TOPOLOGICAL QUANTUM COMPUTATION

MICHAEL H. FREEDMAN, ALEXEI KITAEV, MICHAEL J. LARSEN,
AND ZHENGHAN WANG

ABSTRACT. The theory of quantum computation can be constructed from the abstract study of anyonic systems. In mathematical terms, these are unitary topological modular functors. They underlie the Jones polynomial and arise in Witten-Chern-Simons theory. The braiding and fusion of anyonic excitations in quantum Hall electron liquids and 2D-magnets are modeled by modular functors, opening a new possibility for the realization of quantum computers. The chief advantage of anyonic computation would be physical error correction: An error rate scaling like $e^{-\alpha\ell}$, where ℓ is a length scale, and α is some positive constant. In contrast, the "presumptive" qubit-model of quantum computation, which repairs errors combinatorially, requires a fantastically low initial error rate (about 10^{-4}) before computation can be stabilized.

Quantum computation is a catch-all for several models of computation based on a theoretical ability to manufacture, manipulate and measure quantum states. In this context, there are three areas where remarkable algorithms have been found: searching a data base [15], abelian groups (factoring and discrete logarithm) [19], [27], and simulating physical systems [5], [21]. To this list we may add a fourth class of algorithms which yield approximate, but rapid, evaluations of many quantum invariants of three dimensional manifolds, e.g., the absolute value of the Jones polynomial of a link L at certain roots of unity: $|V_L(e^{\frac{2\pi i}{5}})|$. This seeming curiosity is actually the tip of an iceberg which links quantum computation both to low dimensional topology and the theory of anyons; the motion of anyons in a two dimensional system defines a braid in $2+1$ dimension. This iceberg is a model of quantum computation based on topological, rather than local, degrees of freedom.

The class of functions, **BQP** (functions computable with bounded error, given quantum resources, in polynomial time), has been defined in three distinct but equivalent ways: via quantum Turing machines [2], quantum circuits [3], [6], and modular functors [7], [8]. The last is the subject of this article. We may now propose a "thesis" in the spirit of Alonzo Church: all "reasonable" computational models which add the resources of quantum mechanics (or quantum field theory) to classical computation yield (efficiently) inter-simulable classes: there is one quantum theory of computation. (But alas, we are not so sure of our thesis at Planck scale energies. Who is to say that all the observables there must even be computable functions in the sense of Turing?)

The case for quantum computation rests on three pillars: inevitability— Moore's law suggests we will soon be doing it whether we want to or not, desirability—the above mentioned algorithms, and finally feasibility—which in the past has been

argued from combinatorial fault tolerant protocols. (It is a quirk of human optimism that these petitions are usually pleaded independently: e.g., "feasibility" is not seen as a precondition for "inevitability".)

Focusing on feasibility, we present a model of quantum computation in which information is stored and manipulated in "topological degrees of freedom" rather than in localized degrees. The usual candidates for storing information are either localized in space (e.g., an electron or a nuclear spin) or localized in momentum (e.g., photon polarization.) Almost by definition (see "code subspace" below) a topological degree of freedom is protected from local errors. In the presence of perturbation (of the system Hamiltonian) this protection will not be perfect, but physical arguments suggest undesirable tunneling amplitudes between orthogonal ground states will scale like $e^{-\alpha l}$, where l is a length scale for the system, e.g., the minimum separation maintained between point-like anyonic excitations. We will return to this crucial point.

