Quantum Information Processing, Vol. 2, Nos. 1–2, April 2003 (© 2003)

On Freedman's Lattice Models for Topological Phases

James Brink[1] and Zhenghan Wang[2,3]

Received March 7, 2003; accepted March 18, 2003

Freedman proposes a family of Hamiltonians which define quantum loop gas models on any celluated compact surface. We study the simplest nontrivial cases: celluations of the torus. Our numerical data support Freedman's conjecture about the ground states of the Hamiltonians.

KEY WORDS: topological phase; lattice model.

PACS: 71.10.−w, 71.35.−y

1. INTRODUCTION

The program of topological quantum computation is to realize fault tolerant quantum computation using topological phases of quantum systems.[2] The central open question is whether or not there exist such physical systems which are capable of performing universal quantum computation. In Ref. 1, a family of Hamiltonians $H_{0,l}$ is proposed as candidates for the Chern–Simons phases, which are known to support universal quantum computation for each level $l \geq 3,\ l \neq 4$.[3,4] Freedman conjectures that the perturbed ground states of $H_{0,l}$ are given by the Drinfeld double of the $SO(3)$-Witten–Chern–Simons topological quantum field theories (TQFTs). The approach in Ref. 1 is an algebraic study of the effect of perturbation based on a rigidity result of the picture TQFTs.[5] The picture TQFTs are the Drinfeld double of the $SO(3)$-Witten–Chern–Simons TQFTs. The idea is to treat local relations in picture TQFTs as perturbations. In this paper we will investigate numerically the perturbed ground states of $H_{0,l}$ for some

[1]Department of Math., University of California, Berkeley. E-mail: JBrink@Math.Berkeley.Edu
[2]Department of Math., Indiana University. E-mail: zhewang@indiana.edu
[3]To whom corresspondence should be addressed.

1570-0755/03/0400–0081/0

computationally tractable cases. While numerical study of Hamiltonians is carried out routinely in physics literature, we face a dilemma here. The major motivation of our investigation is quantum computing, but one motivation of quantum computing as Feynman pointed out is to study quantum systems numerically. Therefore, we are in a self-referential situation. This is also manifested in the fact that our numerical computation quickly reaches the limit of present computing power.

Freedman's Hamiltonians $H_{0,l}$ define quantum loop gas models on any celluated compact surface. We study the simplest nontrivial cases: celluations of the torus. Our numerical data support Freedman's conjecture, but the conjectured space of ground states does not come out in full. Study of substantially larger systems is necessary, but computationally intractable now. One new phenomenon we discovered is some lonely states in the ground states of $H_{0,l}$. Those lonely states give rise to unwanted ground state vectors of $H_{0,l}$ which persist in the perturbed ground states. There are several possible explanations of those phenomena: the small size of the system, the choice of our perturbation, or the Euclidean geometry of the torus. We also observe clearly the expected energy gap between ground states and the first excited states.

The Hamiltonians $H_{0,l}$ are 7th order interaction Hamiltonians. It is an open question to find 2nd order or 3rd order interaction Hamiltonians with approximately the same ground states. Note that Ref. 1 also contains a family of 4th order interaction Hamiltonians similar to $H_{0,l}$.

2. FREEDMAN'S HAMILTONIAN

At the heart of this study lie the Hamiltonians $H_{0,l}$, for levels $l \geq 1$, which grants the structure of a TQFT to the hypothetical physical systems. Each Hamiltonian $H_{0,l}$ is described as a sum of local projections, which "implement" the concept of combinatorial generalized isotopy (*g(d)-isotopy*) on a celluated surface.

2.1. Combinatorial Isotopy

Fix a closed oriented surface Σ, and let Δ be a triangulation of Σ with n vertices. The dual graph Δ^* to Δ therefore defines a celluation of Σ by n polygons, with the edges of Δ^* being boundaries of adjacent cells.

A *spin configuration* is an assignment $s : \Delta^* \to \{+, -\}$ of a positive or negative spin to each cell in Δ^*. We denote a "spin flip" by $^-$: on entire spin configurations, $\bar{s}$ denotes a global interchange of $+$ and $-$; for a cell c of

Δ^*, $\bar{s}^c$ denotes the spin configuration which agrees with s away from c, and flips the spin at c.

A spin configuration can be thought of as a 2-coloring of Σ, partitioning Σ into $+$ and $-$ regions. We will be interested in studying the boundary ∂_s of these $+$ and $-$ regions with respect to spin configurations s. We will refer to ∂_s as *domain walls* of a spin configuration. Since Δ^* is dual to a triangulation, the domain walls are all 1-manifolds in Σ. Our goal will be to capture global properties of domain walls ∂_s in terms of local manipulations at each cell.

