Random graphs with finite expected degree and Local algorithms on lattices

Aranka Hrušková

CEU eTD Collection

A thesis submitted in fulfilment of the requirements for the degree
of Practical Mathematics Specialization Program

Central European University

Declaration of Authorship

I, Aranka Hrušková, declare that this thesis entitled *Random graphs with finite expected degree and Local algorithms on lattices* and the work presented in it are my own.

I confirm the following

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.

CEU eTD Collection

Abstract

In the first part of the thesis, motivated by the recent paper [1] by Backhausz and Szegedy which unifies and generalises the concepts of dense (graphon) and local-global (graphing) convergence, we prove a number of inequalities relating the degree sequence of a graph and the operator norms of its adjacency matrix. We then use these to explore the properties of the Erdős-Rényi model $\mathcal{G}(n, \frac{c}{n-1})$, in which for each vertex its expected degree is a constant c independent of n.

The second part of the thesis constructs factor of iid algorithms for obtaining balanced orientation of the triangular lattice and of the square grid.

CEU eTD Collection

1 Random graphs with finite expected degree

The central object of the relatively young field of graph limit theory is a sequence $(G_n)_{n=1}^{\infty}$ of finite graphs, for which we seek to find a limit object. Two main notions of convergence (based on different methods of sampling small subgraphs from large graphs) and their respective limit objects have been described and explored [5]. However, in terms of edge-density, these are useful (or even just defined) in very different environments, that is in a very dense or in a particularly sparse sequence. The lack of applicable convergence notions for graph sequences of intermediate density (e.g. the sequence of hypercubes) has lead to several attempts to bridge the two worlds and fill in the in-between space. In this thesis, we are most interested in the approach taken by Backhausz and Szegedy in [1], some of whose key points we summarise below.

Definition (P-operator)**.** *Let $(\Omega, \mathcal{A}, \mu)$ be a probability space. Then we call a linear operator $A : L^\infty(\Omega) \to L^1(\Omega)$ a P-operator if the norm*

$$\|A\|_{\infty\to 1} = \sup_{v\in L^\infty(\Omega)} \frac{\|vA\|_1}{\|v\|_\infty}$$

is finite.

A real $n \times n$ matrix A is an easy example of a P-operator, where $\Omega = [n]$, μ is the uniform measure, and the L-spaces are formed simply by real vectors. The action of an operator on an L-space determines not only its norms, but gives rise to sets of probability measures as follows.

Definition (k-profiles of matrices)**.** *Let k be a positive integer. The k-profile $\mathcal{S}_k(A)$ of an $n \times n$ matrix A is the set of all discrete probability measures on $\mathbb{R}^{2k}$ of the form*

$$\frac{1}{n}\sum_{j=1}^{n} \delta_{((v_1)_j,\ldots,(v_k)_j,(v_1A)_j,\ldots,(v_kA)_j)}$$

where $v_1, \ldots, v_k$ are real vectors with entries in $[-1, 1]$ and δ_x denotes the Dirac measure concentrated on $x \in \mathbb{R}^{2k}$.

We can associate such k-profiles to any P-operator, not just matrices, and they will always be sets of probability measures on $\mathbb{R}^{2k}$ encoding the dynamical properties of the operator. One of the main advantages of storing information about a P-operator in its profiles is that through them, we can compare even operators acting on completely different L-spaces.

In particular, Backhausz and Szegedy metrise the dynamical properties of P-operators via the Hausdorff distance d_H, which is a general pseudometric for subsets of any metric space (M, d). In particular, for subsets X, Y of M, we have got $d_H(X, Y) = \max\{\sup_{x\in X} \inf_{y\in Y} d(x, y), \sup_{y\in Y} \inf_{x\in X} d(x, y)\}$.

Definition (Metrisation of action convergence)**.** *Let A, B be P-operators. Then their action distance is*

$$d_M(A, B) = \sum_{k=1}^{\infty} \frac{d_H(\mathcal{S}_k(A), \mathcal{S}_k(B))}{2^k},$$

where the Hausdorff distance is based on the Lévy-Prokhorov metric for probability measures $\mathcal{P}(\mathbb{R}^n)$ given by $d_{LP}(\eta, \mu) = \inf\{\epsilon > 0 : \eta(U) \le \mu(U^\epsilon) + \epsilon$ and $\mu(U) \le \eta(U^\epsilon) + \epsilon$ for every Borel set $U \subset \mathbb{R}^n\}$.

Note that for any two probability measures $\eta, \mu \in \mathcal{P}(\mathbb{R}^n)$, we have $d_{LP}(\eta, \mu) \le 1$, hence also for any two P-operators A, B their distance satisfies $d_M(A, B) \le \sum_{k=1}^{\infty} \frac{1}{2^k} = 1$.

The norm $\|\cdot\|_{\infty\to 1}$ is the smallest from the class of operator norms $\|\cdot\|_{p\to q}$ where $p, q \in [1, \infty]$, and so, unsurprisingly, various non-trivial compactness, completeness and other results are proved in [1] assuming certain boundedness conditions on some larger of these norms (but also just on the operator norm $\|\cdot\|_{\infty\to 1}$).

