Percolation, forest-fires and monomer-dimers

(or the hunt for self-organised criticality)

Rachel Minke Brouwer

The research in this thesis has been carried out at the national research institute for Mathematics and Computer Science in the Netherlands (CWI). It has been financially supported by the Netherlands Organisation for Scientific Research (NWO).

VRIJE UNIVERSITEIT

Percolation, forest-fires and monomer-dimers

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor aan
de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof. dr. T. Sminia,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de faculteit der Exacte Wetenschappen
op donderdag 6 oktober 2005 om 13.45 uur
in het auditorium van de universiteit,
De Boelelaan 1105

door

Rachel Minke Brouwer

geboren te Heemskerk

promotor: prof. dr. J. van den Berg

voor opa Willem

I dedicate this thesis to my grandfather who was, and still is, my greatest example of a 'good person'. (Well, he has been caught riding his bike on the sidewalk, thus avoiding driving through a red light...) Unfortunately he has not seen me finish (or even start) this project.

Acknowledgements

The acknowledgements is the first (if not only) part of a thesis that is read. In some sense, it has been the hardest part to write. I have many people to thank for helping me finish; not only for helping me on a mathematical level, but also for making the non-mathematical part of my life much more enjoyable. It is not easy to make the list complete; some people helped me a lot by being interested and others by not asking at all.

First of all, I thank Rob van den Berg, for guiding me into the world of science, sharing his knowledge, and simply refusing to believe that I would not make it. His eye for detail prevented many of my mistakes. I am indebted to Antal Járai, Ronald Meester, Frank Redig and Bálint Tóth, for carefully reading my thesis and (as one of them put it) being 'fusspots'. This thesis has become better because of it. I very much appreciate the help of Misja and Ton, for commenting on parts of my thesis. Special thanks to Vladas Sidoravicius, for inviting me to IMPA in the wonderful city of Rio de Janeiro, and his 'little helpers' Bernardo, Leo and Sacha for making me feel at home. This visit gave me the inspiration to write the final first words. I am grateful to Daniël and Michaël, for (almost) being my personal WSO-ers and enjoyable lunches. Thanks to Adri Steenbeek, for checking my simulations.

As for the defense, Daan, pap, I know you will look wonderful in those penguin suits.

In between business and pleasure, I thank Rob and Bernadet, my roommates. We went through a lot together and it has been weird to be the last one left. Still, I cannot hear a Beatles-song without thinking of you guys.

Thank you,

W94 and associates (in particular Daniël, Misja, Pepijn and Thijs), for nice dinners, bbq's in the park, lovely holidays, movie nights and buurtborrels. I hope many will follow.

Loes, for being my friend for more than 15 years. Your no-nonsense approach always helps.

Marieke, for keeping me interested in other things than my desk and my couch.

Onno and Fenne, for saying the right things at the right time.

The ladies from D1 and the guys from Heeren 3 throughout the years, and all the others who became my friends during my years at UvO. There are too many to name you all.

The best people are always saved for last: my parents Berth and Erica and my sister Susanne gave me their endless support. Thank you for, although I was never able to explain what I was doing, being proud nevertheless. Finally a big 'thank you' to my Cas, for his love and patience, his imitation of a bird and an endless supply of pies. I owe you.

Rachel

Contents

Chapter 1

Introduction

When asked what this thesis is about I usually say either 'forest-fire models' or 'percolation-like problems'. A more precise answer would be 'self-organised criticality'(SOC) and if we concentrate on the problems (and not on the answers) it is mainly about the existence of certain processes[1].

Self-organised criticality has been a very fashionable subject over the last fifteen years. Ever since its introduction by Bak, Tang and Wiesenfeld [BTW88] in 1988 it has attracted a lot of attention. The concept of self-organised criticality is so appealing because it is a way to explain how complex structures arise from simple dynamics. It involves self-similarity, which is often observed in nature. This self-similarity, or scale-free behaviour, manifests itself in the occurrence of power laws: the probability distribution functions of the main quantities in a system obey a power law. The concept of self-organised criticality has been linked to sand-piles, earthquakes, forest-fires, epidemics, biological evolution and the distribution of words in a text. There are many more examples. General introductions on self-organised criticality are [Bak97] and [Jen98].

Traditional models in statistical physics such as the Ising model, or ordinary percolation, also show the power law behaviour which is typical for self-organised criticality. These models display a phase transition: the parameter of the model can take a critical value, above and below which the system behaves essentially different. *At* the critical value, power law behaviour is observed for the main quantities of the system. However, this is a fragile situation because we need to fine-tune the parameter to its critical value. In a self-organised critical model, the system drives *itself* into a critical state. There is no need to fine-tune the parameters of the model so in some sense, the critical behaviour is more robust.

Usually, self-organised critical models are difficult to analyse. Very little is known rigorously and specifically, as in the case of ordinary critical models, the values of the exponents in the power laws are hard to derive. Physicists have been able to predict

[1] Of course, some of the problems in this thesis have a different flavour, but most of them are derived from the above mentioned processes.

exponents for many models but in most cases, rigorous mathematical proofs do not exist. For some models, problems arise at an earlier stage. Before trying to prove power law behaviour and computing the critical exponents, one needs to worry about the existence of the model. Existence is meant in a mathematical sense here: usually it is easy to formulate the model in words, but its mathematical definition should be internally consistent. Problems with existence arise when one looks at infinite structures. In our first attempt to study an alleged self-organised critical model, the problem of existence popped up. It turned out to be a recurrent theme.

From now, this chapter becomes more concrete: it introduces the models that are discussed in this thesis. Being an introductory chapter, it contains an attempt to describe how these problems arose. Therefore the reader will find that their presentation here is in somewhat different order than further on. Chronologically, the chapter on the monomer-dimer model should be first, but since this chapter is least connected to the other chapters, we postpone its discussion till later. Further, since most problems have a strong connection with percolation theory, we present a short introduction to percolation theory first. Readers familiar with the subject are strongly encouraged to read this subsection, because it introduces some notation used later on.

1.1 Introduction to percolation theory

The basic example for percolation theory is the wetting of a stone. If we drop a porous stone into a bucket of water, will the inner part get wet? This question is all about the geometry of the holes in the stone. Broadbent and Hammersley [BH57] were the first to pose this question and they introduced the first so-called percolation model.

A two-dimensional version of this model is the following construction. Consider the graph induced by $\mathbb{Z}^2$. The sites, or vertices of the graph are the points of $\mathbb{Z}^2$. A bond, or edge, exists between two sites whenever their Euclidean distance is 1. Now we declare each bond of the lattice *open* with probability p and *closed* with probability $1-p$, independently of the other bonds. The measure associated with this model is the product measure with parameter p, denoted by $\mathcal{P}_p$. The interesting quantities in the percolation model are the clusters of open bonds. An open *cluster* (sometimes also called open p-cluster, to emphasize the parameter p) of some bond e is the maximal connected set of open bonds containing e. Water can only travel trough open bonds, so the question whether the inside of the stone gets wet becomes equivalent to the question what the clusters of open bonds look like. The behaviour of this model depends clearly on the value of p. Some examples can be found below in Figures 1.1 and 1.2.

The model described above is called bond percolation in two dimensions. We can also define the same model in general dimension d where the underlying graph is $\mathbb{Z}^d$, or instead of considering bonds, we can declare each *site* open with probability p and closed with probability $1-p$. Moreover, one can take other two-dimensional lattices (e.g. the triangular or honeycomb lattice) as the underlying graph, or even trees. In all the modifications above the bonds or sites behave independent of each other and

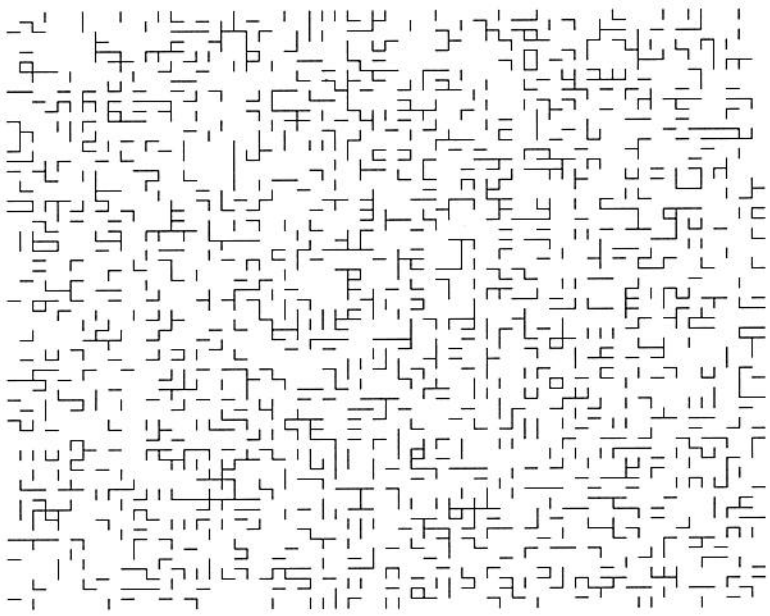

Figure 1.1: Realization of bond percolation where $p = 0.25$; opens bonds are represented by black lines.

hence the process is still governed by a product measure. We refer to this type of percolation by independent percolation. It is a small step to consider mechanisms where there is some dependence. For example, open bonds may be drawn together, in the sense that it is more likely that a bond is open if its neighbours are open. Another famous example is the random-cluster model of Fortuin and Kasteleyn.

One can see that percolation has a very broad scope. Percolation-like models have been used to study ferro-magnetism and electrical networks but also topics less related to physics such as forest-fires, epidemics and the growing of water lilies in a pond. All these models (and many more) have been studied in the literature. Some terminology: instead of 'open' and 'closed', other notions like 'occupied' and 'vacant' or even simpler '1' and '0' are used in the literature, depending on the nature of the model.

Back to bond percolation in two dimensions. As we can see in Figure 1.1 where $p = 0.25$, all open clusters are fairly small. The general behaviour looks very different from Figure 1.2 where we took $p = 0.75$. There it seems that most open bonds are in the same large cluster. At some point there must be a transition between the two

types of behaviour.

To formalise this statement, we first define the *percolation function.*

$$\theta(p) := \mathcal{P}_p(O \leftrightarrow \infty), \tag{1.1.1}$$

where O is the origin $(0,0)$. By $O \leftrightarrow x$ we mean that there exists an open path from O to x: a path $e_1, e_2, \ldots, e_n$ of consecutive disjoint open bonds such that O is an endpoint of e_1 and x is an endpoint of e_n. Similarly, $O \leftrightarrow \infty$ means that there is an infinite open path. Note that since we work on the infinite lattice, all sites behave the same so that the choice of O in the definition of the percolation function is not crucial. An application of the usual zero-one law shows that there is (somewhere) an infinite cluster a.s. whenever $\theta(p) > 0$. In fact, the infinite cluster is unique, in the sense that with probability 1, there is only one infinite cluster. This was first proved by Aizenman, Kesten and Newman in 1987 [AKN87a], [AKN87b]; Burton and Keane provided a more elegant proof in 1989 [BK89].

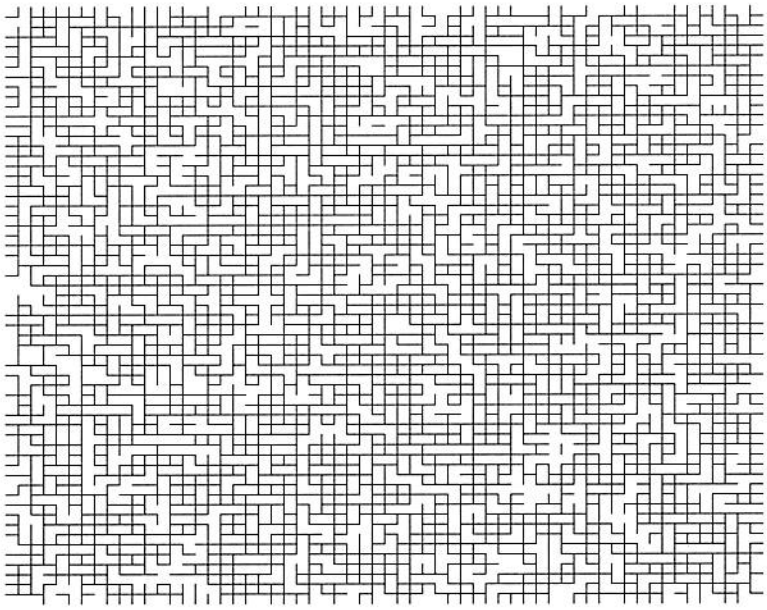

Figure 1.2: Realization of bond percolation where $p = 0.75$; open bonds are represented by black lines.

It is clear that $\theta(0) = 0, \theta(1) = 1$ and that $\theta(p)$ is increasing in p. We need to formalise the observed transition in the behaviour of our model (recall the different types of behaviour in Figures 1.1 and 1.2). To this end, we define the *critical probability*

$$p_c := \inf\{p : \theta(p) > 0\}. \tag{1.1.2}$$

If $p > p_c$ and hence $\theta(p) > 0$, we say that p is in the supercritical regime. Likewise, all $p < p_c$ are in the subcritical regime and p_c itself is called critical. The value of p_c is different for different percolation models and in general not known. We discuss a few.

site or bond percolation on $\mathbb{Z}$	$p_c = 1.$
bond percolation on the square lattice	$p_c = 1/2$
bond percolation on the triangular lattice	$p_c = 2\sin(\pi/18)$
site percolation on the square lattice	$p_c \approx 0.593$
site percolation on the triangular lattice	$p_c = 1/2$

The fact that $p_c = 1$ for bond or site percolation on $\mathbb{Z}$ is a trivial observation. Kesten [Kes80] proved the highly non-trivial fact that $p_c = 1/2$ for bond percolation on the square lattice, using among other things the self-duality of the square lattice. In the same spirit, Russo [Rus81] proved that $p_c = 1/2$ for site-percolation for any self-matching graph (satisfying certain symmetry conditions). This implies in particular that $p_c = 1/2$ for site percolation on the triangular lattice. Kesten's proof was adapted by Wierman [Wie81] who showed that $p_c = 2\sin(\pi/18)$ for bond percolation on the triangular lattice. Until now, no exact value has been found for the critical probability of site percolation on the square lattice. Heuristic arguments and simulations have shown that it is approximately 0.593, see [Hug96] for an overview.

From the definition of p_c it is not clear whether the percolation function is positive at p_c, or equivalently, whether the percolation function is continuous at the point p_c. It has been shown that $\theta(p_c) = 0$ in dimension 2 [Rus81]. This followed from the famous Russo-Seymour-Welsh arguments [Rus78], [SW78] (see Appendix A). Continuity in the critical point is also known in high dimensions ($d \geq 19$) [HS90], [HS94], [BA91].

For dimensions $3 \leq d < 19$ it is believed that $\theta(p_c) = 0$ but is has not been proved (yet). The sketch of $\theta(p)$ on the next page is therefore only known to be valid for $d = 2$ and $d \geq 19$. From this picture the reader may conclude that the percolation function is continuous. This is indeed the case in all dimensions (of course, except maybe in the point p_c in dimensions $3 \leq d < 19$). Right-continuity was first observed by Russo [Rus78], whereas left-continuity for $p > p_c$ was shown by van den Berg and Keane [BK84], under the assumption that the infinite cluster is unique.

Percolation theory has been a much explored area over the last couple of decades. It would go way too far to give an overview of all results in this field. We refer the reader to Grimmett [Gri99] for a thorough mathematical review up to 1999 or Stauffer [Sta85] for a more heuristic introduction in the field. In the last 5 years, an important breakthrough has been made using the so-called stochastic Loewner evolution technique (SLE), see e.g. [LSW02], [TW04]. In Appendix A we present some theorems on independent percolation that are used in later chapters.

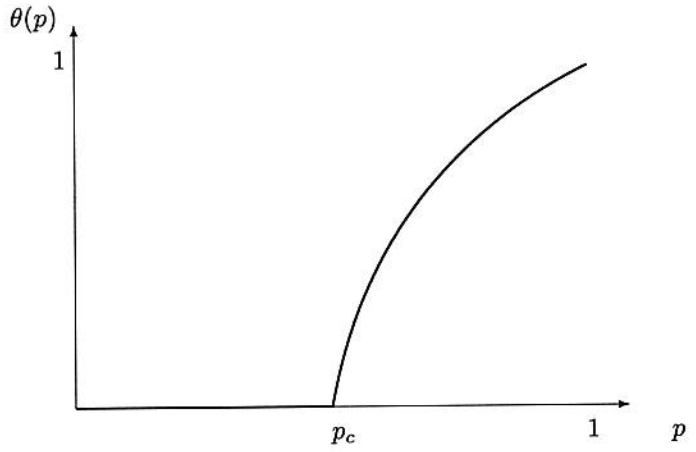

Figure 1.3: Sketch of the percolation function $\theta(p)$ when $d = 2$.

1.2 The permanent self-destructive process

This thesis started with a problem posed by Aldous in [Ald00]. Aldous' paper is on frozen percolation, but he proposes other, similar models. He introduces ([Ald00] Section 5.2) the following intuitive description of a process. Consider the binary tree. Each edge of the tree is either in state 'on' or 'off'. An edge which is 'off' will be turned 'on' at rate 1. When an infinite cluster of edges that are 'on' appears, all edges in this infinite cluster become 'off' instantaneously. Aldous' question is whether there is a (unique) stationary stochastic process satisfying this description. From this the following question arose: does a similar process as described above exist on $\mathbb{Z}^2$?

Instead of considering a stationary process, we chose to study a (maybe more natural) process where at time 0, the system is at rest, i.e. all sites in $\mathbb{Z}^2$ are 'off' or have value 0. A site that is 'on' has value 1. All 0's turn into 1's at rate 1, independently. As soon as there is an infinite cluster of sites with value 1, all sites in this cluster get value 0 simultaneously, and instantaneously. We call this the permanent self-destructive process. One may guess that this process displays self-organised critical behaviour, in the following sense: since we start with an empty configuration (all zeros), up to time t_c, corresponding to the critical probability of ordinary percolation p_c through $p_c = 1 - e^{-t_c}$, we do not have infinite clusters. But directly after t_c an infinite cluster arises and disappears instantaneously. If self-organised critical behaviour occurs, we expect that immediately afterwards, a new

infinite cluster arises and disappears, and so on. The system would be in a critical state at all times after t_c. Before we could answer the question whether such behaviour occurs, we had to show that the process indeed exists. In [BT01] a one-dimensional version of this model had been constructed. However, on $\mathbb{Z}^2$ this turned out to be a too difficult problem at that moment.

There seemed to be two different approaches to simplify the permanent self-destructive process:

(i). Introducing dynamics that govern the destruction of clusters. This turns the model into a forest-fire model and the permanent self-destructive process can be seen as a limiting case. We explain this in Section 1.3.

(ii). Discretising time, firstly by introducing a two-step model. Suppose instead of erasing infinite clusters as soon as they appear, we wait until a fixed time. At that time a catastrophe happens, destroying the infinite cluster. We then enable vacant sites to become occupied for some time. Repeating this procedure, we hope to obtain an approximation of the permanent self-destructive process. We discuss this approach in Section 1.4.

1.3 SOC forest-fire models

Suppose a forest is located on a big square $B(n) = [-n, n]^2$. Trees can grow at a regular distance from each other, on sites in $B(n)$. There are edges between sites when their Euclidean distance is 1. Suppose that the land is bare (empty) when we start at time 0. For a forest-fire model, we need two ingredients: a mechanism to model the growth of trees, and a mechanism to model fires. The growth mechanism is fairly simple. Each site in $B(n)$ has a Poisson clock with rate 1 attached to it, which is independent of the other sites. Whenever a clock rings, a tree tries to grow at its site. We write 'tries' because if a tree is already present, nothing changes. The fire mechanism is slightly more complicated. Suppose the fires are ignited by lightning. Each site has another Poisson clock attached (which is independent of the other clocks), with rate λ, governing lightning. Each time such a clock rings, lightning strikes. If no tree is present at that site, nothing happens. If a tree is present however, it is set on fire. The fire spreads to neighbouring trees and the entire cluster of the tree that was hit by lightning burns down. To make life easy, we assume that all trees in the cluster burn down instantaneously and simultaneously, so that we do not need to worry about trees growing on the boundary of a cluster that is somewhere on fire. This assumption is not that outrageous because in practice, the speed of fire is much larger than the speed at which trees grow. The model described above is a version of the well-known Drossel-Schwabl forest-fire model [DS92].

Intuitively, it seems likely that the forest-fire model resembles the permanent self-destructive process when taking the appropriate limits. First we take the size n of the system to infinity, and then the lightning rate λ to zero. As $\lambda \downarrow 0$, the probability that a finite cluster is hit by lightning goes to zero, so that we may argue that only

infinite[2] clusters are set on fire and moreover, that that will happen as soon as they arise. So intuitively, the behaviour of the forest-fire model should be similar to that of the permanent self-destructive process, in the appropriate limits. Note that when we reverse the order of these limits, we are left with a very boring model.

A simpler model, of which it is maybe less evident that it behaves like the permanent self-destructive process, is the following. Forget about lightning. Instead, suppose that fires are created by careless people, running around the boundary of the forest, throwing their burning cigarettes away. We model this by assuming that as soon as a cluster of trees touches the boundary of the box $B(n)$, it is set on fire. Again clusters of trees are destroyed instantaneously. One might argue that although people do silly things, they are not that stupid; the number of fires thus created could be too large to be realistic. However, adding another layer of sites on the outside of the box (i.e. considering the box $B(n+1)$), should convince the reader that we basically consider a model where trees on the boundary are set on fire with rate 1.

Fires only reach the inside of the box, if there is a connection of trees to the boundary. Near the centre of the box trees are set on fire if they are connected to a boundary far away, thus mimicking the infinite cluster. So intuitively, the behaviour near the centre resembles the permanent self-destructive process if the box is large. Closer to the boundary of course, things look different. It is not clear whether this is a problem when we look at the overall behaviour. Nevertheless, for both forest-fire models described above we can prove the same results.

So far, we have started with an empty box. The interesting behaviour occurs around the critical time t_c when very large (system-spanning) clusters are expected to arise for the first time. Figure 1.4 shows a forest-fire model just before, and just after the critical time. Trees are represented by black dots. We see that a large cluster of trees arises and is destroyed. Results concerning the two-dimensional forest-fire model around the critical time are presented in Chapter 3, Section 3.3.

The problem most studied in the literature is somewhat different. Physicists are usually interested in the behaviour of the system in steady-state: if we let the process run forever, do we reach some kind of equilibrium? And if so, what does the typical configuration in equilibrium look like? In Figure 1.5, one can find a picture of a system that has run for considerable time. Trees are represented by black dots. We see 'patches' of trees that are essentially sub or supercritical.

However, physicists do not agree whether this picture is the correct one, see for example [SDS02], [Gra02]. Limitations on computer speed and memory may influence the outcome. Besides, since we are working with both system size and lightning rate, changing their ratio changes the behaviour observed. On top of that all, it is not clear whether the stationary distribution on finite boxes converges (in some sense) to a steady-state distribution on $\mathbb{Z}^2$. This is an example where things are much simpler

[2]The question is of course, what 'infinite' means when we take a limit of finite boxes. See Open Problem 3.3.3.

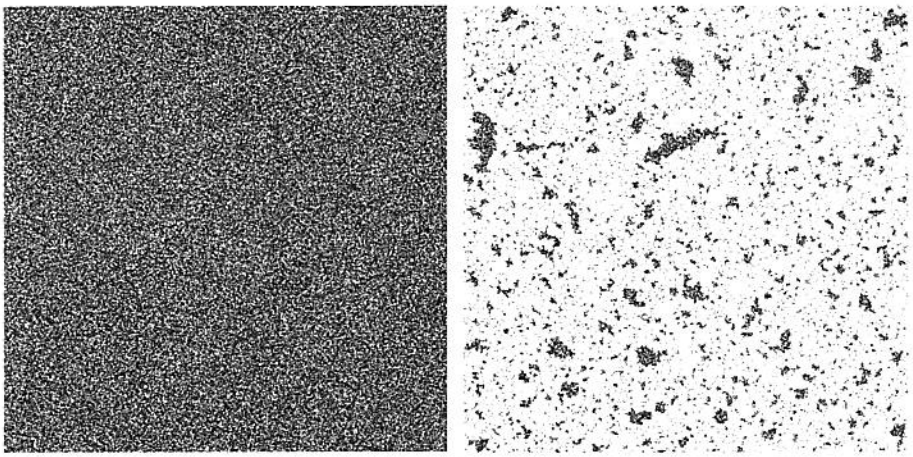

Figure 1.4: Forest-fire model of size 640×640, just before (left) and just after (right) a system-spanning fire. The fire occurred at time 0.94, corresponding to a tree density of 0.61. On the infinite lattice, the critical density is ≈ 0.593.

in one dimension.

On $\mathbb{Z}$, we know that the forest-fire model exists for fixed λ. Further, there exists at least one stationary measure and we have been able to improve a result by physicists on the cluster-size distribution, see Chapter 3, Section 3.2.

The pictures in this section were created with an adjusted version of [Hon97].

1.4 Self-destructive percolation

Maybe the simplest model to study aspects of the permanent self-destructive process is what we call self-destructive percolation. Suppose we perform independent site percolation with parameter $p > p_c$ on $\mathbb{Z}^2$. This gives us a configuration of occupied and vacant sites on $\mathbb{Z}^2$. In the next step, some catastrophe destroys (makes vacant) all sites that are in the infinite occupied cluster. Finally, each vacant site becomes occupied again with probability δ, independently.

If we take the probability p very close to p_c and δ very small, the model may behave just like the permanent self-destructive model very shortly after the critical time. If p is very close to p_c, the infinite cluster is very sparse; it is likely that a small δ is sufficient to reintroduce an infinite cluster. Similarly, we can introduce a three-step (with parameters p, δ_1, δ_2), four-step and in general an n-step model and hope that when $p \downarrow p_c$, $n \to \infty$ and $\delta_1, \delta_2, \ldots \downarrow 0$ the behaviour will resemble the permanent self-destructive process.

The two-step model turned out to be interesting (and complicated) enough in itself. We discuss it in Chapter 2 and prove several properties of the model. Amaz-

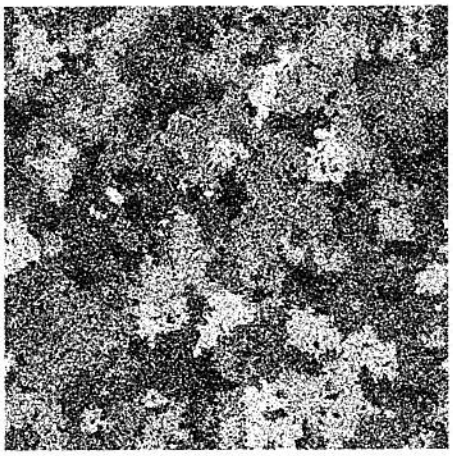

Figure 1.5: Forest fire model of size 640×640, assumed to be approaching steady-state.

ingly, the two-step model gave us a hunch about the existence of the permanent self-destructive process. We formulate a conjecture (see Conjecture 2.4.5) which, if true, implies that the permanent self-destructive process does not exist. This conjecture, if true, also turns out to have consequences for the two-dimensional forest-fire model. Unfortunately, we have not been able to prove or disprove this conjecture.

1.5 Frozen percolation

The step to so-called frozen percolation was easily made. Not only did the original question about the existence of the permanent self-destructive process arise from a paper about frozen percolation, but it also seemed to be a simpler mechanism to freeze clusters than to make them disappear.

Suppose that we attach a uniform $[0, 1]$ variable U_i to each site i of some graph G. At time 0, all sites are at rest. At time U_i, site i becomes activated, but as soon as i satisfies some criterion, it freezes. One should think of this criterion as the site being in an infinite active cluster, being in an active cluster of size at least N, or being connected by a path of active sites to some subset of G. This way, sites change their state at most twice; the model seems simpler than the forest-fire models where sites may change their state infinitely many times.

Unfortunately, there is no obvious coupling between forest-fire models and frozen percolation. One could try to do the following. Suppose that we attach Poisson clocks to all sites in some box $B(n)$. We consider the forest-fire model where fires start on the boundary, using the Poisson clocks to govern the growth of trees. We may couple this forest-fire process and the following frozen percolation process using the same Poisson clocks: if a Poisson clocks rings for the first time, its corresponding site is

activated in the frozen process. Further clock rings are ignored. A cluster freezes as soon as it touches the boundary. However, this direct coupling is of no use. Consider the following example on $B(2)$.

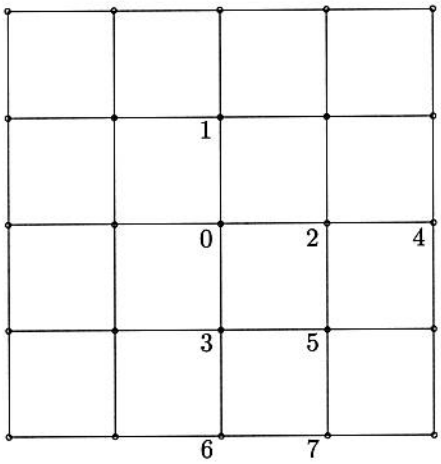

Suppose the Poisson clocks ring in the following order:

site 6 (6 burns, 6 freezes), site 3, site 6 (3,6 burn), site 0, site 5,
site 7 (0,3,5,7 freeze, 5, 7 burn), ...

Here site 0 freezes in the frozen process, but site 0 is *not* set on fire in the forest-fire process. Alternatively, one may think that the following coupling holds: if a site is set on fire, and all the sites in its cluster are set on fire for the first time, then the site freezes as well. A small extension of the previous example shows that this is not the case. Suppose the Poisson clocks ring in the following order:

site 6, site 3, site 6, site 0, site 5, site 7, as before; then,
site 1, site 2, site 4 (0,1,2,4 on fire for the first time, 2,4 freeze),....

In this case site 1 is in a fire, while at that time, all sites in its cluster burn for the first time. However, site 1 does not freeze.

Although lacking a direct connection to the forest-fire models, the frozen percolation process is interesting in itself. Probably the most natural version of frozen percolation is the version where we freeze infinite clusters. However, Benjamini and Schramm have shown (but did not publish) that the frozen percolation process where infinite clusters freeze (in a slightly different setting) does not exist on $\mathbb{Z}^2$. It may be interesting to investigate what happens if we freeze clusters as soon as they have size at least N. Eventually we can send N to infinity. The occurrence of cycles in the graph $\mathbb{Z}^2$ makes this a difficult problem; we turn our attention to $\mathbb{Z}$ and the binary tree. On $\mathbb{Z}$ dependencies can be handled more easily, so that we can do explicit calculations, see Chapter 4, Section 4.2.

On the binary tree, Aldous [Ald00] studied a version of frozen percolation where infinite clusters freeze. He showed that this process exists and that it is a nice example of self-organised criticality (although Aldous does not mention the term SOC): at *all* times, finite clusters behave like (have the same law as) ordinary critical percolation clusters, both conditional on being non-empty. So remarkably, at all times, we observe critical behaviour. We study frozen percolation where infinite clusters freeze on the tree with slightly different dynamics than in [Ald00], see Section 4.3. Most surprisingly, because computations became more complicated, our model also shows critical behaviour.

1.6 The monomer-dimer model

Compared to the rest of this thesis, the section on the monomer-dimer has a different flavour. For one, its emphasis lies on direct application; it presents an algorithm to sample (approximately) from the monomer-dimer distribution.

At a superficial glance, a configuration in the monomer-dimer model looks a bit like frozen percolation when $N = 2$. It is not the same though. In the monomer-dimer model, two neighbouring sites may pair up (in a dimer) but whereas two is company, three is a crowd: no connections of more than two sites are formed. The sites that do not pair up with a neighbouring site stay on their own (being a monomer). Note that for frozen percolation on $\mathbb{Z}^2$ for example, sites may join together in groups of sizes up to five when $N = 2$. A precise description of the monomer-dimer model on a graph $G = (V_G, E_G)$ is as follows: let $M \subseteq E_G$ be a subset of edges, such that no two edges in M share a common endpoint. Such a subset is called a *matching* and we give it a weight proportional to its size. We can define the monomer-dimer distribution μ_λ as follows. For any $M \subseteq E_G$

$$\mu_\lambda(M) := \frac{\lambda^{|M|} I(M \text{ is a matching})}{Z(\lambda)},$$

where $I(\cdot)$ is the indicator function and $Z(\lambda)$ a normalizing constant.

The monomer-dimer model is used in physics, for example to represent the absorption of oxygen on a surface. This is best seen by considering the model on the covering graph (or line graph) of G. The covering graph $\overline{G}$ is obtained from G in an easy manner: let the centre points of edges in E_G be the sites of $\overline{G}$ and add an edge between two sites of $\overline{G}$ whenever their corresponding edges in G share a common endpoint. Now oxygen molecules are absorbed at the sites of $\overline{G}$. Being very large, an oxygen molecule at some site prevents the absorption of another molecule at a neighbouring site. This representation is sometimes called the hard-core model. A few moments of thought should convince the reader that the monomer-dimer model and the hard-core model obtained from it, are the same.

A main question is as always: what does a typical configuration in the monomer-dimer model look like? This obviously depends on the value of λ. A large value of λ puts more weight on configurations with a lot of dimers, whereas a small value

of λ assigns more weight to configurations with mainly monomers. To give a more accurate answer to this question, we would like to sample from the monomer-dimer distribution (and preferably, we would like many samples). Computing the probability distribution and specifically the normalizing constant is not easy. If the graph G has n edges, there are 2^n possible subsets M, a number that becomes large rapidly. Exact computations become very hard. An answer to this problem is approximate sampling. Suppose we can find a Markov chain that has the set of all matchings on G as its state space, and the monomer-dimer distribution μ_λ as its stationary distribution. If we let this chain run forever, the configuration we eventually see looks like a sample from μ_λ. The question we try to solve is to determine what point in time approximates infinity, or equivalently, how long we have to run the chain to obtain an approximate sample. Chapter 5 addresses this issue.

1.7 List of publications

[1]. van den Berg, J. and Brouwer R. *Random sampling for the monomer-dimer model on a lattice,* J. Math. Phys., **41**, 2000, 1585–1597.

[2]. Brouwer R. *A bicategorical approach to Morita equivalence for von Neumann algebras,* J. Math. Phys., **44**, 2003, 2206–2214.

[3]. van den Berg, J. and Brouwer R. *Self-destructive percolation,* Random Structures & Algorithms, **24**, 2004, 480–501.

[4]. van den Berg, J. and Brouwer R. *Self-organised forest-fires near the critical time,* 2004, submitted.

[5]. Brouwer, R. and Pennanen J. *The cluster size distribution for the self-organised critical forest-fire model on* $\mathbb{Z}$, 2005, preprint.

Parts of Chapter 2 have been published as [3]. Chapter 3: Section 3.2 is based on [5]; Section 3.3 is almost identical to [4]. Chapter 5 has been published as [1].

Chapter 2

Self-destructive percolation

2.1 Introduction

Consider the square lattice $\mathbb{Z}^2$, where every site is vacant. Suppose that we make each site occupied with probability p, independent of the other sites. So far, we have ordinary independent percolation. We know that if $p > p_c$ there is a.s. an infinite occupied cluster. Now suppose that, by some catastrophe, all sites that are in the infinite occupied cluster are destroyed. That is, the infinite occupied cluster becomes vacant again. Finally, each site that is vacant after the catastrophe receives an independent 'enhancement': with probability δ each vacant site becomes occupied again, independently of the other sites. For the motivation behind this model, see Chapter 1, in particular Section 1.4.

Let $\mathcal{P}_{p,\delta}$ denote the distribution of the final configuration. Usually we write '1' for occupied and '0' for vacant so that we can consider $\mathcal{P}_{p,\delta}$ as a distribution on $\{0,1\}^{\mathbb{Z}^2}$, with the usual σ-field. We note that the origin O is occupied in the final configuration if either the above mentioned δ-enhancement was successful, or O belonged initially to a non-empty but finite occupied p-cluster. Hence

$$\mathcal{P}_{p,\delta}(O \text{ occupied}) = \delta + (1-\delta)(p - \theta(p)).$$

Recall that $\theta(p)$ is the percolation function for ordinary independent percolation.

As usual, we are interested in whether there is an infinite cluster in the final configuration. We define a percolation function

$$\theta(p,\delta) := \mathcal{P}_{p,\delta}(O \leftrightarrow \infty).$$

It is easy to see that for $p \leq p_c$, there is no infinite occupied p-cluster and hence

$$\mathcal{P}_{p,\delta} = \mathcal{P}_{p+(1-p)\delta}, \tag{2.1.1}$$

where we use the notation $\mathcal{P}_p$ for the product measure with parameter p. In particular, for $p \leq p_c$,

$$\theta(p,\delta) = \theta(p + (1-p)\delta). \tag{2.1.2}$$

Clearly the percolation function is increasing in δ. Further, it is easy to see that $\theta(p,0) = 0$ and that if $\delta > p_c$ then $\theta(p,\delta) > 0$. Hence our model displays a phase transition and we may define a critical probability $\delta_c(p)$ for all $p \in [0,1]$ by

$$\delta_c(p) := \inf\{\delta : \theta(p,\delta) > 0\}.$$

One might be tempted to define a critical probability $p_c(\delta)$ likewise. However, we need monotonicity in p to define this probability unambiguously. It is not at all clear whether the model has this monotonicity (probably not) and we address this issue in Subsection 2.2.4.

If $p > p_c$, we remove a positive fraction $\theta(p)$ of the occupied sites and we may intuitively reason that we need a 'non-negligible' enhancement to reintroduce an infinite occupied cluster. Indeed, we prove $\delta_c(p) > 0$ for $p > p_c$ in Section 2.3. Now what happens if p approaches p_c? The infinite cluster for p close to p_c is very sparse so one may guess that $\delta_c(p)$ is very small. We should be careful that we do not get too enthusiastic and boldly claim that for any $\delta > 0$,

$$\lim_{p \downarrow p_c} \theta(p,\delta) = \theta(p_c,\delta) > 0. \tag{2.1.3}$$

Although this statement is true on the binary tree (Section 2.6) we conjecture that it is not true on the square lattice (Section 2.4.3).

Since we are dealing with a percolation-like model, several other 'natural' questions arise. Is the infinite cluster unique? Does the measure possess the association property? What about exponential decay of cluster sizes? Does our model exhibit an analogue of the so-called RSW-property? In the following sections we answer some of these questions.

2.2 Generalisation and basic properties

We now consider the model in a more general context. Let G be a finite or countably infinite graph. The set of sites of G is denoted by V_G and we let Γ be a subset of V_G or the symbol ∞. Let $p, \delta \in [0,1]$. As we did in the previous section, we make each site $i \in V_G$ occupied with probability p. To this end, we take independent 0-1 variables $(X_i : i \in V_G)$, each X_i being 1 with probability p and 0 with probability $1-p$. Then we change ones into zeros at each site that is connected to Γ by an X-occupied path; an X-occupied path is a path on which each site j has $X_j = 1$. If Γ is the symbol ∞, we destroy all infinite clusters. To describe this formally, we define the $*$-operation as follows.

$$X_i^* = \begin{cases} 1 & \text{if } X_i = 1 \text{ and there is no } X\text{-occupied path from } i \text{ to } \Gamma, \\ 0 & \text{otherwise.} \end{cases} \tag{2.2.1}$$

Finally, let $(Y_i, i \in V_G)$ be 0-1 variables, independent of each other and the X_i's, being 1 with probability δ and 0 with probability $1-\delta$. The final configuration is now

described by the variables $(Z_i := X_i^* \vee Y_i : i \in V_G)$. Let $\mathcal{P}_{p,\delta}^{[G;\Gamma]}$ denote its distribution. We often omit the superscripts when no confusion is possible.

We now discuss an alternative way to describe the model, namely in terms of Poisson processes. This is sometimes useful because the Poisson processes allow us to couple the model at different values of the parameters. We consider the model as if it evolves in time. Again, let G be a finite or countably infinite graph, V_G its set of sites and Γ a subset of V_G or the symbol ∞. To each site, we attach independent Poisson clocks that ring after an exponential waiting time with expectation 1. At time 0, all sites are vacant. Now each time one of the Poisson clocks rings, its corresponding site tries to become occupied. We say 'tries', because if it was occupied already, nothing happens. This continues until time τ, which corresponds to the parameter p through the relation $1 - e^{-\tau} = p$. At time τ, the catastrophe happens: as before, all sites that have an occupied path to Γ are made vacant. Then the process continues and each time a Poisson clock rings, its corresponding site tries to become occupied. At time t, corresponding to the parameter δ through the relation $1 - e^{-(t-\tau)} = \delta$, we consider the final configuration. Note that each site that is vacant after the catastrophe, has time $t - \tau$ to become occupied again, so that that happens with probability δ.

Remark 2.2.1. All results in this chapter are stated for site percolation on the hypercubic lattice $\mathbb{Z}^d$, $d \geq 2$ unless stated explicitly otherwise. However, apart from Proposition 2.3.2, all results are easily adapted to bond percolation on $\mathbb{Z}^d, d \geq 2$. As for the two-dimensional case, the results in this chapter (in particular Subsections 2.4.2 and 2.4.3 and Section 2.5) can be adapted to other common two-dimensional lattices, such as the (bond and site version of the) triangular and honeycomb lattice.

2.2.1 Uniqueness of the infinite cluster

For ordinary percolation on $\mathbb{Z}^d$ the infinite cluster is unique, in the sense that there is a.s. only one infinite cluster, whenever the percolation function is positive. This subsection shows that the same behaviour holds for self-destructive percolation.

Suppose that we work on $\mathbb{Z}^d$. Let $\Gamma = \infty$, i.e. we remove the infinite cluster after independent percolation with parameter p. Let K denote the number of infinite occupied clusters in the final configuration. The question is now whether the infinite cluster is unique, or equivalently, whether $\mathcal{P}_{p,\delta}(K = 1) = 1 - \mathcal{P}_{p,\delta}(K = 0) \in \{0, 1\}$.

Proposition 2.2.2. *Consider the self-destructive percolation model on $\mathbb{Z}^d$. If $\theta(p, \delta) > 0$ then $\mathcal{P}_{p,\delta}(K = 1) = 1$, i.e. the infinite cluster is unique.*

We use the following results. The first theorem is proved[1] by Gandolfi, Keane and Newman [GKN92].

[1] The theorem in [GKN92] is proved in the bond percolation setting, however, the site version can be proved similarly.

Theorem 2.2.3. [Gandolfi-Keane-Newman] *Let μ be a probability measure on $\{0,1\}^{\mathbb{Z}^d}$ satisfying*

(i). μ is translation invariant over any $z \in \mathbb{Z}^d$,

(ii). μ satisfies the positive finite energy condition.

Then

$$\mu(K = 0 \text{ or } K = 1) = 1.$$

Lemma 2.2.4. *The system $(Z_i : i \in \mathbb{Z}^d)$ is a factor of an i.i.d. system.*

Proof. Let $T_z : \Omega \to \Omega$ act as $\mathbb{Z}^d$-shift on the probability space. We define a map $\phi : (\{0,1\} \times \{0,1\})^{\mathbb{Z}^d} \to \{0,1\}^{\mathbb{Z}^d}$ by

$$\phi((X_i, Y_i) : i \in \mathbb{Z}^d)) = ((X_i^* \vee Y_i) : i \in \mathbb{Z}^d),$$

using the notation from the beginning of this section. It is easy to see that ϕ is measurable and that $\phi \circ T_z = T_z \circ \phi$. This makes the system $(Z_i : i \in \mathbb{Z}^d)$ a factor of an i.i.d. system. □

Proof. [Proposition 2.2.2] We first check the conditions of the Gandolfi-Keane-Newman theorem. Translation invariance for $\mathcal{P}_{p,\delta}$ follows immediately from the translation invariance of the underlying Bernoulli measures. As for the positive finite energy condition, it means the following: for all $v \in \mathbb{Z}^d$,

$$\mathcal{P}_{p,\delta}(Z_v = 1 | Z_{v'} : v' \in \mathbb{Z}^d \backslash \{v\}) > 0. \tag{2.2.2}$$

Note that $Z_v = 1$ occurs whenever $Y_v = 1$ and that the latter is independent of everything else. Hence, for all $v \in \mathbb{Z}^d$,

$$\begin{aligned} \mathcal{P}_{p,\delta}(Z_v = 1 | Z_{v'} : v' \in \mathbb{Z}^d \backslash \{v\}) &\geq \mathcal{P}_{p,\delta}(Y_v = 1 | Z_{v'} : v' \in \mathbb{Z}^d \backslash \{v\}) \\ &= \mathcal{P}_\delta(Y_v = 1) = \delta. \end{aligned}$$

We conclude that

$$\mathcal{P}_{p,\delta}(K = 0 \text{ or } K = 1) = 1.$$

By ergodicity (which follows from the fact that the system is a factor of an i.i.d system) it follows that $\mathcal{P}_{p,\delta}(K = 1) \in \{0,1\}$. This shows uniqueness of the infinite cluster for the self-destructive percolation model on $\mathbb{Z}^d$. □

Note that the argument above does not work for the self-destructive percolation model on other graphs. On trees for example, there are infinitely many infinite clusters whenever the percolation function is positive. The argument is very similar to the argument for ordinary percolation.

There are other ways to show uniqueness of the infinite cluster on $\mathbb{Z}^d$, but they require more properties. An application of an earlier result by Gandolfi,Keane and

Russo [GKR88] shows uniqueness of the infinite cluster on $\mathbb{Z}^2$. Here both the association property (Subsection 2.2.3) and ergodicity of the measure under horizontal and vertical translation (separately) are needed. The argument of Burton and Keane [BK89] works for $\mathbb{Z}^d, d \geq 2$. It uses however, besides ergodicity, the *full* finite energy condition: besides satisfying the positive energy condition, the probability that a site has eventually value 0, given the configuration on the other sites, should be positive too. Intuitively this seems to be the case (with $(1-p)(1-\delta)$ as a lower bound), but we have not been able to find a proof.

