

Invariant spanning double rays in amenable groups

Agelos Georgakopoulos* ¹ and Florian Lehner† ²

^{1,2}Mathematics Institute, University of Warwick, CV4 7AL, UK

October 18, 2018

A well-known result of Benjamini, Lyons, Peres, and Schramm states that if G is a finitely generated Cayley graph of a group Γ , then Γ is amenable if and only if G admits a Γ -invariant random spanning tree with at most two ends. We show that this is equivalent to the existence of a Γ -invariant random spanning double ray in a power of G .

1 Introduction

In this paper we study (Cayley) graphs G having a random spanning double ray the distribution of which is invariant under the action of some subgroup of the automorphism group $\text{Aut } G$ of G .

Our first result is

Theorem 1. *The edges of the square grid can be decomposed into a random, \mathbb{Z}^2 -invariant, pair of spanning double rays. Moreover, if G is a connected locally finite Cayley graph of an Abelian group Γ , then G admits a random Γ -invariant spanning double ray.*

A *spanning double ray* is a connected subgraph containing all vertices in which every vertex has exactly two neighbours. The *square grid* is the Cayley graph of \mathbb{Z}^2 with respect to a minimal generating set.

Theorem 1 raises the question of which Cayley graphs admit a random spanning double ray that is invariant under the action of the underlying group, see also Problem 4 below. Benjamini et al. [2] showed that amenability is a necessary condition. The fact that

*Supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 639046).

†Supported by the Austrian Science Fund (FWF), grant J 3850-N32.

the square grid admits a \mathbb{Z}^2 -invariant random spanning double ray was well-known, an example is the Peano UST curve, see e.g. [6].

Alspach [1] conjectured that every Cayley graph of any finite Abelian group admits a decomposition of its edge set into Hamilton cycles. By [5] every Cayley graph of \mathbb{Z}^d admits a decomposition of its edge set into spanning double rays, confirming Alspach's conjecture in spirit for this class of groups. In light of this and the first part of Theorem 1 it is natural to ask whether every Cayley graph of a 1-ended Abelian group admits a random decomposition into double rays which is invariant under the group action.

It is an open problem [4, Problem 3] whether every connected, 1-ended, finitely generated Cayley graph has a spanning double ray, let alone an invariant one. But it is known that a spanning double ray can always be found, if we are allowed to extend the generating set S into $S \cup S^2$ [10]; see also [4, Corollary 6]. The main result of this paper implies that if we extend S further into $S \cup S^2 \cup S^3$, then we can guarantee the existence of an invariant random spanning double ray. It holds in a more general setup of groups acting on graphs:

Theorem 2. *Let G be a countably infinite, connected, amenable graph, and let Γ be a closed subgroup of $\text{Aut } G$ that acts transitively on G . Then there is a Γ -invariant random spanning double ray in G^3 .*

It was known that G^3 contains a (deterministic) spanning double ray for every countable, connected, 1-ended [8] or 2-ended [4] graph G , and this fact motivated our work.

A further source of motivation, and an important tool in the proof of Theorem 2, is a well-known result of Benjamini et al. [2] showing that amenable Cayley graphs admit invariant random spanning trees with 1 or 2 ends. In fact they proved a much more general statement ([2, Theorem 5.3]), that provides probabilistic characterisations of amenability. With Theorem 2 we can add a further probabilistic characterisation, which combined with their statement gives the following:

Corollary 3 (Partly [2, Theorem 5.3]). *Let Γ be a closed unimodular subgroup of $\text{Aut } G$ that acts transitively on a countably infinite, locally finite, connected graph G . Then the following conditions are equivalent:*

1. G is amenable;
2. there is a Γ -invariant random spanning tree of G with at most 2 ends a.s.;
3. there is a Γ -invariant random nonempty connected subgraph ω of G that satisfies $p_c(\omega) = 1$ with positive probability;
4. there is a Γ -invariant random spanning double ray in G^3 ;
5. there is a Γ -invariant random spanning double ray in G^k for some $k \in \mathbb{N}$;
6. Γ is amenable.

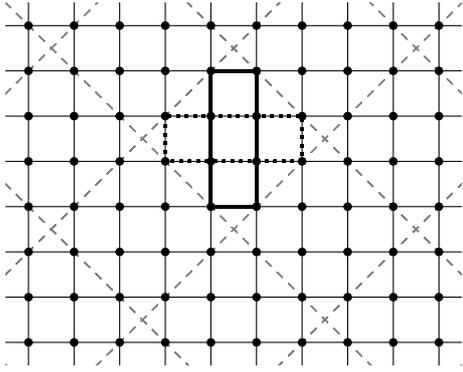


Figure 1: Edge tiling of \mathbb{Z}^2 .

