10/1997

Diamond Ice

Kari Eloranta*

Institute of Mathematics
Helsinki University of Technology
FIN-02015 HUT, Finland

eloranta@janus.hut.fi

Abstract

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 the highly non-trivial cases where results analogous but more complex than the Arctic Circle Theorem for dominoes hold.

Keywords: ice model, six vertex model, symbolic dynamics, cellular automaton, tiling, domino.
AMS Classification: 52C20, 68Q80, 82C20.
Running head: Diamond Ice

* Research partially supported by the Academy of Finland

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 upto -82 can be found in [B].

Our study concentrates on the planar version of the model defined on a finite domain with boundary. Square lattice is used partly because we wish to make comparisons to other models with similar underlying discrete structure. For the same reason we treat only diamond-shaped domains here. Our methods are strictly two-dimensional and do not generalize to other dimensions but are likely to do so to other planar lattices.

Although the results may have some relevance in understanding the geometry of long range order in this model, our principal motivation is not physical. Rather it comes from the context of tilings and symbolic dynamics. In a couple of seminal papers Conway, Lagarias and Thurston ([CL], [T]) investigated the problem of tiling a given planar domain with polyominoes. This was further extended by others, notably for dominoes in [JPS]. On the basis of this 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 (dominoes), ice model, eight-vertex model, color models and yet others. This program is now on its way and our paper is a small contribution.

The first problem to be resolved is the tileability - when is a given domain with a boundary condition possible to tile with the primitives (dominoes, polyominoes, lattice arrows etc.)? Immediately following this is the question of how many such tilings are there? This leads naturally to entropy considerations. An other line of investigation is how the legal tilings/configurations relate to each other - what are the allowed perturbations/transformations? Can the entire set of configurations with a given boundary be generated? What is the generic element like? We provide answers to these and other questions.

Our work reveals a close connection between the ice and domino models. In dominoes there are no boundary arrows to be specified, just the domain shape. Turns out that a diamond domain is highly interesting in that context and that is the case for ice as well. Both dominoes and ice have a non-trivial height function which is of critical importance in the study. Moreover both models exhibit in the case of non-trivial boundary height a rather striking boundary dependency: the Arctic Circle Theorem of [JPS] has a more complex counterpart in ice-context. A few comparisons are also made to the eight-vertex model, a seemingly close relative of the ice model. But it has trivial height and it's properties are quite different as reported in [E2].

1. Basics

The six-vertex/ice model can be defined in any dimension and for all regular lattices. 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 they may be applicable to other types of lattices than the one treated here.

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

Definition 1.1.: *An arrow configuration at a lattice site in* $\mathbf{Z}^2$ *is* **legal** *for the* **six-vertex rule** *if there are exactly two incoming arrows and two outgoing ones. A configuration in legal if it has an allowed vertex configuration at every lattice site.*

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

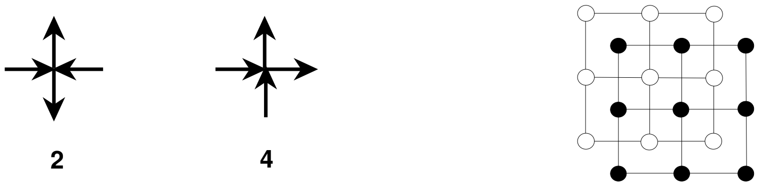

Figure 1a, b. Vertex configurations. Square lattice and its dual.

One can view the rule as incompressibility of a fluid or expressing a conservation law for some other system. It's 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 solved by Lieb for the infinite two-dimensional model in [L].

We concentrate in this paper on the finite case. The model is defined on a bounded domain for which we have to specify the boundary arrows. This is still a rather general problem. For a complicated domain shape it may be impossible to find even one boundary condition that allows a fill-in and to construct it, much less to construct all the configurations for the given boundary configuration.

For simplicity we consider a diamond domain and only occasionally point to other cases. The reason for this is two-fold. Firstly a diamond has the cleanest boundary condition. Secondly we want to relate the ice model to certain other models, in particular to the domino and eight-vertex models, for which diamond-domain is important and has been analyzed. All results of this paper generalize to rectangles standing at the corner but since this adds little insight we record the results for diamonds only.

N**-diamond** is a subset of $\mathbf{Z}^2$ which has N arrows along each of its four diagonal sides, $N \geq 2$, 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 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 ∂ to refer to the boundary.