2.2.2 Weak convergence

Let $S(n)$ be the set of vertices that have graph distance n from O. Recall the definition of $\mathcal{P}_{p,\delta}^{[\mathbb{Z}^d;S(n)]}$ from the beginning of Section 2.2.

Lemma 2.2.5. *Let $d \geq 2$. The measure $\mathcal{P}_{p,\delta}^{[\mathbb{Z}^d;S(n)]}$ convergences weakly (as $n \to \infty$) to $\mathcal{P}_{p,\delta}^{[\mathbb{Z}^d;\infty]} = \mathcal{P}_{p,\delta}$, that is, for all cylinder events A we have*

$$\lim_{n\to\infty} \mathcal{P}_{p,\delta}^{[\mathbb{Z}^d;S(n)]}(A) = \mathcal{P}_{p,\delta}(A). \tag{2.2.3}$$

Proof. Since A is a cylinder event it depends on finitely many sites only. We may assume that $A \subseteq \{0,1\}^W$, for some finite subset W. Take n sufficiently large, such that W is contained in the interior of $S(n)$. We consider the difference

$$\mathcal{P}_{p,\delta}^{[\mathbb{Z}^d;S(n)]}(A) - \mathcal{P}_{p,\delta}(A).$$

Only configurations such that, for some $i \in W$, i is connected by an X-occupied path to $S(n)$ but not to infinity, contribute to this difference. Hence

$$\left|\mathcal{P}_{p,\delta}^{[\mathbb{Z}^d;S(n)]}(A) - \mathcal{P}_{p,\delta}(A)\right| \leq \sum_{i\in W} \left[\mathcal{P}_p(i \leftrightarrow S(n)) - \mathcal{P}_p(i \leftrightarrow \infty)\right].$$

Each term on the right converges to zero as $n \to \infty$. The fact that W is finite finishes the proof. □

Remark 2.2.6. Note that the lemma above holds for any graph which is locally finite.

2.2.3 Association

Recall that a finite collection of 0-1 valued random variables $(\omega_i : 1 \leq i \leq n)$ (or its corresponding probability measure, say μ, on $\{0,1\}^n$) is said to be *associated* (also called *positively associated* in the literature) if for all increasing events $A, B \subset \{0,1\}^n$,

$$\mu(A \cap B) \geq \mu(A)\mu(B).$$

This is equivalent to saying that for all increasing functions f and g on $\{0,1\}^n$,

$$E_\mu(fg) \geq E_\mu(f)E_\mu(g),$$

where

$$E_\mu(f) = \sum_{x \in \{0,1\}^n} \mu(x)f(x).$$

Lemma 2.2.7. *Let G be a finite graph, and let Γ be a subset of V_G. Then the measure $\mathcal{P}_{p,\delta}^{[G;\Gamma]}$ is associated.*

Proof. The following facts (i) - (iii) for 0-1 valued random variables are well-known and we present them without proof.

(i). A collection of independent random variables is associated (FKG-Harris inequality).

(ii). If a collection $(\omega_i : 1 \leq i \leq n)$ is associated, a collection $(\sigma_j : 1 \leq j \leq m)$ is associated and these two collections are independent of each other, then the joint collection $(\omega_i : 1 \leq i \leq n; \sigma_j : 1 \leq j \leq m)$ is associated.

(iii). If a collection $(\omega_i : 1 \leq i \leq n)$ is associated, and $f_1, \ldots, f_k$ are increasing 0-1 valued functions on $\{0,1\}^n$, then the collection

$$(f_1(\omega_1, \ldots \omega_n), \ldots, f_k(\omega_1, \ldots, \omega_n))$$

is associated.

We use the random variables X_i, Y_i, X_i^* and Z_i as defined in the beginning of this section. To prove the lemma we need to show that the collection $(Z_i : i \in V_G)$ is associated. Recall that $Z_i = X_i^* \vee Y_i$, so from properties (i)-(iii) above follows easily that it is sufficient to prove that the collection $(X_i^* : i \in V_G)$ is associated.

First some more notation. Let $\Omega = \{0,1\}^{V_G}$. For $\omega \in \Omega$ we let $C(\Gamma) = C(\Gamma, \omega)$ be the occupied cluster of Γ (i.e. the set of all sites which have a path to Γ on which $\omega \equiv 1$). For $V \subset V_G$, $\overline{V} = V \cup \partial V$ denotes the set which consists of V and all neighbours of V, with the convention that $\overline{\emptyset} = \emptyset$. If $W \subset V_G$ and $A \subset \Omega$ is increasing, the event $\{A$ occurs outside $W\}$ will denote the set of all $\omega \in \Omega$ such that the configuration $\hat{\omega}$, defined by

$$\hat{\omega}_i = \begin{cases} 0 & \text{if } i \in W \\ \omega_i & \text{otherwise,} \end{cases}$$

is in A.

Let $\mathcal{P}_p$ be the distribution of the collection of random variables $(X_i : i \in V_G)$, that is, $\mathcal{P}_p$ is the product measure with parameter p on Ω. In particular, $\mathcal{P}_p$ is associated by property (i) above. Let $\mathcal{P}_p^*$ be the measure on Ω corresponding with the collection $(X_i^* : i \in V_G)$.

We have

$$\begin{aligned}
\mathcal{P}_p^*(A \cap B) &= \sum_W \mathcal{P}_p(\overline{C(\Gamma)} = W, A \cap B \text{ occurs outside } W) \\
&= \sum_W \mathcal{P}_p(\overline{C(\Gamma)} = W)\mathcal{P}_p(A \cap B \text{ occurs outside } W) \\
&\geq \sum_W \mathcal{P}_p(\overline{C(\Gamma)} = W)\mathcal{P}_p(A \text{ occurs outside } W)\ \mathcal{P}_p(B \text{ occurs outside } W),
\end{aligned}$$

where we have summed over all subsets W of V_G and where the second equality uses the fact that the event $\{\overline{C(\Gamma)} = W\}$ depends only on the sites in W. The inequality uses the fact that $\mathcal{P}_p$ is associated: note that, for fixed W, the events $\{A$ occurs outside $W\}$ and $\{B$ occurs outside $W\}$ are increasing.

To simplify notation, let for $W \subset V_G$, $f(W)$ denote $\mathcal{P}_p(A$ occurs outside $W)$ and $g(W)$ denote $\mathcal{P}_p(B$ occurs outside $W)$. It is clear that if $W \subset W'$, then $f(W) \geq f(W')$, and that a similar statement holds for g. So the last expression above equals

$$\begin{aligned}
&\sum_W \mathcal{P}_p(\overline{C(\Gamma)} = W)f(W)g(W) \\
&\quad = \sum_W \sum_{\omega \in \Omega : \overline{C(\Gamma,\omega)} = W} \mathcal{P}_p(\omega)f(W)g(W) \\
&\quad = \sum_{\omega \in \Omega} \mathcal{P}_p(\omega)f(\overline{C(\Gamma,\omega)})g(\overline{C(\Gamma,\omega)}) \\
&\quad \geq \sum_{\omega \in \Omega} \mathcal{P}_p(\omega)f(\overline{C(\Gamma,\omega)}) \sum_{\omega \in \Omega} \mathcal{P}_p(\omega)g(\overline{C(\Gamma,\omega)}) \\
&\quad = \mathcal{P}_p^*(A)\mathcal{P}_p^*(B),
\end{aligned}$$

where the inequality holds because $f(\overline{C(\Gamma,\omega)})$ and $g(\overline{C(\Gamma,\omega)})$ are both decreasing in ω (so that we can apply the association property of $\mathcal{P}_p$ once again), and the last equality follows from similar arguments as before, but now 'working backwards'. □

We say that an infinite collection of 0-1 valued random variables (or its corresponding measure) is associated if every finite subcollection is associated.

Corollary 2.2.8. *Let G be a locally finite graph and Γ be a subset of V_G or the symbol ∞. Then the measure $\mathcal{P}_{p,\delta}^{[G;\Gamma]}$ is associated.*

Proof. Now we know that $\mathcal{P}_{p,\delta}^{[G;\Gamma]}$ is associated for finite graphs G (Lemma 2.2.7), we apply the weak convergence result (Lemma 2.2.5, or rather Remark 2.2.6) to prove the corollary. □

We continue with a second proof of the fact that $(X_i^* : i \in V_G)$ is associated for a finite graph G. We present it here because the method is more powerful and can be used to prove slightly extended results. The proof uses the BKR inequality (conjectured by van den Berg and Kesten [BK85], proved by Reimer [Rei00]), which we will briefly review. Let $A, B \subseteq \{0,1\}^n$. We define an event $A \Box B$ which is, in words, the event that A and B 'occur disjointly'. For a configuration ω and a subset $K \subseteq \{1, \ldots, n\}$ we define the cylinder event $[\omega]_K$ by

$$[\omega]_K = \{\omega' \in \Omega : \omega'_i = \omega_i, \text{ for all } i \in K\}.$$

Now we define the event $A \Box B$ by

$$\begin{aligned} A \Box B \;=\; & \{\omega \in \Omega : \exists K, L \subseteq \{1, \ldots, n\}, \text{ such that } K \cap L = \emptyset, \\ & [\omega]_K \subseteq A, [\omega]_L \subseteq B\}. \end{aligned}$$

Theorem 2.2.9. [BKR inequality] *Let $A, B \subseteq \Omega = \{0,1\}^n$ and $\mathcal{P}$ a product measure on Ω. Then*

$$\mathcal{P}(A \Box B) \leq \mathcal{P}(A)\mathcal{P}(B).$$

First note that events in terms of $(X_i^* : i \in V_G)$ can also be written in terms of $(X_i : i \in V_G)$. If we formulate the events this way, we are dealing with the product measure $\mathcal{P}_p$, which allows us to use the BKR inequality. Keep in mind that an event that is increasing in the collection $(X_i^* : i \in V_G)$ is not necessarily increasing in the collection $(X_i : i \in V_G)$.

Proof. [Second proof of Lemma 2.2.7] Suppose we have an increasing event A and a decreasing event B in terms of $(X_i^* : i \in V_G)$. These events correspond with (not necessarily increasing or decreasing) events A' and B' in terms of $(X_i : i \in V_G)$. If we can show that $A' \cap B' = A' \Box B'$ then we can apply the BKR inequality to obtain

$$\begin{aligned} \mathcal{P}_p^*(A \cap B) \;&=\; \mathcal{P}_p(A' \cap B') \;&=\; \mathcal{P}_p(A' \Box B') \\ &\leq\; \mathcal{P}_p(A')\mathcal{P}_p(B') \;&=\; \mathcal{P}_p^*(A)\, \mathcal{P}_p^*(B). \end{aligned}$$

The inequality above is equivalent to $\mathcal{P}_p^*(A \cap B) \geq \mathcal{P}_p^*(A)\mathcal{P}_p^*(B)$ if both A and B are increasing and we are done. So it is left to show that for A increasing and B decreasing in terms of $(X_i^* : i \in V_G)$ we have

$$A' \cap B' = A' \Box B'.$$

It is immediately clear that (for any two events A' and B'), $A' \Box B' \subseteq A' \cap B'$. For the inclusion the other way we should find disjoint subsets K and L such that $[\omega]_K \subseteq A'$ and $[\omega]_L \subseteq B'$ for any configuration ω in $A' \cap B'$. We start with the simplest case. Let A be the event that $X_v^* = 1$ and B be the event that $X_w^* = 0$ for some vertices v and w in V_G. Then A' is the event that $X_v = 1$ and v is separated from Γ by a set of vertices S such that for all $i \in S, X_i = 0$. By separated we mean that the vertex

v and the set Γ are in different components of $G\backslash S$. Now B' is the event that either $X_w = 0$ or $w \in C(\Gamma)$, where $C(\Gamma)$ denotes the X-occupied cluster of Γ as before. There are many configurations ω in which this could happen, and we have to choose disjoint sets K and L that guarantee the occurrence of A' and B'. How to make this choice? We first distinguish two possibilities.

1. If $w \in \overline{C(\Gamma)}$ we take $K = \overline{C(v)}$ and $L = C(\Gamma)$. Regardless the value of X_w, the values on the set L imply that $X_w^* = 0$.
2. If $w \notin \overline{C(\Gamma)}$ then $X_w = 0$. We take $K = \{v\} \cup \partial C(\Gamma)$ and $L = \{w\}$.

In both cases, it is easy to see that the sets K and L chosen are disjoint. In general, for any configuration ω that is in the intersection of an increasing event A and a decreasing event B we can choose the disjoint sets K and L in a similar way. Suppose that $\omega \in A \cap B$, where A is increasing and B is decreasing in terms of the variables $(X_i^* : i \in V_G)$. Then there are $v_1, \ldots, v_k$ with $X_{v_1} = X_{v_2} = \cdots = X_{v_k} = 1$ and $v_1, v_2, \ldots, v_k$ not in $C(\Gamma)$, implying the event A. Further, there are $w_1, w_2, \ldots, w_m$ such that each w_i has either $X_{w_i} = 0$ or $w_i \in C(\Gamma)$ for $i \in \{1, \ldots, m\}$, implying the event B. We choose

$$\begin{aligned} K &= \partial C(\Gamma) \cup \{v_j : j = 1, \ldots, k\} \\ L &= C(\Gamma) \cup \{w_k : w_k \notin \partial C(\Gamma), k \in \{1, \ldots, m\}\} \end{aligned}$$

Of course we could have taken these sets K and L in the simple case. We hope that treating the simple case extensively gave the reader more feeling for the method. □

Further remarks on association

One can wonder if we could use the BKR inequality to obtain further results concerning association. What if we do not remove the cluster of Γ once, but twice or more? Suppose that for each $i \in V_G$, in addition to X_i and Y_i, we assign another independent Bernoulli variable Y_i' with parameter δ_2. In this case we can look at the collection $((X_i^* \vee Y_i)^* \vee Y_i' : i \in V_G)$. In words, we perform independent percolation with parameter p, remove the cluster of Γ, turn all vertices that now have value zero to one with probability δ, again remove the cluster of Γ and turn all vertices that now have value zero to one with probability δ_2. Is the collection $((X_i^* \vee Y_i)^* \vee Y_i' : i \in V_G)$ again associated? As we have seen before, using the properties (i) - (iii) above, it is sufficient to show that the collection $((X_i^* \vee Y_i)^* : i \in V_G)$ is associated. Let $\mathcal{P}_{p,\delta}^*$ denote the measure corresponding to this collection.

For the collection $((X_i^* \vee Y_i)^* : i \in V_G)$ we cannot prove association in general, but we do have a partial result. We need the following definition.

Definition 2.2.10. A finite collection of 0-1 valued random variables $(\omega_i : 1 \leq i \leq n)$ (or its corresponding probability measure, say μ, on $\{0,1\}^n$) is said to be *pairwise associated* (also called pairwise positively correlated in the literature) if for all i and j in $\{1, \ldots, n\}$,

$$\mu(\omega_i = 1, \omega_j = 1) \geq \mu(\omega_i = 1)\mu(\omega_j = 1).$$

We can show the following.

Lemma 2.2.11. *Let T be any locally finite tree. Let Γ be an arbitrary site. The measure $\mathcal{P}^*_{p,\delta}$ corresponding to the collection $((X_i^* \vee Y_i)^* : i \in V_T)$ is pairwise associated.*

The lemma is proved by a similar method as in the (second) proof of Lemma 2.2.7, but in a slightly more sophisticated way. Note that the lemma includes the line $\mathbb{Z}$.

In general, the BKR inequality cannot be used to prove pairwise association. In the following example we cannot find disjoint subsets that imply the events $\{(X_v^* \vee Y_v)^* = 1\}$ and $\{(X_w^* \vee Y_w)^* = 0\}$, respectively. Let $\Gamma = \{0\}$.

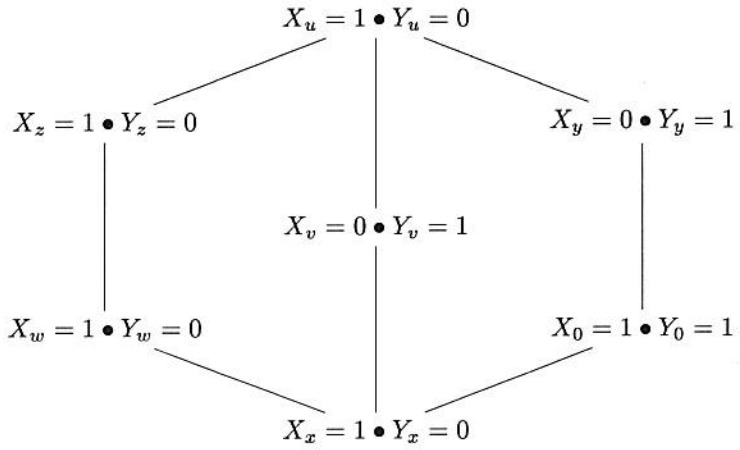

To guarantee the event $\{(X_v^* \vee Y_v)^* = 1\}$ we need both the values X_x and X_z whereas we need either X_x or X_z for the event $\{(X_w^* \vee Y_w)^* = 0\}$.

Open Problem 2.2.12. For general graphs, is the probability measure $\mathcal{P}^*_{p,\delta}$ (pairwise) associated?

Remark 2.2.13. All results in the subsection above are also valid if we assign to each site $i \in V_G$ independent Bernoulli variables with (possibly) different parameter values $p(i), \delta_1(i), \ldots, i \in V_G$.

2.2.4 Monotonicity

As we remarked in the introduction, it is easy to see that the percolation function of the self-destructive percolation model is increasing in δ. For fixed δ, the percolation function is increasing in p on the interval $[0, p_c]$ by (2.1.2). Further, if $\delta < p_c$, $0 = \theta(1, \delta) = \theta(\delta) < \theta(p_c, \delta)$, so that the percolation function cannot be increasing in p on $[0, 1]$. However, we will show that we do have 'some sort' of monotonicity.

Lemma 2.2.14. *Let Γ be a subset of V_G or ∞. For fixed r, the family of probability measures*

$$\left\{\mathcal{P}^{[G;\Gamma]}_{p,\frac{r-p}{1-p}}, 0 < p < r\right\}$$

is stochastically decreasing in p. That is, if $0 < p' < p < r$, then

$$\mathcal{P}^{[G;\Gamma]}_{p',(r-p')/(1-p')} \text{ dominates } \mathcal{P}^{[G;\Gamma]}_{p,(r-p)/(1-p)}.$$

Proof. Recall the 'time evolution' point of view from the beginning of Section 2.2. At time 0 each site of G is vacant. At each site we have a Poisson clock that rings after an exponential waiting time with expectation 1. If a site is vacant at the time its clock rings, it becomes occupied. Now at some deterministic time τ some 'catastrophe' happens: each site that has an occupied path to Γ at time τ becomes vacant instantaneously. Then we move on and each vacant site becomes occupied at rate 1 again.

Fix t and assume that $t > \tau$. We use a coupling argument to show that the configuration at time t is stochastically decreasing in the time τ at which the catastrophe takes place. Let $U_i^{(k)}$, for $k \in \{1,2\}, i \in V_G$ be i.i.d. exponentially distributed variables with mean 1. Then $U_i^{(1)}$ can be seen as the first time that the site i becomes occupied. If $U_i^{(1)} < \tau$ so that the catastrophe at time τ turns i vacant, it becomes occupied again at time $\tau + U_i^{(2)}$.

Now, let $\eta_i(t)$ denote the state of site i in the final configuration at time t. It is clear from the description above that $(\eta_i(t) : i \in V_G)$ has distribution $\mathcal{P}_{1-e^{-\tau},1-e^{-(t-\tau)}}$. We have that $\eta_i(t) = 0$ (vacant) if and only if

$$\left\{U_i^{(1)} > t\right\} \text{ or}$$
$$\left\{U_i^{(2)} > t-\tau,\ U_i^{(1)} < \tau\right.$$
$$\left.\textit{and} \text{ there is a path } \pi \text{ from } i \text{ to } \Gamma \text{ such that for each } j \text{ on } \pi, U_j^{(1)} < \tau\right\}.$$

If this condition holds for some τ, then it clearly holds for all $\tau' \in (\tau, t)$. Hence for fixed t, the family of distributions

$$\left\{\mathcal{P}_{1-e^{-\tau},1-e^{-(t-\tau)}}, 0 \le \tau < t\right\},$$

is stochastically decreasing in τ. We take $p = 1 - e^{-\tau}$ and $r = 1 - e^{-t}$, so that $(r-p)/(1-p) = 1 - e^{-(t-\tau)}$, to finish the proof of the lemma. □

In words, the lemma states the following. If we fix a time t, then the earlier we remove the cluster of Γ, the more occupied sites we see in the final configuration at time t. This should be clear intuitively: if the catastrophe occurs earlier, the cluster of Γ is smaller so we remove less. Furthermore, sites have more time to become occupied again.

An immediate consequence of this lemma and the obvious monotonicity of $\mathcal{P}_{p,\delta}$ in δ is the following

Corollary 2.2.15. *If $p_2 \geq p_1$ and $p_2 + (1-p_2)\delta_2 \leq p_1 + (1-p_1)\delta_1$ then $\mathcal{P}_{p_1,\delta_1}$ dominates $\mathcal{P}_{p_2,\delta_2}$.*

2.2.5 Other properties

A much used property of ordinary percolation is the exponential decay of the radius distribution in the subcritical regime and the radius distribution of finite clusters in the supercritical regime (compare Theorems A.1 and A.2). A natural question is whether the self-destructive percolation model displays the same behaviour. For $p \leq p_c$ this trivially holds by (2.1.1), but this is not the interesting case. Recall that $S(n)$ is the set of sites at graph distance n from the origin.

Open Problem 2.2.16. Is there a positive constant $\sigma_1 = \sigma_1(p,\delta)$ such that for all p, δ such that $\theta(\cdot,\cdot) = 0$ on a neighbourhood of (p,δ),

$$\mathcal{P}_{p,\delta}(0 \leftrightarrow S(n)) \leq e^{-\sigma_1 n} \quad ?$$

Open Problem 2.2.17. Is there a positive constants $\sigma_2 = \sigma_2(p,\delta)$ such that for all p, δ such that $\theta(\cdot,\cdot) > 0$ on a neighbourhood of (p,δ),

$$\mathcal{P}_{p,\delta}(0 \leftrightarrow S(n), |C_O| < \infty) \leq e^{-\sigma_2 n} \quad ?$$

We have not been able to prove or disprove exponential decay. It would be nice to have an affirmative answer to Open Problem 2.2.17, because in two dimensions it implies another much appreciated property: (a form of) the RSW-theorem; more precisely, it implies the analogue of equation (A.4) when ε approaches 1.

The proof of the RSW-theorem breaks down for self-destructive percolation by lack of control on dependencies: the original proof uses independence between the area above and below a lowest occupied crossing. However, in our case 'the infinite cluster could have been anywhere', which creates dependencies between these two areas.

2.3 The critical probability $\delta_c(p)$

A natural question to ask is whether the self-destructive percolation model is meaningful in the sense that $\delta_c(p)$ is non-trivial, i.e. not equal to zero or p_c. From the definition of the model easily follows that $\delta_c(0) = \delta_c(1) = p_c$ and $\delta_c(p_c) = 0$, but these turn out to be the only cases.

Lemma 2.3.1. *For all $p \in (0,1)$, the critical probability $\delta_c(p)$ satisfies $\delta_c(p) < p_c$ on $\mathbb{Z}^d, d \geq 2$.*

Proof. By (2.1.2) we can compute $\delta_c(p)$ explicitly for $p \leq p_c$. We have

$$\delta_c(p) = \frac{p_c - p}{1 - p}. \tag{2.3.1}$$

Further, it is easy to see that $\theta(p, \delta) \geq \theta(\delta)$. Hence $\delta_c(p) \leq p_c$ holds for all p. To prove that strict inequality holds, we use the enhancement method of Aizenman and Grimmett [AG91]. Recall the variables X_i and Y_i from the beginning of Section 2.2. Each site that is occupied in the final configuration is either Y-occupied or X-occupied but in a finite X-cluster. If the remaining finite X-clusters had no effect on the percolation probability, the critical value $\delta_c(p)$ would be equal to p_c. We need to show that the addition of finite X-clusters forms an *essential enhancement* and hence shifts the critical value $\delta_c(p)$ strictly below p_c. An enhancement is called *essential* if there exists a configuration on the graph such that the original configuration does not have a doubly infinite self-avoiding path, but the configuration enhanced at the origin does.

Intuitively, it is clear that adding finite X-clusters would be essential. The problem is however, that 'adding finite X-clusters' is not a properly defined enhancement. Therefore we look at the following. We add a single occupied site with probability $p(1-p)^D$ where D is the number of neighbours of the site, and repeat this at a regular distance. We should take the distance large enough to avoid dependencies. This enhancement can clearly create a doubly infinite path where previously there was none, so it is essential. The addition of finite clusters of arbitrary size dominates the above addition of single sites so we may conclude that $\delta_c(p) < p_c$.

□

The more interesting and more difficult question is whether the critical probability $\delta_c(p)$ is strictly larger than zero. This is clearly not the case for *all* p because $\delta_c(p_c) = 0$. But is the critical probability $\delta_c(p)$ positive for all $p \neq p_c$? Equivalently, for $p \neq p_c$, is there a $\delta > 0$ such that $\theta(p, \delta) = 0$? The exact computation (2.3.1) provides us with such a δ easily for $p < p_c$. Note that this answers our question for bond or site percolation on $\mathbb{Z}$ where $p_c = 1$. On $\mathbb{Z}^d, d \geq 2$, the answer is affirmative, but the proof involves tedious iterations. However, on some two-dimensional lattices we have a more elegant proof.

Proposition 2.3.2. *Suppose $p > p_c$ and $d = 2$. For site percolation on the triangular lattice or the square lattice the following holds.*

(i). $\forall p > p_c \ \exists \delta > 0 \ \ s.t. \ \theta(p, \delta) = 0.$

(ii). In particular,

$$\text{if } p(1-\delta) \geq p_c, \ \text{ then } \theta(p, \delta) = 0.$$

Hence, $\delta_c(p) \geq 1 - \frac{p_c}{p}$.

Proof. It is easy to see that (ii) $\Rightarrow$ (i) so we will prove part (ii). Recall the variables X_i, X_i^*, Y_i and Z_i introduced the beginning of Section 2.2. Suppose we colour all sites i that have $X_i = 1$ and $Y_i = 0$, red. So the probability that a site is red is $p(1-\delta)$, independently of the other sites. If p, δ satisfy the assumption of our claim, there will (a.s.) be arbitrary large red circuits around the origin O. Since $p > p_c$, these red circuits will be part of the infinite X-cluster from some point on. This means that for each j on the red circuits lying in the infinite X-cluster, $X_j^* = 0$ and since also $Y_j = 0$, these circuits eventually have Z-value 0. Moreover, they 'isolate' the origin from infinity. Hence, there is a.s. no infinite Z-occupied path leaving from O, so $\theta(p, \delta) = 0$. □

Note that the argument above is based heavily on the fact that for both the triangular and the square lattice, the original graph is a subgraph of its matching graph, so that a circuit is at the same time a circuit in the matching graph. In fact, the proposition holds for all planar lattices with this property. For two-dimensional lattices lacking this property the argument does not work, but part (i) of Proposition 2.3.2 follows from Proposition 2.3.4.

We proceed with the argument for $\mathbb{Z}^d, d \geq 2$. The proof is based on ideas from Chayes and Chayes [CC84] for ordinary percolation. The essential difference is that we have to control dependencies created by the removal of the infinite cluster. The overall strategy is as follows. Suppose we can find a $\delta(p)$ such that the probability of having an occupied (i.e. with Z-value 1) crossing in a d-dimensional rectangle is very small. An iteration argument will show that then also the probability of having an occupied crossing in a larger rectangle is very small. An application of the Borel-Cantelli lemma then shows that a.s. no infinite cluster exists.

We need the following notation. Let $f_d(n_1, n_2, \ldots, \underline{n_i}, \ldots, n_d)$ denote the probability that in a d-dimensional rectangle $[0, n_1] \times [0, n_2] \times \cdots \times [0, n_d]$ there is a Z-occupied path from the $(d-1)$-dimensional rectangle $[0, n_1] \times \cdots \times [0, n_{i-1}] \times \{0\} \times [0, n_{i+1}] \times \cdots \times [0, n_d]$ to the $(d-1)$-dimensional rectangle $[0, n_1] \times \cdots \times [0, n_{i-1}] \times \{n_i\} \times [0, n_{i+1}] \times \cdots \times [0, n_d]$. That is, $f_d(n_1, n_2, \ldots, \underline{n_i}, \ldots, n_d)$ is the probability that we cross the rectangle in the i-th direction.

Lemma 2.3.3. *For each d, all $p \in [0,1]$ and all $\delta \in [0,1]$, the probabilities $f_d(.)$ satisfy*

(i). For each permutation σ on $(1, \ldots, d)$ we have

$$f_d(n_1, n_2, \ldots, \underline{n_i}, \ldots, n_d) = f_d(n_{\sigma(1)}, n_{\sigma(2)}, \ldots, \underline{n_{\sigma(j)}}, \ldots, n_{\sigma(d)}),$$

where j is such that $\sigma(j) = i$. In particular, we may always rearrange the variables such that we consider an occupied crossing in the first direction.

(ii). For $m_1 \leq n_1$ we have

$$f_d(\underline{n_1}, n_2, \ldots, n_d) \leq f_d(\underline{m_1}, n_2, \ldots, n_d),$$

and for each $j > 1$ and $n_j \leq m_j$ we have

$$f_d(\underline{n_1}, n_2, \ldots, n_j, \ldots, n_d) \leq f_d(\underline{n_1}, n_2, \ldots, m_j, \ldots, n_d).$$

(iii). Suppose $p > p_c$. Then there exist constants $\nu(p) > 0$ and $C(d)$ and such that for all n

$$f_d(\underline{4n}, 8n, 8n, \ldots, 8n) \leq C(d) f_d(\underline{n}, 2n, 2n, \ldots, 2n)^2 + e^{-\nu(p)n}. \quad (2.3.2)$$

Proof. The first part of the lemma is easily seen by symmetry and the second part follows from inclusion. The third part is the hard one. We will prove it first in two dimensions, so that the method of proof becomes clear. For $d = 2$, the claim is the following: for $p > p_c$ and $f_2(\underline{n}, 2n)$ defined as above, we have for all n,

$$f_2(\underline{4n}, 8n) \leq 169 f_2(\underline{n}, 2n)^2 + 36n\, e^{(\frac{-3n}{4}\psi(p))}, \quad (2.3.3)$$

where $\psi(p)$ is the positive constant obtained from Theorem A.2.

Consider the rectangle $[0, 4n] \times [0, 8n]$. We define two disjoint regions containing parts of the rectangle: let A be the set of vertices $\{(x_1, x_2) : x_1 \geq 9n/4\}$ and B the set of vertices $\{(x_1, x_2) : x_1 \leq 7n/4\}$. Let ∂A and ∂B denote the boundary of these regions, respectively. Finally, let $R_A = [3n, 4n] \times [0, 8n]$ and $R_B = [0, n] \times [0, 8n]$. See figure 2.1.

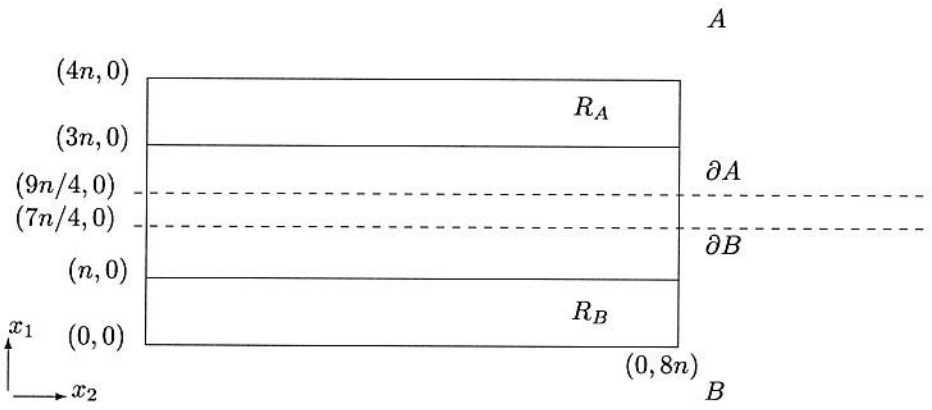

Figure 2.1: A picture of A,B, R_A and R_B.

If a Z-occupied top-down crossing of the rectangle exists, there must be Z-occupied top-down crossings in R_A and R_B. If these crossings were independent, we would obtain $f_2(\underline{4n}, 8n) \leq f_2(\underline{n}, 8n)^2$. Unfortunately, this is not the case but we can show that they are 'almost' independent. Suppose that there is no finite X-cluster connecting R_A and ∂A and that there exists a Z-occupied top-down crossing in R_A. Each site in this crossing is either Y-occupied, or part of a finite X-cluster that lies completely inside A. The existence of a crossing with this property depends on the state of the sites in A only. A similar reasoning can be held for a Z-occupied top-down crossing in R_B, when we assume that there is no finite X-occupied cluster connecting R_B and ∂B. At this point we use independence. Let $C_X(i)$ denote the X-occupied cluster of some site i. We write $\partial A \leftrightarrow_f R_A$ to indicate that there exists a finite X-occupied cluster between ∂A and R_A. We may write

$$\begin{aligned} f_2(\underline{4n}, 8n) &\leq \mathcal{P}_{p,\delta}(\{\partial A \leftrightarrow_f R_A\} \cup \{\partial B \leftrightarrow_f R_B\}) \\ &+ \mathcal{P}_{p,\delta}(\exists\ Z\text{-occupied top-down crossings in } R_A \text{ and } R_B, \\ &\qquad \partial A \not\leftrightarrow_f R_A, \partial B \not\leftrightarrow_f R_B) \\ &\leq \mathcal{P}_{p,\delta}(\exists \text{ top-down crossing } \pi \text{ in } R_A, \text{ s.t. } \forall i \in \pi : Y_i = 1 \text{ or} \\ &\qquad \{X_i = 1 \text{ and } \overline{C_X(i)} \subset A\})^2 + 2\,\mathcal{P}_{p,\delta}(\partial B \leftrightarrow_f R_B) \\ &\leq f_2(\underline{n}, 8n)^2 + 2\,\mathcal{P}_{p,\delta}(\partial B \leftrightarrow_f R_B) \\ &\leq f_2(\underline{n}, 8n)^2 + 36n\, e^{(\frac{-3n}{4}\psi(p))}. \end{aligned} \tag{2.3.4}$$

The exponential decay as well as the constant $\psi(p)$ in the last inequality follow from Theorem A.2. We have used that the number of sites in ∂R_B equals $18n$. It is left to show that

$$f_2(\underline{n}, 8n) \leq 13 f_2(\underline{n}, 2n), \tag{2.3.5}$$

which, together with (2.3.4) proves (2.3.3). Consider the rectangle R_A. A top-down crossing in R_A might lie completely in one of the (seven) rectangles of size n by $2n$. If not, it must cross one of the (six) n by n squares in left-right direction. An example of both possibilities can be found in Figure 2.2. At least one of the crossings indicated

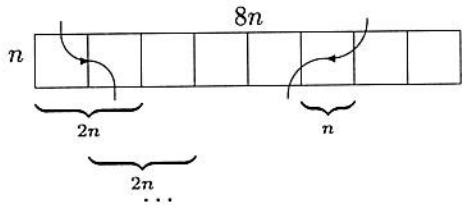

Figure 2.2: Possibilities for a top-down crossing in $[0, n] \times [0, 8n]$.

above must occur if R_A has a top-down crossing. This gives us

$$f_2(\underline{n}, 8n) \leq 7f_2(\underline{n}, 2n) + 6f_2(n, \underline{n}).$$

Applying parts (i) and (ii) of the lemma shows that $f_2(n, \underline{n}) \leq f_2(\underline{n}, 2n)$ which shows (2.3.5) and finishes our proof in two dimensions.

For general dimensions $d \geq 2$ we consider again the rectangle $[0, 4n] \times [0, 8n] \times \cdots \times [0, 8n]$. Completely similar to the two-dimensional argument we define

$$\begin{aligned} A &= \{(x_1, \ldots, x_d) : x_1 \geq 9n/4\}, \\ B &= \{(x_1, \ldots, x_d) : x_1 \leq 7n/4\}, \\ \partial A &= \{(9n/4, x_2, \ldots, x_d)\}, \\ \partial B &= \{(7n/4, x_2, \ldots, x_d)\}, \\ R_A &= A \cap ([3n, 4n] \times [0, 8n] \times \cdots \times [0, 8n]), \\ R_B &= B \cap ([0, n] \times [0, 8n] \times \cdots \times [0, 8n]). \end{aligned}$$

Now we can repeat the computation in (2.3.4) and write

$$f_d(\underline{4n}, 8n, \ldots, 8n) \leq f_d(\underline{n}, 8n, \ldots, 8n)^2 + 2\, \mathcal{P}_{p,\delta}(\partial B \leftrightarrow_f R_B). \tag{2.3.6}$$

We need to bound the second term on the right. Once again applying Theorem A.2, we can take $\nu(p) > 0$ s.t.

$$\begin{aligned} 2\, \mathcal{P}_{p,\delta}(\partial B \leftrightarrow_f R_B) &\leq 2\, |\partial R_B| e^{(\frac{-3n}{4}\psi(p))} \\ &\leq e^{-\nu(p)n}, \end{aligned} \tag{2.3.7}$$

where we used the notation $|\partial R_B|$ for the number of sites in ∂R_B.

We now consider the probability $f_d(\underline{n}, 8n, \ldots, 8n)$, the first term on the right of (2.3.4). Suppose that a Z-occupied crossing in the direction of the first coordinate exists in $[0, n] \times [0, 8n] \times \cdots \times [0, 8n]$. The projection on the second coordinate of any such crossing (which is an interval contained in $[0, 8n]$) lies either completely inside

$$\begin{aligned} &[kn, (k+2)n], \text{ for some } k \in \{0, \ldots, 6\}, \\ &\textit{or} \text{ contains an interval} \\ &[ln, (l+1)n], \text{ for some } l \in \{1, \ldots, 6\}. \end{aligned} \tag{2.3.8}$$

This is the same as saying that an occupied Z-crossing in the first coordinate of $[0, n] \times [0, 8n] \times \cdots \times [0, 8n]$ lies either completely inside $[0, n] \times [kn, (k+2)n] \times [0, 8n] \cdots [0, 8n]$, for some $k \in \{0, \ldots, 6\}$ or crosses a rectangle $[0, n] \times [ln, (l+1)n] \times [0, 8n] \times \cdots \times [0, 8n]$, for some $l \in \{1, \ldots, 6\}$ in the second coordinate. This gives us

$$\begin{aligned} f_d(\underline{n}, 8n, \ldots, 8n) &\leq 7f_d(\underline{n}, 2n, 8n, \ldots, 8n) + 6f_d(n, \underline{n}, 8n, \ldots, 8n) \\ &\leq 13 f_d(\underline{n}, 2n, 8n, \ldots, 8n). \end{aligned}$$

We have used parts (i) and (ii) of the lemma in the second inequality above. We may repeat the argument above for the probability $f_d(\underline{n}, 2n, 8n, \ldots, 8n)$. Consider a Z-occupied crossing in the first coordinate of $[0,n]\times[0,2n]\times[0,8n]\times\cdots\times[0,8n]$. The projection on the third coordinate of such crossing must satisfy a statement similar to (2.3.8) above. It follows that

$$f_d(\underline{n}, 2n, 8n, \ldots, 8n) \leq 13 f_d(\underline{n}, 2n, 2n, 8n, \ldots, 8n).$$

We iterate this argument and obtain

$$f_d(\underline{n}, 8n, \ldots, 8n) \leq 13^{d-1} f_d(\underline{n}, 2n, \ldots, 2n). \tag{2.3.9}$$

Combining this with (2.3.6) and (2.3.7) proves part (iii) of the lemma. □

Now we are ready for the general proof of the claim that $\delta_c(p) > 0$.

Proposition 2.3.4. *Consider the hyper cubic lattice $\mathbb{Z}^d$, $d \geq 2$. If $p > p_c$, the critical probability $\delta_c(p) > 0$.*

Proof. Fix $p > p_c$. We are looking for a $\delta > 0$ such that $\theta(p,\delta) = 0$. The proof uses part 2.3.2 of Lemma 2.3.3. We consider the probability $f_d^\delta(\underline{n}, 2n, \ldots, 2n)$, where we write the superscript to emphasize the dependence on δ. First we define a_n as the last term in (2.3.2).

$$a_n := e^{-\nu(p)n}.$$

Further, define $D = C(d) + 1$ with $C(d)$ as in (2.3.2). We choose n_0 such that for all $n > n_0$, $a_n < (\frac{1}{2D})^2$. Now suppose that $\delta = 0$. Then an occupied crossing can only exist if it is part of a finite X-occupied cluster. This has a probability which is exponentially decaying in the size of the rectangle, so that we can choose $n_1 \geq n_0$ such that for all $n \geq n_1$,

$$f_d^0(\underline{n}, 2n, \ldots, 2n) < \frac{1}{2D}.$$

Note that $f_d^\delta(\underline{n_1}, 2n_1, \ldots, 2n_1)$ is a continuous function of δ, since it depends on the values of finitely many δ-variables. Hence we may choose a $\delta > 0$ small enough such that

$$f_d^\delta(\underline{n_1}, 2n_1, \ldots, 2n_1) < \frac{1}{2D}.$$

We rewrite (2.3.2) as

$$f_d(\underline{4n}, 8n, \ldots, 8n) \leq D \max\{a_n, f_d(\underline{n}, 2n, \ldots, 2n)^2\}. \tag{2.3.10}$$

Iterating (2.3.10), starting with n_1 and using $a_{4n} \leq a_n^4$, we obtain for $i = 1, 2, \ldots$

$$f_d(\underline{4^i n_1}, 2\cdot 4^i n_1, \ldots, 2\cdot 4^i n_1) \leq \frac{1}{2^{2^i} D}.$$

Hence,

$$\sum_{i=1}^{\infty} f_d(\underline{4^i n_1}, 2\cdot 4^i n_1, \ldots, 2\cdot 4^i n_1) \quad < \quad \infty. \tag{2.3.11}$$

Let $B(n)$ be the d-dimensional cube with sides of length n around the origin and consider the annuli $A(4^i n_1, 2\cdot 4^i n_1) = B(2\cdot 4^i n_1)\backslash B(4^i n_1)$, for $i \in \mathbb{N}$. These annuli are clearly disjoint and if an occupied infinite cluster of the origin exists, it must cross all the annuli from $\partial B(4^i n_1)$ to $\partial B(2\cdot 4^i n_1)$. If there is a crossing of the annulus for some i, there must also be a crossing in one of the 'slabs' of size $4^i n_1 \times 2\cdot 4^i n_1 \times \cdots \times 2\cdot 4^i n_1$. Hence the probability that a crossing exists in the annulus $A(4^i n_1, 2\cdot 4^i n_1)$ is smaller than $(2d) f_d(\underline{4^i n_1}, 2\cdot 4^i n_1, \ldots, 2\cdot 4^i n_1)$. The sum of these probabilities over all annuli is finite by (2.3.11) and hence we can apply the Borel-Cantelli lemma, which tells us that with probability one, there are only finitely many annuli that have an occupied crossing going from the inside boundary to the outside boundary. We conclude that $\theta(p,\delta) = 0$, and hence $\delta_c(p) > \delta > 0$. □

2.4 Continuity and discontinuity of the percolation function

We have seen in the previous section that we need a non-trivial δ to reintroduce an infinite occupied cluster in the final configuration. As we mentioned in the introduction, it is easy to believe that $\delta_c(p)$ goes to zero as p goes down to p_c: for every $p > p_c$ there is an infinite cluster, but that cluster becomes very sparse as p approaches p_c and the fraction of sites we remove goes to zero. On the binary tree, the intuition turns out to be true (see Section 2.6). However, the geometry of the infinite cluster seems to play an important role, so that it may be false on other graphs. We will present a conjecture (supported by computer simulations) from which, if true, it follows that on the square lattice, there exists a $\delta > 0$ such that for all $p > p_c$, there is no percolation, i.e. $\theta(p,\delta) = 0$.

In the next subsection, we show continuity of the percolation function outside the points $(p, \delta_c(p))$ and (p_c, δ). Continuity of the percolation function may not come as a surprise, but even for ordinary percolation it is not trivial (and in fact believed but not proved at the critical point in dimensions $3 \le d < 19$). Subsection 2.4.2 introduces the conjecture on which the claimed discontinuity in dimension 2 is based and Subsection 2.4.3 shows how this discontinuity follows from the conjecture.

2.4.1 Continuity

Consider the percolation function $\theta(p,\delta)$ for self-destructive percolation on $\mathbb{Z}^d, d \ge 2$.

Theorem 2.4.1. *Let p be a point of continuity for the ordinary percolation function $\theta(\cdot)$ and let δ be such that either*

1. *$\theta(p,\delta) = 0$ or*
2. *there exists an open $V \ni (p,\delta)$ such that $\theta(q,\varepsilon) > 0$ for all $(q,\varepsilon) \in V$.*

Then $\theta(\cdot,\cdot)$ is continuous at the point (p,δ).

Note that if $\theta(p_c) = 0$, the ordinary percolation function is continuous so that then *all* $p \in [0,1]$ are points of continuity.