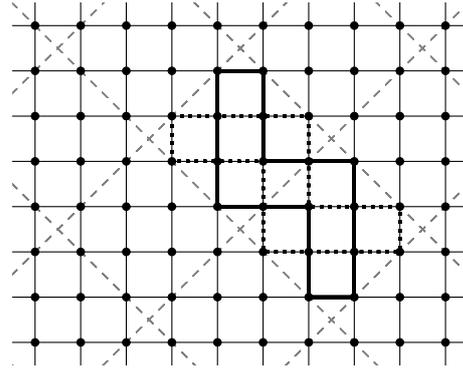


Figure 2: Connecting two tiles.

Our contribution to this is items 4, 5. If the unimodularity condition is dropped, then the following implications still hold: $1 \rightarrow 2 \rightarrow 3 \rightarrow 6$ [2, Theorem 5.3], and $2 \rightarrow 4 \rightarrow 5 \rightarrow 6$: indeed, $2 \rightarrow 4$ is Theorem 2, and $5 \rightarrow 6$ is obtained by applying the implication $2 \rightarrow 6$ on G^k . We do not know whether the implication $5 \rightarrow 2$ holds.

We conjecture that Theorem 2 can be strengthened by replacing G^3 with G^2 , using the techniques of [10, 4] when G is 2-connected. (For this, one would need to produce an invariant ladder-like spanning structure to play the role of the random spanning tree we use in our proof.) The following might be much harder:

Problem 4 ([3, Question 4.5]). Does every amenable, 1-ended Cayley graph $G = \text{Cay}(\Gamma, S)$ admit a Γ -invariant spanning double ray?

We suspect that Corollary 3 still holds when the action is quasi-transitive rather than transitive; in that case, the power G^3 would be optimal, as one could decorate G by attaching an appropriate finite tree to each vertex. If G is assumed to be 2-connected, then G^2 should be the optimal power.

2 Proof of Theorem 1

We first outline a procedure to turn a one-ended spanning tree of \mathbb{Z}^2 into a colouring of the edge set of \mathbb{Z}^2 in which each colour induces a spanning double ray. For this purpose, first consider a decomposition of the edge set of \mathbb{Z}^2 into tiles, see Figure 1 (the bold solid and dotted edges together form one tile). Note that the subgroup Γ of \mathbb{Z}^2 generated by $(\pm 2, \pm 2)$ acts in a natural way on the tiles. Further, 2-colour the edges of each tile as shown in Figure 1: the bold solid edges form one colour class, the bold dotted edges form the other). Colour different tiles so that the colouring is invariant under the action of Γ .

View the tiles as the vertices of the Cayley graph $\text{Cay}(\Gamma, \{(\pm 2, \pm 2)\})$ of Γ , and note that any two adjacent tiles intersect in exactly two vertices. Call these two vertices the *attachment vertices* of those tiles, and call the unique cycle of length 4 containing both of them their *attachment square*. Furthermore, every tile contains four edges of each colour

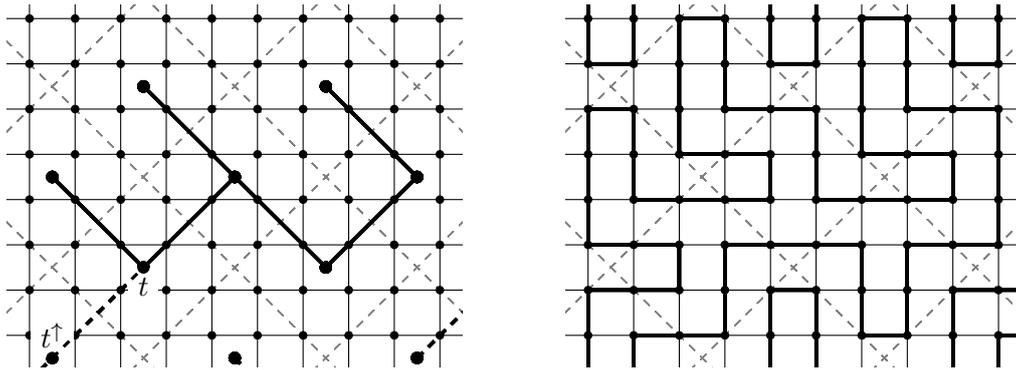


Figure 3: The tree T_t^\downarrow and one of the corresponding paths in \mathbb{Z}^2 .