But let us take a step backward and discuss the standard quantum circuit model and the presumptive path toward its physical implementation. To specify a quantum circuit Γ, we begin with a tensor product $\mathbb{C}_1^2 \otimes \cdots \otimes \mathbb{C}_n^2$ of n copies of $\mathbb{C}^2$, called *qubits*. Physically, this models a system of n non-interacting spin=$\frac{1}{2}$ particles. The circuit then consists of a sequence of K "gates" $U_k, 1 \leq k \leq K$, applied to individual or paired tensor factors. A gate is some coherent interaction; mathematically it is a unitary transformation on either $\mathbb{C}_i^2$ or $\mathbb{C}_i^2 \otimes \mathbb{C}_j^2$, $1 \leq i, j \leq n$, the identity on all remaining factors. The gates are taken from a fixed finite library of unitary 2×2 and 4×4 matrices (with respect to a fixed basis $\{|0\rangle, |1\rangle\}$ for each $\mathbb{C}^2$ factor) and must obey the surprisingly mild condition, called "universality", that the set of possible gate applications generates the unitary group $\mathbb{U}(2^n)$ densely (up to a physically irrelevant overall phase.) Popular choices include a relative phase gate $\begin{pmatrix} 1 & 0 \\ 0 & e^{\frac{2\pi i}{5}} \end{pmatrix}$ and "Controlled NOT" $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$ operating on one and two "particles", respectively. It is known that beyond the density requirement the particular choice of gates is not too important. Let $W_\Gamma = \prod_{i=1}^m U_i$ denote the operator effected by the circuit Γ. It is important for the fault tolerance theory that many gates can be applied simultaneously (to different qubits) without affecting the output of the circuit $W_\Gamma(|0\rangle \otimes \cdots \otimes |0\rangle)$.

Formally, information is extracted from the output by measuring the first qubit. The probability of observing $|1\rangle$ is given according to the usual axioms of quantum mechanics as:

$$(1) \qquad\qquad p(\Gamma) = \langle 0|W_\Gamma^\dagger \Pi_1 W_\Gamma |0\rangle,$$

where Π_1 is the projection to $|1\rangle$, $\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$, applied to the first qubit. Any decision problem, such as finding the k-th binary digit of the largest prime factor of an integer x, can be modeled by a function F on binary strings, $F : \{0,1\}^* \to \{0,1\}$; in our example, the input string would encode (x, k). We say F belongs to **BQP** if there is a classical polynomial-time (in string length) algorithm for specifying a

"quantum circuit" $\Gamma(y)$ (in the example $y = (x, k)$) which satisfies:

$$p(\Gamma(y)) \geq \tfrac{2}{3} \quad \text{if} \quad F(y) = 1 \quad \text{and}$$
$$p(\Gamma(y)) \leq \tfrac{1}{3} \quad \text{if} \quad F(y) = 0.$$

This definition suggests that one needs to make an individual quantum circuit to solve each instance of a computational problem. However, it is possible to construct a single circuit to solve any instance of a given **BQP** problem of bounded size, e.g., factor integers with ≤ 1000 digits. Moreover, by [19] there is a universal circuit, univ.(n,k), which simulates all circuits of size k on n qubits:

$$W_{\text{univ.(n,k)}}\Big(|0\cdots0\rangle \otimes |\Gamma\rangle\Big) = W_\Gamma|0\cdots0\rangle \otimes |\Gamma\rangle.$$

Yet another definition allows one to do measurements in the middle of computation and choose the next unitary gate depending on the measurement outcome. This choice will generally involve some classical Boolean gates. This scheme is called *adaptive quantum computation*. In certain cases, it can squeeze general **BQP** computation out of a gate set which is *not* universal.

IMPLEMENTATION OF A QUANTUM COMPUTER

It is not possible to realize a unitary gate precisely, and even when we do not intend to apply a gate, the environment will interact with the qubits causing decoherence. Both imprecision and decoherence can be considered in a unified way as "errors" which can be quantified by a fidelity distance [17] or a super-operator norm [19]. A crucial step in the theory of quantum computing has been the discovery of error-correcting quantum codes [28] and fault-tolerant quantum computation [25], [29]. These techniques cope with sufficiently small errors. However, the error magnitude must be smaller than some constant (called *an accuracy threshold*) for these methods to work. According to rather optimistic estimates, this constant lies between 10^{-5} and 10^{-3}, beyond the reach of current technologies.