Thus, for a fixed cell c, we will need to consider the boundary ∂c of c, which consists of edges in Δ^*. Color each edge on ∂c according to the neighboring cell; set $\partial_+ c$ to be the edges which bound a $+$-cell, and $\partial_- c$ to be the edges which bound a $-$-cell. We then say that the pair (s, c) is *type-g* if both $\partial_+ c$ and $\partial_- c$ form a connected topological arc. Note that neither $\partial_+ c$ nor $\partial_- c$ can be empty. We define the pair (s, c) to be *type-h* if c and all neighboring cells have the same spin ($\partial_{s(c)} c = \partial c$).

The motivation for these definitions is as follows: if (s, c) is type-g, then the $\pm$ boundary ∂_s meets ∂c along one of the topological arcs $\partial_+ c$ or $\partial_- c$ (the one which is opposite the spin of c). By changing the sign of c, we replace the part of ∂_s which meets ∂c with the other of the topological arcs. But this transformation can be viewed as a fixed endpoint isotopy of one arc to the other, with isotopy occurring over the 2-cell c (see the figure below). So, if (s, c) is type-g, then ∂_s and $\partial_{\bar{s}^c}$ are isotopic.

For type-h pairs (s, c), the cell c and all of its neighbors have the same

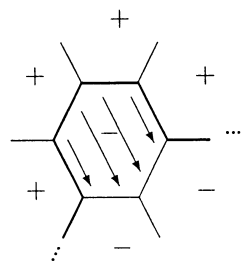

spin. So, if the spin is flipped on cell c, then a closed loop is added to the boundary ∂_s. Type-h cells are therefore responsible for the "generalized" aspect of g(d)-isotopy, namely, the insertion or removal of disc-bounding loops.

Using this terminology, we can consider a combinatorial form of g(d)-isotopy, relative to the celluation Δ^*. Given a spin configuration, we can apply: (i) type-g moves, which consist of flipping the spin of a type-g cell; and (ii) type-h moves, which consist of flipping the spin of a type-h cell. We say that two spin configurations are *combinatorially g(d)-isotopic* if one can be reached from the other via a sequence of such moves. Then we have the following proposition:

Proposition 1. Let s and t be combinatorially g(d)-isotopic spin configurations with respect to the celluation Δ^*. Then ∂_s is g(d)-isotopic to ∂_t.

The comments above show that this result is immediate. The converse is not quite true: it is possible to have non-combinatorially g(d)-isotopic spin configurations s and t such that ∂_s and ∂_t are g(d)-isotopic. The reason for this is that the tiling could be too coarse to allow room for g(d)-isotopy to take place; for example, on a hexagonal tiling of the torus, a spin configuration which assigned different colors to adjacent vertical rows would have no type-g or type-h cells(but its boundary would be g(d)-isotopic to that of its dual).

2.2. Definition of $H_{0,l}$

We are now ready to consider the definition of the Hamiltonian $H_{0,l}$. For a triangulation Δ with n vertices, we associate the 2^n-dimensional Hilbert space $\mathcal{H} = \bigotimes_{i=1}^{n} \mathbb{C}^2$; let $c_1, \ldots, c_n$ denote the dual 2-cells in Δ^*. Then we can express a basis for $\mathcal{H}$ by $\{|s\rangle\}$, where s runs over spin configurations and $|s\rangle = |s(c_1)\rangle \otimes |s(c_2)\rangle \otimes \cdots \otimes |s(c_n)\rangle$. Set the parameter $d = 2\cos \pi/(l+2)$, and put:

$$H_{0,l} = \sum_{\substack{(s,c) \\ \text{type g}}} |s - \bar{s}^c\rangle\langle s - \bar{s}^c| + \sum_{\substack{(s,c) \\ \text{type h}}} \left|s - \frac{1}{d}\bar{s}^c\right\rangle\left\langle s - \frac{1}{d}\bar{s}^c\right|$$

We interpret these terms as follows: $H_{0,l}$ is "indifferent" to g(d)-isotopy, hence the equal weighting on the type-g terms. The factor of $1/d$ expresses that "loops are worth d", as flipping the type-h configuration creates a closed bounding loop in the spin configuration.

The primary observation, to be proved in the next section, is that the ground state $G_{0,l}$ of $H_{0,l}$ consists of equivalence classes of combinatorially g(d)-isotopic spin configurations.

3. THEORETICAL ANALYSIS

First we consider analytic results relevant to the conjecture. The ground state of $H_{0,l}$ is determined explicitly.

3.1. The Ground State $G_{0,l}$

For a spin configuration s, let $n(s)$ denote the number of trivial (disc-bounding) closed loops in s. Thus if $s\prime$ is the configuration s with all trivial closed loops removed, we have $d^{n(s)}s\prime \simeq s$.