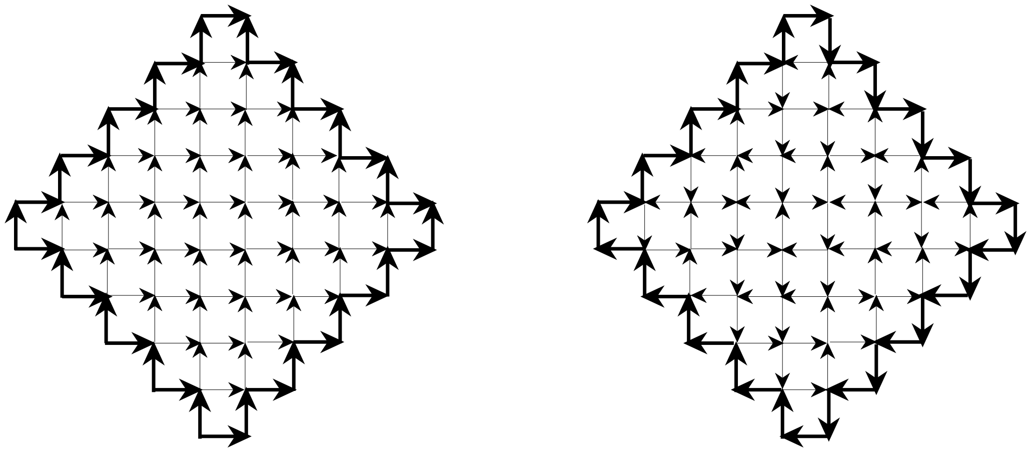

Figure 2a, b. Frozen and temperate diamond configurations.

In Figure 2. the fixed boundary arrows are rendered bold and the interior arrows light. It exhibits two extreme configuration types. The names follow the convention of earlier work; the highly ordered state on the left (the `NE` flow) is thought to correspond to low-temperature regime in Statistical Mechanics models and the one on the right 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 said to be self dual which means that it's dual lattice is the same lattice, only shifted (formally we write $(\mathbf{Z} + \frac{1}{2})^2$). It is illustrated in Figure 1b.

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 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 natural **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: The concept of height in Section 2. will 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 Figure 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 flip 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 of 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**.

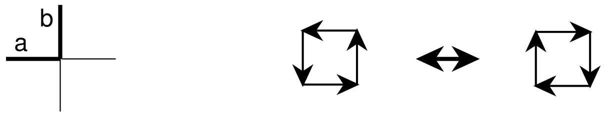

Figure 3. Vertex perturbation and elementary move on a directed 1-cycle.

Call the subsets of legal N-diamond configurations with the same boundary configuration **cosets**. Some sets of configurations are connected under reversal of directed 1-cycles (to simplify wording from this on 1-cycles will always be directed). The natural question then is to characterize these sets i.e. 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 as 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.*

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 Figure 2b. 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.

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-boundary1-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) 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 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 towards the possibility 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 configurations on a diamond with the same legal boundary condition is connected under elementary moves.*

Proof: Given two distinct diamond configurations 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 Figure 4. West of the bold broken line B all arrows agree on the two configurations. But they disagree at the arrow a_d marked by a dot. From a_d we generate the directed disagreement loop, C, exactly as in the Proof on 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.

Next find all the 1-cycle(s) inside C and connect them to C from their NE and SW corners with directed paths travelling in first and third quadrant respectively (directed from NE to SW). Using these paths and the loop C form the minimal directed loop C' that contains the disagreement arrow at a_d and some of the 1-cycles on its boundary. No 1-cycles will remain strictly in its interior. To accomplish this we may have to reverse loops like the one consisting of 1-cycle number 1, its connecting paths to the boundary and the piece of C on its left. Note that there must remain at least one 1-cycle on the inside of the loop C' again by Lemma 2.1. Call this/them c_i (in the Figure they are the ones numbered 1 and 3).

