, and no crossings and no loops. More formally:
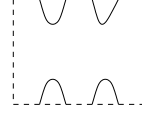
Definition 2.3 *Let D_n be a rectangle with n marked points on the top of the boundary and n marked points on the bottom. A Kauffman n-diagram is a picture sitting inside D_n consisting of n non-intersecting curves that begin and end at distinct marked boundary points. We will consider two such diagrams equal if they are isotopically equivalent (keeping the boundary fixed).*

We define a vector space over these diagrams:

Definition 2.4 *Let $\mathcal{K}_n$ be the vector space of complex linear combinations of Kauffman n-diagrams.*

Multiplication of two Kauffman n-diagrams is done just like in the case of braids. To multiply a diagram k_1 with a diagram k_2 we stack k_1 on top of k_2, and fuse the matching ends of the strands. The resulting diagram is a new Kauffman diagram k_3 with possibly some extra closed loops. We define the product of k_1 and k_2 to be $d^m k_3 \in \mathcal{K}_n$ where m is the number of extra closed loops. For example, a multiplication of the above Kauffman 4-diagram with the Kauffman 4-diagram results in:
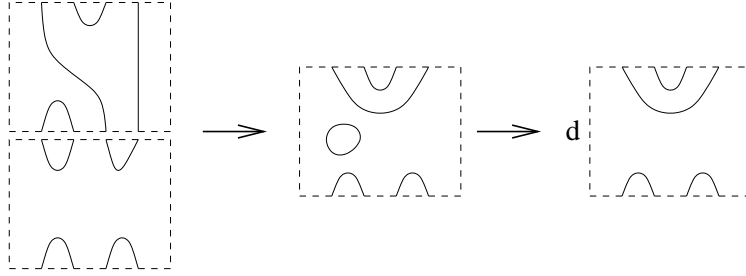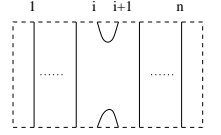

This multiplication rule can be extended linearly to $\mathcal{K}_n$, we call the resulting algebra $gTL_n(d)$.
The algebras $TL_n(d)$ and $gTL_n(d)$ are isomorphic:

Theorem 2.1 *The map $\psi : TL_n(d) \rightarrow gTL_n(d)$ given by the homomorphic extension of $\psi(E_i) =$*

is an isomorphism.

Proof: It is a simple and fun exercise to check that the image of the relations given in Definition 2.2 are relations in $gTL_n(d)$. For the remaining details see [6]. □

We shall refer to the pictures of the form , which generate $gTL_n(d)$, as *capcups.*

2.6 Representing B_n Inside $TL_n(d)$

We define a mapping from the braid group to $TL_n(d)$:

Definition 2.5 *$\rho_A : B_n \mapsto TL_n(d)$ is defined by*

$$\rho_A(\sigma_i) = AE_i + A^{-1}I.$$

Claim 2.1 *For a complex number A which satisfies $d = -A^2 - A^{-2}$, the mapping ρ_A is a representation of the braid group B_n inside $TL_n(d)$.*

Proof: We need to check that the relations of the braid group are satisfied by this mapping. For $|i - j| > 1$, $\rho_A(\sigma_i)$ commutes with $\rho_A(\sigma_j)$ since E_i commutes with E_j. To show that $\rho_A(\sigma_i)\rho_A(\sigma_{i+1})\rho_A(\sigma_i) = \rho_A(\sigma_{i+1})\rho_A(\sigma_i)\rho_A(\sigma_{i+1})$, substitute to get an expression in E_i's. Opening up the first expression we get $A^3 E_i E_{i+1} E_i + A E_{i+1} E_i + A E_i^2 + A^{-1} E_i + A E_i E_{i+1} + A^{-1} E_{i+1} + A^{-1} E_i + A^{-3}$. The second expression gives $A^3 E_{i+1} E_i E_{i+1} + A E_i E_{i+1} + A E_{i+1}^2 + A^{-1} E_{i+1} + A E_{i+1} E_i + A^{-1} E_i + A^{-1} E_{i+1} + A^{-3}$. We remove similar terms, and using the relations of the $TL_n(d)$ it remains to show that $(A^{-1} + Ad + A^3)E_i = (A^{-1} + Ad + A^3)E_{i+1}$. This holds because the constants are 0 due to the relation between d and A. □

2.7 Unitary Representation of B_n

Given a representation τ of $TL_n(d)$, we may use the representation of the braid group inside the $TL_n(d)$ algebra (Definition 2.5) to derive a representation of B_n by composition, as follows. Define the map φ by specifying its operation on the generators σ_i of B_n to be $\varphi(\sigma_i) = \varphi_i = \tau(\rho_A(\sigma_i)) = A\tau(E_i) + A^{-1}I$. This representation is unitary under certain constraints:

Claim 2.2 *If $|A| = 1$ and $\tau(E_i)$ are Hermitian for all i, then the map φ is a unitary representation of B_n.*

Proof: $\tau(\rho_A(\sigma_i))\tau(\rho_A(\sigma_i))^\dagger = (A^{-1}I + A\tau(E_i))((A^{-1})^*I + A^*\tau(E_i)^\dagger) = I + A^{-2}\tau(E_i) + A^2\tau(E_i) + d\tau(E_i) = I.$
$\square$

2.8 Tangles

For this paper, we define a tangle to be a braid in which some of its crossings have been replaced by a picture of the form { $\cup \atop \cap$ }. Braids and Kauffman diagrams are tangles.

2.9 From Braids to Links

We can connect up the endpoints of a braid in a variety of ways to get links. We single out two such ways:

Definition 2.6 *The trace closure of a braid B shall be the link achieved by connecting the strand at the rightmost top peg, around to the right, to the strand at the rightmost bottom peg, then connecting in the same way the next to rightmost top and bottom strands and so on. We denote the resulting link by B^{tr}.*

Definition 2.7 *The plat closure of a $2n$-strand braid shall be the link formed by connecting pairs of adjacent strands (beginning at the leftmost strand), on both the top and bottom. We denote the resulting link by B^{pl}.*

Examples of the trace closure and the plat closure of the same 4-strand braid are:

.
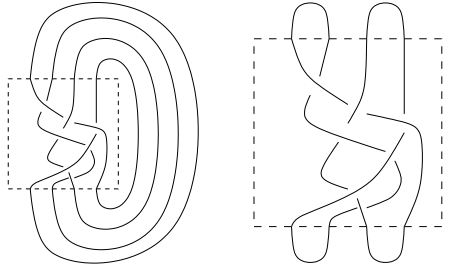
These closures are also well defined for tangles.

