Journal of Statistical Physics, Vol. 96, Nos. 5/6, 1999

Diamond Ice

Kari Eloranta[1]

Received December 11, 1997; final April 28, 1999

The bounded version of the ice model of statistical mechanics is studied. We consider it in a diamond domain on the $\mathbf{Z}^2$-lattice. The configurations sharing a boundary configuration are shown to be connected under simple loop perturbations. This enables an efficient generation of the configurations with a probabilistic cellular automaton. The fill-in from the boundary is critically dependent on the values of the height function along the boundary. We characterize the phenomena at the extrema of this function as well as in some highly nontrivial cases where results analogous to and more complex than the Arctic Circle Theorem for dominoes hold.

KEY WORDS: Ice model; six-vertex model; symbolic dynamics; cellular automaton; tiling; domino.

INTRODUCTION

The subject of this paper is the ice or six-vertex model. It is a classical model in Statistical Mechanics and has been studied extensively. Essentially all the studies have concentrated on the infinite model on some lattice or the finite one with toral boundary condition. A summary of this work up to -82 can be found in ref. 1.

Our study concentrates on the planar version of the model defined on a finite domain with boundary. The square lattice is used partly because we wish to make comparisons to other models having similar underlying discrete structure. For the same reason we treat only diamond-shaped domains here. Our methods, which utilize critically the two-dimensionality, generalize to some other shapes and a few other planar lattices but are unlikely to have higher dimensional counterparts.

[1] Institute of Mathematics, Helsinki University of Technology, FIN-02015 HUT, Finland; e-mail: kve@locus.hut.fi.

0022-4715/99/0900-1091$16.00/0

The main objective is to study how the ensemble of configurations in the interior of the diamond depends on the boundary data, i.e., the fixed configuration on the edge of the diamond. The first question is: when is this set nonempty, i.e., when does the boundary configuration admit a fill-in? Immediately following this is the question of how many such configurations are there? This leads naturally to entropy considerations. Another line of investigation is how the legal configurations relate to each other —what are the allowed perturbations/transformations between them? How can the entire set of configurations with a given boundary be generated? Based on a connectivity result we present a simple algorithm that enables us to simulate the generic element in the configuration space. Utilizing the height function the investigation proceeds to unveil an interesting coexistence of qualitatively different subdomains inside the diamond. This phenomenon—the Arctic Circle separating the domains supporting the highly ordered and the disordered configurations—was first observed in the context of dimers/dominoes.[3]

Aside from having some relevance in understanding the geometry of long range order in ice models, our results have some bearing on tilings and symbolic dynamics. In a couple of seminal papers Conway and Lagarias[4] and Thurston[11] investigated the problem of tiling a given planar domain with polyominoes. On the basis of this and subsequent other work it seems feasible that there could be a unifying theory for planar tilings and two-dimensional symbolic dynamical systems. The latter naturally include many Statistical Mechanics models, e.g., the dimers, ice model, eight-vertex model, color models and yet others on various lattices. The program is now on its way and this paper is a small contribution.

1. BASICS

The ice model can be defined in any dimension and for all regular lattices with even vertex degree. Because of the tiling connection indicated in the introduction we consider only the planar case here. Our methods rely critically on the two-dimensionality but can be applied to, e.g., triangular lattice.

Consider the square lattice in two dimensions, $\mathbf{Z}^2$. Every lattice site has four nearest neighbors. The vertex models do not have any spin variables associated to the lattice points unlike most lattice models. Instead the lattice variables are the orientations of the arrows between nearest neighbor sites.

Definition 1.1. A vertex configuration at a lattice site in $\mathbf{Z}^2$ is *legal* for the *six-vertex rule* if there are exactly two incoming arrows and

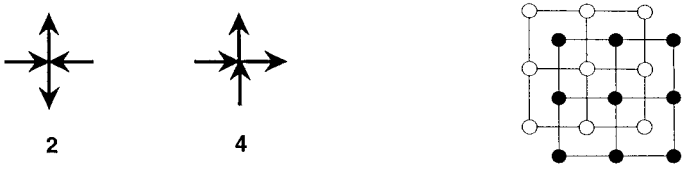

Fig. 1. Vertex configurations. Square lattice and its dual.

two outgoing ones. A configuration is legal if it has an allowed vertex configuration at every lattice site.

The allowed vertex configurations are illustrated in Fig. 1, left. The numbers below indicate the multiplicity of the arrangement. There are six possibilities, hence the name of the model.

One can view the six-vertex rule as incompressibility of a fluid or expressing a conservation law for some other system. Its main physical importance stems however from the modelling of water molecule interaction at low temperature. The key physical quantity, residual disorder at zero temperature, was exactly determined by Lieb for the infinite two-dimensional model.[9]

In this paper we concentrate on the finite case. The model is defined on a bounded domain for which we specify the boundary arrows and then study the possible extensions to the interior. For simple boundaries there can be many extensions. For a complicated domain shape it may be impossible to find boundary conditions other than those for highly ordered states.

For simplicity we consider a diamond domain. The reason for this is two-fold. Firstly we want to relate the ice model to certain other models, in particular to the domino and eight-vertex models, for which the diamond domain is important and has been analyzed. Secondly we will see that by specifying the boundary condition on a diamond we can actually enforce in a few critical cases the boundary condition on an inscribed square and thereby get the square domain results as well. All results of this paper generalize to rectangles rotated 45 degrees but since this adds little insight we record the results for diamonds only.