Theorem 2. Fix a surface Σ, a triangulation Δ of Σ, and let $\simeq_\Delta$ be the equivlance relation of combinatorial g-isotopy on spin configurations with respect to Δ. Then the ground state $G_{0,l}$ of the corresponding Hamiltonian $H_{0,l}$ has as a basis:

$$|v_i\rangle = \sum_{s \simeq_\Delta s_i} d^{n(s)}|s\rangle$$

where $d = 2\cos\pi/(l+2)$ and the $s_1, , s_k$ are representatives of the k equivalence classes determined by $\simeq_\Delta$.

Proof. First, observe that the subspace $V_i = \text{span}\,\{|s\rangle : s \simeq_\Delta s_i\}$ is invariant under $H_{0,l}$; this is immediate from the definition of $H_{0,l}$, as each projector acts invariantly on one of the V_i and trivially on V_j for $i \neq j$. So we can write $\mathcal{H} = \bigoplus_i V_i$, with $H_{0,l}$ acting invariantly on each summand; therefore, the ground state can be likewise decomposed, $G_{0,l} = \bigoplus_i \ker(H_{0,l}|_{V_i})$. So the theorem reduces to the claim that $\ker(H_{0,l}|_{V_i}) = \text{span}|v_i\rangle$.

First, we show that $|v_i\rangle$ lies in $\ker(H_{0,l}|_{V_i})$. For any $s \simeq_\Delta s_i$, we have:

$$\langle s|H_{0,l}|v_i\rangle = \sum_{\substack{\text{cells } c \text{ s.t.} \\ (s^c,c) \text{ type g}}} d^{n(s)} - d^{n(\bar{s}^c)} + \sum_{\substack{\text{cells } c \text{ s.t.} \\ (s,c)\text{type h}}} d^{n(s)} - \frac{1}{d} d^{n(\bar{s}^c)}$$

$$+ \sum_{\substack{\text{cells } c \text{ s.t.} \\ (\bar{s}^c,c)\text{type g}}} d^{n(s)} - d^{n(\bar{s}^c)} + \sum_{\substack{\text{cells } c \text{ s.t.} \\ (\bar{s}^c,c)\text{type h}}} \frac{1}{d^2} d^{n(s)} - \frac{1}{d} d^{n(\bar{s}^c)}$$

For any (s, c) that is type-g, we have $n(s) = n(\bar{s}^c)$, since the two configurations are (strictly) isotopic. Similarly, if (s, c) is of type-h, then $n(\bar{s}^c) = n(s) + 1$, so that $d^{n(s)} - (1/d)d^{n(\bar{s}^c)} = 0$; when $(\bar{s}^c, c)$ is "type h", then

$(1/d^2)d^{n(s)} - (1/d)d^{n(\bar{s}^c)} = 0$. So all terms in the summation cancel; since this holds for all $s \simeq_\Delta s_i$, this shows that $|v_i\rangle \in \ker(H_{0,l}|_{V_i})$.

The converse claim follows similarly; for, suppose some nonzero vector $|v\rangle \in V_i$ is in the kernel of $H_{0,l}|_{V_i}$. Then $\langle s|v\rangle \neq 0$ for some $s \simeq_\Delta s_i$, so that $H_{0,l}|s\rangle\langle s|v\rangle$ will have nonzero components along each $|s'\rangle$ which differs from s by one type-g or type-h move. So in order for $|v\rangle \in \ker(H_{0,l}|_{V_i})$, it must also have nonzero components along these $|s\rangle$. Continuing this argument, we see that $\langle s|v\rangle \neq 0$ for all $s \simeq_\Delta s_i$. Analyzing the above computation, we see that if s and s' differ by a type-g move, then $\langle s|v\rangle = \langle s'|v\rangle$; and if they differ by a type-h move (say, s' is s with a loop removed), then $d\langle s|v\rangle = \langle s'|v\rangle$. These requirements in turn force $|v\rangle$ to be a multiple of $|v_i\rangle$, so that $\ker(H_{0,l}|_{V_i}) = \text{span } |v_i\rangle$.

