

Simulation of Topological Field Theories by Quantum Computers

Michael H. Freedman¹, Alexei Kitaev^{1,*}, Zhenghan Wang²

¹ Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399, USA

² Indiana University, Dept. of Math., Bloomington, IN 47405, USA

Received: 4 May 2001 / Accepted: 16 January 2002

Abstract: Quantum computers will work by evolving a high tensor power of a small (e.g. two) dimensional Hilbert space by local gates, which can be implemented by applying a local Hamiltonian H for a time t . In contrast to this quantum engineering, the most abstract reaches of theoretical physics has spawned “topological models” having a finite dimensional internal state space with no natural tensor product structure and in which the evolution of the state is discrete, $H \equiv 0$. These are called topological quantum field theories (TQFTs). These exotic physical systems are proved to be efficiently simulated on a quantum computer. The conclusion is two-fold:

1. TQFTs cannot be used to define a model of computation stronger than the usual quantum model “BQP”.
2. TQFTs provide a radically different way of looking at quantum computation. The rich mathematical structure of TQFTs might suggest a new quantum algorithm.

1. Introduction

A *topological quantum field theory* (TQFT) is a mathematical abstraction, which codifies topological themes in conformal field theory and Chern–Simons theory. The strictly 2-dimensional part of a TQFT is called a *topological modular functor* (TMF). It (essentially) assigns a finite dimensional complex Hilbert space $V(\Sigma)$ to each surface Σ and to any (self)-diffeomorphism h of a surface a linear (auto)morphism $V(h) : V(\Sigma) \rightarrow V(\Sigma')$. We restrict attention to unitary topological modular functors (UTMF) and show that a quantum computer can efficiently simulate transformations of any UTMF as a transformation on its computational state space. We should emphasize that both sides of our discussion are at present theoretical: the quantum computer which performs our simulation is also a mathematical abstraction – the *quantum circuit model* (QCM) [D, Y].

* On leave from Landau Institute for Theoretical Physics, Moscow.

Very serious proposals exist for realizing this model, perhaps in silicon, e.g. [Ka], but we will not treat this aspect.

There is a marked analogy between the development of the QCM from 1982 Feynman [Fey] to the present, and the development of recursive function theory in the 1930’s and 1940’s. At the close of the earlier period, “Church’s thesis” proclaimed the uniqueness of all models of (classical) calculation: recursive function theory, Turing machine, λ -calculus, etc.... This result was refined in the 1960s, by showing that most “natural” models are *polynomially equivalent* to the Turing machine. The present paper can be viewed as supporting a similar status for QCM as *the* inherently quantum mechanical model of calculation. The modern reconsideration of computation is founded on the distinction between *polynomial time* and *slower* algorithms. Of course, all functions computed in the QCM can be computed classically, but probably not in comparable time. Assigning to an integer its factors, while polynomial time in QCM [Sh] is nearly exponential time, $\exp(O(n^{1/3} \text{poly}(\log n)))$ (an empiric bound, the proved one is even worse) according to the most refined classical algorithms. The origin of this paper is in thought [Fr] that since ordinary quantum mechanics appears to confer a substantial speed up over classical calculations, that some principle borrowed from the early, string, universe might go still further. Each TQFT is an instance of this question since their discrete topological nature lends itself to translation into computer science. We answer here in the negative by showing that for a unitary TQFT, the transformations $V(h)$ have a hidden poly-local structure. Mathematically, $V(h)$ can be realized as the restriction to an invariant subspace of a transformation $\prod g_i$ on the state space of a quantum computer where each g_i is a *gate* and the length of the composition is linear in the length of h as a word in the standard generators, “Dehn twists” of the *mapping class group* = diffeomorphisms (Σ)/identity component. Thus, we add evidence to the unicity of the QCM. Several variants and antecedents of QCM, including quantum Turing machines, have previously been shown equivalent (with and without environmental errors)[Y].

From a physical standpoint, the QCM derives from Schrödinger’s equation as described by Feynman [Fey] and Lloyd [LI]. Let us introduce the model. Given a decision problem, the first or *classical phase* of the QCM is a classical program, which designs a *quantum circuit* to “solve” *instances* of the decision problem of length n . A quantum circuit is a composition \mathbf{U}_n of operators or *gates* $g_i \in \mathbf{U}(2)$ or $\mathbf{U}(4)$ taken from some fixed list of rapidly computable matrices¹, e.g. having algebraic entries. The following short list suffices to efficiently approximate any other choice of gates [Ki]:

$$\left\{ \begin{array}{l} \left| \begin{array}{cc} 0 & 1 \\ 1 & 0 \end{array} \right\rangle, \quad \left| \begin{array}{cc} 1 & 0 \\ 0 & i \end{array} \right\rangle, \quad \text{and} \quad \left| \begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & 0 & \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{2}} \end{array} \right\rangle \end{array} \right\}.$$

The *gates* are applied on some tensor power space $(\mathbb{C}^2)^{\otimes k(n)}$ of “ k qubits” and models a local transformation on a system of k spin $\frac{1}{2}$ particles. The gate g acts as the identity on all but one or two tensor factors where it acts as a matrix as above. This is the middle or *quantum phase* of the algorithm. The final phase is to perform a local von Neumann measurement on a final state $\psi_{\text{final}} = \mathbf{U}_n(\psi_{\text{initial}})$ (or a commuting family of the same) to extract a probabilistic answer to the decision problem. (The initial states’ ψ_{initial} must also be locally constructed.) In this phase, we could declare that observing a certain

¹ The i^{th} digit of each entry should be computable in $\text{poly}(i)$ time.

eigenvalue with probability $\geq \frac{2}{3}$ means “yes”. We are interested only in the case where the classical phase of circuit design and the length of the designed circuit are both smaller than some polynomial in n . Decision problems which can be solved in this way are said to be in the computational class BQP: *bounded-error quantum polynomial*. The use of \mathbb{C}^2 , the “qubit”, is merely a convenience, any decomposition into factors of bounded dimension gives an equivalent theory. We say \mathbf{U} is a quantum circuit over \mathbb{C}^p if all tensor factors have dimension $= p$.

Following Lloyd [Ll], note that if a finite dimensional quantum system, say $(\mathbb{C}^2)^{\otimes k}$, evolves by a Hamiltonian H , it is physically reasonable to assert that H is *poly-local*, $H = \sum_{\ell=1}^L H_\ell$, where the sum has $\leq \text{poly}(k)$ terms and each $H_\ell = \tilde{H}_\ell \otimes \text{id}$, where \tilde{H}_ℓ acts nontrivially only on a bounded number (often just two) qubits and as the identity on the remaining tensor factors. Now setting Plank’s constant $\hbar = 1$, the time evolution is given by Schrödinger’s equation: $\mathbf{U}_t = e^{2\pi i t H}$ whereas gates can rapidly approximate [Ki] any local transformation of the form $e^{2\pi i t H_\ell}$. Only the nonabelian nature of the unitary group prevents us from approximating \mathbf{U}_t directly from $\prod_{\ell=1}^L e^{2\pi i t H_\ell}$. However, by the Trotter formula:

$$(e^{A/n+B/n})^n = e^{A+B} + \mathcal{O}\left(\frac{1}{n}\right),$$

where the error \mathcal{O} is measured in the operator norm. Thus, there is a good approximation to \mathbf{U}_t as a product of gates:

$$\mathbf{U}_t = \left(e^{2\pi i \frac{t}{n} H_1} \dots e^{2\pi i \frac{t}{n} H_L} \right)^n + L^2 \cdot \mathcal{O}\left(\frac{1}{n}\right).$$

Because of the rapid approximation result of [Ki], in what follows, we will not discuss quantum circuits restricted to any small generating set as in the example above, rather we will permit a 2×2 or 4×4 unitary matrix with algebraic number entries to appear as a gate.