An *N-diamond* is a subset of $\mathbf{Z}^2$ which has N arrows along each of its four diagonal sides, $N \geqslant 2$, even. The total number of arrows is hence N^2. One can think it to be made of $N^2/2 - N + 1$ unit squares each of which has four arrows as sides and the neighboring squares sharing an arrow. Such a domain contains $N^2/2 + N$ lattice sites. *The boundary configuration* of the N-diamond, which consists of $4N - 4$ arrows, is fixed. We use the notation a to refer to the boundary (sites or configuration-arrows depending on the context).

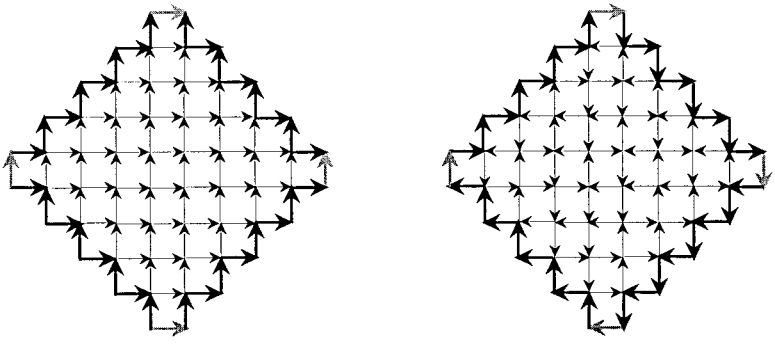

Fig. 2. Frozen and temperate 10-diamond configurations.

In Fig. 2 the fixed boundary arrows are rendered bold and the interior arrows light. We distinguish the corner arrows by a grey shade since they do not influence the fill-in in any fashion. They are included since it will be later more natural to refer to the boundary loop instead of the boundary pieces.

The figure exhibits two extreme configuration types. The highly ordered state on the left (the NE flow) can be thought as "frozen" or to correspond to low-temperature regime in Statistical Mechanics models whereas the one on the right, the "temperate" one, represents a high temperature/disordered regime.

To arrive at the first result characterizing the fill-in it is necessary to invoke the concept of dual lattice. For $\mathbf{Z}^2$ this is particularly easy—it is self dual, which means that its dual lattice is the same lattice, only shifted (formally we write $(\mathbf{Z}^2 + \frac{1}{2})^2$). It is illustrated in Fig. 1, right.

Around each lattice point of $\mathbf{Z}^2$ we can draw a unit square with edges along the lattice lines of the dual. If this is viewed as a clockwise oriented unit loop around the lattice point the six-vertex rule simply says that the flux across the loop has to vanish. If we consider a set of adjacent lattice points and the arrows at them we can form the oriented *boundary loop* for the set by adding up the unit loops. This is the minimal lattice path on the dual lattice that encircles the set. Since an arrow pointing out from a unit loop points either out of the new boundary loop or into an other unit loop we will see that the flux across the boundary loop has to vanish as well. Any legal configuration must therefore have this property and we have arrived to

Lemma 1.2. A necessary but not sufficient condition for a finite domain to fill-in is that the flux across its boundary vanishes.

Remark. With the height function of Section 3. One can resolve the sufficiency.

2. CYCLES

We now proceed to study how the configurations are related to each other. This reveals the topological structure of the set of N-diamond configurations and also leads to a method to generate all legal configurations.

The first observation is that at any vertex we can simultaneously flip the directions of one incoming and one outgoing arrow. This is illustrated on the left in Fig. 3. For any legal vertex configuration we fix one incoming and one outgoing arrow, say a and b (not necessarily oriented as shown), and then reverse each of the non-bold arrows. This yields another legal-vertex configuration at that lattice point. The procedure holds for all six vertex configurations.

In the flipping at a single vertex configuration we violate the rule on two of its neighbors. But if we reverse the arrows along a directed arrow loop (or in the infinite model along a path from infinity to infinity) in the resulting configuration all vertex configurations are again legal.

Let us call the simplest such action, the reversal of the arrows in a directed 1-cycle an *elementary move*.

Call the subset of legal N-diamond configurations with the same boundary configuration a coset. Some sets of configurations are connected under reversal of directed 1-cycles (to simplify wording from now on 1-cycles will always be directed). The natural question then is to characterize the configurations that can by constructed from a given configuration using a finite sequence of elementary moves. Note that in the case of a bounded domain with a fixed boundary a cycle reversal can never reach the other cosets because the path to be reversed cannot contain any boundary arrows.

1-cycles are actually not as rare as one might first believe. To this end we note a simple result.

Lemma 2.1. In a legal ice-configuration inside every directed cycle there is a 1-cycle.

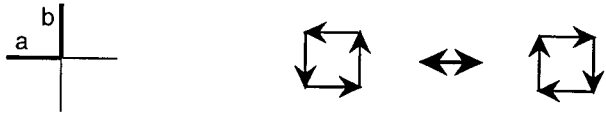

Fig. 3. Vertex perturbation and elementary move on a directed 1-cycle.

