Random Majority Percolation

Paul Balister[*†] Béla Bollobás[†‡§] J. Robert Johnson[¶]
Mark Walters[*¶]

March 2, 2009

Abstract

We shall consider the discrete time synchronous random majority-vote cellular automata on the n by n torus, in which every vertex is in one of two states and, at each time step t, every vertex goes into the state the majority of its neighbours had at time $t-1$ with a small chance p of error independently of all other events. We shall show that, if n is fixed and p is sufficiently small, then the process spends almost half of its time in each of two configurations. Further, we show that the expected time for it to reach one of these configurations from the other is $\Theta(1/p^{n+1})$ despite the actual time spent in transit being $O(1/p^3)$. Unusually, it appears difficult to obtain any results for this regime by simulation.

In this paper we shall consider various cellular automata on graphs. Cellular automata with random initial configuration and deterministic update rules have been studied in numerous papers (see the following section for a discussion of previous work), however, our emphasis in this paper is on automata with probabilistic update rules, i.e., rules that are applied with probability less than one. Such cellular automata are often refered to as

---

[*] Research supported by NSF grant EIA-0130352.

[†] University of Memphis, Department of Mathematics, Dunn Hall, 3725 Noriswood, Memphis, TN 38152, USA.

[‡] Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Cambridge CB3 0WB, UK

[§] Research supported in part by NSF grants DMS-0505550, CNS-0721983 and CCF-0728928, and ARO grant W911NF-06-1-0076

[¶] Queen Mary University of London, London E1 4NS. UK

Probabilistic Cellular Automata (PCA). They have been extensively studied experimentally approximations to some of these models, and a very small number of completely theoretical results (see later for more details). Our aim here is to prove results about a particular probabilistic percolation on the square grid; in fact many of our results will be proved in more generality since then they can be applied more widely (for example to any lattice, in more than two dimensions etc).

The main model we shall consider has a long history: exactly our model was considered by Gray in [14] and very similar models were considered much earlier (e.g. Spitzer [31] in 1970).

To define our models we let G be a graph (which may include loops and multiple edges) with vertex set V and edge set E. We wish to consider several models on G with configuration space $\Omega = \{0, 1\}^V$, so that each vertex is either in state 0 or state 1. The first is the *majority model* which is defined only for graphs in which all the degrees are odd (note that we follow the convention that a loop adds *one* to the degree of a vertex, not two). In this model, at each time step a vertex takes the value of the majority of its neighbours. In the second model, which we call the *random majority model*, each vertex takes on the value of the majority of its neighbours independently with probability $1 - p$ and takes its complement with probability p (we envisage p as small).

A simpler mean field version of the random majority model has been rigorously analysed and many computer simulations have been done of the random majority model on the torus $\mathbb{T}_n^2$ (with single loops added at each vertex so that the degree of each vertex is 5). These both suggest that there is phase transition as p varies: if p is less than some critical probability p_c (computer simulations suggest $p_c \approx 0.134$) then the model spends almost all of its time in either a low density (most vertices in state 0) or high density (most vertices in state 1) configuration. More importantly, computer simulations suggest that the model spends a long time in either low or high density configurations before crossing rapidly to the other density. On the other hand, if $p_c < p \leq 0.5$ the model spends almost all of its time with about half of the vertices in each of the two states.

In this paper we prove rigorously that for very small p the model spends almost all its time in either a low or a high density configuration. A more formal statement is the following. (The exact statement requires some more definitions: see the start of section 4.)

Theorem. *Fix n, and let $\mathbb{T}_n^2$ be the $n \times n$ torus. Then, in the random majority model, the expected time to get from the all zero configuration to the all one configuration is $\Theta(1/p^{n+1})$ as $p \to 0$, whilst the crossing time (the time from when it last leaves the all zero* *configuration before reaching the all one configuration) is* $O(1/p^3)$. *Moreover the proportion of time the model spends in each of the all zero configuration and the all one configuration is* $1/2 - O(p)$.

The theorem only applies for very small values of p; in particular, our proof is only true when p is very much less than 2^{-n^2}, although we believe that our methods can be extended to about $p = o(n^{-6})$. In order to prove anything about the existence of the phase transition mentioned above, this theorem would have to be extended to the case where p is a (very small) constant independent of n. Unfortunately this appears to be well beyond current techniques.

For the case of very small p our theoretical results are ahead of simulations: indeed it appears to be very hard to simulate $p = 0.01$ (let alone the sort of p that we deal with) on even on the 10×10 square torus since our results suggest a time to change from the all zero configuration to the all one configuration of about 10^{22} steps.[1]

The reason our results are ahead of simulations is not because the simulation has to wait a long time for an error to occur: indeed, since the majority model stabilises very quickly you can decide when the errors should occur and work out what the configuration will be immediately prior to the next error *without* calculating the intermediate steps.

A better way of thinking of the difficulty is that a huge number of *errors* (not just steps) must occur. For example, in the case mentioned above of $p = 0.01$ and the 10×10 torus our results suggest that at least 10^{20} *errors* must occur at each vertex. Indeed, one way of visualising how long we have to wait is the following: we wait until all the cells in a single row of the torus (plus one more cell) *simultaneously* have an error!

Relation with Previous Work

General deterministic cellular automata have been very widely studied from Conway's Game of Life to the recent book by Wolfram [33]; there are far too many references to state here. However, our model is very different from these: the complexity in the random majority model arises from the random "noise", not from a complicated update rule: indeed, the update rule is one of the simplest possible.

However, we shall make use of some results about the deterministic majority rule. Goles and Olivos [13], and Poljak and Sûra [27], proved that, for finite graphs, the model either reaches an unchanging situation (i.e., period one) or reaches a cycle of period two.

---

[1] There are some heuristics suggesting that our bounds extend to p of about this size.

Moran [24] extended this to a certain collection of infinite graphs. Poljak and Turzík [28] proved some good bounds on the time the model takes to reach the stable or periodic configuration. For more details see Section 2.

As mentioned above cellular automata with deterministic update rules and a random initial configuration have also been studied extensively, see [1, 2, 3, 4, 6, 12, 18, 25, 29, 30] among many others. However, we should emphasise that this is still completely different from our model: in the random majority model (and PCAs in general) the random noise occurs continually *throughout* the process, not just in the initial configuration.

It is clear, for example, that our model will visit every configuration (assuming the graph G is finite) since at some point any specific collection of errors can occur simultaneously. Periodic behaviour will not occur (almost surely) since the errors will not be periodic. The initial configuration is irrelevant in our model whereas it is obviously of primary importance in models with a deterministic update rule.

There are a small number of papers discussing our model: it is often called the (discrete time, synchronous) majority-vote model. Most of these are by physicists: these are almost all ([26, 23] for example) computer simulations (i.e., Monte-Carlo methods) with a few (e.g. [20, 5]) which prove some results about the mean field approximation with the hope that the results carry over to the actual model. Of course neither of these methods prove anything about the actual model, but they do provide strong evidence that there is a phase transition, and indeed good estimates on the critical probability and critical exponents.

There are very few papers dealing rigorously with models similar to ours: this seems to be due to the inherent difficulty of proving anything about them. For example, in the paper [23] by Molofsky *et al.* mentioned above they explicitly say "Unfortunately, rigorous mathematical analysis is impossible ... Thus we resort to Monte-Carlo simulation ..."

Some PCAs can be related to known models in statistical physics, in which case we can say more about them. For example, Glauber dynamics for the 2-dimensional Ising model can be viewed as a class of PCAs. Since the 2-dimensional Ising model can be solved explicitly, and is known to have a phase transition, the same can be said of these PCAs. However, even in this case, many dynamical properties of these models remain unknown.

Glauber dynamics for the Ising model is very similar to the random majority model: at each time step a random vertex is chosen and then flipped with probability $1/(1+e^{\beta(s-d)})$ where s is the number of neighbours currently in the same state and d is the number of neighbours in the opposite state. Here β is the inverse temperature parameter of the Ising model. For the random majority model, the probabilities are p and $1-p$ depending on whether $s \geq d$ or $d < s$, and all sites are updated synchronously, rather than one at a time. Nevertheless, it is believed that the two models share the same general behaviour, falling in a common universal class of models with, for example, the same critical exponents. Unfortunately none of this is proven.

One of few rigorous results on the random majority model is due to Gray [14]: he proved that the one dimensional version of our model does not have a phase transition (as is the case for the 1-dimensional Ising model). Even this is unexpectedly difficult to prove.

Gács [11] proved the very surprising and impressive result that there is a one-dimensional probabilistic cellular automaton with a phase transition. His model has a very complicated update rule and a huge number of vertex states: at least 2^{30}. Gács's paper is very difficult to read: it consists of over 220 pages of dense argument. Indeed, even Gray's simplified presentation of it [15] runs to nearly 50 pages.

Gács's result contradicts what one would expect from statistical physics: i.e., phase transitions do not occur in one dimension (note, however, that his model is not reversible). In Gray's simplification it is explicitly mentioned that "For the two-dimensional result [i.e., for the random majority model], we only have computer simulations" and that nothing rigorous is known about the phase transition.

Finally, Toom [32] proved that the following majority-vote model does have a phase transition: he took the cubic lattice in three dimensions and took as neighbours to a vertex the three vertices above, to the right and behind it. This model is much simpler than our model as there is no "feedback": i.e., the graph is oriented and any errors will drift away. (This is similar to directed percolation being much simpler to analyse than undirected percolation.)

In this paper we prove results for very small values of p: in particular we require $p \ll 2^{-n^2}$ (however as mentioned above we believe the methods can be extended to $p = o(n^{-6})$). Our results do not prove anything about the phase transition, although they do add further evidence as to its existence. We think that our results are of interest in their own right as some of the very small number of rigorous results about models of this type. Moreover, understanding p is this range may well be a necessary step to proving the existence of the phase transition: indeed, in Gray's proof that there is not a phase transition in one dimension much of the work is done in a smallish part of the space-time domain in which all occurences of three or more errors can be ignored. Our bound on p is restricting us to a similar regime.

The layout of the Paper

The layout of this paper is as follows. In the first section we set up some notation and detail the models that we shall be considering. In the second section we briefly discuss the deterministic majority model and quote some known results about it: in particular that the only periodicity we can have in the majority model is of period two and some bounds on the time it takes to reach such a periodic configuration.

In the third section we show that some aspects of the behaviour of the random r-threshold model are closely related to the $(r-1)$-monotone model (both models defined in the next section).

It is easily seen that the random majority model on the square torus spends most of its time in the all zero configuration or the all one configuration. In the fourth section, which is the heart of the paper, we shall prove that the expected time to change from one to the other is $\Theta(1/p^{n+1})$. In the final section we prove some bounds on how small a set can be if it grows to fill the whole space in the r-monotone model on the cubic lattice in d dimensions.