CEU eTD Collection

Theorem 1 (Sequential compactness)**.** *Let $(A_n)_{n=1}^{\infty}$ be a sequence of P-operators with uniformly bounded $\|\cdot\|_{\infty\to 1}$ norms. Then $(A_n)_{n=1}^{\infty}$ has a Cauchy subsequence with respect to the distance d_M.*

Theorem 2 (Existence of limit object)**.** *Let $p \in [1,\infty)$ and $q \in [1,\infty]$. Let $(A_n)_{n=1}^{\infty}$ be a sequence of P-operators, Cauchy with respect to the distance d_M and with uniformly bounded $\|\cdot\|_{p\to q}$ norms. Then there is a P-operator A such that $\lim_n d_M(A_n, A) = 0$ and $\|A\|_{p\to q} \le \limsup_n \|A_n\|_{p\to q}$.*

Remark 3. *Let us note that Theorem 2 is in a sense as good as it gets, since we can find Cauchy sequences with uniformly bounded $\|\cdot\|_{\infty\to q}$ norms, which, however, do not action converge to any P-operator. Let $(S_n)_n$ be the sequence of the $n \times n$ adjacency matrices of stars, that is n-vertex graphs with one vertex of degree $n-1$ and $n-1$ vertices of degree one. In general, the norm $\|\cdot\|_{\infty\to 1}$ of an adjacency matrix is the average degree of the corresponding graph (see Subsection 1.1), and so here we get $\|S_n\|_{\infty\to 1} = \frac{1}{n}(1\times(n-1)+(n-1)\times 1) < 2$ for any positive integer n. However, since for any probability measure μ on $[-1,1]$ there is a sequence $(\frac{1}{n}\sum_i^n \delta_{x_i})_n$ of discrete probability measures with $\lim d_{LP}(\mu, \frac{1}{n}\sum_i^n \delta_{x_i}) = 0$, we have $\lim_n \mathcal{S}_1(S_n) = \mathcal{P}([-1,1]) \times [-1,1]$, while there cannot be any P-operator A such that $d_H(\mathcal{S}_1(A), \mathcal{P}([-1,1]) \times [-1,1]) = 0$. There are two ways to see this*

1.1 Bounds on the operator norms $\|\cdot\|_{p\to q}$ of adjacency matrices

We have seen above two of the many theorems based on assumptions of uniform boundedness of certain operator norms. The standard inequalities [1, 6] $\max\{\bar{d}, \sqrt{\Delta(G)}\} \le \|A(G)\|_{2\to 2} \le \Delta(G)$, where $A(G)$ is the adjacency matrix of a graph G, $\Delta(G)$ is its maximum and $\bar{d}$ average degree, though, cannot always be used in a desired context. Below, we derive a number of upper and lower bounds on general operator norms $\|\cdot\|_{p\to q}$.

Notation. We denote by $\underline{d}$ the degree sequence $(\deg(i))_{i=1}^n$ of a graph G on vertex set $[n]$, and by $\|\underline{d}\|_q = \sqrt[q]{\frac{1}{n}\sum_i^n \deg(i)^q}$ its q-norm. $A = A(G)$ is the $n \times n$ adjacency matrix of G.

1.1.1 Lower bounds

Lemma 4. *For every $p, q \in [1,\infty)$*

$$\|A\|_{p\to q} \ge \|\underline{d}\|_q \ge \bar{d}.$$

Proof. $\|A\|_{p\to q} \ge \|\underline{d}\|_q$ is certified by the all-ones vector $\mathbb{1}$:

$$\|A\|_{p\to q} \ge \frac{\|\mathbb{1}A\|_q}{\|\mathbb{1}\|_p} = \|\mathbb{1}A\|_q = \sqrt[q]{\frac{1}{n}\sum_{i=1}^n (\mathbb{1}A)_i^q} = \sqrt[q]{\frac{1}{n}\sum_{i=1}^n \Big(\sum_j a_{ij}\Big)^q} = \sqrt[q]{\frac{1}{n}\sum_{i=1}^n \deg(i)^q} = \|\underline{d}\|_q.$$

Also the function $f(x) = x^q$ is convex on $[0,\infty)$, so by Jensen's inequality

$$\|\underline{d}\|_q = \sqrt[q]{\frac{1}{n}\sum_{i=1}^n \deg(i)^q} = \sqrt[q]{\frac{1}{n}\sum_{i=1}^n f(\deg(i))} \ge \sqrt[q]{f\Big(\frac{1}{n}\sum_i^n \deg(i)\Big)} = \frac{1}{n}\sum_{i=1}^n \deg(i) = \bar{d}.$$

□

1.1.2 Upper bounds

1.1.3 Exact results

Lemma 5. *For every $p \in (1,\infty)$*

$$\|A\|_{p\to 1} = \|\underline{d}\|_{\frac{p}{p-1}}.$$

CEU eTD Collection

Proof.

$$\|A\|_{p\to 1} = \sup_{v\in\mathbb{R}^n} \frac{\|vA\|_1}{\|v\|_p}$$

but as A is non-negative, this is equal to

$$\sup_{v\in[0,\infty)^n} \frac{\|vA\|_1}{\|v\|_p} = \sup_{v\in[0,\infty)^n} \frac{\frac{1}{n}\sum_i^n \sum_{ij\in E(G)} v_j}{\sqrt[p]{\frac{1}{n}\sum_i^n v_i^p}} = \sup_{v\in[0,\infty)^n} \Big(\frac{1}{n}\Big)^{1-\frac{1}{p}} \frac{\sum_i^n v_i \deg(i)}{\sqrt[p]{\sum_i^n v_i^p}}$$

$$= \Big(\frac{1}{n}\Big)^{1-\frac{1}{p}} \sup_{v\in[0,\infty)^n} \{\sum_i^n v_i \deg(i) : \sqrt[p]{\sum_i^n v_i^p} = 1\}.$$