Proof. Consider the subset consisting of the directed cycle, the arrows inside the domain it defines and those with end/head at the lattice points on the cycle (the "crossing" arrows). By Lemma 1.2 the flux across the minimal lattice path outside this cycle is zero. Hence at some lattice point on the cycle there is an arrow pointing into the domain. We follow the directed path that this arrow initiates choosing at every vertex the next arrow at random among the two available ones. Since the cycle is finite so is the domain and eventually we either arrive to the boundary or self-intersect. Either way a new domain is formed which encloses fewer lattice points than the original and which has a directed cycle boundary. By induction we conclude the statement. ▌

Remark. To illustrate the Lemma we note that the boundary of the temperate configuration in Fig. 2, right is a cycle and the configuration indeed has a number of 1-cycles. The *frozen* configuration next to it on the other hand has no cycles of any size; this is its defining property.

A 1-cycle is *off-boundary* if it contains no boundary arrows.

Proposition 2.2. A configuration in a N-diamond is the unique fill-in of the boundary configuration iff it does not have off-boundary 1-cycles.

Proof. If a fill-in is unique then clearly there cannot be any off-boundary 1-cycles since reversal of such immediately leads to another configuration with the same boundary.

So suppose that there are two distinct configurations, A and B, which are fill-ins of the same boundary. Pick a pair of neighboring lattice points (x_0, x_1) between which the arrows differ. Say in A this arrow is heading $x_0 \to x_1$. At vertex x_1 there are two other, outgoing arrows in A and in B two ingoing arrows. We can therefore pick a pair (x_1, x_2), $x_2 \neq x_0$ so that again the two arrows are opposite and now form a directed path of length two on each of the two configurations. Continue in this manner and note that the directed path cannot include boundary arrows as they are equal in the two configurations. Since there is a finite number of lattice points, the path will eventually self-intersect. Then we will have two identical loops with opposite orientations in A and B. By Lemma 2.1 they have 1-cycles, which by construction are off-boundary. ▌

The proposition hints that elementary moves might exhaust the set of perturbations by generating all the possible configurations. This is indeed the case.

Theorem 2.3. The set of all configurations on a diamond with a common boundary configuration is connected under elementary moves.

Proof. Given two distinct diamond configurations with a common boundary configuration we will transform one of them to the other through a finite sequence of 1-cycle reversals. We will check the configurations lexicographically as shown in Fig. 4. NW of the bold broken line B all arrows agree on the two configurations. But they disagree at the arrow/edge a_d marked by a dot. From a_d we generate the directed disagreement loop exactly as in the proof of Proposition 2.2. This path can contain arrows from neither the boundaries nor the agreement area. By Lemma 2.1 it will enclose at least one 1-cycle. On one of the configurations the loop is oriented clockwise. Call it C and denote the 1-cycles inside it by C_i.

For each corner of a 1-cycle we can associate a contact sector. By this we mean a closed lattice quadrant rooted at the vertex; first quadrant for the NE corner, second for the NW and so on. The four contact sectors of a 1-cycle are clearly disjoint. One can easily see that by the six-vertex rule there are in a contact sector rooted at a 1-loop both an incoming and an outgoing semi-infinite path through the root vertex.

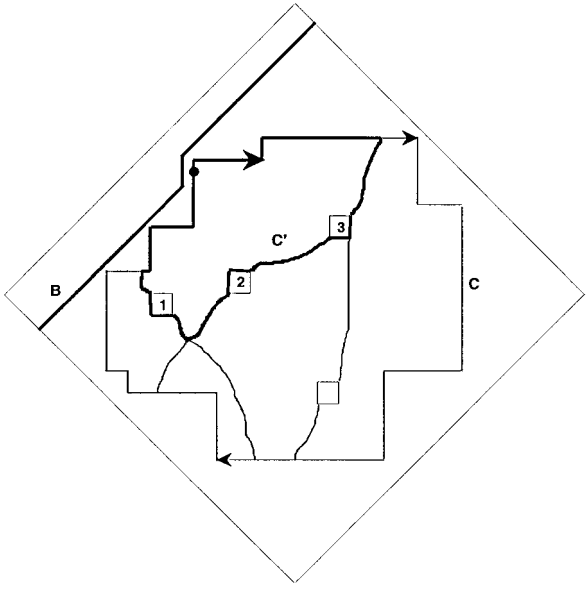

Fig. 4. Proof constructions.

We associate to each C_i two disjoint contact sectors such that neither contains a_d. Then connect each C_i to C with directed paths in the respective contact sectors, incoming path in one sector, outgoing in the other. By choosing the orientations of these paths properly each C_i ends up being on the boundary of a clockwise oriented loop containing the mismatched edge a_d. Using these loops form the minimal directed loop C' that contains the disagreement arrow and some of the 1-cycles on its boundary. No 1-cycles will remain strictly in its interior. Note that there must remain at least one 1-cycle on the inside of the loop C', again by Lemma 2.1.

