– 402 –

Ising models and multiresolution quad-trees

by

W.S. Kendall and R.G. Wilson

July 2002
Revised October 2002

DEPARTMENT OF STATISTICS
UNIVERSITY OF WARWICK

This is quadtree.tex 1.43 2002/11/28 10:23:37 Wilfrid W2k Wilfrid

Contents

List of Figures

Ising models and multiresolution quad-trees

W.S. Kendall* and R.G. Wilson*

November 28, 2002

Abstract: We study percolation and Ising models defined on generalizations of quad-trees used in multiresolution image analysis. These can be viewed as trees for which each mother vertex has 2^d daughter vertices, and for which daughter vertices are linked together in d-dimensional Euclidean configurations. Retention probabilities / interaction strengths differ according to whether the relevant bond is between mother and daughter, or between neighbours. Bounds are established which locate phase transitions and show the existence of a coexistence phase for the percolation model. Results are extended to the corresponding Ising model using the Fortuin-Kasteleyn random-cluster representation.

Keywords: Gibbs state; Ising model; finite island property; image analysis; Markov random field; multiresolution Markov random field; percolation; quad-tree; random-cluster model; segmentation; unique infinite cluster.
AMS Subject Classification (2000): 62M40, 60K35, 82B20

1 Introduction

This paper begins with affectionate birthday greetings to Professor Joseph Mecke from his friend and collaborator WSK. We hope that Professor Mecke will enjoy this account of an investigation into phase transition phenomena on a family of tessellations arising in an applied probability problem with a strongly geometric flavour.

Our aim here is to present a preliminary essay in an investigation of anisotropic Ising models on d-dimensional generalizations $\mathbb{Q}_d$ of the quad-tree structure used in image analysis. There is of course a long-established literature on the behaviour of Ising models on trees, stretching back to Preston [24] and Spitzer [28]. However we *augment* our trees by adding links between neighbours according to some Euclidean structure; the closest results in the literature are therefore those of Newman and Wu [22] concerning anisotropic Ising models on products

*Research supported by EPSRC research grant GR/M75785

of trees with Euclidean spaces. (See also [27, 32, 33], which investigate Ising models on planar transitive hyperbolic graphs.) Our methods borrow much from Newman and Wu, and the precursor paper by Grimmett and Newman [13] concerning percolation, but the special features of the quad-tree set-up require special arguments. In this introductory section we first carefully define the quad-tree structures to be considered, and then discuss how image segmentation motivates the study of behaviour of Ising models defined on these structures.

1.1 Definitions

We first present mathematical definitions of the structures of interest. We borrow from mathematical genetics its conventional terminology for vertices in tree-like structures (mothers, daughters, siblings, cousins).

Definition 1.1 *A* (doubly infinite) generalized quad-tree $\mathbb{Q}_d$ *is the union of a doubly infinite sequence of tessellations of square tiles at increasing levels of resolution:* $\ldots, L_{n-1}, L_n, L_{n+1}, \ldots$ *of* $\mathbb{R}^d$*. The tessellation* L_n *at* resolution level n *divides each square cell* L_{n-1} *of the previous tessellation into* 2^d *square sub-cells. We require that the tessellation* L_n *at resolution level* n *must be based on the cell* $[0, 2^{-n})^d$*. The generalized quad-tree is furnished with a graph structure; vertices are the cells of the various tessellations, with vertices connected by edges as follows: each tessellation cell* $C \in L_n$ *is linked to its immediate* neighbour *cells in the tessellation* L_n *and also to its immediate* mother *cell (the cell in the previous resolution level* L_{n-1} *which contains* C*) and its* daughter *cell (the cells contained in* C *which belong to the next higher resolution level* L_{n+1}*).*

In graph-theoretic terms a generalized quad-tree is an augmentation of a 2^d*-tree* $\mathbb{T}_d$ *(each vertex has* 2^d daughters *and one* mother*). The augmentation endows each vertex with links to* $2d$ neighbours *in a Euclidean configuration, so that the level* L_n *can be viewed as a (scaled) copy of the integer lattice* $\mathbb{Z}^d$*.*

Note that the generalized quad-tree $\mathbb{Q}_1$ is in fact an augmented binary tree.

We endow each resolution level L_n with the metric norm $\|\mathbf{v} - \mathbf{u}\|_{n,\infty} = 2^n \max_i |v_i - u_i|$, where $\mathbf{u}$ and $\mathbf{v}$ are the cells

$$[u_1, u_1 + 2^{-n}) \times \ldots \times [u_d, u_d + 2^{-n})$$
$$[v_1, v_1 + 2^{-n}) \times \ldots \times [v_d, v_d + 2^{-n})$$

respectively. This metric measures distance between two vertices on the same resolution level L_n in terms of separation in the Euclidean lattice $\mathbb{Q}_d \cap L_n$, and is therefore a useful notation when considering percolation issues relating to L_n alone.

For convenience of exposition, we often use the point $u = (u_1, \ldots, u_d)$ to represent the cell $[u_1, u_1 + 2^{-n}) \times \ldots \times [u_d, u_d + 2^{-n})$.

We emphasize that the generalized quad-tree $\mathbb{Q}_d$ is not homogeneous. For each vertex, d of its neighbours are siblings and d are only cousins. Measurement of the degree of consanguinity of the cousins (how far back to the most recent common ancestor) allows us to distinguish an infinite variety of different types of vertices on the basis of their immediate genealogy. The modelling of the graph in terms of successive tessellations is necessary in order to specify the full genealogy of specific vertices. Figure 1 illustrates this for $\mathbb{Q}_1$. Note in particular that the vertex representing the cell $[0, 1)$ is special in level L_0, in that it is not related at all in a tree-like sense to its left neighbour, the vertex representing the cell $[-1, 0)$.

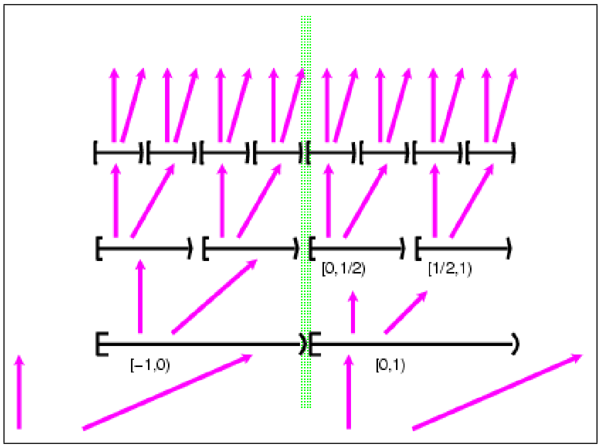

Figure 1: Illustration of $\mathbb{Q}_1$. We reverse the usual order of display for trees: daughters are placed *above* their mothers, to conform to the intuition that they are at a *higher* resolution level. The shaded bar emphasizes the fact that $[0, 1)$ has no tree-like connection to its neighbour $[-1, 0)$.

However there *is* a transitive $\mathbb{Z}^d$ symmetry on the subgraph of $\mathbb{Q}_d$ obtained by considering only the tessellations at resolution levels 0 and higher.

Definition 1.2 *The* (singly infinite) generalized quad-tree *is* $\mathbb{Q}_{d;0}$*: that part of the full generalized quad-tree which is at resolution level* 0 *or higher. We set*

$$\mathbb{Q}_{d;n} \quad = \quad \mathbb{Q}_d \cap (L_n \cup L_{n+1} \cup L_{n+2} \ldots) \ .$$

On $\mathbb{Q}_{d;0}$ *there is a* $\mathbb{Z}^d$*-action induced by the standard additive* $\mathbb{Z}^d$*-action on* $\mathbb{R}^d$*.*

In fact $\mathbb{Q}_{d;0} \cong \mathbb{Q}_{d;n}$ for any integer n; the following proposition describes graph isomorphisms which realize this.

Proposition 1.3 *Given* $\mathbf{o} \in L_0 \subset \mathbb{Q}_d$*, and* $\boldsymbol{u} \in L_n \subset \mathbb{Q}_d$*, there is a graph isomorphism* $\mathcal{S}_{\mathbf{o};\boldsymbol{u}} : \mathbb{Q}_{d;0} \to \mathbb{Q}_{d;n}$ *carrying* $\mathbf{o}$ *to* $\boldsymbol{u}$ *which preserves the generators of* $\mathbb{Z}^d$ *and the generalized quad-tree structure. We describe this in terms of a mapping on the underlying Euclidean space: if* u *is the vector representing the cell* $\boldsymbol{u} = [u_1, u_1 + 2^{-n}) \times \ldots \times [u_d, u_d + 2^{-n})$ *then* $\mathcal{S}_{\mathbf{o};\boldsymbol{u}}$ *can be represented by the affine-linear map*

$$\mathcal{S}_{\mathbf{o};\boldsymbol{u}}(x) \quad = \quad \boldsymbol{u} + 2^{-n} x \,.$$

Remark 1.4 *We define $\mathcal{S}_{\boldsymbol{u};\boldsymbol{v}}$ by $\mathcal{S}_{\boldsymbol{u};\boldsymbol{v}} \circ \mathcal{S}_{\mathbf{o};\boldsymbol{u}} = \mathcal{S}_{\mathbf{o};\boldsymbol{v}}$. Thus $\mathcal{S}_{\boldsymbol{u};\boldsymbol{v}} : \mathbb{Q}_{d;r} \to \mathbb{Q}_{d;s}$ is defined for $\boldsymbol{u} \in L_r$, $\boldsymbol{v} \in L_s$ for all integers r, s, and we can represent it by*

$$\mathcal{S}_{\boldsymbol{u};\boldsymbol{v}}(x) \quad = \quad v + 2^{r-s}(x-u)\,.$$

In particular, $\mathbb{Q}_{d;0}$ is *semi-transitive*: for any vertex $\mathbf{o}$ at the zero level of resolution and any other vertex $\mathbf{v} \in \mathbb{Q}_{d;0}$ (not necessarily in L_0) the map $\mathcal{S}_{\mathbf{o};\mathbf{u}}$ considered as a map into $\mathbb{Q}_{d;0}$ is a graph homomorphism carrying $\mathbf{o}$ into $\mathbf{v}$.

For the purposes of image analysis we are interested in that part of a generalized quad-tree formed by the descendants of a fixed root vertex $\mathbf{o}$ corresponding to the cell $[0,1)^d$, a "pyramid" subset.

Definition 1.5 *The* rooted generalized quad-tree *is* $\mathbb{Q}_d(\mathbf{o})$ *where*

$$\mathbb{Q}_d(\boldsymbol{u}) \quad = \quad \{v \in \mathbb{Q}_d \;:\; v \text{ is a descendant of } \boldsymbol{u}\}\;.$$

Note that

(a) $\mathbb{Q}_d(\mathbf{o}) \subseteq \mathbb{Q}_{d;0}$;

(b) $\mathbb{Q}_d(\mathbf{o}) \cong \mathbb{Q}_d(\mathbf{u})$ for any vertex $\mathbf{u} \in \mathbb{Q}_d$.

The root vertex $\mathbf{o}$ introduces extra inhomogeneity beyond the intrinsic inhomogeneity of the singly infinite generalized quad-tree; clearly the $\mathbb{Z}^d$-symmetry is destroyed. However (*cf* (b) above) $\mathbb{Q}_d(\mathbf{o})$ remains semi-transitive: for any vertex $\mathbf{u} \in \mathbb{Q}_d(\mathbf{o})$ there is a graph-homomorphism $\mathcal{S}_{\mathbf{o};\mathbf{u}}$ mapping $\mathbf{o}$ to $\mathbf{u}$.

An alternative take on this discussion of symmetry is to note it can be developed to show that the vertex set of $\mathbb{Q}_d$ can be viewed as a subset of $\mathbb{R}^+ \times \mathbb{R}^d$ which is invariant under a transformation group generated by d maps of the form $(x_0, \mathbf{x}) \mapsto (x_0, \mathbf{x} + \mathbf{e})$ and 2^d maps of the form $(x_0, \mathbf{x}) \mapsto (x_0/2, \mathbf{x} + (x_0/4)\mathbf{a})$; however the corresponding Cayley graph contains too many edges to be $\mathbb{Q}_d$ (spurious edges arise from inverses to the maps $(x_0, \mathbf{x}) \mapsto (x_0/2, \mathbf{x} + (x_0/4)\mathbf{a})$ – for each vertex just one of these inverses gives rise to a valid edge!).

Where feasible we describe results for general dimension d: however our main interest is in $d = 2$ and we will specialize to this case when necessary or convenient.

1.2 Image Segmentation

Why should we be interested in the peculiar structure of $\mathbb{Q}_d(\mathbf{o})$? Originally introduced to provide an efficient representation of binary image data [11], the structure of $\mathbb{Q}_d$ can be used to provide hierarchical models of simple images as follows: consider $\mathbb{Q}_d \cap (L_0 \cup L_1 \cup \ldots \cup L_{n+1})$ or perhaps its trace on $[0,1)^d \subset \mathbb{R}^d$. Let each of the cells in each of the levels have a *state* which is white or black. States in the boundary $\mathbb{Q}_d \cap L_{n+1}$ are prescribed using the image to be analyzed. States of other cells are modelled by a hierarchical Markov random field, which forms an Ising model on the graph structure of $\mathbb{Q}_d \cap (L_0 \cup L_1 \cup \ldots \cup L_{n+1})$; bond strengths depend on whether the bond is "space-like" (lies within a resolution level) or whether it crosses from one resolution level to the next.

The interest of this paper focuses on phase transition phenomena which are exhibited in the $n \to \infty$ limit.

All this relates to one of the fundamental problems in image analysis: *segmentation* [30]. This is the labelling of the several regions of more or less homogeneous properties of which a "typical" image, such as figure 2, is comprised. While sharing features with conventional classification, it has to contend with the obvious geometrical properties of images. At its simplest, this means that the class of a pixel (x, y) is treated as being dependent on those of its neighbours $(x \pm 1, y \pm 1)$. Clearly, this leads to the use of a Markov random field model for the label field following Geman and Geman [8]. Following this seminal paper, Markov random fields have gained significant attention in the segmentation of regions of more or less uniform colour or texture [8],[17], [20],[23],[29]. For example, Geman *et al.* [7] use the Kolmogorov-Smirnov non-parametric measure of difference between the distributions of spatial features extracted from pairs of blocks of pixel gray levels, with maximum *a posteriori* (MAP) estimation of the boundary, while Panjwani *et al.* [23] characterize textured colour images in terms of spatial interaction within and between colour planes.