1 Notation, Preliminaries and Related Models

Throughout this paper $G = G(V, E)$ will denote a multigraph with vertex set V and edge set E with loops allowed. We will write $|G|$ for the number of vertices in G, and $\Gamma(x)$ for the neighbours of x in G; note that x is contained in $\Gamma(x)$ exactly when x has a loop. It is often convenient to view G as a directed graph $\vec{G} = \vec{G}(V, \vec{E})$ where we we replace each edge xy by *both* the directed edges $\vec{xy}$ and $\vec{yx}$, and replace a loop by the *single* directed edge $\vec{xx}$. The number of edges of G will be denoted $e(G)$ and is defined to be $\frac{1}{2}|\vec{E}|$: this means that each loop only counts as half an edge. Similarly we define the degree of a vertex x to be the out-degree of x in $\vec{G}$, hence loops only count one towards the degree. Although it is often helpful to think of G as a directed graph, we wish to emphasise that, throughout the paper, we only work on directed graphs coming from an undirected graph in this way. Indeed, many of the results we prove are false for a general directed graph.

We will let $\Omega = \{0, 1\}^V$ be the space of configurations. A point $\omega \in \Omega$ will be called a configuration. At times we will identify ω with a subset of V treating ω as a characteristic function. In this spirit we define $\omega^c = 1 - \omega$. The value $\omega_x \in \{0, 1\}$ of ω at a vertex x will be called its state. Vertices with $\omega_x = 1$ will be called active. Vertices will also be called cells or sites.

In addition to the majority model and the random majority model there are some other models that it will be useful to consider in the proofs. The first is the *r-monotone model* in which, at each step a vertex becomes active (state 1) if at least r of its neighbours are active, and once active no cell ever becomes inactive. (In the literature this is also known as *bootstrap percolation*.)

The second model, which we call the *r-threshold model*, is a generalisation of the majority model. At each step a vertex takes the 1 state if at least r of its neighbours are active and otherwise it takes the 0 state. It is important to note that this is not the same as the threshold voter model considered in e.g., [10]: there the vertex changes state if at least r of its neighbours are in the other state, whereas here the vertex changes from 0 to 1 if it has at least r neighbours in state 1, *but* changes from 1 to 0 if at least $d(v) - r + 1$ are in state 0. Observe that, for a regular graph of odd degree Δ, the $(\Delta + 1)/2$-threshold model is the majority model.

The final model is the *random r-threshold model*, defined analogously to the random majority model: i.e., a vertex takes the value that the threshold model would give with probability $1 - p$ and its complement with probability p.

For a fixed graph G, configuration ω and r we define three functions ω_T, $\omega_{R,p}$, ω_M: $\mathbb{N} \to \Omega$ where $\omega_T(t)$ is the configuration reached after time t in the r-threshold model (respectively the r-random-threshold model, the r-monotone model) *starting* from the configuration ω. In the random majority model the event that a vertex does not take the state given by the majority of its neighbours will be called a *flip*. Note that more than one flip can occur at the same time step. We define flip similarly in the random threshold model.

If the model is clear for the context, then for any vertex x we will write x_i for the state of the vertex x at time i.

Finally, for completeness, we recall some standard notation for dealing with the asymptotic behaviour of functions. If f and g are two functions we shall write $f = O(g)$ to mean that there is some constant C with $f(x) \leq Cg(x)$ for all sufficiently large x. Similarly, we shall write $f = \Omega(g)$ to mean that there is some $c > 0$ such that $f(x) \geq cg(x)$ for all sufficiently large x. Finally, we write $f = \Theta(g)$ to mean $f = O(g)$ and $f = \Omega(g)$; i.e., there exist $c, C > 0$ with $cg(x) \leq f(x) \leq Cg(x)$ for all sufficiently large x.

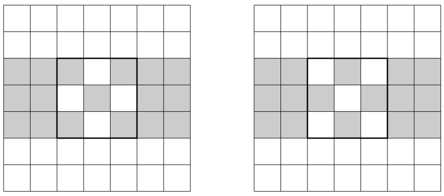

Figure 1: A bounded configuration with period 2. The squares represent the cells.

2 Periodicity of the Majority Model

We start by considering the deterministic majority model. It is obvious that, on a finite graph G, the model must be eventually periodic (possibly with period one). It is easy to see that two-periodic configurations can occur even in very nice graphs such as the square lattice with loops (that is the graph of $\mathbb{Z}^2$ with each point joined to its four nearest neighbours and itself): e.g., take a checker-board pattern. Moreover, we can insist that the periodic region is bounded as shown in Figure 1. This example also shows that the number of active sites need not be constant; indeed by taking many copies of this example we can get an arbitrarily large change in the number of active sites.

Goles and Olivos [13], and independently, Poljak and Sûra [27] proved that two-periodicity is the only non-trivial periodicity that can occur. Formally they proved,

Theorem 1. *Suppose that G is a finite graph with all degrees odd. Then, from any starting configuration, the majority model reaches a configuration which has period one or two.*

In fact they also proved this for any threshold model.

We shall also be interested in the length of the *transient* phase, that is, the time it takes the model to reach a periodic configuation. It is trivial that this is at most $2^{|G|}$. However, Poljak and Turzík proved (Theorem 1 of [28]) the following much stronger result (restated for our situation).

Theorem 2. *Suppose that G is a finite graph with all degrees odd. Then the time T it takes for the majority model to reach a periodic configuration is at most $\frac{1}{2}\sum_{v\in G}(d(v)-1)$. In particular, if $G=\mathbb{T}_n^2$, the $n\times n$ torus with a single loop at each vertex, then we have $T\leq 2n^2$.*

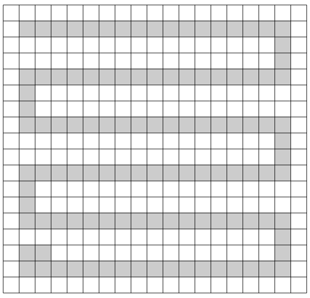

Figure 2: A configuration in $\mathbb{Z}_n^2$ with loops that takes linear time to stabilise.

See Figure 2 for an example that takes time $\Omega(n^2)$ (approximately time $\frac{1}{3}n^2$) to stabilise.

A very simple modification of their proof yields the following: (alternatively a similar but slightly weaker bound is immediate from Theorem 2 of the same paper).

Corollary 3. *Suppose G is a finite graph and that r is fixed. Then the time it takes for the r-threshold model to reach a periodic configuration is at most $4e(G) + |G|$.*

We conclude with the brief observation that many of these results can be extended to certain infinite graphs: in particular to infinite graphs with sub-exponential growth. See Moran [24] for full details.

3 Random Majority and Monotone Models.

In this section we aim to describe some behaviour of the random majority model. For some technical reasons which we will discuss later it is simpler to work on the r-random-threshold model. Throughout this section we will suppose that G is regular of degree Δ and that every vertex has a single loop and that G has no multiple edges. Recall that a loop contributes *one* to the degree (not two).

Fix r with $2 \leq r \leq \Delta - 1$ and ω. We define three functions ω_T, $\omega_{R,p}$, ω_M: $\mathbb{N} \to \Omega$ corresponding to these processes where $\omega_T(t)$ (respectively $\omega_{R,p}(t)$, $\omega_M(t)$) is the configuration reached after time t in the r-threshold model (respectively the r-random-threshold model, the r-monotone model). Throughout this section r will be fixed and we will usually omit it from the notation. Recall that another way of thinking of the random threshold model is to alternately apply one step of the threshold model and and then randomly flip each cell independently with probability p. For ease of notation, $\omega_T(t)$ will be defined as $\omega_T(\lfloor t \rfloor)$ when t is not an integer (and similarly for $\omega_{R,p}$ and ω_M).

Let $\mathbf{1}$ be the all one configuration. We say that a configuration ω *monotone percolates* if $\omega_M(t) = \mathbf{1}$ for some t. We say a configuration ω *threshold-p percolates* if there exist $\varepsilon > 0$ and C such that

$$\mathbb{P}(\omega_{R,p}(C/p) = \mathbf{1}) > \varepsilon$$

for all sufficiently small p (i.e., there is a reasonable chance that the configuration becomes full in a reasonable amount of time in the random threshold model). We wish to show that this is a reasonable definition. In particular that this definition is the same as saying ω gets to the $\mathbf{1}$ configuration sometime before time C/p. We prove a simple lemma first.

Lemma 4. *Suppose that C is fixed and $\mathbf{1}$ is the all one configuration. Then for any $k \leq C/p$*

$$\mathbb{P}(\mathbf{1}_{R,p}(k) = \mathbf{1}) \geq 1 - O(p)$$

uniformly in k.

(In this section we will be concerned with percolation as defined above: thus throughout the section the limiting regime, for example in O notation, will be as p tends to zero.)

Proof. The probability that two or more flips occur at time t is at most $|G|^2p^2$. Therefore the probability that, in the time up until C/p, any two flips are simultaneous is at most $|G|^2Cp$. If no two such events occur then any vertex that is deleted by a flip immediately reappears under the simple threshold rule (since $r \leq \Delta - 1$) at the next time step (unless it is flipped a second time). Thus

$$\begin{aligned}\mathbb{P}(\mathbf{1}_{R,p}(k) \neq \mathbf{1}) &\leq \mathbb{P}(\text{two flips are simultaneous}) + \mathbb{P}(\text{a flip occurred at time } k)\\ &\leq |G|^2Cp + |G|p = O(p)\end{aligned}$$

and the result follows. □

Lemma 5. *Suppose that ω is a configuration and C a constant such that, for all sufficiently small p,*

$$\mathbb{P}(\exists t < C/p : \omega_{R,p}(t) = \mathbf{1}) > \varepsilon.$$

Then for any $C' \geq C$ and sufficiently small p

$$\mathbb{P}(\omega_{R,p}(C'/p) = \mathbf{1}) > \varepsilon/2.$$

Proof. We have (by Lemma 4)

$$\begin{aligned}
\mathbb{P}(\omega_{R,p}(C'/p)=\mathbf{1}) &\geq \sum_{t=1}^{C'/p} \mathbb{P}(\omega_{R,p}(t)=\mathbf{1} \text{ and } \forall t'<t\colon \omega_{R,p}(t')\neq \mathbf{1})\mathbb{P}(\mathbf{1}_{R,p}(C'/p-t)=\mathbf{1})\\
&\geq \sum_{t=1}^{C/p} \mathbb{P}(\omega_{R,p}(t)=\mathbf{1} \text{ and } \forall t'<t\colon \omega_{R,p}(t')\neq \mathbf{1})(1-O(p))\\
&\geq \mathbb{P}(\exists t<C/p : \omega_{R,p}(t)=\mathbf{1})(1-O(p))\\
&\geq \varepsilon/2
\end{aligned}$$

for sufficiently small p. □

A priori, it is possible that, for some C a configuration threshold-p percolates (i.e., $\mathbb{P}(\omega_{R,p}(C/p)=\mathbf{1})>\varepsilon$ for some $\varepsilon>0$) and for other C it does not. The previous lemma tells us that if it threshold-p percolates for C it does for $C'>C$. The following lemma says that the same is true for $C'<C$.