The presumptive path toward a quantum computer includes these steps: (1) build physical qubits and physical gates; (2) minimize error level down below the threshold; (3) implement decoherence-protected logical qubits and fault-tolerant logical gates (one logical qubit is realized by several qubits using an error-correcting code). As a counterpoint, the theme of this article is that implementing physical qubits might be redundant. Indeed, one can "encode" a logical qubit into a physical system which is not quite a set of qubits or even well separated into subsystems. Such a system must have macroscopic quantum degrees of freedom which would be decoherence-protected. A super-conducting phase or anything related to a local order parameter will not work, for if a degree of freedom is accessible by local measurement, it is vulnerable to local error. However, there is another type of macroscopic quantum degree of freedom. It is related to topology and arises in collective electronic systems, e.g. the fractional quantum Hall effect [13] and most recently in $2D$ cuprate superconductors above T_c [12], [23].

Though much studied since the mid-1980's the connection between fractional quantum Hall effect and quantum computation has only recently been realized [11], [20]. It was shown by [4] that the ground state of the $\nu = \tfrac{1}{3}$ electron liquid on the torus is 3-fold degenerate. This follows from the fact that excitations in this system are abelian anyons: moving one excitation around another multiplies the state vector by a phase factor $e^{i\phi}$ (in this case $\phi = \tfrac{2\pi}{3}$). The process of

creating a particle-antiparticle pair, moving one of the particles around the torus, and annihilating it specifies a unitary operator on the ground state. By moving the particle in two different directions, one obtains two different unitary operators A_1 and A_2 with the commutation relation $A_1 A_2 A_1^{-1} A_2^{-1} = e^{i\phi}$, implying a ground state degeneracy. This argument is very robust and only requires the existence of an energy gap or, equivalently, finite correlation length l_0. Indeed, the degeneracy is lifted only by spontaneous tunneling of virtual excitations around the torus. The resulting energy splitting scales as $e^{-\frac{l}{l_0}}$, where l is the size of the system. Interaction with the environment does not change this conclusion, although thermal noise can create actual excitation pairs rather than virtual ones.

Both the ground state degeneracy on the torus and the existence of anyons are manifestations of somewhat mysterious topological properties of the $\nu = \frac{1}{3}$ electron liquid itself. Anyons can be regarded as "topological defects" similar to Abrikosov vortices but without any local order parameter. The presence of a particle enclosed by a loop on the plane can be detected by holonomy—moving another particle around the loop.

At $\nu = \frac{1}{3}$, the electron liquid on the torus could be used as a logical "qutrit" (generalized qubit with 3 states). Unfortunately, this will hardly work in practice. Besides the obvious problem with implementing the torus topology, there is no known way to measure this logical qutrit or prepare it in a pure state.

A more flexible and controllable way of storing quantum information is based on nonabelian anyons. These are believed to exist in the $\nu = \frac{5}{2}$ fractional quantum Hall state. According to the theory [22], [24], there should be charge $\frac{1}{4}$ anyonic particles and some other excitations. The quantum state of the system with $2n$ charge $\frac{1}{4}$ particles on the plane is 2^{n-1}-degenerate. The degeneracy is gradually lifted as two particles come close to each other. More precisely, the 2^{n-1}-dimensional Hilbert space $\mathbb{H}_n$ splits into two 2^{n-2}-dimensional subspaces. They correspond to two different types of charge $\frac{1}{2}$ particles which can result from fusion of charge $\frac{1}{4}$ particles. Thus observing the fusion outcome effects a measurement on the Hilbert space $\mathbb{H}_n$. This model supports adaptive quantum computation when surfaces of high genus are included in the theory and admits a combinatorial description [1] apparently in the same universal class as the fractional quantum Hall fluid.