In contrast to the systems considered by Lloyd, the Hamiltonian in a topological theory vanishes identically, $H = 0$, a different argument - the substance of this paper - is needed to construct a simulation. The reader may wonder how a theory with vanishing H can exhibit nontrivial unitary transformations. The answer lies in the Feynman path-integral approach to QFT. When the theory is constructed from a Lagrangian (functional on the classical fields of the theory), which only involves first derivatives in time, the Legendre transform is identically zero [At], but may nevertheless have nontrivial global features as in the Aharonov-Bohm effect.

Before defining the mathematical notions, we would make two comments. First, the converse to the theorem is also true. It has been shown recently [FLW] that a particular UTMF allows efficient simulation of the universal quantum computer. Second, we would like to suggest that the theorem may be viewed as a positive result for computation. Modular functors, because of their rich mathematical structure, may serve as higher order language for constructing a new quantum algorithm. In [Fr], it is observed that the transformations of UTMF’s can readily produce state vectors whose coordinates are computationally difficult evaluations of the Jones and Tutte polynomials. The same is now known for the state vector of a quantum computer, but the question of whether any useful part of this information can be made to survive the measurement phase of quantum computation is open.

2. Simulating Modular Functors

We adopt the axiomatization of [Wa] or [T] to which we refer for details. Also see, Atiyah [At], Segal [Se], and Witten [Wi].

A *surface* is a compact oriented 2-manifold with parameterized boundaries. Each boundary component has a label from a finite set $\mathcal{L} = \{1, a, b, c, \dots\}$ with involution $\widehat{\cdot}, 1 = \widehat{1}$. In examples, labels might be representations of a quantum group up to a given level or positive energy representations of a loop group, or some other algebraic construct. Technically, to avoid projective ambiguities each surface Σ is provided with a Lagrangian subspace $L \subset H_1(\Sigma; \mathbb{Q})$ and each diffeomorphism $f : \Sigma \rightarrow \Sigma'$ is provided with an integer “framing/signature” so the dynamics of the theory is actually given by a central extension of the mapping class group. Since these extended structures are irrelevant to our development, we suppress them from the notation. We use the letter ℓ below to indicate a label set for all boundary components, or in some cases, those boundary components without a specified letter as label.

Definition 1. A unitary topological modular functor (UTMF) is a functor V from the category of (labeled surfaces with fixed boundary parameterizations, label preserving diffeomorphisms which commute with boundary parameterizations) to (finite dimensional complex Hilbert spaces, unitary transformations) which satisfies:

1. *Disjoint union axiom:* $V(Y_1 \amalg Y_2, \ell_1 \amalg \ell_2) = V(Y_1, \ell_1) \otimes V(Y_2, \ell_2)$.
2. *Gluing axiom:* let Y_g arise from Y by gluing together a pair of boundary circles with dual labels, x glues to \widehat{x} , then

$$V(Y_g, \ell) = \bigoplus_{x \in \mathcal{L}} V(Y, (\ell, x, \widehat{x})).$$

3. *Duality axiom:* reversing the orientation of Y and applying $\widehat{\cdot}$ to labels corresponds to replacing V by V^* . Evaluation must obey certain naturality conditions with respect to gluing and the action of the various mapping class groups.
4. *Empty surface axiom:* $V(\phi) \cong \mathbb{C}$.
5. *Disk axiom:* $V_a = V(D, a) \cong \begin{cases} \mathbb{C}, & \text{if } a = 1 \\ 0, & \text{if } a \neq 1 \end{cases}$
6. *Annulus axiom:* $V_{a,b} = V(A, (a, b)) \cong \begin{cases} \mathbb{C}, & \text{if } a = \widehat{b} \\ 0, & \text{if } a \neq \widehat{b} \end{cases}$
7. *Algebraic axiom:* The basic data, the mapping class group actions and the maps F and S explained in the proof (from which V may be reconstructed if the Moore and Seiberg conditions are satisfied, see [MS] or [Wa] 6.4, 1–14) is algebraic over \mathbb{Q} for some bases in $V_a, V_{a,\widehat{a}}$, and V_{abc} , where V_{abc} denotes $V(P, (a, b, c))$ for a (compact) 3-punctured sphere P . 3-punctured spheres are also called pants.

Comments.

- (1) From the gluing axiom, V may be extended via dissection from simple pieces $D, A,$ and P to general surfaces Σ . But $V(\Sigma)$ must be canonically defined: this looks quite difficult to arrange and it is remarkable that any nontrivial examples of UTMFs exist.

- (2) The algebraic axiom is usually omitted, but holds for all known examples. We include it to avoid trivialities such as a UTMF where action by, say, a boundary twist is multiplication by a real number whose binary expansion encodes a difficult or even uncomputable function: e.g. the i^{th} bit is 0 iff the i^{th} Turing machine halts. If there are nontrivial parameter families of UTMF's, such nonsensical examples must arise – although they could not be algebraically specified. In the context of bounded accuracy for the operation of diffeomorphisms $V(h)$, Axiom 7 may be dropped (and simulation by bounded accuracy quantum circuits still obtained), but we prefer to work in the exact context since in a purely topological theory exactness is not implausible.
- (3) Axiom 2 will be particularly important in the context of a *pants decomposition* of a surface Σ . This is a division of Σ into a collection of compact surfaces P having the topology of 3-punctured spheres and meeting only in their boundary components which we call “cuffs”.

Definition 2. A quantum circuit $\mathbf{U} : (\mathbb{C}^p)^{\otimes k} \rightarrow (\mathbb{C}^p)^{\otimes k} =: W$ is said to simulate on W (exactly) a unitary transformation $\tau : S \rightarrow S$ if there is a \mathbb{C} -linear imbedding $i : S \subset (\mathbb{C}^p)^{\otimes k}$ invariant under \mathbf{U} so that $\mathbf{U} \circ i = i \circ \tau$. The imbedding is said to intertwine τ and \mathbf{U} . We also require that i be computable on a basis in poly(k) time.

Since we prove efficient simulation of the topological dynamics for UTMFs V , it is redundant to dwell on “measurement” within V , but to complete the computational model, we can posit von Neumann type measurement with respect to any efficiently computable frame \mathcal{F} in V_{abc} . The space \mathbb{C}^p above, later denoted $X = \mathbb{C}^p$, is defined by $X := \bigoplus_{(a,b,c) \in \mathcal{L}^3} V_{abc}$ and the computational space $W := X^{\otimes k}$. We have set $S := V(\Sigma)$ and assumed Σ is divided into k “pants”, i.e. Euler class $(\Sigma) = -k$. Any frame \mathcal{F} extends to a frame for $V(\Sigma)$ via the gluing axiom once a pants decomposition of Σ is specified. Thus, measurement in V becomes a restriction of measurement in W . It may be physically more natural to restrict the allowable measurements on $V(\Sigma)$ to cutting along a simple closed curve γ and measuring the label which appears. Mathematically, this amounts to transforming to a pants decomposition with γ as one of its decomposition or “cuff” curves and then positing a Hermitian operator with eigenspaces equal to the summands of $V(\Sigma)$ corresponding under the gluing axiom to labels x on γ^- and \hat{x} on γ^+ , $x \in \mathcal{L}$.