4. NUMERICAL STUDY

Celluations of a surface Σ dual to triangulations have a nice topological property. The domain walls in a spin configuration are all 1-manifolds. The homology class represented by those domain walls is the zero class as the domain walls are bounding 1-manifolds. The linear combinations of domain walls coming from different spin configurations in a fixed celluation forms a finite dimensional vector space. Those vector spaces are combinatorial approximations of a picture TQFT. Local relations in picture TQFTs are the topological realization of perturbations. The rigidity of the picture TQFTs says the only non-trivial local relations for each l is generated by the Jones–Wenzl projector. Therefore, for each level l perturbations of the Hamiltonian $H_{0,l}$ will result in only one possibility: the picture TQFTs at level l. The purpose of the present study then is to perform numerical analysis on computationally tractable cases, in order to test this conjecture. We will focus on perturbations of the form $H_\varepsilon = H + \varepsilon V$. As a preliminary case, we will choose V to the sum of σ_x on each cell. Our study focuses on $l = 3$, so the loop value is $d = (\sqrt{5} + 1)/2$.This corresponds to the double of the $SO(3)$-Witten–Chern–Simons theory at level $= 3$. In Ref. 3, we have shown that the level $l = 3$ theory for both $SU(2)$ and $SO(3)$ supports universal quantum computation.

4.1. On Tilings

The numerical study focused on tilings of the torus. Nine hexagonal tilings of the torus were considered, representing the "obvious" hexagonal tilings of the torus which are computationally tractable. These tilings are

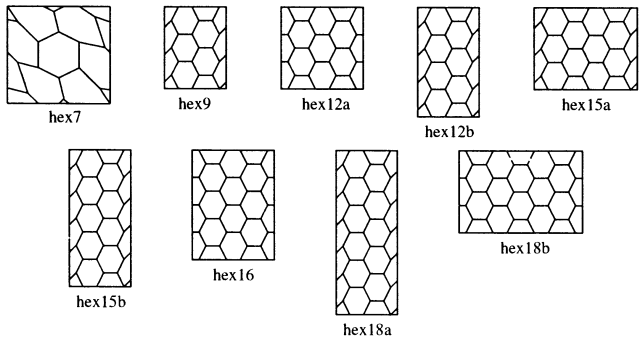

Fig. 1. The hexagonal tilings of the torus used in the numerical study.

depicted in Figure 1, on the torus visualized in the standard way as a rectangle with opposite sides identified.

The tiling **hex7** is the dual of a minimal triangulation of the torus, and thus is a minimal tiling. The other tilings are, loosely speaking, formed by p adjacent vertical columns of q cells; when p is odd, an extra "twist" is needed to align the vertices properly. Thus, tilings with $(p, q) = (3,3)$, $(4,3)$, $(3,4)$, $(5,3)$, $(3,5)$, $(4,4)$, $(3,6)$, $(6,3)$ are represented.

4.2. Ground State Vectors of $H_{0,l}$

First, a combinatorial computation was performed to verify the result of Theorem 2, regarding the form of the ground state vectors of $H_{0,l}$. Spin configurations were grouped into g-isotopy classes, and vectors corresponding to each class were created according to the formula of Theorem 2. Each such vector was found to be in the ground state of $H_{0,l}$; further, the number of such (orthogonal) vectors was equal to the dimension of the ground state, as calculated via an eigenvalue computation, so that the ground state was numerically verified to be exactly the span of such vectors, in all cases considered.

The interesting result of these calculations concerns what we will call *lonely* configurations: spin configurations which are not g-isotopic to any other spin configuration. (In other words, configurations in which there are no type-g or type-h cells.) Some of the tilings considered admit such lonely configurations, and others do not; in considering the numerical analysis of the perturbed Hamiltonians below, it will be necessary to take these lonely configurations into account.

We considered nine different hexagonal tilings of the torus, ranging from 7 cells (the minimal hexagonal tiling of the torus, corresponding to the minimal triangulation of the torus) to 18 cells (the maximum computation-

Table 1. Ground States of $H_{0,1,}$ with Respect to Given Tilings

Tiling	n	dim $G_{0,3}$	Lonely configurations	Non lonely
hex7	7	5	0	5
hex9	9	5	0	5
hex12a	12	8	2	6
hex12b	12	17	12	5
hex15a	15	7	0	7
hex15b	15	8	0	8
hex16	16	24	18	6
hex18a	18	16	8	8
hex18b	18	21	14	7

ally tractable case). The following table presents the results of this preliminary analysis of the tilings:

4.3. Perturbation Ground States $G_{\varepsilon,I}$

For the numerical analysis of the perturbed ground states, we calculated the lowest energy eigenvalues of $H_{\varepsilon,3}$, for ε ranging from 0 to 1. Plots of the eigenvalues as a function of ε are given below, for each of the nine tilings in the table above.

As the collected data simply provides the lowest eigenvalues, the plotted lines simply indicate the trajectories of the various eigenvalues. The lines themselves do not follow a particular eigenvalue, but rather, they connect the lowest eigenvalue, the second-lowest eigenvalue, etc.; in other words, the plots below should be treated more like scatterplots. The connecting lines are included to help elucidate the trajectories of the eigenvalues and make it easier to determine where eigenvalues converge and diverge.