Lemma 6. *Suppose ω threshold-p percolates with constants C and ε. Then, for any $0<C'<C$, there exist $\varepsilon'>0$ such that ω threshold-p percolates with constants C' and ε'.*

Proof. The expected number of flips in time C/p is $C|G|$. Hence the probability that there are more than $2C|G|/\varepsilon$ flips is at most $\varepsilon/2$. Thus,

$$\mathbb{P}(\omega_{R,p}(C/p)=\mathbf{1} \text{ and there are at most } 2C|G|/\varepsilon \text{ flips}) > \varepsilon/2.$$

Consider a particular "route" (i.e., sequence of flips) from ω to $\mathbf{1}$ with at most $2C|G|/\varepsilon$ flips. Suppose it has f flips. Then the probability r of that particular route is $p^f q^{|G|C/p-f}$, where $q=1-p$. Thus the probability of the same route occurring when the probability of a flip is $p'<p$ is at least $(p'/p)^f r \geq (p'/p)^{2C|G|/\varepsilon} r$. Since this is true for every route in the above sum we have

$$\begin{aligned}
\mathbb{P}(\omega_{R,p'}(C/p)=\mathbf{1}) &\geq \left(\tfrac{p'}{p}\right)^{2C|G|/\varepsilon} \mathbb{P}(\omega_{R,p}(C/p)=\mathbf{1} \text{ and there are at most } 2C|G|/\varepsilon \text{ flips})\\
&\geq \left(\tfrac{p'}{p}\right)^{2C|G|/\varepsilon} \varepsilon/2.
\end{aligned}$$

However, if we choose p' so that $C'/p'=C/p$, the left hand side is just $\mathbb{P}(\omega_{R,p'}(C'/p')=\mathbf{1})$, and the result follows with $\varepsilon'=\frac{1}{2}(C'/C)^{2C|G|/\varepsilon}\varepsilon$. □

The next lemma says that the random threshold model is increasing in the sense that if we start in a larger configuration then we will stay in a larger configuration; of course this can only be a probabilistic statement. In particular it shows that the set of configurations that threshold-p percolate is an up-set.

Lemma 7. *Let ω and σ be any two configurations with $\omega \leq \sigma$. Then for any up-set of configurations A, time t and $p < 1/2$ we have*

$$\mathbb{P}(\sigma_{R,p}(t) \in A) \geq \mathbb{P}(\omega_{R,p}(t) \in A).$$

Proof. As mentioned before we can view the random-threshold process as applying the simple threshold algorithm and then flipping each point with probability p. Since $\omega \leq \sigma$ we have $\omega_T(1) \leq \sigma_T(1)$. So all we need to check is that the flipping does not disturb things.

We couple the flipping process on ω and σ. Let x be uniformly distributed on $[0,1]^{V(G)}$. Then for a configuration ϕ define $f(\phi)$ by

$$f(\phi)_v = \begin{cases} 0 & \text{if } \phi_v = 0 \text{ and } x_v < 1-p; \\ 1 & \text{if } \phi_v = 0 \text{ and } x_v \geq 1-p; \\ 0 & \text{if } \phi_v = 1 \text{ and } x_v < p; \\ 1 & \text{if } \phi_v = 1 \text{ and } x_v \geq p. \end{cases}$$

Then $f(\phi)_v$ has the distribution of the random flipping process. Moreover if $\omega \leq \sigma$ we have $f(\omega) \leq f(\sigma)$. Applying this to $\omega_T(1)$ and $\sigma_T(1)$ we see that $f(\omega_T(1)) \leq f(\sigma_T(1))$ or equivalently $\omega_{R,p}(1) \leq \sigma_{R,p}(1)$. A simple induction now yields the result. □

Our aim is to show that, under some fairly weak conditions, a configuration r-threshold-p percolates if and only if it $(r-1)$-monotone percolates. First observe that we cannot hope that this is true in general: the configuration could collapse to the all inactive configuration under the simple threshold rule. Thus with high probability it would collapse to the all empty configuration in the random-threshold rule (this is formalised below). However we can hope that it is true if the configuration is static or periodic under the simple threshold rule. This is slightly complicated by the fact that there are periodic configurations in the simple threshold model; however, as we saw in the previous section, all such configurations must have period two.

Let a *static configuration* be a configuration which is does not change in the threshold model: that is $\omega_T(1) = \omega$. Let a *semi-stable configuration* of w be any configuration that occurs infinitely often in $\omega_T(i)$. In other words we run the simple threshold model starting at ω. Either it reaches a static configuration, in which case this configuration is its unique semi-stable configuration, or it alternates between two configurations, in which case the two configurations are each semi-stable configurations of ω. For a semi-stable configuration ω let ω_{inf} be the set of static active sites of ω (i.e., we replace all periodic sites by inactive sites).

Lemma 8. *A configuration w threshold-p percolates if and only if any of its semi-stable configurations do.*

Proof. Let σ be one of the semi-stable configurations of ω. Then, by Corollary 3, $\sigma = \omega_T(k)$ for some $k \leq 4e(G) + |G|$. The probability that any flip occurs in the first k steps in the random-threshold model is at most $k|G|p$.

Suppose that σ threshold-p percolates. Let ε and C be such that

$$\mathbb{P}(\sigma_{R,p}(C/p) = \mathbf{1}) > \varepsilon$$

for all sufficiently small p. Then with k as above and for sufficiently small p, ($p < \frac{1}{2k|G|}$),

$$\begin{aligned}
\mathbb{P}(\omega_{R,p}((C+1)/p) = \mathbf{1}) &\geq \mathbb{P}(\omega_{R,p}(C/p + k) = \mathbf{1})/2 && \text{(by Lemma 5)}\\
&\geq \mathbb{P}(\omega_{R,p}(k) = \sigma)\mathbb{P}(\sigma_{R,p}(C/p) = \mathbf{1})/2 && \text{(one such route)}\\
&\geq (1 - k|G|p)\varepsilon/2\\
&\geq \varepsilon/4.
\end{aligned}$$

Conversely suppose that $\mathbb{P}(\omega_R(C/p) = \mathbf{1}) > \varepsilon$. Then

$$\begin{aligned}
\mathbb{P}(\sigma_{R,p}(C/p - k) = \mathbf{1}) &= \mathbb{P}(\omega_{R,p}(C/p) = \mathbf{1} \mid \omega_{R,p}(k) = \sigma)\\
&\geq \mathbb{P}(\omega_{R,p}(C/p) = \mathbf{1}) - k|G|p\\
&\geq \varepsilon/2
\end{aligned}$$

and the result follows from Lemma 5. □

Now we come to the key results of this section. For a configuration ω we define the *r-monotone closure* $\mathrm{CL}_r(\omega)$ to be the final configuration reached by ω under the r-monotone rule. We say ω is *r-monotone closed* if $\mathrm{CL}_r(\omega) = \omega$.

Theorem 9. *Suppose that G is a simple graph with a single loop at each vertex. If ω is a semi-stable configuration that r-threshold-p percolates then it $(r-1)$-monotone percolates. If ω_{inf} $(r-1)$-monotone percolates then ω r-threshold-p percolates.*

We prove this result with two lemmas.

Lemma 10. *Suppose that G is a simple graph with a single loop at each vertex and that ω is a semi-stable configuration. Then, for any C, $\mathbb{P}(\omega_{R,p}(t) \leq \mathrm{CL}_{r-1}(\omega)) = 1 - O(p)$ for all $t < C/p$.*

Proof. By Lemma 7 it is sufficient to show that, for any $(r-1)$-monotone closed configuration ω, and for any C

$$\mathbb{P}(\omega_{R,p}(t) \leq \omega) = 1 - O(p)$$

for all $t < C/p$.

As usual we can split the process into two stages: alternately applying the threshold rule and then a random perturbation. The probability that any two flips occur at the same time during the time 1 up to C/p is at most $|G|^2Cp = O(p)$. We show that, provided no two such flips occur, $\omega_{R,p}(t) \leq \omega$ for all t except at time steps when a cell s is flipped when $\omega_{R,p}(t) \leq \omega + s$, ($\omega + s$ denotes the configuration ω with cell s made active; i.e. the set of active cells of $\omega + s$ is $\omega \cup \{s\}$).

The proof is by induction on t. Suppose $\omega_{R,p}(t-1) \leq \omega + s$. Let ϕ be the configuration obtained by applying the simple threshold rule (once) to $\omega_{R,p}(t-1)$. We show that $\phi \leq \omega$.

The vertex s has at most $r-2$ neighbours in ω since ω is monotone-$(r-1)$ closed, and thus at most $r-2$ neighbours in $\omega_{R,p}(t-1) - s$. Since the graph has a single loop at each vertex, s has at most $r-1$ neighbours in $\omega_{R,p}(t-1)$. Thus the simple r-threshold rule removes s at this step. Moreover, the r-threshold rule does not create any other cell in ω^c at this step since for a cell to be created it must have r neighbours in $\omega_{R,p}(t)$ and thus at least $r-1$ in ω and thus the cell is in ω (since ω is $(r-1)$-monotone closed).

The configuration $\omega_{R,p}(t)$ is obtained from $\phi \leq \omega$ by randomly flipping at most one cell (since we assumed that no two get flipped at the same time). Hence $\omega_{R,p}(t) \leq \omega + s'$ and the $+s'$ only occurs if a cell s' was flipped at this time step, which occurs with probability at most $|G|p = O(p)$. ☐

Lemma 11. *Suppose that G is a simple graph with a single loop at each vertex and that ω is a semi-stable configuration. Then, for some $\varepsilon > 0$ and some C,*

$$\mathbb{P}(\exists t < C/p \colon \omega_{R,p}(t) \geq \mathrm{CL}_{r-1}(\omega_{\mathrm{inf}})) \geq \varepsilon$$

for all sufficiently small p.

Proof. Let s_i, $1 \le i \le m$, be a sequence of sites such that s_i has at least $r-1$ neighbours in $\omega_{\inf} + \sum_{j=1}^{i-1} s_j$ and $\omega_{\inf} + \sum_{j=1}^{m} s_j = \mathrm{CL}_{r-1}(\omega_{\inf})$. We actually prove a slightly stronger result: namely that there is a positive probability (bounded away from zero) that the model reaches a configuration containing $\mathrm{CL}_{r-1}(\omega_{\inf})$ by the time m flips have occurred.