A labeled surface (Σ, ℓ) determines a mapping class group $\mathcal{M} = \mathcal{M}(\Sigma, \ell) =$ “isotopy classes of orientation preserving diffeomorphisms of Σ preserving labels and commuting with boundary parameterization”. For example, in the case of an n -punctured sphere with all labels equal (distinct), $\mathcal{M} = \text{SFB}(n)$, the spherical framed braid group ($\mathcal{M} = \text{PSFB}(n)$, the pure spherical framed braid group). To prove the theorem below, we will need to describe a generating set \mathcal{S} for the various \mathcal{M} 's and within \mathcal{S} chains of *elementary moves* which will allow us to prepare to apply any $s_2 \in \mathcal{S}$ subsequent to having applied $s_1 \in \mathcal{S}$.

Each \mathcal{M} is generated by *Dehn-twists* and *braid-moves* (See [B]). A Dehn-twist D_γ is specified by drawing a simple closed curve (s.c.c.) γ on Σ , cutting along γ , twisting 2π to the right along γ and then regluing. A braid-move B_δ will occur only when a s.c.c. δ cobounds a pair of pants with two boundary components of Σ : If the labels of the boundary components are equal then B_δ braids them by a right π -twist. In the case that all labels are equal, there is a rather short list of D and B generators indicated in Fig. 1 below. Also sketched in Fig. 1 is a pants decomposition of diameter $= (\mathcal{O} \log b_1(\Sigma))$,

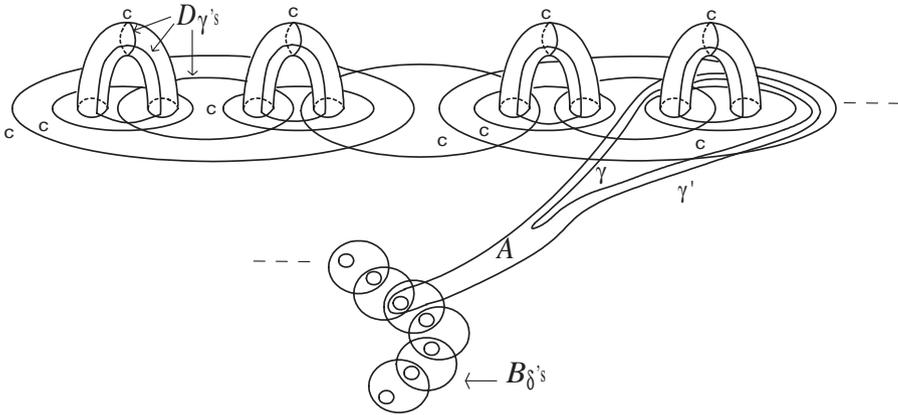


Fig. 1.

meaning the graph dual to the pants decomposition has diameter order log the first Betti number of Σ .

The s.c.c. $\gamma(\delta)$ label Dehn (braid) generators D_γ (and B_δ). Figure 1 contains a punctured annulus A ; note that the composition of oppositely oriented Dehn twists along the two “long” components of ∂A , γ and γ' yield a diffeomorphism which moves the punctures about the loop γ . The figure implicitly contains such an A for each (γ, p) , where p is a preferred puncture. The γ curves come in three types:

- (1) The loops at the top of the handles which are curves (“cuffs”) of the pants decomposition,
- (2) loops dual to type 1, and
- (3) loops running under adjacent pairs of handles (which cut through up to $\mathcal{O}(\log(b_1 \Sigma))$ many cuffs). (See Fig. 1, where cuffs are marked by a “c”.)

Each punctured annulus A is determined as a neighborhood (of a s.c.c. γ union an arc η from γ to p). To achieve general motions of p around Σ , we require these arcs to be “standard” so that for each p , $\pi_1(\widehat{\Sigma}, p)$ is generated by $\{\eta \cdot \gamma \cdot \eta^{-1}\}$, where $\widehat{\Sigma} = \Sigma$ with punctures filled by disks, and the disk corresponding to p serving as a base point. This list of generators is only linear in the first Betti number of Σ .

In the presence of distinct labels, many of the B_δ are illegal (they permute unequal labels). In this case, quadratically many generators are required. Figure 2 displays the replacements for the B ’s, and additional A ’s and D ’s.

Figure 2 shows a collection of B ’s sufficient to effect arbitrary braiding within each commonly-labeled subset of punctures, a quadratically large collection of new Dehn curves $\{\epsilon\}$ allowing a full twist between any pair of distinctly labeled punctures. (If the punctures are arranged along a convex arc of the Euclidean cell in Σ , then each ϵ will be the boundary of a narrow neighborhood of the straight line segment joining pairs of dissimilarly labeled punctures.) Finally a collection of punctured annuli, which enable one puncture p_i from each label – constant subset to be carried around each free homotopy class from $\{\gamma\}$ (respecting the previous generation condition for $\pi_1(\widehat{\Sigma}, p_i)$).

Thus for distinct labels the generating sets are built from curves of type $\gamma, \gamma', \epsilon$ and δ by Dehn twists around $\gamma, \gamma',$ and ϵ , braid moves around δ . Denote by ω , any such curve: $\omega \in \Omega = \{\{\gamma\} \cup \{\gamma'\} \cup \{\epsilon\} \cup \{\delta\}\}$.

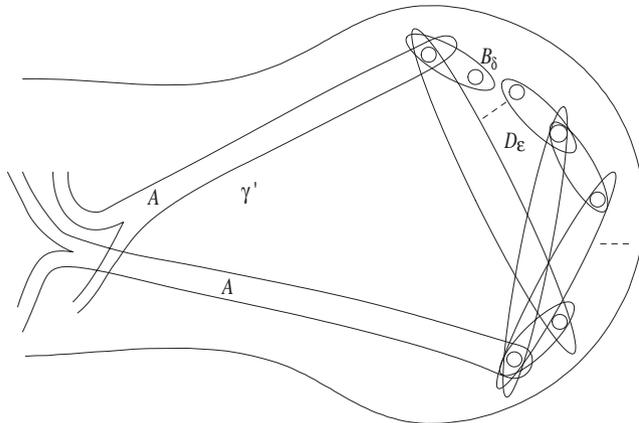


Fig. 2.

Since various ω' 's intersect, it is not possible to realize all ω simultaneously as cuffs in a pants decomposition. However, we can start with the “base point” pants decomposition \mathcal{D} indicated in Fig. 1 (note γ of type(1) are cuffs in \mathcal{D} , but γ of types (2) or (3) are not) and for any ω find a short path of elementary moves: F and S (defined below) to a pants decomposition \mathcal{D}_ω containing ω as a cuff.

Lemma 2.1. *Assume $\Sigma \neq S^2$, disk, or annulus, and \mathcal{D} the standard pants decomposition sketched in Fig. 1. Any ω as above, can be deformed through $\mathcal{O}(\log b_1(\Sigma))$ F and S moves to a pants decomposition \mathcal{D}_ω in which ω is a cuff.*

We postpone the proof of the lemma and the definition of its terms until we are partly into the proof of the theorem and have some experience passing between pants decompositions.