Although they can be effective as models for segmentation, two-dimensional Markov random fields have weaknesses as models of images. In the first place, the typical image consists of a relatively small number of regions, something which is not captured by the equilibrium distribution of the planar discrete Markov random field prior. Although we are more interested in the posterior than the prior, it is discouraging to find that even in the simplest case (the Ising model) no value of coupling parameter will lead to a single, well defined object on a background. Secondly, the high coupling strengths needed to capture large scale structure imply a heavy computational burden in reaching equilibrium of the posterior in many cases.

We can finesse the problem by exploiting a fundamental property of many images, which we may loosely call 'scale invariance'; it is illustrated in Figure 3, which shows that across a range of scales (or more accurately resolutions), image content remains perceptibly 'the same', in terms of the underlying region structure. In this image, each level has a sampling density 4 times that of the next larger scale: there are 4^n samples per unit area on level n. This informal observation, which underlies the wavelet representations now so popular in image processing [19], leads us to conclude that within the corresponding lattice structure, which is a *quad-tree*, neighbouring pixels again should be modelled as having strongly dependent labels. We are thereby led to consider Markov random fields defined on quad-trees, in the expectation that if we can solve the segmentation problem at a low resolution (and low computation), we may use this to 'steer' the solutions at successively higher resolutions. For example, Bouman and Shapiro use sequential maximum a posteriori (SMAP) estimation in conjunction with a multi-scale random field (MSRF) [2], a sequence of random fields at different scales. While such *pure* quad-tree processes can lead to fast 'scale-causal' algorithms for segmentation, they ignore the fundamental translation symmetry of images, since each layer of the tree has a sampling interval half that of the layer above it in resolution. This deficiency manifests itself in priors which favour 'blocky' images, with obvious effects on the estimates.

Consequently in recent years there have been several attempts to implement estimation algorithms based on a model using the full set of neighbours on the quad-tree: four or eight on the same resolution level, combined with one or more *parents* and four or more *children* from adjacent levels. Both deterministic and stochastic methods have been used to find a MAP labelling [15, 16, 21, 18]. This brings us to our quad-tree $\mathbb{Q}_2(\mathbf{o})$, which serves as a good representative of the extension of these models to infinite resolution levels.

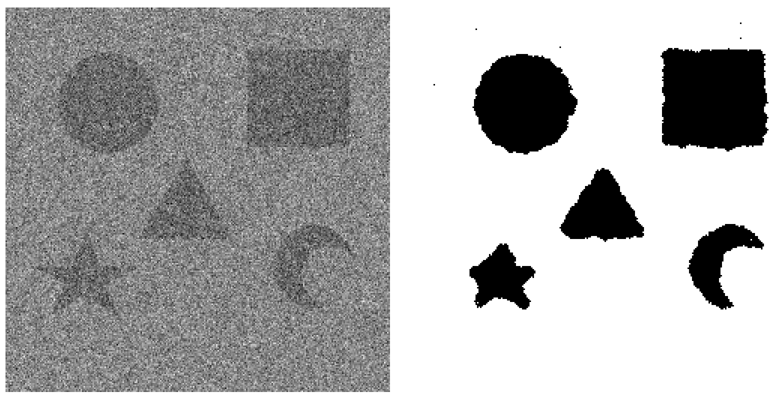

(a) Noisy image. (b) Labelling.

Figure 2: Segmentation of noisy 'Shapes' image. Noise standard deviation is equal to difference between object and background grey levels.

1.3 Image Segmentation Results

A number of experiments have demonstrated that quad-tree models can perform well in image segmentation. For example, consider the image in Figure 2(a), a binary image with added independent mean-zero Gaussian noise variates with unit standard deviation (thus standard deviation is chosen to equal the difference between object and background intensities). The prior model used a second order neighbourhood and a normal observation model. Coupling parameters were tuned by hand and the class means and variances estimated from the data during processing. The estimate at the highest resolution, shown in figure 2(b), using a multiresolution MAP estimation algorithm, [31], had a misclassification rate of 1.3%. Errors of this order have been found across a wide variety of image data using this model [31]. Interestingly, similar error rates are also found in the segmentation of textures having a cell size significantly larger than the error bound; this is a direct consequence of using the multiresolution model. Our paper is motivated by these results. In particular, we wished to understand what one might be able to say about phase transition phenomena which might throw light on simulation results carried out on finite quad-trees with large numbers of layers, and thus permit a better grasp of the way in which practical quad-tree algorithms might behave.

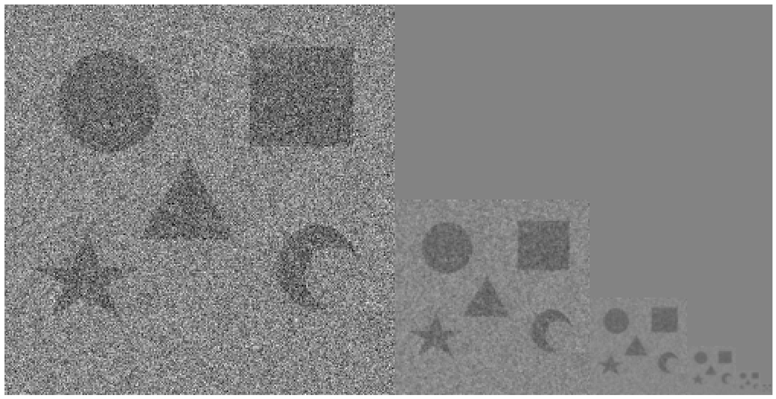

Figure 3: Quad-tree data formed by successive averaging and decimation operations: grey level at a pixel on level N is average of 4 on level $N+1$.

1.4 Plan of paper

The strategy of this paper is firstly to establish phase transition results for anisotropic percolation models, in the following Section 2, and then in Section 3 to derive results for the Ising model by use of the famous Fortuin-Kasteleyn random cluster representation [6, 4, 5]. In conclusion, Section 4 describes some illustrative simulations and discusses prospects for further work.

Acknowledgements

One of us (WSK) is very grateful to his colleague Jonathan Warren for encouragement and discussion.

2 Percolation on generalized quad-trees

Grimmett and Newman [13] introduced (anisotropic) percolation on a transitive non-euclidean graph (in their case, the cartesian product $\mathbb{T}_k \times \mathbb{Z}$ of a regular k-tree with the one-dimensional Euclidean lattice, and also $\mathbb{T}_k \times \mathbb{Z}^d$) and showed the existences of *three* phases:

(a) no infinite clusters;

(b) many (indeed, infinitely many) infinite clusters (the *coexistence* phase);

(c) a single unique infinite cluster.

Much work has followed up both on this and also on the seminal paper of Benjamini and Schramm [1], which describes results and questions on "percolation beyond $\mathbb{Z}^d$", typically in the context of "isotropic" percolation (retention probability parameter the same for each bond) on almost transitive non-amenable graphs. (Note also, Newman and Wu [22] use Fortuin-Kasteleyn random cluster models and the results of [13] to draw out results about anisotropic Ising models on $\mathbb{T}_k \times \mathbb{Z}$, and their procedure will be a template for our work in Section 3.)

Our model for percolation on $\mathbb{Q}_d$ is as follows: we suppose independent bond percolation, with retention probability being λ for bonds between neighbours and τ for bonds between mothers and daughters. This percolation model has non-constant retention probabilities and is not almost-transitive, so the proofs in the literature must be adapted accordingly, sometimes with non-trivial effort. We shall build up schematic phase-transition diagrams for $\mathbb{Q}_{d;0}$-percolation and $\mathbb{Q}_d(\mathbf{o})$-percolation, largely by estimation of locations of phase-transition boundaries in λ-τ space for cases when either λ or τ is small.

2.1 Transition from zero to many infinite clusters for small λ

The work of [13, §3 and §5] applies directly here. We sketch the argument for the sake of completeness.

When deciding whether there can be an infinite cluster, it suffices to consider $\mathbb{Q}_d$.

Theorem 2.1 *There is almost surely no infinite cluster in $\mathbb{Q}_d$ (and consequently in $\mathbb{Q}_{d;0}$, $\mathbb{Q}_d(\mathbf{o})$) if*

$$2^d \tau \mathcal{X}_\lambda \left(1 + \sqrt{1 - \mathcal{X}_\lambda^{-1}}\right) \quad < \quad 1\,,$$

where $\mathcal{X}_\lambda$ is the mean size of the percolation cluster at the origin for bond percolation in $\mathbb{Z}^d$ with bond retention probability λ.

Proof: Consider the mean size of the cluster at $\mathbf{o}$. This is given by

$$\sum_{\mathbf{z}\in\mathbb{Q}_d} \mathbb{P}\left[\,\mathbf{o} \leftrightarrow \mathbf{z}\,\right]\,.$$

Arguing as in the proof of [13, Proposition 1], consideration of open self-avoiding paths shows this is bounded above by

$$\begin{aligned} &\sum_{n=0}^{\infty} \sum_{\underline{\mathbf{t}}:|\underline{\mathbf{t}}|=n} \sum_{\mathbf{o}=\mathbf{z}'_1,\mathbf{z}_1,\ldots,\mathbf{z}'_n,\mathbf{z}_n} \mathbb{P}\left[\,\mathbf{z}'_1 \leftrightarrow \mathbf{z}_1 \text{ in } L_{t_1}\,\right] \times \ldots \times \mathbb{P}\left[\,\mathbf{z}'_n \leftrightarrow \mathbf{z}_n \text{ in } L_{t_n}\,\right] \\ \leq\; & \sum_{n=0}^{\infty} \sum_{\underline{\mathbf{t}}:|\underline{\mathbf{t}}|=n} \mathcal{X}_\lambda \tau^n (\mathcal{X}_\lambda - 1)^{T(\underline{\mathbf{t}})} \mathcal{X}_\lambda^{n-T(\underline{\mathbf{t}})} \;\leq\; \sum_{n=0}^{\infty} \mathcal{X}_\lambda (\tau\mathcal{X}_\lambda)^n \sum_{\underline{\mathbf{t}}:|\underline{\mathbf{t}}|=n} (1-\mathcal{X}_\lambda^{-1})^{T(\underline{\mathbf{t}})}\,. \end{aligned}$$

Here the summation $\sum_{\underline{\mathbf{t}}:|\underline{\mathbf{t}}|=n}$ runs through all length-n sequences $\underline{\mathbf{t}}$ of choices of mother/daughter bonds making up self-avoiding paths starting at $\mathbf{o}$. For each

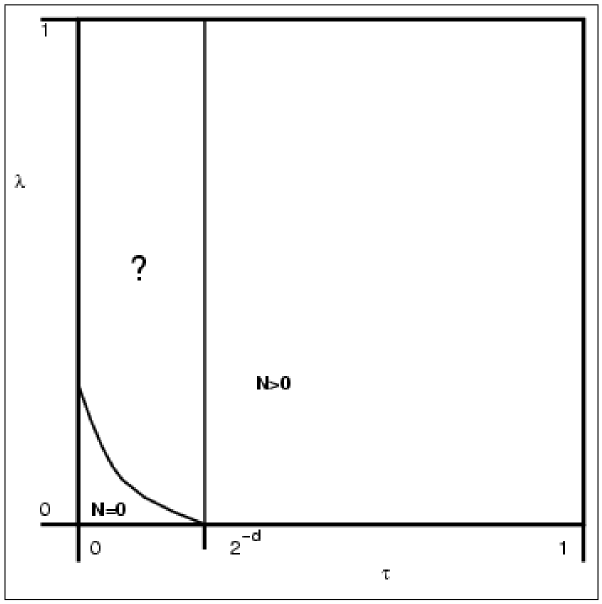

Figure 4: Partial information on existence of infinite clusters: case of small λ. Here N denotes the number of infinite clusters. This does not yet represent good information for large λ and small τ.

---

such sequence representing a possible connection from $\mathbf{o}$ to another location, we count the expected number choices of vertices $\mathbf{o} = \mathbf{z}'_1, \mathbf{z}_1, \ldots, \mathbf{z}'_n, \mathbf{z}_n$ marking the start and end of moves which change level L_i which together could make up an open self-avoiding path (hence requiring each $\mathbf{z}'_i$, $\mathbf{z}_i$ to belong to the same cluster in the $\mathbb{Z}^d$ percolation occurring at their common level). Furthermore, $T(\underline{\mathbf{t}})$ counts the number of times in $\underline{\mathbf{t}}$ that a daughter step (increasing resolution) is directly followed by a mother step (decreasing resolution). In such a case there is at least one state in the relevant $\mathbb{Z}^d$-cluster which has already been counted, and we account for this in the upper bound. Developing this further,

$$\sum_{n=0}^{\infty} \mathcal{X}_\lambda (\tau \mathcal{X}_\lambda)^n \sum_{\underline{\mathbf{t}}:|\underline{\mathbf{t}}|=n} (1-\mathcal{X}_\lambda^{-1})^{T(\underline{\mathbf{t}})} \;\leq\; \sum_{n=0}^{\infty} \mathcal{X}_\lambda (2^d \tau \mathcal{X}_\lambda)^n \sum_{\underline{\mathbf{j}}:|\underline{\mathbf{j}}|=n} (1-\mathcal{X}_\lambda^{-1})^{T(\underline{\mathbf{j}})}$$
$$\approx \sum_{n=0}^{\infty} \mathcal{X}_\lambda (2^d \tau \mathcal{X}_\lambda)^n \left(1 + \sqrt{1-\mathcal{X}_\lambda^{-1}}\right)^n .$$

Here the summation $\sum_{\underline{\mathbf{j}}:|\underline{\mathbf{j}}|=n}$ runs through all length-n sequences $\underline{\mathbf{j}}$ of choices of mother-step versus daughter-step (omitting which kind of daughter), and again $T(\underline{\mathbf{j}})$ counts the number of times in $\underline{\mathbf{j}}$ that a daughter-step is followed by a mother-step. The (large n) asymptotic arises from a spectral analysis of the matrix representation

$$\sum_{\underline{\mathbf{j}}:|\underline{\mathbf{j}}|=n}(1-\mathcal{X}_\lambda^{-1})^{T(\underline{\mathbf{j}})} \quad = \quad \begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1-\mathcal{X}_\lambda^{-1} & 1 \end{bmatrix}^n \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

□

With more work one can derive an upper bound for this phase transition valid for small λ, by considering the effect of adding a single open λ-bond in $\mathbb{Q}_d(\mathbf{o}) \cap (L_0 \cup L_1 \cup \ldots \cup L_n)$. This is discussed for the case $d = 2$ in Appendix A.