By reversing the 1-cycles on the inside of the boundary of C' we obtain a new, smaller directed loop C" inside C' that still contains a_d but none of the c_i's. 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 on 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 arrow 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. ∎

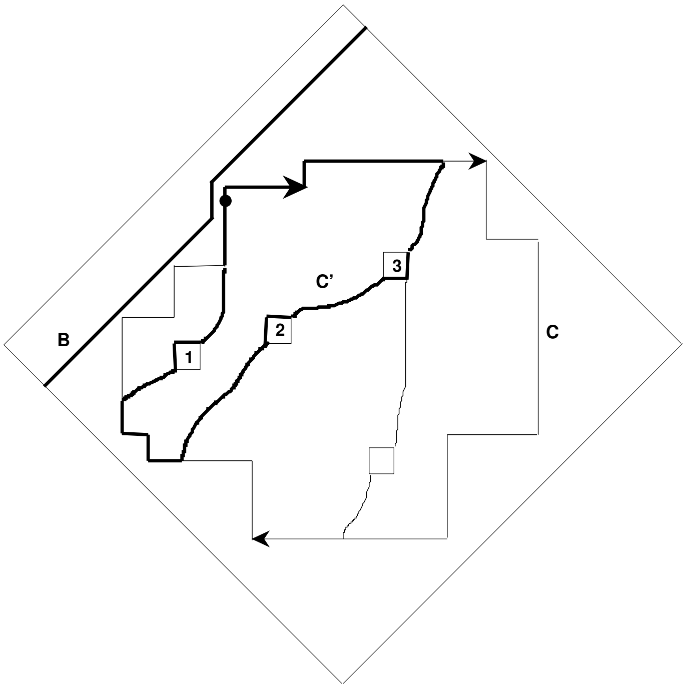

Figure 4. Proof constructions.

The Theorem indicates that indeed the best possible result is true. 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 closely related eight-vertex model, where similar irreducibility holds and the cosets are all of exactly the same size ([E2]), here they are of very different size. Furthermore unlike in the eight-vertex model the internal structure of ice configurations is generically highly non-trivial in some of the cosets.

Finally it may be of interest to note the similarity to a certain two-dimensional dynamical system treated in [E1]. In this model the number of copies of each symbol in a square neightborhood 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 ice height was first introduced using a graphical representation and without connection to flux in a paper by van Beijeren ([vB]). What makes it important notion is its utility in distinguishing different boundary conditions and the fact that it can be defined and is non-trivial for various lattice models. Among them are dominoes (dimer model), certain color models etc. This connection will be discussed in later sections.

Definition 3.1.: **Height**, h, *is an integer-valued function defined on the vertices of the dual lattice. Moving from one such lattice point to its nearest neighbor it 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 everywhere on the configuration upto an additive constant. Giving the height value at any one dual lattice point determines height 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 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. It's restriction to the boundary of the configuration, the **boundary height** is frequently useful. In Figure 5. the underlined entries are the boundary height for the given 4-diamond.
4. Tilt is a number between and including ± 1. In the example in the Figure we have recorded its value over each of the four edges as they are traced to the counterclockwise direction (in bold).

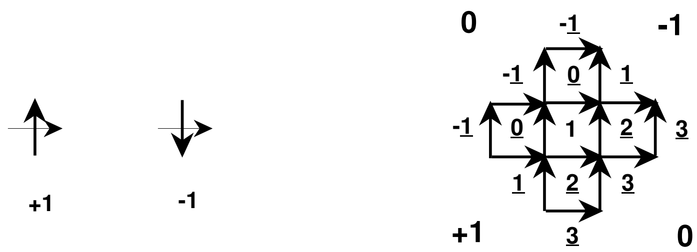

Figure 5. Height. Configuration arrows are bold, height path light.

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

Remarks: 1. The example on the right of Figure 5. illustrates extremal height.
2. Height can be viewed as a (discrete) surface above the configuration. The discrete partial derivative assuming it's 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.

To analyze the boundary dependency it is useful to distinguish **switch** 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 former type: 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.

For n even, $n/2$ adjacent switch blocks of the same sign along a diamond edge correspond to boundary height change by n. Little algebra shows that these arrows uniquely determine $n^2/4 + n/2$ arrows in the interior (the triangular arrangements in Figure 6b.).

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 of Section 1. One can quickly discover that in any N-diamond there are exactly six frozen configurations. They are the four rotations of the example in Figure 5 (the `NE` flow, Figure 2a.) and the two configurations in which the tilts are constant ± 1 alternating from edge to edge. The latter two arise because a diamond with constant edge tilt ± 1 has $N/2 - 1$ switch blocks on each edge. Each side determines $N^2/4 - N/2$ interior arrows. But there is total of $N^2 - 4N + 4$ arrows in the interior so the four edges cannot be determined independently. The arrows forced on the diagonals must agree and as a consequence of this only two such configurations exist.