Also, the plots do not indicate the number of eigenvalues represented by a particular line. For this, the numerical data had to be examined by hand. Thus all lines represent one eigenvalue, except where explicitly labeled otherwise. Not all multiplicities in higher energy eigenstates are labeled.

Finally, the numerical algorithm for calculating eigenvalues did not always converge, and thus provided possibly erroneous results for particular values of . In most cases these were isolated, and therefore can be safely ignored; however, there are a few larger regions in which the algorithm failed to converge. These also are indicated in the plots below, by regions surrounded by dashed lines.

First, we consider the tiling **hex7**:

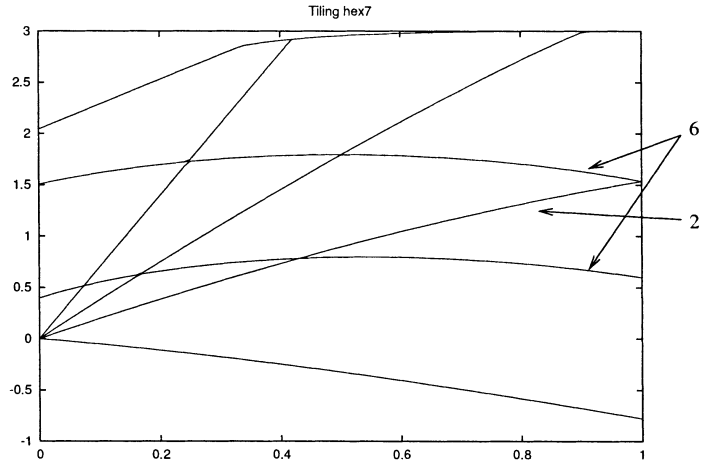

The initial 5-dimensional ground state splits into four seperate states, with the second lowest of these being doubly degenerate, as indicated in the figure.

For **hex9**:

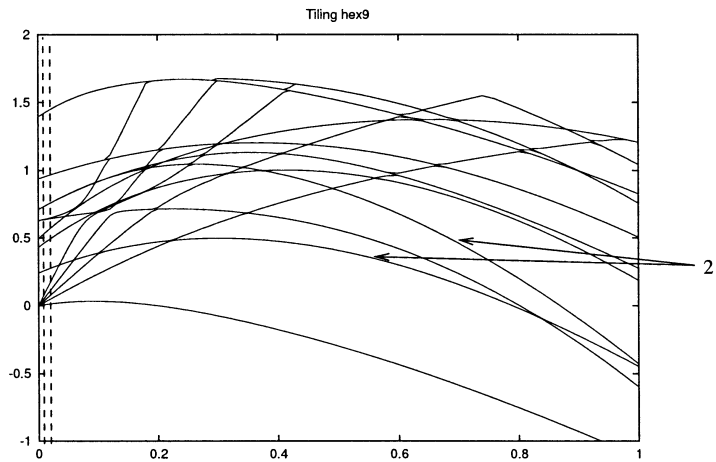

As in **hex7**, the initial 5-dimensional ground state splits,although this time into five seperate states. Again only the lowest of these remains in the ground state as ε grows. (The indicated region of failed convergence is for $0.01 \leq \varepsilon \leq 0.02$.)

The spikes appearing in the plot below for **hex12a** are due to failed convergence of the algorithm; however, the algorithm only failed to converge for some of the higher energy eigenvalues. Therefore it is not indicated as a potential error region(as we are interested mainly in the lowest energy eigenstates).

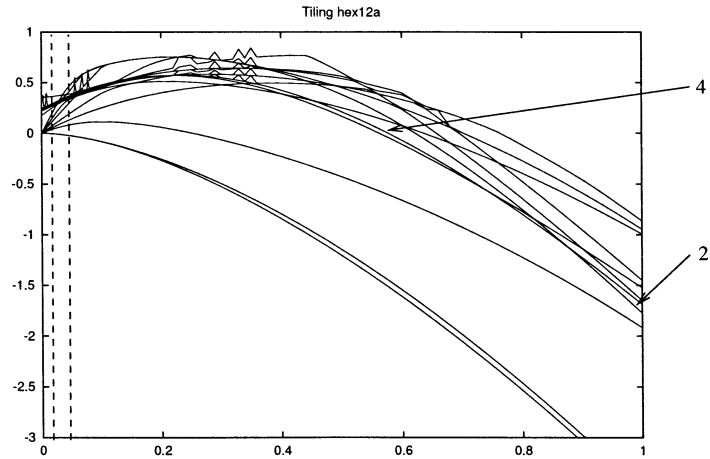

Though it is difficult to see, six of the eight original ground state eigenvalues split in the initial region ($\varepsilon \leq 0.01$), with the two lowest staying together. Then the two lower ones fracture apart around $\varepsilon \approx 0.2$. Recall also that this tiling has two lonely configurations, although it is not possible to tell from these data how the lonely configurations might correspond to the eigenvalue trajectories.