2.10 The Jones Polynomial

A definition of the Jones polynomial $V_L(t)$ due to Kauffman [21] is as follows. We start by defining the Kauffman bracket $< L >$ as a polynomial in A for A such that $A^{-4} = t$.

Definition 2.8 *Consider a link L, given by a link diagram. A state σ of L shall mean a choice, at each crossing $\times$ of L, from the set { $\cup \atop \cap$ $\,||\,$ }. To a state σ of a link L we associate the following expression $\sigma(L)$: Let σ^+ (respectively σ^-) be the number of crossings for which σ chooses $\cup \atop \cap$ (respectively $||$). Let $|\sigma|$ be the number of closed loops in the diagram gotten by replacing each crossing $\times$ by the choice*

8

indicated by the state σ. Define $\sigma(L) = A^{\sigma^+ - \sigma^-} d^{|\sigma|-1}$. The Kauffman bracket polynomial, also called the bracket state sum, for L, is defined to be

$$< L > = \sum_{all\ states\ \sigma} \sigma(L).$$

To define the Jones polynomial, we consider oriented links, namely links with one arrow on each connected component. The connection between the Jones polynomial and the Kauffman bracket is given by a notion called the *writhe*:

Definition 2.9 *For an oriented link L, assign to each crossing that looks like this* the value $+1$, *and*

to each crossing that looks like this: the value -1. *The writhe of L is the sum over all the crossings of these signs.*

Definition 2.10 *The Jones polynomial of an oriented link L is defined to be*

$$V_L(t) = V_L(A^{-4}) = (-A)^{3w(L)} \cdot < L >$$

where $w(L)$ is the writhe of the oriented link L, and $< L >$ is the bracket state sum of the link L, ignoring the orientation.

Thus, the Jones polynomial is a scaled version of the bracket polynomial. Moreover, the writhe of a link can be easily calculated from the link diagram, and hence the problem of calculating the bracket sum polynomial is equivalent in complexity to that of calculating the Jones polynomial.
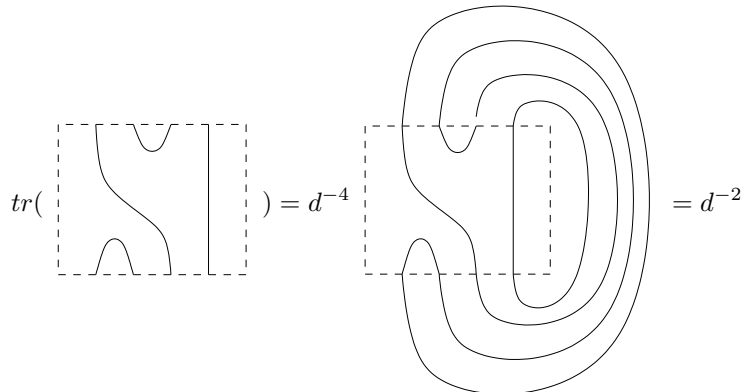
2.11 The Markov Trace

Definition 2.11 *A linear function from an algebra to the complex numbers is called a trace if it satisfies $tr(XY) = tr(XY)$ for every two elements X, Y in the algebra.*

We define the following trace on $gTL_n(d)$.

Definition 2.12 *The Markov trace $tr : gTL_n(d) \to \mathbb{C}$ is defined on a Kauffman n-diagram K as follows. Connect the top n labeled points to the bottom n labeled points of K with non-intersecting curves, as in the trace closure. Let a be the number of loops of the resulting diagram. Define $tr(K) = d^{a-n}$. Extend tr to all of $gTL_n(d)$ by linearity.*

For example:

$$tr(\quad) = d^{-4} \qquad = d^{-2}$$

Since $TL_n(d)$ and $gTL_n(d)$ are isomorphic, tr induces a trace on $TL_n(d)$; for simplicity we shall denote this map by tr as well.

Claim 2.3 *tr satisfies the following three properties:*

1. *$tr(1) = 1$,*

2. *$tr(XY) = tr(YX)$ for any $X, Y \in TL_n(d)$,*

3. *If $X \in TL_{n-1}(d)$ then $tr(XE_{n-1}) = \frac{1}{d}tr(X)$.*

Proof: It is straightforward to verify this by examining the appropriate pictures in $gTL_n(d)$. □

Of particular importance is the third property, which is referred to as the *Markov* property. These three properties uniquely determine a linear map on $TL_n(d)$:

Lemma 2.1 *[19] There is a unique linear function tr on $TL_n(d)$ (and on any representation of it) that satisfies properties $1 - 3$.*

Proof: By a reduced word $w \in TL_n(d)$ we shall mean a word in the set $\{1, E_1, \ldots E_{n-1}\}$ that is not equal to cw' for any c a constant and w' a word of smaller length. Using the relations of $TL_n(d)$ and applying simple combinatorial arguments, we show that a reduced word $w \in TL_n(d)$ contains at most one E_{n-1} term. We induct on n. Clearly the only reduced words in $Tl_2(d)$ are 1 and E_1. Assume the statement is true for reduced words in $Tl_{n-1}(d)$. Suppose there exists a reduced word $w \in TL_n(d)$ containing more than one E_{n-1} term. Write $w = w_1 E_{n-1} w_2 E_{n-1} w_3$ with w_2 a word without E_{n-1}. Since w_2 must be reduced and is in $Tl_{n-1}(d)$, the induction hypothesis implies w_2 contains at most one E_{n-2} term. If w_2 does not contain a E_{n-2} term, $w_2 \in Tl_{n-2}(d)$ and it commutes with E_{n-1} so we have $w = w_1 w_2 E_{n-1} E_{n-1} w_3$ which shows that w was not reduced. Otherwise we can write $w_2 = vE_{n-2}v'$ with v, v' both words in $TL_{n-2}(d)$. It follows therefore that v and v' commute with E_{n-1} and thus $w = w_1 v E_{n-1} E_{n-2} E_{n-1} v' w_3$ which again shows that w was not reduced. We conclude that any reduced word in $TL_n(d)$ contains at most one E_{n-1} term.