By reversing the 1-cycles on the inside of the boundary of C' we obtain a new, smaller directed loop C'' inside C' that contains the mismatched arrow but none of the C_i's inside. Since it must by Lemma 2.1 contain 1-cycles the only possible locations for them are next to C_i's. Hence they are grazing the inside of the boundary C''. Continuing this strictly monotone shrinking of the directed loop containing the disagreement site we arrive after a finite number of steps to the situations where there is a 1-cycle containing the edge a_d. After reversing it the curve B moves to its next lexicographic location. The process can be continued until the entire configuration has been checked and corrected of disagreements. ∎

Remarks. 1. The theorem indicates that the simplest of actions, the elementary move, generates all the configurations it possibly can. The diamond configuration space partitions into cosets inside which the elementary moves are an irreducible action. However unlike in the context of the related eight-vertex model (definition as 1.1 but the number of incoming arrows at a lattice point is 0, 2, or 4, rest are outgoing), where similar irreducibility holds and the cosets are all of exactly the same size,[6] here they are of very different size. Furthermore the internal structure of ice configurations is generically highly nontrivial in some of the cosets unlike in the eight-vertex model.

2. Similar result can be proved for the ice model on triangular and Kagomé lattice.[7] Interestingly when there are several different elementary actions, they are all needed for the connectivity.

3. Finally it may be of interest to note the similarity to a certain two-dimensional dynamical system treated in ref. 5. In this model the number of copies of each symbol in a square neighborhood is fixed—an exact conservation law like our arrow condition. The perturbations work out somewhat similarly. Infinite periodic sequences and their "slide-deformations" play the same roles as directed cycles and their reversals in ice. If the model is defined on a torus this means then rotating periodic sequences of symbols

in finite loops. They generate most of the configurations space but surprisingly not all, i.e., analog to Theorem 2.3 does not quite hold there.

3. HEIGHT

The notion of height is closely related to the concept of flux introduced in the first section. For the ice-model the height was first introduced using a graphical representation and without connection to flux in a paper by van Beijeren.[2] What makes it an important notion is its utility in distinguishing different boundary conditions and the fact that it can be defined for various lattice models. Among them are ice on various lattices, dominoes/dimers, certain color models etc. This connection will be discussed in later sections.

Definition 3.1. *The height function*, h, is an integer-valued function defined on the vertices of the dual lattice. Moving from one such lattice point to its nearest neighbor its value increases by one if the configuration arrow to be crossed points to the left and decreases by one if the arrow points to the right. This determines the value of the function, the *height*, everywhere on the configuration upto an additive constant. Giving the height at any one dual lattice point determines the height function uniquely. Height at the end of a path subtracted by the height at the beginning divided by the number of arrows crossed is the *tilt* of the path.

Remarks. 1. From this on we fix the height to be zero at the leftmost dual lattice point inside the diamond.

2. Note that the flux across a section in the dual lattice is just height computed along with an agreement on what is the positive direction. By the ice rule the flux across any loop vanishes. Hence the definition of height is consistent, i.e., when we loop back to the original dual lattice point from which we started the height computation we recover the initial value.

3. Height is a complete description of the configuration, i.e., determines it uniquely. Its restriction to the boundary of the configuration, the *boundary height* is frequently useful. In Fig. 5 the underlined entries are the boundary height for the given 4-diamond. The values of boundary height inside the diamond are called the *inside boundary height*.

4. Tilt is a number between and including ± 1. In the example in Fig. 4 we have recorded its value over each of the four edges as they are traced to the counterclockwise direction (in bold).

Definition 3.2. The height of a configuration is called *extremal* if the boundary height uniquely defines the entire height.

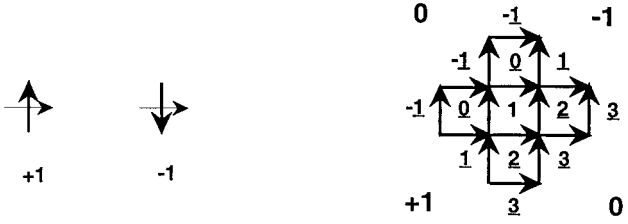

Fig. 5. Height. Configuration arrows are bold, height path are light.

Remarks. 1. The example on the right of Fig. 5 illustrates extremal height.

2. Height can be viewed as a (discrete) surface above the configuration. The discrete partial derivative assuming its extremal value to some direction over the entire configuration corresponds to the extremal height. If we consider a N-diamond but scale it and h by $1/N$ and take the scaling limit $N \to \infty$, we obtain a Lipschitz surface. The height is extremal iff its partial derivative (infinitesimal tilt) is extremal to some direction over the entire unit diamond.

If the boundary height determines the height uniquely then the boundary configuration determines the interior uniquely. This is exactly the case of the frozen configurations like the NE-flow in Fig. 2. One can quickly discover the simplest such configurations. Let the tilt along the entire edge be either constant ± 1 or the height be constant on the inside boundary of the edge. Assign the symbol ± 1 or 0 to the edge accordingly and record the four values in a clockwise run around the boundary. The signature $(1, 0, -1, 0)$ corresponds to the NE-flow, $(1, -1, 0, 0)$ to a L-flow and $(1, -1, 1, -1)$ to a X-flow, where the name describes the geometry of the vector field. Cyclically permuting these signatures (rotating the configurations) results in the 10 most basic frozen configurations.

To analyze the boundary dependency it is useful to distinguish *switch blocks* and *neutral boundary blocks*. These consist of two adjacent boundary arrows, in the former case pointing to the same lattice point and in the latter pointing to different lattice points. Figure 6a illustrates the switch blocks: arrow pairs at boundary height values 1, 3, 5, and 7 above the switch point s_1 are switch blocks contributing $+2$ to the height each and the two below s_1 contribute -2 each.