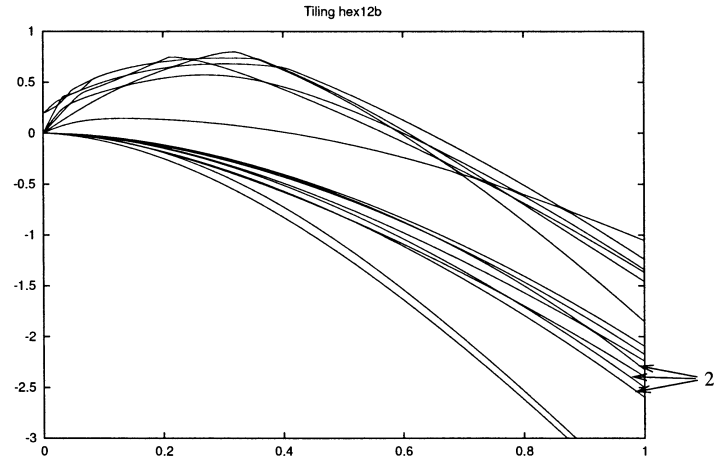

The tiling **hex12b** has 12 lonely configurations; thus the plot below indicates more clearly how these relate to the other ground states.

From this plot it is almost certain that the 12 eigenvalues that initially are negative correspond to the 12 lonely configurations of the tiling.

Neither of the 15-tilings admit any lonely configurations. The plot for **hex15a** follows: All seven of the ground state eigenvalues split, with only one of them remaining in the ground state for $\varepsilon > 0.3$. The area of failed convergence is $0.09 \leq \varepsilon \leq 0.3$.

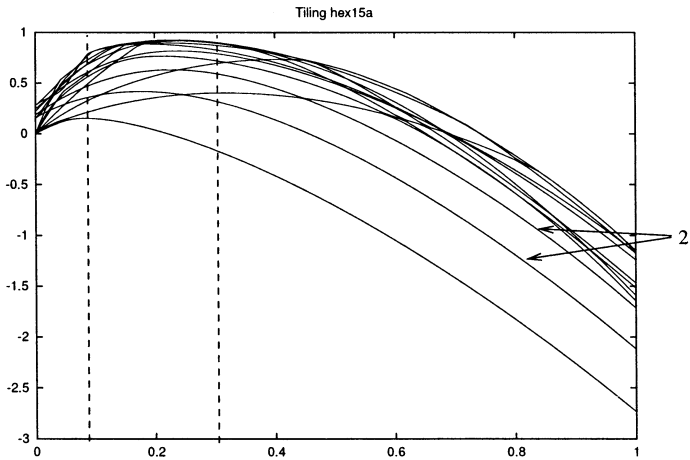

For **hex15b**:

Again, all eight of the ground state eigenvalues split, and only one remains in the ground state beyond $\varepsilon \approx 0.25$. Though it is difficult to see on

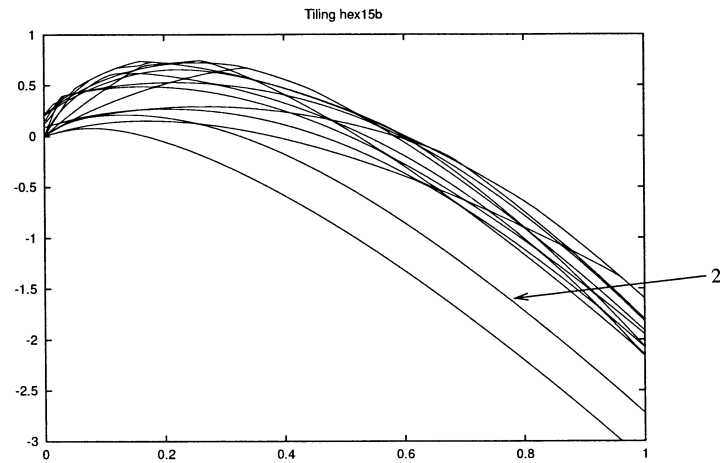

the graph, the line labelled with multiplicity two is not originally in the ground state when $\varepsilon = 0$.