Given $w \in TL_n(d) \backslash Tl_{n-1}(d)$ a reduced word we write $w = w_1 E_{n-1} w_2$ with $w_1, w_2 \in Tl_{n-1}(d)$. Then $tr(w) = tr(w_2 w_1 E_{n-1}) = dtr(w_2 w_1)$, the first equality by property 2, the second by property 3. Thus for any word $w \in TL_n(d)$ we can reduce the trace computation to the trace of a word $w_2 w_1 \in Tl_{n-1}(d)$. Iterating this process, (and using the fact that $tr(1) = 1$), we see that the trace of a word in $TL_n(d)$ is uniquely determined by the relations 1.-3. Since the trace is linear, the result follows. □

We have the following convenient description of the Jones polynomial in terms of the Markov trace.

Lemma 2.2 *Given a braid B, then*

$$V_{B^{tr}}(A^{-4}) = (-A)^{3w(B^{tr})}d^{n-1}tr(\rho_A(B)).$$

Proof: By Definition 2.10, we need to show that $< B^{tr} > = tr(\rho_A(B))d^{n-1}$. We observe that there exists a one to one correspondence between states that appear in the bracket sum $< B^{tr} >$, and Kauffman n-diagrams that appear in $\rho_A(B)$. The weight of an element in the bracket state sum corresponding to the state σ is $A^{\sigma^+ - \sigma^-} d^{|\sigma|-1}$. We observe that the corresponding Kauffman n-diagram appears in $\rho_A(B)$ with the weight $A^{\sigma^+ - \sigma^-}$. Hence, by linearity of the trace, it remains to show that for each σ, the trace of the Kauffman diagram corresponding to σ, times d^{n-1}, equals to the remaining factor in the contribution of σ to the bracket state sum, $d^{|\sigma|-1}$. This is true since by the definition of the trace of a Kauffman diagram, it is exactly $d^{|\sigma|-n}$. □

This lemma also holds if B is replaced by a tangle.

## 2.12   The Path Model Representation of $TL_n(d)$

We describe the path model representation of $TL_n(d)$ due to [20, 18].   The representation will act on a vector space determined by paths on a graph. Specifically, given an integer k (k will be chosen in relation to d later), let G_k be the straight line graph with $k-2$ segments and $k-1$ vertices:
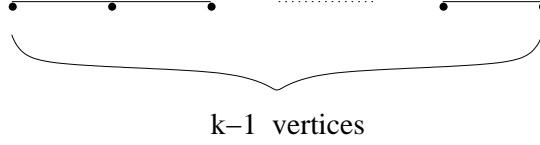
k–1 vertices

Define $Q_{n,k}$ to be the set of all paths of length n on the graph G_k beginning at the leftmost vertex. Given $q \in Q_{n,k}$, we shall denote by $q(0), q(1), \ldots q(n)$ the sequence of vertices of G_k describing q; thus q(0) is the leftmost vertex and $q(i)$ and $q(i+1)$ are adjacent vertices of G_k for all i. We shall think of the elements of $Q_{n,k}$ as an orthonormal basis of a vector space $\mathcal{V}_{n,k}$; hence an element $q \in Q_{n,k}$ shall represent both a path on G_k and a basis element of $\mathcal{V}_{n,k}$. We now construct the path model representation $\tau(TL_n(d)) : \mathcal{V}_{n,k} \to \mathcal{V}_{n,k}$.

Given a Kauffman n-diagram T, to describe $\tau(T)$ it will suffice to give the matrix entry $\tau(T)_{q',q}$ for each pair $q', q \in Q_{n,k}$. To do this, we note that the strands of a Kauffman diagram separate the rectangle into regions; we would like to label the regions by vertices of G_k, such that the labeling of the bottom part of T will correspond to q and the top part to q', and then compute the matrix element $\tau(T)_{q',q}$ from the labels. This is done as follows.

The n marked points of a Kauffman n-diagram divide the top and bottom boundary into $n+1$ segments which we shall refer to as *gaps*. We shall say a set of gaps that bound the same region in the diagram are *connected*. For example, in the following Kauffman 3-diagram:
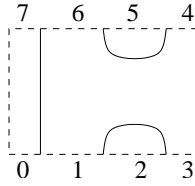
the set of gaps $\{0, 7\}$ are connected as are the set of gaps $\{1, 3, 4, 6\}$. We shall say that the pair (q', q) is *compatible* with T if once we label the gaps on the bottom from left to right by $q(0), q(1), \ldots, q(n)$ and we label the gaps on the top from left to right by $q'(0), q'(1), \ldots, q'(n)$, then any set of connected gaps are all labeled by the same vertex of G_k. Thus in this case each region of T can be thought of as being labeled by a single vertex of G_k.

The matrix entry $\tau(T)_{q',q}$ will only be nonzero in case the pair of paths (q', q) is *compatible* with T. In this case, the regions are indeed labeled by vertices in G_k; we can now do the following. To each local maximum and minimum of the Kauffman diagram T, we associate a complex number that depends on the labeling of the regions that surround them as follows:

$$\underset{l}{\overset{l+1}{\smile}} \mapsto a_\ell \ , \quad \underset{l}{\overset{l-1}{\smile}} , \mapsto b_\ell \quad \underset{l+1}{\overset{l}{\frown}} \ \mapsto c_\ell, \quad \underset{l-1}{\overset{l}{\frown}} \ \mapsto d_\ell$$

The matrix element $\tau(T)_{q',q}$ at a compatible pair (q', q), is defined to be the product of the appropriate complex numbers over all local maxima and minima in T.