Theorem 2.2. *Suppose V is a UTMF and $h : \Sigma \rightarrow \Sigma$ is a diffeomorphism of length n in the standard generators for the mapping class group of Σ described above (see Figs. 1 and 2). Then there are constants depending only on V , $c = c(V)$ and $p = p(V)$ such that $V(h) : V(\Sigma) \rightarrow V(\Sigma)$ is simulated (exactly) by a quantum circuit operating on “qubits” \mathbb{C}^p of length $\leq c \cdot n \cdot \log b_1(\Sigma)$.*

The collection {cuffs} refers to the circles along which the pants decomposition decomposes; the “seams” are additional arcs, three per pant which cut the pant into two hexagons. Technically, we will need each pant in \mathcal{D} to be parameterized by a fixed 3-punctured sphere so these seams are part of the data in \mathcal{D} ; for simplicity, we choose seams to minimize the number of intersections with $\{\omega\}$.

The theorem may be extended to cover a more general form of input. The original algorithm [L] which writes a \mathcal{D}_α , α a s.c.c., as a word in standard generators \mathcal{D}_γ is super-exponential. We define the combinatorial length of α , $\ell(\alpha)$, to be the minimum number of intersections as we vary α by isotopy of α with {cuffs} \cup {seams}. The best upper-bound (known to the authors) to the length L of \mathcal{D}_α as a word in the mapping class group spanned by a fixed generating set is of the form $L(\mathcal{D}_\alpha) <$ super-exponential function $f(\ell)$. For this reason, we consider as input $V(h)$, where h is a composition of k Dehn twists on $\alpha_1, \dots, \alpha_k$ and j braid moves along β_1, \dots, β_j in any order.

Then $V(h)$ is costed as the sum of the combinatorial length of the simple closed curves needed to write h as Dehn twists and braid moves within the mapping class group,

$$\ell(h) := \sum_{i=1}^j \ell(\beta_i) + \sum_{i=1}^k \ell(\alpha_i).$$

We obtain the following extension of the theorem.

Extension.² The map $h_* : V(\Sigma) \rightarrow V(\Sigma)$ is exactly simulated by a quantum circuit QC with length $(QC) \leq 11\ell(h)$ composed of algebraic 1 and 2-qubit \mathbb{C}^p gates.

Pre-Proof. Some physical comments will motivate the proof. $V(\Sigma)$ are quantized gauge fields on Σ (with a boundary condition given by labels ℓ) and can be regarded as a finite dimensional space of internal symmetries. This is most clear when genus $(\Sigma) = 0$, Σ is a punctured sphere, the labeled punctures are “anyons” [Wil] and the relevant mapping class group is the *braid group* which moves the punctures around the surface of the sphere. An internal state $\psi \in V(\Sigma)$ is transformed to $\mathbf{U}(b)\psi \in V(\Sigma)$ under the functorial representation of the braid group. For $\mathbf{U}(b)$ to be defined the braiding must be “complete” in the sense that the punctures (anyons) must return setwise to their initial position. Infinitesimally, the braiding defines a Hamiltonian \overline{H} on $V(\Sigma) \otimes E$, where E is an infinite dimensional Hilbert space which encodes the position of the anyons. The projection of \overline{H} into $V(\Sigma)$ vanishes which is consistent with the general covariance of topological theories. Nevertheless, when the braid is complete, the evolution $\overline{\mathbf{U}}$ of \overline{H} will leave $V(\Sigma)$ invariant and it is $\overline{\mathbf{U}}|_{V(\Sigma)} = \mathbf{U}$ which we will simulate. Anyons inherently reflect nonlocal entanglement so it is not to be expected that $V(\Sigma)$ has any (natural) tensor decomposition and none are observed in interesting examples. Thus, simulation of \mathbf{U} as an invariant subspace of a tensor product $(\mathbb{C}^p)^{\otimes k}$ is the best result we can expect. The mathematical proof will loosely follow the physical intuition of evolution in a super-space by defining, in the braid case (identical labels and genus = 0), two distinct imbeddings “odd” and “even”, $V(\Sigma) \xrightarrow[\text{even}]{\text{odd}} (\mathbb{C}^p)^{\otimes k} = W$ and constructing the local evolution by gates acting on the target space. The imbeddings are named for the fact that in the usual presentation of the braid group, the odd (even) numbered generators can be implemented by restricting an action on W to image odd $V(\Sigma)$ (even $V(\Sigma)$).

Proof. The case genus $(\Sigma) = 0$ with all boundary components carrying identical labels (this contains the classical, uncolored Jones polynomial case [J, Wi]) is treated first. For any number q of punctures ($q = 10$ in the illustration) there are two systematic ways of dividing Σ into pants (3-punctured spheres) along curves $\overrightarrow{\alpha} = \{\alpha_1, \dots, \alpha_{q-3}\}$ or along $\overrightarrow{\beta} = \{\beta_1, \dots, \beta_{q-3}\}$ so that a sequence of q F moves ($6j$ -moves in physics notation) transforms $\overrightarrow{\alpha}$ to $\overrightarrow{\beta}$.

Let $X = \bigoplus_{(a,b,c) \in \mathcal{L}^3} V_{abc}$ be the orthogonal sum of all sectors of the pants Hilbert space. Distributing \otimes over \bigoplus , the tensor power $X^{\otimes(q-2)} := W$ is the sum over all labelings of the Hilbert space for $\bigsqcup(q-2)$ pants. Choosing parameterizations, W is identified with both the label sum space $(\Sigma_{\text{cut}\overrightarrow{\alpha}} \rightarrow)$ and sum $(\Sigma_{\text{cut}\overrightarrow{\beta}} \rightarrow)$. Now Σ is assembled

from the disjoint union by gluing along $\overrightarrow{\alpha}$ or $\overrightarrow{\beta}$ so the gluing axiom defines imbeddings $i(\overrightarrow{\alpha})$ and $i(\overrightarrow{\beta})$ of $V(\Sigma, \ell)$ as a direct summand of $X^{\otimes(q-2)} = W$.

² Lee Mosher has informed us that the existence of the linear bound $f(\ell)$ (but without control of the constants) follows at least for closed and single punctured surfaces from his two papers [M1] and [M2].

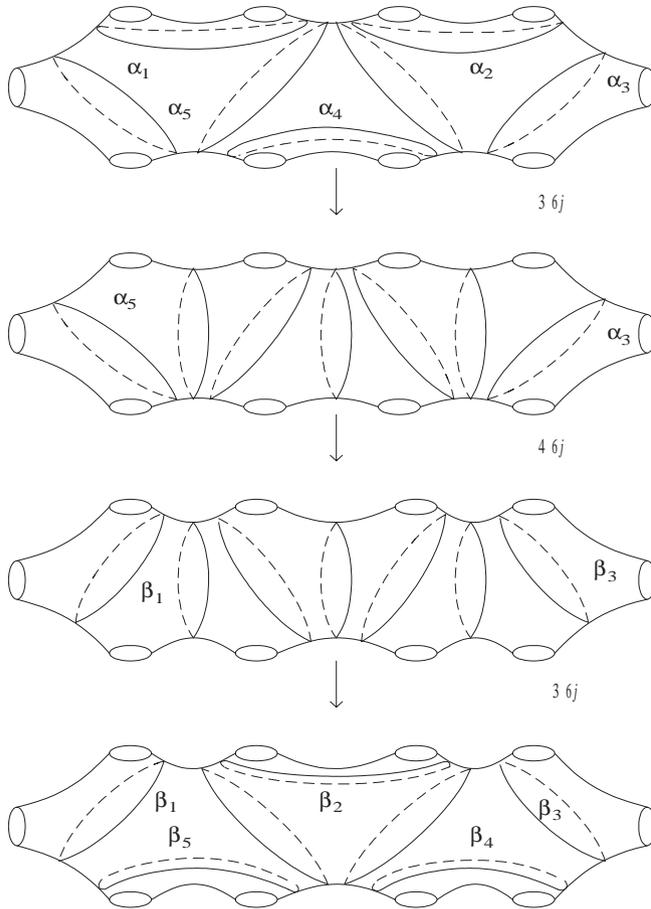


Fig. 3.