We use Lagrangians to find the supremum:

$$P : \text{minimise } -\sum_i^n x_i \deg(i)$$

$$\text{such that } \sum_i^n x_i^p = 1, x \in [0,\infty)^n.$$

The Lagrangian of our problem is

$$L(x,\lambda) = -\sum_i^n x_i \deg(i) - \lambda(\sum_i^n x_i^p - 1)$$

$$= -\sum_i^n (x_i \deg(i) + \lambda x_i^p) + \lambda.$$

Now we wish to minimise $L(x,\lambda)$ for a fixed λ so that we can use the Lagrangian sufficiency theorem. For $\lambda < 0$, $L(x,\lambda)$ is minimised (on $[0,\infty)^n$) by $x_i = \sqrt[p-1]{\frac{\deg(i)}{-\lambda p}}$ for every $i \in [n]$, so let us take $\lambda < 0$ such that

$$1 = \sum_i^n x_i^p = \sum_i^n \Big(\frac{\deg(i)}{-\lambda p}\Big)^{\frac{p}{p-1}} = \frac{1}{(-\lambda)^{\frac{p}{p-1}}} \sum_i^n \Big(\frac{\deg(i)}{p}\Big)^{\frac{p}{p-1}}$$

that is

$$-\lambda = \Big(\frac{1}{p^{\frac{p}{p-1}}} \sum_i^n \deg(i)^{\frac{p}{p-1}}\Big)^{\frac{p-1}{p}}$$

$$= \frac{1}{p}\sum_i^n \deg(i)^{\frac{p}{p-1}}$$

□

1.2 The Erdős-Rényi model $\mathcal{G}(n, \frac{c}{n-1})$

Using the results from the previous subsection, induction on the moments of traces of the adjacency matrices and the Borel-Cantelli lemma, we show that for a fixed c, the sequence $(A_n)_n$ of adjacency matrices is such that each of its subsequences has a further subsequence which is Cauchy with respect to d_M. Also for $c \leq 1$ the graph sequence converges locally-globally, but in fact that doesn't yet necessairly mean that the adjacency matrices converge in d_M because the theorem in [1] relating action and local-global convergence uses uniform boundedness of maximum degree in a sequence.

CEU eTD Collection

2 Local algorithms on lattices

Factor of iid balanced orientation of the triangular lattice

Let G be the infinite triangular grid. We want to construct a balanced orientation of its edges by a factor of iid process.

In the beginning, we have got independent uniformly distributed $[0,1]$-labels on $V(G)$.

Step 1: Use the first digit after the decimal point to two-colour the vertices of G. Colour a vertex red if the first digit is 0, 1, 2, 3 or 4 and colour it blue if the digit is 5, 6, 7, 8 or 9. This is equivalent to site percolation with probability $1/2$, and so by [2] with probability 1 all monochromatic components are finite.
Step 2: Colour all edges with both endpoints red by red colour, all edges with both endpoints blue by blue colour, and leave the edges with one endpoint red and one blue uncoloured.
Step 3, definition: For an edge $uv \in E(G)$, the *guards* of uv are the two common neighbours of u and v.
Step 4, definition: For a connected monochromatic component C, let C^+ be the unique monochromatic component surrounding C.
Step 5, definition: Given a monochromatic component C, let the *reduced boundary* $\partial^- C$ be the set of edges in C which have exactly one guard in C^+.
Step 6, claim: Every vertex $v \in V(G)$ is incident to exactly 0 or 2 or 4 edges in the reduced boundary $\partial^- C(v)$.
Proof. Let us analyse the possible scenarios based on the number $|V(C^+) \cap N(v)|$ of neighbours of v which are in C^+.

v **has 0 neighbours in** C^+ Then no edge incident to v has a guard in C^+, and so no edge incident to v is in the reduced boundary $\partial^- C$.

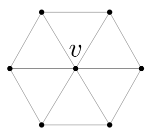

v **has 1 neighbour in** C^+ Then exactly two edges incident to v have this neighbour as a guard, and so exactly two edges incident to v are in $\partial^- C$.

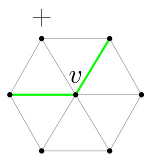

v **has 2 neighbours in** C^+ All configurations of when v has exactly two neighbours in C^+ are sketched below.

v **has 3 neighbours in** C^+ All configurations of when v has exactly three neighbours in C^+ are sketched below. In the last setup no edge has exactly one guard in C^+.

v **has 4 neighbours in** C^+ All configurations for when v has four neighbours in C^+ are sketched below.

CEU eTD Collection

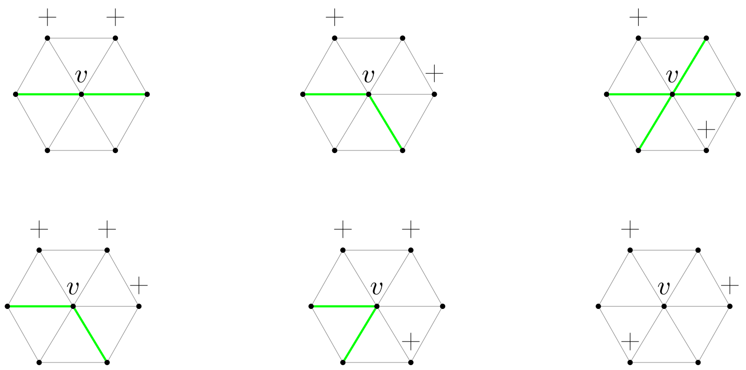

v **has 5 neighbours in** C^+ The only edge incident to v which may be in C has both guards in C^+, so no edge incident to v is in $\partial^- C$.

v **has 6 neighbours in** C^+ Then all edges incident to v have both guards in C^+, and hence none of the edges incident to v is in $\partial^- C$.

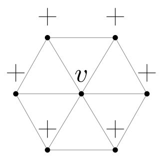

Step 7: For every monochromatic component C, if an edge $e \in E(C)$ has both guards in C^+, recolour e to be purple. Note that no edges in the reduced boundary $\partial^- C$ will be recoloured and that colourless edges are still colourless.