that are not contained in any attachment square; call those the *internal edges* of the tile. Observe that we can connect the colour classes in two adjacent tiles by swapping the colours within their attachment square, see Figure 2. It is not hard to see that colour swaps at attachment squares can be used to connect the colour classes of multiple squares as well. In fact, the following is true:

Proposition 5. *If T is a one-ended spanning tree of $\text{Cay}(\Gamma, \{(\pm 2, \pm 2)\}) \simeq \mathbb{Z}^2$, then swapping the colours at all attachment squares that correspond to edges of T results in a colouring where each colour class comprises a spanning double ray.*

Proof. Clearly, swapping colours within any family of attachment squares does not change the degree of any vertex in either colour class. In particular, the subgraphs induced by the resulting colour classes will be 2-regular, and it only remains to show that they are connected.

For a tile t , let t^\uparrow be the neighbour of t that lies on the unique ray in T starting at t and let T_t^\downarrow be the component of $T - t^\uparrow$ which contains t .

We claim that all internal edges of the squares in T_t^\downarrow in each colour class are connected by a path contained in the union of the squares in T_t^\downarrow , see Figure 3. This is easily proved by induction. If t is a leaf of T_t^\downarrow it is obvious. If we inductively assume that it holds for the neighbours of t in T_t^\downarrow (whose maximum distance to a leaf in their respective subtree T_s^\downarrow is smaller), then it is straightforward to see that the statement also holds for t using the trivial fact that replacing an edge of a cycle by a path results in a connected graph.

Finally note, that every vertex is adjacent to an internal edge of each colour (although these two internal edges might be internal to different tiles). Since for any two tiles t and t' , there is a tile s such that both t and t' are contained in T_s^\downarrow , this implies that each colour class is connected and thus a spanning double ray. \square

It is well known that \mathbb{Z}^2 admits an invariant measure on one ended spanning trees, see for example [7]. If we choose the tree T in the above construction according to such a measure, then it is easy to see that the colour of an edge of \mathbb{Z}^2 only depends on whether

or not a specific edge is contained in T . Hence cylinder sets are measurable and we end up with a Γ -invariant decomposition of \mathbb{Z}^2 into spanning double rays.

To finish the proof of Theorem 1, note that Γ has finite index in \mathbb{Z}^2 . Thus any Γ -invariant measure μ on subgraphs of \mathbb{Z}^2 can be made into a \mathbb{Z}^2 -invariant measure

$$\nu = \frac{1}{[\mathbb{Z}^2 : \Gamma]} \sum_i \mu_i,$$

where μ_i are translations of μ with respect to a system of representatives of the cosets of Γ . This shows that \mathbb{Z}^2 has an invariant decomposition into double rays, and in particular, \mathbb{Z}^2 contains an invariant spanning copy of \mathbb{Z} .

Next, we show that the Cayley graph of \mathbb{Z}^d with respect to the standard generators admits an invariant spanning double ray. First consider the case $d = 3$. Let Z_i be a copy of the standard Cayley graph of \mathbb{Z} for every i . To sample a random spanning double ray R in $X := Z_1 \times Z_2 \times Z_3$, we first sample a pair of independent spanning double rays R_{12}, R_3 in $Z_1 \times Z_2$ and $Z_0 \times Z_3$, respectively, each according to the law of the \mathbb{Z}^2 -invariant spanning double ray we obtained above.¹

Let f be one of the two isomorphisms from Z_0 to R_{12} mapping $0 \in Z_0$ to $0 \in Z_1 \times Z_2$, chosen by a fair coin flip. Extend f canonically to a map $\phi : Z_0 \times Z_3 \rightarrow R_{12} \times Z_3$ by ‘multiplying’ with Z_3 . Then $R := \phi(R_3)$ is a random spanning double ray of X .

Clearly, the law of R is invariant under the canonical action of \mathbb{Z} on X in the Z_3 coordinate. We claim that it is also invariant under the canonical action of \mathbb{Z}^2 in the other two coordinates.