The total frozen area next to the boundary is just the sum of the triangular areas next to the extremal boundary pieces. This can be quantified as follows.

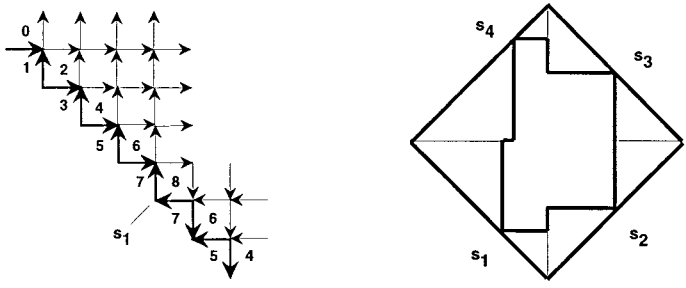

Fig. 6. Extremal blocks, switch point. Frozen periphery.

Lemma 3.3. In a N-diamond let A_f be the number of interior arrows fixed by the boundary pieces of extremal height. Let

$$A_{\pm} = \tfrac{1}{4} \max_{\{d_j\}} \sum_j \left[(h(d_{j+1}) - h(d_j))^2 + (1 \pm 1)\, |h(d_{j+1}) - h(d_j)| \right] \qquad (1)$$

where $\{d_j\}$ is the set of dual lattice sites inside the diamond where the boundary height is computed. Then

$$A_- \leqslant A_f \leqslant A_+ \qquad \text{and} \qquad A_+ - A_- \leqslant 2N$$

Proof. $n/2$ adjacent switch blocks of the same sign along a N-diamond edge (from d_j to d_{j+1}) results in boundary height change $|h(d_{j+1}) - h(d_j)| = n$ ($n \leqslant N$, n even). A little algebra shows that these arrows uniquely determine $n^2/4 + n/2$ arrows in the interior (the triangular arrangements in Fig. 6, right). Hence the upper bound for A_f, A_+, is obtained by summing these contributions over maximally long boundary pieces each of extreme tilt ± 1.

A triangle (again with $n/2$ adjacent switch blocks of the same sign along a diamond edge) sharing a corner with the diamond may share at most n arrows with a similar triangle on the other edge. Removing the possible duplicate arrows from the count and summing as above gives A_-.

The last inequality is immediate. Equality is attained when each edge of the N-diamond consists of $N/2$ switch blocks of same sign. ∎

The diametrically opposite case to the extremal height is the case where each of the diamond edges is a directed path of arrows (ignore the four corner arrows). Call such an arrangement, a *quasicycle*. The cycle boundary of Fig. 2, right is a special case of this. For a quasicycle boundary the inside boundary height is identically zero since the boundary is made of neutral blocks of the same heading (outside the diamond the boundary height is either ± 1). Note that each of the arrow paths along the edges can

be reversed without affecting the fill-in. Since the corner arrows can be chosen at will this means that any quasicycle boundary can be transformed to a cycle boundary without affecting the interior of the configuration. Hence by Lemma 2.1 there is non-uniqueness in the fill-in and indeed this non-uniqueness is maximal. The quasicycle case corresponds to the *temperate* configurations where the boundary forces no interior arrows, i.e., $A_f = 0$.

It is also useful to define *average height*, $H(d) = \mathbf{E}(h)(d)$, where the average is taken over all the possible fill-ins from the boundary with uniform weight. In the frozen case obviously $H(d) \equiv h(d)$.

In the temperate case the average height vanishes inside the diamond. The argument for this is as follows. Consider a dual lattice site d in the inside of the diamond. Take a legal fill-in arrow configuration, call it a_+. Compute the height $h_{a_+}(d)$ by starting at an inside boundary point (there the height is zero since the boundary is a quasicycle. The configuration where all the arrows have been reversed is another legal diamond configuration with a quasicycle boundary. Reversing this boundary arrows gives us a configuration a_- with the same boundary as a_+ but all interior arrows reversed. Following the same dual lattice path that we used for computing $h_{a_+}(d)$ gives now $h_{a_-}(d) = -h_{a_+}(d)$ since all the off-boundary arrows are reversed and height on the inside boundary is zero. Hence for each fill-in there is exactly one "reverse" fill-in and the equally weighted average of the height over any dual-lattice site must therefore equal to zero.

We say that a boundary configuration is non-trivial if it contains segments of constant extremal height which are of non-trivial length (different from 1 and N). In this case the boundary forces some but not all of the interior. Since the configuration is not frozen in the remaining part of the domain there must be 1-cycles.

To summarize these basic findings on height we formulate

Proposition 3.4. In a diamond configuration

(i) h_a is extremal iff a is frozen.

(ii) $h_{a|\partial} = \{(0, 1)\}^*$ or $\{(0, -1)\}^*$ iff a is temperate.

(iii) Temperate and frozen subdomains coexist in a iff $a|\partial$ is non-trivial.