Comparison with the representation of $\lambda = 0$ percolation by the branching process of family size distribution Binomial$(2^d, \tau)$ shows that this bound is sharp at $\lambda = 0$ (when $\mathcal{X}_\lambda = \mathcal{X}_0 = 1$). Moreover a simple monotonicity argument shows that there must be infinite clusters at (τ, λ) whenever there are infinite clusters at $(\tau - \epsilon, \lambda - \delta)$ for non-negative ϵ, δ. This partial information is represented in Figure 4.

The situation for small τ is more involved. We expect infinite clusters once $\lambda > \lambda_c$, where λ_c is the critical retention probability parameter for $\mathbb{Z}^d$ bond percolation. However for a full argument we need to consider $\mathbb{Q}_d(\mathbf{o})$, and the resolution levels $\mathbb{Q}_d(\mathbf{o}) \cap L_n$ are actually *finite*, so the infinite-cluster property cannot simply be inherited from the $\mathbb{Z}^d$ case. (This is in contrast to the Grimmett-Newman setting of $\mathbb{T}_k \times \mathbb{Z}^d$.) There is however an argument based on existence of a *unique* infinite cluster once $\lambda > \lambda_c$, and we now turn to this.

2.2 Transition from many to unique infinite clusters for small τ

It suffices to consider the "pyramid" case $\mathbb{Q}_d(\mathbf{o})$. Here we have to confront the special features of our model, specifically the finiteness of the resolution layers. We restrict to the case $d = 2$, in order to employ planar duality when analyzing the resolution layers.

Theorem 2.2 *Consider the sequence of finite $\mathbb{Z}^2$ percolation problems $\mathbb{Q}_2(\mathbf{o}) \cap L_n$ for $n = 0, 1, \ldots$, when $\lambda > \lambda_c(2) = 1/2$ and $\tau > 0$. Choose $\epsilon > 0$ and set*

$$\ell_n \quad = \quad (n \log 4 + (2+\epsilon)\log n)\,\xi(1-\lambda)\,, \tag{2.1}$$

where $\xi(1-\lambda)$ is the exponent in the nontrivial exponential bound for the connectivity function for the sub-critical bond percolation problem in $\mathbb{Z}^2$ with retention probability $1-\lambda$ [12, Theorem 6.44]. Then, almost surely, for all sufficiently large n there is exactly one $\mathbb{Z}^2$ percolation cluster in $\mathbb{Q}_2(\mathbf{o}) \cap L_n$ containing points separated by ℓ_n or more in the uniform norm.

Proof: We rescale for the sake of notational convenience, and represent $\mathbb{Q}_2(\mathbf{o}) \cap L_n$ as $\{0, 1, \ldots, 2^n - 1\}^2$.

First we show there is at most one such cluster for all large enough n. For there to be more than one such cluster at resolution level n, there must exist a separating self-avoiding path in the dual bond percolation problem for $\{0, 1, \ldots, 2^n - 1\}^2$ with edge retention probability $1 - \lambda$, and this must be in one of four possible forms:

(1) connecting opposite sides of $\{0, 1, \ldots, 2^n - 1\}^2$;

(2) connecting neighbouring sides of $\{0, 1, \ldots, 2^n - 1\}^2$ and separating two pairs $\mathbf{a}$, $\mathbf{b}$ and $\mathbf{c}$, $\mathbf{d}$ of vertices such that $\|\mathbf{a} - \mathbf{b}\|_{n,\infty}$, $\|\mathbf{c} - \mathbf{d}\|_{n,\infty}$ both exceed ℓ_n;

(3) attached twice to one side and separating two pairs of vertices as before;

(4) closed path not attached to sides at all, and separating two pairs of vertices as before.

All of these cases entail existence of at least one path in the dual bond percolation problem stretching between endpoints of separation ℓ_n or more (once n is large enough so that $\ell_n < 2^n$).

Now we apply the exponential bound on the connectivity problem for the dual percolation problem:

$$\begin{aligned}
\mathbb{P}[\text{ dual percolation path as above }] &\leq \sum\sum_{\mathbf{v},\mathbf{u}:\|\mathbf{v}-\mathbf{u}\|=\ell_n} \exp\left(-\frac{\ell_n}{\xi(1-\lambda)}\right) \\
&\leq 4^n \times (4 \times 2\ell_n) \times \exp\left(-\frac{\ell_n}{\xi(1-\lambda)}\right) \\
&\leq 8\ell_n \exp\left(n \log 4 - \frac{\ell_n}{\xi(1-\lambda)}\right) .
\end{aligned}$$

This is summable if (2.1) holds, and so the first Borel-Cantelli lemma may be applied to show that almost surely for all large enough n there will not be more than one of these large clusters in each resolution level.

The existence of such a cluster for all sufficiently large n follows by using the super-criticality of λ-percolation in $\mathbb{Z}^2$: there is therefore a positive lower bound $\theta(\lambda)$ that $(0, 0)$ lies in the infinite cluster for this percolation problem. We choose $4^{n-[n/2]}$ disjoint rectangles in $\{0, 1, \ldots, 2^n - 1\}^2$, each of side-length $2^{[n/2]}$. A point in the centre of such a rectangle has probability at least $\theta(\lambda)$ of percolating to the boundary, and thus establishing an open path between two points separated in uniform norm by $2^{[n/2]}/2$. The probability that none of these $4^{n-[n/2]}$ percolations occur is bounded above by

$$(1 - \theta(\lambda))^{4^{n-[n/2]}}$$

which is summable in n. It follows that almost surely all but finitely many levels must exhibit examples of such a percolation. Since $2^{[n/2]}/2 \geq \ell_n$ for all large enough n it is almost surely the case that the required large cluster must exist for all sufficiently large n. $\square$

We can now apply this to deduce uniqueness (and existence!) of an infinite percolation cluster in $\mathbb{Q}_2(\mathbf{o})$ for $\lambda > \lambda_c(2) = 1/2$ whenever $\tau > 0$. This is illustrated in Figure 5.

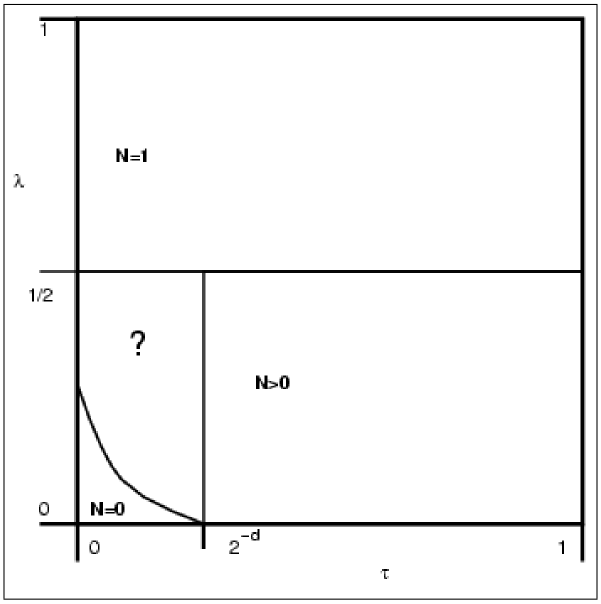

Figure 5: Partial information on uniqueness of infinite clusters, case $d = 2$. Here N denotes the number of infinite clusters. This does not yet represent good information on uniqueness for small λ and large τ.

Theorem 2.3 *Consider percolation in $\mathbb{Q}_2(\mathbf{o})$ for $\lambda > \lambda_c(2) = 1/2$ and $\tau > 0$. Almost surely there is a single infinite cluster.*

Proof: We first show existence of an infinite cluster. Set

$$\ell_n \quad = \quad (n \log 4 + (2+\epsilon)\log n)\,\xi(1-\lambda)$$

as in Equation (2.1), Theorem 2.2. Choose k_n increasing in n such that $2^{2k_n} > 2\ell_n$, and locate 4^{n-2k_n} points in $\mathbb{Q}_2(\mathbf{o}) \cap L_n$ separated horizontally and vertically by 2^{2k_n} and offset horizontally and vertically by 2^{k_n}. Let A_n be the event that none of these 4^{n-2k_n} points $\mathbf{v}$ satisfies the following; that the mother/daughter bond between $\mathbf{v}$ and a specific daughter $\mathbf{v}'$ is open, and $\mathbf{v}$ is also connected in $\mathbb{Q}_2(\mathbf{o}) \cap L_n$ to the outside of a square of side-length 2^{k_n} centred on $\mathbf{v}$, and $\mathbf{v}'$ is also connected in $\mathbb{Q}_2(\mathbf{o}) \cap L_{n+1}$ to the outside of a square of side-length $2^{k_{n+1}}$ centred on $\mathbf{v}'$.

We may deduce, using independence of percolation in the different squares of side-length 2^{k_n},

$$\mathbb{P}\left[\, A_n \,\right] \quad \leq \quad (1 - \theta(\lambda)\tau\theta(\lambda))^{4^{n-k_n}}$$

where $\theta(\lambda) > 0$ is the probability that $(0, 0)$ lies in the infinite cluster for λ-percolation in $\mathbb{Z}^2$.

Since this is summable ($4^{n-k_n} > 4^n/(2\ell_n)^2$), the first Borel-Cantelli lemma implies that A_n will happen for only finitely many n.

Combining this with Theorem 2.2, we see that in all but finitely many levels L_n we may find points $\mathbf{v}_n \in \mathbb{Q}_2(\mathbf{o}) \cap L_n$ which are

(a) connected to the large cluster in $\mathbb{Q}_2(\mathbf{o}) \cap L_n$ which is guaranteed eventually to exist and to be unique by Theorem 2.2;

(b) connected by a single mother/daughter bond to a daughter $\mathbf{v}'_n$ in $\mathbb{Q}_2(\mathbf{o}) \cap L_{n+1}$;

(c) such that this daughter is again connected to the large cluster in $\mathbb{Q}_2(\mathbf{o}) \cap L_{n+1}$ which is also guaranteed eventually to exist and to be unique by Theorem 2.2.

This allows us to deduce that the large clusters guaranteed by Theorem 2.2 must eventually all be connected to each other. This provides the required infinite cluster in $\mathbb{Q}_2(\mathbf{o})$.

We now show uniqueness of this infinite cluster. Pick an initial $\mathbf{v}_n \in \mathbb{Q}_2(\mathbf{o}) \cap L_n$ and let C_0 be the cluster in $\mathbb{Q}_2(\mathbf{o})$ which contains $\mathbf{v}_n$. Let K be the highest resolution level reached by C_0. Define $\mathbf{v}_{n+1}, \mathbf{v}_{n+2}, \ldots, \mathbf{v}_{n+K}$ in C_0 to lie in levels $L_{n+1}, L_{n+2}, \ldots, L_{n+K}$, choosing in such a way that $\mathbf{v}_{n+r+1}$ is measurable with respect to the σ-algebra $\mathfrak{F}_{n+r}$ generated by the neighbour and mother/daughter bonds involving at least one vertex in $\mathbb{Q}_2(\mathbf{o}) \cap (L_0 \cup \ldots \cup L_{n+r})$.

Note that the random resolution level K can be viewed as an optional time with respect to the filtration $\{\mathfrak{F}_{n+1} : n \geq 0\}$, in the sense that $[K \leq n] \in \mathfrak{F}_{n+1}$.

Then

$$\mathbb{P}\left[\,\mathbf{v}_{n+r+1} \text{ percolates at least } \ell_{n+r+1} \text{ into } \mathbb{Q}_2(\mathbf{o}) \cap L_{n+r+1} | \mathfrak{F}_{n+r}\,\right] \;\geq\; \theta_+(\lambda)\,,$$

where $\theta_+(\lambda)$ is the probability that $(0, 0)$ is part of an infinite cluster for λ-percolation in the orthant $\mathbb{Z}^2_+$. (That this is positive for $\lambda > \lambda_2(2) = 1/2$ follows from the exponential bound on the connectivity function for the dual $(1 - \lambda)$-percolation in $\mathbb{Z}^2$, by estimation of the mean number of pairs of boundary vertices $(x, 0)$ and $(0, y)$ which are connected to each other in the dual percolation.) We therefore deduce

$$\mathbb{P}\left[\,\mathbf{v}_{n+r+1} \text{ percolates at least } \ell_{n+r+1} \text{ for some } r\,\right] \;\geq\; 1 - \mathbb{E}\left[\,(1 - \theta_+(\lambda))^K\,\right]\,.$$

Using Theorem 2.2, we deduce that if C_0 is infinite (so K is infinite) then it must be connected to, and so equal, the infinite cluster whose existence was established in the first part of the proof. Hence the infinite cluster in $\mathbb{Q}_2(\mathbf{o})$ (for $\lambda > \lambda_c > 1/2$) is unique. $\square$

Similar techniques work for the case of $\mathbb{Q}_{2;0}$, but eased by the infinite nature of the layers $L_0, L_1, \ldots$:

Corollary 2.4 *Consider percolation in $\mathbb{Q}_{2;0}$ for $\lambda > \lambda_c(2) = 1/2$ and $\tau > 0$. Almost surely there is a single infinite cluster.*

2.3 Transition from many to unique infinite clusters for small λ

Up to this point we have not yet shown that there is any region where there is definitely a coexistence phase, with many (we expect, infinitely many) infinite clusters. It is possible to adapt to our problem the branching random walk comparison used in Benjamini and Schramm [1, Theorem 4], which shows that for small λ there is an interval of τ for which coexistence holds. A more effective argument of Grimmett and Newman [13, §4 and §5] applies to the special case of $\mathbb{T}_k \times \mathbb{Z}^d$; it applies as well to $\mathbb{Q}_{d;0}$. Essentially one modifies the proof of Theorem 2.1 above to bound the probability of a path running between specified vertices at resolution zero in $\mathbb{Q}_{d;0}$. The modification is a simple matter of observing that in such a case there must be exactly as many decrements as increments in resolution. There is just one τ-bond from any vertex which leads to a decreased resolution, so the argument in the proof of Theorem 2.1 can be modified to replace $2^d\tau\mathcal{X}_\lambda$ by $\sqrt{2^d}\tau\mathcal{X}_\lambda$. The consequent deduction is that if $\tau \in (2^{-d}, 2^{-d/2})$ and if λ is small enough then the probability of connection between two such vertices tends to zero as the vertices separate. Arguing as in the start of the proof of Theorem 2.8 below, this prohibits formation of a unique infinite cluster.