Step 8, orient the insides: Let C be a coloured component, that is, the vertices of C are either a) all blue or b) all red and the edges of C are either a) all blue or purple or b) all red or purple. Orient all the edges of C except for the purple ones and except for the ones in $\partial^- C$ as follows. Let x be the vertex in C with the largest label. Let $y \in V(C)$ be the neighbour of x with the largest label among the vertices $V(C) \cap N(x)$. Orient xy to go from x to y. Now propagate this orientation so that all triangles in $C \setminus \partial^- C$ are strongly oriented, i.e. from an onlooker's perspective, either all horizontal edges will go from left to right, all southeast-northwest up and all northeast-southwest edges will go down, *or* all horizontal edges will go from right to left all southwest-northeast edges will go up and all northwest-southeast edges down.

Step 9, orient the uncoloured and purple: Propagate the orientation of C^+ to the purple edges of C and to the uncoloured edges bridging C and C^+.

Step 10, orient reduced boundary: If the reduced boundary $\partial^- C$ of C is non-empty, proceed as follows. If the orientations of C and C^+ are compatible, propagate this orienation to $\partial^- C$. If the orientations are the opposite, fix an oriented Eulerian cycle for every component of $\partial^- C$.
(e.g.: In each component of $\partial^- C$ find the vertex x with the largest label. Let y be the vertex in $V(\partial^- C) \cap N(x)$ with the largest label. Let xy be oriented from x to y. Propagate this orientation until you get back to x or until you reach a vertex z such that z has three incident edges which are in $\partial^- C$ and not yet oriented. In the latter case, let these edges be zu_1, zu_2, and zu_3. Let $u \in \{u_1, u_2, u_3\}$ be the vertex with the largest label of the three. Orient zu from z to u and continue as before.)

Step 11, claim: The resulting orientation is balanced at every vertex $v \in V(G)$.

Proof. If v has 0 or 6 neighbours in $C(v)^+$, then we are done because the pairs of opposite incident edges will

CEU eTD Collection

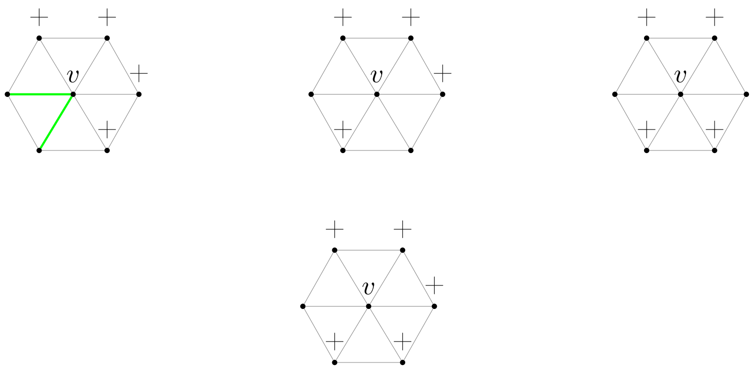

have the opposite orientations with respect to v (i.e. one incoming and one outgoing). If $C(v)$ contains more than one vertex and the inner orientations of $C(v)$ and $C(v)^+$ are the same, we are done exactly as in the previous case. So it remains to analyse what happens when $C(v)$ and $C(v)^+$ have the opposite orientations and v has 1-5 neighbours in $C(v)^+$. We will use the drawings from the proof of Step 6 and we will only need to pay attention to whether the non-boundary edges cancel each other out because we put an Eulerian orientation on ∂^-. The reduced boundary is indicated in green.

v **has 1 neighbour in** C^+ Then v has no incident purple edges.

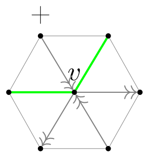

v **has 2 neighbours in** C^+ All configurations of when v has exactly two neighbours in C^+ are sketched below.

v **has 3 neighbours in** C^+ By Step 9 (propagating orientation of C^+ to purple and uncoloured edges), the third setup is the same case as when v has all six neighbours in C^+

v **has 4 neighbours in** C^+ By Step 9 (propagating orientation of C^+ to purple and uncoloured edges), the second and the third setup are the same case as when v has all six neighbours in C^+, while the first is the same as the second in the previous paragraph.

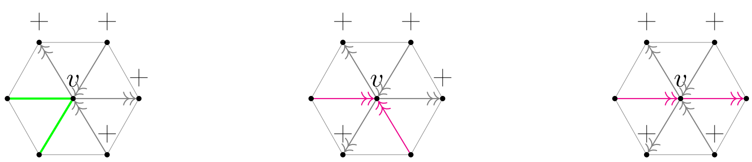

v **has 5 neighbours in** C^+ By Step 9 (propagating orientation of C^+ to purple and uncoloured edges), this is the same case as when v has all six neighbours in C^+.

CEU eTD Collection

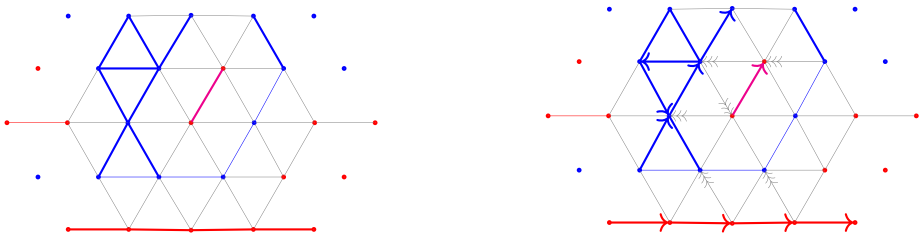

Figure 1: Example of partially orienting coloured triangular grid. Reduced boundaries, which remain to be oriented, are indicated in thin coloured lines.