Beyond this, a discrete family of quantum Hall models exists [26] based on $k+1$-fold hard-core interaction between electrons in a fixed Laudau level which appears to represent the same universality class as Witten-Chern-Simons theory for $SU(2)$ at level$=k$ [33]. Anyons in these models behave as topological defects of a geometric construction [10] and their braiding matrices have been shown to be universal [8], [9] for $k \geq 3, k \neq 4$.

CODE SPACES AND QUANTUM MEDIA

Even after the particle types and positions of anyons are specified, there is an exponentially large (but finite dimensional) Hilbert space describing topological degrees of freedom. In combinatorial models, this Hilbert space becomes a "code subspace" W of a larger "quantum media" Y. Thus a fundamental concept of cryptography is transplanted into physics. Let V be a finite dimensional complex vector space, perhaps $\mathbb{C}^2$, and $Y = V \otimes \cdots \otimes V$ an n-fold tensor product (where n is typically quite large). Let $W \subset Y$ be a linear subspace. We call $W \subset Y$ k-code

if and only if:

$$W \stackrel{\Pi_W \cdot \mathcal{O}}{\longrightarrow} W$$

is multiplication by a scalar whenever $\mathcal{O}$ is a k-local operator (an arbitrary linear map on any k-tensor factors of Y and the identity on the remaining $n - k$ factors) and Π_W is the orthogonal projection onto W. We think of such spaces as resisting local alteration and in the usual interpretation of Y as the Hilbert space of n particles, quantum information stored in W will be relatively secure. It is a theorem [14] that the quantum information in W cannot be degraded by errors operating on fewer than $\frac{k}{2}$ of the n particles.

Let us define a (discrete) quantum medium to be a tensor product $Y = \otimes_i V_i$ as above, where now the set of indices $\{i\}$ consists of points distributed on a geometric surface T, together with a local Hamiltonian $H = \sum H_i$ (each H_i is a Hermitian operator defined only on those tensor factors whose index is within $\epsilon \rangle 0$ of the i-th point in the geometry of the surface). Local Hamiltonians H have been found [10], [20] with highly d-degenerate ground states corresponding to modular functors [31], [32] (and thus braid group representation [16] and link polynomials [18]). In these cases, the ground state G of H will be k-code for $k \sim$ injectivity radius of T $\sim \sqrt{\text{area T}}$. The topological degrees of freedom referred to above reside in G. But we do not attempt a precise mathematical definition of topological degrees of freedom since we would like it to extend beyond discrete system, e.g., to fractional quantum Hall ground states.

In the case when T is a disk D with punctures—physically *anyonic excitations*—a sequence of local modifications to H (see [10]) effects a discrete 1-parameter family H_t of Hamiltonians, where the ground states G_{t_i} and $G_{t_{i+1}}$ at consecutive time steps differ by a $\lfloor \frac{k}{2} \rfloor$-local operator $\mathcal{O}_i$, $\mathcal{O}_i(G_{t_i}) = G_{t_{i+1}}$. Note that if G_{t_i} and $G_{t_{i+1}}$ are both k-code that for $\mathcal{O}_i$ as above, $\mathcal{O}_i|_{G_{t_i}}$ must be unique up to a scalar (for a distinct $\mathcal{O}_i'$, consider the restriction of the k-local operator $\mathcal{O}_i{}^\dagger \circ \mathcal{O}_i'$ projected to G_{t_i}, $\Pi_{G_{t_i}} \circ \mathcal{O}_i{}^\dagger \circ \mathcal{O}_i'$). This uniqueness property forces this discrete-cryptographic transport of code spaces to coincide (up to phase) with the differential geometric notion of adiabatic transport—integration of the canonical connection in the "tautological" bundle of d-planes in $Y = \mathbb{C}^{2^n}$. If the 1-parameter family is a closed loop, a projective representation of the braid group on $\mathbb{U}(\text{code})$ is obtained. By choosing H, these can be engineered to be precisely the Hecke algebra representations $\{\rho_\lambda\}$ associated to the Jones polynomial. From H one can build a concrete model of quantum computation; the model and its connection to the Jones polynomial are described below. Although the H found in [10] is enormously cumbersome, it appears to share the universality class "Witten-Chern-Simons theory of $SU(2)$ at level 3" with a simple 4-body Hamiltonian [26], which has been proposed as a model for the fractional quantum Hall plateaus at $\nu = \frac{13}{5}$ and $\frac{12}{5}$.