Here we establish a larger range of confirmed coexistence phase using a different argument: the gain is slight for large dimension d but significant for $d = 1, 2$. First of all, we establish a preliminary lemma concerning the behaviour of the $\mathcal{S}_{\mathbf{u};\mathbf{v}}$ maps of Proposition 1.3 and Remark 1.4. Let $\mathcal{M}(\mathbf{u})$ denote the mother of vertex $\mathbf{u}$.

Lemma 2.5 *Consider $\boldsymbol{u} \in L_{s+1} \subset \mathbb{Q}_d$ and $\boldsymbol{v} = \mathcal{M}(\boldsymbol{u}) \in L_s \subset \mathbb{Q}_d$. There are exactly 2^d solutions to*

$$\mathcal{M}(\boldsymbol{x}) \quad = \quad \mathcal{S}_{\boldsymbol{u};\boldsymbol{v}}(\boldsymbol{x})$$

which lie in L_{s+1}. They are characterized as follows: one solution is of course $\boldsymbol{x} = \boldsymbol{u}$. The other solutions are given by the remaining $2^d - 1$ vertices $\boldsymbol{y}$ such that the closure of the cell representing $\boldsymbol{y}$ intersects the vertex shared by the closures of the cells representing $\boldsymbol{u}$ and $\mathcal{M}(\boldsymbol{u})$. Finally, if $\boldsymbol{x} \in L_{s+1}$ does not solve $\mathcal{M}(\boldsymbol{x}) = \mathcal{S}_{\boldsymbol{u};\boldsymbol{v}}(\boldsymbol{x})$ then

$$\|\mathcal{S}_{\boldsymbol{u};\boldsymbol{v}}(\boldsymbol{x}) - \mathcal{S}_{\boldsymbol{u};\boldsymbol{v}}(\boldsymbol{u})\|_{s,\infty} \quad > \quad \|\mathcal{M}(\boldsymbol{x}) - \mathcal{M}(\boldsymbol{u})\|_{s,\infty}\,. \tag{2.2}$$

Proof: Since the inequality concerns an L^∞-norm, we can work coordinate-by-coordinate and so reduce to the case $d = 1$. Without loss of generality take $\mathbf{v}$ to represent the cell $[0, 1)$ at resolution level L_0, so that $\mathbf{u}$ represents either $[0, \frac{1}{2})$ or $[\frac{1}{2}, 1)$ at resolution level L_1. In either case $\mathcal{M}(x) = \lfloor x \rfloor$ as x runs through $\{\ldots, -\frac{3}{2}, -1, -\frac{1}{2}, 0, \frac{1}{2}, 1, \frac{3}{2} \ldots\}$ (the half-integers representing cells in L_1) where $\lfloor x \rfloor$ denotes the greatest integer part of x.

Case $[0, \frac{1}{2})$: Then $\mathcal{S}_{\mathbf{u};\mathbf{v}}(x) = 2x$. The equation $\lfloor x \rfloor = 2x$ is solved among $x \in \{\ldots, -\frac{3}{2}, -1, -\frac{1}{2}, 0, \frac{1}{2}, 1, \frac{3}{2} \ldots\}$ by $x = 0$ or $x = -\frac{1}{2}$. Otherwise for such x we have

$$\lfloor x \rfloor > 2x \qquad \text{if } x = \ldots, -\frac{3}{2}, -1\,;$$
$$\lfloor x \rfloor < 2x \qquad \text{if } x = \frac{1}{2}, 1, \frac{3}{2}, \ldots\,.$$

Case $[\frac{1}{2}, 1)$: Then $\mathcal{S}_{\mathbf{u};\mathbf{v}}(x) = 2(x - \frac{1}{2}) = 2x - 1$. The equation $\lfloor x \rfloor = 2x - 1$ is solved among $x \in \{\ldots, -\frac{3}{2}, -1, -\frac{1}{2}, 0, \frac{1}{2}, 1, \frac{3}{2} \ldots\}$ by $x = \frac{1}{2}$ or $x = 1$. Otherwise for such x we have

$$\lfloor x \rfloor > 2x - 1 \qquad \text{if } x = \ldots, -\frac{3}{2}, -1, -\frac{1}{2}, 0\,;$$
$$\lfloor x \rfloor < 2x - 1 \qquad \text{if } x = \frac{3}{2}, 2, \ldots\,.$$

In both cases the criterion of the lemma characterizes the solution set, and Inequality 2.2 holds off the solution set. $\square$

We add a corollary which will be useful during the proof of the main theorem of this sub-section:

Corollary 2.6 *Suppose now that we are given* distinct $\boldsymbol{v}$ *and* $\boldsymbol{y}$ *in the same resolution level of* $\mathbb{Q}_d$*, and we wish to enumerate the pairs of vertices* $\boldsymbol{u}$*,* $\boldsymbol{x}$ *in the resolution level one step higher and such that*

(a) $\mathcal{M}(\boldsymbol{u}) = \boldsymbol{v}$*;*

(b) $\mathcal{M}(\boldsymbol{x}) = \boldsymbol{y}$*;*

(c) $\mathcal{S}_{\boldsymbol{u};\boldsymbol{v}}(\boldsymbol{x}) = \boldsymbol{y}$*.*

There are at most 2^{d-1} *such vertices.*

Proof: From Lemma 2.5 the closure of the cell representing $\mathbf{u}$ must intersect the intersection of the closures of the cells representing $\mathbf{v}$ and $\mathbf{y}$. Since $\mathbf{v}$ and $\mathbf{y}$ are distinct, there can be at most 2^{d-1} such $\mathbf{u}$. Moreover, once $\mathbf{u}$ is specified then its counterpart $\mathbf{x}$ is specified by the condition characterizing $\mathcal{S}_{\mathbf{u};\mathbf{v}}(\mathbf{x}) = \mathbf{y}$ given in Lemma 2.5. $\square$

Remark 2.7 *From Lemma 2.5, the number of pairs described in Corollary 2.6 must vanish unless there is a non-void intersection between the closures of the cells representing* $\boldsymbol{v}$ *and* $\boldsymbol{y}$*.*

Theorem 2.8 *Consider* (λ, τ) *percolation in* $\mathbb{Q}_{d;0}$*. If* $\tau < 2^{-(d-1)/2}$ *and if* $\lambda > 0$ *is sufficiently small then almost surely there can be no unique infinite cluster.*

Proof: Our strategy is to bound the mean number of open self-avoiding paths in $\mathbb{Q}_{d;0}$ running from one vertex $\mathbf{u}$ to another $\mathbf{v}$, both located at at the same resolution level. The tree-like nature of the graph $\mathbb{Q}_{d;0}$ forces most self-avoiding paths which travel to high resolutions to possess a large number of λ-bonds, and thus to have low probability of being open. We will therefore be able to show that the mean number of such paths converges to zero as the distance between $\mathbf{v}$ and $\mathbf{u}$ increases. This is sufficient to rule out the chance of a unique infinite cluster in $\mathbb{Q}_{d;0}$; for otherwise any two vertices generating infinite clusters would have to be interconnected, so that by the FKG inequality the probability of connection would be bounded below by the square of the probability of belonging to an infinite cluster. Using semi-transitivity, we deduce a positive lower bound to this probability so long as infinite clusters are at all possible.

It is convenient to take an algebraic approach, encoding paths as *words* in the symbols δ (for a single step from daughter to mother), $\ell_{\pm i}$ for $i = 1, \ldots, d$ (encoding the $2d$ possible single steps to neighbouring vertices at the same resolution level), and u_x for $x \in \{\pm 1\}^d$ (encoding the 2^d different ways to make a single step from mother to daughter). We consider paths restricted to non-negative resolution levels of $\mathbb{Q}_{d;0}$, and therefore restrict attention to *up-words*, those words for which totals $\# u_x$, $\#\delta$ of various symbols obey

$$\sum_x \# u_x \quad \geq \quad \#\delta$$

for the word itself and also for all initial segments of the word. Moreover, because we consider paths starting at and returning to the zero resolution level L_0 of $\mathbb{Q}_{d;0}$, among up-words we consider the *bridge-words*, those up-words for which

$$\sum_x \# u_x \quad = \quad \#\delta$$

when calculated for the whole word. The *weight* of a bridge-word is simply the probability that a corresponding path is open:

$$\lambda^{\sum_i \#\ell_i} \tau^{\#\delta + \sum_x \# u_x} \quad = \quad \lambda^{\sum_i \#\ell_i} \tau^{2\#\delta} \,.$$

Borrowing from the vocabulary of the analysis of Brownian paths, we decompose a bridge-word (equivalently, the corresponding path) into a family of *excursions*, where an excursion word is a sub-word which is itself a bridge-word starting with a u_x for some x, ending with a δ, and delivering a path with does not revisit the starting resolution level before its end. Note that the family of excursions for a given bridge-word can be viewed as a tree $\mathcal{T}$, under the relationship "excursion ζ_1 is a daughter of excursion ζ_2 if the sub-path used to derive ζ_1 is actually a sub-path of the sub-path used to derive ζ_2" (excursions from L_0 being viewed as daughters of a virtual excursion ζ_0). We label the tree of excursions by attaching to each vertex of $\mathcal{T}$ the start u_x of the relevant excursion.

Consider the possibility of *demoting* a bridge-word by taking one of its excursions and removing the excursion start u_x and the excursion end δ. The result is still a bridge-word, and the tree of excursions of the new bridge-word is obtained from the old tree simply by removing the vertex ζ corresponding to the excursion used in the demotion, and transferring the daughters of ζ to be daughters of the mother of ζ. In general the path produced by the new bridge-word will not have the same end-point as the old bridge-word, unless the excursion subject to demotion possesses a particular property which we now describe.

Observe that we can describe demotion as follows. Let $\mathbf{u}$, $\tilde{\mathbf{u}}$, $\ldots$, $\tilde{\mathbf{v}}$, $\mathbf{v}$ be the sequence of vertices of $\mathbb{Q}_{d;0}$ visited by the part of the path corresponding to the excursion. Suppose for definiteness' sake that $\mathbf{u}$, $\mathbf{v}$ belong to level L_s. Replace the sequence $\mathbf{u}$, $\tilde{\mathbf{u}}$, $\ldots$, $\tilde{\mathbf{v}}$, $\mathbf{v}$ by

$$\mathcal{S}_{\tilde{\mathbf{u}};\mathbf{u}}(\tilde{\mathbf{u}}), \ldots, \mathcal{S}_{\tilde{\mathbf{u}};\mathbf{u}}(\tilde{\mathbf{v}})$$

(where $\mathcal{S}_{\tilde{\mathbf{u}};\mathbf{u}}$ is one of the isomorphisms defined in Proposition 1.3 and Remark 1.4). From Lemma 2.5 we deduce that strict inequality holds in

$$\|\mathcal{S}_{\tilde{\mathbf{u}};\mathbf{u}}(\tilde{\mathbf{v}}) - \mathbf{u}\|_{s,\infty} \quad \geq \quad \|\mathbf{v} - \mathbf{u}\|_{s,\infty} \, ,$$

with just 2^d exceptions, and so in general the path of the bridge-word is therefore broken by the demotion. However this will not hold for the 2^d exceptions for which $\mathcal{S}_{\tilde{\mathbf{u}};\mathbf{u}}(\tilde{\mathbf{v}}) = \mathcal{M}(\mathbf{v})$, and in such cases we say the demotion is *painless*.

Now one of these 2^d possible cells is represented by $\tilde{\mathbf{u}}$, and this is not a possibility for $\tilde{\mathbf{v}}$. The original bridge-word has to produce a self-avoiding path. While painless demotion will in general destroy the self-avoiding property, it will preserve the property that excursions must begin and end at distinct vertices. This rules out $\tilde{\mathbf{u}}$. In the $2^d - 1$ remaining cases painless demotion pulls down the second and second-to-last vertices to the start and end vertices respectively, and the remaining structure of the excursion (in terms of the pattern of u_x, ℓ_i, δ and sub-excursions) is not altered.

Given a bridge-word, we can subject it to all possible painless demotions to produce what we call a *fully-reduced bridge-word*; one for which no excursions can used to provide painless demotion.

The weight contributions of bridge-words which demote painlessly to a given fully-reduced bridge-word of length r can be bounded as follows. Each of the $r + 1$ vertices in the path corresponding to the fully-reduced bridge-word may correspond to the start and/or the end of iterated demoted excursions, except that the start vertex cannot end an excursion and the end vertex cannot start one. Every start/end pair of τ-bonds for an instance of a *painlessly* demoted excursion corresponds to at most 2^{d-1} possible choices, by Corollary 2.6. The relevant excursions can be reconstructed from assignments of powers of τ to the $2r$ possible start and end places for excursions and a choice from 2^{d-1} possibilities for each start/end of an excursion. Finally, the sum of powers of τ so assigned must be equal to twice the number of excursions, since each begins with an up-move and ends with a down-move. So the summed weight multiplier contributed from painlessly demoted bridge-words is bounded above by

$$\sum_{i_0=0}^{\infty}\cdots\sum_{i_{2r}=0}^{\infty}\tau^{i_0}\cdots\tau^{i_{2r}}\left(2^{d-1}\right)^{(i_0+\ldots+i_{2r})/2} \quad = \quad \frac{1}{\left(1-2^{(d-1)/2}\tau\right)^{2r}}$$

so long as $\tau < 2^{-(d-1)/2}$. So the total weight corresponding to a given fully-reduced bridge-word is bounded above by

$$\frac{\lambda^{\sum_i \#\ell_i}\tau^{2\#\delta}}{\left(\left(1-2^{(d-1)/2}\tau\right)^{\sum_i \#\ell_i+2\#\delta}\right)^2}\,.$$

However we can produce an upper bound for $\#\delta$ for a fully-reduced bridge-word in terms of the number of λ-bonds $\sum_i \#\ell_i$. Work through each of the $\#\delta$ from highest resolution level downwards: the corresponding excursion cannot be painlessly demoted. By Lemma 2.5 we may adjust the demoted excursion by removing some ℓ_i symbols so as to ensure the resulting path is connected. This alteration preserves the fully-reduced nature of the bridge-word, since we are working from highest resolution level downwards, and the possibility of painless demotion concerns only the first two and last two vertices of any excursion.