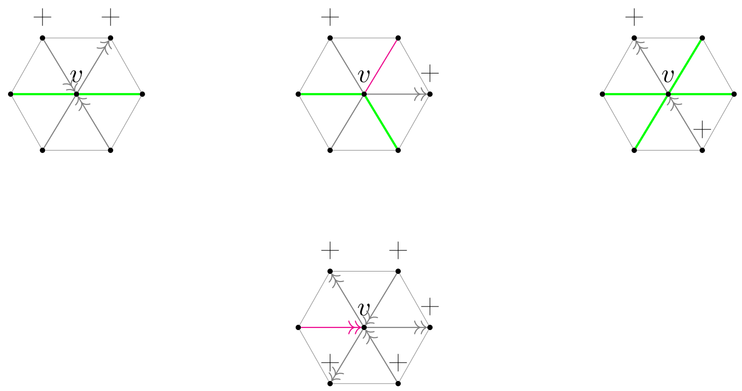

Factor of iid balanced orientation of $\mathbb{Z}^2$

In the beginning, we have got independent uniformly distributed $[0,1]$-labels on $V(\mathbb{Z}^2)$.

Step 1: The key in our construction is that we partition $\mathbb{Z}^2$ into finite clusters such that each cluster C is surrounded by a *unique* cluster C^+. We can get finite clusters simply by site percolation with probability $p = 1/2$, but that does not give us the desired uniqueness of the surrounding component. There are two ways to deal with this drawback. First, we could use the fact that the critical probabilities for site percolations in matching pairs sum up to one (see Theorem 14 in [2]). We would then colour a site blue with the critical probability of the square lattice $\Lambda_\square$ and red with the critical probability of $\Lambda_\boxtimes$. When two red clusters would have Euclidean distance $\sqrt{2}$, they would be merged (following the idea of diagonals silently being available between red sites). It was proven in [7] that at this critical regime, with probability one there is no infinite monochromatic cluster of either colour (though expected sizes of blue as well as red clusters will be infinite). Interestingly, the related conjecture of there being no percolation at criticality (site or bond) in $\mathbb{Z}^d$ has not been settled for $3 \leq d \leq 18$ (see [3] for a partial progress). The advantage of this circumvention of the original obstacle is that diagonal merging of components will occur in one colour only, which may prove helpful for extending the fiid balanced-orienting to fiid two-coloured balanced-orienting (i.e. fiid action of the free abelian group with two generators). The disadvantage is that the critical probability $p_c^s(\Lambda_\square)$ is currently only known approximately.

The second way to overcome the non-uniqueness obstacle is to colour the vertices of $\mathbb{Z}^2$ uniformly randomly red and blue and whenever a square arises whose two opposite vertices are blue while the other opposite

CEU eTD Collection

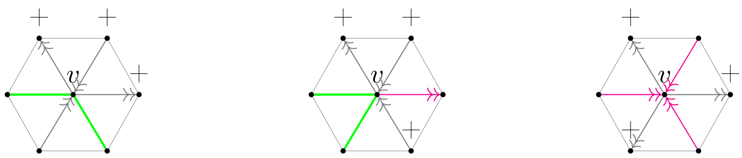

pair is red, choose uniformly randomly to merge either the blue components or the red components. We can describe such a process by site percolation on the augmented square lattice Λ_T called tetrakis square tiling: add the discrete set $\mathbb{Z}^2 + (1/2, 1/2)$ to the set of vertices (i.e. the square faces acquire centres) and connect each such non-integer vertex to the four corners of the square surrounding it. All the faces of the tetrakis square tiling are triangular, so the lattice is self-matching. Hence, using Theorem 14 from [2] again, we see that its critical probability is $1/2$. Applying the proof method of Theorem 8 in [2] (taking a large enough polygon centred at the origin so that infinite paths leave four of its sides and recalling that there is at most one infinite open cluster), we conclude that at the critical probability $p_c^s = 1/2$ all clusters are finite. These new non-integer sites play exactly the role of the random mergers – in all other situations than the surrounding square having (clockwise) red-blue-red-blue vertices, the non-integer site is either a singleton cluster or it makes a neighbouring cluster bigger only by itself.

Step 2: For a vertex x with label $l(x)$, let $l_2(x)$ be obtained from $l(x)$ by keeping only every second digit of the decimal expansion, and let $l_1(x)$ be what remains of $l(x)$ after deleting the digits belonging to $l_2(x)$ (e.g. for $l = 0.1234567891...$, $l_2 = 0.24681...$ and $l_1 = 13579...$). To obtain the first percolation model, colour a site x blue if $l(x) < p_c^s(\Lambda_\square) \doteq 0.59...$ and colour it red if $l(x) \geq p_c^s(\Lambda_\square) = 1 - p_c^s(\Lambda_\boxtimes)$. Finally, add auxiliary red diagonals between any two red vertices in Euclidean distance $\sqrt{2}$ to indicate which clusters count as a single one. To obtain the second model, colour a site x blue if $l_1(x) < p_c^s(\Lambda_T) = 1/2$ and colour it red if $l_1(x) \geq p_c^s(\Lambda_T) = 1/2$. Now use the l_2-s to two-colour the 'diagonal switches'. For each square face of $\mathbb{Z}^2$, let the non-integer centre vertex inherit the colour of that of the four corners with the largest l_2.

Step 3, definition: For an edge $(m, n)(m, n+1) \in E(\mathbb{Z}^2)$, the *guards* of $(m, n)(m, n+1)$ are the four vertices $(m-1, n), (m-1, n+1), (m+1, n)$, and $(m+1, n+1)$. Similarly, the *guards* of $(m, n)(m+1, n) \in E(\mathbb{Z}^2)$ are $(m, n-1), (m, n+1), (m, n-1)$, and $(m, n+1)$.