AN ANYONIC MODEL FOR QUANTUM COMPUTATION

A family of unitary representations of all mapping class and braid groups with certain compatibility properties under fusion is known as a unitary topological modular functor. To define our model, we take only the planar surface portion of the simplest universal modular functor, Witten-Chern-Simons $SU(2)$ modular functor at level 3, and this reduces to the Jones representations of braids $\{\rho_\lambda\}$. For an appropriate family of local Hamiltonians H_t, these representations describe
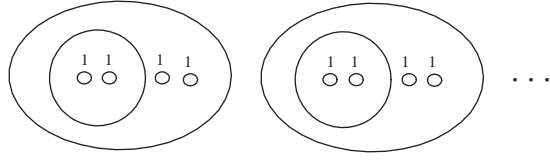
Figure 1.

the adiabatic transport of the lowest energy states with n anyonic excitation pairs; these states form a subspace W of dimension Fibonacci(2n)—as the $2n$ anyonic excitations are "braided" around each other in 2+1 dimensional space time. In this theory (we denote it $CS5$ because of its link to the Jones polynomial at a fifth root of unity), there are 4 label types 0,1,2,3 corresponding to the complete list of irreducible representations of the quantum group $SU(2)_5$ of dimensions 1,2,3 and 4 (or equivalently the irreducible positive energy loop group representations of $\mathcal{L}SU(2)$ at level 3). We initialize our system on the disk D in a known state by pulling anyonic pairs out of the vacuum. This theory is self-dual so the two partners have identical types. Pairs are kept or returned to the vacuum according to the results of local holonomic measurement. Finally we have a known initial state in the disk with $2n$ punctures where each puncture has label=1 and ∂D has label=0. We assume n is even and group the punctures into $n/2$ batches of 4 punctures each. Similar to [8], each batch B is used to encode one qubit $\cong \mathbb{C}^2$: the basis $\{|0\rangle, |1\rangle\}$ is mapped into the type (0 or 2) of the 2-fold composite particle (round circle), which would result from fusing a (fixed) pair of the type 1 particles within B. Initially, both the double and 4-fold composites (ovals) have type 0. By maintaining, at least approximately, this condition on the ovals after the braiding is complete, we define a "computational summand" of the modular functor: it is spanned by n-bit strings of 0's and 2's residing on the 2-fold composites (round circles).

Now as in the quantum circuit model, a classical poly-time algorithm looks at the problem instance (F, y) and builds a sequence of "gates", but this time the gates are braid generators (right half twist between adjacent anyons) σ_i , $1 \le i \le 2n-1$, and a powerful approximation theorem [19], [30] is used to select the braid sequence which approximates the more traditional quantum circuit solving (F, y). So the topological model may be described as:

(1) Initialization of a known state in the modular functor.
(2) Classical computation of braid b effecting a desired unitary transformation X of the computational subspace of the modular functor.
(3) Adiabatic implementation of the braid by (somehow) moving the anyons in D to draw b in space-time. (Here we keep the anyons separated by a scale l.)
(4) Application of a projection operator Π to measure the type (0 or 2) of the "left most" composite particle (as seen in Figure 1).