Recall once again the 'time evolution' version of the model. At each site of our lattice we have a Poisson clock of rate 1. If such a clock rings, its site becomes occupied. If it is already occupied nothing happens. If we remove the infinite cluster at time τ and consider the final configuration at time $t \geq \tau$, we have seen that the configuration at time t has distribution

$$\widehat{\mathcal{P}}_{\tau,t} := \mathcal{P}_{1-e^{-\tau},1-e^{-(t-\tau)}}.$$

We define

$$\widehat{\theta}(\tau,t) := \widehat{P}_{\tau,t}(O \to \infty) = \theta(1-e^{-\tau}, 1-e^{-(t-\tau)}). \tag{2.4.1}$$

The function

$$f : (\tau,t) \mapsto (1-e^{-\tau}, 1-e^{-(t-\tau)})$$

as well as its inverse

$$f^{-1} : (p_1,p_2) \mapsto (-\log(1-p_1), -\log((1-p_2)(1-p_1)))$$

is continuous. So Theorem 2.4.1 is equivalent to

Theorem 2.4.2. *Let τ be such that $1-e^{-\tau}$ is a point of continuity for the ordinary percolation function $\theta(\cdot)$. Further, let t be such that either*

1. *$\widehat{\theta}(\tau,t) = 0$, or*
2. *there exists an open $W \ni (\tau,t)$ such that $\widehat{\theta}(\sigma,s) > 0$ for all $(\sigma,s) \in W$.*

Then $\widehat{\theta}(\cdot,\cdot)$ is continuous at the point (τ,t).

This is the theorem we will prove. The advantage of this point of view is that we now have monotonicity of $\widehat{\theta}(\tau,t)$ in both τ and t by Lemma 2.2.14. First we will show continuity for cylinder events.

Lemma 2.4.3. *Let τ be such that $1-e^{-\tau}$ is a point of continuity of the ordinary percolation function $\theta(\cdot)$. Let A be a cylinder event. Then $\widehat{\mathcal{P}}_{(\tau,t)}(A)$ is a continuous function at the point (τ,t).*

Proof. Let A_S be the (finite) subset of sites determining the cylinder event A. Note that after the removal of the infinite cluster (if any) at time τ, the configuration on A_S depends only on the Poisson clocks inside the set A_S between times τ and t. The conditional probability of A, given the configuration at time τ on A_S is therefore a continuous function in $t-\tau$. Since A_S is finite, it is sufficient to show that for any $\alpha \in \{0,1\}^{A_S}$, $\mathcal{P}_{1-e^{-\tau},0}(\alpha)$ is a continuous function in τ. Suppose $\tau' < \tau$ and

let $\mathcal{P}$ denote the measure governing the Poisson clocks. We couple the processes at parameter values τ' and τ using the same Poisson clocks. This gives us

$$\begin{aligned}\left|\mathcal{P}_{1-e^{-\tau'},0}(\alpha) - \mathcal{P}_{1-e^{-\tau},0}(\alpha)\right| &\leq \mathcal{P}(\text{a Poisson clock in } A_S \text{ rings in the interval } [\tau', \tau]) \\ &+ |A_S|\left[\theta(1-e^{-\tau}) - \theta(1-e^{-\tau'})\right].\end{aligned}$$

If $\tau' > \tau$ a similar inequality can be derived, so that

$$\begin{aligned}\lim_{\tau' \to \tau}\left|\mathcal{P}_{1-e^{-\tau'},0}(\alpha) - \mathcal{P}_{1-e^{-\tau},0}(\alpha)\right| &\leq \lim_{\tau' \to \tau} |A_S||\tau' - \tau| \\ &+ \lim_{\tau' \to \tau} |A_S|\left|\theta(1-e^{-\tau'}) - \theta(1-e^{-\tau})\right| \\ &= 0.\end{aligned}$$

□

Proof. [Theorem 2.4.2.] Continuity of $\widehat{\theta}$ at (τ, t) means that for any sequence $(\tau_i, t_i)_{i\geq 0}$ with $\lim_{i\to\infty}(\tau_i, t_i) = (\tau, t)$,

$$\lim_{i\to\infty} \widehat{\theta}(\tau_i, t_i) = \widehat{\theta}(\tau, t). \tag{2.4.2}$$

Given a point (τ, t) we can divide our parameter space into four quadrants.

$$\begin{aligned}\text{I} &:= \{(x,y) : x \leq \tau, y \geq t\}, \\ \text{II} &:= \{(x,y) : x \geq \tau, y \geq t\}, \\ \text{III} &:= \{(x,y) : x \geq \tau, y \leq t\}, \\ \text{VI} &:= \{(x,y) : x \leq \tau, y \leq t\}.\end{aligned}$$

Note that each of these quadrants includes part of the horizontal and vertical axes trough the point (τ, t). Note that it is sufficient to show that the requested convergence holds along monotone sequences that lie completely inside one of these quadrants.
I) Consider a monotone sequence $(\tau_i, t_i)_{i\geq 0}$ lying inside quadrant I, such that $(\tau_i, t_i) \to (\tau, t)$ as $i \to \infty$ and for all i, $\tau_i \leq \tau_{i+1}$ and $t_i \geq t_{i+1}$. Along this sequence the percolation function is decreasing by Lemma 2.2.14. This implies

$$\widehat{\theta}(\tau, t) \leq \lim_{i\to\infty} \widehat{\theta}(\tau_i, t_i). \tag{2.4.3}$$

We will show that equality holds. Observe that for all $\sigma \leq s$,

$$\widehat{\mathcal{P}}_{\sigma,s}(O \to S(n)) =: \widehat{\theta}_n(\sigma, s) \geq \widehat{\theta}(\sigma, s), \tag{2.4.4}$$

and

$$\lim_{n\to\infty} \widehat{\theta}_n(\sigma, s) = \widehat{\theta}(\sigma, s).$$

Using (2.4.4), for each n,

$$\lim_{i\to\infty} \widehat{\theta}(\tau_i, t_i) \leq \lim_{i\to\infty} \widehat{\theta}_n(\tau_i, t_i) = \widehat{\theta}_n(\tau, t), \tag{2.4.5}$$

where we have applied Lemma 2.4.3. Finally, let $n \to \infty$ in (2.4.5) to show that $\widehat{\theta}(\tau, t) \geq \lim_{i\to\infty} \widehat{\theta}(\tau_i, t_i)$ which together with (2.4.3) shows continuity in quadrant I.

III) Consider a sequence $(\tau_i, t_i) \to (\tau, t)$ lying in the quadrant III. Again it is sufficient that we restrict ourselves to monotone sequences, such that $\tau_i \geq \tau_{i+1}$ and $t_i \leq t_{i+1}$. If $\theta(\tau, t) = 0$, the argument is trivial: by monotonicity

$$0 \leq \lim_{i\to\infty} \widehat{\theta}(\tau_i, t_i) \leq \theta(\tau, t) = 0.$$

So suppose that there exists an neighbourhood of (τ, t) where the percolation function is non-zero. We follow a van den Berg-Keane type argument [BK84]. Let C_{τ_i,t_i} denote the cluster of the origin at time t_i, when we remove the infinite cluster at time τ_i. Monotonicity of the sequence shows

$$C_{\tau_i,t_i} \subseteq C_{\tau_{i+1},t_{i+1}}, \text{ for all } i \geq 0. \tag{2.4.6}$$

Note that for the statement above, we need the 'time evolution' point of view, so that we can couple the process, (using the same Poisson clocks), at all parameter values. Let, for $t \geq \tau$, $\omega(\tau, t) \in \{0,1\}^{\mathbb{Z}^d}$ denote the configuration at time t, when the catastrophe occurs at time τ. We would like to compare this configuration with the configuration we obtain by using the same Poisson clocks, but where no catastrophe takes place. Let $\omega(t) \in \{0,1\}^{\mathbb{Z}^d}$ denote the configuration at time t, when no catastrophe occurs. Further, for $v \in \mathbb{Z}^d, \omega_v(t) \in \{0,1\}$ denotes the value at site v in the configuration $\omega(t)$. We have to show that

$$\mathcal{P}(|C_{\tau,t}| = \infty, |C_{\tau_i,t_i}| < \infty, \text{ for all } i) = 0. \tag{2.4.7}$$

Take j so large that $\widehat{\theta}(\tau_j, t_j) > 0$. Such j exists by the condition in the statement of the theorem. Suppose that $|C_{\tau,t}| = \infty$. We will show that there exists a k such that $|C_{\tau_k,t_k}| = \infty$.

Let I_j be the infinite cluster in $\omega(\tau_j, t_j)$, that exists a.s. by definition of j. If $O \in I_j$ we take $k = j$ and we are done. So suppose that $O \notin I_j$. By (2.4.6) and the uniqueness of $C_{\tau,t}$ (Proposition 2.2.2), I_j is contained in $C_{\tau,t}$. Hence there is a finite path π from O to some site in I_j such that $\omega(\tau, t) \equiv 1$ on π. For every site v on π, at least one of the following holds.

(a). The clock of v rings in the interval (τ, t).

(b). $\omega_v(\tau) = 1$ but the occupied cluster of v in $\omega(\tau)$ is finite.

Note that we have used that almost surely, none of the Poisson clocks ring at exactly time τ or t. If (a) occurs, we define

$$i_v := \min\{i : i \geq j \text{ and the clock of } v \text{ rings in } (\tau_i, t_i)\}.$$

Note that by monotonicity, the clock of v rings in (τ_i, t_i) for all $i \geq i_v$. If (a) does not occur, (b) must occur and there is a finite set K of sites on which $\omega(\tau) \equiv 0$ and which separates v from infinity. Now define

$$i_v := \min\{i : i \geq j \text{ and } \omega(\tau_i) \equiv 0 \text{ on } K\}.$$

Note that we have used the finiteness of K in the definition of i_v. Now let

$$k := \max_{v \in \pi} i_v.$$

Such k exists by finiteness of π. From the above procedure it follows that for each v on π at least one of the following holds.

(a'). The clock of v rings in the interval (τ_k, t_k).

(b'). $\omega_v(\tau_k) = 1$ but the occupied cluster of v in $\omega(\tau_k)$ is finite.

Hence, $\omega(\tau_k, t_k) \equiv 1$ on π. Further, since $k \geq j$, and again by monotonicity, $\omega(\tau_k, t_k) \equiv 1$ on I_j. Since π is a path from O to I_j this implies that I_j is contained in C_{τ_k, t_k} so that $|C_{\tau_k, t_k}| = \infty$. This shows (2.4.7).

II) and IV) Note that, now we have established convergence inside the quadrants I and III, this includes convergence along both axes through (τ, t). We use the monotonicity in τ and t to obtain convergence for paths that lie in the quadrants II and IV. Observe, for $(\tau_i, t_i) \to (\tau, t)$ lying in II,

$$\widehat{\theta}(\tau_i, t) \leq \widehat{\theta}(\tau_i, t_i) \leq \widehat{\theta}(\tau, t_i).$$

The sequence $(\tau_i, t)_i$ lies inside quadrant III and the sequence $(\tau, t_i)_i$ lies in quadrant I, so both the upper and lower bound in the equation above converge to $\widehat{\theta}(\tau, t)$ for $i \to \infty$. In the same way, for $(\tau_i, t_i) \to (\tau, t)$ in IV,

$$\widehat{\theta}(\tau, t_i) \leq \widehat{\theta}(\tau_i, t_i) \leq \widehat{\theta}(\tau_i, t).$$

Again both the upper and lower bound in the equation above converge to $\widehat{\theta}(\tau, t)$. This shows the required convergence in the quadrants II and IV.

□

Remark 2.4.4. We do not know the value of $\theta(p, \delta_c(p))$, so that the argumentation above does not show continuity in the point $(p, \delta_c(p))$. In dimension 2, it can be shown that *if* the self-destructive percolation model satisfies the RSW-property (in particular (A.4) when ε approaches 1) $\theta(p, \delta_c(p)) = 0$ follows for any p. Hence, by the theorem above, the percolation function will then be continuous in the critical points $(p, \delta_c(p))$, for $p \in (1/2, 1)$.

2.4.2 A percolation-like critical value

Finally, it is time to present the conjecture on which our claim of a discontinuity of the percolation function for the self-destructive percolation model on $\mathbb{Z}^2$ is based. Note that the continuity result in the previous subsection holds in all dimensions. The claimed discontinuity highly depends on two-dimensional arguments. So from now on, assume that $d = 2$.

Let $n > 0$. We consider the rectangle $R(n) = [0, 2n] \times [0, n]$ centred inside a box $G(n) = [-n, 3n] \times [-n, 2n]$. Inserting edges inherited from the square lattice, we consider $G(n)$ as a graph. Let $u = u(n)$ be the set $[0, 2n] \times \{n\}$ and $l = l(n)$ the set $[0, 2n] \times \{0\}$, i.e. the upper and lower side of the rectangle $R(n)$. Define, in the notation of Section 2.2,

$$p_n(\delta) := \mathcal{P}_{p_c,\delta}^{[G(n);\partial G(n)]}(l \leftrightarrow u \text{ inside } R(n)). \tag{2.4.8}$$

In words, we have performed critical percolation, then made vacant all sites that have an occupied path to $\partial G(n)$ and finally gave all sites an independent δ-enhancement. In the notation of Section 2.2, the X-variables are independent Bernoulli variables with parameter p_c; the Y-variables are independent Bernoulli variables with parameter δ. Then $p_n(\delta)$ denotes the probability that in the final configuration, there is a Z-occupied top-down crossing inside the rectangle $R(n)$. See Figure 2.3.

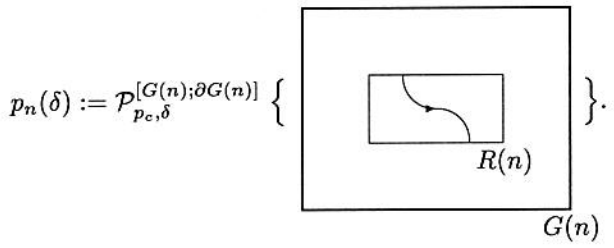

Figure 2.3: A picture of $p_n(\delta)$.

It is straightforward to see that in the final configuration, a site v in $R(n)$ is Z-occupied with probability

$$p_c - \mathcal{P}_{p_c}(v \leftrightarrow \partial G(n)) + \Big(1 - p_c + \mathcal{P}_{p_c}(v \leftrightarrow \partial G(n))\Big)\delta.$$

If n grows, v gets further away from $\partial G(n)$ and the probability above converges to $p_c + (1 - p_c)\delta$. Although this is larger than p_c, there are strong spatial dependencies. Hence, it is not clear whether the configuration is 'essentially supercritical', so that the probabilities $p_n(\delta)$ converge[2] to 1 as $n \to \infty$. Mind that this problem is not

[2]Recall that such crossing probabilities *do* tend to 1 for ordinary supercritical percolation.

answered by the weak convergence result stated in Lemma 2.2.5. Nevertheless, it is clear that $p_n(\delta)$ is increasing in δ so we define a critical value

$$\hat{\delta}_c = \sup\{\delta : p_n(\delta) \text{ is bounded away from 1, uniformly in } n\}. \tag{2.4.9}$$

Conjecture 2.4.5. $\hat{\delta}_c > 0$.

In spite of serious attempts, no proof or disproof of this conjecture has been found. It is supported by computer simulations, but since the size of the box our simulations can handle is limited (up to $n \approx 16000$), one should be careful interpreting such results.

Although it is (in our view) a hard problem to decide whether $p_n(\delta)$ is bounded away from 1, it is easy to see that $p_n(\delta)$ is bounded away from zero for all $\delta > 0$. As the distance between $\partial G(n)$ and $\partial R(n)$ is of order n, there will be a vacant circuit in $G(n)$ surrounding $R(n)$ with positive probability, bounded away from 0. This follows from RSW-arguments for ordinary percolation, compare Corollary A.5. The vacant circuit prevents occupied sites inside $R(n)$ to be connected to $\partial G(n)$ and thus from being removed. Hence the situation in $R(n)$ is like ordinary supercritical percolation after the δ-enhancement. For ordinary supercritical percolation, the probability that there is an occupied top-down crossing in $R(n)$ tends to 1 as n goes to infinity. We conclude that $p_n(\delta)$ is bounded away from zero.

The question remains whether $p_n(\delta)$ is bounded away from 1 for some $\delta > 0$, or equivalently, whether $\hat{\delta}_c > 0$. The rest of this subsection is devoted to other, but similar, critical values. The following lemmas show that the precise form of $R(n)$ and $G(n)$ are of little importance in the statement (and possible proof) of the conjecture.

Instead of removing occupied sites when they are connected to the boundary $\partial G(n)$, we could remove from the boundary of other rectangles. Fix integers k_1, k_2 and let for each n, $Q(n)$ be the rectangle $[0, k_1 n] \times [0, k_2 n]$. Let $\mathcal{G}_n, n \geq 0$ be such that the $Q(n)$ are contained in $\mathcal{G}_n$. Then we define

$$p_n(\delta, \mathcal{G}_n) := \mathcal{P}_{p_c,\delta}^{[\mathcal{G}_n; \partial \mathcal{G}_n]}(l \leftrightarrow u \text{ inside } Q(n)),$$

and

$$\hat{\delta}_c^{\mathcal{G}_n} := \sup\{\delta : p_n(\delta, \mathcal{G}_n) \text{ is bounded away from 1, uniformly in } n\}.$$

Obviously, $\hat{\delta}_c^{\mathcal{G}_n}$ depends on the choice of the graphs $\mathcal{G}_n$. However, it does not seem to matter much where the boundary $\partial \mathcal{G}_n$ is exactly, as long as it is sufficiently far from $Q(n)$. In particular, we prove the following.

Lemma 2.4.6. *Let k_1,k_2 and $Q(n)$ be as above. Let $l = l(n)$ denote the lower boundary $[0, k_1 n] \times \{0\}$ and $u = u(n)$ the upper boundary $[0, k_1 n] \times \{k_2 n\}$ of $Q(n)$. Further, fix positive constants K_1 and K_2. Suppose that $\mathcal{G}_{2,n} \supseteq \mathcal{G}_{1,n}$ are rectangles containing $Q(n)$ such that*

$$\begin{aligned} \min\{|v - w| : v \in \partial Q(n), w \in \partial \mathcal{G}_{1,n}\} &\geq K_1 n, \\ \max\{|v - w| : v \in \partial Q(n), w \in \partial \mathcal{G}_{2,n}\} &\leq K_2 n, \end{aligned} \tag{2.4.10}$$

where $|\cdot|$ denotes graph distance. Then $p_n(\delta, \mathcal{G}_{1,n})$ is bounded away from 1 if and only if $p_n(\delta, \mathcal{G}_{2,n})$ is bounded away from 1. Equivalently, $\hat{\delta}_c^{\mathcal{G}_{1,n}} > 0$ if and only if $\hat{\delta}_c^{\mathcal{G}_{2,n}} > 0$.

Proof. Suppose that $\mathcal{G}_{i,n}, i = 1, 2$ are such that they satisfy (2.4.10). We omit the subscript n for the rest of this proof.

One of the implications is trivial. Suppose that there exists a $\delta > 0$ such that $\sup_n p_n(\delta, \mathcal{G}_2) < 1$. It is easy to see, since $\mathcal{G}_1$ lies inside $\mathcal{G}_2$ and every site that is X-connected to $\partial\mathcal{G}_2$ is also X-connected to $\partial\mathcal{G}_1$, that $p_n(\delta, \mathcal{G}_1) \leq p_n(\delta, \mathcal{G}_2)$. Hence,

$$\sup_n \mathcal{P}_{p_c,\delta}^{[\mathcal{G}_1;\partial\mathcal{G}_1]}(l \leftrightarrow u \text{ inside } Q(n)) \leq \sup_n \mathcal{P}_{p_c,\delta}^{[\mathcal{G}_2;\partial\mathcal{G}_2]}(l \leftrightarrow u \text{ inside } Q(n)) < 1. \quad (2.4.11)$$

Now suppose there exists a $\delta > 0$ such that

$$\sup_n p_n(\delta, \mathcal{G}_1) < 1. \quad (2.4.12)$$

Let E_n denote the event that there is a X-occupied circuit between $\partial Q(n)$ and $\partial\mathcal{G}_1$ *and* an X-occupied path from $\partial Q(n)$ to $\partial\mathcal{G}_2$. The event E_n has positive probability, bounded away from zero and one uniformly in n, by the RSW-property of ordinary percolation and (2.4.10). On the event E_n, each site inside $Q(n)$ that is connected to $\partial\mathcal{G}_1$ by an X-occupied path, is also connected to the X-occupied circuit and $\partial\mathcal{G}_2$. Now we have, for all n,

$$\begin{aligned} p_n(\delta, \mathcal{G}_2) &\leq \mathcal{P}_{p_c,\delta}^{[\mathcal{G}_2;\partial\mathcal{G}_2]}(l \leftrightarrow u \text{ inside } Q(n) \text{ and } E_n \text{ holds}) + \mathcal{P}_{p_c}(E_n^c) \\ &= \mathcal{P}_{p_c,\delta}^{[\mathcal{G}_2;\partial\mathcal{G}_1]}(l \leftrightarrow u \text{ inside } Q(n) \text{ and } E_n \text{ holds}) + \mathcal{P}_{p_c}(E_n^c) \\ &\leq \mathcal{P}_{p_c,\delta}^{[\mathcal{G}_1;\partial\mathcal{G}_1]}(l \leftrightarrow u \text{ inside } Q(n))\mathcal{P}_{p_c}(E_n) + \mathcal{P}_{p_c}(E_n^c). \end{aligned}$$

In the second inequality we have applied the FKG-inequality: the event E_n depends only on, and is increasing in, the X-variables inside $\mathcal{G}_2 \backslash Q(n)$ whereas the event that there finally exists an occupied crossing in $Q(n)$ is decreasing in those variables. Finally we use the assumption (2.4.12) and the fact that $\mathcal{P}_{p_c}(E_n)$ is bounded away from zero and one to conclude that $\sup_n p_n(\delta, \mathcal{G}_2) < 1$. □

The previous lemma showed that, given the inner rectangle $Q(n)$, the conjecture does not change in a crucial way when we move the outer boundary from which occupied sites are destroyed. Similarly, for the statement (and possible proof) of Conjecture 2.4.5, the shape of the inner rectangle is not essential as we will prove in the following lemma. We compare the rectangle $R = [0, 2n] \times [0, n]$ and the corresponding probability $p_n(\delta)$ with a rectangle of different size. In the following lemma, we take this rectangle of size n by $6n$ (it will be used in the following subsection) but clearly the lemma can be proved for other sizes as well.

Lemma 2.4.7. *Let $R_1 = R_1(n)$ be the rectangle $[0, 6n] \times [0, n]$ contained in a rectangle $R_2 = R_2(n)$ where the distance between ∂R_1 and ∂R_2 is of order n. Let $l = l(n)$ be*

the set $[0,6n]\times\{0\}$ *and* $u=u(n)$ *be the set* $[0,6n]\times\{n\}$. *Then*

$$p'_n(\delta) := \mathcal{P}_{p_c,\delta}^{[R_2(n);\partial R_2(n)]}(l \leftrightarrow u \text{ inside } R_1(n))$$

is bounded away from 1, *uniformly in* n *if and only if* $p_n(\delta)$ *is bounded away from* 1, *uniformly in* n.

Proof. One of the implications is quite simple. Suppose that $\sup_n p'_n(\delta) < 1$. An application of Lemma 2.4.6 shows that the probabilities

$$\mathcal{P}_{p_c,\delta}^{[G(3n);\partial G(3n)]}(l \leftrightarrow u \text{ inside } R_1(n)), \tag{2.4.13}$$

are bounded away from 1. Note that it is easier to have an occupied top-down crossing in $R_1(n)=[0,6n]\times[0,n]$ than in $R(n)=[0,2n]\times[0,n]$. This implies that the probabilities

$$\mathcal{P}_{p_c,\delta}^{[G(3n);\partial G(3n)]}(l \leftrightarrow u \text{ inside } R(n))$$

are bounded by the probabilities in (2.4.13) and hence, bounded away from 1. Another application of Lemma 2.4.6 shows that then also $\sup_n p_n(\delta) < 1$.

The other implication requires a bit more work. Suppose that $\sup_n p_n(\delta) < 1$. In fact, we will show that

$$1 - \sup_n p'_n(\delta) = \inf_n \mathcal{P}_{p_c,\delta}^{[R_2(n);\partial R_2(n)]}(l \not\leftrightarrow u \text{ inside } R_1(n)) > 0. \tag{2.4.14}$$

We apply the same trick as in the proof of Lemma 2.3.3. An occupied crossing from $l(n)$ to $u(n)$ lies either completely inside one of the $2n$ by n rectangles $[kn,(k+2)n]\times[0,n]$ for $k=0,\ldots,4$, or if it is not inside one of these rectangles, at least one of the n by n squares $[ln,(l+1)n]\times[0,n]$ for $l=1,\ldots,4$ has an occupied left-right crossing. We should show that the probability that none of these events occur, is bounded away from zero. Applying the association property (Lemma 2.2.7) we get

$$\begin{aligned}
&\mathcal{P}_{p_c,\delta}^{[R_2(n);\partial R_2(n)]}(l \not\leftrightarrow u \text{ inside } R_1(n)) \geq \\
&\quad \prod_{k=0}^{4} \mathcal{P}_{p_c,\delta}^{[R_2(n);\partial R_2(n)]}(l \not\leftrightarrow u \text{ inside } [kn,(k+2)n]\times[0,n]) \\
&\quad \times \prod_{l=1}^{4} \mathcal{P}_{p_c,\delta}^{[R_2(n);\partial R_2(n)]}(\{ln\}\times[0,n] \not\leftrightarrow \{(l+1)n\}\times[0,n] \text{ inside } [ln,(l+1)n]\times[0,n])
\end{aligned} \tag{2.4.15}$$

A direct application of Lemma 2.4.6 (and a translation) shows, by the assumption $\sup_n p_n(\delta) < 1$, that the terms in the first product on the right of (2.4.15) are bounded away from zero. Another application of Lemma 2.4.6 (and translation plus a rotation of the whole system over 90 degrees) shows that for $l \in \{1,\ldots,4\}$,

$$\inf_n \mathcal{P}_{p_c,\delta}^{[R_2(n);\partial R_2(n)]}(\{ln\}\times[0,n] \not\leftrightarrow \{(l+1)n\}\times[0,n] \text{ inside } [ln,(l+1)n]\times[0,n]) > 0$$

if and only if

$$\inf_n \mathcal{P}_{p_c,\delta}^{[G(n);\partial G(n)]}([n,2n]\times\{0\} \not\leftrightarrow [n,2n]\times\{n\} \text{ inside } [n,2n]\times[0,n]) > 0. \tag{2.4.16}$$

Further we observe that the probability of crossing a square is smaller than the probability of crossing a rectangle in the shortest direction, so that

$$\begin{aligned}
&\inf_n \mathcal{P}_{p_c,\delta}^{[G(n);\partial G(n)]}([n,2n]\times\{0\} \not\leftrightarrow [n,2n]\times\{n\} \text{ inside } [n,2n]\times[0,n]) \\
&\quad\geq \inf_n \mathcal{P}_{p_c,\delta}^{[G(n);\partial G(n)]}([0,2n]\times\{0\} \not\leftrightarrow [0,2n]\times\{n\} \text{ inside } [0,2n]\times[0,n]) \\
&\quad\geq 1-\sup_n p_n(\delta) > 0.
\end{aligned} \tag{2.4.17}$$

Combining (2.4.17), (2.4.16) and (2.4.15) shows (2.4.14). □

Finally, we present a stronger conjecture, which we used in [BB04b]. Consider the graph $R(n) = [0,3n]\times[0,2n]$ with the edges inherited from the square lattice. Let l be the lower side of the rectangle, $l = l(n) = [0,3n]\times\{0\}$, $m = m(n)$ the horizontal middle line $m(n) = [0,3n]\times\{n\}$ and $u = u(n)$ the upper side, $u(n) = [0,3n]\times\{2n\}$. Then we define

$$a_n(\delta) := \mathcal{P}_{p_c,\delta}^{[R(n);u(n)]}(l \leftrightarrow m).$$

At first we thought that $\sup_n a_n(\delta) < 1$ if and only if $\sup_n p_n(\delta) < 1$. That the former implies the latter is easily seen from an easy comparison argument and Lemma 2.4.6. However, we have not been able to show the reversed implication.

2.4.3 A discontinuity at p_c on $\mathbb{Z}^2$?

Recall the definitions of $p_n(\delta)$ and $\hat{\delta}_c$ from the previous subsection.

Theorem 2.4.8. *Let $d = 2$ and suppose that $\hat{\delta}_c > 0$. Then there exists a $\delta > 0$ such that*

$$\forall p > p_c, \quad \theta(p,\delta) = 0.$$

In words, the theorem states that if Conjecture 2.4.5 holds, then no matter how small the fraction of sites we remove, no matter how sparse the infinite cluster becomes, its geometry prevents a new infinite cluster to arise 'soon'. The critical probability $\delta_c(p)$ stays bounded away from zero, as p goes to p_c. In particular this implies that there exists a $\delta > 0$ such that

$$\lim_{p\downarrow p_c} \theta(p,\delta) \neq \theta(p_c,\delta).$$

In other words, the function $\theta(\cdot,\cdot)$ is not continuous at the point (p_c,δ). The proof of this theorem is split up in several lemmas. First we need some notation. Let

$n \geq 0$, $B(n) = [-n, n]^2$ and $\partial B(n)$ the boundary of the box $B(n)$. We define a new probability $q_n(p, \delta)$ as follows. Consider the annulus $A(n, 5n) = B(5n) \backslash B(n)$. First make each site in $A(n, 5n)$ occupied with probability p. Then make each site whose occupied cluster inside $A(3n, 5n)$ contains a contour (circuit) around $B(3n)$ vacant. Finally enhance each vacant site with probability δ. Now $q_n(p, \delta)$ is the probability that in the final configuration we cross the annulus $A(n, 3n)$, or in other words, $\partial B(n) \leftrightarrow \partial B(3n)$. For simplicity, we often write 'contour in $A(3n, 5n)$' instead of 'contour in $A(3n, 5n)$ around $B(3n)$'.

Lemma 2.4.9. *If δ is such that the sequence $p_n(\delta)$ is uniformly bounded away from 1, then the sequence $q_n(p_c, \delta)$ is bounded away from 1, uniformly in n.*

Proof. Suppose that $p_n(\delta)$ is bounded away from 1. From Lemma 2.4.7 we conclude that also $p'_n(\delta)$, the analogue of $p_n(\delta)$ on a $6n$ by n box, is bounded away from 1. In this proof we write p'_n for $p'_n(\delta)$ and q_n for $q_n(p_c, \delta)$. We will show

$$\begin{aligned} q_n &\leq \mathcal{P}_{p_c}(\text{there is no occupied contour in } A(3n, 5n)) \\ &+ (1 - (1 - p'_n)^4)\mathcal{P}_{p_c}(\exists \text{ occupied contour in } A(3n, 5n)). \end{aligned} \tag{2.4.18}$$

From this the lemma easily follows, once again using the RSW-property.

Analogously to what we did earlier, we let $(X_i : i \in A(n, 5n))$ be independent Bernoulli variables with parameter p and $(Y_i : i \in A(n, 5n))$ be independent Bernoulli variables with parameter δ. We now define for $i \in A(n, 5n)$

$$\begin{aligned} X_i^{*c} &:= I(X_i = 1 \text{ and the } X\text{-occupied cluster of } i \text{ in } A(n, 5n) \\ &\qquad \text{does not contain a contour around } B(3n)), \end{aligned}$$

$$\begin{aligned} X_i^{*B} &:= I(X_i = 1 \text{ and } \nexists\, X\text{-occupied path in } A(n, 5n) \\ &\qquad \text{from } i \text{ to } \partial B(5n)), \end{aligned}$$

where $I(\cdot)$ denotes the indicator function. The introduction of the following variables should not come as a surprise.

$$\begin{aligned} Z_i^c &:= X_i^{*c} \vee Y_i, \\ Z_i^B &:= X_i^{*B} \vee Y_i, \end{aligned}$$

for all $i \in A(n, 5n)$. It is clear that for each $i \in A(n, 3n)$, each X-occupied path from i to $\partial B(5n)$ intersects all contours in $A(3n, 5n)$ around $B(3n)$. So on the event that there exists an X-occupied contour in $A(3n, 5n)$ we have for all $i \in A(n, 3n)$,

$$X_i^{*c} \leq X_i^{*B}.$$

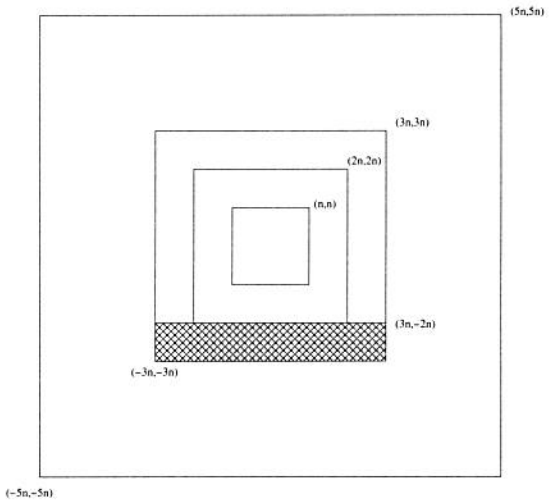

Figure 2.4: The annulus $A(n, 5n)$. The shaded region is the rectangle R_1.

Now we compute

$$\begin{aligned} q_n &= \mathcal{P}(\exists Z^c\text{-occupied path from } \partial B(n) \text{ to } \partial B(3n)) \\ &\leq \mathcal{P}(\nexists X\text{-occupied contour in } A(3n, 5n)) \\ &+ \mathcal{P}(\exists X\text{-occupied contour in } A(3n, 5n) \text{ and} \\ &\qquad \exists Z^B\text{-occupied path from } \partial B(2n) \text{ to } \partial B(3n)) \\ &\leq \mathcal{P}(\nexists X\text{-occupied contour in } A(3n, 5n)) \\ &+ \mathcal{P}(\exists X\text{-occupied contour in } A(3n, 5n)) \\ &\qquad \times \mathcal{P}(\exists Z^B\text{-occupied path from } \partial B(2n) \text{ to } \partial B(3n)). \end{aligned} \tag{2.4.19}$$

In the last inequality we have applied the FKG-inequality: the event $\{\exists\ X$-occupied contour in $A(3n, 5n)\}$ is completely determined by, and increasing in, the variables $(X_i : i \in A(3n, 5n))$, whereas the event $\{\exists Z^B$-occupied path from $\partial B(2n)$ to $\partial B(3n)\}$ is clearly decreasing in those variables.

Let $R_1, \ldots, R_4$ be the four $n \times 6n$ rectangles whose union is $A(2n, 3n)$. More explicitly, let R_1 be the rectangle $[-3n, 3n] \times [-2n, -3n]$. Let l denote the set $[-3n, 3n] \times \{-3n\}$ and u the set $[-3n, 3n] \times \{-2n\}$. See Figure 2.4. Each occupied path from $\partial B(2n)$ to $\partial B(3n)$ must contain an occupied crossing of one of the rectangles $R_1, \ldots, R_4$ in the shortest direction. Furthermore, by symmetry each of

these rectangles play the same role. In the usual notation, we have

$$\begin{aligned}
&\mathcal{P}(\exists Z^B\text{-occupied path from } \partial B(2n) \text{ to } \partial B(3n)) \\
&\quad = \mathcal{P}_{p_c,\delta}^{[A(n,5n);\partial B(5n)]}(\partial B(2n) \leftrightarrow \partial B(3n)) \\
&\quad \leq 1 - \left[1 - \mathcal{P}_{p_c,\delta}^{[A(n,5n);\partial B(5n)]}(l \leftrightarrow u \text{ inside } R_1)\right]^4,
\end{aligned} \tag{2.4.20}$$

where we have used the association property (Lemma 2.2.7). The rectangle R_1 is a $6n$ by n rectangle so we may apply Lemma 2.4.6, or rather the remark directly before equation (2.4.11), to the sequence p'_n and the probabilities on the right of (2.4.20). It follows that the term on the right hand side of (2.4.20) is at most $1 - (1 - p'_n)^4$. Combining this with (2.4.19) shows (2.4.18). □

Lemma 2.4.10. *If δ is such that $q_n(p_c, \delta)$ is bounded away from 1, uniformly in n, then there exists $\varepsilon > 0$ such that the family $q_n(p, \varepsilon), n = 1, 2, \ldots$ and $p \in [p_c, p_c + \varepsilon]$, is bounded away from 1.*

Proof. Note that (in a similar way as in the proof of Lemma 2.2.14),

$$q_n\left(p, \frac{r-p}{1-p}\right) \text{ is decreasing in } p \text{ for fixed } r. \tag{2.4.21}$$

Take ε so small that $(p_c + \varepsilon) + (1 - p_c - \varepsilon)\varepsilon < p_c + (1 - p_c)\delta$. Then apply (2.4.21) and the fact that $q_n(p, \frac{r-p}{1-p})$ is increasing in r. □

Lemma 2.4.11. *If δ is such that the sequence $p_n(\delta)$ is bounded away from 1 uniformly in n, then, with ε as in the previous lemma,*

$$\theta(p, \varepsilon) = 0, \forall p \in (p_c, p_c + \varepsilon].$$

Remark 2.4.12. In fact we prove a much stronger result: let δ be such that the sequence $p_n(\delta)$ is bounded away from 1, uniformly in n. Then, with ε as in Lemma 2.4.10,

$$\exists \alpha > 0 \text{ s.t. } \forall p \in (p_c, p_c + \varepsilon] \; \exists C \text{ s.t. } \forall n \; \mathcal{P}_{p,\varepsilon}(O \leftrightarrow \partial B(n)) \leq C n^{-\alpha}. \tag{2.4.22}$$

Proof. [Lemma 2.4.11] Let δ and ε be as in the statement of the lemma. Then, by Lemma 2.4.10, we can find a $q < 1$ such that for all n and all $p \in [p_c, p_c + \varepsilon]$, $q_n(p, \varepsilon) < q$. Let $n \geq 1$ and choose $p \in (p_c, p_c + \varepsilon]$. Let C_1 be a positive constant, depending on p. Later on, we will make this more precise. Take the variables X_i, Y_i, X_i^* and Z_i, $i \in \mathbb{Z}^2$ as in the beginning of Section 2.2. In addition we now take for all $i \in \mathbb{Z}^2$,

$$X_i^{*C} := I(X_i = 1 \text{ and its } X\text{-occupied cluster contains no contours around } B(C_1 \log n)),$$

and

$$Z_i^C := X_i^{*C} \vee Y_i.$$

Note that on the event that there exists an infinite X-occupied path from $\partial B(C_1 \log n)$, we have $X_i^* \leq X_i^{*C}$ for all $i \in \mathbb{Z}^2 \backslash B(C_1 \log n)$, since this infinite path intersects all contours around $B(C_1 \log n)$. So we have

$$\begin{aligned} \mathcal{P}_{p,\varepsilon}(O \leftrightarrow \partial B(n)) &= \mathcal{P}(\exists Z\text{-occupied path from } O \text{ to } \partial B(n)) \\ &\leq \mathcal{P}(\exists Z\text{-occupied path from } \partial B(C_1 \log n) \text{ to } \partial B(n)) \\ &\leq \mathcal{P}(\exists Z^C\text{-occupied path from } \partial B(C_1 \log n) \text{ to } \partial B(n)) \\ &+ \mathcal{P}(\nexists \text{ infinite } X\text{-occupied path from } \partial B(C_1 \log n)) \end{aligned} \tag{2.4.23}$$

From standard percolation theory we know that the last term on the right in (2.4.23) is at most

$$C_2(p)e^{-\beta_1(p)C_1 \log n},$$

where $C_2(p)$ and $\beta_1(p)$ are positive constants depending only on p. Now we handle the other term on the right of (2.4.23). Let $I = I(n)$ denote the set of all positive even integers i with $C_1 \log n < 3^i < 5 \times 3^i < n$ and consider the annuli $A(3^i, 5 \times 3^i), i \in I$. These annuli are pairwise disjoint so using the obvious domination of X_j^{*C} by X_j^{*c} for $j \in A(3^i, 5 \times 3^i)$, $i \in I$,

$$\mathcal{P}(\exists Z^C\text{-occupied path from } \partial B(C_1 \log n) \text{ to } \partial B(n)) \leq \prod_{i \in I} q_{3^i}(p, \varepsilon) < q^{|I|}, \tag{2.4.24}$$

where $|I|$ denotes the size of the set I. From the definition of I it is clear that there exist constants $\beta_2 > 0$ and $C_3 > 0$, (where the latter is a function of C_1) such that the right hand side of (2.4.24) is at most $C_3 n^{-\beta_2}$. So for all $C_1 > 0$, we can find a $C_3 > 0$ such that for all $p \in (p_c, p_c + \varepsilon]$ and for all sufficiently large n,

$$\mathcal{P}_{p,\varepsilon}(O \leftrightarrow \partial B(n)) \leq C_2(p)e^{-\beta_1(p)C_1 \log n} + C_3 n^{-\beta_2}.$$

The claim (2.4.22) now follows by choosing $C_1 = C_1(p)$ so large that $C_1\beta_1(p) \geq \beta_2$ and finally choosing the multiplicative constant C so large that the inequality (2.4.22) holds for all n. This proves Lemma 2.4.11. □

Theorem 2.4.8 now follows easily from this lemma, combined with Proposition 2.3.4 (applied with $p = p_c + \varepsilon$).

2.5 The permanent self-destructive process on $\mathbb{Z}^2$

Informally, consider the following process in time on the square lattice: At time 0 each site is vacant. Each vacant site becomes occupied at rate 1, independent of the other sites. But every time an infinite cluster occurs, it is destroyed immediately. After destruction the process continues and each vacant site becomes occupied at rate 1 as before. An interesting question is whether this description makes sense.

The permanent self-destructive process can be seen as a limit of a forest fire model on a finite box in $\mathbb{Z}^2$. As above, initially everything is vacant and sites become occupied by a tree at rate 1. Further, ignition attempts are made at rate λ at each site, independent of the growth of trees. If at the time of such an ignition attempt, the corresponding site is occupied, the tree is set on fire. A tree burns for an exponentially distributed time with expectation μ, during which it spreads fire to its (occupied) neighbours at rate κ. When a tree is burned down its site is vacant again and becomes occupied at rate 1 again etc. Special attention is given in the literature to what happens (to the equilibrium distribution) when we send the size of the box and the fire spread rate κ to infinity, and the ignition rate λ and burning time μ to zero, in some suitable way. It seems, although most of the results are non-rigorous, that this leads to so-called self-organised critical behaviour, see e.g. [Jen98]. The model described above is a version of the so-called Drossel-Schwabl forest fire model [DS92]. We deal with this and similar models in Chapter 3.

From a mathematical point of view it is natural to ask whether the forest-fire process is still well-defined if we take all the parameters equal to the above mentioned limiting values. We can interpret $\lambda = 0$ by saying that only infinite clusters will be ignited. A natural interpretation of $\kappa = \infty$ and $\mu = 0$ could be that these clusters are burnt instantaneously. These considerations lead to the question above. Is there a mathematically consistent definition of the permanent self-destructive process that satisfies the informal description above? The question of existence is closely related to a problem posed by Aldous [Ald00]. An analogue of this question in the one-dimensional case is answered affirmatively by van den Berg and Tóth in [BT01]. In that paper some further background can be found.

Formally, our question is as follows:

Question 2.5.1. Do there exists processes $\chi_i : \mathbb{R}_+ \to \{0,1\}, i \in \mathbb{Z}^2$, defined jointly on some probability space, with the following properties:

(i). Almost surely, for all $i \in \mathbb{Z}^2, \chi_i(0) = 0$.

(ii). Almost surely, for all $i \in \mathbb{Z}^2, t \mapsto \chi_i(t)$ is continuous from the right having left limits (c.a.d.l.a.g.).

(iii). Let $U_i^{(k)}$ denote the length of the kth interval during which $\chi_i(.)$ equals 0. Then the random variables $(U_i^{(k)})_{i \in \mathbb{Z}^2, k \in \mathbb{N}}$ are independent, exponentially distributed with mean 1.

(iv). Almost surely, for all $t \in \mathbb{R}_+$ and $i \in \mathbb{Z}^2$ with $\chi_i(t^-) = 1$; if there is an infinite path π from i having $\chi_j(t^-) = 1$ for all j on π then $\chi_i(t) = 0$, else $\chi_i(t) = 1$.

Theorem 2.5.2. *If Conjecture 2.4.5 holds, then there are* no *processes* $\chi_i : \mathbb{R}_+ \to \{0,1\}, i \in \mathbb{Z}^2$ *with properties (i)-(iv) above.*

The reason for this is, loosely said, the following. We have seen in the previous section that if Conjecture 2.4.5 is true, it takes an essential probability, or essential time, to form a new occupied infinite cluster after removal of an infinite cluster. Let $\hat{t}$ be slightly larger than t_c. With positive probability, at some time $t \in (t_c, \hat{t})$, an infinite cluster arises, and is hence destroyed instantaneously. By ordinary percolation results, there must have been an infinite cluster at some time between t_c and t, so there must have been an earlier fire. But it takes an essential time to reintroduce an infinite cluster, so the fire at time t could not have been there.

Before we prove Theorem 2.5.2 we state a lemma that we need in the proof. Let $k \in \mathbb{N}$ and $\delta > 0$. Suppose initially all sites of the square lattice are vacant. We occupy each site independently with probability p_c and then remove all sites whose occupied cluster contains a contour around the box $B(m)$. Finally we occupy all vacant sites with probability δ. Let $\mathcal{P}^{(m)}_{p_c,\delta}$ denote the distribution of the final configuration. Recall the definition of $p_n(\delta)$ from (2.4.8).

Lemma 2.5.3. *If* δ *is such that* $\sup_n p_n(\delta) < 1$ *(i.e. Conjecture 2.4.5 holds) then*

$$\forall m, \quad \mathcal{P}^{(m)}_{p_c,\delta}(\exists \text{ an infinite } Z\text{-occupied cluster}) = 0.$$

The proof of this lemma is very similar to (and in fact simpler than) the proof of Theorem 2.4.8, so we do not present it here.

Proof. [Theorem 2.5.2] Suppose we have $\delta > 0$ such that the sequence $(p_n(\delta),\ n \in \mathbb{N})$ is bounded away from 1. Assume we do have processes $\chi_i(\cdot)$ with the properties (i)-(iv). We will show that this leads to a contradiction. Let the $U_i^{(k)}$'s be as in property (iii).