Consider the action of braid move about α . This acts algebraically as $\theta(\alpha_i)$ on a single X factor of W and as the identity on other factors. This action leaves $i(\vec{\alpha})$ ($V(\Sigma, \ell)$) invariant and can be thought of as a “qubit” gate:

$$\theta(\alpha_i) = V(\text{braid}_{\alpha_i}) : X \rightarrow X,$$

where dimension $\dim(X) = p$. Similarly the action of $V(\text{braid}_{\beta_i})$ leaves $i(\vec{\beta})$ invariant.

It is well known [B] that the union of loops $\vec{\alpha} \cup \vec{\beta}$ determines a complete set of generators of the braid group. The general element ω , which we must simulate by an action on W is a word in braid moves on α 's and β 's. Part of the *basic data* – implied by the gluing axiom for a UTMF is a fixed identification between elementary gluings:

$$F_{abcd} : \bigoplus_{x \in \mathcal{L}} V_{xab} \otimes V_{xcd} \longrightarrow \bigoplus_{y \in \mathcal{L}} V_{ybc} \otimes V_{yda}$$

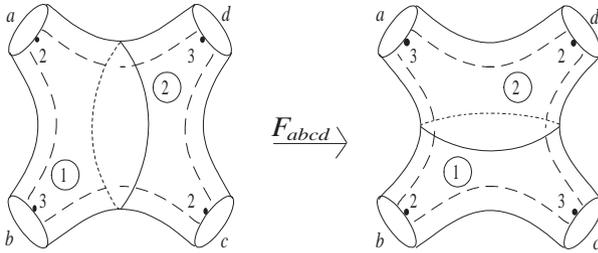


Fig. 4.

corresponding to the following two decompositions of the 4-punctured sphere into two pairs of pants (the dotted lines are pant “seams”, the uncircled number indicate boundary components, the letters label boundary components, and the circled numbers order the pairs of pants.):

For each F , we choose an extension to a unitary map $F' : X \otimes X \rightarrow X \otimes X$. Then extend F' to \bar{F} by tensoring with identity on the $q - 4$ factors unaffected by F . The composition of q F 's, extended to q \bar{F} 's, corresponding to the q moves illustrated in the case $q = 10$ by Fig. 3. (For $q > 10$ imagine the drawings in Fig. 3 extended periodically.) These define a unitary transformation $T : W \rightarrow W$ with $T \circ i(\vec{\alpha}) = i(\vec{\beta})$. The word ω in the braid group can be simulated by τ on W , where τ is written as a composition of the unitary maps $T, T^{-1}, \theta(\alpha_i)$, and $\theta(\beta_j)$. For example,

$$\beta_5 \alpha_1 \beta_2^{-1} \alpha_1 \alpha_3$$

would be simulated as

$$\tau = T^{-1} \circ \theta(\beta_5) \circ T \circ \theta(\alpha_1) \circ T^{-1} \circ \theta(\beta_2^{-1}) \circ T \circ \theta(\alpha_1) \circ \theta(\alpha_3).$$

As described τ has length $\leq 2q$ length ω . The dependence on q can be removed by dividing Σ into $\frac{q}{2}$ overlapping pieces Σ_i , each Σ_i a union of 6 consecutive pants. Every loop of $\vec{\alpha} \cup \vec{\beta}$ is contained well within some piece Σ_i so instead of moving between two fixed subspaces $i_\alpha(V)$ and $i_\beta(V) \subset W$, when we encounter a β_j , do constantly many \bar{F} operations to find a new pants decomposition modified locally to contain β_j . Then $\theta(\beta_j)$ may be applied and the \bar{F} operations reversed to return to the α pants decomposition. The resulting simulation can be made to satisfy length $\tau \leq 7$ length ω . This completes the braid case with all bounding labels equal - an important case corresponding to the classical Jones polynomial [J]. \square

Proof of Lemma. We have described the F -move on the 4-punctured sphere both geometrically and under the functor. The S -move is between two pants decompositions on the punctured torus T^- . (Filling in the puncture, a variant of S may act between two distinct annular decomposition of T^2 . We suppress this case since, without topological parameter, there can be no computational complexity discussion over a single surface.)

By [Li] or [HT] that one may move between any two pants decompositions via a finite sequence of moves of three types: F, S , and diffeomorphism M supported on the interior of a single pair of pants (see the Appendix [HT]). To pass from \mathcal{D} , our “base point” decomposition, to \mathcal{D}_ω , F and S moves alone suffice and the logarithmic count is a consequence of the log depth nest of cuff loops of \mathcal{D} on the planar surface obtained by

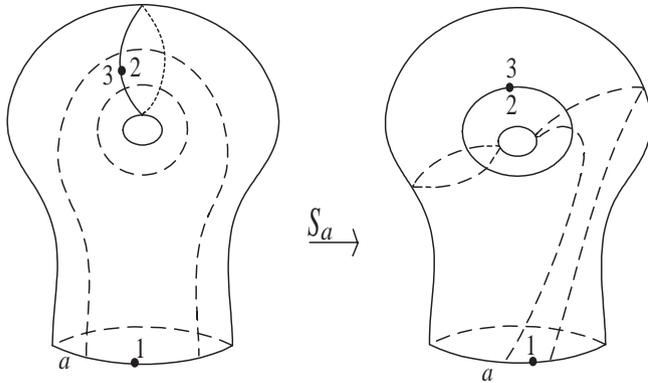


Fig. 5. $V(S): \bigoplus_{x \in \mathcal{L}} V_{ax\hat{x}} \longrightarrow \bigoplus_{y \in \mathcal{L}} V_{ay\hat{y}}$

cutting Σ along type (1) γ curves. Below we draw examples of short paths of F and S moves taking \mathcal{D} to a particular \mathcal{D}_ω .

The logarithmic count is based on the proposition.

Proposition 2.3. *Let K be a trivalent tree of diameter $= d$ and f be a move, which locally replaces $\{ \succ \leftarrow \}$ and with $\{ \leftarrow \succ \}$, then any two leaves of K can be made adjacent by $\leq d$ moves of type f . (Here we consider abstract trees rather than ones imbedded in the plane.)*

Passing from K to a punctured sphere obtained by imbedding $(K, \text{univalent vertices})$ into $(\frac{1}{2}R^3, R^2)$, thickening and deleting the boundary R^2 , the f move induces the previously defined F move. \square

Some example of paths of F, S moves (Fig. 6).