Lieb computed the residual entropy of the infinite ice-model on the $\mathbf{Z}^2$ lattice.[9] This quantity, the average uncertainty per arrow, is the same as *topological entropy* which can be computed as

$$h_{\text{top}} = \lim_{N \to \infty} \frac{1}{N^2} \ln\{\text{number of } N\text{-diamond configurations}\}$$

They key here is that there is no boundary condition on the diamond. From ref. 9 we get that $h_{\text{top}} = \frac{3}{4} \ln \frac{4}{3}$.

Using this we immediately get an asymptotic upper bound for the number of legal configurations in the case of a boundary condition. By formula (1) the fraction of the arrows in a N-diamond that is fixed by the boundary is $A_f(N)/N^2$. In the frozen case $A_f(N) = N^2 + O(N)$ and in the case of quasicycle boundary $A_f(N) \equiv 0$. Inside the diamond but off the frozen area the entropy per site obviously cannot exceed that of the free model. Hence in the scaling limit (lattice spacing equals to $1/N$, $N \to \infty$, so the limiting domain will be the unit diamond $|x| + |y| \leqslant 1/2$) we obtain

Proposition 3.5. For the diamond ice the average topological entropy over the unit diamond for a given boundary condition is bounded from above by

$$\left(1 - \lim_{N \to \infty} \frac{A_f(N)}{N^2}\right) \frac{3}{4} \ln \frac{4}{3} \tag{2}$$

whenever the limit exists. In particular in the frozen case entropy vanishes whereas in the temperate case it is bounded by Lieb's number.

Remark. Although we do not have a lower bound for the entropy already from the propositions and the analysis above it is plain that the boundaries fill-in in very different ways. This is in notable contrast with the eight-vertex model, where all legal boundaries fill-in in exactly the same number of ways.[6]

4. CONFIGURATIONS

4.1. Generating Algorithm

We now show a way of computing the configurations satisfying a given boundary condition. The reason is two-fold. Firstly the underlying principle is simple yet interesting and utilizes the results derived upto now. Secondly to analyze the case of nontrivial boundary we need an efficient way of computing the configurations.

The method is based on the elementary moves and Theorem 2.3. Consider the set of all N-diamond configurations for a given boundary condition. Suppose that we are given one of them. From that we form the first *even configuration* in the following way. Since every (not necessarily directed) 1-cycle consists of four arrows there are 16 different ones. Form the symbol

set $S = \{0,..., 15\}$ from them. The arrow configuration is completely determined if we specify the symbols at every other 1-cycle site, i.e., on a checkerboard pattern which includes the boundary arrows. Take the boundary cycles to be dark/even and denote the configuration of all the dark/even symbols by $C^{(e)}$.

Consider now four adjoining dark 1-cycles in a cross formation. The local *rule* is simply to read off from them the fifth (light) 1-cycle at the center. This is then *reversed with probability p if it is a directed* 1*-cycle*. If the reversal takes place the adjoining 1-cycles are updated as well, since one arrow in each of them was reversed. This local operation performed at every neighborhood centered at a light 1-cycle gives the new *odd configuration* $C^{(o)}$, the symbols on light/odd squares. The local rule immediately gives the global map, the *probabilistic cellular automaton* $F_p: C^{(e)} \to C^{(o)}$. The map $C^{(o)} \to C^{(e)}$ that updates the even symbol array works essentially the same way. There we have to augment the image with the symbols on the dark boundary squares. They can never be reversed since the boundary arrows are fixed.

Alternating the two maps generates the infinite forward orbit of even and odd configurations all of which correspond to legal configurations for the given boundary. If the initial configuration is frozen, there are no directed 1-cycles to reverse and the orbit is trivial. But other cases are less so. Note that by Theorem 2.3 the action of the cellular automaton is irreducible; every legal configuration associated to the given boundary condition can be reached. In fact the local updates are done independently and non-trivially, i.e., $0 < p < 1$ this orbit reaches every allowed configuration almost surely in finite time. The automaton relaxes a legal initial configuration to the equilibrium distribution on all legal configurations. This distribution is uniform (the measure of maximal entropy). At $p = 1/2$ the relaxation rate is maximal and the Markov Chain is expected to be rapidly mixing.[10]

As the rule only uses integer operations it can be implemented as a fast lookup table with a random mutation on two symbols (the directed 1-cycles). This is indeed a very efficient way of generating all the configurations associated with a given boundary conditions.

4.2. Simulation Results

For a legal boundary with non-trivial height the geometry and statistics in the interior of the configuration are highly nontrivial. We now investigate this regime for a few boundary types for which the scaling limit (2) exists.

Recall that the non-triviality of the boundary meant that the boundary configuration has segments of constant extremal tilt whose length is neither

minimal (1) nor maximal (N, the entire side of the N-diamond). The simplest case for this is the one where two opposite sides are split into two equally long extremal pieces. If the switch points s_i are the midpoints of the edges the height surface is geometrically a ridge roof. The unforced part of the interior, a square, is also indicated in Fig. 7a. Next to the diamond sides we have noted the tilts (when extremal, i.e., ± 1) and inside boundary heights (when constant 0). Note that the tilts along edges of the inscribed square are now forced to be maximal ± 1.

The middle and right plots are results from a simulation with a $F_{1/2}$-cellular automaton on a 102-diamond. The middle plot is at the iterate 2×10^4. The 16 different symbols are rendered in different grades of grey. In the interior of the "free" square there is a clear demarcation between the frozen and disordered domains. One can also discern ribbon-like structures on the boundary region. These are randomly fluctuating 1-dimensional defects in the ordered domain.