For the map $\tau(T)$ described above to be well-defined, it has to give the same result for isotopic Kauffman diagrams. An isotopic move can be seen to only create or eliminate local maxima and minima in pairs; we see that the conditions

$$b_{\ell+1} c_\ell = 1 = a_{\ell-1} d_\ell \tag{1}$$

are necessary and sufficient for the map to be isotopically invariant. A single extra constraint is needed to produce a representation of $TL_n(d)$:

Claim 2.4 *If the coefficients* $a_\ell, b_\ell, c_\ell, d_\ell$ *satisfy (1) and*

$$d = b_\ell d_\ell + a_\ell c_\ell, \tag{2}$$

then the resulting map τ defined as above is a representation of $TL_n(d)$.

Proof: To prove the result we need only verify that the matrices $\tau(E_i)$ satisfy the relations of Definition 2.2. This amounts to verifying that the matrix elements of the operators on both sides of each relation are equal. Pictorially, a matrix element in a product of operators is given by stacking the operators together, and summing up over all possible labeling inside loops, such that the label inside the loop is different by exactly one from the label outside the loop. This summation corresponds to the summation over the intermediate index in matrix multiplication.

   We now check for the different relations in Definition 2.2. For the generator E_i, there are only four types of non-zero elements, namely, four types of compatible pairs corresponding to the following types of labeling of the regions near the i-th strand:

$$\begin{array}{cccc} \overset{\overbrace{l-1}}{l} & \overset{\overbrace{l+1}}{l} & \overset{\overbrace{l-1}}{l} & \overset{\overbrace{l+1}}{l} \\ \underset{\underbrace{l+1}}{} & \underset{\underbrace{l+1}}{} & \underset{\underbrace{l-1}}{} & \underset{\underbrace{l-1}}{} \end{array}$$

   In the first two relations no loops are created when the operators are multiplied, so the verification follows from the isotopy invariance of τ (i.e., Equation (1)). The third relation follows from Equation (2), using the fact that one loop was created and there are two possible ways to label the region inside. □

   We would like $\tau(E_i)$ to be Hermitian, so that the induced representation on B_n is unitary by Claim 2.2. For this we add the constraints

$$a_\ell = c_\ell^*, b_\ell = d_\ell^* \tag{3}$$

Claim 2.5 *If the coefficients $a_\ell, b_\ell, c_\ell, d_\ell$ satisfy Equations (1)-(3) then $\tau(E_i) = \tau(E_i)^\dagger$ for all i.*

Proof: We need to prove that $\tau(E_i)_{q,q'} = (\tau(E_i)_{q',q})^*$. We need to check only for compatible pairs, namely for q, q' which on the $i-1, i, i+1$ entry are equal to $\ell, \ell \pm 1, \ell$ for some ℓ. This amounts to checking that $a_\ell d_\ell = b_\ell^* c_\ell^*$, $b_\ell d_\ell = b_\ell^* d_\ell^*$, and $a_\ell c_\ell = a_\ell^* c_\ell^*$. This follows from (3). □

   It is now left to solve Equations $1-3$ to derive the definition of τ.

Claim 2.6 *Define $\lambda_\ell = sin(\pi\ell/k)$ for $\ell \in \{1, ..., k-1\}$. Then $a_\ell = c_\ell^* = \sqrt{\frac{\lambda_\ell}{\lambda_{\ell-1}}}$, $b_\ell = d_\ell^* = \sqrt{\frac{\lambda_\ell}{\lambda_{\ell+1}}}$ satisfy Equations (1)-(3), with $d = 2cos(\pi/k)$.*

Proof: We solve the equations directly. Define $x_\ell = \frac{a_\ell}{b_{\ell+1}}$, and solve for x_ℓ. (x_ℓ is defined for $\ell \in \{1, \ldots, k-2\}$). We have $\frac{1}{x_{\ell-1}} + x_\ell = d$, and if we set $x_0 \ldots x_\ell = y_\ell$ (and $y_\ell = 0$ whenever undefined) we have $y_{\ell-2} + y_\ell = dy_{\ell-1}$. This is the familiar equation, defining the eigenvector of the adjacency matrix of the graph G_k

$$\begin{pmatrix} 0 & 1 & \cdots & & & 0 \\ 1 & 0 & 1 & & & \\ & 1 & & & & \\ \vdots & & & \ddots & & \\ & & & & 0 & 1 \\ 0 & & \cdots & & 1 & 0 \end{pmatrix},$$

with eigenvalue d. It is not difficult to check that the $(k-1)$-dimensional vector λ is exactly this eigenvector. We have $x_\ell = y_\ell / y_{\ell-1}, x_\ell = a_\ell^2$, which gives $a_\ell = \sqrt{\lambda_\ell / \lambda_{\ell-1}}$. The Equations $b_\ell = 1/a_{\ell-1}$ and (3) determine b_ℓ, c_ℓ, d_ℓ. $\square$

Using these coefficients, we get the definition of $\tau(E_i)$ as follows. $\tau(E_i)_{q,q'} = 0$ if (q, q') is not compatible with E_i. For a compatible pair (q, q'), $\tau(E_i)_{q,q'}$ is the product of two coefficients, one corresponding to the maximum and the other to the minimum in E_i. For example, if $q = q'$ both move from the site ℓ to $\ell+1$ in the ith step, and then return to the site ℓ in the $i+1$th step, we have $\tau(E_i)_{q,q} = a_\ell c_\ell$. Given $\tau(E_i)$, we can extend our definition of τ to all elements in $TL_n(d)$.

2.13   Unitary Path Model Representation of B_n

The previous section provided a representation τ of $TL_n(d)$ defined for $d = 2cos(\pi/k)$. We have that $d = -A^2 - A^{-2}$ for $A = ie^{-\pi/2k}$, and also that $\tau(E_i)$ are Hermitian. Thus, the conditions of Subsection 2.7 are satisfied. We define:

Definition 2.13 *The unitary path model representation of B_n is defined to be $\varphi(B) = \tau(\rho_A(B))$.*

The map φ can be extended to operate on tangles by letting ρ_A be applied only to the crossings in the tangle.

3   The Quantum Algorithm

We are now ready to prove Theorems 1.1 and 1.2. We first translate the path model representation to work on qubits, and show that it can be implemented efficiently. We use this to design the algorithms, and then prove their correctness.