Now first consider, for each m, the following process: initially all sites are vacant. Vacant sites become occupied at rate 1. At time t_c (defined by the relation $1 - e^{-t_c} = p_c$), each site whose occupied cluster contains a contour around $B(m)$, become vacant. After that, vacant sites become occupied again at rate 1. Let $\eta_i^{(m)}(t)$ be the state of site i at time t. All these processes with different m can be simply coupled together, and with the χ process, by using the U_i variables mentioned in property (iii) of the χ-process. In fact, for this coupling we only use the $U^{(1)}$'s and $U^{(2)}$'s exactly in the same way as in the Lemma 2.2.14, with the τ now equal to t_c, and with an obvious modification of 'catastrophe': a site i is initially vacant and becomes occupied at time $U_i^{(1)}$. If this time is larger than t_c, it then remains occupied forever. However, if at time t_c its occupied cluster contains a contour around $B(m)$, it becomes vacant, and remains vacant until time $t_c + U_i^{(2)}$, after which it is occupied forever.

It is clear that, for $t > t_c$, the collection $(\eta_i^{(m)}(t),\ i \in \mathbb{Z}^2)$ has distribution $\mathcal{P}_{p_c,\delta}^{(m)}$, defined above Lemma 2.5.3, with $\delta = 1 - e^{-(t-t_c)}$. Hence, by that lemma, the assumption in the beginning of this proof, and because each $\eta_i^{(m)}(t)$ is increasing in $t > t_c$, $\exists \varepsilon > 0$ such that for all m we have

$$\mathcal{P}(\exists \text{ infinite } \eta^{(m)}(t)\text{-occupied cluster for some } t < t_c + \varepsilon) = 0. \tag{2.5.1}$$

Now back to the $\chi_i(t)$'s. Let ε be as in (2.5.1). Clearly, by property (i) and (iii) of the χ process, there is a.s. no infinite χ-occupied cluster before or at time t_c, and hence, by property (iv), there are a.s. no site i and time $t \le t_c$ with $\chi_i(t^-) = 1$ and $\chi_i(t) = 0$. If, in addition, there is no $t \in (t_c, t_c + \varepsilon)$ and site i with $\chi_i(t^-) = 1$ and $\chi_i(t) = 0$, then $\chi_j(t) = I(U_j^{(1)} \le t)$, $j \in \mathbb{Z}^2$, $t \in (t_c, t_c + \varepsilon)$. However, with probability 1, for all $t > t_c$ there is an infinite cluster of sites $j \in \mathbb{Z}^2$ with $U_j^{(1)} < t$. This gives a contradiction with property (iv). So, almost surely, there is a $t \in (t_c, t_c + \varepsilon)$ and a site i s.t. $\chi_i(t^-) = 1$ and $\chi_i(t) = 0$. And hence, there is a $\hat{t} \in (t_c, t_c + \varepsilon)$ such that

$$\mathcal{P}\big(\exists i \in \mathbb{Z}^2 \text{ and } s \in (\hat{t}, t_c + \varepsilon) \text{ with } \chi_i(s^-) = 1,\ \chi_i(s) = 0\big) > 0. \tag{2.5.2}$$

Fix such $\hat{t}$. Now, with exactly the same argument we can prove that, almost surely, the following event occurs:

$$\exists s \in (t_c, (t_c + \hat{t})/2) \text{ and an } i \in \mathbb{Z}^2 \text{ with } \chi_i(s^-) = 1,\ \chi_i(s) = 0.$$

But if that event occurs, then, according to the required properties of the χ-process, there is an infinite path on which the χ-values changes from 1 to 0 at that time s. But such path intersects, for some n, all contours around $B(n)$. By similar arguments as in the proof of Lemma 2.2.14 and Lemma 2.4.11 we conclude that a.s.

$$\exists m \text{ s.t. } \chi_j(t) \le \eta_j^{(m)}(t), \text{for all } j \in \mathbb{Z}^2 \text{ and } t > (t_c + \hat{t})/2. \tag{2.5.3}$$

Combining (2.5.1) and (2.5.3) contradicts (2.5.2). □

2.6 The binary tree

2.6.1 The rooted binary tree

We turn our attention to the rooted binary tree $\mathcal{T}$. This is the infinite tree in which one site, which we call the root, has two neighbours, and all other sites have three neighbours. The root will simply be denoted by O and its two neighbours by 1 and 2. The sites 1 and 2 can be considered as roots of the subtrees $\mathcal{T}[1]$ and $\mathcal{T}[2]$, respectively, see Figure 2.5.

For ordinary site percolation on $\mathcal{T}$ with parameter p it is well-known that the critical probability equals $p_c = 1/2$ and that for $p \ge 1/2$

$$\theta(p) = \frac{2p-1}{p}.$$

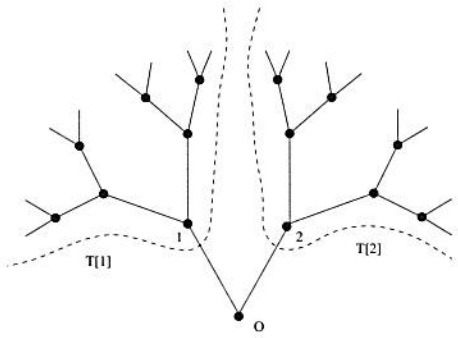

Figure 2.5: The rooted binary tree $\mathcal{T}$ and its subtrees $\mathcal{T}[1]$ and $\mathcal{T}[2]$.

In particular $\theta(p_c) = 0$. In fact, using that $\mathcal{T}[1]$ and $\mathcal{T}[2]$ are disjoint copies of $\mathcal{T}$ one has that $\theta(p)$ satisfies

$$\theta(p) = p(1 - (1 - \theta(p))^2). \tag{2.6.1}$$

We study the probability measures $\mathcal{P}_{p,\delta}^{[\mathcal{T};\infty]}$ and the function $\theta(p,\delta) = \mathcal{P}_{p,\delta}^{[\mathcal{T};\infty]}(0 \leftrightarrow \infty)$ for the self-destructive process on the tree. One noticeable difference between the binary tree and the square lattice is that on the binary tree there is no unique infinite cluster for ordinary percolation. In fact, there are infinitely many infinite clusters, which are all removed in self-destructive percolation model.

We show that for $\mathcal{T}$ the analogue of Proposition 2.3.2 is true, but also that the intuitive picture is correct, in the sense that $\delta_c(p) \downarrow 0$ as $p \downarrow p_c$. In other words, the limit (2.1.3) holds for $\mathcal{T}$; this is in contrast to what we conjectured for the square lattice (Section 2.3 and Subsection 2.4.3).

Theorem 2.6.1. *On the rooted binary tree $\mathcal{T}$ we have*

(i). $\forall p > p_c \ \exists \delta > 0$ *s.t.* $\theta(p,\delta) = 0$.

(ii). *If* $p(1-\delta) \geq p_c$, *then* $\theta(p,\delta) = 0$.

(iii). *For all* $\delta > 0$, *there exists a* $\hat{p} > p_c$ *and* $\alpha > 0$ *such that for all* $p \in [p_c, \hat{p}]$, $\theta(p,\delta) > \alpha$.

(iv). *Moreover,*

$$\lim_{p \downarrow p_c} \theta(p,\delta) = \theta(p_c,\delta) \quad \Big(= \theta(p_c + (1-p_c)\delta) > 0\Big). \tag{2.6.2}$$

Proof. As for ordinary percolation it is natural to use the fact that $\mathcal{T}[1]$ and $\mathcal{T}[2]$ are copies of $\mathcal{T}$. The removal of infinite clusters introduces dependencies between these copies, which complicates things. By defining suitable classes of random variables

these complications can be handled quite smoothly: let $(X_i : i \in V_{\mathcal{T}})$, and $(Y_i : i \in V_{\mathcal{T}})$ be as in the beginning of Section 2.2. So the X_i's, are i.i.d. Bernoulli variables with parameter p, and the Y_i's i.i.d. Bernoulli variables with parameter δ. The family of X_i's is independent of the family of Y_i's. We also define the X_i^* and Z_i, $i \in V_{\mathcal{T}}$ as in Section 2.2 with $\Gamma = \infty$. In addition to our $*$-operation we define a $*+$-operation on our variables by

$$X_i^{*+} = \begin{cases} 1 & \text{if } X_i = 1 \text{ and there is no } X\text{-occupied path from } i \text{ to } \infty \text{ or to } O, \\ 0 & \text{otherwise.} \end{cases}$$

It is clear from this definition that $X_0^{*+} = 0$.

We also introduce the 'natural analogues' of these variables for $\mathcal{T}[1]$ and $\mathcal{T}[2]$. Let $k \in \{1, 2\}$. Then for $i \in V_{\mathcal{T}[k]}$ we define

$$X_i^*[k] = \begin{cases} 1 & \text{if } X_i = 1 \text{ and there is no } X\text{-occupied path in } \mathcal{T}[k] \text{ from } i \text{ to } \infty, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$X_i^{*+}[k] = \begin{cases} 1 & \text{if } X_i = 1 \text{ and there is no } X\text{-occupied path in } \mathcal{T}[k] \text{ from} \\ & i \text{ to } \infty \text{ or } k \\ 0 & \text{otherwise.} \end{cases}$$

Further, for $i \in V_{\mathcal{T}[k]}$ and $k \in \{1, 2\}$,

$$Z_i[k] = X_i^*[k] \vee Y_i,$$

and

$$Z_i^+[k] = X_i^{*+}[k] \vee Y_i.$$

We proceed with the proof of Theorem 2.6.1. Note that if the Z-occupied cluster of O is infinite then the following event A must occur:

$$A := \{X_O \vee Y_O = 1 \text{ and } \exists i \in \{1, 2\} \text{ s.t. } i \text{ is on an infinite } Z[i]\text{-occupied path in } \mathcal{T}[i]\}.$$

Hence

$$\theta(p, \delta) \leq \mathcal{P}(A) = (p + (1 - p)\delta)\,[2\theta(p, \delta) - \theta(p, \delta))^2]. \tag{2.6.3}$$

On the other hand, if A occurs and the X-occupied cluster of O is finite then the Z-occupied cluster of O is infinite. This gives

$$\theta(p, \delta) \geq \mathcal{P}(A) - \theta(p) = (p + (1 - p)\delta)\,[2\theta(p, \delta) - \theta(p, \delta))^2] - \theta(p). \tag{2.6.4}$$

We can improve the quite trivial upper bound (2.6.3) as follows. First define the event

$$B := \{X_O = 1, Y_O = 0 \text{ and } \exists i \in \{1, 2\} \text{ s.t. } i \text{ is on an infinite } Z[i]\text{-occupied path in } \mathcal{T}[i] \text{ and } 3 - i \text{ is on an infinite } X\text{-occupied path in } \mathcal{T}[3 - i]\}.$$

It is easily seen that $B \subset A$, but also that if B-occurs, then $Z_O = 0$ so that O is not on an infinite Z-occupied path. Hence

$$\theta(p,\delta) \leq \mathcal{P}(A) - \mathcal{P}(B). \tag{2.6.5}$$

Now,

$$\begin{aligned}
\mathcal{P}(B) &= p(1-\delta) \times \\
&\quad \big[2\theta(p,\delta)\theta(p) - \mathcal{P}(i \text{ is on both an infinite } Z[i]\text{-occupied path} \\
&\quad \text{and on an infinite } X\text{-occupied path in } \mathcal{T}[i], i \in \{1,2\})\big] \\
&\geq p(1-\delta) \times \\
&\quad \big[2\theta(p,\delta)\theta(p) - \mathcal{P}(i \text{ is on an } Z[i]\text{-occupied path in } \mathcal{T}[i], i \in \{1,2\})\big] \\
&= p(1-\delta)[2\theta(p,\delta)\theta(p) - \theta(p,\delta)^2].
\end{aligned} \tag{2.6.6}$$

This, together with (2.6.5) and the expression for $\mathcal{P}(A)$ in (2.6.4) gives

$$\theta(p,\delta) \leq (p+(1-p)\delta)[2\theta(p,\delta) - \theta(p,\delta)^2] - p(1-\delta)[2\theta(p,\delta)\theta(p) - \theta(p,\delta)^2]. \tag{2.6.7}$$

Hence, if $\theta(p,\delta) > 0$,

$$1 \leq (p+(1-p)\delta)[2-\theta(p,\delta)] - p(1-\delta)[2\theta(p) - \theta(p,\delta)]. \tag{2.6.8}$$

Using (2.6.1) and elementary manipulations, the inequality above is equivalent to

$$1 \leq 2(1-p+p\delta) - \delta\theta(p,\delta). \tag{2.6.9}$$

Summarizing, $\theta(p,\delta) > 0$ implies that $1 < 2(1-p+p\delta)$, or equivalently, that $p(1-\delta) < 1/2 = p_c$. This proves part (ii) (and hence also part (i)) of the lemma.

Before we prove parts (iv) and (iii), it is useful to note the following. When p is larger than but very close to p_c, then $\theta(p)$ is close to 0 so that both (2.6.4) and (2.6.3) are similar to the equation (2.6.1), with $\theta(p+(1-p)\delta)$ as one of its solutions. One would like to conclude from this that if $p \downarrow p_c$, then $\theta(p,\delta) \to \theta(p_c+(1-p_c)\delta)$, which as we observed earlier, equals $\theta(p_c,\delta)$. The trouble is of course that when we take $\theta(p,\delta) = 0$ the above inequalities are also satisfied. However, if we can prove that $\liminf_{p\downarrow p_c} \theta(p,\delta) > 0$, then the above argument works and we get part (iv) of the theorem. The positivity of $\liminf_{p\downarrow p_c} \theta(p,\delta)$ follows immediately from part (iii) which we will prove now:

Let, for $n \geq 1$, A_n be the event that there is a Z-occupied path of length n from O, and B_n the event that there is a Z^+-occupied self-avoiding path of length n from O. Let E be the event that the X-occupied cluster of O is infinite. Similarly $A_n[k]$, $B_n[k]$ and $E[k]$ are defined for $k = 1,2$, in the obvious way. For instance, $B_n[1]$ is the event that there is a $Z^+[1]$-occupied path of length n from 1 in $\mathcal{T}[1]$. We consider, for each n, the probabilities of the combinations of occurrence or non-occurrence of these events:

$$\begin{aligned} f_1(n) &:= \mathcal{P}(\neg E, \neg A_n, \neg B_n). \\ f_2(n) &:= \mathcal{P}(E, \neg A_n, \neg B_n). \\ f_3(n) &:= \mathcal{P}(\neg E, A_n, B_n). \\ f_4(n) &:= \mathcal{P}(\neg E, A_n, \neg B_n). \\ f_5(n) &:= \mathcal{P}(E, A_n, B_n). \end{aligned}$$

These probabilities depend on p and δ of course, but we have omitted these in our notation. One might expect eight combinations, but some do not occur. For instance, B_n implies A_n. One can check that, for $p > p_c$, $\delta > 0$, and fixed n, each of $f_1(n), \ldots, f_5(n)$ is positive and that their sum equals 1.
Also note that, for fixed n, the occurrences/non-occurrences of A_{n+1}, B_{n+1}, E are completely determined by those of $A_n[k]$, $B_n[k]$, $E[k]$, for $k = 1, 2$ and X_O and Y_O. To illustrate this, suppose the events $\neg E[1]$, $\neg A_n[1]$, $\neg B_n[1]$ occur and also the events $\neg E[2]$, $A_n[2]$, $\neg B_n[2]$, and $X_O = 1$. This has probability $pf_1(n)f_4(n)$. It is not hard to check that then $\neg E$, A_{n+1} and $\neg B_{n+1}$ occur. The same holds when we exchange the roles of $\mathcal{T}[1]$ and $\mathcal{T}[2]$. This gives a contribution $2pf_1(n)f_4(n)$ to $f_4(n+1)$. Another contribution to $f_4(n+1)$ comes from the simultaneous occurrence of the events $\neg E[1], \neg A_n[1], \neg B_n[1], \neg E[2], A_n[2], B_n[2], X_O = 1$ and $Y_O = 0$, or the same events with the roles of '1' and '2' reversed. This contribution is $2p(1-\delta)f_1(n)f_3(n)$. We have

$$f_4(n+1) \geq 2p(1-\delta)f_1(n)f_3(n) + 2pf_1(n)f_4(n). \tag{2.6.10}$$

It is easy to find *all* contributions but the above is enough for our purpose.

It is easy to see from the definitions that the functions $f_1(n)$ and $f_2(n)$ are increasing in n and that $f_3(n)$ and $f_5(n)$ are decreasing in n. By this monotonicity their limits (as $n \to \infty$), which we denote by f_1, f_2, f_3, f_5, exist.
Clearly these limits depend on the values of p and δ. Since the sum $\sum_i f_i(n) = 1$, also the limit f_4 of the sequence $f_4(n)$ exists. Moreover, the limits satisfy the analogue of equation (2.6.10).

We now proceed with the proof of the theorem. Let $\delta > 0$ be given. Suppose that p is such that $\theta(p, \delta) = 0$. Hence $f_3(n) + f_4(n) + f_5(n) \to 0$, $f_1(n) \to 1 - \theta(p)$ and $f_2(n) \to \theta(p)$ as $n \to \infty$. Choose $\varepsilon > 0$ such that $(1+\delta(1-\delta))(1-\varepsilon) > 1$. The reason for this choice of ε will become clear later. Because of the convergence of $f_1(n)$ to $1 - \theta(p)$, we can find an n_0 such that $\forall n > n_0, f_1(n) > 1 - \theta(p) - \varepsilon$. Applying this to (2.6.10) we get for all $n > n_0$:

$$f_4(n+1) \geq \big[2p(1-\delta)f_3(n) + 2pf_4(n)\big](1 - \theta(p) - \varepsilon). \tag{2.6.11}$$

Now we use

$$f_3(n) \geq \delta \frac{1-p}{p} f_4(n),$$

which can be obtained easily from the definitions.

Applying this to (2.6.11) we obtain for all $n > n_0$,

$$f_4(n+1) \geq (2p + 2\delta(1-\delta)(1-p))\,(1-\theta(p)-\varepsilon)\,f_4(n).$$

If now

$$(2p + 2\delta(1-\delta)(1-p))(1-\theta(p)-\varepsilon) > 1, \tag{2.6.12}$$

we obtain $f_4(n+1) > f_4(n)$ for all $n > n_0$, which together with $f_4(n) > 0$ for all n, contradicts the assumption that, for the chosen value of p, $f_4(n) \to 0$.

At $p = p_c = 1/2$ inequality (2.6.12) is satisfied (we have chosen our ε accordingly) and hence, by continuity, there is a $\hat{p} > p_c$ such that the inequality is satisfied for all $p \in [p_c, \hat{p}]$. By the above mentioned contradiction we conclude that $\theta(p,\delta) > 0$ for p in this interval. Summarizing, we have proved the following:

$$\forall \delta > 0, \exists p' > p_c \text{ such that } \forall p \in [p_c, p'],\ \theta(p,\delta) > 0. \tag{2.6.13}$$

From this we get part (iii) of the theorem as follows: let $\delta > 0$. By (2.6.13) we can find a $\hat{p} > p_c$ that satisfies $\theta(\hat{p}, \delta/2) > 0$ and further (by taking $\hat{p}$ sufficiently close to p_c),

$$\hat{p} + (1-\hat{p})\delta/2 < p_c + (1-p_c)\delta,$$

and hence $\hat{p} + (1-\hat{p})\delta/2 < p + (1-p)\delta$ for all $p \in [p_c, \hat{p}]$. Corollary 2.2.15 now implies that for all $p \in [p_c, \hat{p}], \theta(p,\delta) \geq \theta(\hat{p}, \delta/2)$, which is strictly positive.

□

Remark 2.6.2. Gathering all the terms, it is possible to obtain a full iterative scheme in the sense that we can write $f_i(n+1)$, in terms of $f_i(n), i = 1, \ldots, 5$. Hence, for fixed p and δ, we can find an approximation of $\theta(p,\delta)$ using the iterative scheme. The limits $f_1, \ldots, f_5$ satisfy similar relations. Unfortunately, since these relations contain quadratic terms, it is not possible to determine $f_1, \ldots, f_5$ for general p and δ uniquely.

2.6.2 The regular binary tree

Now we have Theorem 2.6.1 for the rooted binary tree, it is easy to obtain similar results for the Bethe lattice $\mathcal{B}$, the tree where *each* site has exactly three neighbours. That graph is in fact nicer because all sites are 'equivalent'. Note that if we delete a site and its three edges from $\mathcal{B}$, three separate copies of $\mathcal{T}$ are left. For ordinary percolation this gives immediately that the probability that a given site is in an infinite occupied cluster equals p (the probability that the site itself is occupied) times the probability that at least one of its three neighbours is in an infinite occupied cluster in its corresponding copy of $\mathcal{T}$. So

$$\theta^{\mathcal{B}}(p) = p(1 - (1 - \theta^{\mathcal{T}}(p))^3),$$

where we write the superscript to indicate the graph under consideration. In particular $\mathcal{B}$ has the same critical probability $1/2$ as $\mathcal{T}$. For our model 'with destruction'

the situation is more complicated, but we can still immediately get an inequality as follows. Using the description with X and Y variables in Section 2.2, it is clear that if a site i has $X_i = 0$, then each infinite X-occupied cluster belongs to exactly one of the three above mentioned copies of $\mathcal{T}$. A few seconds thought then yields

$$\theta^{\mathcal{B}}(p, \delta) \geq (1-p)\delta(1 - (1 - \theta^{\mathcal{T}}(p, \delta))^3).$$

The factor δ comes from the fact that Y_i has to be 1 when $X_i = 0$ and $Z_i = 1$. This, combined with part (iii) of Theorem 2.6.1, immediately gives an analogue for $\mathcal{B}$ of that part of the theorem. The analogue for $\mathcal{B}$ of the second part of the theorem can then be obtained from the analogue for $\mathcal{B}$ of the first part, in exactly the same way in which we have derived part (iv) for $\mathcal{T}$ from part (iii) for $\mathcal{T}$. Part (i) and (ii) for $\mathcal{B}$ also follow easily from their analogues for $\mathcal{T}$.

Chapter 3

Self-organised forest-fire models

3.1 Introduction

Consider the following, informally described, forest-fire model. Each site of the lattice $\mathbb{Z}^d$ is either vacant or occupied by a tree. Vacant sites become occupied at rate 1, independently of anything else. Further, sites are hit by lightning at rate λ, the parameter of the model. When a site is hit by lightning, its entire occupied cluster instantaneously burns down, that is, becomes vacant.

This is a continuous-time version of the Drossel-Schwabl model which has received much attention in the physics literature. See e.g. [BJ05], [DS92], [Gra02], [SDS02], and sections in the book by Jensen [Jen98]. For comparison with real forest-fires see [MMT98]. The most interesting questions are related to the asymptotic behaviour when the lightning rate tends to 0 in steady state. It is believed that this behaviour resembles that of 'ordinary' statistical mechanics systems at criticality. In particular, it is believed that, asymptotically, the cluster size distribution displays power-law behaviour. Heuristic results confirming such behaviour have been given in the literature, but the validity of some of these results is debatable [Gra02] and apart from the one-dimensional case, very little is known rigorously. Section 3.2 handles the cluster size distribution in steady state on $\mathbb{Z}$. Although it was known that this distribution follows a power law [DCS93], [Dro96] no rigorous results had been proved.

Most literature handles the stationary distribution of forest-fire models on finite volumes in dimension 2, specifically when the volume tends to infinity. It would be interesting to study the stationary distributions directly on $\mathbb{Z}^2$, but usually one does not know whether a stationary distribution exists or even if it does, when it is reached. Our goal is more modest and we address some basic problems which, surprisingly, have so far been practically ignored, although their solution is crucial for a beginning of rigorous understanding of these models. In Section 3.3 we investigate what happens

if in the two-dimensional case, we start our process from an empty lattice. Around the critical time, corresponding to the critical probability for ordinary percolation, interesting behaviour may occur. It can be argued intuitively that the system stays in a critical state from the critical time on, meaning that we see large fires at all times. Nevertheless, the rigorous results suggest different behaviour. We remind the reader of the permanent self-destructive process (Section 2.5) which can be seen as the limiting case of the forest-fire model when the lightning rate tends to zero. For the permanent self-destructive process, similar intuition as mentioned above applies, but the rigorous results suggested otherwise.

3.2 The cluster size distribution on $\mathbb{Z}$

3.2.1 Introduction and background

This section discusses the forest-fire model in one dimension. In time, trees can grow, or disappear by fire. Let $\Omega = \{0,1\}^{\mathbb{Z}}$ be the state space. A '1' represents a tree and a '0' represents a vacant space. To each site of $\mathbb{Z}$ we assign a Poisson process with rate 1, independently of the other sites. If such a Poisson clock rings, there is a birth attempt. If the site is empty, a tree grows. If a tree is already present, nothing happens. Further, to each site we assign a Poisson process with rate λ, independently of the other sites. If such a Poisson clock rings, there is a lightning attempt. The lightning destroys instantaneously all the trees that are in the cluster of the site that is hit. If the site was vacant, nothing happens. In what follows, a Poisson event is the ring of a Poisson clock (either a birth attempt or a lightning attempt). It is expected that this model displays so-called self-organised critical behaviour as the lightning rate λ goes to zero.

If we start from any configuration with infinitely many vacant sites on both half-lines, with probability one, infinitely many of those sites in $\mathbb{Z}$ stay vacant during the time interval $[0,t]$. This 'breaks up' the line in finite pieces, and hence the process is easily constructed (for example by using a graphical representation, see [Dur88], [Lig85]). Note that the existence of the process on an infinite volume is not clear at all in higher dimensions, since a site may in principle be influenced by infinitely many Poisson events in finite time. Although the existence of the process on $\mathbb{Z}$ is relatively easy, it is not immediately clear whether there exists a stationary (in time) translation-invariant (in space) measure. We come back to this issue later. For now, let μ_λ denote any stationary translation-invariant measure.

For any $\omega \in \Omega$, we write $\omega = (\cdots \omega_{-1}\omega_0\omega_1 \cdots)$ where $\omega_i \in \{0,1\}$ denotes the state of the site at position i. In short notation, we write $\mu_\lambda(0) = \mu_\lambda(\omega_0 = 0)$ for the probability that the origin is empty in the steady-state, and $\mu_\lambda(1)$ for the probability that the origin is occupied. It seems likely that the probability that a site is empty goes to zero as the lightning rate goes to zero. This is indeed the case and the speed at which this happens is known: it has been proved in [BJ05] that there exist positive constants A_1 and A_2 such that for $\lambda < 1$ and any measure μ_λ invariant under the

dynamics,

$$\frac{A_1}{\ln(1/\lambda)} \leq \mu_\lambda(0) \leq \frac{A_2}{\ln(1/\lambda)}. \tag{3.2.1}$$

This forest-fire model is closely related to the well-known Drossel-Schwabl forest-fire model [DS92]. In that model the state space is large but finite, the speed of fires is finite and time is discrete. The self-organised critical behaviour is expected when the volume and the speed of fires go to infinity, and the lightning rate to zero in a suitable way. In dimension 1, the behaviour of this model has been studied in [DCS93] but the results are not rigorous and some need significant correction. Van den Berg and Jarai [BJ05] have studied our version of the model rigorously. For the overall behaviour, it does not matter much which model we consider. The proof of the main result can easily be adjusted to the Drossel-Schwabl model.

The forest-fire process has 'natural scales'. If we consider a string of length n, there is a non-trivial (i.e. bounded away from zero and one) probability that all sites grow a tree in a time interval of length $\ln(n)$. If the lightning parameter is of order $1/(n\ln(n))$, there is also a non-trivial probability that a lightning attempt occurs in this string. This and other considerations ([BJ05], [DS92]) lead to the following definition of a characteristic length: $s_{max} = s_{max}(\lambda)$ is the integer satisfying

$$s_{max} \ln s_{max} \leq \frac{1}{\lambda} < (s_{max}+1)\ln(s_{max}+1). \tag{3.2.2}$$

Our result concerns $p(s)$, where

$$p(s) := \mu_\lambda(\omega_0 = 0, \omega_1 = \cdots = \omega_s = 1, \omega_{s+1} = 0).$$

The probabilities $p(s), s \geq 0$ are called the cluster size distribution in [DCS93], although strictly speaking, they are not a distribution but the probabilities of the event that the origin is the left-boundary of a cluster of size s. From $p(s)$, it is easy to recover the true cluster size distribution. By translation invariance, the probability that a fixed site is in a cluster of size s is $sp(s)$.

In [DCS93] it was shown that for fixed s, the probabilities $p(s)$ satisfy

$$p(s) \simeq \mu_\lambda(0) s^{-2}, \tag{3.2.3}$$

where the symbol '$\simeq$' means that the quotient of the left and right side is bounded from above and below as $\lambda \downarrow 0$. Further, it is an 'ansatz' in their paper that (3.2.3) holds for s up to s_{max}. Although this ansatz led to a correct prediction in [DCS93] of the asymptotic behaviour (3.2.1) of $\mu_\lambda(0)$, it was shown in [BJ05] that (3.2.3) does not hold for s of the order s_{max}. This raises the question for which $s < s_{max}$ (3.2.3) does hold. In this section we partly answer this question by showing that, loosely formulated, (3.2.3) holds for s up to $s_{max}^{1/3}$ and hence (by (3.2.1))

$$p(s) \simeq s^{-2}/\ln(1/\lambda) \text{ for all } s \text{ up to } s_{max}^{1/3}.$$

A precise formulation is stated in the following section. Section 3.2.3 handles some preliminaries and in Section 3.2.4 we give the proof of the main result. In Section 3.2.5, we address the issue of the existence of a stationary translation-invariant measure for the one-dimensional forest-fire process.

3.2.2 Statement of the main result

The main result is as follows.

Theorem 3.2.1. *Let $\alpha < 1/3$. There exists positive constants B_1 and B_2 such that for all $\lambda < 1$ and any stationary, translation-invariant measure μ_λ of the forest-fire model with parameter λ on $\mathbb{Z}$,*

$$\frac{B_1}{s^2 \ln(1/\lambda)} \leq \mu_\lambda(\omega_0 = 0, \omega_1 = \cdots = \omega_s = 1, \omega_{s+1} = 0) \leq \frac{B_2}{s^2 \ln(1/\lambda)}, \tag{3.2.4}$$

for all $s \leq s^\alpha_{max}$.

The theorem above is only useful when there is at least one stationary translation-invariant measure for the one-dimensional forest-fire model. We will show that this is the case in Section 3.2.5.

3.2.3 Preliminaries

Note that on a finite interval, the total rate of the Poisson processes assigned to the sites in this interval is finite. This implies that the probability of two Poisson events occurring inside this finite space interval in a time interval of length t is $o(t)$, as $t \downarrow 0$. We repeatedly use this fact when proving statements about the stationary measure. However, when we consider $\mathbb{Z}$, infinitely many Poisson events occur in any time interval; far away a tree could be hit by lightning and the fire thus started could travel over a very large distance, creating long-range dependencies. We have no a priori bound on these dependencies. Unfortunately, this complicates the argument. In the rest of this section, we show that in some sense 'there are enough vacant sites at all times', which gives us a bound on the size of clusters, and hence a bound on the size of fires. We make this precise in Proposition 3.2.4. We need an auxiliary model where destructions are local.

The model with local destructions is coupled to the original one by using the same Poisson clocks. The state space is $\{0,1\}^{\mathbb{Z}}$ as before. Each time a growth clock rings, a tree tries to grow. As for the lightning attempts: a lightning attempt destroys a tree (if present) instantaneously. The difference to the original model is that the rest of its cluster remains intact. Each site thus behaves independently of the other sites. Hence the distribution of the configuration at time t converges to ν_λ, which is the product measure with density $1/(\lambda+1)$, as $t \to \infty$. This happens for any initial configuration. Now suppose that we take a configuration x in $\{0,1\}^{\mathbb{Z}}$ and start both the processes from that configuration. Using the same Poisson events for both models gives us a natural coupling. Let $\eta^x(t)$ denote the configuration at time t when we start in configuration x for the model with local destructions. For the original model we define $\omega^x(t)$ likewise. Then from the definition of the processes it should be clear that for all initial configurations $x \in \{0,1\}^{\mathbb{Z}}$, and all times $t \geq 0$,

$$\omega^x(t) \leq \eta^x(t). \tag{3.2.5}$$

Consider the interval $I = [i_1, i_m] \subset \mathbb{Z}$ of length m. We define for any configuration $x \in \{0,1\}^{\mathbb{Z}}$,

$$\begin{aligned} i_l = i_l[x] &:= \max\{j \leq i_1 - 2 : x_j + x_{j+1} = 0\}, \\ i_r = i_r[x] &:= \min\{j \geq i_m + 2 : x_{j-1} + x_j = 0\}, \\ S_I(x) &:= [i_l, i_r]. \end{aligned} \tag{3.2.6}$$

In words, $S_I(x)$ is the smallest set of consecutive sites containing the interval $[i_1, i_m]$, two consecutive zeros to the left of i_1 and two consecutive zeros to the right of i_m. In many of the applications below, I will be the set determining some cylinder event. We prove the following lemma.

Lemma 3.2.2. *For any $\lambda < 1$, any stationary measure μ_λ and any interval $I = [i_1, i_m]$, there exists a $D = D(\lambda) \in (0,1)$ such that for all $s \geq m + 4$,*

$$\mu_\lambda(|S_I| \geq s) \leq 2D^{\lfloor (s-m) \rfloor}.$$

Proof. Let I be as in the statement of the lemma. Let μ_λ be a stationary measure and let $x \in \{0,1\}^{\mathbb{Z}}$. Let $\mathcal{P}_\lambda$ denote the measure governing the Poisson clocks. By (3.2.5), for any $t \geq 0$,

$$\mathcal{P}_\lambda(|S_I(\omega^x(t))| \geq s) \leq \mathcal{P}_\lambda(|S_I(\eta^x(t))| \geq s) \tag{3.2.7}$$

This inequality remains valid if we integrate over x with respect to μ_λ. The left side of (3.2.7) is then simply $\mu_\lambda(|S_I| \geq s)$. Recall that the η-process converges to its stationary measure ν_λ for every initial configuration, when we take the limit $t \to \infty$. Taking this limit, we obtain

$$\mu_\lambda(|S_I| \geq s) \quad \leq \quad \nu_\lambda(|S_I| \geq s). \tag{3.2.8}$$

Suppose the size of S_I is at least s. Then there are at least $\lfloor (s-m)/2 \rfloor$ sites in S_I directly to the right of i_m (or directly to the left of i_1). These sites can be divided into disjoint pairs, where each pair, apart from the right-most (or left-most, respectively), has at least one occupied site, by definition of S_I. This gives us

$$\nu_\lambda(|S_I| \geq s) \leq 2\Big(1 - \big(\frac{\lambda}{1+\lambda}\big)^2\Big)^{\lfloor (s-m)/4 \rfloor - 1}. \tag{3.2.9}$$

Combining (3.2.8) and (3.2.9) proves the lemma. □

Note that in the proof above, we did not need translation invariance for μ_λ.

Remark 3.2.3. An immediate consequence of Lemma 3.2.2 is that occupied clusters are a.s. finite.

Lemma 3.2.2 gives us the tools to bound influences from far away. When the initial configuration is drawn from μ_λ, we define $\mathcal{P}^{\mu_\lambda}$ to be the measure governing this initial configuration and the Poisson clocks, hence determining the forest-fire process. Recall that $\omega(t)$ denotes the state of the model at time t. To make the statement at the beginning of this subsection precise we prove the following:

Proposition 3.2.4. *Let I be an interval $[i_1, i_m] \subset \mathbb{Z}$ and let $t > 0$. Recall the definition of $S_I(\omega)$ and let $M(I, \omega(0), t)$ denote the number of Poisson events occurring in the set $S_I(\omega(0))$ in the time-interval $[0, t]$. For any stationary measure μ_λ and interval I,*

$$\mathcal{P}^{\mu_\lambda}(M(I, \omega(0), t) > 1) = o(t),$$

as $t \downarrow 0$.

Proof. For fixed sets S_I it is immediately clear that the probability of two or more Poisson events in S_I is $o(t)$ as $t \downarrow 0$. But now S_I is random; at this point we use that the size distribution of S_I decays exponentially. Further, we use that for a Poisson process $X(t)$ with $EX(t) = \alpha t$ we have $\mathcal{P}(X(t) > 1) \leq (\alpha t)^2$.

$$\begin{aligned}
&\mathcal{P}^{\mu_\lambda}(M(I, \omega(0), t) > 1) \\
&\quad \leq \sum_{j \geq m+4} \mathcal{P}^{\mu_\lambda}(M(I, \omega(0), t) > 1 \mid |S_I(\omega(0))| = j) \mu_\lambda(|S_I(\omega(0))| = j) \\
&\quad \leq \sum_{j \geq m+4} j^2 (\lambda + 1)^2 t^2 \mu_\lambda(|S_I(\omega(0))| = j).
\end{aligned}$$

The sum over j is finite by Lemma 3.2.2, which proves the proposition. □

3.2.4 Proof of the main result

The proof of our main result is based on ideas from [Pen00]. We introduce some notation and state some lemmas first. From now, let $\lambda > 0$ be fixed and suppose that μ_λ is a stationary translation-invariant measure. In what follows, when we write $o(t)$, we implicitly take the limit $t \downarrow 0$.

In the proof of the main theorem, we need to bound the probability of the event that a tree on the edge of a cluster is burnt. The next lemma shows that although large clusters may arise, the probability that a boundary tree is on fire is not that large.

Define for $i \in \mathbb{Z}$ and $t > 0$,

$$\begin{aligned}
B_i^+(t) &:= \{\omega_i(0) = 0, \omega_{i+1}(0) = 1, \omega_{i+1}(t) = 0\}, \\
B_i^-(t) &:= \{\omega_i(0) = 0, \omega_{i-1}(0) = 1, \omega_{i-1}(t) = 0\}.
\end{aligned}$$

Note that by translation invariance, the probability of $B_i^+(t)$ does not depend on i. The same holds for B_i^-.

Lemma 3.2.5. *Let $B_i^+(t), B_i^-(t)$ be defined as above. Then for all i and $t > 0$,*

$$\begin{aligned} \mathcal{P}^{\mu_\lambda}(B_i^+(t)) &\leq \lambda t + o(t), \\ \mathcal{P}^{\mu_\lambda}(B_i^-(t)) &\leq \lambda t + o(t). \end{aligned}$$

Proof. We prove the lemma only for $B_i^+(t)$. The proof for $B_i^-(t)$ is completely similar. Note first that we can bound the possibility that more than one Poisson event influences the event $B_i^+(t)$. This will give us for example, that the event that the site $i+1$ is occupied, becomes vacant, again becomes occupied and again vacant (which is in principle possible by the definition of $B_i^+(t)$) has probability $o(t)$. We take $I = \{i+1\}$ and apply Proposition 3.2.4:

$$\mathcal{P}^{\mu_\lambda}(B_i^+(t)) \leq \mathcal{P}^{\mu_\lambda}(B_i^+(t) \cap M(\{i+1\}, \omega(0), t) \leq 1) + o(t) \tag{3.2.10}$$

On $\{M(\{i+1\}, \omega(0), t) \leq 1\}$, only Poisson events occurring inside $S_{\{i+1\}}$ can cause $B_i^+(t)$ to occur: the two zeros on the left and right boundary of $S_{\{i+1\}}$ (and the fact that at most one of these turns into a one) prevents fires from the outside to reach $i+1$. For the same reason,

$$\mathcal{P}^{\mu_\lambda}(B_i^+(t) \cap M(\{i+1\}, \omega(0), t) = 0) = 0. \tag{3.2.11}$$

Now, using that the cluster of site i is finite a.s. we consider all possibilities that cause $B_i^+(t)$ to occur. Let L_j denote the event that site j is hit by lightning in the time interval $[0, t]$. We get

$$\begin{aligned} &\mathcal{P}^{\mu_\lambda}(B_i^+(t) \cap M(\{i+1\}, \omega(0), t) = 1) \\ &\leq \sum_{j=0}^{\infty} \mathcal{P}^{\mu_\lambda}(\omega_i(0) = 0, \omega_{i+1}(0) = \cdots = \omega_{i+j}(0) = 1 \text{ and } L_j) \\ &= \sum_{j=0}^{\infty} \mu_\lambda(\omega_{-j} = 0, \omega_{-j+1} = \cdots = \omega_0 = 1) \mathcal{P}^{\mu_\lambda}(L_j) \\ &\leq \mu_\lambda(1)\lambda t + o(t) \leq \lambda t + o(t). \end{aligned} \tag{3.2.12}$$

Once more, we have used translation invariance in the equality above. Combining (3.2.10), (3.2.11) and (3.2.12) proves the lemma. □

We now concentrate on what happens in a finite string. Define

$$\Omega_n[k] := \{\omega \in \Omega : \sum_{i=0}^{n-1} \omega_i = k\},$$

We consider the event that in a string of n consecutive sites there are exactly k occupied sites, and the ends of the string are empty. We define for $k \leq n-1$,

$$A_n^k = \{\omega \in \Omega_n[k] : \omega_0 = \omega_{n-1} = 0\}. \tag{3.2.13}$$

Note that we are particularly interested in A^n_{n+2}. The proof of the main theorem is based heavily on the following relation for the A^k_n's.

Lemma 3.2.6.

$$\left|\mu_\lambda(A^0_n) - \frac{\mu_\lambda(0)}{n}\right| \leq \lambda n \tag{3.2.14}$$

$$\left|\mu_\lambda(A^k_n) - \frac{n-k-1}{n-k}\mu_\lambda(A^{k-1}_n)\right| \leq \frac{4\lambda n}{n-k}. \tag{3.2.15}$$

Note that the event A^0_n contains only configurations that are equal to zero to on $[0, n-1]$.

Proof. For any measurable $A \subset \{0,1\}^{\mathbb{Z}}$ and any $t > 0$, we have

$$\mathcal{P}^{\mu_\lambda}(\omega(0) \in A) = \mathcal{P}^{\mu_\lambda}(\omega(t) \in A),$$

and hence

$$\mathcal{P}^{\mu_\lambda}(\omega(0) \notin A, \omega(t) \in A) = \mathcal{P}^{\mu_\lambda}(\omega(0) \in A, \omega(t) \notin A). \tag{3.2.16}$$

We refer to equation (3.2.16) as the 'steady-state equation'. We refer to the l.h.s. of (3.2.16) as 'going in' side and to the r.h.s. as 'going out' side, for obvious reasons.

We first show (3.2.14). To this end, we apply (3.2.16) to $A = \{\omega \in \Omega : \omega_0 = \omega_1 = \cdots = \omega_{n-2} = 0, \omega_{n-1} = 1\}$. Now on the 'going in' side of the steady-state equation for A we get the following contributions.

- At time 0 we see the configuration $\{\omega_0(0) = \cdots = \omega_{n-1}(0) = 0\}$ and a tree grows at site $n-1$ in the time interval $[0,t]$.
- We see a configuration at time 0 where there is a cluster of trees with rightmost site between 0 and $n-3$ and this cluster is burnt during the time interval $[0,t]$. This has probability at most $\lambda(n-2)t + o(t)$ by Lemma 3.2.5.

All other possibilities have probability $o(t)$ by Proposition 3.2.4. From the above we conclude that the l.h.s. of (3.2.16) is at least

$$\mu_\lambda(\omega_0 = \cdots = \omega_{n-1} = 0)t + o(t)$$

and at most

$$\mu_\lambda(\omega_0 = \cdots = \omega_{n-1} = 0)t + (n-2)\lambda t + o(t). \tag{3.2.17}$$

On the 'going out' side of the steady-state equation we get contributions from growing trees and fires as well.

- A tree grows on one of the vacant sites, which gives us a factor $(n-1)\mu_\lambda(\omega_0 = \cdots = \omega_{n-2} = 0, \omega_{n-1} = 1)t + o(t)$.
- The tree on site $n-1$ is burnt. This gives us a contribution of exactly $\mathcal{P}^{\mu_\lambda}(B^+_{n-2}(t))$, which is at most $\lambda t + o(t)$ by Lemma 3.2.5.

All other possibilities have probability $o(t)$ by Proposition 3.2.4. We conclude that the r.h.s of (3.2.16) is at least

$$(n-1)\mu_\lambda(\omega_0 = \cdots = \omega_{n-2} = 0, \omega_{n-1} = 1)t + o(t)$$

and at most

$$(n-1)\mu_\lambda(\omega_0 = \cdots = \omega_{n-2} = 0, \omega_{n-1} = 1)t + \lambda t + o(t). \tag{3.2.18}$$

Combining (3.2.16), (3.2.17) and (3.2.18) we obtain

$$\left|\mu_\lambda(\omega_0 = \cdots = \omega_{n-1} = 0)t - (n-1)\mu_\lambda(\omega_0 = \cdots = \omega_{n-2} = 0, \omega_{n-1} = 1)t\right| < \lambda n t + o(t).$$

We divide by t and subsequently let $t \downarrow 0$:

$$\left|\mu_\lambda(\omega_0 = \cdots = \omega_{n-1} = 0) - (n-1)\mu_\lambda(\omega_0 = \cdots = \omega_{n-2} = 0, \omega_{n-1} = 1)\right| < \lambda n. \tag{3.2.19}$$

Note that

$$\begin{aligned} \mu_\lambda(\omega_0 = \cdots = \omega_{n-2} = 0, \omega_{n-1} = 1) \;+\; & \mu_\lambda(\omega_0 = \cdots = \omega_{n-1} = 0) \\ = \; & \mu_\lambda(\omega_0 = \cdots = \omega_{n-2} = 0). \end{aligned} \tag{3.2.20}$$

Now, combining (3.2.20) with (3.2.19) we obtain

$$\left|\mu_\lambda(A_n^0) - \frac{n-1}{n}\mu_\lambda(A_{n-1}^0)\right| < \lambda. \tag{3.2.21}$$

Iterating (3.2.21) we get (3.2.14).