The rightmost plot from the same run represent the density of 1-cycle reversals in the configuration at the equilibrium. Here we have recorded the number of 1-cycle reversals at every site during the iterates $2 - 2.5 \times 10^4$ and converted this to grey level. Dark cells are the sites of most activity. In the forced corner triangles there is obviously no such activity hence they are rendered white.

The relaxations was performed to several different initial configurations with the given boundary configuration in diamonds of different size. The results were all essentially as above except that in a larger diamond the boundary of the disk obviously appeared smoother.

Alternatively we could split each of the four diamond sides into segments of equal length and maximal tilt so that the tilts it would alternate as in Fig. 8a. A simple variational argument shows that with this choice of switch points the "free" area in the center is maximal. Note

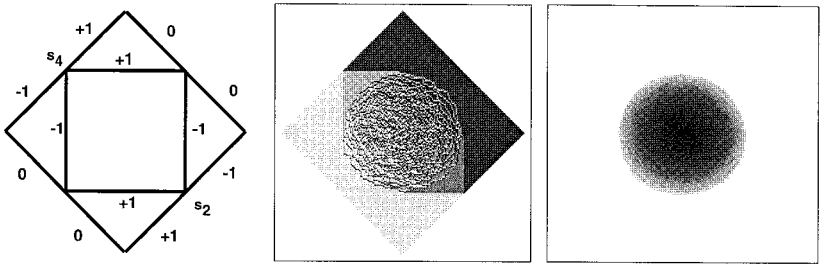

Fig. 7. Boundary tilts, configuration and the 1-cycle density.

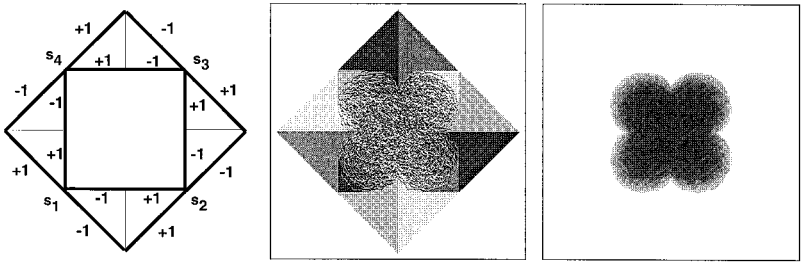

Fig. 8. Less restricted ice domain and the equilibrium 1-cycle density.

also that the forced area is exactly the same now that it was in the first case above. However the tilt along the boundary of the inscribed square is different as indicated.

It is fairly easy to construct some fill-ins for this boundary condition. There is obviously some freedom in doing this since this boundary belongs to type (iii) of Proposition 3.4, i.e., there will be 1-cycles in the interior. If one then lets the cellular automaton relax from any of these the result is generically as in the middle illustration.

The plots are from a 150-diamond. After some 10^4 iterates the configuration has reached an equilibrium. The middle plot is at the iterate 1.5×10^4. There is again a clear demarcation between the frozen and temperate domains which results in the flower-like boundary curve. The rightmost plot from the same run represent the density of 1-cycle reversals in the configuration at the equilibrium (we have recorded the number of 1-cycle reversals at every site during the iterates $1.1 - 1.5 \times 10^4$ and converted this to grey level).

Note that Proposition 3.5 gives for both of the cases above the same upper bound for the average topological entropy over the unit diamond (in the scaling limit), namely $3/8 \ln(4/3)$, half of that for the free model. It seems obvious that the latter case has higher entropy but proving this is difficult.

Boundary conditions with shorter pieces of extremal tilt force less area inside by Lemma 3.3. Hence one expects the boundary curve between frozen and disordered domains move closer to the "hard boundary" (the forced triangles in the examples above). This was verified in the simulations as well. In all the cases smooth curve pieces are separated from each other by cups as in the second case above.

Finally we point out that ice on triangular and Kagomé lattices have height functions. The coexistence of the frozen and the temperate domains is qualitatively similar to the cases presented here.[7]

5. COMPARISON TO DOMINOES

Suppose that one wishes to tile a finite planar domain with dominoes, i.e., 1×2 and 2×1 pieces. The success depends on trivial things like evenness of the area, but also on subtle things related to the shape of the domain. One can define a height function for dominoes analogously to Definition 3.1.[11] Note however that since there are no arrows, only the shape of the domain determines the boundary height. it turns out that in some domains like the square the boundary does not influence the interior of the tiling much. For example the orientation of the tile at any given interior site is quite uniform over the different tilings of the domain. This is due to the fact that the boundary height for this domain is essentially zero. In some domains with non-trivial boundary height there is however quite striking boundary dependency.

Figure 9a shows one such domain, an *Aztec diamond* (of order 6, i.e., 2×6 rows), and one of its domino tilings. In a string of papers Propp *et al.* investigated this set-up and found a particularly clean geometric result. The authors proved that generically the temperate subtiling (disordered domino tiling) is separated from the frozen one (brickwall tiling) by a curve which is a circle that grazes the diamond.[3] We illustrate this result in Fig. 9b, where the density of elementary moves in dominoes are plotted at every (dual lattice) site in an Aztec diamond of order 122 (between iterates $8 - 12 \times 10^3$). Some of the basic correspondences between ice and dominoes are given in Table 1.