First we show that if $\sigma \ge \omega + \sum_{j=1}^{i} s_j$ then $\sigma_T(t) \ge \omega_T(t) + \sum_{j=1}^{i} s_j$ for all t. Indeed we have $\sigma_T(t) \ge \omega_T(t)$ trivially. Also any point s_k, $k \le i$, has at least $r-1$ neighbours in $\omega_{\inf} + \sum_{j=1}^{k-1} s_j$ and s_k is a neighbour of itself (since G has loops). Thus in total it has at least r active neighbours in $\omega_{\inf} + \sum_{j=1}^{i} s_j$. Since $\omega_T(t) \ge \omega_{\inf}$, s_k stays active in $\sigma_T(t)$ for all $t \ge 0$.

Let t_i be the time of the ith flip and $t_0 = 0$. We show by induction on i that with positive probability either $\omega_{R,p}(t_i) \ge \omega_T(t_i) + \sum_{j=1}^{i} s_j$, or $\omega_{R,p}(t_j) = \mathbf{1}$ for some $j < i$. Indeed, suppose that $\omega_{R,p}(t_{i-1}) \ge \omega_T(t_{i-1}) + \sum_{j=1}^{i-1} s_j$ and $\omega_{R,p}(t_{i-1}) \ne \mathbf{1}$. Then set $s = s_i$ if s_i is not active in $\omega_{R,p}(t_{i-1})$, and let s be an arbitrary fixed inactive site of $\omega_{R,p}(t_{i-1})$ if s_i is active. Then

$$\mathbb{P}(\omega_{R,p}(t_i) \ge \omega_T(t_i) + \textstyle\sum_{j=1}^{i} s_j \mid \mathbf{1} \ne \omega_{R,p}(t_{i-1}) \ge \omega_T(t_{i-1}) + \sum_{j=1}^{i-1} s_j) \ge \frac{1}{|G|},$$

since above event happens if the cell s is the cell flipped at time t_i, and this happens with probability $\frac{1}{|G|}$.

A simple induction yields

$$\mathbb{P}(\omega_{R,p}(t_m) \ge \omega_{\inf} + \textstyle\sum_{j=1}^{m} s_j \text{ or } \exists t < t_m \colon \omega_{R,p}(t) = \mathbf{1}) \ge \frac{1}{|G|^m}.$$

In a time period $3/p$ the expected number of flips is $(3/p)|G|p = 3|G|$, and the variance of the number of flips is $(3/p)|G|p(1-p) \le 3|G|$. Thus by the Chebyshev inequality, the probability that we have less than m ($\le |G|$) flips is at most $3|G|/(2|G|)^2 \le \frac{3}{4}$. Thus with probability at least $\frac{1}{4}$ at least m flips occur in time C/p where $C = 3$.

Thus, we have that, with probability at least a constant, in time C/p at least m cells are flipped; and conditional on at least m cells being flipped the probability the model reaches a configuration containing $\mathrm{CL}_{r-1}(\omega_{\inf})$ is at least $1/|G|^{|G|}$. Thus

$$\mathbb{P}(\exists t < C/p \colon \omega_{R,p}(t) \ge \mathrm{CL}_{r-1}(\omega_{\inf})) \ge \frac{1}{4|G|^{|G|}}.$$

$\square$

Lemmas 10 and 11 together with Lemma 5 prove Theorem 9.

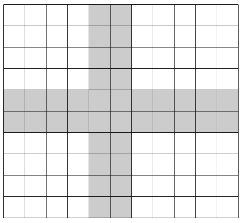

Figure 3: A small configuration on $\mathbb{T}_n^2$ that almost certainly reaches $\mathbf{1}$ before $\mathbf{0}$.

The important feature in the above proofs is that the graph has loops and that they have "weight" one. We needed the fact that if a vertex becomes active and it had almost enough neighbours already active then it will stay active. We also needed that a vertex with far too few active neighbours does not stay alive if it is randomly made active. Indeed this is expected: loops have no effect in the monotone model. In an extreme example if we weight each loop by $\Delta + 1$ then any cell that becomes active will stay active (unless it is flipped a second time) and thus even the empty state would threshold-p percolate. This was why we worked on r-threshold percolation not just majority percolation: we cannot just add loops to break ties.

Corollary 12. *There is a configuration ω of size $4n - 4$ on the $n \times n$ torus $\mathbb{T}_n^2$ (with single loops at each vertex), such that, in the random majority model $\omega_{R,p}$ with $r = 3$, $\omega_{R,p}$ reaches $\mathbf{1}$ before it reaches $\mathbf{0}$ with probability $1 - O(p)$.*

Proof. Let ω be the union of a band of height two horizontally around the torus and a band of width two vertically around the torus as shown in Figure 3. This configuration is static and 2-monotone percolates, so by Theorem 9 it majority-p percolates, say with parameters C and ε. However ω^c is 2-monotone closed and semi-stable, so by Lemma 10, $\omega^c_{R,p}(t) \leq \omega^c$ for all $t < C/p$ with probability $1 - \delta$ where $\delta = O(p)$. The probability that $\omega_{R,p}$ reaches $\mathbf{1}$ without ever reaching a configuration less than ω in a block of length C/p is at least $\varepsilon - \delta$. By monotonicity, the probability it reaches $\mathbf{1}$ before $\mathbf{0}$ is at least

$$\sum_{i=0}^{\infty}(1 - \varepsilon - \delta)^i(\varepsilon - \delta) = \tfrac{\varepsilon-\delta}{\varepsilon+\delta} \geq 1 - \tfrac{2\delta}{\varepsilon} = 1 - O(p)$$

as required. □

Theorem 9 is not quite an if and only if statement: there is a gap when the configuration is periodic. We consider some examples on the square torus with loops. One example of a configuration which does 2-monotone percolate but does not majority-p percolate is the checker board pattern. Indeed, if a cell appears it will not immediately die but will die one step later. Thus to leave this configuration two flips must occur at the same time. This implies that we need time $\Omega(1/p^2)$ to leave this pair of periodic configurations. Another example to consider is the following: let ω be the all active configuration apart from a band of width two around the torus which is checkered. Then ω does 2-monotone percolate and does majority-p percolate but ω_{inf} does not 2-monotone percolate. Hence, in general we cannot strengthen the previous theorem.

However, on specific graphs we may be able to say more. We illustrate this on the square torus with loops with $r = 3$, i.e., the majority model.

Observe that any periodic cell must have at least two periodic neighbours (if not it has two static neighbours in the same state and thus when it goes to that state it will stay there). Thus the periodic cells split into components with minimum degree 2. We will say that a component K of a configuration's periodic cells *threshold-p freezes* if there exist ε and C such that $\mathbb{P}(\omega_{R,p}(C/p)|_K = \mathbf{1}|_K) > \varepsilon$ for all sufficiently small p.

Lemma 13. *If ω is a semi-stable configuration on the square torus $\mathbb{T}_n^2$ and K is a component of ω's periodic states and K has a static active neighbour then K threshold-p freezes with $r = 3$ (i.e., the majority model).*

Proof. We prove the slightly stronger statement

$$\mathbb{P}(\exists t < C/p\colon \omega_{R,p}(t)|_K = \mathbf{1}|_K \text{ and all flips up to time } t \text{ make states active}) > \varepsilon, \qquad (1)$$

Note that if this event occurs, it requires at least two simultaneous flips to "unfreeze" K, since by monotonicity, a single flipped site p in K would revert to being active at the next time step, and neighbours of p in K need two adjacent elements of K to be inactive before they become inactive. Then as in the proof of Lemma 5, $\mathbb{P}(\omega_{R,p}(C'/p)|_K = \mathbf{1}|_K) > \varepsilon(1 - O(p))$ for any $C' \geq C$ and the lemma follows.

Let ω be a periodic configuration which is a counter example to (1) with the minimum number of periodic sites in K. For any site x, the probability that x is the site of the first flip is $\frac{1}{|G|}$, and with reasonable probability this occurs before time C/p. If flipping x to an active state results in a site of K becoming static active, then it is unlikely there will be another flip within time $2n^2$, during which the configuration will become semi-stable (by

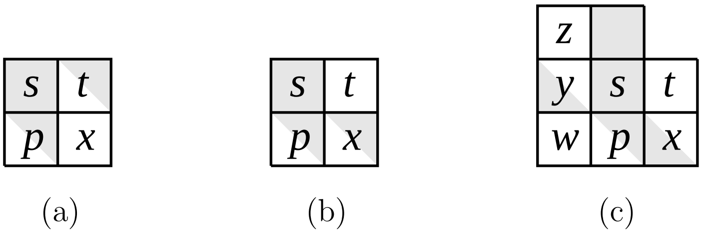

Figure 4: Freezing configurations. The shaded squares represent cells constantly in state 1, the unshaded squares cells constantly in state 0, the squares with the top right shaded alternate state 0101…, and the squares with the bottom left shaded alternate 1010… (i.e., they have the opposite phase of those shaded in the top right.)

Theorem 2) with the remaining sites of K either periodic or static active. (By monotonicity they cannot become static inactive.) Thus either the whole of K becomes static active, or we obtain a counterexample to (1) with a smaller number of periodic sites in K. Hence there are no sites which can be flipped to an active state causing sites in K to become static active.

Suppose s is a static active cell s neighbouring a periodic cell $p \in K$, and there is a neighbour t of s diagonal from p that is active when p is (so t is either periodic and in phase with p, or static active). Then if the fourth cell x of the 2×2 square containing s, t and p flips to an active state when the cell p is active all four cells will freeze (see Figure 4a). Note it does not matter whether the cell x is periodic or static. Similarly, if x is active when p is active, then flipping t to an active state when p is active freezes all four cells (Figure 4b). Thus we may assume x and t are both either static inactive or periodic and out of phase with p. Moreover, they cannot both be periodic, since t needs three inactive neighbours when it is active to change to the inactive state.

Since p has at least two periodic neighbours, at least one cell in position equivalent to x must be periodic, so at least one cell equivalent to t must be static inactive. Since s never has three inactive neighbours, the neighbour y opposite from t cannot be static inactive, and thus must be periodic, out of phase with p. The last neighbour of s must then be static active (see Figure 4c). Now flipping z when y is active freezes y. If the next flip occurs at w when p is active, then p will freeze. This occurs with reasonable probability, indeed within time C/p with reasonable probability, so we are done. □

Thus a periodic component which does not threshold-p freeze has no static active neighbour; i.e., all of its neighbours are static inactive. This implies that K has no vertex of degree two (since a vertex of degree two would have two static neighbours both of which are inactive and so when it goes inactive it would stay inactive). However the only components with minimum degree greater than two loop entirely around the torus.

4 Transition Time

In this section we show that, for very small p, the random majority model on the $n \times n$ square torus spends most of its time in either the completely full or the completely empty configuration. We also provide good bounds on the time it takes to transit between these two configurations. Let T_0^1 be the first time at which the configuration is all one starting from the all zero configuration. There are two possible quantities for the crossing time: namely $T_{\mathbb{E}} = \mathbb{E}(T_0^1)$ or $T_\varepsilon = \min\{t : \mathbb{P}(T_0^1 < t) > \varepsilon\}$ (i.e., either the expected time or the first time by which there is a "reasonable" chance of having crossed).