Arguing this way, we see that the number of excursions (which is to say, $\#\delta$) cannot exceed the number of ℓ_i symbols of all kinds. Accordingly $\#\delta \leq \sum_i \#\ell_i$ and $\sum_i \#\ell_i + 2\#\delta \leq 3\sum_i \#\ell_i$ for a given fully-reduced bridge-word, yielding an upper bound for the number of fully-reduced bridge-words with $\sum_i \#\ell_i = s$. The total number of symbols must be bounded above by $3s$, and they are drawn from an alphabet of length $1+2d+2^d$. It follows that the total weight for bridge-words with $\sum_i \#\ell_i = s$ is bounded above by

$$\left(\frac{\left(1+2d+2^d\right)^3\lambda}{\left(1-2^{(d-1)/2}\tau\right)^6}\right)^s,$$

since $\tau^{2\#\delta} \leq 1$. Under our condition $\tau < 2^{-(d-1)/2}$, this is summable in s for small enough λ, allowing us to deduce that $\tau < 2^{-(d-1)/2}$ implies that for well-separated $\mathbf{v}$ and $\mathbf{u}$ (in which case the number s of $\ell_{\pm i}$ symbols in connecting bridge-words cannot be too small) the probability of both lying in the same cluster tends to zero as their separation increases. This forces the conclusion that there can be no unique infinite cluster for $\tau < 2^{-(d-1)/2}$ and small enough λ, as required. $\square$

Corollary 2.9 *We can extract from the proof, almost surely there is no unique infinite cluster in $\mathbb{Q}_{d;0}$ if $\tau < 2^{-(d-1)/2}$ and*

$$\lambda \quad < \quad \left(\frac{1-2^{(d-1)/2}\tau}{\sqrt{1+2d+2^d}}\right)^6 .$$

Remark 2.10 *Note that in case $d = 1$, and in contrast to the $\mathbb{T}_k \times \mathbb{Z}$ case of [13, §4] we obtain that for any fixed $\tau < 1$ almost surely there is no unique infinite cluster in $\mathbb{Q}_{1;0}$ so long as λ is small enough. Of course the planarity of the graph $\mathbb{Q}_{1;0}$ should allow a more direct proof; we leave the task of establishing this to an interested reader.*

Theorem 2.8 extends to $\mathbb{Q}_d(\mathbf{o})$ by use of an argument reminiscent of ladder variables for random walks.

Corollary 2.11 *Consider (λ, τ) percolation in $\mathbb{Q}_d(\mathbf{o})$. If $\tau < 2^{-(d-1)/2}$ and if $\lambda > 0$ satisfies*

$$\lambda \quad < \quad (1-\tau)\left(\frac{1-2^{(d-1)/2}\tau}{\sqrt{1+2d+2^d}}\right)^6$$

then almost surely there can be no unique infinite cluster.

Proof: We will show that for $\mathbf{u} \in \mathbb{Q}_{d;0} \cap L_0$ we have

$$\sup_{\mathbf{v}\in L_n} \mathbb{P}\left[\,\mathbf{u} \text{ connects to } \mathbf{v} \text{ in } \mathbb{Q}_{d;0}\,\right] \quad \to \quad 0$$

as $n \to \infty$. This suffices to establish the corollary.

Observe that

$$\mathbb{P}\left[\,\mathbf{u} \text{ connects to } \mathbf{v} \text{ in } \mathbb{Q}_{d;0}\,\right] \quad \leq \quad \sum_{\substack{\pi\,:\,\mathbf{u}\leftrightarrow\mathbf{v}\\ \text{self-avoiding path}}} \mathbb{P}\left[\,\pi \text{ open }\,\right].$$

Any such self-avoiding path π can be decomposed, working *backwards* from its end-point $\mathbf{v}$, into a concatenation of paths

$$\pi \quad = \quad \pi_0\nu_0\pi_1\nu_1\ldots\nu_{n-1}\pi_n$$

where

(a) π_r is a self-avoiding path $\mathbf{u}_r \leftrightarrow \mathbf{v}_r$,

(b) $\mathbf{v}_r$ is the last site in L_r to be visited by π, so $\mathbf{v}_n = \mathbf{v}$,

(c) $\mathbf{u}_r$ is the immediate successor of $\mathbf{v}_{r-1}$ if $r > 0$, and $\mathbf{u}_0 = \mathbf{u}$,

(d) ν_r is a path comprising a single step up in resolution.

We note that the map $\pi \mapsto \{\pi_0, \pi_1, \ldots, \pi_n\}$ is actually 1:1 (though not onto), since the procedure of building π by working backwards from $\mathbf{v}$ shows that the choices of the ν_r are all forced choices. Moreover the π_r all correspond to bridge-words, to which we may apply the bound obtained in the proof of Theorem 2.8.

We deduce

$$\mathbb{P}\left[\,\pi \text{ open }\,\right] \quad = \quad \tau^n \prod_{r=0}^{n} \mathbb{P}\left[\,\pi_r \text{ open }\,\right]$$

and therefore

$$\sum_{\substack{\pi\,:\,\mathbf{u}\leftrightarrow\mathbf{v}\\ \text{self-avoiding path}}} \mathbb{P}\left[\,\pi \text{ open }\,\right] \;\leq\; \tau^n \left(\sum_{\pi_0} \mathbb{P}\left[\,\pi_0 \text{ open }\,\right]\right)^{n+1}$$

$$\leq \qquad \tau^{-1}\left(\tau \sum_{s=0}^{\infty}\left(\frac{\left(1+2d+2^d\right)^3\lambda}{\left(1-2^{(d-1)/2}\tau\right)^6}\right)^s\right)^{n+1}$$

$$= \qquad \tau^{-1}\left(\frac{\tau\left(1-2^{(d-1)/2}\tau\right)^6}{\left(1-2^{(d-1)/2}\tau\right)^6-\left(1+2d+2^d\right)^3\lambda}\right)^{n+1}$$

where the last inequality uses the bound obtained in the proof of Theorem 2.8. For sufficiently small λ (as given in the statement of this corollary) not only does the geometric sum over n converge, but also the term in the $n+1^{\text{st}}$ power is smaller than 1, so the required convergence to zero is obtained. □

We conclude this section by giving a simple upper bound on the threshold probability τ at which the infinite cluster becomes unique for all positive λ. The idea is to compare with independent bond percolation on $\mathbb{Z}^d$ by considering connectivity between vertices in L_0. To simplify the exposition we work with the case $d = 2$ only, though the method clearly generalizes to all dimensions $d > 1$ (albeit changing the bound on τ).

Theorem 2.12 *If $\tau > \sqrt{2/3}$ then the infinite cluster of $\mathbb{Q}_{2:0}$ is almost surely unique for all positive λ.*

Proof: We prune bonds in $\mathbb{Q}_{2;0}$ as follows. In each cell of L_0 we retain all τ-bonds at top level, but at lower levels (L_1, L_2, $\ldots$) remove τ-bonds which lead to internal vertices. We also remove all τ-bonds pointing into alternating 45^o sectors, changing the parity of alternation in adjacent cells as illustrated in Figure 6. Finally we retain only those λ-bonds reaching across boundaries of cells in L_0 and not contained in deleted 45^o sectors.

The effect of this is that a direct connection is certainly established across the boundary between the cells corresponding to two neighbouring vertices $\mathbf{u}$, $\mathbf{v}$ in L_0 if

(a) the τ-bond leading from $\mathbf{u}$ to the relevant boundary is open (probability τ);

(b) a τ-branching process (formed by using τ-bonds mirrored across the boundary) survives indefinitely, where this branching process has family-size distribution Binomial$(2, \tau^2)$;

(c) the τ-bond leading from $\mathbf{v}$ to the relevant boundary is open (probability τ);

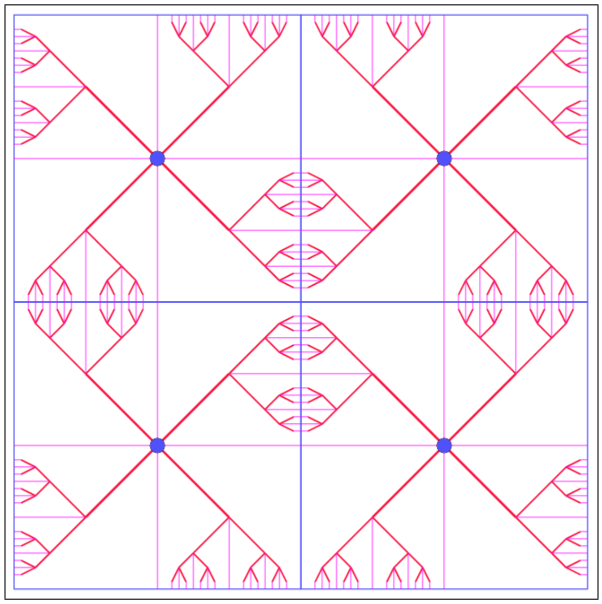

Figure 6: Construction of pruned percolation problem for $\mathbb{Q}_{2;0}$. For the sake of pictorial clarity we depart from the convention in the text, and represent cells by vertices placed at their centroids.

For then there will be infinitely many independent chances to complete the connection by using single λ-bonds.

The probability that (b) occurs is readily computed using branching process theory. The family-size generating function is

$$\tau^4 s^2 + 2\tau^2(1-\tau^2)s + (1-\tau^2)^2$$

and so the extinction probability is the least non-negative root of

$$s \quad = \quad \tau^4 s^2 + 2\tau^2(1-\tau^2)s + (1-\tau^2)^2 .$$

The solution is

$$s \quad = \quad \frac{(1-\tau^2)^2}{\tau^4} .$$

Consequently the known theory for bond percolation in $\mathbb{Z}^2$ tells us that a sufficient condition for there to be just one infinite cluster for this pruned percolation problem is that

$$\tau \left(1 - \frac{(1-\tau^2)^2}{\tau^4} \right) \tau \quad = \quad \frac{2\tau^2 - 1}{\tau^2} \quad > \quad \frac{1}{2} .$$

This leads to the condition $\tau^2 > 2/3$.

We now must argue that the resulting unique infinite cluster is connected to all connected regions in the *full* percolation problem, and therefore establishes uniqueness of the infinite cluster for the full problem also.

First note that if $\tau^2 > 2/3$ and $\lambda > 0$ then we will retain a unique infinite cluster if we further prune the percolation problem by deleting all bonds except those contained in $L_0 \cup L_1 \cup \ldots \cup L_{n-1}$, so long as n is sufficiently large (depending on τ and λ). Indeed, it suffices to ensure that the probability of connecting across the boundary between cells corresponding to two neighbouring vertices in L_0 still exceeds $1/2$. We can therefore identify a whole infinite sequence of pruned percolation problems, based on $L_{nr} \cup L_{nr+1} \cup \ldots \cup L_{n(r+1)-1}$ for $n = 0$, $1, \ldots$, each of which contains a unique infinite cluster.

Further, note that the τ-branching process is super-critical when $\tau^2 > 2/3$. A 0-1-law argument can then be deployed to show, in the full problem the infinite clusters for $L_{nr} \cup L_{nr+1} \cup \ldots \cup L_{n(r+1)-1}$ must all be connected to each other, thus forming a large infinite cluster C.

Furthermore τ-super-criticality implies that any infinite cluster in the full percolation problem cannot be confined to a finite number of resolution levels. A suitable adaptive enumeration of the vertices in such an infinite cluster, delivering a sequence of vertices lying in the successive layers L_0, L_n, L_{2n}, $\ldots$, shows that almost surely eventually such a cluster must intersect C.

We have therefore shown there is just one infinite cluster when $\tau^2 > 2/3$ and $\lambda > 0$. □

The techniques of Theorems 2.2 and 2.3 can be applied to extend this result to the case of $\mathbb{Q}_2(\mathbf{o})$, working with $\mathbb{Z}^2$ bond percolation for layers L_0, L_n, L_{2n}, $\ldots$, and noting that the FKG inequality allows us to combine percolation within $L_{nr} \cup L_{nr+1} \cup \ldots \cup L_{n(r+1)-1}$ with connection along τ-bonds between L_{nr} and $L_{n(r+1)}$. We obtain

Corollary 2.13 *If $\tau > \sqrt{2/3}$ and $\lambda > 0$ then almost surely there is a unique infinite cluster for $\mathbb{Q}_2(\mathbf{o})$.*

Figure 7 provides a graphical summary of the information we have obtained concerning uniqueness and existence of infinite clusters for dimension $d = 2$.

In their treatment of percolation on $\mathbb{T}_k \times \mathbb{Z}$, Grimmett and Newman [13] remark that they have not established whether the λ-τ region of uniqueness of a infinite cluster is an increasing subset of $[0, 1]^2$. However recent work by Häggström, Peres and Schonmann [14, Definition 2.2, Theorem 2.3] shows how to use invasion percolation to establish the corresponding fact for independent bond percolation with constant retention probability on infinite connected graphs of bounded degree which are semi-transitive (or, more generally, exhibit uniform percolation at supercritical levels of the retention probability). This class of percolation problems includes our graphs $\mathbb{Q}_d(\mathbf{o})$ except for the constancy of retention probability; however the proof is easily modified to allow for two different levels of retention probability.

This is discussed in Appendix B.

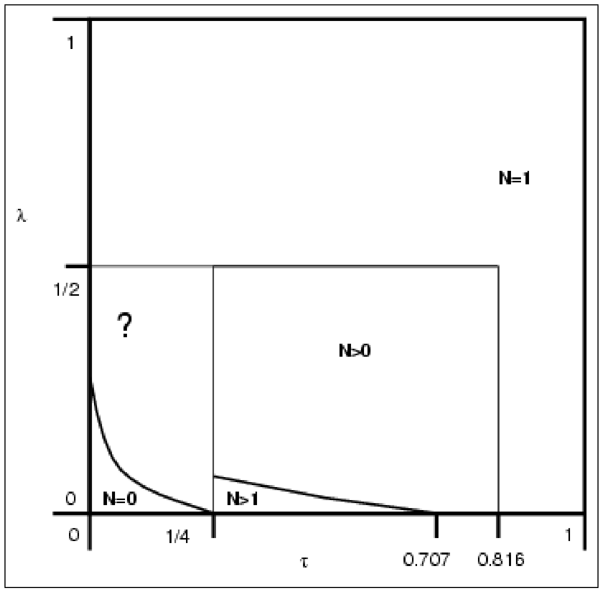

Figure 7: Information on transition concerning existence and uniqueness of infinite clusters, for $\mathbb{Q}_{2;0}$ or $\mathbb{Q}_2(\mathbf{o})$. Here N denotes the number of infinite clusters.