3.1   Moving to Qubits

The adaptation of the path model representation to qubits is fairly straightforward. We simply switch from presenting paths by the list of their locations, to a binary representation which indicates the *direction* of each step. Thus, we shall interpret a string of n bits to be a sequence of instructions, where a 0 shall mean take one step to the left and a 1 shall mean take one step to the right. We shall restrict our attention to those n bit strings that describe a path that starts at the leftmost vertex of G_k and remains inside G_k at each step. From here on when we say "path" we actually mean the bit string that represents the path.

Definition 3.1 *We define $P_{n,k,\ell}$ to be the set of all paths p on G_k of n steps which start at the left most site and end at the ℓ's site. We define the subspace $\mathcal{H}_{n,k,\ell}$ to be the span of $|i\rangle$ over all $i \in P_{n,k,\ell}$. In a similar way, we define $P_{n,k}$ to be all paths with no restriction on the final point, i.e., $P_{n,k} = \cup_{l=1}^{k} P_{n,k,l}$, and we define $\mathcal{H}_{n,k}$ to be the span of the corresponding computational basis states.*

We define a representation Φ as a homomorphism from $TL_n(d)$ to matrices operating on $\mathcal{H}_{n,k}$. To define Φ it suffices to specify the images of the E_i's, $\Phi(E_i) = \Phi_i$. The operators Φ_i are defined so that they correspond to the operators $\tau(E_i)$ (see Subsection 2.12) via the natural isomorphism between $\mathcal{V}_{n,k}$ and $\mathcal{H}_{n,k}$. Thus, the transition from τ to Φ is merely a change of language. This is done as follows.

To uniquely define Φ_i on $\mathcal{H}_{n,k}$, it suffices to define what it does to each basis element, namely, to $|p\rangle$ for $p \in P_{n,k}$. We need the following notation:

Definition 3.2 *Let $p|_i$ denote the restriction of a path p to its first $i-1$ coordinates. Given a path p on G_k, we denote by $\ell(p) \in \{1, ..., k-1\}$ the location in G_k that the path p reached in its final site. Denote $z_i = \ell(p|_i)$.*

We can now define the operation of Φ_i. Φ_i is defined as an operation on the first $i+1$ coordinates in a path p:

$$
\begin{aligned}
\Phi_i|p|_i00\rangle &= 0 \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad (4)\\
\Phi_i|p|_i01\rangle &= \frac{\lambda_{z_i-1}}{\lambda_{z_i}}|p|_i01\rangle + \frac{\sqrt{\lambda_{z_i+1}\lambda_{z_i-1}}}{\lambda_{z_i}}|p|_i10\rangle\\
\Phi_i|p|_i10\rangle &= \frac{\lambda_{z_i+1}}{\lambda_{z_i}}|p|_i10\rangle + \frac{\sqrt{\lambda_{z_i+1}\lambda_{z_i-1}}}{\lambda_{z_i}}|p|_i01\rangle\\
\Phi_i|p|_i11\rangle &= 0
\end{aligned}
$$

To apply Φ_i on the n-bit string $|p\rangle$ we tensor the above transformation with identity on the last $n-i-1$ qubits. For dealing with the edge cases, we use the convention $\lambda_j = 0$ for any $j \notin \{1,..,k-1\}$.

Once we have defined Φ_i, Φ is then extended to the entire algebra by the multiplication property of a representation, and by linearity.

Claim 3.1 Φ *is a representation of* $TL_n(d)$. Φ_i *are Hermitian. Moreover,* Φ *induces a unitary representation of the Braid group* B_n, *operating on* $\mathcal{H}_{n,k}$.

Proof: The proof follows from the corresponding properties of τ, namely Subsections 2.12, 2.13, and the natural isomorphism between paths on G_k presented by their sequence of locations, $\mathcal{V}_{n,k}$, and paths presented as bit strings, namely $\mathcal{H}_{n,k}$. $\square$

With the risk of confusion, we denote the unitary representation of B_n induced by Φ, also by φ as in Definition 2.13. The only difference is that now $\varphi(B)$ operates on $\mathcal{H}_{n,k}$ rather than on $\mathcal{V}_{n,k}$.

3.2   Efficient application of one crossing

The matrices φ_i are defined so far only on $\mathcal{H}_{n,k}$ which is a subspace of the Hilbert space of n qubits; we arbitrarily define their extension to the rest of the Hilbert space to be the identity.

Claim 3.2 *For all* $i \in \{1,...,n\}$, φ_i *can be implemented on the Hilbert space of* n *qubits using* $\mathrm{poly}(n,k)$ *gates.*

Proof: We note that the application of φ_i on $p \in P_{n,k}$ modifies only the $i,i+1$ bits of p, and the modification depends on the location up to the ith step, namely on $z_i = \ell(p|_{i-1})$. z_i is a number which can be calculated efficiently and written on $O(log(k))$ ancilla qubits, using the following standard technique: Initialize a counter register of, say, $log(2k)$ qubits to the value 1. Then move along the qubits of the path p from left to right, and for each of the n qubits update the current state of the counter, ℓ, by applying

$$|b\rangle|\ell\rangle \mapsto |b\rangle|\ell + (-1)^b \mathrm{mod} 2k\rangle$$

where b is the state of the currently read qubit. Since this is a unitary operation on $log(2k)+1$ qubits, it can be applied using polynomially in k many elementary quantum gates (this is a standard result in quantum computation). We end up with the extra register carrying $\ell(p|_{i-1})$.

Now, φ_i depends only on the location $\ell(p|_{i-1})$ and on the i and $i+1$ qubits. Hence, once again we have a unitary transformation which operates on logarithmically in k many qubits, and so we can implement it in polynomially in k many quantum gates.