Lemma 14. *Let $T_{\mathbb{E}}$ and T_ε be defined as above. Then, for any fixed positive ε and ε' we have $T_{\mathbb{E}} = \Theta(T_\varepsilon)$ and $T_{\varepsilon'} = \Theta(T_\varepsilon)$ independently of p.*

Proof. It is immediate that $T_{\mathbb{E}} \geq (1-\varepsilon)T_\varepsilon$.

By monotonicity (Lemma 7) $\mathbb{P}(T_0^1 > kT_\varepsilon) \leq (1-\varepsilon)^k$. Thus, for any fixed $\varepsilon, \varepsilon' < 1$, $T_{\varepsilon'} = \Theta(T_\varepsilon)$. Furthermore

$$T_{\mathbb{E}} = \mathbb{E}(T_0^1) \leq T_{\frac{1}{2}} \sum_{k=0}^{\infty} \mathbb{P}(T_0^1 > kT_{\frac{1}{2}}) \leq T_{\frac{1}{2}} \sum_{k=0}^{\infty} \frac{1}{2^k} = 2T_{\frac{1}{2}}.$$

so $T_{\mathbb{E}} = \Theta(T_{\frac{1}{2}}) = \Theta(T_\varepsilon)$. □

We will only be interested in the order of the bounds so it does not matter which of these possibilities we use.

We prove two theorems in this section. The second of these is rather harder.

Theorem 15. *Let $\mathbb{T}_n^2$ be the $n \times n$ torus and let ω be any configuration. Then, for any $t = \Omega(1/p^3)$, we have*

$$\mathbb{P}(\omega_{R,p}(t) = \mathbf{1} \text{ or } \omega_{R,p}(t) = \mathbf{0}) = 1 - O(p).$$

Theorem 16. *Let $\mathbb{T}_n^2$ be the $n \times n$ torus and suppose $n \geq 3$. Then the expected time $T_{\mathbb{E}}$ to get from the all zero configuration to the all one configuration is $\Theta(1/p^{n+1})$.*

4.1 Proof of Theorem 15

We now show that all configurations reach one of the constant configurations ($\mathbf{0}$ or $\mathbf{1}$) quickly.

Lemma 17. *Suppose that ω is any configuration. Then there exist C and $\varepsilon > 0$ such that*

$$\mathbb{P}(\exists t < C/p^2 \colon \omega_{R,p}(t) = \mathbf{1} \text{ or } \mathbf{0}) > \varepsilon.$$

Proof. We may assume that the configuration is semi-stable. Indeed, by Theorem 2, we know that it reaches a semi-stable in time $2n^2$ if no flips occur in this period. This happens with probability $1 - O(p)$. Hence, if the theorem is true for all semi-stable configurations then it is true for all configurations.

Further we may assume that the configuration is static or that every state is periodic. This follows by Lemma 13 since all periodic components of a configuration which have a static cell threshold-p freeze (in either ω or ω^c).

If every state of the configuration is periodic then, since every state has period two, every 2×2 square always contains two active site and two inactive sites. Suppose that the two inactive sites flip at the same time-step: then this square freezes, and by Lemma 13 the whole configuration freezes. This occurs with positive probability in time C/p^2.

Since ω is static, Lemmas 10 and 11 imply that there exist C' and ε' such that $\mathbb{P}(\exists t < C'/p \colon \omega_{R,p}(t) = \mathrm{CL}_2(\omega)) > \varepsilon'$. Applying this result to the complement of ω, we see that with probability at least ε' there is a $t < C'/p$ such that $\omega_{R,p}(t)^c$ is 2-monotone closed. The result now follows from the next lemma. □

Lemma 18. *Suppose ω is a non-empty configuration and ω^c is 2-monotone closed. Then there exist C and $\varepsilon > 0$ such that*

$$\mathbb{P}(\exists t < C/p^2 \colon \omega_{R,p}(t) = \mathbf{1}) > \varepsilon \tag{2}$$

Proof. Suppose ω is a maximal configuration with ω^c 2-monotone closed and for which (2) does not hold. We aim to construct a larger configuration for which it also fails.

Run the 2-monotone model $\omega_M(t)$ starting with ω. We shall show that any cell in $\omega_M(t) \setminus \omega$ has degree four in $\omega_M(t+1)$. Suppose not. Let t be the first time at which it fails and let x be a cell for which this is not true. Let y, z be two of the neighbours of x that were active at time $t-1$. Assume y is in ω. Since ω^c is 2-monotone closed, y can have at most one inactive neighbour in ω. Thus y has three active neighbours in ω, and so four in $\omega_M(t)$ (since x is active). If y is not in ω, then by our choice of t, all four of its neighbours are active at time t. The same argument applies to z, so both y and z have four active neighbours at time t. Hence the remaining neighbours of x have at least two active neighbours at time t (x and a neighbour of either y or z), and so become active at time $t+1$. Since $\omega_M(t)$ eventually equals $\mathrm{CL}_2(\omega)$, both $\mathrm{CL}_2(\omega)$ and its complement are 2-monotone closed.

Hence, we may assume that both ω and ω^c are monotone-2 closed, i.e., every cell has at least three neighbours in the same state. This implies that ω is a band around the torus of width at least two. To complete the proof we show that such a band satisfies (2).

Suppose that ω is such a band. Then any single flip will be immediately reversed by the simple majority rule. Suppose that the first time at which two flips occur the two cells flipped are adjacent to each other and to the band. This gives a stable configuration, with monotone-2 closure the band one wider. Then, (2) is true for ω if and only if it is true for the band one wider and the result follows. $\square$

Proof of Theorem 15. First note that, if we do not have at least 4 flips in a time interval of length 4 in the time up until T, then $\mathbf{1}_{R,p}(t) = \mathbf{1}$ for all $t < T$ unless there was a flip at time $t-1$ or $t-2$. Indeed this is just the observation that any state with at most three inactive sites returns to the all one configuration in two steps under the simple majority rule. The corresponding result holds if we replace the $\mathbf{1}$ configuration by the all zero configuration $\mathbf{0}$.

Now fix some large T and pick any ω. Suppose that no cell was flipped at time $T-2$ or $T-1$: this occurs with probability $1 - O(p)$. Let T_0 be the last time before T at which a time interval of length four contained four flips (or 0 if no such interval exists). Then, by the above we know that if $\omega_{R,p}(t) = \mathbf{1}$ for any $T_0 < t < T$ then $\omega_{R,p}(T) = \mathbf{1}$.

Suppose that $T - T_0 > 1/p^3$: this occurs with probability $1 - O(p)$. Then

$$\mathbb{P}(\exists t,\, T_0 < t < T \colon \omega_{R,p}(t) = \mathbf{1} \text{ or } \omega_{R,p}(t) = \mathbf{0}) = 1 - O(p)$$

(by Lemma 17) and thus

$$\mathbb{P}(\omega_{R,p}(T) = \mathbf{1} \text{ or } \omega_{R,p}(t) = \mathbf{0}) = 1 - O(p)$$

as required. $\square$

4.2 Upper Bound in Theorem 16

We construct a reasonable path from $\mathbf{0}$ to $\mathbf{1}$. First we construct a path from $\mathbf{0}$ to a band of width two around the torus.

For $r \neq 1$ let σ_r denote a configuration consisting of a $2 \times r$ rectangle of active states. Since $r \neq 1$, the configuration σ_r is majority stable. Thus the probability that we leave configuration σ_r at the next time step is at most Np, where $N = n^2$ is the number of sites of the torus.

From the state σ_r the probability that the two cells on the short end of the rectangle both flip in the next time step is p^2. Thus the probability that the configuration becomes σ_{r+1} at the next time step is at least p^2. However, in the case $r = n - 1$ (i.e., σ_r is a band of width two around the torus minus a single column) this can be improved. Indeed, if we add either of the two missing squares then we get a configuration which will not decrease unless a flip occurs (i.e., it's stable with respect to decreases). Define $\sigma_{n-\frac{1}{2}}$ to be this configuration (i.e., σ_n minus a single square). The probability the configuration moves from σ_{n-1} to $\sigma_{n-\frac{1}{2}}$ in a single time step is at least p. If no cell is flipped then $\sigma_{n-\frac{1}{2}}$ moves to σ_n in a single time step. This occurs with probability at least $1-Np$. Finally the probability that we leave σ_{n-1} in a single time step is at most Np (since the configuration is stable we must have at least one flip). From the $\mathbf{0}$ configuration the probability that in the next time-step we move to σ_2 is at least p^4.

Suppose the initial configuration is σ_r. Then the probability that the configuration first changes within time C/p and that the new configuration is σ_s is at least

$$\begin{aligned}\sum_{i=0}^{C/p} \mathbb{P}(\sigma_r(1) = \sigma_s)\mathbb{P}(\sigma_r(1) = \sigma_r)^i &\geq \mathbb{P}(\sigma_r(1) = \sigma_s)\sum_{i=0}^{C/p}(1 - Np)^i \\ &\geq \mathbb{P}(\sigma_r(1) = \sigma_s)\frac{1-(1-Np)^{C/p}}{1-(1-Np)} \\ &= \mathbb{P}(\sigma_r(1) = \sigma_s)\frac{1-(1-Np)^{C/p}}{Np} \\ &\geq \frac{\mathbb{P}(\sigma_r(1) = \sigma_s)}{2Np}\end{aligned}$$

for large C. In particular, during a time interval of length C/p the probability the configuration goes from σ_r to σ_{r+1} is at least $p/2N$, the probability it goes from $\mathbf{0}$ to σ_2 is at least $p^3/2N$, the probability that it goes from σ_{n-1} to $\sigma_{n-\frac{1}{2}}$ is at least $1/2N$. We also know that the probability that it goes from $\sigma_{n-\frac{1}{2}}$ to σ_n in a single time step is at least $1 - Np$. Combining all of these the probability that it goes from $\mathbf{0}$ to σ_n within time nC/p is $\Omega(p^n)$. Call this probability p_1.

Thus, by Lemma 7 (monotonicity) the probability the configuration $\mathbf{0}$ reaches a configuration containing σ_n in time Cn/p^{n+1} is at least

$$1-(1-p_1)^{1/p^n} \geq \varepsilon$$

for some $\varepsilon > 0$ provided that p is sufficiently small. However, by Lemma 18, the configuration σ_n has a positive chance (independent of p), ε' say, of reaching the all one configuration in time C/p^2. Thus, with probability $\varepsilon\varepsilon'$ the zero configuration $\mathbf{0}$ reaches the one configuration $\mathbf{1}$ in time $O(1/p^{n+1})$: i.e., $T_{\varepsilon\varepsilon'} = O(1/p^{n+1})$ and hence the transition time (any of the definitions) of the all one configuration from the all zero configuration is $O(1/p^{n+1})$.