For this, let r_{12} be the (random) labelling of the pairs of vertices (x, y) of $Z_1 \times Z_2$ where $r_{12}(x, y)$ is the distance between x and y along R_{12} . Let $B = I_1 \times I_2 \times I_3$, where I_i is a finite subpath of Z_i , be a ‘box’ in X , to be thought of as a cylinder set of our σ -algebra. Notice that conditioning on the values of r_{12} inside $I_1 \times I_2$ uniquely determines the distribution of $R \cap B$ by the construction of R . As R_{12} is invariant under the canonical action of \mathbb{Z}^2 on $Z_1 \times Z_2$, so is r_{12} . Thus the distribution of $R \cap B$ coincides with that of $R \cap gB$ for every $g \in \mathbb{Z}^2$, which completes the proof that R is invariant under the canonical action of \mathbb{Z}^3 on X .

This construction generalises to $\mathbb{Z}^d, d > 3$ by induction as follows. Assume we have proved that the standard Cayley graph of \mathbb{Z}^{d-1} admits both an invariant spanning double ray as well as an invariant spanning copy H of the standard Cayley graph of \mathbb{Z}^{d-2} . We repeat the above construction with $X := Z_1 \times Z_2 \cdots \times Z_d$, with R_{12} as before, but with R_3 being a sample of H in $Z_0 \times Z_3 \times Z_4 \times \cdots \times Z_d$. Repeating the above ideas yields now an invariant spanning double ray in X . We now also need an invariant spanning copy of \mathbb{Z}^{d-1} in X , to prove our inductive hypothesis, but this is easier: we can just take $R_{12} \times H$.

Finally, let Γ be an arbitrary finitely generated Abelian group, let S be a finite generating set, and let G be the Cayley graph of Γ with generators S . Let $S' \subseteq S$ be a maximal set of linearly independent generators with infinite order. Then the subgroup

¹Alternatively, we could just use the Peano UST curve from [6].

Γ' of Γ generated by S' has finite index in Γ , and the Cayley graph G' of Γ' with respect to S' is the standard Cayley graph of \mathbb{Z}^d for some d .

From now on, view G' as a subgraph of G with the obvious embedding. We will use a spanning double ray in G' to define a spanning double ray in G as follows. Let $M_1 \cup M_2$ be the decomposition of the edge set of the spanning double ray into two perfect matchings. In order to make the choice canonical, decide by a fair coin flip which of the two sets is M_1 . It is well known that every Cayley graph of a finite Abelian group admits a Hamilton path, see [9]. From such a path in the Cayley graph of Γ/Γ' (with generators corresponding to S), it is easy to construct a path P in G which starts at 0 and visits every coset of Γ' precisely once. Pick such a path P arbitrarily, let p be the other endpoint of P and let f_p be the automorphism of G mapping x to $x + p$.

We claim that the subgraph of G consisting of $M_1, f_p(M_2)$, and all translates of P by elements of Γ' is the desired spanning double ray. Since P visits every coset exactly once, every vertex of G lies on a unique translate of P . Vertices in G' and $f_p(G')$ are incident to an additional matching edge in M_1 and $f_p(M_2)$ respectively, whence R is 2-regular. To see that it is connected, observe that contracting all edges in translates of P gives the spanning double ray of G' that we started with.

To finish the proof, note that if the spanning double ray in G' was taken from an invariant distribution on spanning double rays, then the result of the above construction is invariant under Γ' . Since $[\Gamma : \Gamma']$ is finite, we can then use the same argument as in the case of \mathbb{Z}^2 to turn this into a Γ -invariant spanning double ray.

3 Proof of Theorem 2

Let (\mathcal{T}_G, μ) be the probability space arising from the implication $1 \rightarrow 2$ of Corollary 3 (the condition of unimodularity is not needed for that implication, see the remark after the statement), so that the elements of \mathcal{T}_G are the spanning trees of G , and consider the product probability space

$$\mathcal{T} := (\Omega, \Sigma, \mathbb{P}) = \mathcal{T}_G \times \prod_{v \in V} \mathcal{U}_{<_v},$$

where $\mathcal{U}_{<_v}$ denotes the uniform distribution on total orders on the neighbours of v in G .

We will construct a map $\phi: \Omega \rightarrow 2^{E(G^3)}$, mapping each $\omega = (T, (<_v)_{v \in V}) \in \Omega$ to a spanning double ray of G^3 . The main idea behind the definition of this map ϕ stems from [4], and is illustrated in Figure 4. To explain it, let us first pretend that T is a spanning tree of a finite graph G , rooted at a vertex v , and we want to use it to obtain a Hamilton cycle of G^3 . We do this by induction on the size of T as follows. Suppose that for each subtree T_i of T obtained by removing v , we have a path P_i in G^3 from the root v_i of T_i to one of its neighbours v_i^\dagger in T_i visiting each vertex of T_i exactly once (unless v_i is the unique vertex of T_i , in which case P_i is the trivial path). Then we can join the P_i using an edge from G^3 as shown in Figure 4, to obtain a path from v to a neighbour v^\dagger , proving the inductive hypothesis. We can then produce a Hamilton cycle of G^3 by adding the $v-v^\dagger$ edge to that path.