Now take $n-1 \geq k > 0$. As before, we consider the steady-state equation (3.2.16), but now for the event A_n^k to prove (3.2.15). Again we ignore multiple Poisson events using Proposition 3.2.4.

On the 'going in' side of the equation we obtain the following contributions.

- A tree grows in a configuration with $k-1$ trees. There are $n-2-(k-1)$ possible locations; recall that there are no trees allowed at site 0 or $n-1$ in a configuration in A_n^k.
- The possibility to get into a configuration in A_n^k by a fire is contained in $\cup_{i=0}^{n-2}(B_i^+(t) \cup B_{i+1}^-(t))$ so applying Lemma 3.2.5 gives a contribution of at most $2(n-1)\lambda t + o(t)$.

The total contribution on the 'going in' side is at least

$$(n-k-1)\mu_\lambda(A_n^{k-1})t + o(t)$$

and at most

$$(n-k-1)\mu_\lambda(A_n^{k-1})t + 2(n-1)\lambda t + o(t). \tag{3.2.22}$$

On the 'going out' side of the equation we get the following contributions.

- In a configuration in A_n^k a tree grows. There are $n-k$ possible locations.
- The possibility to leave a configuration in A_n^k by a fire is contained in $\cup_{i=0}^{n-3}(B_i^+(t) \cup B_{i+2}^-(t))$. As before, Lemma 3.2.5 bounds this probability by $2(n-2)\lambda t + o(t)$.

The total contribution on the 'going out' side is at least

$$(n-k)\mu_\lambda(A_n^k)t + o(t)$$

and at most

$$(n-k)\mu_\lambda(A_n^k)t + 2(n-2)\lambda t + o(t). \tag{3.2.23}$$

Combining (3.2.16), (3.2.22) and (3.2.23) and subsequently dividing by t (and $n-k$) we obtain as $t \downarrow 0$,

$$\left|\mu_\lambda(A_n^k) - \frac{n-k-1}{n-k}\mu_\lambda(A_n^{k-1})\right| \leq \frac{4(n-1)\lambda}{n-k}.$$

This proves (3.2.15). □

Now we are ready for the proof of Theorem 3.2.1.

Proof. [Thm 3.2.1] Note that $\{\omega \in \Omega : \omega_0 = 0, \omega_1 = \cdots = \omega_n = 1, \omega_{n+1} = 0\} = A_{n+2}^n$. We apply (3.2.14) and (3.2.15) repeatedly to obtain

$$\left|\mu_\lambda(A_{n+2}^n) - \frac{1}{2}\frac{2}{3}\cdot\cdots\cdot\frac{n}{n+1}\frac{\mu_\lambda(0)}{n+2}\right| \leq \frac{n\lambda}{n+1} + \sum_{j=1}^{n+1}\frac{4n\lambda}{j},$$

which tells us

$$\left|\mu_\lambda(A_{n+2}^n) - \frac{\mu_\lambda(0)}{(n+1)(n+2)}\right| \leq (4n[\ln(n+1)+1]+1)\lambda.$$

The question is now for which cluster sizes the error can be bounded uniformly in λ. Using (3.2.1), the relative error is

$$\frac{(n+1)(n+2)(4n[\ln(n+1)+1]+1)\lambda}{\mu_\lambda(0)} \simeq n^3\ln(n)\lambda\ln(1/\lambda).$$

Now choose $\alpha < 1/3$ and suppose that $n \leq s_{max}^\alpha$. The right hand side is dominated by $\alpha s_{max}^{3\alpha-1}\ln(1/\lambda)$. By definition of s_{max} (3.2.2) this dominating factor goes to zero as $\lambda \downarrow 0$ and this proves Theorem 3.2.1. □

3.2.5 Existence of a stationary translation-invariant measure

Proposition 3.2.7. *For the forest-fire model with parameter λ on $\mathbb{Z}$, there exists at least one stationary, translation-invariant measure.*

Since the forest-fire process is not a Feller process, the proposition does not follow from the standard theory, see for example [Lig85].

Proof. Fix $\lambda > 0$. Let $\omega(\cdot)$ denote the forest-fire process on $\mathbb{Z}$. Let $k \in \mathbb{N}$ and let $\omega^{(k)}(\cdot)$ be an auxiliary forest-fire process where on $[-k, k]$ we have the dynamics of the ordinary forest-fire process, but with the understanding that we consider sites $-k$ and k to be neighbours i.e., we consider the forest-fire process on a one-dimensional torus embedded in $\mathbb{Z}$. For instance if $-k$ and k are both occupied we consider them to be in the same occupied cluster for this process. Outside the interval $[-k, k]$ nothing happens, i.e. all sites are kept vacant.

Essentially, $\omega^{(k)}$ is just the forest-fire process on a circle with $2k+1$ sites. This is a Markov chain with a unique stationary distribution which we denote by $\mu^{(k)}$. Note that by the above description,

$$\mu^{(k)}(\omega_i^{(k)} = 0) = 1 \text{ for all } i \text{ with } |i| > k.$$

By standard arguments the sequence $\mu^{(k)}, k = 1, 2, \ldots$ has a weakly convergent subsequence $\mu^{(k_i)}, i = 1, 2 \ldots$. We denote its limit by μ. By the above mentioned correspondence to circles, where we have rotation invariance for $\mu^{(k)}, k > 0$, it follows immediately that μ is translation invariant. We will show that μ is a stationary distribution for the process $\omega(\cdot)$.

First some notation. Let ν be a distribution on $\{0,1\}^{\mathbb{Z}}$, assigning infinitely many zeros to both half-lines (so that the process starting in a configuration drawn from this distribution exists). Let $\mathcal{P}^{\nu}$ denote the law governing all Poisson processes and the initial configuration. In particular, $\mathcal{P}^{\nu}(\omega(t) \in \cdot)$ is the distribution of the configuration at time t for the forest-fire process starting in a configuration drawn from ν. Similarly, when $\nu(k)$ is a distribution on $\{0,1\}^{[-k,k]}$ we write $\mathcal{P}^{\nu(k)}(\omega^{(k)}(t) \in \cdot)$ for the auxiliary process $\omega^{(k)}$, starting in a configuration drawn from $\nu(k)$. For a distribution ν on $\{0,1\}^{\mathbb{Z}}$ and a subset J of $\mathbb{Z}$ we denote by ν_J the restriction of ν to J, i.e. its projection on J.

To show that μ is a stationary distribution for the process $\omega(\cdot)$, it is sufficient to show that for all cylinder events A and all $t < 1$,

$$\mathcal{P}^{\mu}(\omega(t) \in A) = \mu(A). \tag{3.2.24}$$

So let A be a cylinder event and $t < 1$. Let $I = [i_1, i_m]$ be an interval such that A is determined by the configuration on I. Take positive integers L and N. We define $J = J(L) = [i_1 - L, i_m + L]$ and let $k > \max\{|i_1 - L|, |i_m + L|\}$. If the ordinary forest-fire process $\omega(\cdot)$ and the auxiliary process $\omega^{(k)}(\cdot)$ start with initial configurations which agree on J and we use the same Poisson clocks, then the only way to have disagreement between $\omega(t)$ and $\omega^{(k)}(t)$ on I is by influences (fires) from outside J.

These fires can only reach I if all sites in the interval $[i_1 - L, i_1)$ that were vacant at time 0, become occupied before time t, or if all sites in the interval $(i_m, i_m + L]$ that were vacant at time 0 become occupied before time t. If in both intervals the number of zeros is at least N, the above event clearly has probability at most $2t^N$.

Now we couple the process $\omega(\cdot)$ with initial distribution μ and the auxiliary process $\omega^{(k)}(\cdot)$ with initial configuration $\mu^{(k)}$ by using the same Poisson clocks and by optimally coupling μ_J and $\mu_J^{(k)}$. Using such coupling and the argument from the previous paragraph, gives

$$\begin{aligned} \left|\mathcal{P}^{\mu}(\omega(t) \in A) - \mu^{(k)}(A)\right| &= \left|\mathcal{P}^{\mu}(\omega(t) \in A) - \mathcal{P}^{\mu^{(k)}}(\omega^{(k)}(t) \in A)\right| \\ &\leq d_V(\mu_J, \mu_J^{(k)}) + 2t^N + \mu(B_I(L,N)), \quad (3.2.25) \end{aligned}$$

where d_V denotes variational distance[1] and

$$B_I(L,N) = \{ \text{ less than } N \text{ vacant sites in } [i_1 - L, i_1) \text{ or in } (i_m, i_m + L]\}.$$

Now let $k \to \infty$ along the subsequence mentioned in the beginning of this section. Since $\mu^{(k)}$ converges weakly to μ along that subsequence, (3.2.25) then becomes

$$\left|\mathcal{P}^{\mu}(\omega(t) \in A) - \mu(A)\right| \leq 2t^N + \mu(B_I(L,N)). \quad (3.2.26)$$

Now we let $L \to \infty$. By exactly the same reasons as in the beginning of Section 3.2.3, μ is dominated by a product measure with a positive density of zeros. Hence, the last term in (3.2.26) goes to zero as $L \to \infty$. Finally, we let $N \to \infty$ to finish the proof. □

It has been pointed out by F. Redig, that the proposition may follow from the more general theory of 'almost Feller' processes, a notion developed by him and C. Maes.

3.3 Forest-fires near the critical time on $\mathbb{Z}^2$

3.3.1 Motivation

In this section we restrict ourselves to dimension 2, that is, we consider the forest-fire process on $\mathbb{Z}^2$. It seems to be taken for granted in the literature that, informally speaking, as we let λ tend to 0, the steady-state probability[2] that a given site is vacant, stays away from 0. But is this really obvious? (Is it even true?). In the one-dimensional case, we have seen that this probability goes to zero. However, it is believed in the literature that the system behaves differently in dimension 2. The intuitive reasoning seems to be roughly as follows:

[1] See also Section 5.2.1 for results on couplings and variational distance.

[2] The steady-state measure is not necessarily unique, so that strictly speaking, it would be appropriate to speak from *a* steady-state probability.

"If the limit of the probability to be occupied was 1, then the system would have an 'infinite occupied cluster'. But that cluster would be immediately destroyed, bringing the occupation density away from 1: contradiction".

Of course this reasoning is, mildly speaking, quite shaky. A rigorous solution of this problem is necessary for a clear understanding of the forest-fire model.

The problems investigated in this section are, although not the same as the one just described, of the same spirit. Instead of looking at the steady-state distribution, we start with all sites vacant and look at the time t_c at which, in the modified model where there is only growth but no ignition, an infinite cluster starts to form. Intuitive reasoning similar to that above makes plausible that, informally speaking, for every $t > t_c$, the probability that the origin O burns before time t stays away from 0 as λ tends to 0. Continuing such intuitive reasoning then leads to the 'conclusion' that, again informally speaking, if we take m sufficiently large and replace the event $\{O$ burns before time $t\}$ by the event

$$\{\text{some vertex at distance} \leq m \text{ from } O \text{ burns before time } t\},$$

the corresponding probability will be, as λ tends to 0, as close to 1 as we want. We relate this to problems which are closer to ordinary percolation. In particular we show that, under a percolation-like assumption (which we believe to be true), the above 'conclusion' is false. We hope our results will lead to further research and clarification of the problems above.

3.3.2 Formal statement of the problems

So far, we have not yet defined our model precisely. We now give this more precise description, formulate some of the above mentioned problems more formally, and introduce much of the terminology used in the rest of this section.

We work on the square lattice $\mathbb{Z}^2$. To each site we assign two Poisson clocks: one (which we call the 'growth clock') having rate 1, and the other (the 'ignition clock') having rate λ. All Poisson clocks behave independently of each other. A site can be occupied by a tree or vacant. These states are denoted by '1' and '0', respectively. Initially all sites are vacant. We restrict ourselves to a finite box $B(n) := [-n, n]^2$. In our theorems we consider the behaviour as $n \to \infty$. The dynamics is as follows: when the growth clock of a site rings, that site becomes occupied (unless it already was occupied, in which case the clock is ignored); when the ignition clock of a site v rings, each site that has an occupied path in $B(n)$ to v, becomes vacant instantaneously. Note that this means that if v was already vacant, nothing happens. Now let $\eta_v^n(t) = \eta_v(t) \in \{0, 1\}$ denote the state of site v at time t, and define $\eta(t) = \eta^n(t) := (\eta_v^n(t), v \in B(n))$. Note that, for each n, $(\eta^n(t), t \geq 0)$ is a finite-state (continuous-time) irreducible Markov chain with state space $\{0, 1\}^{B(n)}$. The assignment of Poisson clocks to *every* site of the square lattice provides a natural coupling of the processes $\eta^n(\cdot), n \geq 1$ with each other, and with other processes (see below).

For $m \leq n$, we often use the informal phrase "η^n has a fire in $B(m)$ before time t" for the event

$$\{\exists v \in B(m) \text{ and } \exists s \leq t \text{ such that } \eta_v^n(s^-) = 1 \text{ and } \eta_v^n(s) = 0\}.$$

Similarly, we use "η^n has at least two fires in $B(m)$ before time t" for the event

$$\{\exists v, w \in B(m) \text{ and } \exists s < u \leq t \text{ s.t. } \eta_v^n(s^-) = \eta_w^n(u^-) = 1 \text{ and } \eta_v^n(s) = \eta_w^n(u) = 0\}.$$

Note that we allow v and w to be equal.

Let $\mathcal{P}_\lambda$ be the measure that governs all the underlying Poisson processes mentioned above and hence, for all n simultaneously, the processes $\eta^n(\cdot)$. Often, when there is no need to explicitly indicate the dependence on λ, or when we consider events involving the growth clocks only, we will omit this subscript.

It is trivial that for all times t and all n, m the probability that η^n has a fire in $B(m)$ before time t goes to 0 as $\lambda \downarrow 0$, and hence for all $m \geq 0$ and all $t \geq 0$,

$$\lim_{n\to\infty} \lim_{\lambda\downarrow 0} \mathcal{P}_\lambda(\eta^n \text{ has a fire in } B(m) \text{ before time } t) = 0.$$

A much more natural (and difficult!) question is what happens when we reverse the order of the limits. For the investigation of such questions it turns out to be very useful to consider the modified process $\sigma(t)$ on the infinite lattice, which we obtain, loosely speaking, if we obey the above mentioned growth clocks but ignore the ignition clocks:

$$\sigma_v(t) = I_{\{\text{the growth clock at } v \text{ rings in } [0,t]\}},$$

where I denotes the indicator function. It is clear that, for each time t, the $\sigma_v(t)$, $v \in \mathbb{Z}^2$, are Bernoulli random variables with parameter $1 - \exp(-t)$. So, if we define t_c by the relation $p_c = 1 - \exp(-t_c)$, where p_c is the critical value for ordinary site percolation on the square lattice, we see that $\sigma(t)$ has no infinite occupied cluster for $t \leq t_c$ but does have an infinite cluster for $t > t_c$. To illustrate the usefulness of comparison of η with σ (and as an introduction to more subtle comparison arguments), we show that

$$\limsup_{\lambda\downarrow 0} \limsup_{n\to\infty} \mathcal{P}_\lambda(\eta^n \text{ has a fire in } O \text{ before time } t) \leq \theta(1 - e^{-t}), \tag{3.3.1}$$

where $\theta(\cdot)$ is the percolation function for ordinary site percolation. The argument is as follows: let $C_t(O)$ denote the occupied cluster of O in the configuration $\sigma(t)$. It is easy to see from the process descriptions above that in order to have a fire in O before time t in η^n, it is necessary (but not sufficient) that at least one of the ignition clocks in the set $C_t(O)$ has rung before time t. Note that this gives us a bound on the probability that η^n has a fire in O before time t which is independent of n. Using the independence of the different Poisson clocks, we have

$$\begin{aligned}
&\mathcal{P}_\lambda(\eta^n \text{ has a fire in } O \text{ before time } t)\\
&\quad\le \sum_{k=1}^{\infty} \mathcal{P}_\lambda(|C_t(O)| = k \text{ and } \exists v \in C_t(O) \text{ that has ignition before time } t)\\
&\quad\quad + \quad \mathcal{P}(|C_t(O)| = \infty)\\
&\quad= \sum_{k=1}^{\infty} \mathcal{P}(|C_t(O)| = k)(1 - e^{-\lambda t k}) + \theta(1 - e^{-t}).
\end{aligned}$$

Note that, in the r.h.s. above, the first term clearly goes to 0 as $\lambda \to 0$ (by bounded convergence). The desired result follows.

In particular, we have for each m and each $t \le t_c$,

$$\lim_{\lambda \downarrow 0} \lim_{n \to \infty} \mathcal{P}_\lambda(\eta^n \text{ has a fire in } B(m) \text{ before time } t) \le |B(m)|\theta(1 - e^{-t}) = 0, \quad (3.3.2)$$

where $|B(m)|$ denotes the number of sites in $B(m)$. Up to and including time t_c the forest-fire model shows no surprising behaviour, but what happens right after t_c? Intuitively one might argue as follows:

"If the l.h.s. of (3.3.2) is 0 for some $t > t_c$, then roughly speaking, the system at time t looks as in ordinary percolation with parameter $1 - \exp(-t)$, so that an infinite occupied cluster has built up, and this cluster intersects $B(m)$ with positive probability. But an infinite cluster has an infinite total ignition rate and hence catches fire immediately: contradiction. Hence for each $t > t_c$ the l.h.s. of (3.3.2) is strictly positive."

As we said before, such reasoning is very shaky. Its conclusion is correct for the directed binary tree, see Lemma 3.4.1. We have some inclination to believe that the conclusion also holds for the square lattice, but prefer to formulate this as an open problem, rather than a conjecture:

Open Problem 3.3.1. Is, for all $t > t_c$,

$$\liminf_{\lambda \downarrow 0} \liminf_{n \to \infty} \mathcal{P}_\lambda(\eta^n \text{ has a fire in } O \text{ before time } t) > 0\,? \quad (3.3.3)$$

Believing the answer to the above problem is affirmative, it is intuitively very tempting to go further and 'conclude' that also the answer to the following problem is affirmative:

Open Problem 3.3.2. Is it true that for all $t > t_c$ and each $\varepsilon > 0$ there exists an m such that

$$\liminf_{\lambda \downarrow 0} \liminf_{n \to \infty} \mathcal{P}_\lambda(\eta^n \text{ has a fire in } B(m) \text{ before time } t) > 1 - \varepsilon\,? \quad (3.3.4)$$

The intuitive and again shaky reasoning here is, roughly speaking, that if the answer to Problem 3.3.1 is affirmative, there will be a positive density of sites that burn before time t, and hence the probability of having such a site in B(m) will tend to 1 as $m \to \infty$. Our main result, Theorem 3.3.4, indicates that the behaviour of the process may be considerably different from what the above intuition suggests.

At this point, one could wonder whether it is really necessary to first restrict to finite n, so that we have the annoying 'extra' limit $n \to \infty$ in our theorems and problem formulations.

Open Problem 3.3.3. Is the forest-fire model well-defined on the infinite lattice for each $\lambda > 0$?

For sufficiently large λ one can easily see that this is true. Using domination by suitable Bernoulli processes one can, for such λ, make a standard graphical construction. However, for our investigation (where we let λ approach 0) this is of no use. In Section 3.5.2 we consider a slightly modified process that *is* defined on the infinite lattice. In this modified process occupied clusters with size larger than or equal to L (the parameter of the model) are instantaneously removed. For that model we have results very similar to those for the original one.

3.3.3 Statement of the main results

We remind the reader of the percolation-like critical value, denoted by $\hat{\delta}_c$, which plays a major role in the statement of our main results. See Subsection 2.4.2 for its definition (2.4.9) and more background.

We are now ready to state our main results:

Theorem 3.3.4. *If $\hat{\delta}_c > 0$, then there exists a $t > t_c$ such that for all m,*

$$\limsup_{\lambda \downarrow 0} \liminf_{n \to \infty} \mathcal{P}_\lambda(\eta^n \text{ has a fire in } B(m) \text{ before time } t) \leq 1/2. \tag{3.3.5}$$

The key to Theorem 3.3.4 is the following proposition:

Proposition 3.3.5. *If $\hat{\delta}_c > 0$, then there exists a $t > t_c$ such that for all m,*

$$\lim_{\lambda \downarrow 0} \limsup_{n \to \infty} \mathcal{P}_\lambda(\eta^n \text{ has at least 2 fires in } B(m) \text{ before time } t) = 0. \tag{3.3.6}$$

The proofs of the above proposition and theorem are given in Section 3.3.4.

Remark 3.3.6. In the model at hand the square lattice plays the role of space. Completely analogous results can be proved in the same way, for all planar graphs for which the results in Subsection 2.4.3 (in particular Lemma 2.4.9) hold.

3.3.4 Proofs

The proof of our main theorem (Theorem 3.3.4) depends heavily on Proposition 3.3.5. For the proof of the proposition we need two auxiliary models. One of these, the 'pure growth' model $\sigma(t)$, was already introduced in Section 3.3.2. The other, which has the same growth mechanism but where removal of trees takes place at time t_c only, is described below.

Removal at t_c only

Let I denote the set of all positive *even* integers i and consider the annuli $A(3^i, 5{\cdot}3^i) := B(5 \cdot 3^i) \backslash B(3^i), i \in I$. Note that these annuli are pairwise disjoint. In the process we are going to describe, again every site can be vacant (have value 0) or occupied (value 1). By a 'surrounding i cluster' we will mean an occupied circuit C around O in the annulus $A(3^i, 5{\cdot}3^i)$, together with all occupied paths in $A(3^i, 5{\cdot}3^i)$ that contain a site in C. The process is completely determined by the Poisson growth clocks introduced in Section 3.3.2, in the following way. Initially each site is vacant. Whenever the growth clock of a site rings, the site becomes occupied. When it is already occupied, the clock is ignored. Destruction ($1 \to 0$ transitions) only takes place at time t_c: at that time, for each positive even integer i, each 'surrounding i cluster' is instantaneously made vacant. After t_c the growth mechanism proceeds as before. Note that this model is closely related to the model associated with the variables $(Z_i^c : i \in \mathbb{Z}^2)$ from Subsection 2.4.3. Let $\xi_v(t)$ denote the value of site v at time t. Earlier in this paper we mentioned an obvious but useful relation (comparison) between the pure growth process $\sigma(\cdot)$ and the forest-fire process $\eta(\cdot)$. There is also a useful relation between $\xi(\cdot)$ and $\eta(\cdot)$, but its statement and proof are less straightforward, see Lemma 3.3.8 later in this section. Another lemma involving the process $\xi(\cdot)$ that will be important for us is the following.

Lemma 3.3.7. *If $\hat{\delta}_c > 0$, then there exist $\gamma < 1$ and $\varepsilon > 0$ such that for all $i \in I$,*

$$\mathcal{P}(\partial B(3^i) \to \partial B(3 \cdot 3^i) \text{ in the configuration } \xi(t_c + \varepsilon)) < \gamma. \qquad (3.3.7)$$

The proof of this lemma is very similar to that of Lemma 2.4.9. The probabilities above correspond to $q_{3^i}(p_c, \delta)$, with $\delta = 1 - e^{-\varepsilon}$. The modifications needed to go from discrete to continuous time are straightforward.

Proof of Proposition 3.3.5

Fix m. Since the probability in the statement of the proposition is monotone in m, we may assume that m is of the form 3^l for some even positive integer l. So each annulus $A(3^i, 5 \cdot 3^i)$, $i \in I$, defined above, is either contained in $B(m)$ or disjoint from $B(m)$.

Let $\tau = \tau(n,m)$ be the first time that η^n has a fire in $B(m)$, more precisely,

$$\tau := \inf\{t : \exists v \in B(m) \text{ s.t. } \eta_v^n(t) = 0 \text{ and } \eta_v^n(t^-) = 1\}.$$

Next, define for $\lambda \in (0,1)$,

$$\begin{aligned} K(\lambda) &:= \frac{1}{\sqrt[3]{\lambda}} \\ k(\lambda) &:= \frac{1}{\sqrt[4]{\lambda}} \\ A(k(\lambda), K(\lambda)) &:= B(K(\lambda)) \backslash B(k(\lambda)). \end{aligned} \tag{3.3.8}$$

Further, define the following events:

$$\begin{aligned} B_1 = B_1(\lambda) &:= \{\text{no ignitions in } B(K(\lambda)) \text{ before or at time } \tau\} \\ B_2 = B_2(\lambda) &:= \{\sigma(t_c) \text{ has a vacant *-contour in } A(k(\lambda), K(\lambda))\}, \end{aligned}$$

where by '*-contour' we mean a contour (surrounding 0) in the matching lattice (i.e. the lattice obtained from the square lattice by adding the two 'diagonal edges' in each face of the square lattice).

We will use the following relation between the forest-fire process $\eta(\cdot)$ and the auxiliary process $\xi(\cdot)$ described in the previous subsection.

Lemma 3.3.8. *Let $\lambda \in (0,1)$. On $B_1 \cap B_2$ we have, for all $t > \tau$, all $v \in B(k(\lambda)) \setminus B(m)$ and all n, that*

$$\eta_v^n(t) \leq \xi_v(t).$$

Proof. [Lemma 3.3.8]. Suppose $B_1 \cap B_2$ holds. Take n, t and v as in the statement of the lemma. Obviously, we may assume that $k(\lambda) > m$. To simplify notation we will, during the proof of this lemma, omit the superscript n from η, and the argument λ from k and K. Suppose $\xi_v(t) = 0$. We have to show that then also $\eta_v(t) = 0$. Since $\xi_v(t) = 0$, the growth clock of v does not ring in the interval $(t_c, t]$, and we assume that just before t_c the occupied ξ-cluster of v surrounds $B(m)$. If not, the desired conclusion follows trivially. From the definitions of the processes it then follows that the occupied σ-cluster of v at time t_c, which we will denote by $C_v^\sigma(t_c)$, surrounds $B(m)$. By B_2 we have that $C_v^\sigma(t_c)$ is in the interior of a vacant (that is, having $\sigma(t_c) = 0$) *-circuit in $A(k,K)$. Clearly, $\eta \equiv 0$ on this circuit during the time interval $(0, t_c]$, which prevents fires starting in its exterior to reach its interior. From this, and the event B_1, we conclude that $\tau > t_c$, and that $\eta(t_c)$ and $\sigma(t_c)$ agree in the interior of the circuit. In particular, the occupied η-cluster of v at time t_c, denoted by $C_v^\eta(t_c)$, equals the above mentioned set $C_v^\sigma(t_c)$. From B_1 it follows that at time τ a connected set is burnt which contains sites in $B(m)$ as well as in the complement of $B(K)$. But then it also contains a site in $C_v^\eta(t_c)$, because $C_v^\eta(t_c)$ surrounds $B(m)$ and lies inside

$B(K)$. So the whole set $C_v^\eta(t_c)$, and in particular v, burns at some time $s \in (t_c, \tau]$. Since the growth clock of v does not ring between time t_c and t, it follows that indeed $\eta_v(t) = 0$. This completes the proof of Lemma 3.3.8. □

Now we go back to the proof of the proposition. Assume $\hat{\delta}_c > 0$. Choose ε and γ as in Lemma 3.3.7. By (3.3.2) it is sufficient to prove that

$$\lim_{\lambda \downarrow 0} \limsup_{n \to \infty} \mathcal{P}_\lambda(\eta^n \text{ has at least 2 fires in } B(m) \text{ in } (t_c, t_c + \varepsilon)) = 0. \tag{3.3.9}$$

Define, in addition to B_1 and B_2 above, the event

$$\tilde{B}_1 = \{\text{no ignitions in } B(K(\lambda)) \text{ in the time interval } (0, t_c + \varepsilon)\}.$$

We have

$$\begin{aligned} &\mathcal{P}(\text{at least 2 fires in } B(m) \text{ in } (t_c, t_c + \varepsilon)) \\ &\quad \leq \mathcal{P}(\{\text{at least 2 fires in } B(m) \text{ in } (t_c, t_c + \varepsilon)\} \cap \tilde{B}_1 \cap B_2) \\ &\quad + \mathcal{P}(\tilde{B}_1^c) + \mathcal{P}(B_2^c). \end{aligned} \tag{3.3.10}$$

Now note that $\tilde{B}_1$ does not depend on n, and that

$$\mathcal{P}(\tilde{B}_1^c) \leq \lambda \, |B(K(\lambda))| \, (t_c + \varepsilon) \to 0, \quad \text{as } \lambda \downarrow 0, \tag{3.3.11}$$

by the definition of $K(\lambda)$ (3.3.8).

Next, note that the probability of B_2 does not depend on n either:

$$\mathcal{P}(B_2^c) \leq \mathcal{P}\{\partial B(k(\lambda)) \leftrightarrow \partial B(K(\lambda)) \text{ in } \sigma(t_c)\} \to 0, \text{ as } \lambda \downarrow 0, \tag{3.3.12}$$

by a well-known result from ordinary percolation and the fact that $K(\lambda)/k(\lambda) \to \infty$ as $\lambda \downarrow 0$.

Finally, we handle the event in the first term on the right hand side of (3.3.10). Since we will take limits as $\lambda \downarrow 0$, we may restrict to λ's for which $k(\lambda) > m$. Then we have the following relation between events:

$$\begin{aligned} &\Big\{\{\text{at least 2 fires in } B(m) \text{ in } (t_c, t_c + \varepsilon)\} \cap \tilde{B}_1 \cap B_2\Big\} \\ &\quad = \Big\{\{\tau \in (t_c, t_c + \varepsilon) \text{ and at least 1 fire in } B(m) \text{ in } (\tau, t_c + \varepsilon)\} \cap \tilde{B}_1 \cap B_2\Big\} \\ &\quad \subset \{\{\partial B(m) \leftrightarrow \partial B(k(\lambda)) \text{ in } \eta^n(s) \text{ for some } s \in (\tau, t_c + \varepsilon)\} \cap B_1 \cap B_2\} \\ &\quad \subset \{\partial B(m) \leftrightarrow \partial B(k(\lambda)) \text{ in } \xi(t_c + \varepsilon)\}, \end{aligned} \tag{3.3.13}$$

where the second inclusion follows from Lemma 3.3.8 (and the monotonicity of $\xi(t)$ for $t > t_c$), and the first inclusion holds because, by the event $\tilde{B}_1$, fires in $B(m)$ before time $t_c + \varepsilon$, can only arrive from outside $B(K(\lambda))$. To handle the probability of the last event in (3.3.13), first observe that, for each i, the random variables $\xi_v(t)$, $t \geq 0$, $v \in A(3^i, 5 \cdot 3^i)$ are completely determined by the Poisson clocks assigned to the sites

inside the annulus $A(3^i, 5 \cdot 3^i)$. We use the notation $I(\lambda)$ for the set of all positive even integers j with $m < 3^j < 5 \cdot 3^j \leq k(\lambda)$. Since the annuli $A(3^i, 5 \cdot 3^i)$, $i \in I(\lambda)$ are disjoint, we get from Lemma 3.3.7 that

$$\mathcal{P}\{\partial B(m) \leftrightarrow \partial B(k(\lambda)) \text{ in } \xi(t_c + \varepsilon)\} \leq \gamma^{|I(\lambda)|}. \tag{3.3.14}$$

Combining (3.3.13) and (3.3.14), and using that $k(\lambda)$, and hence also $|I(\lambda)|$ goes to ∞ as $\lambda \downarrow 0$, we get

$$\lim_{\lambda \downarrow 0} \limsup_{n \to \infty} \mathcal{P}(\{\text{at least 2 fires in } B(m) \text{ before time } (t_c + \varepsilon)\} \cap \tilde{B}_1 \cap B_2) = 0. \tag{3.3.15}$$

Combining (3.3.10), (3.3.11), (3.3.12) and (3.3.15) completes the proof of Proposition 3.3.5.

□

Proof of Theorem 3.3.4

Suppose $\hat{\delta}_c > 0$ and that for all $t > t_c$ there exists an $m = m(t)$ such that

$$\limsup_{\lambda \downarrow 0} \liminf_{n \to \infty} \mathcal{P}_\lambda(\text{fire in } B(m) \text{ before time } t) > 1/2. \tag{3.3.16}$$

We will show that this leads to a contradiction. Choose t as in Proposition 3.3.5. Now take $u \in (t_c, t)$. By (3.3.16) there exist an integer m_0 and an $\alpha(u) > 0$ such that

$$\limsup_{\lambda \downarrow 0} \liminf_{n \to \infty} \mathcal{P}_\lambda(\text{fire in } B(m_0) \text{ before time } u) > 1/2 + \alpha(u). \tag{3.3.17}$$

By (3.3.1) and the continuity of θ we can choose an $s \in (t_c, u)$ with

$$\limsup_{\lambda \downarrow 0} \liminf_{n \to \infty} \mathcal{P}_\lambda(\text{fire in } B(m_0) \text{ before time } s) \leq \alpha(u)/2. \tag{3.3.18}$$

By (3.3.16) there exists an m_1 (which we can choose larger than m_0) such that

$$\limsup_{\lambda \downarrow 0} \liminf_{n \to \infty} \mathcal{P}_\lambda(\text{fire in } B(m_1) \text{ before time } s) > 1/2. \tag{3.3.19}$$

Clearly,

$$\begin{aligned}
&\mathcal{P}(\text{fire in } B(m_0) \text{ before time } u) \\
&\quad \leq \mathcal{P}(\text{fire in } B(m_1) \text{ before time } s \text{ and fire in } B(m_0) \text{ between times } s \text{ and } u) \\
&\quad + \mathcal{P}(\text{fire in } B(m_0) \text{ before time } s) \\
&\quad + \mathcal{P}(\textit{no} \text{ fire in } B(m_1) \text{ before time } s).
\end{aligned} \tag{3.3.20}$$

This inequality stays valid if we take the $\liminf_{n \to \infty}$ on both sides. Then we replace the 'lim inf' by a 'lim sup' in front of the first and last term on the right hand side. Finally, we take the $\limsup_{\lambda \downarrow 0}$ on both sides.

Then, by Proposition 3.3.5, the first term on the r.h.s. will vanish. Using this, and applying (3.3.18) and (3.3.19) to the second and the third term, respectively, yields

$$\limsup_{\lambda \downarrow 0} \liminf_{n \to \infty} \mathcal{P}(\text{fire in } B(m_0) \text{ before time } u) \leq 1/2 + \alpha(u)/2,$$

which contradicts (3.3.17). This completes the proof of Theorem 3.3.4.

□

3.4 The binary tree

In this section we consider the same dynamics as for the process η in the previous section, but now we take the binary tree instead of the square lattice. Moreover, we restrict ourselves to the *directed* binary tree. On the directed binary tree we can make use of iteration arguments which made it possible to obtain some results.

By the infinite binary directed tree, denoted by $\mathcal{T}$, we mean the tree where one site (called the root) has two edges, all other sites have three edges, and where all edges are oriented in the direction of the root. The root will be denoted by O. By the children of a site v we mean the two sites *from* which there is an edge *to* v. (And we say that v is the parent of these sites). By the first generation of v we mean the set of children of v, by the second generation the children of the children of v etc.. The subgraph of $\mathcal{T}$ containing O and its first n generations will be denoted by $\mathcal{T}(n)$.

Let us now describe the model in detail. Because it is not clear whether the forest-fire process exists on $\mathcal{T}$, we restrict ourselves to finite subtrees $\mathcal{T}(n)$. Initially all sites are vacant. As in the original model vacant sites become occupied at rate 1 and occupied sites are ignited at rate λ. When a site v is ignited, instantaneously each site on the occupied path from v in the direction of the root is made vacant.

The forest-fire interpretation is not very natural here. More natural is the interpretation in terms of a nervous system: replace the word 'site' by 'node', 'occupied' by 'alert', 'vacant' by 'recovering' and 'ignition' by 'arrival of a signal from outside the system'. Then the above description says that whenever an alert node v receives a signal (either from a child, or from outside the system), it immediately transmits it to its parent (except when $v = O$, in which case it 'handles' the signal itself), after which it needs an exponentially distributed recovering time to become alert again.

As before we use '1' to represent an occupied ('alert') and '0' to represent a vacant ('recovering') vertex. Let $\zeta_v(t) \in \{0, 1\}$ denote the state of vertex v at time t. If we want to stress dependence on n we write $\zeta_v^n(t)$.

As we have seen before, the processes $\zeta^n(\cdot)$ can be completely described in terms of independent Poisson growth and ignition clocks, assigned to the sites of $\mathcal{T}$.

Recall that site percolation on the binary tree has critical probability $1/2$, and percolation probability function $\theta(p) = (2p - 1)/p$, for $p \geq 1/2$. Combining this with the same arguments that led to (3.3.1) shows that, if we first let $n \to \infty$ and then

$\lambda \downarrow 0$, the probability that the root burns before time $\ln(2)$ goes to 0, and, moreover, that for $t > \ln(2)$

$$\limsup_{\lambda \downarrow 0} \limsup_{n \to \infty} \mathcal{P}_\lambda(O \text{ burns before time } t) \leq \frac{1 - 2e^{-t}}{1 - e^{-t}}. \tag{3.4.1}$$

A nice feature of the binary tree is that we can (quite simply in fact) also prove a lower bound (compare with Open Problem 3.3.1 for the square lattice):

Lemma 3.4.1. *For all $t > \ln(2)$,*

$$\liminf_{\lambda \downarrow 0} \limsup_{n \to \infty} \mathcal{P}_\lambda(\zeta^n \textit{ has a fire in } O \textit{ before time } t) \geq \frac{1}{2} \frac{1 - 2e^{-t}}{1 - e^{-t}}. \tag{3.4.2}$$

Note that this lower bound is half the upper bound (3.4.1).

Proof. Define the functions

$$f_n^\lambda(t) := \mathcal{P}_\lambda(\zeta^n \text{has a fire in } O \text{ before time } t), \text{ for } t > 0,$$

and

$$g_n^\lambda(s,t) := f_n^\lambda(t) - f_n^\lambda(s), \text{ for } 0 < s < t,$$

i.e. the probability that the first time that O burns is between s and t.

Fix a $t > \ln(2)$ and take $\tilde{t} \in (\ln(2), t)$. Suppose that

$$\liminf_{\lambda \downarrow 0} \limsup_{n \to \infty} f_n^\lambda(t) < \frac{1}{2} \frac{1 - 2e^{-\tilde{t}}}{1 - e^{-\tilde{t}}}. \tag{3.4.3}$$

We will show that this leads to a contradiction. By (3.4.3) there exist an $\alpha > 0$ and a sequence $(\lambda_i,\ i = 1, 2, \cdots)$, which is decreasing, converges to 0 and has, for all i

$$\limsup_{n \to \infty} f_n^{\lambda_i}(t) < \frac{1}{2} \frac{1 - 2e^{-\tilde{t}}}{1 - e^{-\tilde{t}}} - \alpha. \tag{3.4.4}$$

Fix j large enough such that

$$e^{-\lambda_j \tilde{t}}[1 + 2\alpha(1 - e^{-\tilde{t}}]) > 1. \tag{3.4.5}$$

The reason for this choice will become clear later.

Observe that, if v and w are the children of O, the processes $\zeta_v^{n+1}(\cdot)$ and $\zeta_w^{n+1}(\cdot)$ are independent copies of $\zeta_O^n(\cdot)$ and are also independent of the Poisson clocks at O. Also observe that, to ensure that the first fire at the root occurs between times $\tilde{t}$ and t, it is sufficient that the growth clock of O rings before time $\tilde{t}$, no ignition occurs at the root before time $\tilde{t}$, at least one of its children burns between times $\tilde{t}$ and t and none of its children burn before time $\tilde{t}$. Hence, by these observations,

$$g_{n+1}^\lambda(\tilde{t}, t) \geq (1 - e^{-\tilde{t}})\, e^{-\lambda \tilde{t}} \left[g_n^\lambda(\tilde{t}, t)^2 + 2 g_n^\lambda(\tilde{t}, t)\,(1 - f_n^\lambda(\tilde{t})) \right]. \tag{3.4.6}$$

Now we take λ equal to λ_j in (3.4.6), and apply (3.4.4), noting that $f_n^\lambda(t) \geq f_n^\lambda(\tilde{t})$. This gives that (with the abbreviation g_k for $g_k^{\lambda_j}(\tilde{t}, t)$, $k = 1, 2, \cdots$) for all sufficiently large n

$$\begin{aligned} g_{n+1} &\geq (1 - e^{-\tilde{t}})\, e^{-\lambda_j \tilde{t}}\, 2 g_n \,(1 - f_n^{\lambda_j}(\tilde{t})) \\ &\geq g_n \cdot \left[e^{-\lambda_j \tilde{t}} [1 + 2\alpha(1 - e^{-\tilde{t}})] \right]. \end{aligned} \tag{3.4.7}$$

However, the factor behind g_n in the r.h.s. of (3.4.7) does not depend on n and is, by (3.4.5), strictly larger than 1, so that the sequence of g_n's 'explodes': a contradiction. Hence

$$\liminf_{\lambda \downarrow 0} \limsup_{n \to \infty} f_n^\lambda(t) \geq \frac{1}{2} \frac{1 - 2e^{-\tilde{t}}}{1 - e^{-\tilde{t}}}. \tag{3.4.8}$$

This holds for each $\tilde{t} \in (t_c, t)$. Letting $\tilde{t} \uparrow t$ in (3.4.8) completes the proof of Lemma 3.4.1. □

By a similar 'independent copies' observation as used a few lines above (3.4.6) (now for all sites in the $m-$th generation of the root), Lemma 3.4.1 immediately gives the following corollary (compare with Theorem 3.3.4 and Proposition 3.3.5):

Corollary 3.4.2. *For all $t > \ln(2)$, all $\varepsilon > 0$ and all k, there exists an integer m such that*

$$\liminf_{\lambda \downarrow 0} \limsup_{n \to \infty} \mathcal{P}_\lambda(\zeta^n \text{ has at least } k \text{ fires in } \mathcal{T}(m) \text{ before time } t) > 1 - \varepsilon.$$

3.5 Modified models

As we mentioned in the introduction of this section, the double limit $n \to \infty$ and $\lambda \downarrow 0$ is somewhat annoying. Therefore we study modified models that involve only one parameter. Although these modified models are seemingly simpler than the original one we think the problems are, essentially, as hard as before.

3.5.1 Ignition of sufficiently large clusters

Consider the following alternative forest-fire model. Again we work on the square lattice. In this model the growth mechanism is the same as before (that is, vacant sites become occupied at rate 1), but the ignition mechanism is different: instead of the ignition rate λ we have an (integer) parameter L. The ignition rule now is that whenever a cluster of size $\geq L$ occurs, it is instantaneously ignited and burnt down (that is, each of its sites becomes vacant). A very pleasant feature of this model is that, since the interactions now have finite range, it can be defined on the infinite lattice using a standard graphical construction. This frees us from the necessity to first work on $B(n)$ and later take limits as $n \to \infty$ so that we are left with the single limit $L \to \infty$. As before, we start at time 0 with all sites vacant. Let $\eta_v^{[L]}(t)$ denote the value (0 or 1) of site v at time t. The analogue of Open Problem 3.3.1 is:

Open Problem 3.5.1. Is, for all $t > t_c$

$$\liminf_{L\to\infty} \mathcal{P}(\eta^{[L]} \text{ has a fire in } O \text{ before time } t) > 0\,? \tag{3.5.1}$$

Similarly, there is a straightforward analogue of Open Problem 3.3.2.

We have, with $\hat{\delta}_c$ as before, (2.4.9), analogues of Theorem 3.3.4 and Proposition 3.3.5.

Theorem 3.5.2. *If $\hat{\delta}_c > 0$, then there exists a $t > t_c$ such that for all m,*

$$\liminf_{L\to\infty} \mathcal{P}(\eta^{[L]} \textit{ has a fire in } B(m) \textit{ before time } t) \leq 1/2. \tag{3.5.2}$$

Proposition 3.5.3. *If $\hat{\delta}_c > 0$, then there exists $t > t_c$ such that for all m,*

$$\lim_{L\to\infty} \mathcal{P}(\eta^{[L]} \textit{ has at least 2 fires in } B(m) \textit{ before time } t) = 0. \tag{3.5.3}$$

Theorem 3.5.2 follows from Proposition 3.5.3 in the same way as Theorem 3.3.4 from Proposition 3.3.5. The proof of Proposition 3.5.3 is very similar to that of Proposition 3.3.5 and we only indicate the main modifications: instead of (3.3.8) we define

$$\begin{aligned} K_L &:= L^{1/3}, \\ k_L &:= L^{1/4}. \end{aligned}$$

Next, the events B_1, B_2 are replaced by the single event

$$B_3 := \{\sigma(t_c) \text{ has a vacant } *\text{-circuit in } A(k_L, K_L)\},$$

and Lemma 3.3.8 is replaced by the following lemma, whose proof is a straightforward modification of that of the former.

Lemma 3.5.4. *On B_3 we have, for all $t > \tau$ and all $v \in B(k_L) \setminus B(m)$ that*

$$\eta_v^{[L]}(t) \leq \xi_v(t). \tag{3.5.4}$$

Of course we take m as before, and $\tau = \tau(L, m)$ is now defined as the first time that $\eta^{[L]}$ has a fire in $B(m)$. The proof of Proposition 3.5.3 now proceeds as before.

3.5.2 Fires that start on the boundary

Yet another model that serves as a forest-fire model is the following. On a box $B(n)$ we let trees grow with rate 1, independently, as in our original model. At time 0 every site is vacant. As soon as an occupied cluster touches the boundary, it is destroyed instantaneously. Although this process is, like the original one in Section 3.3.2, defined on *finite* boxes, whose size tends to infinity, it has only one parameter, like the model in the previous subsection. Let $\eta_v^{[\partial n]}(t)$ denote the state of site v at time t. Again we can pose similar questions and prove analogues of the results as for the original forest-fire model.