It is interesting to observe that there is a second completely different route that also takes time $O(1/p^{n+1})$. The first step is a *much* less likely step: namely a whole row flips at the same time: explicitly that is all n cells in a row flip simultaneously. We build up from here in a similar way to the above argument. However a little care is needed since this single row can be "destroyed" by a single flip.

Since this is not necessary for the proof of the theorem we only give a brief sketch of the argument. Let σ_0 denote a complete row: i.e., a band completely around the torus and let σ_r denote the row together with r adjacent cells in one of the neighbouring rows. Then σ_i is stable except when $r = 1$. The probability that we go from configuration $\mathbf{0}$ to σ_0 in a single step is $\Theta(p^n)$. The probability that we go from σ_0 to σ_2 before visiting any other state is $p^2/Np = \Omega(p)$. Finally, for $r \geq 2$, the probability that we go from σ_r to σ_{r+1} before any other state is $p/Np = \Omega(1)$. Using a similar argument to before we see that the configuration has a positive chance of reaching a two wide band in time $O(1/p^{n+1})$. Since we know we have a positive chance of reaching the $\mathbf{1}$ configuration from the two wide band in time $\Omega(1/p^2)$ the result follows.

4.3 Lower Bound

We need to show that the expected time to flip from the all zero configuration to the all one configuration is $\Omega(1/p^{n+1})$. Let $\omega(t) = \mathbf{0}_{R,p}(t)$ denote the configuration under the random majority rule at time t starting from the all zero configuration.

First we have several definitions. A *row* is a set of cells of the form $\{(x,y) : y = a\}$ and a *column* a set of the form $\{(x,y) : x = a\}$. We shall write R_a (respectively C_a) for these sets. An *interval* of rows is a collection of rows $\{R_a : a_0 \leq a \leq a_1\}$ and similarly with an interval of columns. A row or column is *active* if it contains an active cell. An interval of rows is active if every one of its rows is active. An interval of rows is inactive if none of its rows is active. An active interval is *maximal* if it is not a proper subset of any other active interval.

For a configuration ω lets α denote the number of active rows in ω; let β denote the number of maximal inactive row intervals in ω; finally define χ_c to be one if ω has at most one active *column* and zero otherwise (i.e., χ_c is one if ω is a subset of a single column and zero otherwise). Note that the number of maximal inactive row intervals is precisely the number of maximal row active intervals unless all rows are active or all rows are inactive (this exception is an irritation).

Now, for a configuration ω define its *row rank* to be $\alpha + \beta - \chi_c$. Define the *column rank* similarly. Finally define the *rank* of ω to be the maximum of the row rank and the column rank.

Next we prove some simple lemmas about the rank.

Lemma 19. *Suppose ω and σ are configurations with $\sigma \leq \omega$. Then* $\operatorname{rank}(\sigma) \leq \operatorname{rank}(\omega)$.

Proof. Since $\sigma \leq \omega$ the number of active rows of σ cannot be more than that of ω, i.e., $\alpha(\sigma) \leq \alpha(\omega)$. Similarly $-\chi_c(\sigma) \leq -\chi_c(\omega)$. The number of inactive row intervals can be larger in σ: however the number of additional inactive row intervals is at most the number of additional inactive rows, i.e., the increase in β is smaller than the decrease in α. Thus the row rank of σ is at most that of ω. Since the same is true for the column rank, the result follows. $\square$

Lemma 20. *Suppose ω is a configuration and let $\bar{\omega}$ be the monotone-2 closure of ω. Then* $\operatorname{rank}(\bar{\omega}) = \operatorname{rank}(\omega)$.

Proof. Since $\bar{\omega}$ contains ω we have $\operatorname{rank}(\bar{\omega}) \geq \operatorname{rank}(\omega)$. If ω is a subset of a single column then so is $\bar{\omega}$. Hence $\chi_c(\bar{\omega}) = \chi_c(\omega)$.

Consider an inactive row interval of length at least two. Then these rows form a set of inactive cells, in which each cell has at least three neighbours in the set. Thus no cell in this interval ever becomes active (under the 2-monotone rule) and hence the interval stays inactive. A maximal inactive interval of length one (i.e., a single row) can become active, but, if it does, it reduces the number of maximal inactive intervals by one. Thus the row rank does not change. Since the same is true for the column rank the result follows. $\square$

Since the 2-monotone model dominates the majority model, the preceding two lemmas imply that the majority model never increases the rank. Moreover, if ω is periodic, then both semi-stable states of ω have the same rank.

Lemma 21. *Suppose ω is a configuration and that* $\text{rank}(\omega) < n$*. Then ω^c majority-p percolates.*

Proof. Since $\text{rank}(\omega) < n$ then either $\alpha(\omega) < n$ or $\chi_c(\omega) = 1$: i.e., either there is an inactive row and an inactive column, or ω is a subset of a single active row or column (this includes the case when ω is empty).

In the first case, ω^c contains a row and a column of active sites. In the second case, ω^c contains the complement of a row. In both cases, ω^c contains a majority stable, 2-monotone percolating set. Thus by Theorem 9 and monotonicity, ω^c majority-p percolates. $\square$

Next we have some lemmas about how the rank changes with time.

Lemma 22. *Suppose $n \geq 3$, ω is a semi-stable configuration on $\mathbb{T}_n^2$ and that $\sigma = \omega_{R,p}(T)$ is also semi-stable. Further suppose that* $\text{rank}(\sigma) > \text{rank}(\omega)$ *and that k cells are flipped by time T. Then*

$$k \geq \text{rank}(\sigma) - \text{rank}(\omega) + 1.$$

Proof. We define a new process by flipping each cell at the same times as it is flipped in the original process except when it is already active or the flip occurs after time T. Let σ' be a semi-stable configuration that is eventually reached in the new process, say at time $T' \geq T$.

Then by monotonicity of the majority rule, either $\sigma' \geq \sigma$ or $\sigma'_T(1) \geq \sigma$ and either $\sigma' > \omega$ or $\sigma'_T(1) > \omega$ (we may need the second cases if ω or σ is periodic). Let σ'' be the one that is bigger than ω. The number of flips for σ'' is smaller than the number for σ and $\text{rank}(\sigma'') \geq \text{rank}(\sigma)$. Thus, if the lemma is true for σ'' it is true for σ. From this point on we replace σ by σ'' and thus have that $\sigma > \omega$.

Now, since an empty row is majority stable, every newly active row in σ must have a cell in it that flipped.

Suppose that some newly active row R contains two active sites. We claim that either two flips occurred in this row, or that both its neighbours had two active cells before it first had two active cells. Consider the first time t at which R contains two active sites, say r_1 and r_2. If R had no active cells at time $t-1$ then both r_1 and r_2 must have been created by flips. If R had a single active cell then R must have contained a flip at some time before $t-1$. Hence if either r_1 or r_2 were created by a flip at time t we have the two flips we claimed. Otherwise both of these cells must be present by the simple majority rule. This means that both r_1 and r_2 must have had three active neighbours (possibly including themselves) at times $t-1$. Since both r_1 and r_2 had at most one active neighbour in R (since there was only one active cell in R) they both must have had active neighbours in both adjacent rows; hence both these rows contained two active neighbours at time $t-1$.

Now suppose that R is a newly active row adjacent to an inactive row S. Call such a pair a *change pair*. Since σ is semi-stable, R must contain at least two active sites. Then, by the above, either two flips occurred in R, or S had two active sites before R did, but then applying the claim to S we see that two flips must have occurred in S.

Combining these two observations we see that every change pair R, S requires at least one extra flip in addition to the one we know occurs in R anyway. Thus, the total number of flips is at least the number of newly active rows plus the number of change pairs.

We split into several cases. First, suppose that $\chi_c(\omega) \leq \chi_c(\sigma)$. Fix a maximal inactive row interval I of ω. I contributes 1 to the rank of ω since it is a maximal inactive interval. Consider the contribution to the rank of σ coming from rows in I, and suppose that this is greater than 1. Thus, in particular, σ must have an active row in I. The contribution to the rank of the rows in I is the number of newly active rows in I plus the number of maximal inactive row intervals in σ which are contained in I. Every maximal inactive interval $J \subset I$ can be associated with a change pair; indeed, unless J shares an endpoint with I, we can associate two change pairs with J. Hence the total number of flips is at least the contribution to the rank of I. However I previously contributed 1 to the rank of ω so the number of flips is at least the change in rank from I plus one.

Otherwise we may assume $\chi_c(\omega) = 1$ and $\chi_c(\sigma) = 0$. Then, since ω is semi-stable, either it is empty or a complete column. If ω is empty then every maximal σ-inactive interval J can be identified with two change pairs. Hence, provided σ has an inactive interval, the number of flips is at least $\alpha(\sigma) + 2\beta(\sigma) \geq \alpha(\sigma) + \beta(\sigma) + 1$. Since $\operatorname{rank}(\sigma) \leq \alpha(\sigma) + \beta(\sigma)$ this case is proved. If σ does not have an inactive row interval then at least n cells must be flipped. Since $\chi_c(\sigma) = 0$ and σ is semi-stable then, some row contains two cells. This implies that at least $n+1$ are flipped and the change in rank is bounded by n.

If ω is a single column then it has rank $n-1$. For the rank of σ to be larger it must contain an active site in another column. Since σ is semi-stable either some such a column contains two cells in which case at least two flips occur by the argument above; or σ contains all the cells in a row which requires $n-1 \geq 2$ flips. □

Note that we do need the semi-stable condition. However, by Theorem 2, we know that if no cell was flipped "shortly" before time T, then the configuration is semi-stable. Let T_{settle} be the maximum time it takes for a state to stop changing in the (deterministic) majority model. Theorem 2 then implies $T_{\text{settle}} \leq 2n^2$. Thus, if no cell is flipped between time $T - T_{\text{settle}}$ and time T then Lemma 22 applies.

Whilst Lemma 22 is true for any time period, for longer time periods it can often be strengthened. In particular the state stabilising "costs" something. Thus we divide the time period up into several small blocks with gaps between them big enough for the state to stabilise. This is important because the probability that several flips occur close together is very low.

Suppose that the first M flips of the random majority process occur at times $1 \leq t_1 \leq t_2 \leq \cdots \leq t_M$ (they need not be distinct).