After we apply φ_i, we erase the calculation of $\ell(p|_{i-1})$ by applying the inverse of the first transformation which wrote the location down. $\square$

As a corollary, we can deduce that

14

Corollary 3.1 *For every braid $B \in B_n$, with m crossings, there exists a quantum circuit $Q(B)$ that applies $\varphi(B)$ on n qubits, using $\mathrm{poly(m, n, k)}$ elementary gates.*

Proof: Order the crossings in the braid in topological order, and apply the corresponding unitary matrix of each crossing, φ_i, one by one, in that order. Each crossing takes $\mathrm{poly}(n, k)$ elementary gates by Claim 3.2, and there are m of them. □

3.3 The Algorithms

We can now describe the algorithms. The input for both is a braid of n strands and m crossings, and an integer k.

Algorithm Approximate-Jones-Plat-Closure

- Repeat for $j = 1$ to $poly(n, m, k)$:

    1. Generate the state $|\alpha\rangle = |1, 0, 1, 0, \ldots, 1, 0\rangle$

    2. Output a random variable x_j whose expectation value is $\mathcal{R}e\langle\alpha|Q(B)|\alpha\rangle$ using the Hadamard test.

- Do the same but for random variables y_j whose expectation value is $\mathcal{I}m\langle\alpha|Q(B)|\alpha\rangle$ using the appropriate variant of the Hadamard test.

- Let r be the average over all $x_j + iy_j$ achieved this way. Output $(-A)^{3w(B^{pl})}d^{3n/2-1}\lambda_1 r/N$.

Algorithm Approximate-Jones-Trace-Closure

- Repeat for $j = 1$ to $poly(n, m, k)$:

    1. Classically, pick a random path $p \in P_{n,k}$ with probability $Pr(p) \propto \lambda_\ell$, where ℓ is the index of the site which p ends at.

    2. Output a random variable x_j whose expectation value is $\mathcal{R}e\langle p|Q(B)|p\rangle$ using the Hadamard test.

- Do the same but for random variables y_j whose expectation value is $\mathcal{I}m\langle p|Q(B)|p\rangle$ using the appropriate variant of the Hadamard test.

- Let r be the average over all $x_j + iy_j$. Output $(-A)^{3w(B^{tr})}d^{n-1}r$.

Claim 3.3 *The above two quantum algorithms can be performed in time polynomial in n, m, k.*

Proof: Algorithm Approximate-Jones-Plat-Closure is clearly efficient, because the Hadamard test can be applied efficiently using Corollary 3.1. To perform the first step of the second algorithm efficiently, we pick a random $\ell \in \{1, \ldots, k\}$, with probability proportional to λ_ℓ, and then use Claim 3.4 below. We note that our calculations involve irrational numbers, λ_ℓ, but these can be approximated to within an exponentially good precision efficiently. □

Claim 3.4 *Given n, k, ℓ, there exists a classical probabilistic algorithm which runs in time polynomial in n and k, and outputs a random path in $P_{n,k,\ell}$ according to a distribution which is exponentially close to uniform.*

Proof: (We thank O. Regev for a discussion that lead to this variant.) We briefly sketch the proof. We first note that the following $n \times k$ array S, defined by $S_{i,j} = |P_{i,k,j}|$ (the number of path on G_k of i steps that start at 1 and end at j) can be easily calculated efficiently using the recursion $|P_{n,k,\ell}| = |P_{n-1,k,\ell-1}| + |P_{n-1,k,\ell+1}|$. To pick a random path ending at ℓ, we pick the values of the sites in p one by one in reverse order. If $p(j) = \ell$, we randomly decide if $p(j-1) = \ell - 1$ or $\ell + 1$ according to the ratio $S_{j-1,\ell-1} : S_{j-1,\ell+1}$. Given $p(j-1)$ we can continue to pick $p(j-2)$ in a similar manner, and so on. $\square$

Theorem 3.1 *Algorithm* Approximate-Jones-Trace-Closure *approximates the Jones polynomial of B^{tr} at $A^{-4} = e^{2\pi i/k}$, to within the precision specified in Theorem 1.1.*

Proof: We will need the following definition.

Definition 3.3 *Define $Tr_n(W)$ for every W in the image of $\Phi(TL_n(d))$ to be:*

$$Tr_n(W) = \frac{1}{N} \sum_{\ell=1}^{k-1} \lambda_\ell Tr(W|_\ell)$$

where $W|_\ell$ denotes the restriction of W to the subspace $\mathcal{H}_{n,k,\ell}$, and Tr denotes the standard trace on matrices. The renormalization is $N = \sum_\ell \lambda_\ell dim(\mathcal{H}_{n,k,l})$ where the sum is taken over all ℓ's such that $P_{n,k,\ell}$ is non empty.

This definition makes sense because matrices in the image of φ are block diagonal, with the blocks indexed by ℓ, the last site of the paths:

Claim 3.5 *For any $T \in TL_n(d)$, $\Phi(T)\mathcal{H}_{n,k,\ell} \subseteq \mathcal{H}_{n,k,\ell}$.*

Proof: Φ_i cannot change the final point of a path since it only moves 01 to 10 and vice versa. $\square$

Hence, the above trace function simply gives different weights to these blocks (and gives zero weights on strings that aren't paths). We claim that Tr_n is a Markov trace.

Claim 3.6 *The function $Tr_n(\cdot)$ satisfies the three properties in Claim 2.3.*

Proof: That $Tr_n(\Phi(1)) = 1$ follows from the renormalization. The second property follows from Claim 3.5 plus the fact that the standard trace on matrices satisfies this property, so $Tr_n(\cdot)$ satisfies it on each block separately. To show the Markov property, we have to show that if $X \in TL_{n-1}(d)$) then $Tr_n(\Phi(X)\Phi(E_{n-1})) = \frac{1}{d}Tr_n(\Phi(X))$. We note that for any $X \in TL_{n-1}(d)$, $\Phi(X)$ can be written as a linear combination of terms of the form $|p\rangle\langle p'| \otimes I$, with $p, p' \in P_{n-1,k}$ and the identity operates on the last qubit. By linearity, it suffices to prove the Markov property on such matrices. Writing $|p\rangle\langle p'| \otimes I = |p0\rangle\langle p'0| + |p1\rangle\langle p'1|$ we require: $Tr_n(|p0\rangle\langle p'0|\Phi_{n-1} + |p1\rangle\langle p'1|\Phi_{n-1}) = \frac{1}{d}Tr_n(|p0\rangle\langle p'0| + |p1\rangle\langle p'1|)$. This can be easily verified using the definition of Φ by checking the two cases $p = p'$ and $p \neq p'$.