Note that we only list the simplest frozen states. For ice there are more frozen states than just the ten indicated in Section 3. The same holds for dominoes, e.g., the herring bone pattern whenever it fits the domain. The fact that rotation of 2×2 domino pairs generates the set of allowed domino tilings of a given domain is in exact correspondence to our Theorem 2.3.

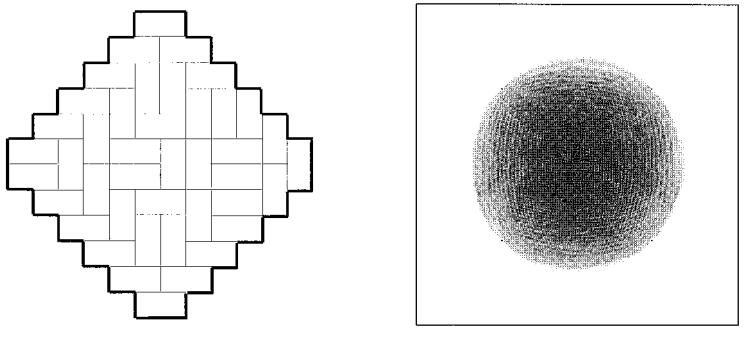

Fig. 9. Aztec diamond and the equilibrium density of 2×2-moves.

Table 1. Ice and Domino Composition

	Ice	Domino
simplest frozen states	NE, L ja X-flows and their rotations	brickwall, its shift and their rotations
maximal disorder	quasicycle boundary	square domain
allowed perturbations	directed loop reversals	rotation or reflection of a symmetric subdomain
elementary move	directed 1-cycle reversal	domino pair rotation

Moreover the same procedure that was outlined in Section 4 can be applied to dominoes as well. The resulting simple probabilistic cellular automaton, that does the random flipping of 2×2 domino pairs, gives all possible domino covers to a given domain along its orbit from any legal initial tiling. This was the method with which we made the illustration in Fig. 9, right.

The domino tilt along the Aztec diamond is extremal on each side, alternatively ± 1 as we trace the boundary around. Every tiling of the Aztec diamond has exactly one 2×2 domino pair touching each of the sides. Generically it is in the middle of the side—this is the reason the boundary of the temperate zone just grazes the diamond (in the scaling limit).

Both the dominoes and the ice-model indicate similar boundary dependency which in the simplest non-trivial case manifests as an Arctic Circle that separates the frozen and disordered regimes. The appearance of it is due to the existence of height function. The fundamental difference between the two models is that while the diamond domain forces the boundary height in dominoes, in ice it doesn't. But as soon as we choose for ice the height on a diamond boundary to be like in the domino case, the circle result follows.

The common phenomena in both models seem to be related to the non-differentiability of the average height (in the scaling limit). The physical principles have been noted that seem to imply this pushing away of the boundary curve from the diamond edge. In domino case these tilt discontinuities are only at corners. In our second example the tilt is discontinuous at the centerpoints of the diamond edges. Hence average height must be non-differentiable and indeed the same phenomenon takes place at these locations as well (Fig. 8b, c). How to formulate this rigorously remains an open problem as does the exact shapes in the scaling limit.

ACKNOWLEDGMENTS

The author would like to thank Bob Burton for early discussions on the model and the referee for useful criticism. A referee pointed out that

related connectivity and convergence rate results for the square lattice case have appeared in ref. 10. This research was partially supported by the Academy of Finland.

REFERENCES

1. R. J. Baxter, *Exactly Solvable Models in Statistical Mechanics* (Academic Press, 1982).
2. H. van Beijeren, Exactly solvable model for the roughening transformation of a crystal surface, *Phys. Rev. Lett.* **38**(18):993–996 (1977).
3. H. Cohn, N. Elkies, and J. Propp, Local statistics for random domino tilings of the Aztec diamond, *Duke Math. J.* **85**:117–166 (1996).
4. J. H. Conway, and J. C. Lagarias, Tilings with polyominoes and combinatorial group theory, *J. Combin. Theory*, *Ser. A* **53**:183–208 (1990).
5. K. Eloranta, A note on certain rigid subshifts, *Ergodic Theory of* $\mathbf{Z}^2$*-Actions*, London Mat. Soc. Lect. Notes, Vol. 228 (Cambridge University Press, 1996), pp. 307–317.
6. K. Eloranta, The bounded eight-vertex model, *Res. Report*., *Inst. of Math*., *Helsinki Univ. of Tech* (1997), submitted.
7. K. Eloranta,The bounded triangular and Kagomé ice, *Res. Report*., *Inst. of Math*, *Helsinki Univ. of Tech* (1999).
8. B. Grünbaum and G. C. Shephard, *Tilings and Patterns* (Freeman, 1987).
9. E. H. Lieb, Residual entropy of square ice, *Phys. Rev.* **162**(1):162–172 (1967).
10. M. Luby, D. Randall, and A. Sinclair, Markov chain algorithms for planar lattice structures, *36th Ann. Symp. for Found. of Comp. Sci.*, pp. 150–9 (1995).
11. W. P. Thurston, Conway's tiling groups, *Am. Math. Monthly*, pp. 757–773 (1990).