Lemma 23. *Fix $t' < t''$ and suppose $\omega(t')$ and $\omega(t'')$ are both semi-stable. If $\operatorname{rank}(\omega(t'')) - \operatorname{rank}(\omega(t')) \geq k$ then at least k flips during time t' to t'' occur within times T_{settle} of the preceding flip.*

Proof. We can find intervals of time, *blocks*, $B_i \subset \{t', t'+1, \dots, t''\}$ such that: the blocks are disjoint, every t_j with $t' \leq t_j \leq t''$ is contained in some block, and $\max B_i > \max_{\{t_j \in B_i\}} t_j + T_{\text{settle}}$ except possibly for the last block where we allow $\max B_j = t''$. The last condition ensures that the configuration is semi-stable at the end of any block (since $\omega(t'')$ is semi-stable it is true for the last block too), and, hence, that Lemma 22 applies. We may further assume that all such blocks are minimal: i.e., we know that every flip in a block after the first flip in that block occurs within time T_{settle} of the previous flip.

Suppose that the rank increases by k_1 during a block B. Then at least $k_1 + 1$ flips occur during B and each one after the first must be within time T_{settle} of the previous one (by the definition of block).

Since the rank increases by at least k between t' and t'' we can find at most k blocks such that the rank increases by at least k during the union of these blocks. Applying the preceding argument to each of them and summing gives the result. $\square$

For a configuration ω define its *stable rank*, $\operatorname{rank}_{\text{stab}}(\omega)$ to be the rank of either of its semi-stable configurations (since the rank does not increase under the simple majority rule this is well defined).

Lemma 24. *Fix a configuration ρ with $\operatorname{rank}_{\text{stab}}(\rho) < n$. Let E be the event that after $N = n^2$ flips the configuration is $\mathbf{0}$ and that no two of these flips occur within time T_{settle} of each other and the first flip occurs after time T_{settle}. Then there exists an $\varepsilon > 0$ such that $\mathbb{P}(E) \geq \varepsilon$ for all sufficiently small p.*

Proof. Let $\sigma = \rho(T_{\text{settle}})$. Since the first flip occurs after this time the configuration σ is semi-stable and has rank less than n. Since σ^c majority-p percolates (Lemma 21) the probability that we reach $\mathbf{0}$ by the Nth flip is at least ε for some $\varepsilon > 0$. The probability that none of the first N flips are within T_{settle} of each other is $1-O(p)$. Thus the probability that both of these occur is at least $\varepsilon - O(p)$. Moreover, if these both occur the configuration will stay in the $\mathbf{0}$ configuration after it first reaches it (since no two flips occur at the same time). The result follows. ☐

Call the event E a *decay* event. Define an *up event* to be the event that either the next flip occurs within time T_{settle} or two flips from the first N occur within time T_{settle} of each other.

Finally we prove the lower bound on the transition time. Suppose $t_0, t_1, \dots$ are the (not necessarily distinct) flip times. Define a subsequence $t_{n_0}, t_{n_1}, \dots$ so that n_i is the first $n > n_{i-1}$ such that $t_n - t_{n-1} < T_{\text{settle}}$, or $n_i = n_{i-1} + N$ if no such flip occurs within N flips of $t_{n_{i-1}}$. Then each $(t_{n_{i-1}}, t_{n_i}]$ is independent and gives rise to a decay event (with probability $\geq \varepsilon$) and an up event (with probability $\leq NT_{\text{settle}}p$). The probability of obtaining n up events before any decay event is then $\leq (NT_{\text{settle}}p/\varepsilon)^n$.

Hence, starting from the $\mathbf{0}$ configuration the probability that we reach the $\mathbf{1}$ configuration before the next decay event is at most $(NT_{\text{settle}}p/\varepsilon)^n$, (since we need to have n up events before any decay event). Thus, the probability that we reach the $\mathbf{1}$ configuration within the first M flips is at most $M(NT_{\text{settle}}p/\varepsilon)^n$ so if this probability is not $o(1)$ then $M = \Omega(1/p^n)$. Finally, with probability $1 - o(1)$ the Mth flip occurs after time $M/(2|G|p) = \Omega(1/p^{n+1})$ and the proof is complete. ☐

Let the *crossing time* be $T_0^1 - \max_{t<T_0^1}\{\mathbf{0}_{R,p}(t) = \mathbf{0}\}$; i.e., the time it takes to reach $\mathbf{1}$ from the last time it leaves $\mathbf{0}$. Although the expected transition time is $\Theta(1/p^{n+1})$ the crossing time is $O(-(\log p)/p^2)$. Indeed, by Lemma 17 the probability that the configuration reaches $\mathbf{0}$ or $\mathbf{1}$ in a time interval of length C/p^2 is at least ε for some $C, \varepsilon > 0$. Hence the probability that the configuration does not reach $\mathbf{0}$ or $\mathbf{1}$ in an interval of length kC/p^2 is at most $(1-\varepsilon)^k$. There are at most $1/p^{n+1}$ possible such intervals in time up to $1/p^{n+1}$. Hence, with probability $1-(1-\varepsilon)^k/p^{n+1}$ the configuration reaches $\mathbf{0}$ or $\mathbf{1}$ in *every* interval of length kC/p^2. This probability tends to one if k is some constant times $\log(1/p)$.

5 Small Percolating Sets

In the previous section we saw that one way the random majority model can cross from **0** to **1** is for many cells to flip in the same time step taking the configuration to one which reaches **1** easily. In the second section we gave a criterion for a configuration to reach **1** quickly: namely if the configuration monotone percolated. The unlikeliness of this configuration spontaneously appearing is order p^s where s is the number of active sites in the configuration; this motivates the study of small monotone percolating or threshold percolating sets. In this section we prove some bounds on their size.

First we prove a lower bound.

Theorem 25. *Let G be a regular graph of degree Δ. Suppose that a state ω percolates in the r-monotone model. Then ω has at most $\frac{\Delta}{2r}|G|$ inactive vertices.*

Proof. Fix a state ω. Let G_0 be the subgraph induced by the inactive vertices. If G_0 has a subgraph H with minimum degree $\Delta - r + 1$ then ω does not percolate (no vertex in H will ever appear). So every subgraph of G_0 has a vertex of degree at most $\Delta - r$. Thus $e(G_0) \leq (\Delta - r)|G_0|$. Summing over the degrees of vertices in G_0 we see that $e(G_0, G_0^c) \geq \Delta|G_0| - 2e(G_0) \geq (2r - \Delta)|G_0|$. However counting from G_0^c we have $e(G_0, G_0^c) \leq \Delta|G_0^c|$. Combining we have $|G_0^c|\Delta \geq (2r - \Delta)|G_0|$. Thus $\Delta|G| \geq 2r|G_0|$ or equivalently $|G_0| \leq \frac{\Delta}{2r}|G|$. □

Note that the trivial bound this gives if $r \leq \Delta/2$ is often correct as a leading term: for example on the square torus the state which is active on the diagonal monotone-2 percolates and has size $\sqrt{|G|} = o(|G|)$.

The following theorem shows that it is also tight for $r > \Delta/2$ even on the lattice $\mathbb{Z}_n^d$. For these bounds it does not matter whether we work on the d dimensional torus $\mathbb{T}_n^d = \mathbb{Z}_n^d$ or on a subset of $\mathbb{Z}^d$. Indeed suppose we have a percolating subset of $\mathbb{T}_n^d$ then we can place it in $\mathbb{Z}_n^d$ and make all the boundary sites active. Since the boundary has size $O(n^{d-1})$ this does not effect our bounds.

Theorem 26. *Let $\mathbb{Z}_n^d$ be the cubic lattice in d dimensions. For any r, $d \leq r \leq 2d$, there exists a percolating set for the r-monotone model of size $(1 - \frac{d}{r})n^d + O(n^{d-1})$.*

Proof. It is enough to find such a percolating set on $\mathbb{T}_n^d$ when n is even and divisible by r. Indeed, we may increase n to $n' > n$ with these properties. Then taking any percolating set on $\mathbb{T}^d_{n'}$, intersecting with $\mathbb{Z}^d_n$ and making all the boundary sites active, we obtain a percolating set for $\mathbb{T}^d_n$ with at most $2dn^{d-1}$ more active sites.

Let A_t be the set of points $x = (x_1, x_2, \dots, x_d) \in \mathbb{T}^d_n$ for which

$$x_1 + x_2 + \cdots + x_d \equiv t \mod n$$

and write $A_{i \bmod 2} = \bigcup_{t \equiv i \bmod 2} A_t$. The sets $A_{0 \bmod 2}$ and $A_{1 \bmod 2}$ form complementary d-dimensional checker board patterns in $\mathbb{T}^d_n$.

Assume that r is even and define

$$f(x) = x_1 + 2x_2 + \cdots + (\tfrac{r}{2} - 1)x_{r/2-1} \bmod \tfrac{r}{2}$$

$$S = \{x \in \mathbb{Z}^d_n : f(x) \in \{1, 2, \dots, r - d\}\}.$$

The numbers $1, 2, \dots, \frac{r}{2} - 1$ run over all non-zero residue classes mod $\frac{r}{2}$. Thus if $x \notin S$ then there are exactly $r - d$ coordinates i such that if we increase x_i by one then we obtain an element of S. Similarly there are $r - d$ coordinates such that if we decrease x_i by one then we obtain an element of S. Thus if $x \in A_t \setminus S$ then x has $r - d$ neighbours in $A_{t+1} \cap S$ and $r - d$ neighbours in $A_{t-1} \cap S$. If $x \in S$ then by a similar argument, x has precisely $2(\frac{r}{2} - (r - d)) = 2d - r$ neighbours that are not in S, and hence r neighbours that are in S (not including x itself). In particular, if $x \in A_{i \bmod 2} \cap S$ then x has r neighbours in $A_{(i+1) \bmod 2} \cap S$.

Now assume r is odd. In this case define

$$f(x) = 2x_1 + 4x_2 + 6x_3 + \cdots + (r - 1)x_{(r-1)/2} \mod r$$

$$S = \{x \in \mathbb{Z}^d_n : f(x) \in \{1, 2, \dots, 2(r - d)\}\}.$$

If $x \notin S$ then there are exactly $r-d$ indices i, $1 \le i \le (r-1)/2$, such that $f(x)+2i \bmod r$ lies in $\{1, 2, \dots, 2(r-d)\}$ and $r-d$ indices such that $f(x)-2i \bmod r$ lies in $\{1, 2, \dots, 2(r-d)\}$. Hence if $x \in A_t \setminus S$ then x has $r-d$ neighbours in $A_{t+1} \cap S$ and $r-d$ neighbours in $A_{t-1} \cap S$. If $x \in S$, then $\pm 2, \pm 4, \dots, \pm(r - 1)$ runs over all the non-zero residue classes mod r, so there are exactly $r - 2(r - d) = 2d - r$ neighbours of x that do not lie in S. Once again, if $x \in A_{i \bmod 2} \cap S$ then x has r neighbours in $A_{(i+1) \bmod 2} \cap S$.