The last step is the direct analogue of measuring the first qubit in the quantum circuit model and the same formula (1) applies: the probability of observing type 0 (the null particle) is:

$$\text{prob}(0) = \langle 0|X^\dagger \Pi X|0\rangle,$$

where the 0's on the right hand side represent our carefully prepared initial state with $2n$ type 1 excitations.

$$\text{L} = \qquad b^{-1}$$
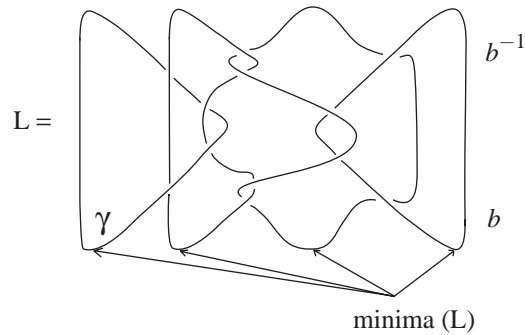$$\gamma \qquad b$$
$$\text{minima (L)}$$

FIGURE 2.

To close the topological discussion, we note that the previous formula can be translated (using the S matrix) into a plat closure (see Figure 2) of a braid $L = \mathrm{plat}(b^{-1}\gamma b)$, where γ is a small loop inserted to measure the left-most qubit, and now the outcome of the quantum circuit calculation, $\mathrm{prob}(0)$, becomes a Jones evaluation

$$\mathrm{prob}(0) = \frac{1}{1 + [2]_5^2} \left(1 + \frac{(-1)^{c(L)+w(L)}(-a)^{3w(L)}V_L(e^{2\pi i/5})}{[2]_5^{m(L)-2}} \right),$$

where $[2]_5 = \frac{1+\sqrt{5}}{2}$; c, m, and w are the number of components, number of local minima, and writhe of L respectively; and $a = e^{\pi i/10}$. Details of this calculation will be posted at URL: www.tqc.iu.edu.

The braiding disturbs and reforms composite particle types with sufficient subtlety to effect universal computation. To reduce this model to engineering, very significant obstacles must be overcome: stable quantum media must be maintained in a suitable phase, e.g., $CS5$; excitations must be readily manipulated; and electrical neutral particle types 0 and 2 must be distinguished, presumably by holonomy experiments. Although these challenges are daunting, they are, perhaps, less difficult than a head-on assault on the accuracy threshold in the quantum circuit model.

REFERENCES

1. S. Bravyi, A. Kitaev, private communications.
2. D. Deutsch, *Quantum theory, the Church-Turing principle, and the universal quantum computer,* Proc. Roy. Soc. London, **A400** (1985), 97-117. MR **87a:**81017
3. D. Deutsch, *Quantum computational networks,* Proc. Roy. Soc. London, **A425** (1989), 73-90. MR **90k:**81023
4. T. Einarsson, *Fractional statistics on a torus,* Phys. Rev. Lett. **64** (1990), 1995-1998. MR **91b:**81051
5. R. Feynman, *Simulating physics with computers,* Int. J. Theor. Phys. **21** (1982), 467-488. MR 658311
6. R. Feynman, *Quantum mechanical computers,* Found. Phys., **16** (1986), 507-531. MR **88i:**81022
7. M. Freedman, A. Kitaev, and Z. Wang, *Simulation of topological field theories by quantum computers*, Comm. Math. Phys. **227** (2002), no. 3, 587–603.
8. M. Freedman, M. Larsen, and Z. Wang, *A modular functor which is universal for quantum computation*, Comm. Math. Phys. **227** (2002), no. 3, 605–622.

9. M. Freedman, M. Larsen, and Z. Wang, *The two-eigenvalue problem and density of Jones representation of braid groups*, Comm. Math. Phys. **228** (2002), no. 1, 177–199. CMP 2002:14

10. M.H. Freedman, *Quantum computation and the localization of modular functors*, Found. Comput. Math. **1** (2001), no. 2, 183–204.

11. M.H. Freedman, *P/NP, and the quantum field computer*, Proc. Natl. Acad. Sci. USA **95** (1998), no. 1, 98–101. MR **99b:**68064

12. T. Senthil and M.P.A. Fisher, *Fractionalization, topological order, and cuprate superconductivity*, cond-mat/0008082.