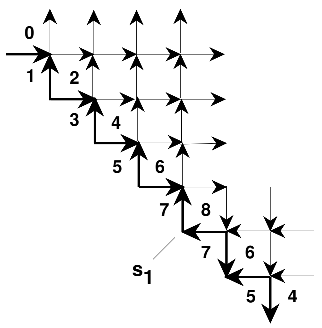

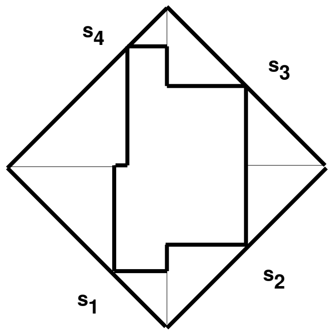

Figure 6a, b. Extremal blocks, switch point. Frozen periphery.

It is useful to quantify this a bit further. The total frozen area next to the boundary is just the sum of the triangular areas next to the extremal boundary pieces. Formally then

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

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

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

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

Proof: As observed above $n/2$ adjacent switch blocks of the same sign determine $n^2/4 + n/2$ interior arrows in a triangle formation. However the maximum of n of these (the ones on the boundary) may be shared with two adjacent triangle edges. Hence only half of these shared arrows will be accounted in the lower bound. This gives the expression $n^2/4$. The rest is summing up the contributions over maximally long boundary pieces each of extremal height. ∎

The diametrically opposite case to the extreme height is the case of a directed cycle boundary. For such boundary the boundary height is identically zero in the diamond since the boundary is made of neutral blocks of the same heading (outside the diamond the boundary height is either ± 1. This is the temperate case illustrated in Figure 2b. By Lemma 2.1. there is non-uniqueness in the fill-in and indeed this non-uniqueness is maximal. This is due to the fact that 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. This is argued 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 cycle). The configuration where all the arrows have been reversed is another legal diamond configuration with a cycle boundary. Reversing this boundary cycle 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 boundary is zero. Hence to 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 is non-trivial if it contains segments of 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 i.e.* $a|_\partial$ *is a cycle.*

(iii) *Temperate and frozen phases coexist in* a *iff* $a|_\partial$ *is nontrivial.*

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

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

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

Using this we immediately get an asymptotic upper bound for the number of legal configurations in the case of a boundary condition. By (1) the fraction of the arrows in a N-diamond that is fixed by the boundary is $A_f(N)/N^2$. 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 $1/N$, $N \to \infty$, hence the limiting domain will be the unit diamond $|x| + |y| \leq 1/2$) 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*

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

whenever the limit exists.

Remarks 1.: The frozen cases where (1) gives $A_f(N) = N^2 + O(N)$ are clearly zero entropy. On the other hand the cycle boundary implies $A_f(N) \equiv 0$. Hence the temperate case seems to be the maximum entropy case. The intermediate range where the bound is non-trivial is most interesting and the subject of the rest of the paper.

2. 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 ([E2]).

4. Dynamics

We now indicate 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 $S = \{0, \ldots, 15\}$ 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 from 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.

As the rule only uses integer operations it can be implemented as a fast look-up 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.

5. Flower and comparison

For a legal boundary with non-trivial height the geometry and statistics in the interior of the configuration are highly non-trivial. We now investigate this regime and compare it to an other statistical mechanics model, the dominoes (the dimer model).

Recall that the non-triviality of the boundary meant that the boundary configuration has segments of extremal height 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 each side is split into two equally long extremal pieces. If the switch points s_i are the midpoints of the edges the unforced part of the interior is a square as indicated in Figure 7a. Next to the sides we have noted the tilts. A simple variational calculation also shows that with this choice of s_i's the "free" area in the center is maximal.