Thus the schematic of Figure 7, for example, may be redrawn to include boundaries between the various phases which are curves defined as non-increasing functions of τ. However we are *not* able to guarantee that the coexistence phase (existence of many infinite clusters) intersects all levels of λ up to the critical level $\lambda = 1/2$.

2.4 Finite islands phenomenon

We need one more result to facilitate our discussion in Section 3 of a reasonably complete schematic phase transition diagram for the Ising model. Consider the connected clusters of sites formed by supercritical percolation in $\mathbb{Q}_{d;0}$ or $\mathbb{Q}_d(\mathbf{o})$, and remove all infinite clusters. The remaining clusters form *islands* under the connectivity relation of adjacency. When are there no infinite islands? This question has been investigated for $\mathbb{T}_k \times \mathbb{Z}$ by Newman and Wu [22], and (focusing on isotropic bond percolation on more general graphs) by Schonmann [26]. The case of $\mathbb{Q}_{d;0}$ and $\mathbb{Q}_d(\mathbf{o})$ can be treated by an easy variation on the methods of Newman and Wu.

Theorem 2.14 *For fixed λ, for all sufficiently large τ almost surely there are no infinite islands in any of $\mathbb{Q}_d(\mathbf{o})$, $\mathbb{Q}_{d;0}$, or $\mathbb{Q}_d$.*

Proof: It suffices to find an upper bound for the mean size of the island at some $\mathbf{u}_0$.

Let $\eta = \eta(\lambda, \tau) > 0$ be the probability that $\mathbf{o}$ is not part of an infinite cluster in $\mathbb{Q}(\mathbf{o})$. (We suppose $\tau < 1$.) For fixed $\lambda > 0$ we know by comparison with the extinction probability η_{br} for the branching process formed from τ-bonds that $\eta(\lambda, \tau) \leq \eta_{\mathrm{br}} \to 0$ as $\tau \to \infty$.

We follow Newman and Wu [22, Lemma 3.3] in noting a bound of isoperimetric type. Define the "cone boundary" $\partial_c(S)$ of a finite subset S of vertices in $\mathbb{Q}_d$ as the collection of *daughters* $\mathbf{v}$ of S such that $\mathbb{Q}_d(\mathbf{v}) \cap S = \emptyset$. Since S is finite we may suppose $S \subseteq \mathbb{Q}_d(0)$. Using induction on construction of S, layer L_n after layer L_{n-1}, we obtain the isoperimetric bound $\#(\partial_c(S)) \geq (2^d - 1)\#(S)$. (Adding a vertex at the lowest layer certainly introduces 2^d new daughters into the cone boundary: it may cloak one older member of the cone boundary but no more than that.)

It follows, the probability that a self-avoiding path S of length n lies entirely in the island at $\mathbf{u}_0$ is bounded above by the probability of the intersection of independent events corresponding to failure to create open infinite paths corresponding to each of the $\mathbb{Q}_d(\mathbf{v}) \in \partial_c(S)$:

$$\mathbb{P}\left[\, S \text{ in island at } \mathbf{u} \,\right] \quad \leq \quad (1 - \tau(1-\eta))^{(2^d-1)n} \,.$$

(Note: each vertex in the cone boundary must by definition have its mother in S.)

On the other hand the number $N(n)$ of self-avoiding paths of length n beginning at $\mathbf{u}_0$ is bounded above by

$$N(n) \quad \leq \quad (1 + 2d + 2^d)(2d + 2^d)^n$$

and so (using $\eta(\lambda, \tau) \leq \eta_{\mathrm{br}}$ and the probability generating function relationship for η_{br}) the mean size of the island is bounded above by

$$\sum_{n=0}^{\infty} (1 + 2d + 2^d)(2d + 2^d)^n \eta_{\mathrm{br}}^{n(1-2^{-d})} \,.$$

For large enough $\tau < 1$ we have $\eta_{\mathrm{br}} < 1/(2d + 2^d)^{2^d/(2^d-1)}$ and therefore convergence for the above sum, and so we can deduce that the island is almost surely finite. This establishes the proof. $\square$

We summarize the information obtained about the finite island property for the case $d = 2$ in Figure 8. The critical value for τ for small λ is obtained as follows. We have noted $\eta = \eta(\lambda, \tau)$ is bounded above by the least non-negative root $\eta_{\mathrm{br}} = \eta_0$ of

$$\eta_0 \quad = \quad (1-\tau)^4 + 4(1-\tau)^3 \tau\eta_0 + 6(1-\tau)^2(\tau\eta_0)^2 + 4(1-\tau)(\tau\eta_0)^3 + (\tau\eta_0)^4 \,.$$

We need to solve this for τ when $\eta_0^{3/4} = 1/8$, in order to identify a threshold above which propagation to infinite islands has zero probability. The relevant root is $\tau = 0.533333$.

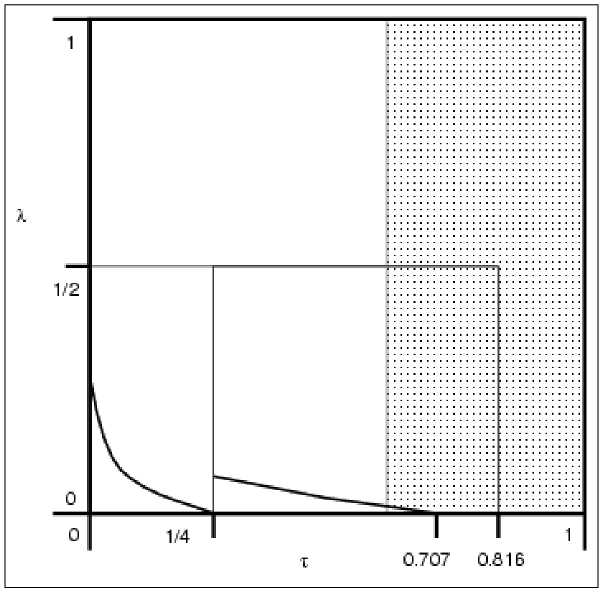

Figure 8: The finite island property holds for $\mathbb{Q}_2(\mathbf{o})$ in the shaded region.

3 Ising models on generalized quad-trees

The motivation of this paper is primarily to gain a better understanding of Ising models defined on generalized quad-trees, particularly $\mathbb{Q}_{d;0}$ and $\mathbb{Q}_d(\mathbf{o})$, and with particular reference to the case when λ-bonds and τ-bonds have different bond strengths J_λ and J_τ, typically with one of J_λ, J_τ being small. We follow Newman and Wu [22] very closely, so this section simply sets out the general reasoning and refers to [22] for some details. The plan is to use the representation of Ising models in terms of random-cluster models, and the Fortuin-Kasteleyn comparison inequalities, to relate the resulting dependent bond percolation to results concerning (independent) bond percolation. This allows exploitation of the percolation results gained in Section 2.

3.1 The random-cluster model representation

Recall the now classical representation of the Ising model on a finite graph G in terms of a random cluster model. Suppose the Ising model is based on the Hamiltonian

$$H \quad = \quad -\frac{1}{2} \sum_{\langle x,y\rangle \in \mathcal{E}(G)} J_{\langle x,y\rangle}(S_x S_y - 1) \tag{3.1}$$

where $S_x \in \pm 1$ is the spin at site $x \in G$ and the sum runs over the (undirected) bonds of G, namely $\langle x, y\rangle \in \mathcal{E}(G)$. Thus the probability of a given configuration

$\{S_x : x \in G\}$ is proportional to

$$\exp(-H) \quad = \quad \exp\left(\frac{1}{2}\sum_{\langle x,y\rangle\in\mathcal{E}(G)} J_{\langle x,y\rangle}(S_x S_y - 1)\right) . \tag{3.2}$$

We now set

$$p_{\langle x,y\rangle} \quad = \quad 1 - \exp(-J_{\langle x,y\rangle}) . \tag{3.3}$$

In our case, of course, $J_{\langle x,y\rangle}$ is set equal to J_λ or to J_τ depending on whether $\langle x, y\rangle$ is a λ-bond or a τ-bond.

Denote a configuration of open/closed bonds by $\{b_{\langle x,y\rangle} : \langle x, y\rangle \in \mathcal{E}(G)\}$, where $b_{\langle x,y\rangle} = 1$ if the bond is open and $b_{\langle x,y\rangle} = 0$ if it is closed. A configuration $\{b_{\langle x,y\rangle} : \langle x, y\rangle \in \mathcal{E}(G)\}$ forming C clusters of sites has a probability under the q-random cluster model which is proportional to

$$q^C \times \prod_{\langle x,y\rangle\in\mathcal{E}(G)} \left((p_{\langle x,y\rangle})^{b_{\langle x,y\rangle}} \times (1 - p_{\langle x,y\rangle})^{1-b_{\langle x,y\rangle}}\right) .$$

From the work of Fortuin and Kasteleyn [4, 5, 6], the Ising model can be obtained by choosing spins ± 1 uniformly at random, one spin for each cluster of sites connected by open bonds of the $(q = 2)$-random cluster model on $\mathcal{E}(G)$.

Of course the case $q = 1$ yields bond-percolation.

Furthermore Fortuin and Kasteleyn establish comparison inequalities of which the following is a special case:

Lemma 3.1 *For $q \geq 1$, let $\mathbb{P}_{q,p}$ denote the q-random cluster measure on the bonds of G using the bond probability parameters $p_{\langle x,y\rangle}$. Let A be an increasing event concerning the configuration of bonds on G (if a configuration lies in A then so does any configuration derived by making more bonds open). Then*

$$\mathbb{P}_{q,p}(A) \quad \leq \quad \mathbb{P}_{1,p}(A) \tag{3.4}$$

$$\mathbb{P}_{q,p}(A) \quad \geq \quad \mathbb{P}_{1,p'}(A) \tag{3.5}$$

where

$$p'_{\langle x,y\rangle} = \frac{p_{\langle x,y\rangle}}{p_{\langle x,y\rangle} + (1 - p_{\langle x,y\rangle})q} = \frac{p_{\langle x,y\rangle}}{q - (q - 1)p_{\langle x,y\rangle}} .$$

Accessible proofs of these assertions in the case of constant $p_{\langle x,y\rangle}$ may be found for example in [9]. They may be used to establish the existence of limiting random-cluster measures on infinite graphs, though care has to be taken to distinguish between possibly different limits arising for free boundary conditions and "wired" boundary conditions (all components intersecting the boundary are viewed as connected into one wired cluster). The monotonicity results (3.4), (3.5) persist in the limit.

The representation of the Ising model using a $(q = 2)$-random cluster model, together with these comparison inequalities, now allow us to address questions about phase transition for Ising models on $\mathbb{Q}_{d;0}$ and $\mathbb{Q}_d(\mathbf{o})$ using our results about percolation on $\mathbb{Q}_{d;0}$ and $\mathbb{Q}_d(\mathbf{o})$.

3.2 Uniqueness and non-uniqueness of Gibbs states

Consider first what we may discover using the percolation results of Section 2.1 on transition from zero to many infinite clusters. The event that a bond percolation model has infinite clusters is increasing, in the sense used in Lemma 3.1. Simple comparison arguments using the FKG inequality show that we can use a continuous contour to delimit the (λ, τ)-region corresponding to the absence of infinite clusters for the $(q = 2)$-random cluster model. We shall use the percolation results to get information as to where this contour meets the λ- and τ-axes.

In passing, we note that the proof of the 0:1 law given as [22, Lemma 3.2] will work in the case of $\mathbb{Q}_2(\mathbf{o})$, so the probability that the $(q = 2)$-random cluster model possesses infinite clusters must be zero or one depending on the λ and τ parameters.

The random cluster representation applies most simply to Ising models on finite graphs, so we consider the intersections of $\mathbb{Q}_{d;0}$ and $\mathbb{Q}_d(\mathbf{o})$ with the first $n+1$ levels $L_0 \cup L_1 \cup \ldots \cup L_n$. Let A_n be the event that the bond configuration allows percolation from $\mathbf{o}$ through to resolution level n in this intersection. As $\lambda \to 0$, Theorem 2.1 shows that the probability of A_n under (λ, τ)-independent bond percolation is bounded away from zero, or not, as $n \to \infty$ according as to whether $\tau > 2^{-d}$, or not. The comparison inequalities (3.4), (3.5) then show that the limiting probability $\lim_{n\to\infty} A_n$ under (λ, τ)-$(q = 2)$-random cluster percolation is positive if $\tau > 2^{-d} \times 2/(1 + 2^{-d})$, and is zero for all small enough $\lambda > 0$ if $\tau < 2^{-d}$. Lemma 3.1 allows us to argue to the limit for the monotonic event A_∞, and to show that under (λ, τ)-$(q = 2)$-random cluster percolation the probability of A_∞, the event that the bond configuration builds an infinite cluster based on $\mathbf{o}$, is positive if $\tau > 2^{-d} \times 2/(1 + 2^{-d})$, and is zero for all small enough $\lambda > 0$ if $\tau < 2^{-d}$.

A similar argument can be deployed in the case of dimension $d = 2$ (which is the case of most relevance to image analysis), using Theorem 2.2 to show that the probability of A_∞ under (λ, τ)-$(q = 2)$-random cluster percolation is positive for all small $\tau > 0$ if $\lambda > 2/3$, and is zero for all small enough $\tau > 0$ if $\lambda < 1/2$.

Consequently the contour which delimits the region of no infinite clusters will meet the λ-axis at a value of τ lying in the interval $[2^{-d}, 2^{-d} \times 2/(1+2^{-d})]$. Moreover, in the case of dimension $d = 2$ the contour must meet the τ-axis at a value of λ lying in the interval $[1/2, 2/3]$.