The plots for the 16- and 18-tilings are restricted to $0 \leq \varepsilon \leq 0.5$, as they are computationally expensive, and the trajectories of the eigenvalues seem clear. For **hex16**, the general trend regarding the lonely configurations

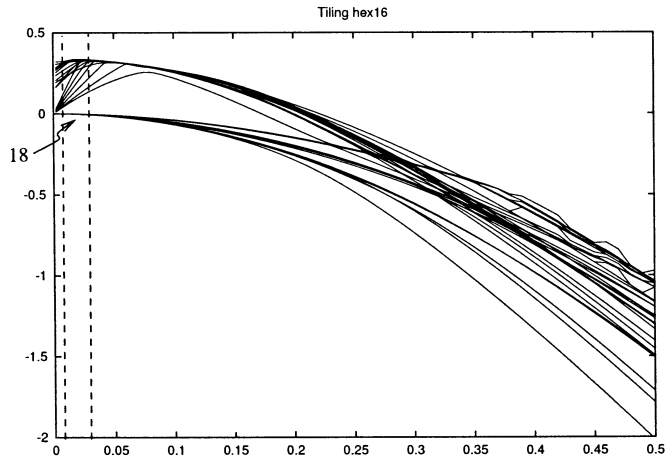

remains consistent; interestingly, the other pattern seems to hold, although in a slighty more dramatic fashion. The eigenvalue corresponding to the lowest non-lonely configuration first rises sharply, then hits a peak around $\varepsilon \approx 0.07$, and then begins to decline, remaining below the other non-ground state trajectories until $\varepsilon \approx 0.3$. Although it appears to nearly "merge" with the non-ground state trajectories at $\varepsilon \approx 0.07$, manual inspection of the data indicates that in fact it remains clearly distinct.

The plot for **hex18a** follows: Again, eight of the ground state eigenvalues separate distinctly from the rest, almost certainly corresponding to the eight lonely configurations. And again, only one of the original ground state eigenvalues remains consistently low, with only the lonely configurations below it.

Finally, the tiling **hex18b** demonstrates the same pattern. The fourteen dimensions of the ground state corresponding to the lonely configurations initially split off and remain consistently lower than the rest, and one of the remaining ground state configurations seperates and remains in a lower state than everything but the lonely configurations.

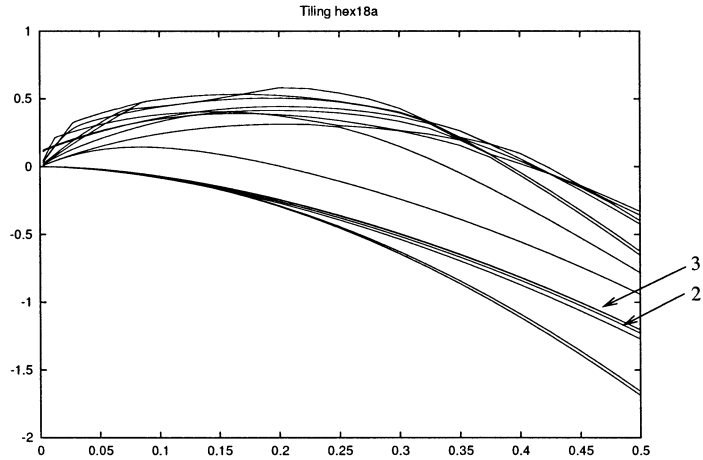

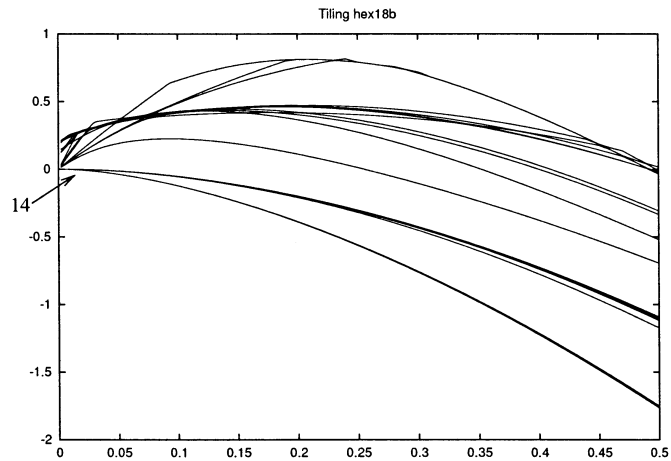

5. DISCUSSION

- The expected dimension of the perturbed ground states for $l = 3$ is 4. Ignoring the lonely states represented by lonely configurations, the un-perturbed ground states form a space of dimensions 5,6,7, and 8. There was consistently one branch of the original ground state which split off from the rest and remained in the ground state. This is an indication of what we're looking for, but not coming out in full because of the small size of the tilings. Only two of the tilings (**hex18a** and **hex15b**) even admitted four parallel essential circles,

and even then only in one direction (e.g., just horizontally, not vertically).

- The problem of lonely configurations might be only a technical issue. If they exist, then they will certainly be the lowest energy eigenstates under perturbation. It seems likely the lonely configurations are "isolated" in the sense that, if the system is prepared in some suitable initial state, then it will never "jump" to one of these lonely configurations, but rather slowly fall into one of the ground state configurations corresponding to a large g-isotopy class.