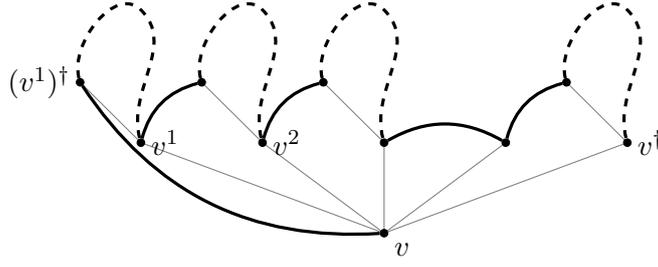


Figure 4: Edges of types (i) and (ii) together with $v^i-(v^i)^\dagger$ -paths form a $v-v^\dagger$ -path.

It is possible to produce such a Hamilton cycle by making local decisions at each vertex of T simultaneously, rather than working recursively: given a fixed ordering of the neighbours of v in Figure 4, we can decide which additional edges to join the P_i with without seeing the P_i . This is the key observation that will allow us to use the same idea in our situation, where T is infinite, with one or two ends. These ‘local’ decisions are encoded in the aforementioned map ϕ . This is where the orders $(<_v)_{v \in V}$ will be used.

We define ϕ as follows. Let T be a 1- or 2-ended spanning tree of G . For a vertex v , denote by N_v^∞ the (one or two) neighbours of v in T that lie in an infinite component of $T - v$, and let N_v^\times be the remaining neighbours of v in T . Clearly $<_v$ gives rise to a total order on N_v^\times . Denote by v^i the i -th largest element in this order. If $N_v^\times \neq \emptyset$, then denote by v^\dagger the maximal element, otherwise let $v^\dagger = v$.

Now let $\phi(\omega) \subseteq E(G^3)$ consist of the following edges:

- (i) for $v \in V(G)$ with $N_v^\times \neq \emptyset$ the edge connecting v and $(v^1)^\dagger$
- (ii) for $v \in V(G)$ with $N_v^\times \neq \emptyset$ and $v^i <_v v^\dagger$ the edge connecting v^i to $(v^{i+1})^\dagger$, and
- (iii) for $uv \in E(G)$ where $u \in N_v^\infty$ and $v \in N_u^\infty$
 (This case only materialises when T contains a double ray, and that double ray contains uv .)
 - if $u = \max N_v^\infty$ and $v = \max N_u^\infty$ the edge uv ,
 - if $u = \max N_v^\infty$ and $v \neq \max N_u^\infty$ the edge $u^\dagger v$,
 - if $u \neq \max N_v^\infty$ and $v \neq \max N_u^\infty$ the edge $u^\dagger v^\dagger$,

As we will prove below (Lemma 7), items (i) and (ii) implement the intuition of Figure 4, while item (iii) ensures that if R is a double ray in T^3 (which is unique if it exists), then for each edge uv in R , we choose exactly one edge joining the subtree of $T - uv$ rooted at u to the subtree rooted at v , and we choose it in such a way to ensure all vertices have degree 2. If G is 1-ended, then we may assume T to be 1-ended too by a result of Timar [11], in which case item (iii) can be dropped and the proof of Lemma 7 accordingly simplified.

We would like to interpret the pushforward of \mathcal{T} via ϕ as a measure on subgraphs of G^3 . In order to do so, it is desirable that cylinder sets (i.e. events telling us what happens

on finite subgraphs of G^3) are measurable. The following proposition says that this is the case.

Lemma 6. *For every cylinder set C of $2^{E(G^3)}$, the set $\phi^{-1}(C)$ is measurable in \mathcal{T} .*

Proof. It is sufficient to show that for every edge $e \in E(G^3)$ the event $[e \in \phi(\omega)]$ is measurable.