Open Problem 3.5.5. Is, for all $t > t_c$

$$\liminf_{n\to\infty} \mathcal{P}(\eta^{[\partial n]} \text{ has a fire in } O \text{ before time } t) > 0\,? \tag{3.5.5}$$

We have

Theorem 3.5.6. *If $\hat{\delta}_c > 0$, then there exists a $t > t_c$ such that for all m,*

$$\liminf_{n\to\infty} \mathcal{P}(\eta^{[\partial n]} \textit{ has a fire in } B(m) \textit{ before time } t) \leq 1/2. \tag{3.5.6}$$

Proposition 3.5.7. *If $\hat{\delta}_c > 0$, then there exists an $t > t_c$ such that for all m,*

$$\lim_{n\to\infty} \mathcal{P}(\eta^{[\partial n]} \textit{ has at least 2 fires in } B(m) \textit{ before time } t) = 0. \tag{3.5.7}$$

Given Proposition 3.5.7 the proof of Theorem 3.5.6 is a straightforward analogue of that of Theorem 3.3.4. The proof of Proposition 3.5.7 is similar to that of Proposition 3.3.5; modifications of the same nature as for Proposition 3.5.3 are needed. This time replace the pair of events B_1, B_2 by the event

$$B_3 = \{\exists \text{ vacant } *\text{-circuit in } A(\sqrt{n}, n) \text{ in } \sigma(t_c)\}. \tag{3.5.8}$$

Analogously to Lemma 3.3.8 we then obtain, with $\tau = \tau(n, m)$ now the first time that $\eta^{[\partial n]}$ has a fire in $B(m)$:

Lemma 3.5.8. *On the event B_3 we have for all $v \in B(\sqrt{n})\backslash B(m)$ and all $t > \tau$, that*

$$\eta_v^{[\partial n]}(t) \leq \xi_v(t).$$

Using this lemma, we can practically repeat the proof of Proposition 3.3.5.

Chapter 4

Frozen percolation

4.1 Introduction

Let G be an arbitrary graph and let $N \geq 2$ be an integer. We attach a random uniform $[0,1]$ number U_i to each site i of G, independently. Informally, the frozen percolation model can be described as follows. At time 0 all sites are at rest, or coloured white. At time U_i, the site i becomes active, or coloured green. If a site is in a cluster of green sites that has size at least N, this cluster *freezes*, or becomes red, instantaneously. The term frozen percolation was firstly introduced by Aldous [Ald00] in a slightly different setting[1]. We discuss his results briefly in Section 4.3.

We use both the terminology of freezing and of the colours in what follows, whichever seems more appropriate. So, how does our model evolve in time? Some sites get activated and never freeze, others will freeze at the moment that they are activated and some after they were activated. It follows from the above description that at time 1 all sites are green or red.

4.1.1 Existence of frozen percolation models

It should be clear that the frozen percolation model exists for finite graphs G and any N. In particular, when N is larger than the number of sites in G, each site turns green at some point, and stays green forever. As soon as we take the underlying graph infinite, existence of the model is not immediately clear.

The existence of the frozen percolation model is relatively easy on $\mathbb{Z}$. For any time $t < 1$, there will be infinitely many sites j on the left and right half lines that have

[1] Aldous freezes infinite clusters. Readers familiar with Aldous' work should note that in his description, the boundary of a frozen cluster is by definition at rest at the time a cluster freezes, and stays at rest forever. In our description however, at the time when a frozen cluster arises, its boundary can be either at rest (and allowed to get activated and even frozen later on), or already frozen.

$U_j \geq t$. This 'breaks up' the line in finite pieces; the state of any site at time t then depends on finitely many variables only and this is sufficient to show existence.

More or less the same reasoning can be used to show that the process can be constructed on $\mathbb{Z}^d$, for fixed N. We use the standard construction for finite range particle systems, see e.g. [Lig85]. Suppose that a configuration $\omega \in \{\text{white, green, red}\}^{\mathbb{Z}^d}$ at time s is given. We should show that for all times $t > s$, we can determine the state of an arbitrary site, say O, at time t, by looking at ω on a finite subset $V \subset \mathbb{Z}^d$ and the U-values on V. Consider the box $B(N) = [-N, N)^d$ and let for $x \in (2N\mathbb{Z})^d$

$$B_x = x + B(N).$$

We define for $x \in (2N\mathbb{Z})^d$,

$$Y_x = \begin{cases} 1 & \text{if } \exists z \in B_x : U_z \in [s,t], \\ 0 & \text{otherwise.} \end{cases}$$

Note that the Y-variables are independent. We define x_1, $x_2 \in (2N\mathbb{Z})^d$ to be neighbours if there exist $z_1 \in B_{x_1}$ and $z_2 \in B_{x_2}$ such that the graph distance between z_1 and z_2 is smaller than $2N$. The thus constructed block variables $(Y_x : x \in (2N\mathbb{Z})^d)$ form a percolation process on the graph G, where G is the graph induced by $\mathbb{Z}^d$, with extra connections between two point $z_1 = (z_{1_1}, z_{1_2}, \ldots, z_{1_d})$ and $z_2 = (z_{2_1}, z_{2_2}, \ldots, z_{2_d})$ if $\max_{j=1,\ldots,d} |z_{1_j} - z_{2_j}| \leq 1$.

Now take $t - s$ so small that $1 - (1 - t + s)^{(2N)^d}$ is smaller than the critical probability of this percolation process on G. Mind that the length of the interval $[s, t]$ does not depend on s. The choice of t makes our process subcritical, so that any cluster of blocks with Y-value 1 (in particular the cluster of Y_O) is finite, and separated from infinity by a boundary of blocks with Y-value 0. The set V mentioned above is the set of all sites within the blocks in this boundary around the cluster of Y_O and its interior. Given that for all sites z in any block in the boundary, $U_z \notin [s, t]$, the sites outside V cannot influence the colour of site O in the time interval $[s, t]$. Repeating this argument starting from the configuration at time t, we construct the process for all times.

Now we consider infinite graphs, it is tempting to take N equal to infinity, i.e. only infinite clusters are allowed to freeze. On $\mathbb{Z}$, this does not lead to interesting behaviour: all sites in $\mathbb{Z}$ stay green up to time 1 and at time 1, every site becomes red. On a (regular) binary tree, the process does exist in the slightly different form presented by Aldous [Ald00]. Benjamini and Schramm have shown (but not published) that on $\mathbb{Z}^2$, the frozen percolation process where infinite clusters freeze (in Aldous' setting), does not exist. See also [BT01]. See Subsection 4.3 for a frozen percolation model where infinite clusters freeze on the tree, using our dynamics.

4.2 Frozen percolation on $\mathbb{Z}$

This subsection is meant as an appetizer for the frozen percolation process on the tree, in the next subsection.

As was mentioned in the introduction, the frozen percolation process exists on $\mathbb{Z}$ for any value of N. By modifying on a null set, we assume that all $U_i, i \in \mathbb{Z}$ are different. For each $i \in \mathbb{Z}$, we define $Y_i = Y_i(N)$ as the time that site i freezes, with $Y_i = \infty$ if i never freezes. Fix N and let the set $\mathcal{W}(t)$ be defined as the set of all sites that are at rest at time t. Further, let the set $\mathcal{G}(t)$ be defined as the set of all activated (but not yet frozen) sites at time t. We are particularly interested in $\mathcal{R}(t)$ that is defined as the set of all frozen sites at time t. Clearly, the behaviour of this set again depends on the parameter N. From the above description we have for each $i \in \mathbb{Z}$ with $Y_i < \infty$,

$$\begin{aligned} &i \in \mathcal{W}(t) \text{ for } t \in [0, U_i), \\ &i \in \mathcal{G}(t) \text{ for } t \in [U_i, Y_i), \\ &i \in \mathcal{R}(t) \text{ for } t \in [Y_i, 1]. \end{aligned} \tag{4.2.1}$$

In the next subsection, we explicitly compute the probability that a given site is in $\mathcal{R}(t)$, when $N = 2$.

In the final configuration, a frozen cluster has size between N and $2N-1$ and it is separated from other frozen clusters by 0 to $N-1$ activated sites[2]. So by translation invariance,

$$1 \geq \mathcal{P}(0 \in \mathcal{R}(1)) \geq \frac{N}{2N-1}, \tag{4.2.2}$$

so that at least half of the sites are eventually frozen if we let the parameter N to infinity. As was mentioned before, taking N equal to infinity gives us no interesting behaviour: every site i becomes activated at time U_i and at time 1 all sites turn red instantaneously. The question is whether the behaviour is different if we start with fixed N, compute the probability that a site is red and then send N to infinity. Can we find an upper bound in (4.2.2) that is strictly smaller than 1 if N goes to infinity, or equivalently, do we see green sites in the final configuration as N grows large? We will answer that question affirmatively in Subsection 4.2.2. As a bonus, we also obtain a better lower bound than the one in (4.2.2), see Remark 4.2.4.

4.2.1 Exact computations for $N = 2$

In this subsection we perform some explicit computations when $N = 2$. In particular, we compute the probability that a given site, say 0, is frozen before (or at) time t. Let

$$F(t) := \mathcal{P}(0 \in \mathcal{R}(t)).$$

Lemma 4.2.1. *The probability $F(t)$ that a given site is frozen before (or at) time t is given by*

$$F(t) = 4t + 4e^{-t} - \frac{1}{3}t^3 + 4te^{-t} + 2e^{-2t} - 6.$$

[2]Note that if a site freezes at time t_1, and a neighbouring site freezes at time $t_2 > t_1$ they are considered to be in two different red clusters, even though in the final configuration it may not be possible to distinct between the two.

In particular, the probability that a given site is eventually frozen equals $F(1) = \frac{8}{e} + \frac{2}{e^2} - 2\frac{1}{3} \approx 0.88$.

In the proof of this lemma, we turn to the half-line first. Instead of considering $\mathbb{Z}$ we restrict the process to $0, 1, \ldots$. The dynamics stays the same. This way, we only need to consider events where site 0 freezes together with site 1 (and possibly site 2). We have a combinatorial argument to find the probability that site 0 is eventually frozen on the half-line. This result also follows from the formula on the half-line for general t, obtained in the proof of Lemma 4.2.1, but we have included it here to give the reader more feeling for the model. We write $\mathcal{P}_h$ to indicate that we work on the half-line.

Lemma 4.2.2. *On the half-line,* $\mathcal{P}_h(0 \in \mathcal{R}(1)) = \frac{2}{e}$.

Proof. Now that we look at the final configuration, the actual times at which sites are activated are not important any more. The only thing that matters is the order in which sites are activated. Since the variables $(U_i : i \geq 0)$ are uniform and independent, it now becomes a matter of counting. Below we list possible events that cause site 0 to freeze and their respective probabilities.

event	probability
$U_0 < U_1$	$1/2$
$U_1 < U_0$ and $U_0 < U_2$	$1/3!$
$U_2 < U_3 < U_1 < U_0$	$1/4!$
$U_3 < U_2 < U_1 < U_0$ and $U_2 < U_4$	$3/5!$
$U_4 < U_5 < U_3 < U_2 < U_1 < U_0$	$1/6!$
$U_5 < U_4 < U_3 < U_2 < U_1 < U_0$ and $U_4 < U_6$	$5/7!$
....	....

The events above are disjoint and it is not hard to see that, going on like this, we obtain all possibilities. Hence

$$P_h(0 \in \mathcal{R}(1)) \quad = \quad \sum_{k=1}^{\infty} \frac{1}{(2k)!} + \frac{2k-1}{(2k+1)!} = \frac{2}{e}.$$

□

Proof. [Lemma 4.2.1]. For the result on the full line $\mathbb{Z}$, we turn to the simpler case on the half-line first. The main observation is that Y_0, the time that site 0 freezes on $\mathbb{Z}$, is the minimum of the time that 0 freezes together with 1 and the time 0 freezes together with -1. So, let for all $i \in \mathbb{Z}$, $\vec{Y}_i$ denote the time that i freezes on the right half-line $i, i+1, \ldots$ when the left half-line $\ldots, i-2, i-1$ is ignored. Likewise, let $\overleftarrow{Y}_i$ denote the time that i freezes on the left half-line when the right half-line is ignored. The time Y_i that site i freezes can be written as

$$Y_i = \min\{\vec{Y}_i, \overleftarrow{Y}_i\}. \tag{4.2.3}$$

We establish a recursive relation:

$$\begin{aligned} F_h(t) &:= \mathcal{P}_h(\vec{Y_0} \leq t) \\ &= \mathcal{P}_h(U_0 \leq \vec{Y_1} \leq t) + \mathcal{P}_h(\vec{Y_1} > t, U_0 \leq t, U_1 \leq t) \\ &= \mathcal{P}_h(U_0 \leq \vec{Y_1} \leq t) + \mathcal{P}_h(U_0 \leq t) \left[\mathcal{P}_h(U_1 \leq t) - \mathcal{P}_h(U_1 \leq t, \vec{Y_1} \leq t)\right]. \quad (4.2.4) \end{aligned}$$

By symmetry, we can establish a similar equation for $\mathcal{P}_h(\overleftarrow{Y_0} \leq t)$ on the half-line $\ldots, -1, 0$. Evaluating the r.h.s. of (4.2.4) we obtain, using $\mathcal{P}_h(\vec{Y_1} \leq t, U_1 \leq t) = \mathcal{P}_h(\vec{Y_1} \leq t)$,

$$\begin{aligned} F_h(t) &= \int_0^t [F_h(t) - F_h(s)]ds + t(t - F_h(t)) \\ &= t^2 - \int_0^t F_h(s)ds. \quad (4.2.5) \end{aligned}$$

Differentiating with respect to t we obtain the following differential equation.

$$\frac{dF_h(t)}{dt} = 2t - F_h(t). \quad (4.2.6)$$

Using $F_h(0) = 0$ as a boundary condition, equation (4.2.6) has

$$F_h(t) = 2t - 2 + 2e^{-t} \quad (4.2.7)$$

as its solution. By symmetry, $\overleftarrow{Y_0}$ also has distribution function $F_h(t)$. Note that $F_h(1) = \frac{2}{e}$ was already computed in Lemma 4.2.2.

Now, using (4.2.3) and (4.2.7) we compute

$$\begin{aligned} F(t) = \mathcal{P}(Y_0 \leq t) &= \mathcal{P}(\min\{\vec{Y_0}, \overleftarrow{Y_0}\} \leq t) \\ &= 2F_h(t) - \mathcal{P}(\vec{Y_0} \leq t, \overleftarrow{Y_0} \leq t). \quad (4.2.8) \end{aligned}$$

We use (4.2.4) to compute the last term on the right:

$$\begin{aligned} \mathcal{P}(\vec{Y_0} \leq t, \overleftarrow{Y_0} \leq t) &= \int_0^t [F_h(t) - F_h(s)]^2 ds + t[t - F_h(t)]^2 \\ &+ 2(t - F_h(t)) \int_0^t [F_h(t) - F_h(s)]ds. \quad (4.2.9) \end{aligned}$$

Combining (4.2.8), (4.2.9) and (4.2.7) proves the lemma.

□

Remark 4.2.3. A question that comes to mind immediately is the following: can we repeat the computation above for $N > 2$? In principle we could compute $\mathcal{P}_h(0 \in \mathcal{R}(1))$ for $N > 2$, analogously to Lemma 4.2.2. An integral equation like (4.2.5) is easily established for $N > 2$. Provided we can solve the equation (in practice, a problem already for $N = 3$) a result like (4.2.7) can be obtained. Nevertheless, it is much harder to take the final step from the half-line to the full line: if the site 0 freezes on the full line together with sites both on the left and on the right, it does not necessarily freeze on one of the half-lines, so that (4.2.3) does not hold. There are more possibilities to take into account. In theory, one can obtain a result like Lemma 4.2.1, but the computations get much more involved.

4.2.2 Green clusters have positive mass as $N \to \infty$

The frozen percolation model on $\mathbb{Z}$ with N *equal* to infinity, behaves differently from the model where we take N *approaching* infinity. If we freeze infinite clusters, there are no green sites in the final configuration. If we start with finite N and let N approach infinity, are there any green sites in the final configuration? More precisely, is the fraction of green sites bounded away from 0 as N goes to infinity? There is a simple argument to answer this question affirmatively, using the same basic idea as the argument below. However, it is less precise and therefore gives a much worse lower bound for the fraction of green sites than the one obtained in (4.2.13) below. Therefore, we present the more involved computation. We bound $\mathcal{P}^N(0 \in \mathcal{G}(1))$, where we use the superscript to emphasize the dependence on N. By translation invariance,

$$\begin{aligned} \mathcal{P}^N(0 \in \mathcal{G}(1)) &= \sum_{i=1}^{N-1} i\mathcal{P}^N([0, i-1] \subset \mathcal{G}(1), \{-1, i\} \subset \mathcal{R}(1)) \\ &= \sum_{i=1}^{N-1} i\mathcal{P}^N(-1 \in \mathcal{R}(U_0^-), i \in \mathcal{R}(U_{i-1}^-)). \end{aligned} \tag{4.2.10}$$

By definition, if $j \in \mathcal{R}(U_i^-)$ then $U_j < U_i$ and j is frozen before i is activated. Furthermore, $i \notin \mathcal{R}(U_i^-)$.

The probability that site -1 freezes before site 0 is activated, is bounded from below by the probability that site -1 freezes while at that time both the external boundary points (site 0 and some other site in $[-N-1, -2N]$) of its red cluster are not activated yet:

$$\mathcal{P}^N(-1 \in \mathcal{R}(U_0^-)) \geq \sum_{j=0}^{N-1} \mathcal{P}^N([-N-j, -1] \subset \mathcal{R}(\min\{U_0, U_{-N-j-1}\}^-)). \tag{4.2.11}$$

If j = 0, the probability above is equal to

$$\mathcal{P}^N(\min\{U_0, U_{-N-1}\} > \max\{U_{-N}, \dots, U_{-1}\}) = \frac{2}{(N+1)(N+2)}.$$

However, the site -1 can also freeze in a larger cluster of size $N+j$, where $j \in [1, \dots, N-1]$. The order at which the sites $U_{-N-j}, \dots, U_{-1}$ are activated is now important: the cluster on $[-N-j, -1]$ freezes (with both its external boundary points at rest) if and only if

$$\min\{U_0, U_{-N-j-1}\} > \max\{U_{-N-j}, \dots, U_{-1}\} \text{ and}$$

the site i such that $U_i = \max\{U_j : j \in [-N-j, -1]\}$ is in $[-N, -j-1]$.

So for $j \in [0, N-1]$, the corresponding term in (4.2.11) is

$$\frac{2}{(N+j+2)(N+j+1)} \frac{N+j-2j}{N+j},$$

so that

$$\mathcal{P}^N(-1 \in \mathcal{R}(U_0^-)) \geq \sum_{j=0}^{N-1} \frac{2}{(N+j+2)(N+j+1)} \frac{N-j}{N+j} = \frac{1}{2N+1}. \tag{4.2.12}$$

By symmetry, the same lower bound holds for $\mathcal{P}^N(i \in \mathcal{R}(U_{i-1}^-))$, for all i. Note that the event $\{-1 \in \mathcal{R}(U_0^-)\}$ depends on variables $(U_j : j \leq 0)$ only, and the event $\{i \in \mathcal{R}(U_{i-1}^-)\}$ depends on the variables $(U_j : j \geq i-1)$. So for $i > 1$ these events are independent and we apply this in (4.2.10) to compute

$$\begin{aligned}
\mathcal{P}^N(0 \in \mathcal{G}(1)) &= \sum_{i=1}^{N-1} i\mathcal{P}^N(-1 \in \mathcal{R}(U_0^-), i \in \mathcal{R}(U_{i-1}^-)) \\
&\geq \sum_{i=2}^{N-1} i\mathcal{P}^N(-1 \in \mathcal{R}(U_0^-)) \cdot \mathcal{P}^N(i \in \mathcal{R}(U_{i-1}^-)) \\
&\geq \sum_{i=2}^{N-1} i\Big(\frac{1}{2N+1}\Big)^2 = \frac{(N-2)(N+1)}{2(2N+1)^2} \longrightarrow 1/8, \qquad (4.2.13)
\end{aligned}$$

if $N \to \infty$. We have used (4.2.12) in the last inequality. This shows that green clusters indeed have positive mass if N grows large.

Remark 4.2.4. Note that by translation invariance, multiplying each term in the sum in (4.2.11) by $N+j$ gives us a lower bound for the probability that a site freezes, by assuming that it freezes while at that time both boundary points of its cluster are at rest:

$$\mathcal{P}^N(0 \in \mathcal{R}(1)) \geq \sum_{j=0}^{N-1} \frac{2(N-j)}{(N+j+1)(N+j+2)} \longrightarrow 2 - 2\ln(2) \approx 0.61,$$

as $N \to \infty$, which is already a better lower bound than (4.2.2) for the probability that a given site eventually freezes.

4.3 Frozen percolation on the tree

It is a priori not clear that the frozen percolation process where infinite clusters freeze exists on the regular binary tree. In a slightly different setting Aldous [Ald00] proved that a form of this process indeed exists and we will discuss his result briefly in the next subsection. Later we discuss the existence of the process on the binary tree for our version of frozen percolation.

4.3.1 Description of Aldous' result

Informally, the model considered by Aldous is as follows. Assign a uniform, independent 0-1 variable U_e to each edge e of the regular binary tree. Let $\mathcal{A}_0 = \emptyset$. For each edge e, at time $t = U_e$ set $\mathcal{A}_t = \mathcal{A}_{t-} \cup \{e\}$ if each end-vertex of e is in a finite cluster of $\mathcal{A}_{t-}$; otherwise set $\mathcal{A}_t = \mathcal{A}_{t-}$. That is, the boundaries of infinite clusters will never join the process. In the final configuration, there are infinite clusters, finite clusters and boundary edges in between. Aldous shows that this process exists by first 'guessing' what the distribution function of the time that an edge joins the process should be on the directed tree. Heuristic arguments suggest that this distribution function should be

$$G(t) = \begin{cases} 0 & t \in [0, 1/2) \\ 1 - \frac{1}{2t} & t \in [1/2, 1] \\ \frac{1}{2} & 1 < t < \infty \\ 1 & t = \infty. \end{cases} \tag{4.3.1}$$

Then (4.3.1) is used to construct a process on the undirected tree and finally Aldous shows that the model indeed meets with its intuitive description. We follow this line of reasoning below.

The model proposed by Aldous is meant as a model for polymerisation. In this model a polymer is made up of molecular units, where each unit is capable of forming three bonds. Before the critical time, there are only finite polymers (the sol) and later on infinite polymers (the gel) are observed. Aldous' model displays the behaviour observed by chemical physicists: beyond the gel point, the number of small polymers decreases but their average size retains a constant value. Formally, this means that *at any time* $t > t_c$(where t_c corresponds to the critical time of ordinary percolation on the tree), finite clusters of edges that have joined the process have the same distribution as critical percolation clusters. See Proposition 11 of [Ald00]. It is a perfect example of self-organised criticality (although Aldous does not mention SOC).

A priori, we see no reason why the boundary of infinite clusters should not join the process. Computations turn out to be harder if they do and we do not get such nice results as Aldous did. Nevertheless, it seems worth to try to construct the frozen percolation process where boundaries of infinite clusters are allowed to join the process. From now on, we refer to frozen percolation using our dynamics as modified frozen percolation to distinguish between the two versions. It is interesting to see whether the self-organised critical behaviour observed in Aldous' model is 'a coincidence', in the sense that a small perturbation of the dynamics causes the critical

behaviour to disappear. In the same spirit (see also Remark 4.3.11), we chose to study the site version of the modified frozen percolation model. We will see in Subsection 4.3.4 that, although the perfect correspondence to ordinary critical percolation is lost, the modified frozen percolation process still behaves critically.

4.3.2 Heuristics and computational arguments on the directed tree

Consider the directed rooted binary tree $\mathcal{T}$, whose root O has degree 2 and all other vertices have degree 3. Let each edge be directed away from the root. To each site, we assign a random uniform $[0,1]$ number U_i independently. The dynamics is as before, but we now freeze infinite directed rays: at time 0 all sites are at rest (coloured white). Each site i becomes activated (green) at time $t = U_i$, but as soon as i is in an infinite (directed) active cluster, this cluster freezes (becomes red) instantaneously. In this subsection, we argue what the behaviour of the directed modified frozen percolation process should be like. Later on, we use these heuristic arguments to formally construct the modified frozen percolation process.

Let 1 and 2 be the children of O. We use $\mathcal{T}[i]$, $i = 1, 2$ to denote the subtree that has site i as its root, see Figure 4.1. If the root O is activated, it freezes as soon as one

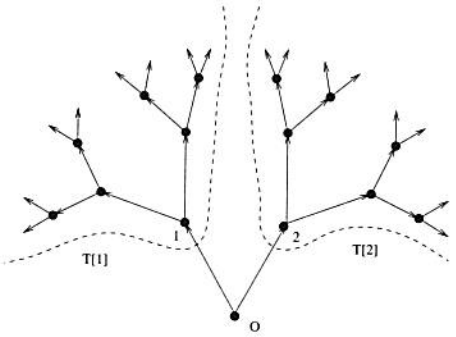

Figure 4.1: The rooted binary tree $\mathcal{T}$ and its subtrees $\mathcal{T}[1]$ and $\mathcal{T}[2]$.

of its children freezes. Suppose Y is the time that site O freezes. If O never freezes, we take Y equal to infinity. Since $\mathcal{T}[1]$ and $\mathcal{T}[2]$ are isomorphic to $\mathcal{T}$, we require that the time Y_1 (Y_2) that site 1 (2) freezes in $\mathcal{T}[1]$ ($\mathcal{T}[2]$, respectively) is distributed like Y. Furthermore, a site cannot freeze before it is in an infinite cluster of activated sites so Y should be at least $1/2$, by ordinary percolation results. To formalise these statements, define I to be $[1/2, 1] \cup \{\infty\}$. Let $\Phi(x, y, z)$ be the following function on $I \times I \times [0, 1]$.

$$\Phi(x,y,z) = \begin{cases} x \text{ if } x \geq z \\ y \text{ if } x < z \leq y \\ \infty \text{ otherwise.} \end{cases} \tag{4.3.2}$$

The heuristic arguments above require that

$$Y \stackrel{d}{=} \Phi(\min\{Y_1, Y_2\}, \max\{Y_1, Y_2\}, U_O). \tag{4.3.3}$$

The following lemma shows which distribution functions satisfy equality (4.3.3).

Lemma 4.3.1. *Let F be a (possibly defective) probability distribution function, of an I-valued r.v. with the following additional properties: F is continuous and differentiable with strictly positive derivative on $[1/2, 1]$ and has $F(1/2) = 0$. Let (Y_1, Y_2, U) be independent random variables, Y_1, Y_2 each having probability distribution function F and U having the uniform distribution on $[0, 1]$. Then,*

$\Phi(\min(Y_1, Y_2), \max(Y_1, Y_2), U)$ *again has probability distribution function F,*

if and only if

$$F(t) = \begin{cases} 0 & t \leq 1/2 \\ \ln(2t) & 1/2 < t \leq 1 \\ \ln(2) & 1 < t < \infty \\ 1 & t = \infty. \end{cases} \tag{4.3.4}$$

Proof. The equation $F(t) = \mathcal{P}(\Phi(\min(Y_1, Y_2), \max(Y_1, Y_2), U) \leq t)$ is equivalent to

$$F(t) = \int_0^t \mathcal{P}(t \geq \min(Y_1, Y_2) > s)ds \; + \int_0^t \mathcal{P}(\min(Y_1, Y_2) < s < \max(Y_1, Y_2) \leq t)ds.$$

Since both Y_1 and Y_2 have probability distribution function F this is equivalent to

$$\begin{aligned} F(t) &= \int_0^t \left[2F(t) - F(t)^2 - 2F(s) + F(s)^2\right] ds + 2\int_0^t (F(t) - F(s))F(s)ds \\ &= t[2F(t) - F(t)^2] + 2F(t)\int_0^t F(s)ds - \int_0^t 2F(s) + F(s)^2 ds. \end{aligned} \tag{4.3.5}$$

Differentiating with respect to t we obtain

$$\frac{dF(t)}{dt} = \frac{dF(t)}{dt}[2t(1 - F(t))] + 2\frac{dF(t)}{dt}\int_0^t F(s)ds. \tag{4.3.6}$$

Since F has positive derivative on $[1/2, 1]$, we may divide (4.3.6) by $\frac{dF(t)}{dt}$. The remaining integral equation has $F(t) = \ln(t) + C$, for some constant C, as its solution. Finally, we use $F(1/2) = 0$ to obtain $C = \ln(2)$.

The reversed implication follows from straightforward calculation. □

Note that the solution above is defective, i.e. has mass in infinity. We can interpret this distributions as the probability that a vertex freezes before time t, as we have seen above.

When we release the restriction that F must be differentiable with positive derivative on $[1/2,1]$ in the lemma above, we obtain more solutions. One of these is

$$F(t) = \begin{cases} 0 & t < \infty \\ 1 & t = \infty, \end{cases} \tag{4.3.7}$$

We have the following heuristic argument to rule out the zero solution (4.3.7) as a candidate. We approximate the directed process by considering processes on finite directed trees. As before, an independent uniform $[0,1]$ variable U_i is attached to each site $i \in \mathcal{T}$. Let $\mathcal{P}$ denote the measure governing the U-variables in $\mathcal{T}$. The root of the tree O is said to be at level 0 and a site is at level (depth) n if its distance to the root is n. For each n, the dynamics is as follows. Let $\mathcal{T}(n)$ denote the first $n+1$ levels of the tree. At time 0 all sites are at rest. Each site becomes activated at time U_i and as soon as there is a directed path of activated sites to the leaves of the tree (i.e. level n) this path freezes. For fixed n this process exists and as we did on $\mathbb{Z}$, we can define sets evolving in time to formalize the dynamics. Let $\mathcal{W}_n(0) := V_{\mathcal{T}(n)}$, all vertices of $\mathcal{T}(n)$, $\mathcal{G}(0) := \emptyset$ and $\mathcal{R}(0) := \emptyset$. For $t > 0$,

$$\begin{aligned} \mathcal{W}_n(t) &:= \{i \in \mathcal{T}(n) : U_i > t\}, \\ \mathcal{G}_n(t) &:= \{i \in \mathcal{T}(n) : U_i \le t \text{ and for all paths } \pi = i, i_2, \ldots, i_k \text{ with } i_k \text{ on level } n, \\ &\quad \exists j \in \{2, \ldots, k\} \text{ with } U_j > t \text{ or } U_{i_k} < \max\{U_i, U_{i_2}, \ldots, U_{i_{k-1}}\}) \\ \mathcal{R}_n(t) &:= \{i \in \mathcal{T}(n) : U_i \le t \text{ and } i \notin \mathcal{G}(t)) \end{aligned} \tag{4.3.8}$$

Analogously to the previous subsection we define

$$F_n(t) := \mathcal{P}(O \in \mathcal{R}_n(t)).$$

Lemma 4.3.2. *For all $t \in [1/2,1]$, we have*

$$\limsup_{n\to\infty} F_n(t) \ge 1 - \frac{1}{2t}. \tag{4.3.9}$$

Note that the lower bound in (4.3.9) is exactly half the ordinary percolation function on the binary tree. Further, it is equal to the distribution G (4.3.1) on the directed tree that drops out in Aldous model, corresponding to the distribution F (4.3.4) in our case. We have seen a lower bound like this before, recall Lemma 3.4.1. The proof of this lemma uses the same ideas as the proof of Lemma 3.4.1 and is in fact simpler so we do not present it here.

At this point, it is not clear whether $F_n(t)$ converges as n tends to infinity. Even if it does, it is not clear that its distribution converges to the infinite process we described above. Nevertheless, Lemma 4.3.2 suggests that we can rule out the solution (4.3.7). In view of the heuristic arguments, we propose the distribution F (4.3.4). as

a candidate for the directed modified frozen percolation process. From now on, let F denote the distribution function in (4.3.4).

So far, our arguments (apart from Lemmas 4.3.1 and 4.3.2) are non-rigorous. We have inferred what the distribution of the time that a site freezes should be, if the modified frozen percolation process exists on the directed binary tree. However, there are still many questions. Does the modified frozen percolation process exist on the directed binary tree? If so, can we use the directed process to construct the modified frozen percolation process on the Bethe lattice $\mathcal{B}$? On the Bethe lattice, where each site has degree three, we can consider each site as being attached to the roots of three directed subtrees. It is not clear however, whether the directed processes exist simultaneously: does there exists a unique translation invariant law of the states of the sites at time $t, 0 \leq t \leq 1$, for *all* directions simultaneously? Finally, if the modified frozen percolation process can be constructed on $\mathcal{B}$, does it display critical behaviour?

In the following subsection we construct the modified frozen percolation model on $\mathcal{B}$. We follow Aldous [Ald00] but we need to make several non-trivial adjustments.

4.3.3 The frozen percolation process on $\mathcal{B}$

From now on we consider the Bethe lattice $\mathcal{B}$ where each site has degree 3. There is no natural sense of direction as there was on the rooted directed tree, so that we cannot apply the results from the previous subsection directly. We solve this by directing the adjoined edges of a site i outward. The children of i, called i_1,i_2 and i_3 here, form the roots of directed subtrees $\mathcal{T}[i_1]$, $\mathcal{T}[i_2]$ and $\mathcal{T}[i_3]$, where the direction of the edges is inherited from the edges leaving from i, see Figure 4.2.

We define $Y_{i\to i_1}$ to be the time that i_1 freezes in the directed subtree $\mathcal{T}[i_1]$. We use the subscript $i \to i_1$ to indicate that we consider i_1 to be the root of the directed subtree leaving from i_1, where the direction of the edges is inherited from the edge $i \to i_1$. The variables $Y_{i\to i_2}$ and $Y_{i\to i_3}$ should be likewise. In what follows, if we write $Y_{i\to j}$ it is implicit that i is a neighbour of j.

The following lemma shows that we can define the variables $Y_{i\to j}$ for all $i \in \mathcal{B}$ and all $j \in \partial\{i\}$, in a consistent manner. Recall that $\partial\{i\}$ denotes the set of children of i.

Lemma 4.3.3. *There exists a joint law for $((U_i, Y_{i\to j}) : i \in \mathcal{B}, j \in \partial\{i\})$ which is invariant under automorphisms of the tree and such that for each $i \in \mathcal{B}$ and each $j \in \partial\{i\}$ we have*

(i).

$$Y_{j\to i} \text{ has distribution function } F.$$

(ii).

$$Y_{j\to i} = \Phi(\min(Y_{i\to k}, Y_{i\to l}), \max(Y_{i\to k}, Y_{i\to l}), U_i) \text{ a.s.}, \tag{4.3.10}$$

where $\{k, l\} = \partial\{i\}\backslash\{j\}$.

(iii). For each finite connected set $S \subset \mathcal{B}$, the variables $(Y_{i\to j} : i \in S, j \notin S)$ are independent of each other and independent of the collection $(U_i : i \in S)$.

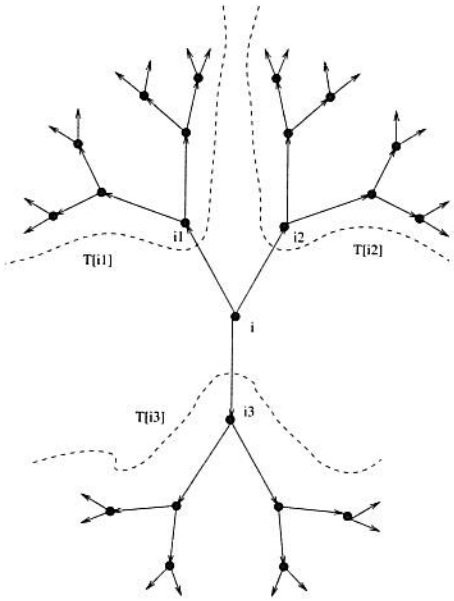

Figure 4.2: The site i, and the three subtrees $\mathcal{T}[i_1]$, $\mathcal{T}[i_2]$ and $\mathcal{T}[i_3]$ leaving from its children.

Proof. Let O be the root and let $n \geq 1$. Define $V_{\leq n}$ to be the set of sites at distance at most n from O. Let V_n be the set of sites at distance n. Take $(U_i : i \in V_{\leq n})$ independent of each other. Then take $(Y_{j\to i} : j \in V_n, i \in V_{n+1})$ independent of the above mentioned U_i's and independent of each other, with each $Y_{j\to i}$ having probability distribution F. We apply equation (4.3.10) successively to define $Y_{i\to j}$ for all sites $i, j \in V_{\leq n}$ and by Lemma 4.3.1 all these have law F. As n increases, the joint laws are consistent and hence we can apply the Kolmogorov consistency theorem to obtain a joint law for $((U_i, Y_{i\to j}) : i \in \mathcal{B}, j \in \partial\{i\})$. Automorphism invariance of this law is straightforward. Part (iii) of the lemma easily follows from the construction. □

Corollary 4.3.4.

$$\mathcal{P}(\exists i, k \in \mathcal{B}, l \in \partial\{k\} \text{ s.t. } U_i = Y_{k\to l}) = 0.$$

Proof. Since the number of sites in $\mathcal{B}$ is countable it is sufficient to show that for fixed i, k and l with $l \in \partial\{k\}$, $U_i \neq Y_{k\to l}$ almost surely. By automorphism invariance, we may assume that $i = O$. We take n so large that the site k is in $V_{\leq n}$. By the construction in the proof of Lemma 4.3.3, $Y_{k\to l}$ must be equal to one of the Y-values in $(Y_{v\to w} : v \in V_n, w \in V_{n+1})$ or infinity. Since n is fixed, we can modify on a null set so that U_i is different from the values $(Y_{v\to w} : v \in V_n, w \in V_{n+1})$ almost surely. It follows that $U_i \neq Y_{k\to l}$, almost surely. □

The process on $\mathcal{B}$ can now be described by the following: We define for each $i \in \mathcal{B}$,

$$\begin{aligned} Z_i &:= \min\{Y_{j\to i} : j \in \partial\{i\}\} \\ &= \min\{Y_{i\to j} : j \in \partial\{i\}, Y_{i\to j} \geq U_i\}, \end{aligned} \tag{4.3.11}$$

where we take $Z_i = \infty$ if the minimum does not exist. The equality follows from (4.3.10). The last expression seems more complicated but we use it later on, since we want to make use of the fact that U_i is independent of $Y_{i\to j}$.

We define $\mathcal{W}(0) := V_{\mathcal{B}}, \mathcal{G}(0) := \emptyset$ and $\mathcal{R}(0) := \emptyset$. These evolve in time as follows: for $t \in [0,1]$,

$$\begin{aligned} \mathcal{W}(t) &:= \{i \in \mathcal{B} : U_i > t\}, \\ \mathcal{G}(t) &:= \{i \in \mathcal{B} : U_i \leq t, Z_i > t\}, \\ \mathcal{R}(t) &:= \{i \in \mathcal{B} : Z_i \leq t\}. \end{aligned} \tag{4.3.12}$$

One should think of the set $\mathcal{W}(t)$ as the set of white (at rest) sites; of $\mathcal{G}(t)$ as the set of green (activated but in a finite cluster) sites and $\mathcal{R}(t)$ as the set of sites that are red (frozen infinite clusters) at time t. We use this terminology from now on.

From the definition immediately follows that for all sites i, with $Z_i < \infty$,

$$\begin{aligned} &i \in \mathcal{W}(t) \text{ for } t \in [0, U_i), \\ &i \in \mathcal{G}(t) \text{ for } t \in [U_i, Z_i), \\ &i \in \mathcal{R}(t) \text{ for } t \in [Z_i, 1]. \end{aligned}$$

If $Z_i = \infty$, clearly

$$\begin{aligned} &i \in \mathcal{W}(t) \text{ for } t \in [0, U_i), \\ &i \in \mathcal{G}(t) \text{ for } t \in [U_i, 1]. \end{aligned}$$

This description gives us the following:

Lemma 4.3.5. *For every vertex v and for every pair of neighbouring vertices $v \sim w$,*

(i). The probability that a vertex is eventually not frozen equals

$$\mathcal{P}(v \in \mathcal{G}(1)) = \frac{3}{2}\ln(2)^2 - \frac{1}{2} \approx 0.22,$$

so that the probability that a vertex is eventually frozen equals

$$\mathcal{P}(v \in \mathcal{R}(1)) = 1 - \mathcal{P}(v \in \mathcal{G}(1)) \approx 0.77.$$

(ii). The probability that both v and w are eventually frozen but in different frozen infinite clusters equals

$$\mathcal{P}(v, w \in \mathcal{R}(1), Z_v \neq Z_w) = 3\ln(2) - 2 \approx 0.079.$$

Proof. The lemma follows easily from the definition of Z_i's (4.3.11) and $Y_{\cdot\to\cdot}$'s (4.3.10) and the distribution F (4.3.4). For example, part (i):

$$\begin{aligned}\mathcal{P}(v \in \mathcal{G}(1)) &= \mathcal{P}(Z_v = \infty) \\ &= \mathcal{P}(\forall j \in \partial\{v\} : Y_{v\to j} < U_i \text{ or } Y_{v\to j} = \infty) \\ &= \int_{1/2}^{1} [\ln(2s) + 1 - \ln(2)]^3 ds + \frac{1}{2}(1 - \ln(2))^3 = \frac{3}{2}\ln(2)^2 - \frac{1}{2}.\end{aligned}$$

The other statement is proved similarly. □

In contrast to the finite tree where it was clear that the sets $\mathcal{W}(t), \mathcal{G}(t)$ and $\mathcal{R}(t)$ fit their intuitive description, we need to prove this in some detail here. So far, we have constructed *a* process, but it is not clear that it acts like the informally described modified frozen percolation process. The following lemmas show that only infinite clusters join $\mathcal{R}(\cdot)$ and that there are no infinite green clusters in $\mathcal{G}(\cdot)$, almost surely.

Lemma 4.3.6. *Let $S_s(i)$ denote the cluster of i in $\mathcal{G}(s)$, considered as a set of sites. Almost surely, $\forall i \in \mathcal{B}$ with $Z_i < \infty$,*

(i). $S_{Z_i^-}(i) \subseteq \mathcal{R}(Z_i)$

(ii). $|S_{Z_i^-}(i)| = \infty$

Proof. Since the number of sites in $\mathcal{B}$ is countable, it is sufficient to show that a.s. (i) and (ii) hold for some fixed i. Choose an arbitrary site i, with $Z_i < \infty$. By definition of Z_i (4.3.11), there exists a site $j \in \partial\{i\}$ such that $Y_{i\to j} \geq U_i$ and $Z_i = Y_{i\to j}$. Using (4.3.10) repeatedly there exists an infinite self-avoiding path $i =: j_0, j =: j_1, j_2, \ldots$ such that

$$Z_i = Y_{i\to j} = Y_{j\to j_2} = Y_{j_2\to j_3} = \cdots, \tag{4.3.13}$$

and

$$U_{j_k} < Z_i \text{ for } k \geq 0, \text{ almost surely.} \tag{4.3.14}$$

Note that we have used Corollary 4.3.4 to obtain a strict inequality in (4.3.14). In the directed subtree where the direction is given by $i \to j$, each site j_k has only two children. One of them is j_{k+1}; let v_{k+1} be the other child, for $k > 0$. By (4.3.10) we have for all $k > 0$,

$$Y_{j_k\to v_{k+1}} > Z_i \text{ or } Y_{j_k\to v_{k+1}} < U_{j_k}. \tag{4.3.15}$$

We need to show that the path $j_1, j_2, \ldots$ constructed above has $Z_{j_k} = Y_{j_k\to j_{k+1}} (= Z_i)$ for all $k > 0$. By definition

$$Z_{j_k} = \min\{Y_{j_k\to w} : w \in \{j_{k+1}, j_{k-1}, v_{k+1}\}, Y_{j_k\to w} \geq U_{j_k}\}.$$

By (4.3.13) and (4.3.15), the minimum cannot be achieved taking w equal to v_{k+1}. So suppose that $U_{j_k} \leq Y_{j_k \to j_{k-1}} < Z_i$. Applying (4.3.10) and (4.3.15) repeatedly we obtain

$$Y_{j_k \to j_{k-1}} = Y_{j_{k-1} \to j_{k-2}} = \cdots = Y_{j \to i} = Y_{i \to w} < Z_i, \tag{4.3.16}$$

for some $w \in \partial\{i\}, w \neq j$. From (4.3.10) also follows that $Y_{i \to w} \geq U_i$, which implies $Z_i \leq Y_{i \to w} < Z_i$, a contradiction. This shows that $Z_i = Z_j = Z_{j_2} = \cdots$ which together with (4.3.14) shows that an infinite path contained in $S_{Z_i^-}(i)$ joins $\mathcal{R}(Z_i)$ at time Z_i. This proves part (ii) of the lemma.

Now suppose that we have a site v_k in $S_{Z_i^-}(i)$ that is adjacent to the infinite path $j_0, j_1, \ldots$ at the site j_{k-1} for some $k > 0$. Because $v_k \in S_{Z_i^-}(i)$, we have that $U_{v_k} < Z_i$ and further that $Z_{v_k} \geq Z_i$. We need to show that equality holds here. But this follows easily from the above, since $Y_{v_k \to j_{k-1}} = Y_{j_{k-1} \to j_k} = Z_{j_{k-1}} = Z_i$ by the construction of the infinite path above, and hence $Y_{v_k \to j_{k-1}} \geq U_{v_k}$ and

$$Z_i \leq Z_{v_k} \leq Y_{v_k \to j_{k-1}} = Z_{j_{k-1}} = Z_i.$$

We repeat this argument 'working to the outside of $S_{Z_i^-}(i)$' so that we obtain that $S_{Z_i^-}(i) \subseteq \mathcal{R}(Z_i)$, which proves part (i) of the lemma. □

Lemma 4.3.7. *Almost surely, there is no infinite component in $\mathcal{G}(t)$ for $t \in [0,1)$.*

Proof. Suppose there exist $t \in [0,1)$ and $i \in \mathcal{B}$ such that

$$S_t(i) \subseteq \mathcal{G}(t) \text{ and } |S_t(i)| = \infty.$$

Then $Z_i > t$, by definition. We will show that almost surely, all $j \in S_t(i)$ have $Z_j = Z_i$, so that if an infinite green cluster exists at some time-point, it will exist in a time-interval of positive length. Suppose that there is a $k \in S_t(i)$ with $Z_k \in (t, Z_i)$. This means that $S_t(i) = S_t(k) \subseteq S_{Z_k^-}(k) \subseteq \mathcal{R}(Z_k)$ almost surely, by Lemma 4.3.6. So $i \in \mathcal{R}(Z_k)$ and this contradicts the fact that $Z_i > Z_k$. We conclude that almost surely, all $j \in S_t(i)$, $Z_j \geq Z_i$ so that $S_t(i) \subseteq \mathcal{G}(Z_i^-)$.