Using the correspondence (3.3), and arguing as in [22], we deduce

Theorem 3.2 *Consider the Ising model on $\mathbb{Q}_{d;0}$ or $\mathbb{Q}_d(\mathbf{o})$, using interactions J_λ and J_τ, and with spin $+1$ at the boundary "at infinite resolution level".*

> ***Case of small J_λ:*** *If $J_\tau < \ln(1/(1 - 2^{-d}))$, and J_λ is sufficiently small, then the spin at $\mathbf{o}$ is equally likely to be ± 1 and indeed there is just one Gibbs measure regardless of boundary conditions. If $J_\tau > \ln((1+2^{-d})/(1-2^{-d}))$ then the spin at $\mathbf{o}$ is more likely to be $+1$ than -1, and moreover there is more than one Gibbs state for the Ising model.*

Case of small J_τ*: If dimension $d = 2$, $J_\lambda < \ln 2$ and J_τ is sufficiently small then the spin at* $\mathbf{o}$ *is equally likely to be* ± 1 *and indeed there is just one Gibbs measure regardless of boundary conditions. If $J_\lambda > \ln 3$ then the spin at* $\mathbf{o}$ *is more likely to be* $+1$ *than* -1*, and moreover there is more than one Gibbs state for the Ising model.*

3.3 Free Ising model and mixtures of extreme Gibbs states

We now consider the implications of the transition from many infinite clusters to a unique infinite cluster, as described in Sections 2.2 and 2.3. Standard considerations show now that the probability under independent bond percolation of any two vertices $\mathbf{x}$ and $\mathbf{y}$ being connected, whether in $\mathbb{Q}_{d;0}$ or $\mathbb{Q}_d(\mathbf{o})$, is bounded away from zero if and only if there may be a unique infinite cluster. Uniqueness of the infinite cluster does not constitute an increasing event, but connectedness of two vertices does, and this allows us to deduce the following using the same line of argument as that of Section 3.2, and applying Theorems 2.3, 2.8, and 2.12.

Theorem 3.3 *Consider the* free *Ising model on $\mathbb{Q}_{d;0}$ or $\mathbb{Q}_d(\mathbf{o})$, using interactions J_λ and J_τ (so spins at the boundary "at infinite resolution level" are unspecified).*

> ***Case of small*** J_λ*: If $J_\tau < \ln(1/(1-2^{-(d-1)/2}))$, and J_λ is sufficiently small, then the correlation between spins at $\boldsymbol{x}$ and $\boldsymbol{y}$ decay to 0 as the distance between the two vertices increases. In the case of dimension $d = 2$, if $J_\tau > \ln((1+\sqrt{2/3})/(1-\sqrt{2/3}))$ then the correlation is positive and bounded below away from zero.*
>
> ***Case of small*** J_τ*: If dimension $d = 2$ and $J_\lambda > \ln 3$ then the correlation is positive and bounded below away from zero.*

Corollary 3.4 *Suppose $\ln((1+2^{-d})/(1-2^{-d})) < J_\tau < \ln(1/(1-2^{-(d-1)/2}))$, and J_λ is sufficiently small. Then Theorems 3.2 and 3.3 combine to show that the free Ising model cannot be expressed as a mixture of the two extreme Gibbs states determined by spin $+1$ and spin -1 at the boundary "at infinite resolution level".*

Remark 3.5 *For small J_τ the picture presented by Theorems 3.2 and 3.3 does not guarantee an intermediate stage of the kind presented in Corollary 3.4: in this part of the phase-transition diagram the model appears to have more in common with the planar Ising model than elsewhere.*

While uniqueness of the infinite cluster does not in itself constitute an increasing event, it *does* constitute an increasing event when combined with the finite-island property described in Section 2.4, as Newman and Wu observe [22]. Following their arguments, and using Theorem 2.14, we can show the following:

Theorem 3.6 *Consider the* free *Ising model on $\mathbb{Q}_{d;0}$ or $\mathbb{Q}_d(\mathbf{o})$, using interactions J_λ and J_τ (so spins at the boundary "at infinite resolution level" are unspecified). For large enough J_τ, and certainly for $J_\tau > \ln((1+\sqrt{2/3})/(1-\sqrt{2/3}))$ in the case of dimension $d=2$, the free Ising model is the $\frac{1}{2}:\frac{1}{2}$-mixture of the extreme Gibbs states determined by spin $+1$ and spin -1 at the boundary "at infinite resolution level".*

Figure 9 sketches out the regions established above for these different phases for $\mathbb{Q}_2(\mathbf{o})$. Note that the middle phase (root influenced by boundary values all set to a single spin, but dependence between specified vertices will decay to zero with distance) is drawn as extending to all levels of λ less than the limiting critical level for small τ; however we do not know whether this is in fact the case.

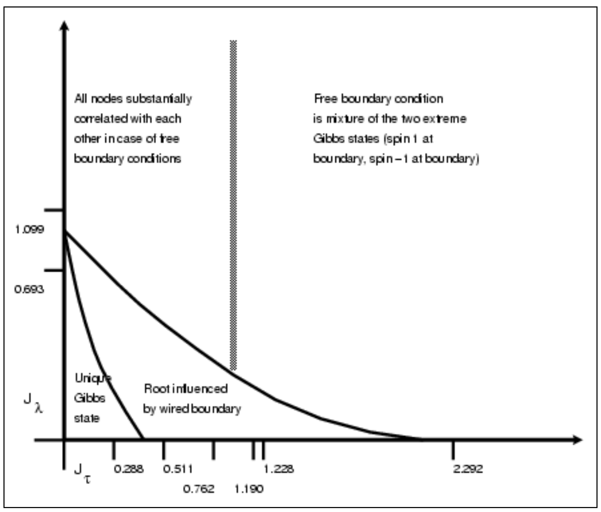

Figure 9: Schematic phase-transition diagram for Ising model on $\mathbb{Q}_2(\mathbf{o})$. (Figure not drawn to scale.) Note that the middle phase (root influenced by boundary values all set to a single spin) may not in fact extend up to the limiting critical level of λ for small τ.

4 Simulations and further work

The above results describe the free Ising model on $\mathbb{Q}_{d;0}$ or $\mathbb{Q}_d(\mathbf{o})$. The case of $\mathbb{Q}_d(\mathbf{o})$ is amenable to simulation, since the global Markov property for Markov random fields allows us to view each successive resolution level L_n as being produced by simulation of an Ising model on $\{0,1,\ldots,2^n-1\}^2$ influenced by a magnetic field representing the interaction with the realization of the Ising model on the previous level L_{n-1}. Convergence at each level is typically fast for small J_λ (as can be assessed using simple CFTP methods following [25]), and the main computational issue is one of handling the exponential growth in memory at increasing resolution levels.

We have run a large number of simulations of an approximate version of the model in order to explore its behaviour, the results of which are summarized in figure 10. An animated version of this figure is to be found at

http://www.dcs.warwick.ac.uk/~rgw/sira/sim.html.

The hyperbolic growth of the tree structure obliges us to restrict the simulations in the following ways:

(1) Only 200 resolution levels are simulated;

(2) At each resolution level the simulation consists of 1000 sweeps through the image in scan order, giving a total of some 4 million site visits, after which we assume approximate equilibrium is attained;

(3) At each resolution level a square sub-region of 128×128 pixels is simulated, this being conditioned by the corresponding 64×64 pixel region at the mother level;

(4) We approximate by imposing periodic boundary conditions on each 128×128 square region;

(5) At the coarsest resolution level (resolution level 0), all pixels are set white. At each subsequent resolution level the 'all black' state is used as the initial condition before applying the 1000 sweeps through the image.

The results in figure 10 show broad agreement with the predictions, in that three distinct phases are observable, occurring in the expected regions of the $J_\lambda - J_\tau$ plane: (i) a single Gibbs distribution at low $J_\lambda - J_\tau$ values, leading to (ii) a phase with many large clusters and lastly (iii) a single large cluster at high values of J_λ or J_τ. Note that for high J_λ the single large cluster is white; this reflects the strong influence of the pixels at the bottom level for these parameter settings.

However real images (such as those to be found in [31], or even the artificial images in Figures 2 and 3) are appropriately modelled neither by free Ising models nor by either of the two extreme Gibbs states, but by specification of specific mixtures of ± 1 on the ideal boundary at $n = \infty$ representing images (with, we would expect, some kind of smoothness of contours between $+1$ and -1). A random cluster model approach to this requires us to specify boundary conditions which wire together two or more clusters of sites at infinity (black versus white image) and *condition* on these clusters not being connected. The question is whether such conditioning produces an interface which reaches down to low (finite!) resolution levels. We defer investigation of this to a future project, and merely note for now (a) that the analogous problem for the random cluster model on $\mathbb{Z}^d$ is treated by Gierls and Grimmett [10], developing Dobrushin's classic interface phenomenon for the Ising model on $\mathbb{Z}^3$ [3] and (b) the case of $d = 1$ with $\lambda = \tau$ is related to the work of Series and Sinai [27], which considers

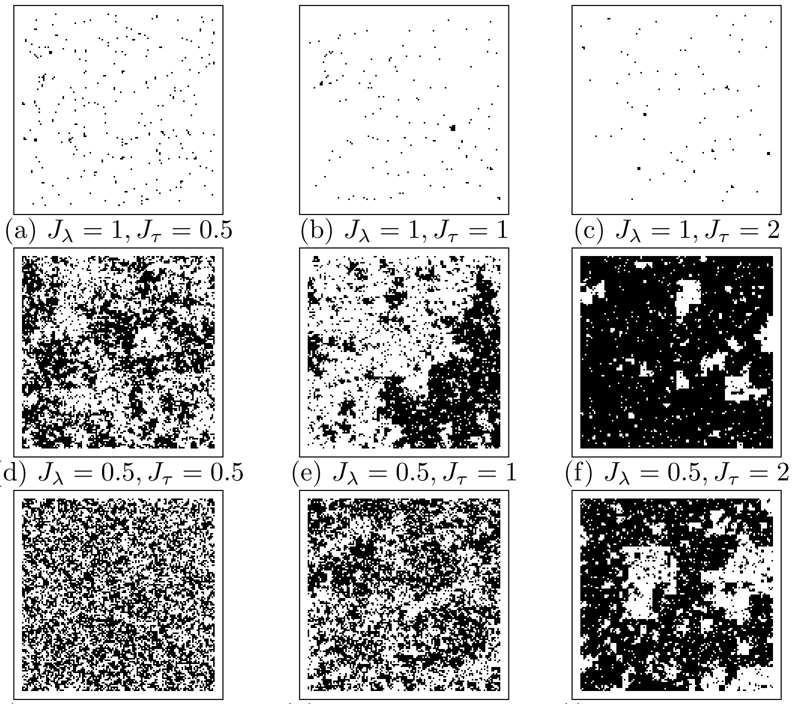

(a) $J_\lambda = 1, J_\tau = 0.5$ (b) $J_\lambda = 1, J_\tau = 1$ (c) $J_\lambda = 1, J_\tau = 2$

(d) $J_\lambda = 0.5, J_\tau = 0.5$ (e) $J_\lambda = 0.5, J_\tau = 1$ (f) $J_\lambda = 0.5, J_\tau = 2$

(g) $J_\lambda = 0.25, J_\tau = 0.5$ (h) $J_\lambda = 0.25, J_\tau = 1$ (i) $J_\lambda = 0.25, J_\tau = 2$

Figure 10: Samples from the multiresolution Ising simulation. The coupling parameters are shown for each sample.

Ising models defined by Cayley graphs of finitely generated co-compact groups of isometries of the hyperbolic plane, and which establishes exactly that such interfaces then exist.

Finally we note that the discussion in Section 3 generalizes to Potts models, again following the methods of [22]. However we omit discussion of this here, leaving the details as an exercise for the interested reader.

References

[1] I. Benjamini and O. Schramm. Percolation beyond $\mathbf{Z}^d$, many questions and a few answers. *Electronic Communications in Probability*, 1:no. 8, 71–82 (electronic), 1996.

[2] C. A. Bouman and M. Shapiro. A Multiscale Random Field Model for Bayesian Image Segmentation. *IEEE Transactions on Image Processing*, 3:162–176, 1994.

[3] R. Dobrushin. Coexistence of phase. *Theory of Probability and Applications*, 17:582–600, 1972.

[4] C.M. Fortuin. On the random-cluster model. II. The percolation model. *Physica*, 58:393–418, 1972.

[5] C.M. Fortuin. On the random-cluster model. III. The simple random-cluster model. *Physica*, 59:545–570, 1972.

[6] C.M. Fortuin and P.W. Kasteleyn. On the random-cluster model. I. Introduction and relation to other models. *Physica*, 57:536–564, 1972.

[7] D. Geman, S. Geman, C. Graffigne, and P. Dong. Boundary Detection by Constrained Optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:609–628, 1990.

[8] S. Geman and D. Geman. Stochastic relaxation, Gibbs distribution, and Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.

[9] H. Georgii, O. Häggström, and C. Maes. The random geometry of equilibrium phases. In C. Domb and J.L. Lebowitz, editors, *Phase Transitions and Critical Phenomena*, pages 1–142, London, 2001. Academic Press.

[10] G. Gielis and G.R. Grimmett. Rigidity of the interface for percolation and random-cluster models. *Journal of Statistical Physics*, 109:1–37, 2002. See also arXiv:math.PR/0109103.

[11] R.C. Gonzalez and R.E. Woods. *Digital Image Processing*. Prentice Hall, 2nd edition, 2002.

[12] G. Grimmett. *Percolation*. Springer-Verlag, Berlin, second edition, 1999.

[13] G. Grimmett and C. Newman. Percolation in $\infty+1$ dimensions. In *Disorder in physical systems*, pages 167–190. The Clarendon Press Oxford University Press, New York, 1990.

[14] O. Häggström, Y. Peres, and R.H. Schonmann. Percolation on transitive graphs as a coalescent process: relentless merging followed by simultaneous uniqueness. In *Perplexing problems in probability*, pages 69–90. Birkhäuser Boston, Boston, MA, 1999.

[15] Z. Kato, M. Berthod, and J. Zerubia. A Hierarchical Markov Random Field Model and Multi-Temperature Annealing for Parallel Image Classification. Technical Report Tech. Rept. 1938, INRIA, 1993.

[16] J-M. Laferte, P. Perez, and F Heitz. Discrete Markov Image Modeling and Inference on the Quadtree. *IEEE Transactions on Image Processing*, 9:390–404, 2000.