13. S. Girvin, The quantum Hall effect: novel excitations and broken symmetries, in Topological aspects of low dimensional systems, Edited by A. Comtet, T. Jolicoeur, S. Ouvry, and F. David EDP Sci., Les Ulis, 1999.

14. D. Gottesman, *Theory of fault-tolerant quantum computation,* Phys. Rev. Lett. **A57** (1998), 127-137.

15. L. Grover, *Quantum Mechanics helps in search for a needle in a haystack,* Phys. Rev. Lett., **79** (1997), 325-328.

16. V.F.R. Jones, *Hecke algebra representations of braid groups and link polynomials,* Ann. Math., **126** (1987), 335-388. MR **89c:**46092

17. R. Jozsa, *Fidelity for mixed quantum states,* Journal of Modern Optics, **41** (1994), no. 12, 2315-2323.

18. L. Kauffmann and S. Lins, *Temperley-Lieb recoupling theory and invariants of 3-manifolds*, Ann. Math. Studies, vol 134, Princeton Univ. Press, 1994. MR **95c:**57027

19. A. Kitaev, *Quantum computations: algorithms and error correction,* Russian Math. Surveys, **52:61** (1997), 1191-1249. MR **99a:**68061

20. A. Kitaev, *Fault-tolerant quantum computation by anyons,* quant-ph/9707021.

21. S. Lloyd, *Universal quantum simulators,* Science, **273** (1996), 1073-1078. MR **97h:**81031

22. G. Moore and N. Read, *Nonabelians in the fractional quantum Hall effect,* Nucl. Phys, **B360** (1991), 362-396. MR **92j:**81291

23. C. Nayak, and K. Shtengel, *Microscopic models of two-dimensional magnets with fractionalized excitations*, Phys. Rev. B, **64**:064422 (2001).

24. C. Nayak and F. Wilczek, $2n$-quasihole states realize 2^{n-1}-dimensional spinor braiding statistics in paired quantum Hall states, Nucl. Phys **B479** (1996), 529-553. MR **97m:**81053

25. J. Preskill, *Fault tolerant quantum computation,* quant-ph/9712048.

26. N. Read, and E. Rezayi, *Beyond paired quantum Hall states: parafermions and incompressible states in the first excited Laudau level*, cond-mat/9809384.

27. P. Shor, *Algorithms for quantum computation, discrete logarithms and factoring,* Proc. 35th Annual Symposium on Foundations of Computer Science, IEEE Computer Society Press, Los Alamitos, CA, 1994, 124-134.

28. P. Shor, *Scheme for reducing decoherence in quantum computer memory*, Phys. Rev. A **52**, 2493 (1995).

29. P. Shor, *Fault-tolerant quantum computation,* Proc. 37th Annual Symposium on Foundations of Computer Science, IEEE Computer Society Press, Los Alamitos, CA, 1996.

30. R. Solvay, private communication.

31. V. Turaev, *Quantum invariants of knots and 3-manifolds*, de Gruyter Studies in Math., vol. 18, 1994. MR **95k:**57014

32. K. Walker, *On Witten's 3-manifold invariants*, preprint, 1991 (available at http://www.xmission.com/~kwalker/math/).

33. E. Witten, *Quantum field theory and the Jones polynomial,* Comm. Math. Phys., **121** (1989), 351-399. MR **90h:**57009

(M. H. Freedman) Microsoft Research, One Microsoft Way, Redmond, Washington 98052

(A. Kitaev) On leave from L.D. Landau Institute for Theoretical Physics, Kosygina St. 2, Moscow, 117940, Russia
Current address: Microsoft Research, One Microsoft Way, Redmond, Washington 98052

(M. J. Larsen and Z. Wang) Indiana University, Department of Mathematics, Bloomington, Indiana 47405