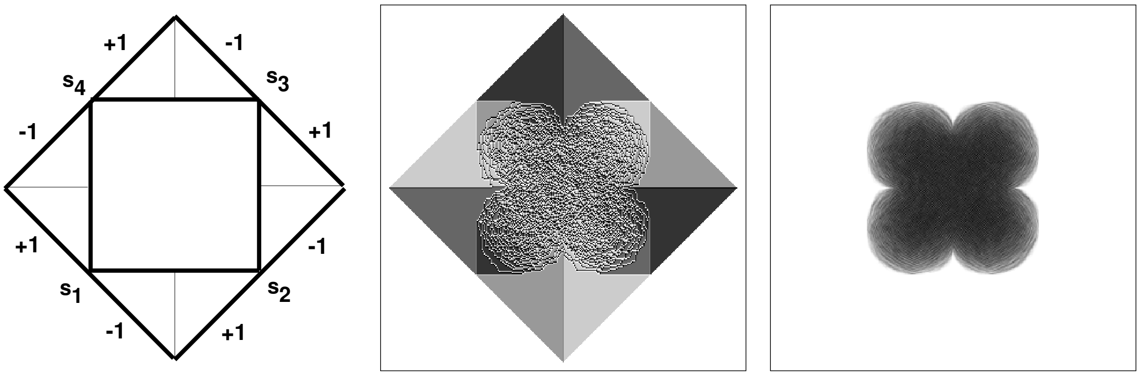

Figure 7a, b, c. Boundary tilts, configuration and the 1-cycle density.

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 $F_{1/2}$ of the previous Section to relax from any of these the result is generically as in the middle illustration.

The middle and right 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. The 16 different symbols are rendered in different grades of grey. The frozen area forced by the boundary is clearly visible. But more interesting is the interior of the “free” square. There is a clear demarcation between the frozen and temperate domains which results in the flower-like boundary curve. 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 $1.1 - 1.5 \times 10^4$ and converted this to grey level. Dark cells are the sites of most activity. In the frozen triangles there is obviously no such activity hence they are rendered white.

The relaxation was performed to several different initial configurations with the given boundary structure in diamonds of different size. The results were all essentially as above except that in a larger diamond the boundary of the flower appears smoother.

Four copies of the boundary above can be glued together to form a $2N$-diamond (remove the arrows from the interior). It has therefore three switch points along each edge, evenly distributed. This relaxes in the expected fashion. Its configuration as well as 1-cycle density plot are in Figure 8 (156-diamond, configuration at iterate 8×10^3, 1-cycles collected between iterates 4×10^3 and 8×10^3).

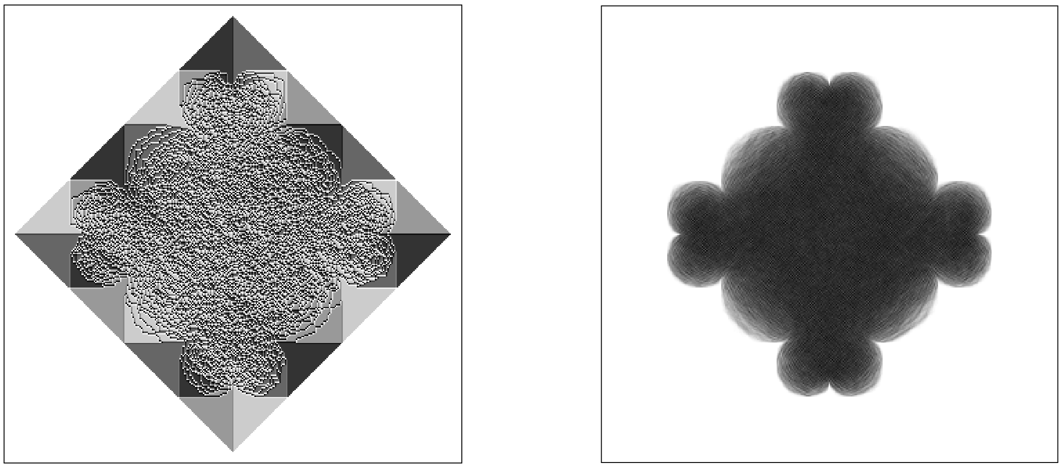

Figure 8a, b. Less restricted ice domain and the equilibrium 1-cycle density

Suppose 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 height function for dominoes analogously to Definition 3.1. (see [JPS]). Note however that since there are no arrows, only the shape of the domain determines the boundary height. 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. In [JPS] 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. We illustrate this result in Figure 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 $8 - 12 \times 10^3$).