From the above argument follows that if $\mathcal{G}(t)$ contains an infinite component, then there exist two rationals $t_1 < t_2$ such that $\mathcal{G}(s)$ contains an infinite component for all $s \in [t_1, t_2]$. Since the number of such pairs t_1, t_2 is countable, it is sufficient to show that for fixed $t_1 < t_2$,

$$\mathcal{P}(\exists \text{ an infinite component which, } \forall s \in [t_1, t_2], \text{ is in } \mathcal{G}(s)) = 0. \tag{4.3.17}$$

For $t_1 \in [0, 1/2]$ the above easily follows from ordinary percolation results.

Now suppose $t_1 \in [1/2, 1)$ and let $t_2 \in (t_1, 1]$. Then fix a site i_0 and a path $\pi = i_0, i_1, \ldots, i_{n-1}$ of length n. Add direction to the subtree leaving from i_1 consistent with $i_0 \to i_1$. Let for $j \in [2, n-1]$, v_j denote the child of i_{j-1} that is not on π. Consider

$$\begin{aligned}
&\mathcal{P}(\forall s \in [t_1, t_2]\ \forall i \in \pi : i \in \mathcal{G}(s)) \\
&\quad = \mathcal{P}(\forall j \in [0, n-1] : U_{i_j} \leq t_1 \text{ and } Z_{i_j} \geq t_2) \\
&\quad \leq \mathcal{P}(\forall j \in [1, n-2] : U_{i_j} \leq t_1 \text{ and either } Y_{i_j \to v_{j+1}} \geq t_2 \text{ or } Y_{i_j \to v_{j+1}} < U_{i_j}) \\
&\quad = \left[\mathcal{P}(U_{i_1} \leq t_1 \text{ and either } Y_{i_1 \to v_2} \geq t_2 \text{ or } Y_{i_1 \to v_2} < U_{i_1})\right]^{n-2}.
\end{aligned} \tag{4.3.18}$$

The inequality follows from (4.3.11) and the equality from automorphism invariance and the independence property, see Lemma 4.3.3, part (iii). Now using the knowledge about the distribution of $Y_{i_1 \to v_2}$ (recall that it is equal to F, (4.3.4)) the last term in (4.3.18) equals

$$\Big[\int_{1/2}^{t_1} \ln(2s)ds + t_1[1 - \ln(2t_2)]\Big]^{n-2} = \big[1/2 + t_1 \ln(t_1/t_2)\big]^{n-2}$$

Summing over all $3 \cdot 2^n$ possible paths leaving from i_0 we obtain

$$\begin{aligned} &\mathcal{P}(\exists \text{ path } \pi \text{ of length } n \text{ leaving from } i_0 \text{ such that } \forall j \text{ on } \pi : U_j \le t_1 \text{ and } Z_j \ge t_2) \\ &\le \quad 12\big(1 + 2t_1 \ln(t_1/t_2)\big)^{n-2}. \end{aligned} \tag{4.3.19}$$

Note that $t_2 > t_1$ so that $\ln(t_1/t_2) < 0$. Applying (4.3.19), we get

$$\begin{aligned} &\mathcal{P}(\exists \text{ infinite path } \pi \text{ from } i_0 \text{ which, } \forall s \in [t_1, t_2], \text{ is contained in } \mathcal{G}(s)) \\ &\le \lim_{n\to\infty} 12\big(1 + 2t_1 \ln(t_1/t_2)\big)^{n-2} = 0. \end{aligned}$$

We sum over all i_0 to show (4.3.17) for $t_1 \in (1/2, 1)$. This finishes the proof of the lemma. □

Summarising, we have constructed a process on the binary tree with the following properties. At time 0, all sites are at rest. Each site i becomes activated at time U_i, where U_i is a uniform $U[0,1]$ random variable (chosen independently for each i). As soon as a site is in an infinite activated cluster at some time $t < 1$, it freezes. In the following subsection we will see that the modified frozen percolation process displays some features of critical behaviour.

4.3.4 Critical behaviour

Aldous shows for his dynamics that finite non-empty clusters (which are the analogue of green clusters in our terminology) are distributed as ordinary critical percolation clusters conditional on being non-empty. Remarkably this holds for *all* $t \ge 1/2$. His frozen percolation process behaves like a critical system at all times (after the critical time), so that it is a nice example of self-organised criticality.

In our case, the distribution of green clusters is different: suppose that S is an arbitrary finite connected set of sites such that $|S| > 1$; let ∂S denote its boundary. If v is a neighbour of w we write $v \sim w$.

Lemma 4.3.8. *$S \subseteq \mathcal{G}(t)$ if and only if*

(a). $\forall w \in S : U_w \le t$ *and*

(b). $\forall v \notin S$ *such that* $v \sim s$ *for some* $s \in S$, *we have either* $Y_{s\to v} < U_s$ *or* $Y_{s\to v} > t$.

Proof. One of the implications is trivial. Suppose $S \subseteq \mathcal{G}(t)$. Then (a) and (b) hold by definition of $\mathcal{G}(t)$ (4.3.12). The reversed implication is only slightly more difficult. Suppose that (a) and (b) hold. Then, 'working from the outside in', we can recursively find all $Y_{w_1 \to w_2}$ for all $w_1 \in S$, $w_2 \in \partial\{w_1\}$ using (4.3.10). It is easy to see that all Y-values thus obtained are in $(t, 1]$ or equal to ∞, so that $S \subseteq \mathcal{G}(t)$. □

If we want to compute the probability that $S \subseteq \mathcal{G}(t)$, i.e. by Lemma 4.3.8 the joint probability of (a) and (b) above, the actual geometry of the set S may be of importance. Not surprisingly, for $t \leq 1/2$ we find

$$\mathcal{P}(S \subseteq \mathcal{G}(t)) = t^{|S|}.$$

Now suppose $t > 1/2$ and consider (a) and (b). Every $w \in S$ that has two neighbours v_1, v_2 in ∂S, contributes a factor

$$\int_{1/2}^{t} \ln(2s)^2 ds + 2(1 - \ln(2t)) \int_{1/2}^{t} \ln(2s) ds + t(1 - \ln(2t))^2 = t - \ln(2t).$$

Note that the contribution of such w to the joint probability of (a) and (b) only depends on $U_w, Y_{w \to v_1}$ and $Y_{w \to v_2}$, by part (iii) of Lemma 4.3.3. Intuitively, this factor can be explained by considering w to be activated, but not frozen in the directed subtree where w is the root, and v_1 and v_2 are its children.

Further, every w that has only one neighbour v in ∂S contributes a factor

$$\int_{1/2}^{t} \ln(2s) ds + t(1 - \ln(2t)) = 1/2.$$

The contribution of such w to the joint probability of (a) and (b) only depends on U_w and $Y_{w \to v}$. The occurrence of the factor '$1/2$' is quite surprising. We have no intuitive explanation, but it is the reason for the critical behaviour observed (see for example equation (4.3.21) and Lemma 4.3.9) later on. Finally, every $w \in S$ that has no neighbours in ∂S contributes a factor t. Define

$$\begin{aligned} n(0) &:= \#\{w \in S : \#\{v \in \partial S, v \sim w\} = 0\}, \\ n(1) &:= \#\{w \in S : \#\{v \in \partial S, v \sim w\} = 1\}, \\ n(2) &:= \#\{w \in S : \#\{v \in \partial S, v \sim w\} = 2\}. \end{aligned}$$

Note that for any connected set S we have $n(0) + 2 = n(2)$. For any connected set S such that $|S| > 1$,

$$\mathcal{P}(S \subseteq \mathcal{G}(t)) = (t - \ln(2t))^{n(2)} \left(\frac{1}{2}\right)^{n(1)} t^{n(0)}. \tag{4.3.20}$$

From this we can see that the distribution of green clusters is not the same as distribution of ordinary percolation clusters (for any parameter value), because not

only the size but also the geometry of a cluster plays a role. Nevertheless, by (4.3.20) we have for all $t \geq 1/2$, and all sites v and w

$$\mathcal{P}(\exists \text{ green path from } v \text{ to } w \text{ at time } t) = \left(\frac{1-\ln(2)}{2}\right)^2 \left(\frac{1}{2}\right)^{\pi(v,w)}, \qquad (4.3.21)$$

where $\pi(v,w)$ is the number of sites in the path from v to w. This is the most remarkable similarity to ordinary critical percolation; there the probability that the path from v to w is contained in an open cluster equals $(1/2)^{\pi(v,w)}$.

Arbitrary green clusters are also critical in some sense:

Lemma 4.3.9. *Let E denote the expectation with respect to the law of the frozen percolation process. Fix a site O and let $\mathcal{G}_O(t)$ denote the green cluster of O at time t. We consider V_n, the set of sites at distance n from O. Then there exist positive constants A_1 and A_2 such that for all n,*

$$A_1 \leq E|\mathcal{G}_O(t) \cap V_n| \leq A_2. \qquad (4.3.22)$$

Proof.

$$\begin{aligned} E|\mathcal{G}_O(t) \cap V_n| &= \sum_{v \in V_n} \mathcal{P}(v \in \mathcal{G}_O(t)) \\ &= \sum_{v \in V_n} \mathcal{P}(\text{the unique path from } O \text{ to } v \text{ is in } \mathcal{G}(t)) \\ &= 3 \cdot 2^{n-1} \left(\frac{1}{2}\right)^{n-1} (t - \ln(2t))^2 = 3(t - \ln(2t))^2. \end{aligned} \qquad (4.3.23)$$

The last factor is bounded from above and below for $t \in [1/2, 1]$, which proves the lemma. □

Similar behaviour is observed for ordinary critical percolation (see [Gri99], Section 10.1).

4.3.5 Dependencies

Any two sites v and w in the tree, are connected by a unique path. The dependence between the colours of sites v and w at time t, decays exponentially in $\pi(v,w)$, the number of sites between v and w. Although we did not mention this in Section 4.2, this is also true on the line $\mathbb{Z}$.

Lemma 4.3.10. *There exists a constant $A \in (0,1)$ such that for all sites v and w and all $t \in [0,1]$, and for any $\mathcal{C}_k(t) \in \{\mathcal{W}(t), \mathcal{G}(t), \mathcal{R}(t)\}, k = 1, 2,$*

$$\left|\mathcal{P}(v \in \mathcal{C}_1(t), w \in \mathcal{C}_2(t)) - \mathcal{P}(v \in \mathcal{C}_1(t))\mathcal{P}(w \in \mathcal{C}_2(t))\right| \leq 5A^{\lfloor \frac{\pi(v,w)}{12} \rfloor}. \qquad (4.3.24)$$

Proof. For $t \leq 1/2$, sites can only be white or green and the transition from white to green happens independent of the other sites. In fact, the left hand side of (4.3.24) equals zero. So suppose $t > 1/2$ and let $v = i_1, i_2, \ldots, i_n = w$ denote the unique path between v and w. We may assume that $n \geq 12$, otherwise there is nothing to prove. Define the following events:

$$\begin{aligned} E_1 &:= \{\exists j \in (3n/4, n) : \min\{U_{i_j}, U_{i_{j+2}}\} > Z_{i_{j+1}}\} \\ E_2 &:= \{\exists k \in (0, n/4) : \min\{U_{i_k}, U_{i_{k-2}}\} > Z_{i_{k-1}}\} \end{aligned}$$

If $\min\{U_{i_j}, U_{i_{j+2}}\} > Z_{i_{j+1}}$ for some j, it follows that $Z_{i_{j+1}} = Y_{i_{j+1} \to a_{j+1}}$, where a_{j+1} is the child of i_{j+1} not lying on the path from v to w. So if E_1 occurs for some j, it only depends on $U_{i_j}, U_{i_{j+1}}, U_{i_{j+2}}$ and the U-variables in the subtree leaving from a_{j+1}. Fix $j \in (3n/4, n)$. Taking disjoint sets of three consecutive sites on the path from v to w, we arrive at

$$\mathcal{P}(E_1^c) \leq \left(1 - \mathcal{P}(Y_{i_{j+1} \to a_{j+1}} < \min\{U_{i_j}, U_{i_{j+2}}\})\right)^{\lfloor \frac{n}{12} \rfloor}, \tag{4.3.25}$$

by automorphism invariance. A similar reasoning holds for E_2. Further, on $E_1 \cap E_2$, the colours of v and w are independent: suppose we cut the tree in two parts by removing any edge (i_l, i_{l+1}) with $n/4 < l < 3n/4 - 1$ on the path from v to w. The occurrence of E_1 and E_2 ensure that the U-variables on the part containing v (w) determine the colour of v (w, respectively) uniquely. If E_1 occurs for some j, at least one of the sites i_j and i_{j+2} is at rest at time t, or site i_{j+1} is frozen without i_j or i_{j+2} at time t. In both cases, the configuration on i_j, i_{j+1}, i_{j+2} prevents anything happening on the part of the tree containing v to influence the colour of w. We compute

$$\begin{aligned} &\mathcal{P}(v \in \mathcal{C}_1(t), w \in \mathcal{C}_2(t)) \\ &\quad = \mathcal{P}(v \in \mathcal{C}_1(t), w \in \mathcal{C}_2(t), E_1, E_2) + \mathcal{P}(v \in \mathcal{C}_1(t), w \in \mathcal{C}_2(t), E_1^c \cup E_2^c) \\ &\quad = \mathcal{P}(v \in \mathcal{C}_1(t) \cap E_2)\mathcal{P}(w \in \mathcal{C}_2(t) \cap E_1) + \mathcal{P}(v \in \mathcal{C}_1(t), w \in \mathcal{C}_2(t), E_1^c \cup E_2^c) \\ &\quad = \left[\mathcal{P}(v \in \mathcal{C}_1(t)) - \mathcal{P}(v \in \mathcal{C}_1(t) \cap E_2^c)\right]\left[\mathcal{P}(w \in \mathcal{C}_2(t)) - \mathcal{P}(w \in \mathcal{C}_2(t) \cap E_1^c)\right] \\ &\qquad + \mathcal{P}(v \in \mathcal{C}_1(t), w \in \mathcal{C}_2(t), E_1^c \cup E_2^c). \end{aligned} \tag{4.3.26}$$

From (4.3.26) we obtain

$$\left|\mathcal{P}(v \in \mathcal{C}_1(t), w \in \mathcal{C}_2(t)) - \mathcal{P}(v \in \mathcal{C}_1(t))\mathcal{P}(w \in \mathcal{C}_2(t))\right| \leq 5\mathcal{P}(E_1^c). \tag{4.3.27}$$

The only thing we need to show now, is that the probability on the left side of (4.3.25) is bounded away from 1. We use that by (4.3.4),

$$\mathcal{P}(Y_{i_{j+1} \to a_{j+1}} < \min\{U_{i_j}, U_{i_{j+2}}\}) = 2\int_{1/2}^{1} (1-s)\ln(2s)ds > 0.$$

This allows us to define the required constant A. □

Remark 4.3.11. One may wonder why we took the site version of the model, whereas Aldous' model is stated in terms of edges. The first reason is that we studied the frozen percolation model on $\mathbb{Z}$ using the site version and we wanted to keep the models consistent. The second reason is that we wanted to see whether a perturbation in the dynamics would cause the critical behaviour to disappear. The third (and in practice maybe the most important) reason is that the edge version introduces dependencies, which make it harder (if not impossible) to construct the model.

Up to a certain point, we could just as well have taken the edge version of the process. Then, we replace each edge e of the tree by two directed edges $\vec{e}$ and $\overleftarrow{e}$. We attach an independent $U[0,1]$ variables $U_e = U_{\vec{e}} = U_{\overleftarrow{e}}$ to each edge. Intuitively, we need variables $Y_{\vec{e}}$ and $Y_{\overleftarrow{e}}$ that represents the time a directed edge freezes in the directed subtree, where the direction is inherited from $\vec{e}$ or $\overleftarrow{e}$. We can define a relation similar to (4.3.10) which $Y_{\vec{e}}$ and $Y_{\overleftarrow{e}}$ have to satisfy (recall the definition of Φ in (4.3.2)):

$$Y_{\vec{e}} = \Phi(\min(Y_{\vec{e_1}}, Y_{\vec{e_2}}), \max(Y_{\vec{e_1}}, Y_{\vec{e_2}}), U_e), \tag{4.3.28}$$

where $\vec{e_1}$ and $\vec{e_2}$ are the children of the directed edge $\vec{e}$. We follow the heuristic arguments from Subsection 4.3.2, using (4.3.28). The distribution $F(t) = \ln(2t)$ for $t \geq 1/2$ that drops out (compare Lemma 4.3.1) for the edge version is the same as for the site version. The analogue of Lemma 4.3.3 can be proved similarly. However, problems arise because the edge version of frozen percolation carries more dependencies that the site version. Particularly, the analogue of Z_i (see equation (4.3.11)) for the edge version of the model is not easily defined.

Open Problem 4.3.12. Is it possible to construct the edge version of the frozen percolation process?

Chapter 5

Random sampling for the monomer-dimer model on a lattice

5.1 Introduction

The monomer-dimer model, described below, originates from statistical physics where it has been used to study the absorption of oxygen molecules on a surface and the properties of a binary mixture. See Heilmann and Lieb [HL72] for further background and references. More recently, the model has also drawn much attention in the fields of operations research, combinatorics and graph theory [JS96], [KK97].

Consider a finite, undirected graph $G = (V_G, E_G)$, where V_G is the set of vertices of G and E_G is the set of edges. A *matching* on G is a subset $M \subset E_G$ such that no two edges in M have a common endpoint. Suppose that the model parameter λ is positive. Now assign to each matching a probability proportional to $\lambda^{|M|}$, where $|M|$ is the size of M.

Alternatively, define the state space $\Omega = \{0,1\}^{E_G}$. Elements of Ω, called configurations on E_G, are typically denoted by $\omega = (\omega_e : e \in E_G)$. The monomer-dimer distribution with parameter λ for G is then defined as

$$\mu(\omega) = \frac{\lambda^{|\omega|} I(\omega \text{ is allowable})}{Z(\lambda)},$$

where 'ω is allowable' means that the set $\{e : \omega_e = 1\}$ is a matching and $Z(\lambda)$ is the normalizing factor, or partition function. Consider a '1' as presence and a '0' as absence of an edge. It should be clear that the two descriptions (one with state space the set of all matchings, the other with state space $\{0,1\}^{E_G}$) are equivalent, and both descriptions will be used in this chapter.

To continue, we need more terminology and notation. If two vertices i and j are adjacent, we write $i \sim j$. Note that also $i \sim i$. The degree $deg(v)$ of a vertex v is defined as the number of edges that have v as an endpoint. If two edges e_1 and e_2 share a common endpoint, we write $e_1 \sim e_2$. Let in the rest of this section, Δ be a subset of E_G. We denote the set $\{0,1\}^\Delta$ by Ω_Δ. Similarly, if $\omega \in \Omega$, then ω_Δ denotes the restriction of ω to Δ, i.e. the element $(\omega_e : e \in \Delta)$ of Ω_Δ. If $\omega, \omega' \in \Omega_\Delta$ we call e an *edge of disagreement* with respect to (ω, ω') if $\omega_e \neq \omega'_e$, and we denote the set of all such edges by $D_\Delta(\omega, \omega')$. The boundary $\partial\Delta$ of Δ consists of all elements $e \in E_G \backslash \Delta$ such that $e \sim e'$ for some $e' \in \Delta$.

Let $\alpha \in \Omega_{\partial\Delta}$, so α is a configuration on the boundary of Δ. The monomer-dimer distribution for Δ *with boundary condition* α is defined as follows:

$$\mu_\Delta^\alpha(\omega) = \frac{\lambda^{|\omega|} I(\omega \text{ is allowable w.r.t. } \alpha)}{Z_\Delta(\lambda)},$$

where 'allowable with respect to α' means that the set $\{e \in \Delta : \omega_e = 1\} \cup \{e \in \partial\Delta : \alpha_e = 1\}$ is a matching.

It is easy to check that the monomer-dimer model satisfies the following spatial Markov-property: let σ denote a random configuration on E_G and let $\Delta \subset E_G$. Then the conditional distribution of σ_Δ, given $\sigma_{E_G \backslash \Delta}$ equals $\mu_\Delta^{\sigma_{\partial\Delta}}$, and hence depends only on $\sigma_{\partial\Delta}$.

A paper by van den Berg [Ber99] shows that the monomer-dimer model on a lattice has certain very strong spatial mixing properties. In Section 5.3 we explicitly show how this can be used to improve, for 'nice' subgraphs of a lattice, earlier results by Jerrum and Sinclair [JS96] concerning the generation of (approximate) random samples.

5.2 Preliminaries

In this section we give the background needed in Section 5.3. First we present some general and quite well-known results on coupling and variational distance. Then we will state the earlier mentioned result by Jerrum and Sinclair. Finally we present and prove a result which is very similar to, but more convenient for our purpose than, a result by van den Berg [Ber99]. Throughout this section Ω denotes an arbitrary finite set.

5.2.1 Coupling and variational distance

Suppose we have two probability distributions μ_1 and μ_2 on Ω. Their *variational distance* $d_V(\mu_1, \mu_2)$ is defined by:

$$d_V(\mu_1, \mu_2) = \frac{1}{2} \sum_{\omega \in \Omega} |\mu_1(\omega) - \mu_2(\omega)|. \tag{5.2.1}$$

Another but equivalent definition of variational distance is the following:

$$d_V(\mu_1, \mu_2) = \max_{A \subset \Omega} |\mu_1(A) - \mu_2(A)|.$$

This equivalence is quite easy to check.

Suppose we have two probability distributions μ_1 and μ_2 on Ω. A *coupling* $\mathcal{P}$ of μ_1 and μ_2 is a distribution on $\Omega \times \Omega$ which has the following properties:

$$\sum_{\omega_1 \in \Omega} \mathcal{P}(\omega_1, \omega_2) = \mu_2(\omega_2) \qquad \text{for all } \omega_2 \in \Omega,$$

and

$$\sum_{\omega_2 \in \Omega} \mathcal{P}(\omega_1, \omega_2) = \mu_1(\omega_1) \qquad \text{for all } \omega_1 \in \Omega,$$

i.e. the marginal distributions of $\mathcal{P}$ are μ_1 and μ_2. Similarly, one can define couplings of more than two probability distributions. A trivial example of a coupling is the product coupling $\mu_1 \times \mu_2$.

Define the event '*unequal*' as the set $\{(\omega_1, \omega_2) \in \Omega \times \Omega : \omega_1 \neq \omega_2\}$. Likewise we define the complementary event '*equal*' as $\{(\omega_1, \omega_2) \in \Omega \times \Omega : \omega_1 = \omega_2\}$. The following results, Propositions 5.2.1 and 5.2.4 and Lemma 5.2.3 are quite standard and not difficult to prove.

Proposition 5.2.1. *Let μ_1, μ_2 and μ_3 be probability distributions on Ω and let $\mathcal{P}_{1,2}$ and $\mathcal{P}_{2,3}$ be couplings of μ_1 and μ_2, and of μ_2 and μ_3, respectively. Then there exists a coupling $\mathcal{P}_{1,3}$ of μ_1 and μ_3 with the following property:*

$$\mathcal{P}_{1,3}(unequal) \leq \mathcal{P}_{1,2}(unequal) + \mathcal{P}_{2,3}(unequal). \tag{5.2.2}$$

Proof. Suppose we first sample an element x_0 from the distribution $\mu_1(\cdot)$, then sample an element y_0 from the conditional distribution $\mathcal{P}_{1,2}(x_0, \cdot)/\mu_1(x_0)$ and finally an element z_0 from the conditional distribution $\mathcal{P}_{2,3}(y_0, \cdot)/\mu_2(y_0)$. Intuitively, the distribution of the pair (x_0, z_0) is a candidate for the coupling we are looking for. More formally, define the following '3-coupling' $\mathcal{P}_{1,2,3}$ of μ_1, μ_2 and μ_3.

$$\mathcal{P}_{1,2,3}(x, y, z) = \begin{cases} \frac{\mathcal{P}_{1,2}(x,y) \cdot \mathcal{P}_{2,3}(y,z)}{\mu_2(y)} & \text{if } \mu_2(y) > 0 \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to check that $\sum_{y,z} \mathcal{P}_{1,2,3}(x, y, z) = \mu_1(x)$, $\sum_{x,z} \mathcal{P}_{1,2,3}(x, y, z) = \mu_2(y)$ and $\sum_{x,y} \mathcal{P}_{1,2,3}(x, y, z) = \mu_3(z)$, so that $\mathcal{P}_{1,2,3}$ is indeed a coupling of μ_1, μ_2 and μ_3. From this immediately follows that $\mathcal{P}_{1,3}$ defined by

$$\mathcal{P}_{1,3}(x, z) = \sum_y \mathcal{P}_{1,2,3}(x, y, z),$$

is a coupling of μ_1 and μ_3.

We also have $\sum_x \mathcal{P}_{1,2,3}(x,y,z) = \mathcal{P}_{2,3}(y,z)$ and $\sum_z \mathcal{P}_{1,2,3}(x,y,z) = \mathcal{P}_{1,2}(x,y)$. Property (5.2.2) easily follows from the above, using

$$\{(x,y,z) : x \neq z\} \subset \big\{\{(x,y,z) : x \neq y\} \cup \{(x,y,z) : y \neq z\}\big\}. \tag{5.2.3}$$

□

Remark 5.2.2. Note that for *any* '3-coupling' $\mu_{1,2,3}$ of μ_1, μ_2 and μ_3, the two-dimensional marginals satisfy

$$\sum_y \mu_{1,2,3}(x \neq z) \leq \sum_z \mu_{1,2,3}(x \neq y) + \sum_x \mu_{1,2,3}(y \neq z),$$

by (5.2.3).

We proceed with a lemma that states the basic properties of variational distance.

Lemma 5.2.3. *Let μ, μ' and ν be probability distributions on Ω. We have:*

(i). $d_V(\mu,\nu) \geq 0$

(ii). $d_V(\mu,\nu) = d_V(\nu,\mu)$

(iii). $d_V(\mu,\nu) \leq d_V(\mu,\mu') + d_V(\mu',\nu)$

(iv). $d_V(\gamma\mu + (1-\gamma)\mu',\nu) \leq \gamma d_V(\mu,\nu) + (1-\gamma)d_V(\mu',\nu)$ *for all* $\gamma \in [0,1]$

Proof. The first two properties follow directly from the definition of variational distance, the third property follows from applying the triangle inequality on each term in the sum in (5.2.1). The fourth property can be seen easily by writing out the definition. □

The following proposition relates the notions of variational distance and couplings. Recall the notions '*equal*' and '*unequal*' defined earlier.

Proposition 5.2.4. *For all probability distributions μ and ν on Ω:*

$$d_V(\mu,\nu) = \min_{\mathcal{P}} \mathcal{P}(unequal),$$

where the minimum is taken over all couplings $\mathcal{P}$ of μ and ν.

Proof. Let $\mathcal{P}$ be a coupling of μ and ν. It is a useful (for those not yet familiar with these notions) and fairly easy exercise to show that $d_V(\mu,\nu) \leq \mathcal{P}(unequal)$. To complete the proof, we need to construct a coupling $\widetilde{\mathcal{P}}$ for which

$$d_V(\mu,\nu) = \widetilde{\mathcal{P}}(unequal). \tag{5.2.4}$$

To achieve this, define

$$\widetilde{\mathcal{P}}(x,x) = \min\{\mu(x),\nu(x)\}. \tag{5.2.5}$$

With this definition, the probability $\widetilde{\mathcal{P}}(equal)$ is fixed and hence its complement $\widetilde{\mathcal{P}}(unequal)$ is fixed too. Any coupling satisfying (5.2.5) also satisfies (5.2.4). We still have to define $\widetilde{\mathcal{P}}(x,y)$ for $x \neq y$ in such way that $\widetilde{\mathcal{P}}$ is truly a coupling. In general, there are several ways to do this. For example, we can define

$$\widetilde{\mathcal{P}}(x,y) = \frac{[\mu(x) - \min\{\mu(x),\nu(x)\}]\,[\nu(y) - \min\{\mu(y),\nu(y)\}]}{\widetilde{\mathcal{P}}(unequal)},$$

for all $x \neq y$. Straightforward calculation shows that $\widetilde{\mathcal{P}}$ is indeed a coupling. □

A coupling that reaches the minimum in Proposition 5.2.4 is called *optimal*. For an extensive treatment of coupling methods, see [Tho00] or [Lin92].

5.2.2 Mixing times and the Jerrum-Sinclair result

Suppose we have an ergodic Markov chain on Ω. Let π be the stationary distribution of this chain and let $x \in \Omega$. Let $\mu^{x,t}$ be the distribution of the Markov chain at time t when it has started in initial state x. Let $\varepsilon > 0$. Define the *mixing time with respect to initial state* x of the Markov chain as follows:

$$\tau_x(\varepsilon) = \min_t\{d_V(\mu^{x,t'},\pi) \leq \varepsilon \text{ for all } t' \geq t\}. \tag{5.2.6}$$

The *(total) mixing time* of the Markov chain is then defined by

$$\tau(\varepsilon) = \max_{x\in\Omega} \tau_x(\varepsilon). \tag{5.2.7}$$

Jerrum and Sinclair [JS96] have studied the mixing time of a suitable Markov chain for the monomer-dimer model. More precisely, they have proved the following: let $G = (V_G, E_G)$ be a finite graph and let $\Omega = \{\text{all matchings on } G\}$. Consider the monomer-dimer distribution with parameter $\lambda > 0$ on Ω. Denote this distribution by π_λ. To sample from this distribution, Jerrum and Sinclair apply the Metropolis algorithm implemented to the monomer-dimer model: they study a specific Markov chain $mc(\lambda)$ that has stationary distribution π_λ. A transition $M \to M'$ in the Markov chain $mc(\lambda)$ is described as follows:

(i). With probability $\frac{1}{2}$ let $M' = M$; otherwise

(ii). Choose uniformly at random an edge $e = (u,v) \in E_G$.
Define M' as follows:

$$M' = \begin{cases} M \cup e & \text{if } u \text{ unmatched in } M \text{ and } v \text{ unmatched in } M, \\ M \backslash e & \text{if } e \in M, \\ (M \cup e)\backslash e' & \text{if either } u \text{ or } v \text{ (but not both) is matched} \\ & \text{and } e' \text{ is the matching edge,} \\ M & \text{otherwise.} \end{cases}$$

(iii). Move to M' with probability $\min\{1, \frac{\pi_\lambda(M')}{\pi_\lambda(M)}\}$.

Note that $mc(\lambda)$ is aperiodic because $\mathcal{P}(M, M) \geq 1/2$ for all matchings M. The factor $1/2$ is not essential, any number in $(0, 1)$ will do in (i). It is also clear that $mc(\lambda)$ is irreducible (because all matchings communicate through the empty matching), and it is easy to check that $mc(\lambda)$ satisfies the detailed balance condition

$$\pi_\lambda(M)P(M, M') = \pi_\lambda(M')P(M', M),$$

for all matchings M, M'. We conclude that $mc(\lambda)$ has stationary distribution π_λ and the process is reversible under this distribution. Furthermore, for any initial state, the distribution of the state converges to π_λ. By a clever application of the so-called canonical path method, Jerrum and Sinclair obtained the following bound for the mixing time of $mc(\lambda)$.

Theorem 5.2.5. [Jerrum, Sinclair] *The mixing time of $mc(\lambda)$ satisfies*[1]*:*

$$\tau(\varepsilon) \leq 4|E_G|n\lambda'(n\ln(4n) + n\ln(\lambda') + \ln(\varepsilon^{-1})), \tag{5.2.9}$$

where $\lambda' = \max\{1, \lambda\}$ and $n = \lceil |V_G|/2 \rceil$.

5.2.3 A result on spatial dependencies

The following theorem is very similar to a result in Section 3 of [Ber99], the ideas of which go back to [Ber93] and [BM94], but slightly stronger and more convenient for our purposes. Therefore we give a fairly detailed proof. Recall the definitions of $D(\omega, \omega')$ and $deg(v)$ from Section 5.1.

Theorem 5.2.6. *Let μ, for a given value of λ, be the monomer-dimer distribution on a graph $G = (V_G, E_G)$. Let $\Delta \subset E_G$ and let $\alpha, \beta \in \Omega_{\partial\Delta}$. Then a coupling $\mathcal{P}_{\Delta,\alpha,\beta}$ of μ_Δ^α and μ_Δ^β exists such that*

$$E_{\Delta,\alpha,\beta}(\#\{e \in \Delta : e \text{ edge of disagreement}\}) \leq 2c\lambda|D_{\partial\Delta}(\alpha, \beta)|, \tag{5.2.10}$$

where $E_{\Delta,\alpha,\beta}$ denotes expectation with respect to $\mathcal{P}_{\Delta,\alpha,\beta}$ and c equals $\max_{v\in\Delta}\{deg(v)\} - 1$.

Proof. Let Δ, α and β be as in the statement of the theorem. We construct the desired coupling $\mathcal{P}_{\Delta,\alpha,\beta}$ on $\Omega_\Delta \times \Omega_\Delta$ as follows. Let x and y be independent random configurations with distribution μ_Δ^α and μ_Δ^β, respectively. Modify these configurations in the following way: for every $\tilde{e} \in D_{\partial\Delta}(\alpha, \beta)$ define the set

$$\begin{aligned} S(\tilde{e}, x, y) \quad = \quad & \{e \in D_\Delta(x, y) : \exists \text{ a sequence } \tilde{e} \sim e_1 \sim e_2 \sim \cdots \sim e_n = e \\ & \text{of distinct edges in } \Delta, \text{ with } \forall i \in \{1, \ldots, n\} : x_{e_i} \neq y_{e_i}\}. \end{aligned}$$

[1] In fact, Proposition 2.14 of Jerrum and Sinclair [JS96] states

$$\tau(\varepsilon) \leq 4|E_G|n\lambda'(n(\ln(n) + \ln(\lambda')) + \ln(\varepsilon^{-1})). \tag{5.2.8}$$

However, we could only verify the proof when the factor $\ln(n)$ is replaced by $\ln(4n)$, in (5.2.8).

We call such a sequence a *path of disagreement* of length n from v to e, where v is the common endpoint of $\tilde{e}$ and e_1. Let $\mathcal{D}_{x,y}^{\alpha,\beta}$ be the union of paths of disagreement leaving from $D_{\partial\Delta}(\alpha,\beta)$:

$$\mathcal{D}_{x,y}^{\alpha,\beta} := \cup_{\tilde{e}\in D_{\partial\Delta}(\alpha,\beta)} S(\tilde{e},x,y).$$

The modified configurations $\hat{x}$ and $\hat{y}$ are defined by:

$$\begin{aligned}\hat{x}_e &= \begin{cases} x_e & \text{if } e \in \mathcal{D}_{x,y}^{\alpha,\beta}, \\ y_e & \text{otherwise,}\end{cases} \\ \hat{y}_e &= y_e \text{ for all } e \in \Delta.\end{aligned}$$

Note that the configurations $\hat{x}$ and $\hat{y}$ only differ from each other on $\mathcal{D}_{\hat{x},\hat{y}}^{\alpha,\beta}$ and furthermore,

$$\mathcal{D}_{\hat{x},\hat{y}}^{\alpha,\beta} \text{ equals } \mathcal{D}_{x,y}^{\alpha,\beta}.$$

We define $\mathcal{P}_{\Delta,\alpha,\beta}$ as the distribution of the pair $(\hat{x},\hat{y})$ constructed as above.

Lemma 5.2.7. *The distribution of $\mathcal{P}_{\Delta,\alpha,\beta}$ defined above is indeed a coupling of μ_Δ^α and μ_Δ^β.*

Proof. It is sufficient to show that $\hat{x}$ has distribution μ_Δ^α and that $\hat{y}$ has distribution μ_Δ^β. The latter follows trivially since $\hat{y} = y$. As for $\hat{x}$, we introduce auxiliary configurations $\tilde{x}$ and $\tilde{y}$. In words $(\tilde{x},\tilde{y})$ is the pair of configurations obtained from (x,y) by exchanging x and y on the set of edges that do not have a path of disagreement to $D_{\partial\Delta}(\alpha,\beta)$. More precisely,

$$\begin{aligned}\tilde{x}_e &= \begin{cases} x_e & \text{if } e \in \mathcal{D}_{x,y}^{\alpha,\beta}, \\ y_e & \text{otherwise.}\end{cases} \\ \tilde{y}_e &= \begin{cases} y_e & \text{if } e \in \mathcal{D}_{x,y}^{\alpha,\beta}, \\ x_e & \text{otherwise.}\end{cases}\end{aligned}$$

By appropriate use of the Markov property (see Lemma 5 in [Ber99]) the pair $(\tilde{x},\tilde{y})$ has the same distribution as the pair (x,y). Finally from the definitions it follows that $\tilde{x} = \hat{x}$. Hence $\hat{x}$ has distribution μ_Δ^α. □

We now show that the coupling $\mathcal{P}_{\Delta,\alpha,\beta}$ has property (5.2.10). First recall (see the remark just before Lemma 5.2.7) that the left side of (5.2.10) is equal to the expected size of $\mathcal{D}_{x,y}^{\alpha,\beta}$, where x and y are drawn independently from μ_Δ^α and μ_Δ^β, respectively. Therefore we study the paths of disagreement for the pair (x,y). So consider an edge $\tilde{e} \in D_{\partial\Delta}(\alpha,\beta)$, say $\tilde{e} = (v_1,v_2)$. Observe that if a path of disagreement of length k leaving from v_1 exists, then this path is unique. Otherwise, as one can easily check, there would be three distinct edges which share a common endpoint, and on each of them $x \neq y$. But then at least two of these edges have either $x = 1$ or $y = 1$, which contradicts the fact that x and y are allowable configurations. Similarly, a path of disagreement leaving from v_2 is unique. Define $l_i(\tilde{e})$ for $i = 1,2$ as the path

of disagreement of maximal length, starting from v_i. From the above observations we conclude that the left side of (5.2.10) is at most

$$\sum_{\tilde{e}\in D_{\partial\Delta}(\alpha,\beta)} E_{\Delta,\alpha,\beta}[|l_1(\tilde{e})| + |l_2(\tilde{e})|]$$

$$= \sum_{\tilde{e}\in D_{\partial\Delta}(\alpha,\beta)} \sum_{k=1}^{\infty} 2\,\mathcal{P}_{\Delta,\alpha,\beta}(|l_1(\tilde{e})| \geq k) \tag{5.2.11}$$

To complete the proof of Theorem 5.2.6 we must, in view of (5.2.11), bound the probability

$$P(l_1(\tilde{e}) \text{ has length } \geq k).$$

Before we do this, we state a simple general lemma. Consider the monomer-dimer model on the very special 'star-shaped' graph, which consists of n edges and $n+1$ vertices, one of which (the 'centre of the star') has one edge to each of the other vertices. It is clear that each allowable configuration on this graph has either no edges or one edge with value 1, and the latter has probability $\lambda n/(1+\lambda n)$. This probability is increasing in n. That observation, together with the Markov property mentioned in Section 5.1 implies immediately the following:

Lemma 5.2.8. *Consider the monomer-dimer model with parameter λ on an arbitrary finite graph G. Let v be a vertex of G and let A be a subset of the edges adjacent to v. Then the conditional probability that there exists an edge in A with value 1, given the values of all edges outside A, is at most*

$$\frac{\lambda|A|}{1+\lambda|A|}.$$

We now proceed with the proof of Theorem 5.2.6. Suppose a path of disagreement of length k exists. What is the conditional probability that a path of length $k+1$ exists? Let $e_1 \sim \cdots \sim e_k = e$ be the unique path of length k leaving from v_1, so that $\tilde{e} \sim e_1$. By the uniqueness property mentioned before, we have that the path of disagreement of length $k+1$ (if it exists) is an extension of the path of length k. Define

$$Adj(e) = \{b \in \Delta : b \sim e \text{ and } b \not\sim e_{k-1}\}.$$

Note that $|Adj(e)| \leq c$, with c as in the statement of the theorem.

By assumption, $x_e = 0$ and $y_e = 1$ or vice versa. Without loss of generality we assume the former. Since y is a matching, we have $y_b = 0$ for every edge $b \in Adj(e)$. Hence we have a path of disagreement of length $k+1$ if and only if an edge $a \in Adj(e)$ exists with $x_a = 1$. By Lemma 5.2.8 above, the (conditional) probability of this event is at most $\lambda c/(\lambda c+1)$. Iterating the above we get

$$\mathcal{P}_{\Delta,\alpha,\beta}(l_1(\tilde{e}) \text{ has length} \geq k) \leq (\frac{\lambda c}{\lambda c+1})^k. \tag{5.2.12}$$

Combining (5.2.12) with (5.2.11), it follows that the left side of (5.2.10) is at most

$$2 \sum_{\bar{e} \in D_{\partial\Delta}(\alpha,\beta)} \sum_{k=1}^{\infty} \left(\frac{\lambda c}{\lambda c + 1}\right)^k = 2\lambda c \cdot |D_{\partial\Delta}(\alpha, \beta)|.$$

This completes the proof of Theorem 5.2.6. □

Remark 5.2.9. In Section 5.3, we will only work with d-dimensional hypercubes Δ. For such sets Δ, each edge on the boundary $\partial\Delta$ has exactly one vertex in common with an edge in the box Δ. For these special cases, the above result is improved by a factor 2, so that

$$E_{\Delta,\alpha,\beta}(\#\{e \text{ edge of disagreement}\}) \leq \lambda(2d-1)|D_{\partial\Delta}(\alpha,\beta)|, \tag{5.2.13}$$

for every hypercube Δ.

5.3 Random sampling on subgraphs of the d-dimensional lattice

5.3.1 Description and motivation of the method

In Subsection 5.2.2 we stated the Jerrum-Sinclair result. This result holds for general graphs. In the present section we study certain specifically 'nice' graphs, say a d-dimensional torus, described below. Suppose we want to sample (approximately) from the monomer-dimer model for such a graph. According to the Jerrum-Sinclair result, Theorem 5.2.5, we can do this by running the Markov chain $mc(\lambda)$ a number of steps given by (5.2.9). For the torus this is, for fixed λ and ε, asymptotically of the order $(\text{volume})^3 \ln(\text{volume})$. Here 'volume' is the number of edges in the graph, or the number of vertices, since for these graphs these differ by a constant factor.

Can the bound on the number of steps of the Markov chain, for these special graphs, be improved? There are several possibilities. One approach is to use logarithmic Sobolev inequalities; the results on spatial dependencies in Subsection 5.2.3 imply a mixing condition which in turn, following a quite general theory developed by Stroock and Zeglarinski [SZ92], could lead to a mixing time of order $(\text{volume}) \ln(\text{volume})$. We write *could* because there is an extra, quite subtle condition which has to be checked to obtain such a bound from the Stroock-Zegarliński theory. See Theorem 1 in the survey paper by Frigessi, Martinelli and Stander [FMS97]. This result would be very interesting, but when one really wants to generate random samples, it is not sufficient to know the asymptotic order of the mixing time, but one needs an explicit upper bound to carry out the algorithm. To get reasonable explicit bounds from the Stroock-Zeglarinski theory is probably a lot of work though in our opinion, certainly worth the effort.

However, in this work we follow a somewhat different approach, which is based on a small modification of coupling and rescaling arguments, that have become quite

standard, see [AH87], [MO94]. This approach has the advantage that it gives, with relatively simple and few computations, an explicit bound whose asymptotic order is only a little bit worse than the above mentioned (volume) ln(volume); we get an extra factor of the order ln(volume).

Our approach is to combine (using rescaling and coupling arguments) the result of Jerrum and Sinclair (Theorem 5.2.5) with the result on spatial dependencies from Subsection 5.2.3. Although this approach applies to a larger class of graphs, we concentrate for simplicity on a graph Γ, which corresponds to a d-dimensional torus. See Remark (iii) in Subsection 5.3.4. More precisely, let N be a positive integer, and define Γ as the pair (V_Γ, E_Γ), where the set of vertices is defined as

$$V_\Gamma = \{0, \ldots, N\}^d,$$

and the set edges E_Γ is

$$E_\Gamma = \{(v_1, v_2) : v_1, v_2 \in V_\Gamma \text{ and } |v_1 - v_2| = 1 \mod (N-1)\}. \tag{5.3.1}$$

We would like to sample from the monomer-dimer distribution π_Γ (with parameter λ) on this graph.

One way of approximately sampling from this distribution on Γ is the following. Let Δ be a d-dimensional cube of length l. Here l depends on d and λ; a suitable value will be determined later. More precisely, Δ is the following set of edges.

$$\Delta = \{(v_1, v_2) : v_1, v_2 \in \{0, \ldots, l\}^d \text{ and } |v_1 - v_2| = 1\}.$$

Let $\widehat{X}(t), t = 0, 1, \ldots$ be a Markov chain with state space $\{0,1\}^{E_\Gamma}$, which starts in some configuration x_0 on Γ. Its transitions are described as follows. Suppose $\widehat{X}(t) = x$. Choose uniformly at random a vertex $i \in V_\Gamma$. Let $\Delta(i)$ be the box Δ shifted by i in the torus, i.e.

$$\Delta(i) = \{((v_1 + i) \mod N, (v_2 + i) \mod N) : (v_1, v_2) \in \Delta\}.$$

Consider the monomer-dimer distribution on $\Delta(i)$ with boundary condition $x_{\partial\Delta(i)}$ and parameter λ, denoted by $\mu_{\Delta(i)}^{x_{\partial\Delta(i)}}$. Now sample a configuration $\tilde{x}$ from this distribution. Then at time $t+1$ the configuration becomes

$$\widehat{X}(t+1)_e = \begin{cases} \widehat{X}(t)_e & \text{if } e \notin \Delta(i), \\ \tilde{x}_e & \text{if } e \in \Delta(i). \end{cases} \tag{5.3.2}$$

It can be proved, using the spatial mixing properties mentioned before, that for sufficiently large l the mixing time of this Markov chain is of order $\mathcal{O}(|V_\Gamma| \cdot \ln |V_\Gamma|)$ for fixed λ, i.e. of the same order as we mentioned above in connection with logarithmic Sobolev inequalities.