[17] S. M. Lavalle and S. A. Hutchinson. A Bayesian Segmentation Methodology For Parametric Image Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:211–217, 1995.

[18] J. Li, R.M. Gray, and R.A. Olshen. Multiresolution Image Classification by Hierarchical Modeling with Two-Dimensional Hidden Markov Models. *IEEE Transactions on Information Theory*, 46:1826–1840, 2000.

[19] S.G. Mallat. Multifrequency Channel Decompositions of Images and Wavelet Models. *IEEE Transactions on Acoustics, Speech, Signal Processing*, 37(12), December 1989.

[20] B. S. Manjunath and R. Chellappa. Unsupervised Texture Segmentation Using Markov Random Field Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13:478–482, 1991.

[21] M. Mignotte, C. Collet, P. Perez, and P. Bouthemy. Sonar Image Segmentation Using an Unsupervised Hierarchical MRF Model. *IEEE Transactions on Image Processing*, 9:1216–1231, 2000.

[22] C. Newman and C. Wu. Markov fields on branching planes. *Probability Theory and Related Fields*, 85(4):539–552, 1990.

[23] D. K. Panjwani and G. Healey. Markov Random Field Models For Unsupervised Segmentation Of Textured color Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:939–954, 1995.

[24] C.J. Preston. Spatial birth-and-death processes. *Bull. Inst. Internat. Statist.*, 46:371–391, 1977.

[25] J.G. Propp and D.B. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, 9:223–252, 1996.

[26] R.H. Schonmann. Multiplicity of phase transitions and mean-field criticality on highly non-amenable graphs. *Comm. Math. Phys.*, 219(2):271–322, 2001.

[27] C. Series and Y. Sinaĭ. Ising models on the Lobachevsky plane. *Comm. Math. Phys.*, 128(1):63–76, 1990.

[28] F. Spitzer. Markov random fields on an infinite tree. *The Annals of Probability*, 3(3):387–398, 1975.

[29] R. Szeliski. *Bayesian Modelling of Uncertainty in Low-level Vision*. Kluwer Academic Publishers, 1989.

[30] R. Wilson and M. Spann. *Image Segmentation and Uncertainty*. Research Studies Press, 1988.

[31] R.G. Wilson and C.T. Li. A class of discrete multiresolution random fields and its application to image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(11):Accepted for publication, 2002.

[32] C. Wu. Ising models on hyperbolic graphs. *J. Statist. Phys.*, 85(1-2):251–259, 1996.

[33] C. Wu. Ising models on hyperbolic graphs. II. *J. Statist. Phys.*, 100(5-6):893–904, 2000.

The following appendices summarize work which was carried out during the investigation reported here, but which has not found its way into the version submitted for publication.

A Percolation: More on transition from zero to many infinite clusters

Theorem 2.1 gives a sharp lower bound on the $\lambda \downarrow 0$ point at which the transition occurs between zero and many infinite clusters. We can also derive a lower bound for the correction to the mean number of daughters at resolution level n if just one horizontal bond is allowed at some resolution level between 0 and n. For the sake of simplicity we describe this only for $\mathbb{Q}_2(\mathbf{o})$, though the argument generalizes directly to other dimensions d.

Consider first the mean number of sites at level n which are connected to $\mathbf{o}$ in $\mathbb{Q}_2(\mathbf{o}) \cap (L_0 \cup L_1 \cup \ldots \cup L_n)$. We know we must have many infinite clusters in $\mathbb{Q}_2(\mathbf{o})$ when this strictly exceeds 1. For small λ we may estimate the excess of this mean over $(4\tau)^n$ (the mean value if $\lambda = 0$) to order λ^2 by adding the contributions to the excess arising from adding each of the various λ-bonds singly.

The ingredients of the calculation are illustrated in Figure 11.

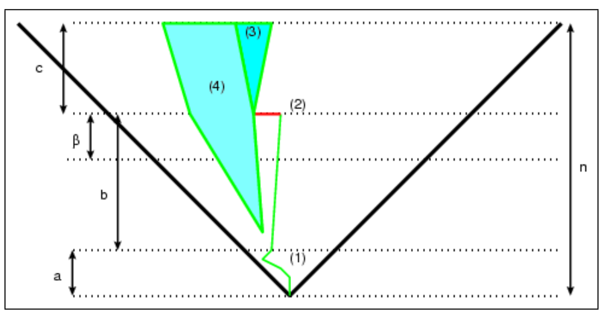

Figure 11: Illustration of calculations used in generating upper bound for transition from zero to many infinite clusters

(1) Now the addition of a single λ-bond corresponds to increased daughters at resolution level n if and only if exactly one of its vertices is descended from the most recent common ancestor of both vertices. This suggests we sum the contributions over the $1+4+4^2+\ldots+4^{n-1}$ vertices which might serve as a most recent common ancestor to an added λ-bond:

$$c_{\lambda,\tau}(n) \quad = \quad c(n) \quad = \quad \sum_{a=0}^{n-1} (4\tau)^a \times c_2(n,a)\,.$$

(2) For each such vertex we must now sum over all possible λ-bonds with the vertex as most recent common ancestor, allowing for the two different ways in which just one vertex of the λ-bond is connected to the ancestor, and including the probability that the λ-bond in question is open:

$$c_2(n,a) \quad = \quad \sum_{b=1}^{n-a} 2\tau^b \lambda \left(c_3(n,a,b) + c_4(n,a,b)\right) .$$

(3) The vertex *not* connected to the ancestor now contributes a further number of direct descendants at level n. We must add a factor ensuring that there is no connection between vertex and ancestor:

$$c_3(n,a,b) \quad = \quad \left(\sum_{\beta=0}^{b-1} \tau^\beta (1-\tau)\right) (4\tau)^{n-a-b} \quad = \quad \left(1-\tau^b\right) (4\tau)^{n-a-b} .$$

(4) Furthermore we must add in descendants at level n deriving from ancestors of this vertex on the (broken) line of descent from the most recent common ancestor. Again we must add a factor ensuring that there is no connection between vertex and ancestor:

$$c_4(n,a,b) \quad = \quad \left(\sum_{\beta=1}^{b-1} \tau^\beta (1-\tau) \sum_{\gamma=1}^{\beta} 3\tau (4\tau)^{\gamma-1}\right) (4\tau)^{n-a-b}$$

(the sum vanishes if $b = 1$).

Combining these formulae gives an expression for $c_{\lambda,\tau}(n)$, the excess over $(4\tau)^n$. This can be evaluated in closed form, but is complicated and unenlightening. However we may now consider the smooth critical path $\{(\lambda(s), \tau(s)) : s \geq 0\}$ such that $(4\tau)^n + c_{\lambda(s),\tau(s)}(n) = 1$. Taking $\lambda(s) = s$, we deduce using calculus that the slope of this path at $\lambda = 0$, $\tau = 1/4$ is given by

n	100	200	300	400	500
slope	-7.1752	-7.15899	-7.15361	-7.15092	-7.1493

This provides us with an indication of an upper bound for the infinite-cluster region.

B Percolation: More on transition from many to unique infinite clusters

Following [14, §3], it is possible to use invasion percolation to establish that if the infinite cluster in $\mathbb{Q}_{d,0}$, respectively $\mathbb{Q}_d(\mathbf{o})$, is unique at parameters λ, τ then it is also unique at parameters λ', τ' whenever $\lambda' \geq \lambda$ and $\tau' \geq \tau$. Here is a sketch of the argument.

We first realize the percolation by attaching independent uniformly distributed marks to each bond. Supposing (for example) that $\lambda' < \tau'$, for τ-bonds the relevant mark M is uniformly distributed over $(0, 1)$ whereas for λ-bonds it is uniformly distributed over $0, \tau'/\lambda'$, and the bond is open in the (λ', τ')-percolation exactly when $M < \tau'$.

The invasion cluster $\mathcal{I}_{\mathbf{u}}$ at a vertex $\mathbf{u}$ is constructed by setting $\mathcal{V}^0_{\mathbf{u}} = \{\mathbf{u}\}$ and $\mathcal{I}^0_{\mathbf{u}} = \emptyset$, and constructing $\mathcal{I}^{n+1}_{\mathbf{u}}$ by adding the bond with least mark which is not in $\mathcal{I}^n_{\mathbf{u}}$ but is incident to $\mathcal{V}^n_{\mathbf{u}}$, and adding the non-incident vertex to $\mathcal{V}^n_{\mathbf{u}}$ to form $\mathcal{V}^{n+1}_{\mathbf{u}}$. Finally set $\mathcal{I}_{\mathbf{u}} = \mathcal{I}^{\infty}_{\mathbf{u}}$.

Semi-transitivity is used to show that any percolation allowing positive chance of infinite clusters is *uniform*: sufficiently large metric balls have probability uniformly close to 1 of intersecting an infinite cluster.

If $\mathbf{u}$ has an infinite cluster C' under (λ', τ')-percolation then C' must contain $\mathcal{I}_{\mathbf{u}}$.

However we can use uniform percolation to argue that with arbitrarily high probability one of the $\mathcal{I}^n_{\mathbf{u}}$ will eventually intersect an infinite cluster C under (λ, τ)-percolation (adapt the proof of [14, Lemma 3.2] and see [14, Proof of Proposition 3.1]).

The invasion cluster is thus used as an intermediary to show that C' intersects (and so contains) the infinite cluster C established at the lower level of percolation.

If the (λ, τ)-percolation exhibits only a single infinite cluster then it follows C' must intersect and therefore be equal to all other infinite clusters under (λ', τ')-percolation, and so the (λ', τ')-percolation must also exhibit only a single infinite cluster.

Other University of Warwick Department of Statistics Research Reports authored or co–authored by W.S. Kendall.

161: The Euclidean diffusion of shape.
162: Probability, convexity, and harmonic maps with small image I: Uniqueness and fine existence.
172: A spatial Markov property for nearest–neighbour Markov point processes.
181: Convexity and the hemisphere.
202: A remark on the proof of Itô's formula for C^2 functions of continuous semimartingales.
203: Computer algebra and stochastic calculus.
212: Convex geometry and nonconfluent Γ-martingales I: Tightness and strict convexity.
213: The Propeller: a counterexample to a conjectured criterion for the existence of certain convex functions.
214: Convex Geometry and nonconfluent Γ-martingales II: Well–posedness and Γ-martingale convergence.
216: *(with E. Hsu)* Limiting angle of Brownian motion in certain two–dimensional Cartan–Hadamard manifolds.
217: Symbolic Itô calculus: an introduction.
218: *(with H. Huang)* Correction note to "Martingales on manifolds and harmonic maps."
222: *(with O.E. Barndorff-Nielsen and P.E. Jupp)* Stochastic calculus, statistical asymptotics, Taylor strings and phyla.
223: Symbolic Itô calculus: an overview.
231: The radial part of a Γ-martingale and a non-implosion theorem.
236: Computer algebra in probability and statistics.
237: Computer algebra and yoke geometry I: When is an expression a tensor?
238: *Itovsn3*: doing stochastic calculus with *Mathematica*.
239: On the empty cells of Poisson histograms.
244: *(with M. Cranston and P. March)* The radial part of Brownian motion II: Its life and times on the cut locus.
247: Brownian motion and computer algebra (Text of talk to BAAS Science Festival '92, Southampton Wednesday 26 August 1992, with screenshots of illustrative animations).
257: Brownian motion and partial differential equations: from the heat equation to harmonic maps (Special invited lecture, 49th session of the ISI, Firenze).
260: Probability, convexity, and harmonic maps II: Smoothness via probabilistic gradient inequalities.
261: *(with G. Ben Arous and M. Cranston)* Coupling constructions for hypoelliptic diffusions: Two examples.
280: *(with M. Cranston and Yu. Kifer)* Gromov's hyperbolicity and Picard's little theorem for harmonic maps.
292: Perfect Simulation for the Area-Interaction Point Process.
293: *(with A.J. Baddeley and M.N.M. van Lieshout)* Quermass-interaction pro-

cesses.
295: On some weighted Boolean models.
296: A diffusion model for Bookstein triangle shape.
301: COMPUTER ALGEBRA: an encyclopaedia article.
308: Perfect Simulation for Spatial Point Processes.
319: Geometry, statistics, and shape.
321: From Stochastic Parallel Transport to Harmonic Maps.
323: *(with E. Thönnes)* Perfect Simulation in Stochastic Geometry.
325: *(with J.M. Corcuera)* Riemannian barycentres and geodesic convexity.
327: Symbolic Itô calculus in *AXIOM*: an ongoing story.
328: *Itovsn3* in *AXIOM*: modules, algebras and stochastic differentials.
331: *(with K. Burdzy)* Efficient Markovian couplings: examples and counterexamples.
333: Stochastic calculus in *Mathematica*: software and examples.
341: Stationary countable dense random sets.
347: *(with J. Møller)* Perfect Metropolis-Hastings simulation of locally stable point processes.
348: *(with J. Møller)* Perfect implementation of a Metropolis-Hastings simulation of Markov point processes
349: *(with Y. Cai)* Perfect simulation for correlated Poisson random variables conditioned to be positive.
350: *(with Y. Cai)* Perfect implementation of simulation for conditioned Boolean Model via correlated Poisson random variables.
353: *(with C.J. Price)* Zeros of Brownian Polynomials.
371: *(with G. Montana)* Small sets and Markov transition densities.
382: *(with A. Brix)* Simulation of cluster point processes without edge effects.
391: *(with E. Thönnes, A. Bhalerao, R.G. Wilson)* A Bayesian approach to inferring vascular tree structure from 2D imagery.
392: *(with A. Bhalerao, E. Thönnes, R.G. Wilson)* Inferring vascular tree structure from 2D and 3D imagery.
402: *(with R.G. Wilson)* Ising models and multiresolution quad-trees.

Also see the following related preprints

317: E. Thönnes: Perfect Simulation of some point processes for the impatient user.
334: M.N.M. van Lieshout and E. Thönnes: A Comparative Study on the Power of van Lieshout and Baddeley's J-function.
359: E. Thönnes: A Primer on Perfect Simulation.
366: J. Lund and E. Thönnes: Perfect Simulation for point processes given noisy observations.

If you want copies of any of these reports then please email your requests to the secretary using `statistics@warwick.ac.uk` (mail address: the Department of Statistics, University of Warwick, Coventry CV4 7AL, UK).