Continuation of the proof of the theorem. For the general case, we compute on numerous imbeddings of $V(\Sigma)$ into W (rather than on two: $i_\alpha(V(\Sigma))$ and $i_\beta(V(\Sigma))$ as in the braid case). Each imbedding is determined by a pants decomposition and the imbedding changes (in principle) via the lemma every time we come to a new literal of the word ω . Recall that $\omega \in \mathcal{M}$, the mapping class group, is now written as a word in the letters (and their inverses) of type $D_\gamma, D'_\gamma, D_\epsilon,$ and B_δ . Pick as a home base a fixed pants decomposition \mathcal{D}_0 corresponding to $i_0(V(\Sigma)) \subset W$. If the first literal is a twist or braid along the s.c.c. ω , then apply the lemma to pass through a sequence of F and S moves from \mathcal{D}_0 to \mathcal{D}_1 containing ω as a ‘‘cuff’’ curve. As in the braid case, choose extensions \bar{F} and \bar{S} to unitary automorphisms of W and applying V to the composition gives a transformation T_1 of W such that $i_1 = T_1 \circ i_0, i_1$ being the inclusion $V(\Sigma) \rightarrow W$ associated with \mathcal{D}_1 . Now execute the first literal ω_1 of ω as a transformation $\theta(\omega_1)$, which leaves $i_1(V(\Sigma))$ invariant and satisfies: $\theta(\omega_i) \circ i_1 = i_1 \circ V(\omega_1)$. Finally apply T_1^{-1} to return to the base inclusion $i_0(V(\Sigma))$. The previous three steps can now be repeated for the second literal of ω : follow $T_1^{-1} \circ \theta(\omega_1) \circ T_1$ by $T_2^{-1} \circ \theta(\omega_2) \circ T_2$. Continuing in this way, we construct a composition τ which simulates ω on W :

$$\tau = T_n^{-1} \circ \theta(\omega_n) \circ T_n^{-1} \cdot \dots \cdot T_1^{-1} \circ \theta(\omega_1) \circ T_1.$$

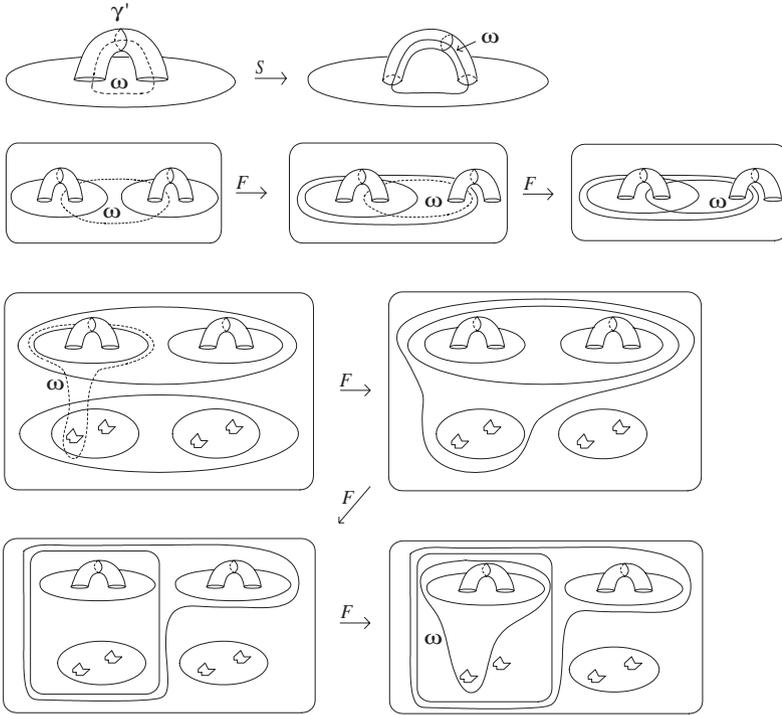


Fig. 6.

From Lemma 2.1 the length of this simulation by one (corresponding to S and $\theta(\omega_i)$) and two (corresponding to F moves) qubit gates is proportional to $n = \text{length } \omega$ and $\log b_1(\Sigma)$, where $p = \dim(X)$.

Proof of Extension. What is at issue is the number of preparatory moves to change the base point decomposition \mathcal{D} to \mathcal{D}_γ containing $\gamma = \alpha_i$ or β_j as a cuff curve $1 \leq i \leq k$ or j . We have defined the F and S moves rigidly, i.e. with specified action on the seams. This was necessary to induce a well defined action on the functor V . Because of this rigid choice, we must add one more move – an M move – to have a complete set of moves capable of moving between any two pants decompositions of a surface (compare [HT]). The M move is simply a Dehn twist supported in a pair of pants of the current pants decomposition; it moves the seams (compare Chapter 5 [Wa]). Note that if M is a $+1$ Dehn twist in a s.c.c. ω then, under the functor, $V(M)$ is a restriction of $\theta(\omega)$ in the notation above.

As in [HT], the cuff curves of \mathcal{D} may be regarded as level curves of a Morse function $f : \Sigma \rightarrow \mathbb{R}^+$, constant on boundary components which we assume to have minimum complexity (= total number of critical points) satisfying this constraint. Isotope α (we drop the index) on Σ to have the smallest number of local maximums with respect to f and is disjoint from critical points of f on Σ .

Now generically deform f in a thin annular neighborhood of γ so that γ becomes a level curve. Consider the graphic G of the deformation $f_t, 0 \leq t \leq 1$. For regular t the Morse function f_t determines a pants decomposition: let the 1- complex K consist of Σ / \sim where $x \sim y$ if x and y belong to the same component of a level set of f_t , and let