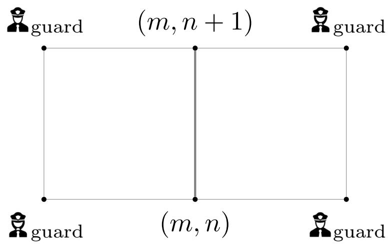

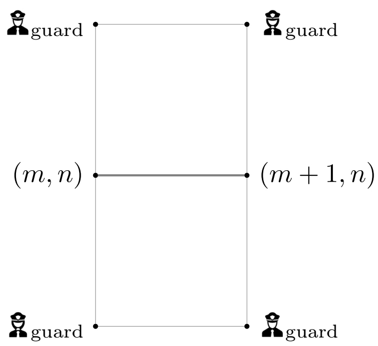

Step 4, definition: For a connected monochromatic component C, let C^+ be the unique monochromatic component (of the other colour than C) surrounding C.

Step 5, definition: Given a monochromatic component C, let its *boundary* ∂C be the following set

CEU eTD Collection

of edges.

$$\begin{aligned}\partial C = \{uv \in \mathbb{Z}^2 | \{u,v\} \cap V(C) \neq \emptyset, \{u,v\} \cap V(C^+) = \emptyset, \text{ and} \\ \text{either } uv \text{ has exactly one of its guards in } C^+ \\ \text{or it has exactly two guards in } C^+ \text{ and these two guards are adjacent}\}\end{aligned}$$

Note that the boundary of C does not touch the surrounding component C^+, but it may touch a component C^- which is itself being surrounded by C. Likewise, the boundary of C^+ may touch a vertex of C.

Step 6, claim: Every vertex $v \in V(\mathbb{Z}^2)$ is incident to exactly 0 or 2 or 4 edges in the boundary $\partial C(v)$. It is also incident to exactly 0 or 2 or 4 edges in the boundary $\partial C(v)^+$.
Proof. Let us first prove the statement about incidence to $\partial C(v)$. Suppose on the contrary that v is incident to exactly three edges from $\partial C(v)$. Say these are vb, uv, and vw as pictured in Figure 2. vb must have at

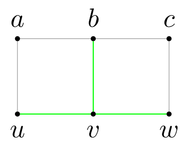

Figure 2: The boundary $\partial C(v)$ is indicated in green.

least one guard in $C(v)^+$. Note that vb, uv and vw being in $\partial C(v)$ means that b, u, and w are not in $C(v)^+$. Hence, by definition of ∂C, vb must have exactly one guard in $C(v)^+$. Wlog, let a be in $C(v)^+$ and c not. Then considering $uv \in \partial C$, we conclude that neither x nor y can be in $C(v)^+$.

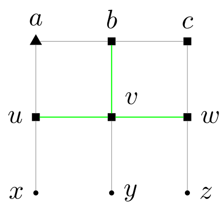

Figure 3: Vertices in $C(v)^+$ are drawn as triangles, while vertices not in $C(v)^+$ are drawn as squares.

Then z also isn't in $C(v)^+$, otherwise vy would be in $\partial C(v)$. However, that means that vw has no guards in $C(v)^+$, and so it isn't in $\partial C(v)$. This contradiction proves that v cannot be incident to exactly three edges in $\partial C(v)$.

Now suppose that of the four edges incident to v, exactly one is from $\partial C(v)$, say vb. Then there must be two adjacent guards of vb which are not in $C(v)^+$. Wlog, let c and w be such guards.

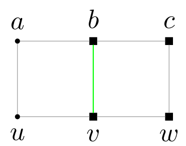

Figure 4: The boundary $\partial C(v)$ is indicated in green.

CEU eTD Collection

By assumption, vw is not in $\partial C(v)$, so neither y nor z can be in $C(v)^+$. Continuing similarly, we conclude that neither x nor u can be in $C(v)^+$.

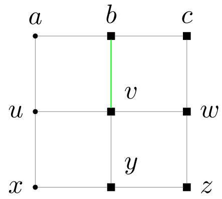

Figure 5: Vertices not in $C(v)^+$ are drawn as squares.

Hence, a must be in $C(v)^+$ in order for bv to be in $\partial C(v)$. But then uv is in $\partial C(v)$ too, contradicting that only one of the edges incident to v is in the boundary $\partial C(v)$.

We have proven that v is incident to exactly 0 or 2 or 4 edges in $\partial C(v)$. The proof for the same statement with $\partial C(v)$ replaced by $\partial C(v)^+$ carries through virtually unchanged, except for any usage of $C(v)^+$ being replaced by C^++.

Step 7, definition: Given a monochromatic component C, define its set $P(C)$ of purple edges to be the edges which have some guards in C^+ but are not in ∂C.

$$P(C) = \{uv \in E(\mathbb{Z}^2) | \{u, v\} \cap V(C) \neq \emptyset, \{u, v\} \cap V(C^+) = \emptyset, \text{ and}$$
$$\text{either } uv \text{ has at least three guards in } C^+$$
$$\text{or it has two guards in } C^+ \text{ and these are not adjacent}\}$$

Let the *inner* edges of C be those with no guards in C^+ and either both endpoints in C or one endpoint in C and one in a component C^- surrounded by C.

Step 8, orient the insides: Let C be a monochromatic component, that is, the vertices of C are either a) all blue or b) all red. Orient all the inner edges of C (if there are any) as follows. Let xy be the inner edge of C with the largest sum of the labels of its endpoints. Now orient xy to go from the endpoint with the larger label to the endpoint with the smaller label. Propagate this orientation so that in the (possibly not connected) subgraph induced by the inner edges of C, all squares are strongly oriented. If an edge uv we need to orient is not a part of a full square, we proceed as if it were, i.e. for any perpendicular edge vw we require that exactly one of uv, vw is oriented into v and exactly one from v. (So at the end, from an onlooker's perspective, the horizontal edges alternate between going left and right, and the vertical edges alternate between going up and down. Also squares alternate in a chessboard-like manner between having their border oriented clockwise and counterclockwise.)