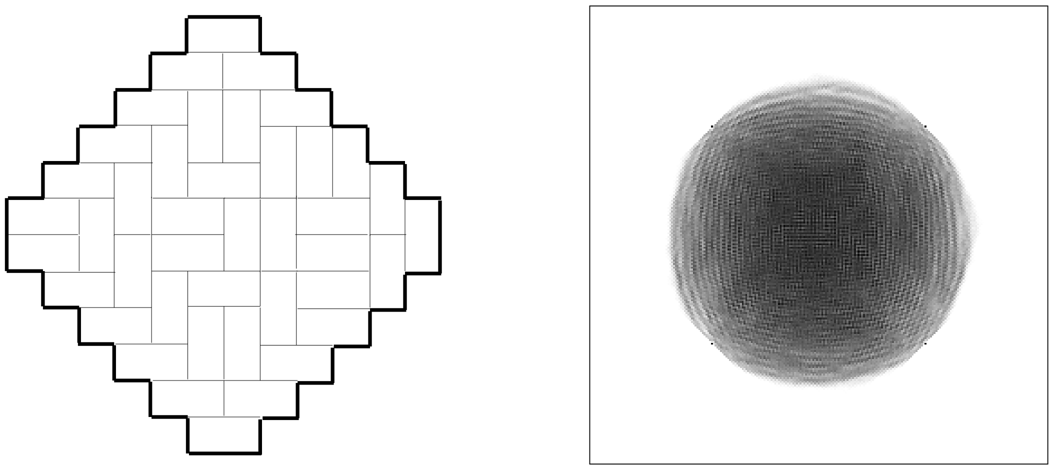

Figure 9a, b. Aztec diamond and the equilibrium density of 2×2-moves

Some of the basic correspondences between ice and dominoes are as follows.

	Ice	Dominoes
frozen states	`NE, NW, SE, SW` flows and two exotic (6 total)	brickwall, its shift and rotations (4 total)
allowed perturbations	directed loop reversals	rotation or reflection of symmetric subdomain
elementary move	directed 1-cycle reversal	domino pair rotation

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. (for dominoes this result seems a Folklore Theorem). 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 illustration in Figure 9b.

The table also explains why the Aztec diamond result is of interest in ice context. The domino height along the Aztec diamond is extremal on each side, alternatively increasing and decreasing as we trace the boundary around. The tilt is constant ± 1 on each side. This boundary height would yield a frozen configuration in ice - one of the two "exotic" ones. However 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). Our boundary condition for ice allows exactly one 1-cycle at the center of each side, at the switch points $\{s_i\}$ (see Figure 7a.). Elsewhere the height is extremal and alternating. So it is the simplest ice-boundary that yields non-trivial interior behavior and in terms of height and elementary moves it is as close to an Aztec diamond as possible.

The common phenomenon in both models seems to be related to the non-differentiability of the average height (in the scaling limit). In [JPS] the physical principles were 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 the ice-context the tilt is discontinuous at the centerpoints of the diamond edges as well hence average height must be non-differentiable and indeed the same phenomenon takes place at these locations as well (Figure 7b, c.). How to formulate this rigorously remains an open problem as does the exact shape of the flower figure in the scaling limit.

Acknowledgement

The author would like to thank Bob Burton for early discussions on the model.

References

[B] Baxter, R.J.:*Exactly solvable models in statistical mechanics*, Academic Press, 1982.

[vB] van Beijeren, H.: Exactly solvable model for the roughening transformation of a crystal surface, *Phys. Rev. Lett.*, **38**, 18, pp. 993-996, 1977.

[CL] Conway, J.H., Lagarias, J.C.: Tilings with polyominoes and combinatorial group theory, *J. Combin. Theory*, Ser. A, **53**, pp. 183-208, 1990.

[E1] Eloranta, K.: A note on certain rigid subshifts, *Ergodic Theory of* $\mathbf{Z}^d$*-Actions*, London Mat. Soc. Lect. Notes **228**, Cambridge Univ. Press, pp. 307-317, 1996.

[E2] Eloranta, K.: The bounded eight-vertex model, 11 p., preprint, 1997.

[GS] Grünbaum, B., Shephard, G. C.: *Tilings and patterns*, Freeman, 1987.

[JPS] Jockush, W., Propp, J., Shor, P.: Random domino tilings and the Arctic Circle Theorem, preprint, 1995.

[L] Lieb, E.H.: Residual entropy of square ice, *Phys. Rev.*, **162**, 1, pp. 162-172, 1967

[T] Thurston, W.P.: Conway's tiling groups, *Am. Math. Montly*, pp. 757-773, Oct. 1990.