We start with the case $p \neq p'$. In this case the right hand side is 0. As for the left hand side, $\langle p'0|\Phi_{n-1}$ has a zero component on $\langle p0|$. To see this, we check the two cases: if p' ends with 0 then $\langle p'0|\Phi_{n-1} = 0$, otherwise p' ends with 1. $\langle p'0|\Phi_{n-1}$ is then a sum of two terms, one equals to $\langle p'|$ and is therefore different than $\langle p0|$, and the other ends with 01 and is thus also different from $\langle p0|$. The same argument works to show that $\langle p'1|\Phi_{n-1}$ has a zero component on $\langle p1|$. Hence, the left hand side is also 0.

It is left to check the equality in the case $p = p'$. We require

$$Tr_n(|p0\rangle\langle p0|\Phi_{n-1} + |p1\rangle\langle p1|\Phi_{n-1}) = \frac{1}{d}Tr_n(|p0\rangle\langle p0| + |p0\rangle\langle p0|).$$

Suppose $\ell(p) = \ell$. Then the right hand side is equal to $\frac{1}{d}(\lambda_{\ell-1} + \lambda_{\ell+1}) = \lambda_\ell$, using the properties of the eigenvector λ, as in Claim 2.6. To see that the left hand side is the same, we again divide to cases. Suppose first that p ends with 0. In this case $\langle p0|\Phi_{n-1} = 0$. As for the other term, $\langle p1|\Phi_{n-1} = \frac{\lambda_\ell}{\lambda_{\ell+1}}\langle p1|$, using the

16

definition of Φ_{n-1} and the fact that p without its last step ends in $\ell + 1$. The weight in the trace of the left hand side is $\lambda_{\ell+1}$, and so the left hand side is equal to λ_ℓ too. The argument is similar in the case that p ends with 1. $\square$

By the uniqueness of the Markov trace, Lemma 2.1, we have that $Tr_n(\varphi(B)) = tr(\rho_A(B))$. Hence, using Lemma 2.2, we have that for any braid $B \in B_n$

Lemma 3.1
$$V_{B^{tr}}(A^{-4}) = (-A)^{3w(B^{tr})} d^{n-1} Tr_n(\varphi(B)).$$

Due to Lemma 3.1, the correctness of the algorithm follows trivially from the following claim.

Claim 3.7 *With all but exponentially small probability, the output r satisfies $|r - Tr_n(\varphi(B))| \leq \varepsilon$ for ε which is inverse polynomial in n, k, m.*

Proof: The Hadamard test indeed implies that the expectation of x_j for a fixed p is exactly $\mathcal{R}e\langle p|\varphi(B)|p\rangle$. The expectation of the variable x_j taken over a random p is thus

$$\frac{\sum_{\ell,p \in P_{n,k,\ell}} \lambda_\ell \mathcal{R}e(\langle p|\varphi(B)|p\rangle)}{\sum_{\ell,p \in P_{n,k,\ell}} \lambda_\ell} = \frac{\sum_\ell \lambda_\ell \mathcal{R}e(Tr(\varphi(B)|_\ell))}{\sum_\ell \lambda_\ell dim(H_{n,k,l})} =$$

$$= \mathcal{R}e(Tr_n(\varphi(B))).$$

The same argument works for the imaginary part. Since r is the sum of two averages of polynomially many $i.i.d$ random variables, each taking values between 1 and -1, the result follows by the Chernoff-Hoeffding bound. $\square$

This completes the proof of Theorem 3.1. $\square$

Theorem 3.2 *Algorithm* Approximate-Jones-Plat-Closure *approximates the Jones polynomial of B^{pl} at $A^{-4} = e^{2\pi i/k}$, to within the desired precision as in Theorem 1.2.*

Proof: By the correctness of the Hadamard test, and by the Chernoff-Hoeffding bound, the variable r which the algorithm computes is, with exponentially good confidence, within $\delta = 1/poly(n, m, k)$ from $\langle \alpha|\varphi(B)|\alpha\rangle$, which is equal to $Tr(\varphi(B)|\alpha\rangle\langle\alpha|)$. We need to connect this expression to the Jones polynomial of the plat closure of B. The main observation here is that the plat closure of a braid B is isotopic to the trace closure of a tangle C achieved by applying the braid on $n/2$ capcups, as in the following picture:
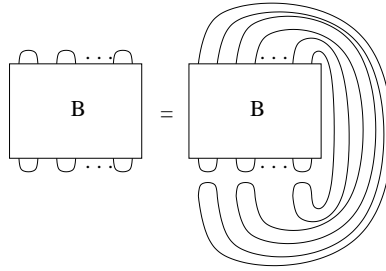
It therefore suffices to relate the Jones polynomial of C^{tr} to $Tr(\varphi(B)|\alpha\rangle\langle\alpha|)$. Since the question is now stated in terms of trace closures and traces, there is hope to be able to apply Lemma 2.2 and Claim 3.6 as in the proof of Theorem 3.1. But we first need to make the connection between the projection on $|\alpha\rangle$ and capcups.

Claim 3.8 $|\alpha\rangle\langle\alpha| = \Phi_1\Phi_3\ldots\Phi_{n-1}/d^{n/2}$.