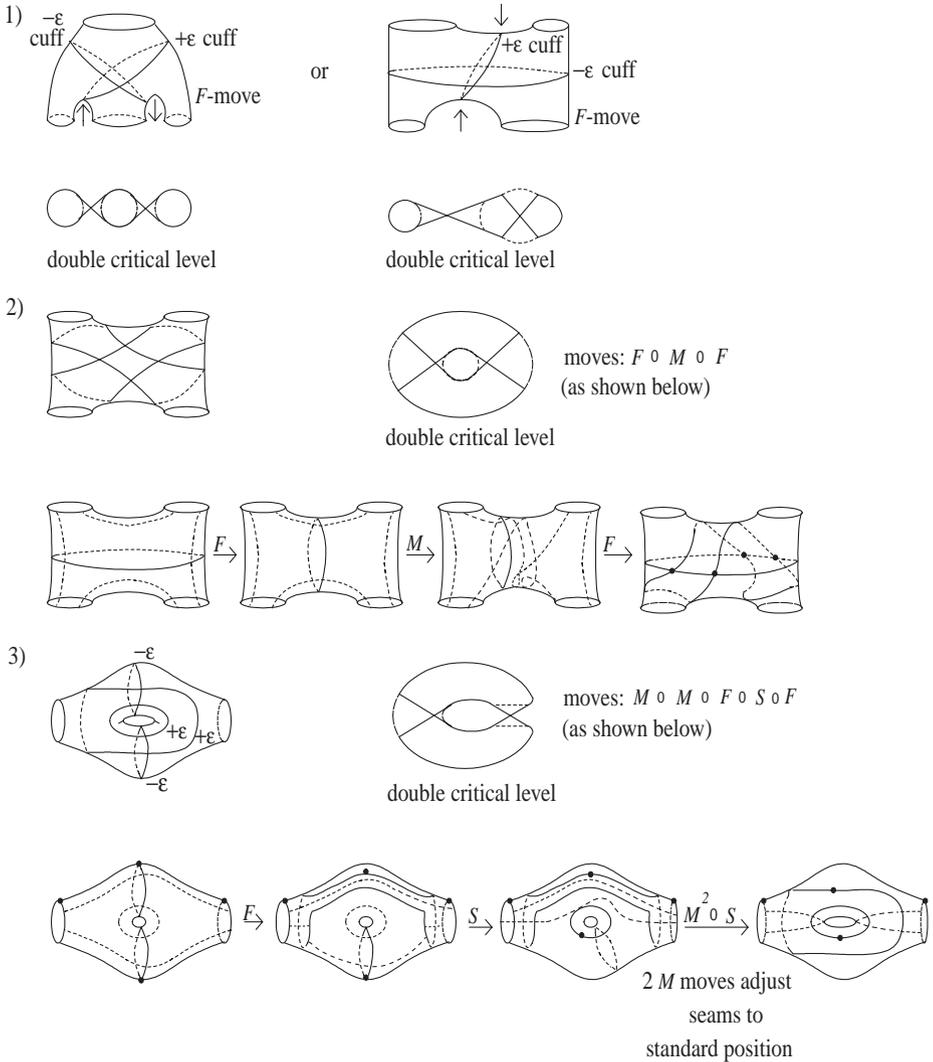


Fig. 7.

$L \subset K$ be the smallest complex to which K collapses relative to endpoints associated to boundary components. For example in Fig. 8, the top tree does not collapse at all while in the lower two trees the edge whose end is labeled, “local max” is collapsed away. The preimage of one point from each intrinsic 1-cell of L not containing a boundary point constitutes a {cuffs} determining a pants decomposition \mathcal{D}_t . For singular t_0 , let $\mathcal{D}_{t_0-\epsilon}$ and $\mathcal{D}_{t_0+\epsilon}$ may differ or may agree up to isotopy. The only change in \mathcal{D} occurs when t is a crossing point for index= 1 handles where the two critical points are on the same connected component of a level set $f_t^{-1}(r)$. There are essentially only three possible “Cerf-transitions” and they are expressible as a product of 1, 2, or 3 F and S moves together with braid moves whose number we will later bound from above. The Cerf

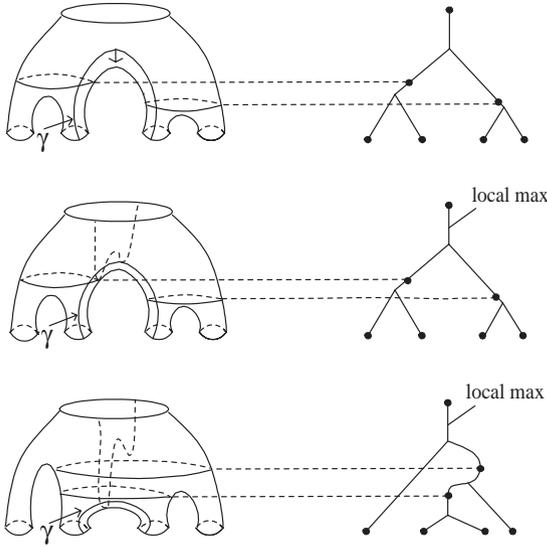


Fig. 8. Pulling γ down yields $F \circ F$

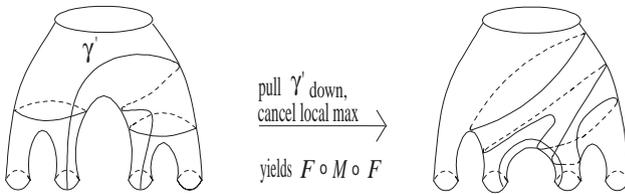


Fig. 9.

transitions on \mathcal{D} are shown in Fig. 7, together with their representation as compositions of elementary moves.

Critical points of $f|_\gamma$ become critical points of f_t of the same index once the deformation as passed an initial $\epsilon_0 > 0$, and before any saddle-crossings have occurred. Let P be a pant from the composition induced by f and $\delta \subset \gamma \cap P$ an arc. Applying the connectivity criterion of the previous paragraph, we can see that flattening a local maxima can effect at most the two cuff circles which δ meets, and these by elementary Cerf transition shown in Fig. 8.

If γ crosses the seam arcs then the transitions are of the Cerf type, precomposed with M -moves to remove these crossings as shown in Fig. 9. Dynamically seam crossings by γ produce *saddle connections* in the Cerf diagram.

The total number of these twists is bounded by length (γ) . The number of flattening moves as above is less than or equal $|\gamma \cap \text{cuffs}| \leq \text{length}(\gamma)$. The factor of 11 in the statement allows up to 5 F , S , and M moves for expressing each Cerf singularity which arises in passing from \mathcal{D}_o to \mathcal{D}_γ and the same factor of 5 to pass back from \mathcal{D}_γ to \mathcal{D}_o . again, while saving at least one step to implement the twist or build move along γ . This completes the proof of the extension. \square

We should emphasize that although, we have adopted an “exact” model for the operation of the UTMF, faithful simulation as derived above does not depend on a perfectly accurate quantum circuit. Several authors have proved a threshold theorem [Ki, AB], and [KLZ]: If the rate of large errors acting on computational qubits (or qupits) is small enough, the size of ubiquitous error small enough, and both are uncorrelated, then such a computational space may be made to simulate with probability $\geq \frac{2}{3}$ an exact quantum circuit of length $= L$. The simulating circuit must exceed the exact circuit in both number of qubits and number of operations by a multiplicative factor $\leq \text{poly}(L)$.

3. Simulating TQFT’s

We conclude with a discussion about the three dimensional extension, the TQFT of a UTMF. In all known examples of TMF’s there is an extension to a TQFT meaning that it is possible to assign a linear map $V(\Sigma) \xrightarrow{b_*} V(\Sigma')$ subject to several axioms [Wa] and [T] whenever Σ and Σ' cobounds a bordism b (with some additional structure). The case of bordisms with a product structure is essentially the TMF part of the theory. Unilarity is extended to mean that if the orientation of the bordism b is reversed to \bar{b} , we have $b_*^\dagger = (\bar{b})_*$. It is known that a TMF has at most one extension to a TQFT and conjectured that this extension always exists. Non-product bordisms correspond to some loss of information of the state. This can be understood by factoring the bordism into pieces consisting of a product union 2-handle: $\Sigma \times I \cup h$. The 2-handle h has the form $(D^2 \times I, \partial D^2 \times I)$ and is attached along the subspace $\partial D^2 \times I$. The effect of attaching the handle will be to “pinch” off an essential loop ω on Σ and so replace an annular neighborhood of ω by two disks turning Σ into a simpler surface Σ' . It is an elementary consequence of the axioms that if $b = \Sigma \times I \cup h$ then b_* is a projector as follows: Let \mathcal{D} be a pants decomposition containing ω as a dissection curve. There are two cases:

- (1) ω appears as the first and second boundary components of a single pant called P_0 or
- (2) ω appears as the first boundary component on two distinct pants called P_1 and P_2 .

$$\begin{aligned}
 V(\Sigma) &= \\
 &= \bigoplus_{c \in \mathcal{L}} \left(\left(\bigoplus_{a \in \mathcal{L}} V_{a\hat{a}c} \right) \otimes V(\Sigma \setminus P_0, \text{with label } c \text{ on } \partial_3 P_0) \right), \quad \text{case (1),} \\
 &\text{or} \\
 &= \bigoplus_{\text{labels}} \left(\bigoplus_{a \in \mathcal{L}} V_{abc} \otimes V_{\hat{a}de} \right) \otimes V(\Sigma \setminus (P_1 \cup P_2), \text{appropriate labels}), \quad \text{case (2).}
 \end{aligned}$$