Step 9, orient the purple edges: Propagate the orientation of (the inner edges of) C^+ to the purple edges of C.

Step 10, orient the boundary: If the boundary ∂C of C is non-empty, proceed as follows. If the orientations of C and C^+ are compatible, propagate this orienation to ∂C. If the orientations are the opposite, fix an oriented Eulerian cycle for every component of ∂C.
(e.g.: In each component of ∂C find the vertex x with the largest label. Let y be the vertex in $V(\partial C) \cap N(x)$ with the largest label. Let xy be oriented from x to y. Propagate this orientation until you get back to x or until you reach a vertex z such that z has three incident edges which are in ∂C and not yet oriented. In the

CEU eTD Collection

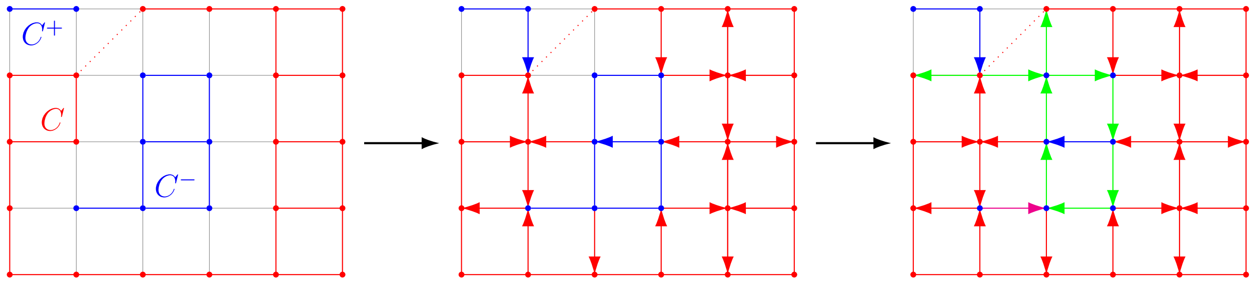

Figure 6: Example of orienting coloured square lattice when the inner orientations of C and C^+ agree and those of C and C^- disagree. Boundaries are indicated in green.

latter case, let these edges be zu_1, zu_2, and zu_3. Let $u \in \{u_1, u_2, u_3\}$ be the vertex with the largest label of the three. Orient zu from z to u and continue as before.)

Step 11, claim: The resulting orientation is balanced at every vertex $v \in V(\mathbb{Z}^2)$.
Proof. Let us analyse the possible scenarios based on how many of the four edges incident to v are elements of the boundary $\partial C(v)$. By Step 6 we know that this number must be even. We will continue to indicate vertices in $C(v)^+$ by triangles and vertices not in $C(v)^+$ by squares. Let us also note in advance that any edge bridging a component C and its surrounding component C^+ must be either in ∂C^+ or in $P(C^+)$ or be an inner edge of C^+. Similarly, any edge bridging C and a vertex not in C^+ must belong either to ∂C or to $P(C)$ or it must be an inner edge of C.

all the four edges incident to v are in $\partial C(v)$ If the (inner) orientations of $C(v)$ and $C(v)^+$ are not compatible, an oriented Eulerian cycle is fixed on the boundary $\partial C(v)$, so the orientation at v is balanced. Otherwise, the four corners (i.e. pairs of perpendicular edges) involving v satisfy that one is oriented into v and one from v (see Step 8), so partitioning the four edges into two corners, we see that the orientation at v is again balanced.

(a) incompatible orientations

(b) compatible orientations

Figure 7: Examples of the neighbourhood of v in case all incident edges are in the boundary of $C(v)$

two of the edges incident to v are in $\partial C(v)$ These two edges (let us call them $v\alpha$ and $v\beta$) can be either perpendicular or aligned. In the first case, we further distinguish whether the fourth vertex of the square induced by $v\alpha$ and $v\beta$ is in $C(v)^+$ or not, while in the second case we note that either both of the two neighbours of v which are neither α nor β are not in $C(v)^+$ or one is and one is not. All four variants with their implications are sketched below.

CEU eTD Collection

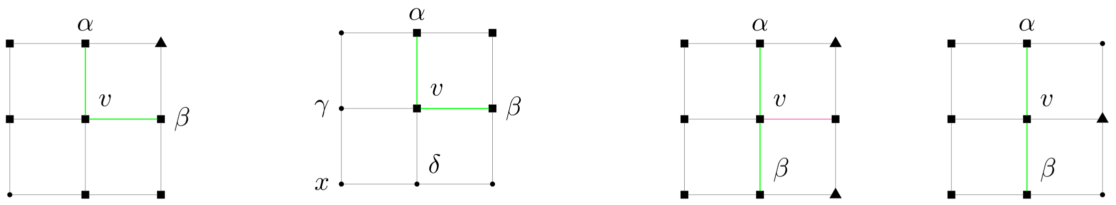

(a) perpendicular boundary (b) aligned boundary

Figure 8: Neighbourhood of v when exactly two incident edges are in the boundary of $C(v)$

In the first of the four cases, the vertex on the opposite side of v than the one in $C(v)^+$ cannot be in $C(v)^+$ too, otherwise all four of the edges incident to v would be in $\partial C(v)$. Thus both the remaining edges are $C(v)$'s inner edges, and recalling that pairs of perpendicular edges within a single inner orientation are balanced at the corner vertex, we see that together with the boundary either having an Eulerian orientation or being part of the same inner orientation, the orientation at v is balanced in the first of the four cases.