First note that the event

$$[u \in N_v^\infty] = [uv \in E(T) \text{ and } \exists V_u \subseteq V(G): u \in V_u, |V_u| < \infty, \text{ no edge leaving } V_u \text{ is in } T - uv]$$

is \mathcal{T} -measurable for any pair u, v . Since there are only countably many finite sets containing u , this is a countable union of cylinder sets, and thus \mathcal{T} -measurable. Consequently the events $[u \in N_v^\infty]$, $[u = v^\dagger]$, $[u = v^\ddagger]$, and $[u = \max N_v^\infty]$ are measurable because they can be obtained by finite intersections and unions of events of the form $[xy \in E(T)]$, $[x \in N_y^\infty]$, $[x \notin N_y^\infty]$ and $[x <_z y]$. Finally, the event $[e \in \phi(\omega)]$ can be obtained by finite intersections and unions of the above events, hence it is measurable in \mathcal{T} . \square

Theorem 2 is now an easy consequence of the following lemma.

Lemma 7. *Let H be the random subgraph of G^3 defined by the pushforward of $(\Omega, \Sigma, \mathbb{P})$ via ϕ . Then H almost surely is a spanning double ray.*

Proof. We prove the statement of the proposition for a fixed tree T . This is sufficient since $\mathbb{P}(A) = \int \mathbb{P}(A | T) d\mu(T)$ where μ is the measure on the space \mathcal{T}_G .

For a vertex v denote by T_v^∞ the subtree of T induced by v and all finite components of $T - v$. Let $h(v)$ the maximum distance from v to a vertex in T_v^∞ . We prove by induction on h that the subgraph of H induced by the vertices of T_v^∞ is a spanning path from v to v^\dagger . If $h(v) = 0$, this is trivially true. For the induction step, note that there can't be an edge of type (iii) inside T_v^∞ . Indeed, if xy is such an edge, then there is an edge e separating x and y in T such that neither component of $G - e$ is finite which implies that $x \notin T_y^\infty$. Now the induction hypothesis together with the definition of edges of type (i) and (ii) prove the claim.

In particular, this claim shows that all vertices of T_v^∞ lie in the same component of H and that apart from v and v^\dagger all of them have degree 2.

If T has one end, then for any vertex x there is a vertex v such that $x \in T_v^\infty \setminus \{v, v^\dagger\}$, showing that x has degree 2. Furthermore, for any two vertices x, y there is a vertex v such that T_v^∞ contains both x and y , showing that H is connected.

If T has two ends, then the edges for which both components of $T - e$ are infinite form a double ray. Note that these are precisely the edges uv for which $u \in N_v^\infty$ and $v \in N_u^\infty$. Let x be a vertex on this double ray, and let y and z be its neighbours. Then the edges of type (iii) corresponding to xy and xz connect T_x^∞ to T_y^∞ and T_z^∞ respectively. In particular, this shows that H is connected. Finally note that one of the edges of type (iii) attaches to x and the other one attaches to x^\dagger (depending on whether $y <_x z$ or $z <_x y$), thus making sure that both x and x^\dagger have degree 2 in H . \square

Acknowledgements

We thank Gabor Pete for a discussion that triggered this paper.

References

- [1] B. Alspach. Unsolved problem 4.5. *Annals of Discrete Mathematics*, 27:464, 1985.
- [2] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999.
- [3] D. Beringer, G. Pete, and A. Timar. On percolation critical probabilities and unimodular random graphs. *Electron. J. Probab.*, 22(106):1–26, 2017.
- [4] A. Georgakopoulos. Infinite Hamilton cycles in squares of locally finite graphs. *Advances in Mathematics*, 220:670–705, 2009.
- [5] F. Lehner J. Erde and M. Pitz. Hamilton decompositions of one-ended cayley graphs. *arXiv:1709.09463*.
- [6] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*. Cambridge University Press, New York, 2016. Available at <http://pages.iu.edu/~rdlyons/>.
- [7] Robin Pemantle. Choosing a spanning tree for the integer lattice uniformly. *Ann. Probab.*, 19:1559 – 1574, 1991.
- [8] M. Sekanina. On an ordering of the set of vertices of a connected graph. *Publ. Fac. Sc. Univ. Brno*, 412:137–142, 1960.
- [9] Dragan Marušič. Hamiltonian circuits in cayley graphs. *Disc. Math.*, 46(1):49–54, 1983.
- [10] C. Thomassen. Hamiltonian paths in squares of infinite locally finite blocks. *Annals of Discrete Mathematics*, 3:269–277, 1978.
- [11] A. Timar. One-ended spanning trees in amenable unimodular graphs. *arXiv:1805.10690*.