  Given a triangulation of a surface, we can always subdivide the surface to get a new triangulation without any lonely configurations. The problem with this solution is that we have to introduce much more tilings into the problem, which is definitely computationally intractable.
- Recall that in the Chern–Simons theory, the Lagrange has a quadratic term $A \wedge \mathrm{d}A$ and a cubic term $\frac{2}{3}A^3$. If the cubic term is treated as a perturbation, then it is a 3rd order term. So maybe a better perturbation should be at least of order 3.

6. GAP IN THERMODYNAMIC LIMIT

Inspecting the data, we observe that the lowest eigenvalue and the next one has a clear gap. The question is whether this gap persists if we go to finer and finer tilings. Further study will be carried out.

ACKNOWLEDGMENTS

Research of Z.W. is supported by NSF grant CISE/EIA-0130388 and Army Research Office.

APPENDIX A. IMPLEMENTATION DETAILS

The open source numerical package ARPACK[7] was the main tool used for finding the ground state and low energy eigenstates of the Hamiltonians. ARPACK is optimized to find certain eigenvalues(e.g., those of the lowest magnitude) of large, sparse, symmetric,real-valued matrices, and is thus well-suited to our problem. A set of templates interface the Fortran ARPACK code with C++ was used; all of the custom code for our problem was done in C++.

For a fixed Hamiltonian, let n denote the number of tiles in the associated tiling; thus, the dimension of the associated H_0 will be 2^n.

The sparseness of the Hamiltonian can be analyzed by determining the number of type-g and type-h spin configuration pairs associated with a given tiling. It is relatively easy to determine that, for a hexagonal tiling with n tiles, there are $62n2^{n-7}$ type-g and type-h spin-configuration pairs, and therefore twice as many nonzero entries in the matrix (total of 2^{2n} entries). While this is relatively sparse, it is still inefficient to compute and store the matrix explicitly (using the "Compressed Sparse Column" format of ARPACK), as the number of required entries grows exponentially. Computing the matrix-vector product "on the fly" is much more efficient. This algorithm is given in pseudo-code below:

```
//helper method
Bar(c,t); // returns the configuration c with tile t flipped

H0_product(v[],w[]) { //computes H_0|v>=|w>
   w=(0,0,...,0);
   foreach tile t {
      foreach type-h cfg (c,t) {
         w[c]+=v[c]-(1/d)v[Bar(c,t)];
         w[Bar(c,t)]+=(-1/d)v[c]+(1/d^2)v[Bar(c,t)];
      }
      foreach type-g cfg (c,t) {
         w[c]+=v[c]-v[Bar(c,t)];
         w[Bar(c,t)]+=v[Bar(c,t)]-v[c];
      }
   }
}

He_product(v[],w[]) { //computes H_\epsilon|v>=|w>
   H0_product(v,w); //first do the H_0 part
   foreach cfg c {  // then add the perturbation
      foreach tile t {
         w[Bar(c,t)]+=v[c]*epsilon;
      }
   }
}
```

The algorithm requires $O(n2^{n-7})$ floating-point operations (flops) per matrix-vector product; compared with the standard product using the CSC-stored format, which would required $O(2^{2n+1})$ flops. The configurations that

are type-h and type-g, with respect to each tile, are calculated beforehand and stored. Therefore, the storage requirements are $62n2^{n-7}$ integers, which is on the same order as the storage requirements for storing the matrix in CSC format.

An attempt was made to parallelize the code, as there is a parallel version of ARPACK available. However, the code only allows for parallelization of the matrix-vector product; and the lag in communicating results across nodes would be significantly more costly than the actual computation of the results. Thus the only way in which parallelization might be useful is if the actual eigenvalue algorithm was parallelized. This approach was not pursued.

REFERENCES

1. M. Freedman, *Commun. Math. Phys.* **234**, 129–183 (2003).
2. M. Freedman, A. Kiatev, M. Larsen, and Z. Wang, *Bull. AMS* **40**, 31–38 (2003).
3. M. Freedman, M. Larsen, and Z. Wang, *Commun. Math. Phys.* **227**, 605–622 (2002).
4. M. Freedman, M. Larsen, and Z. Wang, *Commun. Math. Phys.* **228**, 177–199 (2002).
5. M. Freedman, C. Nayak, K. Walker, and Z. Wang, *Picture TQFTs*, in preparation.
6. L. Kaufmann and S. Lins, *Temperley-Lieb Recoupling Theory and Invariants of 3-manfiolds*, Ann. Math. Studies, Vol 134, Princeton Univ. Press.
7. ARPACK.