In case (2), there may be a relation $b = \hat{c}$ and/or $d = \hat{e}$ depending on the topology of \mathcal{D} . The map b_* is obtained by extending linearly from the projections onto summands:

$$\begin{aligned}
 \bigoplus_{a, c \in \mathcal{L}} V_{a\hat{a}c} &\longrightarrow V_{111} \stackrel{\text{canonically}}{\cong} V_1, & \text{(case 1)} \\
 &\text{or} \\
 \bigoplus_{a, b, c, d, e \in \mathcal{L}} V_{abc} \otimes V_{\hat{a}de} &\longrightarrow V_{1b\hat{b}} \otimes V_{1d\hat{d}} \stackrel{\text{canonically}}{\cong} V_{b\hat{b}} \otimes V_{d\hat{d}}. & \text{(case 2)}
 \end{aligned}$$

If the orientation on b is reversed the unitarity condition implies that \bar{b} determines an injection onto a summand with a formula dual to the above. Thus, any bordism's morphism can be systematically calculated.

In quantum computation, as shown in [Ki], a projector corresponds to an intermediate binary measurement within the quantum phase of the computation, one outcome of which leads to cessation of the other continuation of the quantum circuits operation. Call such a probabilistically abortive computation a *partial computation* on a *partial quantum circuit*. Formally, if we write the identity as a sum of two projectors: $\text{id}_V = \Pi_0 + \Pi_1$, and let \mathbf{U}_0 and \mathbf{U}_1 be unitary operators on an ancillary space A with $\mathbf{U}_0(|0\rangle) = |0\rangle$ and $\mathbf{U}_1|0\rangle = |1\rangle$. The unitary operator $\Pi_0 \otimes \mathbf{U}_0 + \Pi_1 \otimes \mathbf{U}_1$ on $V \otimes A$ when applied to $|v\rangle \otimes |0\rangle$ is $|\Pi_0 v\rangle \otimes |0\rangle + |\Pi_1 v\rangle \otimes |1\rangle$ so continuing the computation only if the indicator $|0\rangle \in A$ is observed simulates the projection Π_0 .

It is clear that the proof of the theorem can be modified to simulate 2-handle attachments as well as Dehn twists and braid moves along s.c.c.'s ω to yield:

Scholium 3.1. *Suppose b is an oriented bordism from Σ_0 to Σ_1 , where Σ_i is endowed with a pants decomposition \mathcal{D}_i . Let complexity (b) be the total number of moves of four types: F , S , M , and attachment of a 2-handle to a dissection curve of a current pants decomposition that are necessary to reconstruct b from $(\Sigma_0, \mathcal{D}_0)$ to $(\Sigma_1, \mathcal{D}_1)$. Then there is a constant $c'(V)$ depending on the choice of UTQFT and $p(V)$ as before (for the TQFTs underlining TMF) so that $b_* : V(\Sigma_0) \rightarrow V(\Sigma_1)$ is simulated (up to a non-topological factor of the form v^{n_2} , where n_2 is the number of 2-handles attached) by a partial quantum circuit over \mathbb{C}^p of length $\leq c'$ complexity (b).*

In general, the difference between topological objects (such as b_* or closed 3-manifold invariants) and quantum mechanical ones (the evolution and probability) is related to critical points of a Morse function. A similar phenomenon for links in R^3 has been mentioned in [FKLW]. This subject will be addressed in detail in a forthcoming paper by S. Bravyi and A. Kitaev, "Quantum invariants of 3-manifolds and quantum computation".

Acknowledgements. We would like to thank Greg Kuperberg and Kevin Walker for many stimulating discussions on the material presented here.

References

- [At] Atiyah, M.: Topological quantum field theories. Publ. Math. IHES **68**, 175–186 (1989)
- [AB] Aharonov, D. and Ben-Or, M.: Fault-tolerant quantum computation with constant error. LANL e-print quan-ph/9611025
- [B] Birman, J.: Braids, links, and mapping class groups. Ann. Math. Studies, Vol. **82**
- [D] Deutsch, D.: Quantum computational networks. Proc. Roy. Soc. London, **A425**, 73–90 (1989)
- [Fey] Feynman, R.: Simulating physics with computers. Int. J. Theor. Phys. **21**, 467–488 (1982)
- [FKLW] Freedman, M.H., Kitaev, A., Larsen, M.J. and Wang, Z.: Topological Quantum Computation; LANL e-print quant-ph/0101025
- [FLW] Freedman, M.H., Kitaev, A., Larsen, M.J. and Wang, Z.: A modular functor which is universal for quantum computation. LANL e-print quant-ph/0001108
- [Fr] Freedman, M.H.: P/NP, and the quantum field computer. Proc. Natl. Acad. Sci., USA **95**, 98–101 (1998)
- [HT] Hatcher, A. and Thurston, W.: A presentation for the mapping class group of a closed orientable surface. Topology **19**, no. 3, 221–237 (1980)
- [J] Jones, V.F.R.: Hecke algebra representations of braid groups and link polynomial. Ann. Math. **126**, 335–388 (1987)
- [Ka] Kane, B.: A silicon-based nuclear spin quantum computer. Nature **393**, 133–137 (1998)
- [Ki] Kitaev, A.: Quantum computations: algorithms and error correction. Russian Math. Survey **52:61**, 1191–1249 (1997)

- [KLZ] Knill, E., Laflamme, R. and Zurek, W.: Threshold Accuracy for Quantum Computation. LANL e-print quant-ph/9610011, 20 pages, 10/15/96
- [Li] Lickorish, W.: A representation of orientable, combinatorial 3-manifolds. *Ann. Math.* **76**, 531–540 (1962)
- [Ll] Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996)
- [M1] Mosher, L.: Mapping class groups are automatic. *Ann. of Math.* **142**, 303–384 (1995)
- [M2] Mosher, L.: Hyperbolic extensions of groups. *J. of Pure and Applied Alg.* **110**, 305–314 (1996)
- [MS] Moore, G. and Seiberg, N.: Classical and quantum conformal field theory. *Commun. Math. Phys.* **123**, 177–254 (1989)
- [Se] Segal, G.: The definition of conformal field theory. Preprint (1999)
- [Sh] Shor, P.W.: *Algorithms for quantum computers: Discrete logarithms and factoring*. Proc. 35th Annual Symposium on Foundations of Computer Science. Los Alamitos: IEEE Computer Society Press, CA: pp. 124–134
- [T] Turaev, V.G.: *Quantum invariant of knots and 3-manifolds*. de Gruyter Studies in Math., Vol. **18**
- [Wa] Walker, K.: *On Witten's 3-manifold invariants*. Preprint, 1991
- [Wil] Wilczek, F.: *Fractional statistics and anyon superconductivity*. Teaneack, NJ: World Scientific Publishing Co., Inc., 1990
- [Wi] Witten, E.: Quantum field theory and the Jones polynomial. *Commun. Math. Phys.* **121**, 351–399 (1989)
- [Y] Yao, A.: *Quantum circuit complexity*. Proc. 34th Annual Symposium on Foundations of Computer Science. Los Alamitos, CA: IEEE Computer Society Press, pp. 352–361

Communicated by P. Sarnak