However, a problem arises when one actually tries to execute this algorithm. How to compute the distribution $\mu_{\Delta(i)}^{x_{\partial\Delta(i)}}$? Even for relatively small l, this is a huge problem.

If for example $d = 2$ and the length of the hypercube is 10, the state space has already more than 2^{100} elements. In practice, this algorithm cannot be used.

One way to proceed now would be to use certain comparison theorems to obtain a bound on the mixing time of the Markov chain $mc(\lambda)$ for this this model from the bound on the mixing time of the above 'block dynamics' (see [DSC93], [RT98] and [Mar97]). However, these comparison arguments do not involve the two mixing times directly, but indirectly via the spectral gap or logarithmic Sobolev constant. Since the relation between the mixing time ((5.2.6)-(5.2.7)) and these quantities is not tight, this method would introduce a factor of the order 'volume', so the final result would be of the order $(\text{volume})^2 \ln(\text{volume})$.

Therefore we do the following: Instead of drawing a configuration *exactly* from the distribution $\mu_{\Delta(i)}^{x_{\partial\Delta(i)}}$ mentioned before, we will sample *approximately* from this distribution. In other words, we replace each (macro) step in the Markov chain $\widehat{X}(t)$ by a number of (micro) steps, where each micro step corresponds to a transition of the Markov chain $mc(\lambda)$ on $\Delta(i)$ with boundary condition $x_{\partial\Delta(i)}$, studied by Jerrum and Sinclair. It will turn out that for fixed δ and λ, the total number of steps needed to obtain a 'δ-close' approximate sample from π_Γ is at most of order $(\text{volume}) \ln(\text{volume})^2$. See Corollary 5.3.3 at the end of this section.

More precisely, the modified Markov chain, which we denote by $X(t), t = 0, 1, \ldots$ has the same state space and initial state as $\widehat{X}(t)$, but the transitions are now as follows. Let $\delta > 0$ and suppose $X(t) = x$. As before, choose uniform at random a vertex $i \in \Gamma$; determine the box $\Delta(i)$, and consider the monomer-dimer distribution $\mu_{\Delta(i)}^{x_{\partial\Delta(i)}}$. We will approximate this distribution. To do this we first define

$$\varepsilon = \frac{\delta}{2}\Big(\frac{|\Delta| - \lambda(2d-1)|\partial\Delta|}{|\Delta||E_\Gamma|}\Big). \tag{5.3.3}$$

The choice of this value will become clear later. Now consider the (auxiliary) Markov chain $mc(\lambda)$ with respect to the monomer-dimer model on $\Delta(i)$, with boundary condition $x_{\partial\Delta(i)}$ as described in Subsection 5.2.2. Although the initial state does not matter to the computation below, it is natural to take it equal to $x_{\Delta(i)}$. Denote the distribution of this chain at time t by $\nu_{\Delta(i)}^{t,x}$. Let Δ_V be the set of vertices which are endpoints of edges in Δ. Theorem 5.2.5 now gives us that $\nu_{\Delta(i)}^{t,x}$ converges to $\mu_{\Delta(i)}^{x_{\partial\Delta(i)}}$ and that

$$d_V(\nu_{\Delta(i)}^{T,x}, \mu_{\Delta(i)}^{x_{\partial\Delta(i)}}) \leq \varepsilon, \tag{5.3.4}$$

where T is the minimal number of steps that the Markov chain has to take, given by

$$T = |\Delta||\Delta_V|\lambda' \left[|\Delta_V| \ln(2|\Delta_V|) + |\Delta_V| \ln(\lambda') + 2\ln(\varepsilon^{-1})\right], \tag{5.3.5}$$

where $\lambda' = \min\{1, \lambda\}$. Let $\tilde{x}$ be the configuration on $\Delta(i)$ after T transitions of $mc(\lambda)$; this is a sample of $\nu_{\Delta(i)}^{T,x}$. Now take

$$X(t+1)_e = \begin{cases} X(t)_e & \text{if } e \notin \Delta(i), \\ \tilde{x}_e & \text{if } e \in \Delta(i). \end{cases}$$

This completes our description of a (macro) step in the Markov chain $X(t)$. In the next section we will give an upper bound on the number of macro steps after which the variational distance between μ^t, the distribution of $X(t)$ and π_Γ becomes smaller than δ. The total number of micro steps needed then simply follows from multiplying this number by T from (5.3.5).

5.3.2 A bound on the number of steps

In this subsection we will bound the number of steps of the Markov chain $X(t)$ to approximately reach the stationary distribution π_Γ. First, we define a suitably coupled system $(X(t), Y(t)), t = 0, 1, \cdots$ where $X(t)$ is the Markov chain introduced in the previous subsection, and $Y(t)$ is a Markov chain with the same transition probabilities as $\widehat{X}(t)$, also introduced in the previous subsection, but which starts, and hence stays, in the stationary distribution π_Γ. Using the results in Section 5.2, we will obtain an upper bound for the variational distance between the distributions of $X(t)$ and $Y(t)$ for every time t. This is done by studying the number of edges of disagreement $|D_\Gamma(X(t), Y(t))|$.

More precisely, let $X(0) = x_0$ and let $Y(0)$ be drawn from the distribution π_Γ. Suppose at time t, $X(t) = x$ and $Y(t) = y$. Now we follow the description of a transition of $X(t)$ given in the previous subsection. However, instead of sampling a single configuration $\tilde{x}$ on $\Delta(i)$, we now sample a pair $(\tilde{x}, \tilde{y})$ as follows. First consider the following three distributions on $\Omega_{\Delta(i)}$: $\nu^{T,x}_{\Delta(i)}, \mu^{x_{\partial\Delta(i)}}_{\Delta(i)}$ and $\mu^{y_{\partial\Delta(i)}}_{\Delta(i)}$. From now on, we omit the subscript $\Delta(i)$. Let

$$\mathcal{P}_{opt,\nu^{T,x},\mu^{x_{\partial\Delta(i)}}} \tag{5.3.6}$$

be an optimal coupling of $\nu^{T,x}$ and $\mu^{x_{\partial\Delta(i)}}$. Further let

$$\mathcal{P}_{\mu^{x_{\partial\Delta(i)}},\mu^{y_{\partial\Delta(i)}}}$$

be a coupling of $\mu^{x_{\partial\Delta(i)}}$ and $\mu^{y_{\partial\Delta(i)}}$ which satisfies Theorem 5.2.6. Finally, let

$$\mathcal{P}_{\nu^{T,x},\mu^{y_{\partial\Delta(i)}}} \tag{5.3.7}$$

be a coupling of $\nu^{T,x}$ and $\mu^{y_{\partial\Delta(i)}}$ obtained from the previous two as described in the proof of Proposition 5.2.1. The expectation with respect to the distribution in (5.3.6) is denoted by $E_{opt,\nu^{T,x},\mu^{x_{\partial\Delta(i)}}}$. The expectations for the other couplings are denoted likewise. Now, sample a pair $(\tilde{x}, \tilde{y})$ from the last coupling (5.3.7). Take

$$X(t+1)_e = \begin{cases} X(t)_e & \text{if } e \notin \Delta(i) \\ \tilde{x}_e & \text{if } e \in \Delta(i) \end{cases} \qquad Y(t+1)_e = \begin{cases} Y(t)_e & \text{if } e \notin \Delta(i) \\ \tilde{y}_e & \text{if } e \in \Delta(i). \end{cases}$$

This completes the description of the transitions of the pair $(X(t), Y(t))$. Note that $\tilde{x}$ has been drawn from $\nu^{T,x}$ so that the Markov chain $X(t)$ has indeed the same transition probabilities as in Subsection 5.3.1. Similarly, note that $Y(t)$ has indeed the distribution π_Γ for each t.

Recall that μ^t denotes the distribution of $X(t)$. Let $N(t)$ denote the expectation of $|D_\Gamma(X(t), Y(t))|$. Using Proposition 5.2.4 we have

$$d_V(\mu^t, \pi_\Gamma) \leq \mathcal{P}(X(t) \neq Y(t)) \leq N(t). \tag{5.3.8}$$

Therefore we will study $N(t)$. In particular, we study the change in this quantity after one macro step of the coupled Markov chain. Using a property analogous to equation (5.2.2) in Proposition 5.2.1 we get

$$\begin{aligned} & E_{\nu^{T,x},\mu^{y_{\partial\Delta(i)}}}(\#\{e \in \Delta(i) : e \text{ edge of disagreement}\}) \\ & \quad \leq E_{opt,\nu^{T,x},\mu^{x_{\partial\Delta(i)}}}(\#\{e \in \Delta(i) : e \text{ edge of disagreement}) \\ & \quad + E_{\mu^{x_{\partial\Delta(i)}},\mu^{y_{\partial\Delta(i)}}}(\#\{e \in \Delta(i) : e \text{ edge of disagreement}). \end{aligned} \tag{5.3.9}$$

So we need upper bounds for the expectations in (5.3.9). By Theorem 5.2.6 and equation (5.2.13)

$$E_{\mu^{x_{\partial\Delta(i)}},\mu^{y_{\partial\Delta(i)}}}(\#\{e \in \Delta(i) : e \text{ edge of disagreement}) \leq \lambda(2d-1)|D_{\partial\Delta(i)}(x_{\partial\Delta(i)}, y_{\partial\Delta(i)})|. \tag{5.3.10}$$

Because the coupling $\mathcal{P}_{opt,\nu^{T,x},\mu^{x_{\partial\Delta(i)}}}$ is optimal, we have

$$\begin{aligned} & E_{opt,\nu^{T,x},\mu^{x_{\partial\Delta(i)}}}(\#\{e \in \Delta(i) : e \text{ edge of disagreement}) \\ & \quad \leq |\Delta(i)| \cdot \mathcal{P}_{opt,\nu^{T,x},\mu^{x_{\partial\Delta(i)}}}(unequal) \\ & \quad = |\Delta(i)| \cdot d_V(\nu^{T,x}, \mu^{x_{\partial\Delta(i)}}) \\ & \quad \leq \varepsilon \cdot |\Delta(i)|. \end{aligned} \tag{5.3.11}$$

The last inequality follows from (5.3.4). Together, inequalities (5.3.9) to (5.3.11) yield

$$\begin{aligned} & E_{\nu^{T,x},\mu^{y_{\partial\Delta(i)}}}(\#\{e \in \Delta(i) : e \text{ edge of disagreement}\}) \\ & \quad \leq \varepsilon \cdot |\Delta(i)| + \lambda(2d-1)|D_{\partial\Delta(i)}(x_{\partial\Delta(i)}, y_{\partial\Delta(i)})|. \end{aligned} \tag{5.3.12}$$

We now state and prove the following lemma concerning $N(t)$.

Lemma 5.3.1.

$$N(t+1) \leq b\ N(t) + \varepsilon|\Delta|, \tag{5.3.13}$$

where

$$b = 1 - \frac{|\Delta| - \lambda(2d-1)|\partial\Delta|}{|E_\Gamma|}. \tag{5.3.14}$$

Proof. Let $n(t) = |D_\Gamma(X(t), Y(t))|$. Note that the expectation of $n(t)$ is equal to $N(t)$. Suppose that $X(t), Y(t), i$ and hence $\Delta(i)$ are known. Consider the conditional expectation of the number of edges of disagreement that disappear during the transition $t \to t+1$:

$$\begin{aligned} & E[n(t) - n(t+1)|X(t) = x, Y(t) = y, \Delta(i) = A] \\ & \qquad = |D_A(x_A, y_A)| - E_{\nu^{T,x},\mu^{y_{\partial A}}}(|D_A(\tilde{x}, \tilde{y})|). \end{aligned}$$

By (5.3.12) this is larger than or equal to

$$|D_A(x_A, y_A)| - \varepsilon|A| - \lambda(2d-1)|D_{\partial A}(x_{\partial A}, y_{\partial A})|.$$

Averaging over A we get

$$\begin{aligned} &E[n(t) - n(t+1) | X(t) = x, Y(t) = y] \\ &\quad \geq E(|D_{\Delta(i)}(x_{\Delta(i)}, y_{\Delta(i)})| - \varepsilon|\Delta| - \lambda(2d-1)E(|D_{\partial\Delta(i)}(x_{\partial\Delta(i)}, y_{\partial\Delta(i)})|), \end{aligned} \tag{5.3.15}$$

where the expectation on the right refers to the distribution of $\Delta(i)$. Recall that this is the uniform distribution, so that by symmetry each edge of E_Γ has the same probability, $|\Delta|/|E_\Gamma|$, to belong to $\Delta(i)$. Similarly, the probability that a given edge belongs to $\partial\Delta(i)$ equals $|\partial\Delta|/|E_\Gamma|$. Filling this in in (5.3.15) we obtain

$$E[n(t) - n(t+1) | X(t) = x, Y(t) = y] \geq N(t)\Big(\frac{|\Delta| - \lambda(2d-1)|\partial\Delta|}{|E_\Gamma|}\Big) - \varepsilon|\Delta|.$$

Taking expectations and dividing by $N(t)$ we get

$$\frac{N(t) - N(t+1)}{N(t)} \geq \Big(\frac{|\Delta| - \lambda(2d-1)|\partial\Delta|}{|E_\Gamma|}\Big) - \varepsilon|\Delta|,$$

from which the lemma follows immediately. □

For the moment we assume that the following inequalities hold.

$$0 < 1 - \frac{|\Delta| - \lambda(2d-1)|\partial\Delta|}{|E_\Gamma|} < 1. \tag{5.3.16}$$

Iterating equation (5.3.13), we get

$$\begin{aligned} N(t+1) &\leq b^{t+1}|E_\Gamma| + \varepsilon|\Delta| \sum_{i=0}^{t} b^i \\ &= b^{t+1}|E_\Gamma| + \frac{1 - b^{t+1}}{1-b}\varepsilon|\Delta| \\ &= b^{t+1}|E_\Gamma| + (1 - b^{t+1})\frac{\delta}{2}. \end{aligned}$$

Here we used the definitions of ε and b, equations (5.3.3) and (5.3.14), and the fact that $N(0) \leq |E_\Gamma|$.

With (5.3.8) this gives

$$d_V(\mu^t, \pi_\Gamma) \leq b^t|E_\Gamma| + (1 - b^t)\frac{\delta}{2}. \tag{5.3.17}$$

If we want to find t such that the above mentioned variational distance is smaller than δ, it suffices to solve

$$b^t|E_\Gamma| \leq \frac{\delta}{2} \tag{5.3.18}$$

Taking logarithms on both sides of (5.3.18) and using that $\ln(1-x) \leq -x$ for $0 < x < 1$, we find that (5.3.18) holds if

$$t \geq \frac{|E_\Gamma| \ln(2|E_\Gamma|\delta^{-1})}{|\Delta| - \lambda(2d-1)|\partial\Delta|}. \tag{5.3.19}$$

Recall that every step of the Markov chain $X(t)$ is in fact a macro step which corresponds with T micro steps in some box $\Delta(i)$, where T is given by (5.3.5). Hence the total number of micro steps $\tau(\delta)$ after which the distribution of $X(t)$ has variational distance $\leq \delta$ from π_Γ is at most T times the right side of (5.3.19), i.e.

$$\begin{aligned} \tau(\delta) \quad \leq \quad & |\Delta||\Delta_V|\lambda' \cdot \Big[|\Delta_V| \ln(2|\Delta_V|) + |\Delta_V| \ln(\lambda') + 2\ln(\varepsilon^{-1})\Big] \times \\ & \Big(\frac{|E_\Gamma| \ln(2|E_\Gamma|\delta^{-1})}{|\Delta| - \lambda(2d-1)|\partial\Delta|}\Big), \end{aligned} \tag{5.3.20}$$

where $\varepsilon = \varepsilon(\delta)$ is defined as in (5.3.3) and $\lambda' = \max\{1, \lambda\}$. Optimization considerations on a simplified modification of the right side of (5.3.20) lead to the following choice of the length l of Δ:

$$l := \lceil \lambda(4d+2) \rceil. \tag{5.3.21}$$

Note that $|\Delta_V| = (l+1)^d, |\Delta| = dl(l+1)^{d-1}$ and $|\partial\Delta| = 2d(l+1)^{d-1}$ so that with l as in (5.3.21),

$$|\Delta_V| = (\lceil \lambda(4d+2) + 1 \rceil)^d, \tag{5.3.22}$$

$$|\Delta| = d\lceil \lambda(4d+2) \rceil (\lceil \lambda(4d+2) \rceil + 1)^{d-1}, \tag{5.3.23}$$

$$|\partial\Delta| = 2d(\lceil \lambda(4d+2) \rceil + 1)^{d-1}. \tag{5.3.24}$$

Using (5.3.22)-(5.3.24), it is easy to check that for every $\lambda > 0$ and every $d \geq 2$, the above choice of l implies the upper bound in (5.3.16). The lower bound in (5.3.16) is satisfied if $|E_\Gamma| > |\Delta|$. Using (5.3.22),(5.3.23), (5.3.24) and (5.3.3), we can now express the upper bound (5.3.20) on $\tau(\delta)$ completely in terms of δ, λ, d and $|E_\Gamma|$.

5.3.3 Summary of the algorithm and the main result

Concluding, we can state the following. Let $0 < \delta < 1$ and $\lambda > 0$. Consider the monomer-dimer distribution π_Γ with parameter λ on the d-dimensional torus Γ, as described in Subsection 5.3.1. Take $l = \lceil (4d+2)\lambda \rceil$ and let Δ be the hypercube with length $l+1$ as described in Subsection 5.3.1. Compute ε from (5.3.3). Finally, compute T for the above choice of l as in (5.3.5). Consider the Markov chain $X(t)$ with state space $\{0,1\}^{E_\Gamma}$ with transitions as follows: choose uniformly at random a

vertex $i \in \Gamma$ and consider the box $\Delta(i) = (i + \Delta)$. On this box, run the Markov chain $mc(\lambda)$ described in Section 5.2.2 for T steps, with the current $X(t)$ values on the boundary fixed. These steps are called micro steps. This completes one transition, or macro step, in the chain $X(t)$.

Theorem 5.3.2. *In the algorithm described above, the number of micro steps $\tau(\delta)$ after which the distribution of $X(t)$ has variational distance at most δ from the stationary distribution π_Γ satisfies*

$$\begin{aligned} \tau(\delta) &\leq T \cdot \Big(\frac{|E_\Gamma| \ln(2|E_\Gamma|\delta^{-1})}{|\Delta| - \lambda(2d-1)|\partial\Delta|} \Big) \\ &= |\Delta||\Delta_V|\lambda' \cdot \Big[|\Delta_V| \ln(2|\Delta_V|) + |\Delta_V| \ln(\lambda') + 2\ln(\varepsilon^{-1}) \Big] \times \\ &\quad \Big(\frac{|E_\Gamma| \ln(2|E_\Gamma|\delta^{-1})}{|\Delta| - \lambda(2d-1)|\partial\Delta|} \Big), \end{aligned} \tag{5.3.25}$$

where $|\Delta_V|, |\Delta|, |\partial\Delta|$ and ε are given by (5.3.22),(5.3.23), (5.3.24) *and* (5.3.3), *respectively, and $\lambda' = \max\{1, \lambda\}$.*

This result gives immediately (note the dependence of ε on $|E_\Gamma|$):

Corollary 5.3.3. *For the algorithm above, if λ, d and δ are fixed, $\tau(\delta)$ satisfies*

$$\tau(\delta) = \mathcal{O}\Big(|E_\Gamma| \ln(|E_\Gamma|)^2 \Big). \tag{5.3.26}$$

5.3.4 Remarks

(i). Since the definition of the Markov chain $X(t)$ depends on δ, it is strictly speaking not correct to call $\tau(\delta)$ in (5.3.20) its mixing time.

(ii). From (5.3.26) it follows that for fixed λ and δ, on a large torus our bound is considerably better than the bound of Jerrum and Sinclair in Theorem 5.2.5. Note that on a torus, the number of edges equals the dimension times the number of vertices, so $|E_\Gamma| = d|V_\Gamma|$. However, our bound in (5.3.20) involves a factor $\lambda^{2d}\lambda'$, whereas the bound of Jerrum and Sinclair is linear in λ, which is important in certain applications [JS96]. Hence if the size of the torus is relatively small with respect to λ, their bound is better than ours.

(iii). The algorithm in the previous subsection was described on a torus Γ. A similar result is still valid when the algorithm is applied to a sufficiently nice finite subset of $\mathbb{Z}^d$, for instance a hypercube $H = (V_H, E_H)$ where $V_H = \{0, \ldots, m\}^d$ and $E_H = \{(v_1, v_2) : v_1, v_2 \in V_H \text{ and } |v_1 - v_2| = 1\}$. Since H is not a torus, the box $\Delta(i)$ must now be defined as $\Delta(i) \cap E_H$, where i is now the center of the box $\Delta(i)$. The fact that in some cases $\Delta(i)$ consists of roughly $|\Delta|/2^d$ elements

(when i is 'close to a corner') leads to an increase of the size of a suitable Δ. This in turn leads to a number of micro steps needed in the procedure which is a constant (depending on the dimension d) larger than for the torus.

(iv). One may think of several modifications of our computations to improve the right hand side of (5.3.25). For instance it would be interesting and worth trying to improve Theorem 5.2.6. As to alternative methods, see the remark about logarithmic Sobolev inequalities in the beginning of this section.

Appendix A

Some theorems from percolation theory

This appendix contains a short discussion of some well-known theorems in the field of percolation theory on $\mathbb{Z}^d$. We assume that $d \geq 2$. For $p < p_c$, all clusters are a.s. finite by definition. Even stronger, it can be shown that the distribution of the radius of clusters decays exponentially. Let $S(n)$ be the sphere of radius n, i.e. $S(n)$ contains all sites that are at graph distance n from the origin O. The following theorem was proved independently by Menshikov [Men87] and Aizenman and Barsky [AB87].

Theorem A.1. *Let $p < p_c$. There exists a $\phi(p) > 0$ such that*

$$\mathcal{P}_p(O \to S(n)) < e^{-n\phi(p)},$$

for all n.

Clearly, for $p > p_c$ such a result cannot hold. However, if we restrict ourselves to connections that do reach $S(n)$ but do not reach infinity, we have a similar result. Let C_O denote the cluster of the origin and let $|C_O|$ denote its size.

Theorem A.2. *Let $p > p_c$. There exists a $\psi(p) > 0$ such that*

$$\mathcal{P}_p(O \to S(n), |C_O| < \infty) < e^{-n\psi(p)},$$

for all n.

This theorem is due to Chayes, Chayes, Grimmett, Kesten and Schonmann [CCG+89] and Chayes,Chayes and Newman [CCN87].

Much used features of ordinary independent percolation in dimension 2 are the Russo-Seymour-Welsh theorem and its consequences [Rus78], [Rus81], [SW78]. Let $R(n,m)$ denote the rectangle $[0,n] \times [0,m]$ and let $LR\,[R(n,m)]$ denote the event that

there exists an open left-right crossing in the rectangle. An open left-right crossing is an open path of edges from $\{0\} \times [0, m]$ to $\{n\} \times [0, m]$ inside the rectangle. We use the notation $B(n)$ for $R(n, n)$.

Theorem A.3. [RSW] *There exists a function $\pi : [0,1] \to [0,1]$ independent of n, satisfying $\pi(\varepsilon) \to 1$ as $\varepsilon \to 1$ and $\pi(\varepsilon) > 0$ when $\varepsilon > 0$ such that*

$$\mathcal{P}_p(LR[B(n)]) > \varepsilon \Longrightarrow \mathcal{P}_p(LR[R(3n, n)]) > \pi(\varepsilon). \tag{A.4}$$

This theorem is very useful since it provides lower bounds for crossing probabilities that do not depend on n. For example, if $\theta(p) > 0$ it is not so hard to show that the probability of crossing a square goes to 1 using uniqueness of the infinite cluster. The RSW theorem now shows that also the probabilities of crossings in rectangles goes to 1. A much used consequence of the RSW-theorem is the following: Let $O(n, 3n)$ denote the event that inside the annulus $A(n, 3n) = B(3n) \backslash B(n)$ there is an open circuit (path returning to its starting point) around $B(n)$.

Corollary A.5. *There exists a function $\pi' : [0,1] \to [0,1]$ independent of n, satisfying $\pi'(\varepsilon) \to 1$ as $\varepsilon \to 1$ and $\pi'(\varepsilon) > 0$ when $\varepsilon > 0$ such that*

$$\mathcal{P}_p(LR[B(n)]) > \varepsilon \Longrightarrow \mathcal{P}_p(O(n, 3n)) > \pi'(\varepsilon).$$

An application of the Borel-Cantelli lemma shows that at parameter value p_c, there will be infinitely many open circuits around each point in the lattice.

Bibliography

[AB87] M. Aizenman and D. J. Barsky. Sharpness of the phase transition in percolation models. *Comm. Math. Phys.*, 108(3):489–526, 1987.

[AG91] M. Aizenman and G. Grimmett. Strict monotonicity for critical points in percolation and ferromagnetic models. *J. Statist. Phys.*, 63(5-6):817–835, 1991.

[AH87] M. Aizenman and R. Holley. Rapid convergence to equilibrium of stochastic Ising models in the Dobrushin Shlosman regime. In *Percolation theory and ergodic theory of infinite particle systems (Minneapolis, Minn., 1984–1985)*, volume 8 of *IMA Vol. Math. Appl.*, pages 1–11. Springer, New York, 1987.

[AKN87a] M. Aizenman, H. Kesten, and C. M. Newman. Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range percolation. *Comm. Math. Phys.*, 111(4):505–531, 1987.

[AKN87b] M. Aizenman, H. Kesten, and C. M. Newman. Uniqueness of the infinite cluster and related results in percolation. In *Percolation theory and ergodic theory of infinite particle systems (Minneapolis, Minn., 1984–1985)*, volume 8 of *IMA Vol. Math. Appl.*, pages 13–20. Springer, New York, 1987.

[Ald00] D.J. Aldous. The percolation process on a tree where infinite clusters are frozen. *Math. Proc. Cambridge Philos. Soc.*, 128(3):465–477, 2000.

[BA91] D. J. Barsky and M. Aizenman. Percolation critical exponents under the triangle condition. *Ann. Probab.*, 19(4):1520–1536, 1991.

[Bak97] P. Bak. *How Nature Works: the science of self-organized criticality.* Oxford University Press, 1997.

[BB00] J. van den Berg and R. Brouwer. Random sampling for the monomer-dimer model on a lattice. *J. Math. Phys.*, 41(3):1585–1597, 2000. Probabilistic techniques in equilibrium and nonequilibrium statistical physics.

[BB04a] J. van den Berg and R. Brouwer. Forest-fires near the critical time. 2004. Submitted.

[BB04b] J. van den Berg and R. Brouwer. Self-destructive percolation. *Random Structures Algorithms*, 24(4):480–501, 2004.

[Ber93] J. van den Berg. A uniqueness condition for Gibbs measures, with application to the 2-dimensional Ising antiferromagnet. *Comm. Math. Phys.*, 152(1):161–166, 1993.

[Ber99] J. van den Berg. On the absence of phase transition in the monomer-dimer model. In *Perplexing problems in probability*, volume 44 of *Progr. Probab.*, pages 185–195. Birkhäuser Boston, Boston, MA, 1999.

[BH57] S. R. Broadbent and J. M. Hammersley. Percolation processes. I. Crystals and mazes. *Proc. Cambridge Philos. Soc.*, 53:629–641, 1957.

[BJ05] J. van den Berg and A.A. Jarai. On the asymptotic density in a one-dimensional self-organized critical forest-fire model. *Comm. Math. Phys.*, 253(3):633–644, 2005.

[BK84] J. van den Berg and M. Keane. On the continuity of the percolation probability function. In R.T. Durrett, editor, *Particle Systems, Random Media and Large Deviations*, volume 26 of *Contemporary Mathematics Series*, pages 61–65. AMS, 1984.

[BK85] J. van den Berg and H. Kesten. Inequalities with applications to percolation and reliability. *J. Appl. Probab.*, 22(3):556–569, 1985.

[BK89] R. M. Burton and M. Keane. Density and uniqueness in percolation. *Comm. Math. Phys.*, 121(3):501–505, 1989.

[BM94] J. van den Berg and C. Maes. Disagreement percolation in the study of Markov fields. *Ann. Probab.*, 22(2):749–763, 1994.

[BT01] J. van den Berg and B. Tóth. A signal-recovery system: asymptotic properties, and construction of an infinite-volume process. *Stochastic Process. Appl.*, 96(2):177–190, 2001.

[BTW88] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality. *Phys. Rev. A (3)*, 38(1):364–374, 1988.

[CC84] J.T. Chayes and L. Chayes. Percolation and random media. In *Phénomenes critiques, systèmes aléatoires théories de jauge*, Session XLIII, Les Houches, 1984.

[CCG+89] J. T. Chayes, L. Chayes, G. R. Grimmett, H. Kesten, and R. H. Schonmann. The correlation length for the high-density phase of Bernoulli percolation. *Ann. Probab.*, 17(4):1277–1302, 1989.

[CCN87] J. T. Chayes, L. Chayes, and C. M. Newman. Bernoulli percolation above threshold: an invasion percolation analysis. *Ann. Probab.*, 15(4):1272–1287, 1987.

[DCS93] B. Drossel, S. Clar, and F. Schwabl. Exact results for the one-dimensional self-organized critical forest-fire model. *Phys. Rev. Lett.*, 71(23):3739–3742, 1993.

[Dro96] B. Drossel. Self-organized criticality and synchronization in the forest-fire model. *Phys. Rev. Lett.*, 76:736–739, 1996.

[DS92] B. Drossel and F. Schwabl. Self-organized critical forest-fire model. *Phys. Rev. Lett.*, 69:1629–1632, 1992.

[DSC93] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.*, 3(3):696–730, 1993.

[Dur88] R. Durrett. *Lecture notes on particle systems and percolation.* Wadsworth & Brooks/Cole, Pacific Grove, California, 1988.

[FMS97] A. Frigessi, F. Martinelli, and J. Stander. Computational complexity of Markov chain Monte Carlo methods for finite Markov random fields. *Biometrika*, 84(1):1–18, 1997.

[GKN92] A. Gandolfi, M. S. Keane, and C. M. Newman. Uniqueness of the infinite component in a random graph with applications to percolation and spin glasses. *Probab. Theory Related Fields*, 92(4):511–527, 1992.

[GKR88] A. Gandolfi, M. Keane, and L. Russo. On the uniqueness of the infinite occupied cluster in dependent two-dimensional site percolation. *Ann. Probab.*, 16(3):1147–1157, 1988.

[Gra02] P. Grassberger. Critical behaviour of the Drossel-Schwabl forest-fire model. *New Journal of Physics*, 4:17.1–17.5, 2002.

[Gri99] G. Grimmett. *Percolation*, volume 321 of *Grundlehren der mathematischen Wissenschaften.* Springer-Verlag, second edition, 1999.

[HL72] O. Heilmann and E. Lieb. Theory of monomer-dimer systems. *Comm. Math. Phys.*, 25:190–232, 1972.

[Hon97] A. Honecker. Program for simulating the two-dimensional forest-fire model, 1997. http://www-public.tu-bs.de:8080/~honecker/software/forest2d.html.

[HS90] T. Hara and G. Slade. Mean-field critical behaviour for percolation in high dimensions. *Comm. Math. Phys.*, 128(2):333–391, 1990.

[HS94] T. Hara and G. Slade. Mean-field behaviour and the lace expansion. In *Probability and phase transition (Cambridge, 1993)*, volume 420 of *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, pages 87–122. Kluwer Acad. Publ., Dordrecht, 1994.

[Hug96] B. D. Hughes. *Random walks and random environments. Vol. 2.* Oxford Science Publications. The Clarendon Press Oxford University Press, New York, 1996. Random environments.

[Jen98] H. Jensen. *Self-Organized Criticality*, volume 10 of *Cambridge Lecture Notes in Physics.* Cambridge University Press, 1998.

[JS96] M. Jerrum and A. Sinclair. The Markov chain Monte Carlo method: an approach to approximate counting and integration. In D. Hochbaum, editor, *Approximation Algorithms for NP-hard Problems.* PWS Publishing, Boston, 1996.

[Kes80] H. Kesten. The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$. *Comm. Math. Phys.*, 74(1):41–59, 1980.

[KK97] J. Kahn and P.M. Kayll. On the stochastic independence properties of hard-core distributions. *Combinatorica*, 17(3):369–391, 1997.

[Lig85] T. M. Liggett. *Interacting Particle Systems*, volume 276 of *Grundlehren der mathematischen Wissenschaften.* Springer-Verlag, New York, 1985.

[Lin92] T. Lindvall. *Lectures on the coupling method.* Wiley series in Probability and Mathematical Statistics. Wiley, New York, 1992.

[LSW02] G. F. Lawler, O. Schramm, and W. Werner. One-arm exponent for critical 2D percolation. *Electron. J. Probab.*, 7:no. 2, 13 pp. (electronic), 2002.

[Mar97] F. Martinelli. Lecture notes on glauber dynamics for discrete spin models. Saint Flour Summer School in Probability. 1997.

[Men87] M. V. Menshikov. Percolation theory. Coincidence of the critical values $P_H = P_T$. In *Proceedings of the 1st World Congress of the Bernoulli Society, Vol. 1 (Tashkent, 1986)*, pages 231–234, Utrecht, 1987. VNU Sci. Press.

[MMT98] B.D. Malamud, G. Morein, and D.L. Turcotte. Forest fires: An example of self-organized critical behaviour. *Science*, 281:1840–1841, 1998.

[MO94] F. Martinelli and E. Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region. I. The attractive case. *Comm. Math. Phys.*, 161(3):447–486, 1994.

[Pen00] J. Pennanen. Self-organized criticality of the forest fire model. Master's thesis, University of Helsinki, apr 2000.

[Rei00] D. Reimer. Proof of the van den Berg-Kesten conjecture. *Combin. Probab. Comput.*, 9(1):27–32, 2000.

[RT98] D. Randall and P. Tetali. Analyzing Glauber dynamics by comparison of Markov chains. In *LATIN'98: theoretical informatics (Campinas, 1998)*, volume 1380 of *Lecture Notes in Comput. Sci.*, pages 292–304. Springer, Berlin, 1998.

[Rus78] L. Russo. A note on percolation. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 43(1):39–48, 1978.

[Rus81] L. Russo. On the critical percolation probabilities. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 56(2):229–237, 1981.

[SDS02] K. Schenk, B. Drossel, and F. Schwabl. Self-organised critical forest-fire model on large scales. *Phys. Rev. E*, 65:026135–1–8, 2002.

[Sta85] D. Stauffer. *Introduction to percolation theory.* Taylor and Francis Ltd., London and Philadelphia, 1985.

[SW78] P. D. Seymour and D. J. A. Welsh. Percolation probabilities on the square lattice. *Ann. Discrete Math.*, 3:227–245, 1978. Advances in graph theory (Cambridge Combinatorial Conf., Trinity College, Cambridge, 1977).

[SZ92] D. W. Stroock and B. Zegarliński. The logarithmic Sobolev inequality for continuous spin systems on a lattice. *J. Funct. Anal.*, 104(2):299–326, 1992.

[Tho00] H. Thorisson. *Coupling, stationarity, and regeneration.* Probability and its Applications (New York). Springer-Verlag, New York, 2000.

[TW04] B. Tsirelson and W. Werner. *Lectures on probability theory and statistics*, volume 1840 of *Lecture Notes in Mathematics.* Springer-Verlag, Berlin, 2004. Lectures from the 32nd Probability Summer School held in Saint-Flour, July 7–24, 2002, Edited by Jean Picard.

[Wie81] J. C. Wierman. Bond percolation on honeycomb and triangular lattices. *Adv. in Appl. Probab.*, 13(2):298–313, 1981.

Samenvatting

Percolatie, bosbranden en monomer-dimers

Dit proefschrift behandelt een aantal kanstheoretische modellen: Zelf-vernietigende percolatie, het zelf-organiserende bosbrand model, bevroren percolatie en het monomer-dimer model. Hoofdstuk 5, over het monomer-dimer model staat min of meer op zichzelf en we bespreken het later. De overige hoofdstukken hebben een aantal overkoepelende thema's. Een daarvan is zelf-organiserend kritiek gedrag. Zelf-organiserend kritiek gedrag is een populair onderwerp omdat het kan verklaren hoe complexe structuren ontstaan uit simpele basisprincipes. Het geeft een wiskundige verklaring voor een in de natuur veelvuldig voorkomend gegeven, namelijk dat veel biologische en geologische systemen er op elke schaal hetzelfde uitzien (een fractale structuur hebben). Denk bijvoorbeeld aan de fjorden van Noorwegen, of de vorm van een slakkenhuis. Wiskundig gezien uit zich dit in het voorkomen van machtswetten: De dichtheden van de belangrijkste grootheden in de modellen volgen een machtswet. Als voorbeeld nemen we het bosbrandmodel. Een belangrijke grootheid is daar de grootte van dicht bij elkaar gelegen bomen, de grootte van een *cluster*. Laat $p(s)$ de kans zijn dat het cluster van een bepaald punt s bomen bevat. Als zelf-organisered kritiek gedrag voorkomt, dan moet er een constante τ zijn zodat

$$p(s) \simeq s^{-\tau}.$$

In het algemeen zijn dit soort uitspraken moeilijk te bewijzen, maar voor het één-dimensionale bosbrandmodel kunnen we laten zien dat

$$p(s) \simeq s^{-1},$$

voor waarden van s die niet te groot zijn, zie Sectie 3.2. Zelf-organiserend kritiek gedrag wordt geassociëerd met o.a. bosbranden, aardbevingen, epidemieën, zandhopen, evolutie en de verdeling van woorden in een text. Niet alles hiervan is wiskundig bewezen. Vaak is het observeren van een machtswet al genoeg om zelf-organiserend kritiek gedrag te suggereren.

De term zelf-organiserend kritiek gedrag valt uiteen in 'zelf-organiserend' en 'kritiek gedrag'. In sommige traditionele modellen in de statistische fysica, als het Ising

model of onafhankelijke percolatie, vinden we ook de bovengenoemde machtswetten. Het gedrag van deze modellen hangt af van de waarde van een parameter (in het Ising model de temperatuur, in het percolatiemodel de kans dat een punt bezet is). We zien een fase-overgang bij een kritieke waarde van de parameter. Dat betekent dat het gedrag van het model essentieel anders is bij het vastleggen van de parameter boven en onder deze kritieke waarde. In het percolatiemodel bijvoorbeeld, is er boven de kritieke waarde met kans één een oneindig cluster van bezette punten; onder de kritieke waarde zijn alle clusters van bezette punten met kans één eindig. Wanneer de parameter deze kritieke waarde *aanneemt*, zien we de machtswetten verschijnen. Dit is echter een instabiele toestand, want met een kleine verandering van de parameter gelden de machtswetten niet meer. Deze instabiele toestand noemen we kritiek. Een zelf-organiserend kritiek systeem is robuuster. Het is niet nodig de parameter heel nauwkeurig te kiezen om de machtswetten te observeren, deze zijn inherent aan de dynamica van het systeem. Het systeem stuurt *zichzelf* naar een kritieke toestand, vandaar de term zelf-organiserend.

Een ander overkoepelend thema in dit proefschrift is het bestaan van processen. 'Bestaan' is hier bedoeld in zijn wiskundige context: De regels die het proces definiëren moeten consistent met elkaar zijn. Zodra we het over oneindige structuren hebben is het bestaan van een proces vaak een probleem. Hoewel het in dit proefschrift misschien niet altijd even duidelijk naar voren komt, zijn de problemen rond het bestaan van processen vaak van invloed geweest op de (al dan niet beantwoorde) vraagstukken.

Het probleem waar vele andere uit voort zijn gekomen is het permanent zelf-vernietigende percolatie proces. We kunnen dit intuitief als volgt beschrijven. We bekijken het twee-dimensionale rooster $\mathbb{Z}^2$ in de tijd. Alle punten zijn 'inactief' op tijdstip 0. Onafhankelijk van elkaar worden de punten 'actief' en zodra een punt in een actief oneindig cluster zit, maken we dat actieve cluster in zijn geheel weer inactief. Het is de vraag of dit proces bestaat. Dit is een interessante vraag omdat de intuitieve beschrijving doet vermoeden dat dit process zelf-organiserend kritiek gedrag vertoont. Tot de kritieke tijd (corresponderend met de kritieke waarde van gewone onafhankelijke percolatie) zijn er geen oneindige actieve clusters. Meteen daarna ontstaat een oneindig actief cluster en dit cluster wordt inactief gemaakt. Omdat het oneindige cluster erg 'dun' was op dit tijdstip, lijkt het waarschijnlijk dat meteen daarna opnieuw een oneindig actief cluster ontstaat dat inactief wordt gemaakt, enzovoorts. Het systeem zou dan in een kritieke toestand blijven; een typisch voorbeeld van zelf-organiserend kritiek gedrag. Het permanent zelf-vernietigende proces is de inspiratiebron geweest voor de hoofstukken 2,3 en 4.

Hoofdstuk 2 bespreekt zelf-vernietigende percolatie. In dit model wordt eenmaal onafhankelijke percolatie uitgevoerd, dan wordt het oneindige bezette cluster vernietigd en vervolgens wordt elk punt weer onafhankelijk bezet met een kleine kans. We bespreken algemene eigenschappen als monotoniciteit, associativiteit en continuïteit van de percolatie functie. Verder bestuderen we een eigenschap in dimensie 2 (die we noch kunnen bewijzen, noch ontkrachten) waaruit volgt dat de percolatie functie een discontinuïteit heeft. Verder volgt uit deze eigenschap dat het permanent zelf-

vernietigende percolatie proces niet bestaat. De eigenschap heeft ook gevolgen voor de bosbrandmodellen, die we bestuderen in hoofdstuk 3.

Zoals gezegd handelt hoofdstuk 3 over bosbrandmodellen. We bekijken een (eindig) vierkant bos, waar de roosterpunten door bomen bezet kunnen worden. Af en toe slaat de bliksem in waardoor een bosbrand ontstaat. We zijn het meest geïnteresseerd in het limietgedrag als de grootte van het systeem naar oneindig gaat, en de intensiteit van de blikseminslagen naar nul. Dit zou een benadering kunnen zijn van het permanent zelf-vernietigende proces (hoewel we in hoofdstuk 2 hebben gezien dat dat waarschijnlijk niet bestaat in dimensie 2). Fysici hebben bosbrandmodellen uitgebreid bestudeerd, hoewel de resultaten wiskundig gezien meestal discutabel zijn. Zij concentreren zich veelal op het gedrag van het bosbrandmodel in de evenwichtsverdeling. Wij bestuderen deze evenwichtsverdeling in dimensie 1, en berekenen expliciet de verdeling van de clustergrootte. In dimensie 2 concentreren we ons niet op de evenwichtsverdeling, maar op het gedrag van het bosbrandmodel als we beginnen met een leeg bos. Het interessante gedrag moet dan te zien zijn rond de kritieke tijd, corresponderend met het kritieke punt van gewone onafhankelijke percolatie. De bovengenoemde eigenschap impliceert voor de bosbrandmodellen dat na een grote brand, het een tijdje duurt voordat er opnieuw een grote brand voorkomt. Dit zou juist een teken zijn dat het twee-dimensionale bosbrandmodel *geen* zelf-organisered kritiek gedrag vertoont.

Hoofdstuk 4 gaat over bevroren percolatie. In tegenstelling tot bosbrandmodellen waar punten oneindig vaak van toestand kunnen wisselen, gebeurt dit in een bevroren percolatie model maximaal twee keer. In het bevroren percolatie model zijn punten van een willekeurige graaf in rust op tijd 0. Op een gegeven moment worden ze geactiveerd (dit gebeurd onafhankelijk van de andere punten) en zodra een punt in een cluster van oneindig veel actieve punten zit, bevriest het gehele cluster. Omdat daar oneindig veel punten bij betrokken zijn is het bestaan van bevroren percolatie modellen niet duidelijk. Aldous [Ald00] heeft laten zien dat een vorm van het bevroren percolatie proces bestaat als we de binaire boom als onderliggende graaf nemen. Het mooie van dit model is dat het zelf-organiserend kritiek gedrag vertoont: Op alle tijdstippen na de kritieke tijd zijn hebben de eindige niet-lege clusters dezelfde verdeling als kritieke onafhankelijke percolatie clusters. We zien ongeveer hetzelfde gedrag voor onze versie van bevroren percolatie. We kunnen de dynamica zo aanpassen dat clusters bevriezen zodra ze in een actief cluster van grootte N zitten. Dit maakt dat de modellen bestaan, en in dimensie 1 kunnen we zelfs concrete berekeningen doen.

In hoofdstuk 5 ligt de nadruk op directe toepassing. Het bespreekt een methode om bij benadering te trekken uit de monomer-dimer verdeling. Het is niet mogelijk exact te trekken uit deze verdeling voor grote of middelgrote systemen, omdat de rekentijd te lang wordt. We beschrijven een Markov-keten die de monomer-dimer verdeling als stationaire verdeling heeft. In dit hoofdstuk bepalen we wanneer er genoeg stappen in de Markov-keten gedaan zijn; met andere woorden, wanneer de verdeling bij benadering de stationaire verdeling is.