In the second case, neither of the two remaining perpendicular edges can be an inner edge of $C(v)$ as both bridge C and C^+ or/and have a guard in C^+. So if both these edges bridge C and C or a C^-, they must both be in $P(C)$, and we are done. If one edge bridges C and C or a C^+ (and is hence purple of C) and the other edge bridges C and C^+, Step 6 tells us the latter one cannot be in ∂C^+, and so it is either inner of C^+, which is orienting-wise the same as if it was in $P(C)$, or it is in $P(C^+)$. That is, the last remaining subcase is when one of the remaining perpendicular edges bridges C and C or a C^-, the other bridges C and C^+, and the latter is in $P(C^+)$. Say wlog that γ, as indicated in Figure 7, is not in C^+ while δ is in C^+. Then x must be in $C^+ +$ in order for $v\delta$ to be in $P(C^+)$. Hence, γ cannot be in C and must be in a C^-, otherwise C and C^{++} would not have been two separate components. But then the four vertices v, γ, δ, and x are each in a different component, which is not possible in our construction.

In the third case, we wlog chose a guard of α to be in C^+, which forced the fourth guard of $v\alpha$ to not be in C^+, which in turn forced the corner-vertex opposite to the one in C^+ to also not be in C^+ (else we would have had three edges in $\partial C(v)$). That means the last corner-vertex must also be in C^+ (because $v\beta$ must have a guard in C^+), and so of the two remaining edges incident to v, one is an inner edge of $C(v)$ and the other is a purple edge of $C(v)$. This is orienting-wise the same situation as when one of these edges is inner of $C(v)$ and the other one is inner of $C(v)^+$ to which we will arrive in the next last case.

In the last case, the left edge of v is an inner edge of $C(v)$, while the right edge bridges $C(v)$ and $C(v)^+$. In general, an edge between C and C^+ is either an inner edge of C^+ or a purple edge of C^+ or element of the boundary ∂C^+. However, by Step 6 we know that v cannot be incident to exactly one edge from ∂C^+, so in this scenario the edge is either inner or purple of C^+. Moreover, if it were to be a purple edge of C^+, both of the two remaining corner vertices in the sketch would have to be in C^{++}. That would in turn mean that both α and β are from a C^-, else C and C^{++} would be merged through α and/or β. But then the four vertices of the square induced by v, the vertex from C^+ and α (or β) each belong to a different monochromatic component, which is impossible by construction of the components, so in fact the edge has to be an inner edge of C^+. Then we fall into the two (by now classic) scenarios of compatible/incompatible inner orientations of C and C^+.

none of the edges incident to v are in $\partial C(v)$ Let us break this down further by how many of the four incident edges are inner edges of C.

CEU eTD Collection

Suppose first that at least one of the four edges is an inner edge of C. Then all the four guards of this edge are not in C^+. But if any of the three remaining vertices in Euclidean distance 1 or $\sqrt{2}$ from v *are* from C^+, then at least one of the edges perpendicular to the inner one is in $\partial C(v)$, contradicting our assumption that none of the four edges are in $\partial C(v)$. So those three vertices must also not be in C^+, meaning that all the eight vertices closest to v, and hence all the guards of the edges incident to v are not in C^+. So *all* the four edges incident to v are in fact inner edges of C.

Now suppose that none of the edges incident to v are inner edges of $C(v)$. Then let's focus on how many of them are in $\partial C(v)^+$. By Step 6, this must be an even number. If all the four edges are in $\partial C(v)^+$, then orientation-wise the situation is the same as in the first case of the whole proof (see Figure 6). If none of the four edges is in $\partial C(v)^+$, then by now they can only be inner edges of C^+ or in $P(C^+)$ or in $P(C)$. As C and C^{++} cannot neighbour and no edge incident to v is in ∂C^+, this only leaves two possibilities, both of which are balanced at v: either all the four edges are in $P(C^+)$ or all the four edges are inner of C^+ or purple of C (which means the same orientation).

Finally, we consider what happens when none of the edges incident to v are inner edges of $C(v)$ and exactly two of them are in $\partial C(v)^+$. Similarly to the second case of the proof, these two edges are either perpendicular or aligned (see Figure 7). If they are perpendicular, we again have that either both of the two remaining edges are in $P(C^+)$ or both of the two are inner of C^+ or in $P(C)$. If they are aligned, then for each of them it must be the case that exactly one of their guards is in C^{++}, because C (here represented by v) and C^{++} cannot neighbour. Combined with there being no more than two boundary edges, this means that one of the two remaining edges must be in $P(C^+)$ and the other one is inner of C^+ or purple of C. This produces balanced orientation at v both when the (inner) orientations of C^+ and C^{++} are compatible (and so propagated to the boundary) and when they are not (in which case an Eulerian orientation is fixed on the boundary).

References

[1] Á. Backhausz and B. Szegedy (2018). Action convergence of operators and graphs. arXiv 1811.00626

[2] B. Bollobás and O. Riordan (2009). *Percolation*, Chapter 5.

[3] M. Damron, C.M. Newman, V. Sidoravicius (2015). Absence of site percolation at criticality in $\mathbb{Z}^2 \times \{0, 1\}$. *Random Structures and Algorithms.*

[4] B.Q. Feng (2003). Equivalence constants for certain matrix norms. *Linear Algebra and its Applications.*

[5] L. Lovász (2012). *Large Networks and Graph Limits*, Providence, R.I.: American Mathematical Society.

[6] L. Lovász (2019). *Graphs and Geometry*, Providence, R.I.: American Mathematical Society.

[7] L. Russo (1981). On the Critical Percolation Probabilities. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete.*

CEU eTD Collection