Proof: It is easy to verify that $\Phi_1\Phi_3\ldots\Phi_{n-1}$ applied to any path except for $|\alpha\rangle$ gives 0, and when applied to $|\alpha\rangle$ it gives the desired factor. To do this we use the fact that Φ_i commute if their indices are more than one apart, and so we can first apply Φ_1, then Φ_3 and so on. Since the path starts at the left most site, Φ_1 on p simply applies the following rescaled projection: $d|10\rangle\langle10|$ on the first two coordinates. This projection forces the first two coordinates to be 10, and so $\Phi_1|p\rangle$ returns to the starting point after two steps. Therefore a similar argument applies when we apply Φ_3 on the next two coordinates, and so on. By induction, we get the desired result. $\square$

We thus have, using Definition 3.3 and Claim 3.8:

$$\langle\alpha|\varphi(B)|\alpha\rangle = Tr(\varphi(B)|\alpha\rangle\langle\alpha|) = \frac{N}{\lambda_1}Tr_n(\varphi(B)|\alpha\rangle\langle\alpha|) =$$

$$= \frac{N}{\lambda_1}Tr_n(\varphi(B)\Phi_1\Phi_3\ldots\Phi_{n-1}/d^{n/2}) = \frac{N}{\lambda_1 d^{n/2}}Tr_n(\varphi(C))$$

By the uniqueness of the Markov trace, Lemma 2.1, and by Claim 3.6, we have that $Tr_n(\varphi(C)) = tr(\rho_A(C))$. Using Lemma 2.2, we have:

$$V_{B^{pl}}(A^{-4}) = V_{C^{tr}}(A^{-4}) = (-A)^{3w(C^{tr})}d^{n-1}Tr_n(\varphi(c)).$$

which we substitute in the previous equation. This completes the proof of Theorem 3.2. $\square$

4   Acknowledgements

References

[1] Aharonov D. and Arad I., On the BQP-hardness of Approximating the Jones Polynomial, preprint, 2006

[2] Alexander, J. W. Topological invariants of knots and links. Trans. Amer. Math. Soc. 30 (1928), no. 2, 275–306.

[3] Artin, E. (1947). Theory of braids. *Annals of Mathematics*, **48** 101–126.

[4] Bernstein E and Vazirani U, Quantum complexity theory, SIAM Journal of Computation 26 5 pp 1411-1473 October, 1997

[5] Birman, J. (1974). Braids, links and mapping class groups. *Annals of Mathematical Studies*, **82**.

[6] D. Bisch and V. Jones, *A*lgebras associated to intermediate subfactors, Invent. Math. 128 (1997), 89-157.

[7] Bordewich M, Freedman M, Lovasz L, and Welsh D., Approximate counting and Quantum computation, Combinatorics, Probability and Computing, **14**, Issue 5-6, pp: 737 − 754, 2005

[8] A. Childs and R. Cleve and E. Deotto and E. Farhi and S. Gutmann and D. Spielman, Exponential algorithmic speedup by quantum walk, STOC 2003

[9] J.H. Conway *An enumeration of knots and links,and some of their algebraic properties.* Computational Problems in Abstract Algebra (Proc. Conf., Oxford, 1967) (1970) 329–358

[10] W. van Dam and S. Hallgren, Efficient quantum algorithms for shifted quadratic character problems. In quant-ph/0011067.

[11] M. Freedman, *P*/NP and the quantum field computer, Proc. Natl. Acad. Sci., USA, **95**, (1998), 98–101

[12] M.Freedman, A.Kitaev, M. Larsen, Z. Wang, *Topological quantum computation. Mathematical challenges of the 21st century* (Los Angeles, CA, 2000). Bull. Amer. Math. Soc. (N.S.) 40 (2003), no. 1, 31–38

[13] M. H. Freedman, A. Kitaev, Z. Wang *S*imulation of topological field theories by quantum computers Commun.Math.Phys. 227 (2002) 587-603

[14] M. H. Freedman, M. Larsen, Z. Wang *A modular Functor which is universal for quantum computation* Commun.Math.Phys. 227 (2002) no. 3, 605-622

[15] S. Hallgren, Polynomial-time quantum algorithms for Pell's Equation and the principal ideal problem. STOC 2002, pp. 653–658.

[16] Jaeger, F.; Vertigan, D. L.; Welsh, D. J. A. On the computational complexity of the Jones and Tutte polynomials. Math. Proc. Cambridge Philos. Soc. 108 (1990), no. 1, 35–53.

[17] Mark Jerrum, Alistair Sinclair and Eric Vigoda, A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries, STOC 2001, 712-721.

[18] V.F.R Jones, *A polynomial invariant for knots via von Neumann algebras.* Bull. Amer. Math. Soc. 12 (1985), no. 1 103–111.

[19] V.F.R. Jones, *Index for subfactors*, Invent. Math **72** (1983), 1–25.

[20] V.F.R. Jones, Braid groups, Hecke Algebras and type II factors, in *G*eometric methods in Operator Algebras, Pitman Research Notes in Math., **123** (1986), 242–273

[21] L.Kauffman, *State models and the Jones polynomial.* Topology **26**,(1987),395-407.

[22] G. Kuperberg, A subexponential-time quantum algorithm for the dihedral hidden subgroup problem, arXiv:quant-ph/0302112.

[23] Neilsen, Chuang, Quantum Computation and Quantum Information, Cambridge press, 2000

[24] A. Podtelezhnikov, N. Cozzarelli, A. Vologodskii, Equilibrium distributions of topological states in circular DNA: interplay of supercoiling and knotting. Proc. Natl. Acad. Sci. USA 96 (1999), no. 23, 12974–12979.

[25] Preskill J., Topological quantum computation, Lecture notes for Caltech course # 219 in Physics, *http* : *//www.theory.caltech.edu/ preskill/ph*229*/#lecture*

[26] V. Subramaniam, P. Ramadevi, Quantum Computation of Jones' Polynomials, quant-ph/0210095

[27] Simon D, On the power of quantum computation, SIAM J. Comp., 26, No. 5, pp 1474-1483, October 1997

[28] P. W. Shor: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM J. Comput. 26(5) 1997, pp. 1484–1509.

[29] P. Vogel, representation of links by braids: A new algorithm, In *Comment. math. Helvetici*, **65** (1) pp. 104–113, 1990

[30] J. Watrous, Quantum algorithms for solvable groups. STOC 2001, pp. 60–67.

[31] D. J. A. Welsh, "The Computational Complexity of Some Classical Problems from Statistical Physics," Disorder in Physical Systems, Clarendon Press, Oxford, 1990, pp. 307-321.

[32] E. Witten, Quantum field theory and the Jones polynomial. Comm. Math. Phys. 121 (1989), no. 3, 351–399.

[33] F. Y. Wu, Knot Theory and statistical mechanics, Rev. Mod. Phys. **64**, No. 4., October 1992