Now in both the r even and r odd cases, consider the sets

$$T_t = (A_{t \bmod 2} \cap S) \cup \bigcup_{i=0}^{t} A_i.$$

If $x \in A_{(t+1) \bmod 2} \cap S$ then x has at least r neighbours in $A_{t \bmod 2} \cap S \subseteq T_t$, so become active at the next time step. If $x \in A_{t+1} \setminus S$, then x has at least $r-d$ neighbours in $A_{t+2} \cap S \subseteq A_{t \bmod 2} \cap S$ and d neighbours in A_t, so these points also become active. Thus every site in T_{t+1} becomes active at the next step. Hence by induction, if we start with all the cells in T_0 active, then after n steps all cells in $T_n = \mathbb{T}_n^d$ will be active. In other words, T_0 r-monotone percolates.

Finally we estimate $|T_0|$. Fixing $x_2, \dots, x_d$, there are at most $\lceil n\frac{2(r-d)}{r}\rceil$ choices of x_1 giving a point $x \in S$ in both the r even and r odd cases. Thus $|S| \le \frac{2(r-d)}{r}n^d + n^{d-1}$. Since the condition $x \in S$ never depends on x_d, exactly half of these points lie in $A_{0 \bmod 2}$. Finally, for any $0 \le t \le n-1$, we have $|A_t| = n^{d-1}$, and thus $|T_0| \le (1-\frac{d}{r})n^d + O(n^{d-1})$. □

Since $\mathbb{Z}^d$ is $2d$-regular Theorem 25 gives a lower bound on the size of a r-monotone percolating set of $n^d(1-\frac{d}{r})$. Thus Theorem 26 is best possible up to the $O(n^{d-1})$ term.

The sets defined in Theorem 26 have a special property. Almost no site has any neighbours in the set itself. Indeed, this is required as can be seen from the proof of Theorem 25. This may lead one to believe that the bound for r-threshold percolation must be higher. This is not always true. Indeed the next theorem shows that for $r = d$ and $r = d+1$ the bound in Theorem 25 can still be attained for threshold percolation.

Theorem 27. *For $r \in \{d, d+1\}$, there exist r-threshold percolating sets in $\mathbb{T}_n^d$ of size $(1-\frac{d}{r})n^d + O(n^{d-1})$.*

Proof. For $r = d$ the set $T = A_0 \cup A_1$ will d-threshold percolate. This can be seen because T contains a d-monotone percolating set (e.g. A_0) and every point in T has d neighbours in T (so no point of T ever gets removed). This set has size $O(n^{d-1})$.

For $r = d+1$ the set T_1 defined in Theorem 26 ($(d+1)$-threshold) percolates. As in the proof of Theorem 26 if the sites in T_t are active at time t then the sites in $A_{(t+1) \bmod 2} \cap S$ and $A_{t+1} \setminus S$ become active at time $t+1$. But each $x \in A_i$, $0 \le i \le t$, has d active neighbours in either A_{i-1} or A_{i+1}, plus itself, so remains active. Thus all sites in T_{t+1} are active at time $t+1$. Hence T_1 majority percolates. □

For the torus $\mathbb{T}_n^2$, the configuration T_1 consists of a checker board pattern with one in every three columns removed (plus a couple of diagonal lines $A_0 \cup A_1$); see Figure 5. In three dimensions T_1 consists of alternating empty planes and planes that are checker board patterned (plus two “diagonal” planes $A_0 \cup A_1$).

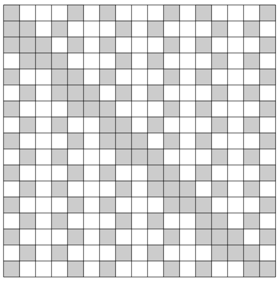

Figure 5: A small configuration that majority percolates in $\mathbb{T}_n^2$.

For $r > d+1$ the situation is unclear, but one can prove using arguments similar to Theorem 25 that for $r = 2d$ one needs density at least $1 - \frac{1}{2d}$ in any r-threshold percolating set, which is strictly larger than the bound of $1 - \frac{d}{2d} = \frac{1}{2}$ of Theorem 25 when $d > 1$.

6 Conjectures and Open Questions

At the moment, the most important question concerning this random cellular automaton is the following: Does there exist a critical constant $p_c > 0$ *independent of* n such that for $p < p_c$ the model spends almost all of its time near $\mathbf{0}$ or $\mathbf{1}$, and for $p > p_c$ the model spends most of its time with about half of its cells active. As expected this is much simpler for the mean field model and an exact result $p_c = 7/30$ can be proved (see [5]). That result, and this paper strongly suggest an affirmative answer to this question. Based on this and computer simulations we conjecture the following.

Conjecture. *There is a critical probability* p_c.

This conjecture seems to be very difficult. Numerically, it appears to exist with a value of approximately 0.134. In fact, even proving that there is some constant $\varepsilon > 0$ below which the former behaviour occurs seems difficult.

There are some other questions which appear to be rather simpler.

Question. *Suppose that* G *is a (finite)* Δ*-regular graph and* ω *a configuration. What is the maximum time* ω *can take to stabilise in the (deterministic) threshold model?*

In Theorem 2 Poljak and Turzík showed that the time to stabilise, was at most $\frac{1}{2}\sum_{v \in G}(d(v) - 1)$. The question is how tight is this bound. There are easy constructions to show that their *proof* is tight, but these do not show that the result itself is tight.

Our last question concerns threshold percolation on the torus $\mathbb{T}^d$.

Question. *Suppose that $r > d+1$ and that ω is an r-threshold percolating set in $\mathbb{T}_n^d$. How small can ω be?*

The results in section 5 provided some bounds but the complete picture is still unclear.

References

[1] J. Adler, *Bootstrap percolation*, Physica A **171** (1991) 453–470.

[2] J. Adler, van Enter, and J.A. Duarte, *Finite-size effects for some bootstrap percolation models*, J. Statist. Phys., **60** (1990) 322–332.

[3] M. Aizenman and Lebowitz, *Metastability effects in bootstrap percolation*, J. Phys. A **21** (1988) 3801–3813.

[4] E. Andjel, T. Mountford, R. H. Schonmann, *Equivalence of exponential decay rates for bootstrap-percolation-like cellular automata.* Ann. Inst. H. Poincar Probab. Statist. **31** (1995), no. 1, 1325.

[5] P. Balister, B. Bollobás and R. Kozma *Mean field models of probabilistic cellular automata*, submitted.

[6] R. Cerf and E.N. Cirillo, *Finite size scaling in three-dimensional bootstrap percolation*. Ann. Probab., **27** (1999) 1837–1850.

[7] H. Chaté, A. Lemaitre, Ph. Marcq, P. Manneville, *Non-trivial collective behavior in extensively chaotic dynamical systems: an update*. Physica A, **224** (1996) 447–457.

[8] H. Chaté and P. Manneville, *Evidence of collective behavior in cellular automata*, Europhys. Lett., **14** (1991) 409–413.

[9] H. Chaté and P. Manneville, *Collective behaviors in spatially extended systems with local interactions and synchronous updating*, Progr. Theoretical Phys., **87** (1992) 1–60.

[10] R. Durrett and J. E. Steif *Fixation Results for Threshold Voter Systems* Annals of Probability, **21**, (1993), 232–247

[11] P. Gács, *Reliable cellular automata with self-organization.* J. Statist. Phys. **103** (2001), 45–267.

[12] J. Gravener and E. McDonald, *Bootstrap percolation in a polluted environment*, J. of Statist. Phys., **87** (1997) 915–927.

[13] E. Goles, J. Olivos. *Comportement périodique des fonctions à seuil binaires et applications.* Discrete Appl. Math. **3** (1981), 93–105.

[14] L. F. Gray, *The behavior of processes with statistical mechanical properties. Percolation theory and ergodic theory of infinite particle systems (Minneapolis, Minn., 1984–1985)*, 131–167,

[15] L. F. Gray, *A Reader's Guide to Gacs's "Positive Rates" Paper"*, J. Statist. Phys. **103** (2001), nos. 1-2, 45–267.

[16] G. Grimmett, *Percolation, in Fundamental Principles of Mathematical Sciences*, Springer-Verlag, Berlin, (1999) *xiv*+444pp.

[17] R. Holley and T.M. Liggett, *The survival of contact processes*, Ann. Probability **6** (1978) 198–206.

[18] A. Holroyd,*The metastability threshold for modified bootstrap percolation in d dimensions.* Electron. J. Probab. **11** (2006), 418-433

[19] R. Kozma et al., P.N. Balister, B. Bollobás, and W.J. Freeman, *Dynamical behavior of percolation models of cortical phase transitions*, Proc. 2001 Int. Conf. Nonlinear Theory and Applications NOLTA2001, Sendai, Japan, October 2001 (accepted).

[20] J. L. Lebowitz, C. Maes, E. R. Speer, *Statistical mechanics of probabilistic cellular automata.* J. Statist. Phys., **59** (1990), 117–170

[21] T.M. Liggett, *Survival of discrete time growth models, with applications to oriented percolation*, Ann. Appl. Probability **5** (1995) 613–636.

[22] P. Marcq, H. Chaté, and P. Manneville, *Universality in Ising-like phase transitions of lattices of coupled chaotic maps*, Phys. Review E, **55** (1997) no. 3 2606–2627.

[23] J. Molofsky, R. Durrett, J. Dushoff, D. Griffeath, and S. Levin. *Local frequency dependence and global coexistence.* Theoret. Pop. Biol., **55** (1999), 270-282

[24] G. Moran, sl On the period-two-property of the majority operator in infinite graphs. Trans. Amer. Math. Soc. **347** (1995), no. 5, 1649–1667.

[25] T. Mountford, *Rates for the probability of large cubes being non-internally spanned in modified bootstrap percolation*, Probability Theory and Related Fields, **93** (1992) 159–167.

[26] M. J. de Oliviera *Isotropic majority-vote model on a square lattice.* J. Stat. Phys. **66**, (1992) 273–285

[27] S. Poljak, M. Sûra. *On periodical behaviour in societies with symmetric influences.* Combinatorica **3** (1983), no. 1, 119–121.

[28] S. Poljak, D. Turzik. *On pre-periods of discrete influence systems,* Discrete Applied Mathematics. **13**, Issue 1, January 1986, 33–39.

[29] R. H. Schonmann, *Critical points of two-dimensional bootstrap percolation-like cellular automata.* J. Statist. Phys. **58** (1990), no. 5-6, 1239-1244.

[30] R. H. Schonmann, *On the behaviour of some cellular automata related to bootstrap percolation*, Ann. Probability, **20** (1992) 174–193.

[31] F. Spitzer *Interaction of Markov processes.* Advances in Math. **5** (1970) 246–290

[32] A. L. Toom, *Stable and attractive trajectories in multicomponent systems.* Multicomponent random systems, pp 549–575, Adv. Probab. Related Topics **6** Dekker, New York, 1980.

[33] S. Wolfram, *A New Kind of Science.* Wolfram Media, Inc., 2002