

Locally infinite graphs and symmetries



SÉBASTIEN MARTINEAU

Abstract

When one studies geometric properties of graphs, local finiteness is a common implicit assumption, while that of transitivity is a more explicit one. By compactness arguments, local finiteness guarantees several regularity properties (see Section 1.2). It is generally easy to find counterexamples to such regularity results when the assumption of local finiteness is dropped. The present work focuses on the following problem: determining whether these regularity properties still hold when local finiteness is replaced by an assumption of transitivity.

After recalling the locally finite situation, we show that there are Cayley graphs of $\bigoplus_{n \geq 2} \mathbb{Z}/n\mathbb{Z}$ and \mathbb{Z} (with infinite generating systems) that have infinite diameter but do not contain any infinite geodesic ray. We also introduce a notion of generalised diameter. The generalised diameter of a graph is either an ordinal or ∞ and captures the extension properties of geodesic paths. It is a finite ordinal if and only if the usual diameter is finite, and in that case the two notions agree. Besides, the generalised diameter is ∞ if and only if the considered graph contains an infinite geodesic ray. We show that there exist Cayley graphs of abelian groups of arbitrary generalised diameter.

Finally, we build Cayley graphs of abelian groups that have isomorphic balls of radius n for every n but are not globally isomorphic. This enables us to construct a non-transitive graph in which any two balls of the same radius are isomorphic.

MSC 2010. Primary: 05C63, 05C12; Secondary: 05C10.

1 Introduction

Graphs are fundamental objects in discrete mathematics. This paper will focus on their geometric aspect. From this point of view, it is common to work under the assumption of local finiteness: a graph is **locally finite** if each of its vertices has finitely many neighbours. Another important notion is that of transitive graphs: a graph is **(vertex-)transitive** (or **homogeneous**) if its automorphism group acts transitively on its vertex-set. See e.g. [5].

It is well-known that local finiteness implies some form of compactness, which guarantees that the geometric study of locally finite graphs is in some sense well-behaved. This paper aims at showing that the situation gets wilder when the local finiteness assumption is dropped, *even if* one makes the stringent assumption that the graphs under study are transitive (or weakly transitive, see p. 44).

In Section 1.1, for completeness and in order to fix conventions, we recall the vocabulary of graph theory that we will need. In Section 1.2, we recall the locally finite situation. In Section 1.3, after a brief presentation of our motivations, we state our results and open questions. These results are established in the subsequent sections.

1.1 Vocabulary of graph theory

We denote by \mathbb{N} the set of all non-negative integers and by \mathbb{N}^* that of positive integers.

Given a set X , denote by $\binom{X}{2}$ the set whose elements are the subsets of X that contain exactly 2 elements. A “**graph**” is an ordered pair $\mathcal{G} = (V, E)$, where $E \subset \binom{V}{2}$. The elements of V are called the **vertices** of \mathcal{G} , and the elements of E are called the **edges** of \mathcal{G} . Two vertices u and v are said to be **adjacent** (or **neighbours**) if $\{u, v\}$ is an edge. The elements of an edge are called its **endpoints**. A **path** is a map $\kappa : I \rightarrow V$ such that

- I is a non-empty interval of \mathbb{Z} ,
- $\forall n \in I, n + 1 \in I \implies \{\kappa(n), \kappa(n + 1)\} \in E$.

The **length** of $\kappa : I \rightarrow V$ is “ $\sup I - \inf I$ ” $\in \mathbb{N} \cup \{\infty\}$. If this length is finite, then the path is said to be **finite**, and it is said to **connect** $\kappa(\min I)$ and $\kappa(\max I)$. A “graph” is **connected** if any two vertices are connected by some finite path. In this paper, a **graph** is a connected “graph” with a non-empty vertex set. A graph is **locally finite** if every vertex has finitely many neighbours; this number need not be bounded. A **rooted graph** is the data of a graph together with that of a vertex of this graph, which is referred to as the **root**.

Given a graph $\mathcal{G} = (V, E)$, one defines a distance on V — the **graph distance** — by setting $d(u, v)$ to be the minimal length of a finite path connecting u and v . Given $r \in \mathbb{R}_+$ and $u \in V$, the **ball** $B := B_{\mathcal{G}}(u, r)$ of radius r and with centre u is the set of the elements v satisfying $d(u, v) \leq r$. It can be seen as a graph by setting its edge-set to be $E \cap \binom{B}{2}$. It can be seen as a rooted graph by rooting it at u . The **diameter** of $\mathcal{G} = (V, E)$ is $\sup\{d(u, v) : u, v \in V\} \in \mathbb{N} \cup \{\infty\}$.

A **(graph) isomorphism** from a graph $\mathcal{G} = (V, E)$ to a graph $\mathcal{G}' = (V', E')$ is a bijection $\varphi : V \rightarrow V'$ such

that

$$\forall u, v \in V, \{u, v\} \in E \iff \{\varphi(u), \varphi(v)\} \in E'.$$

This notion coincides with that of isometry: an isomorphism from \mathcal{G} to \mathcal{G}' is the same as a distance-preserving bijection from V to V' . Two graphs are **isomorphic** if there is an isomorphism from the first one to the second one. This defines an equivalence relation on the class of graphs. A **(graph) automorphism** of a graph is an isomorphism from itself to itself. These notions extend to rooted graphs by adding the condition that φ maps the root of \mathcal{G} to that of \mathcal{G}' . Graph automorphisms of a graph form a group under composition, which is called the **automorphism group** of this graph.

A graph \mathcal{G} is said to be **(vertex-)transitive** (or **homogeneous**) if the natural action of its automorphism group on its vertex-set is transitive, i.e. if for any vertices u and v , there is a graph automorphism φ of \mathcal{G} such that $\varphi(u) = v$. One says that \mathcal{G} is **quasi-transitive** if its automorphism group acts on its vertices with finitely many orbits. Given a group G and a generating subset S of G , the **Cayley graph** $\text{Cay}(G, S)$ associated with (G, S) is the “graph” defined by taking G to be the vertex-set and by declaring two *distinct* vertices g and h to be adjacent if and only if $g^{-1}h \in S \cup S^{-1}$. This “graph” is actually a transitive *graph*. The action of G on itself by left-multiplication is free, transitive, and by graph automorphisms of $\text{Cay}(G, S)$. Notice that *we do not* assume S to be finite.

A path $\kappa : I \rightarrow V$ is **geodesic** if it satisfies

$$\forall m, n \in I, d(\kappa(m), \kappa(n)) = |m - n|.$$

A path is **infinite** if its length is infinite. A path is **bi-infinite** if it satisfies $I = \mathbb{Z}$. A **ray** (or **infinite ray**) is a path $\kappa : \mathbb{N} \rightarrow V$. A path κ **starts at** some vertex u if $\kappa(0)$ is well-defined and equal to u .

Given a graph $\mathcal{G} = (V, E)$, we see the set of its paths as a subset of $(V \cup \{\emptyset\})^{\mathbb{Z}}$: just extend every path $\kappa : I \rightarrow V$ to a map defined on \mathbb{Z} by setting its value to be \emptyset for every $n \notin I$. Endowing $V \cup \{\emptyset\}$ with the discrete topology, the product topology on $(V \cup \{\emptyset\})^{\mathbb{Z}}$ induces a topology on the space of the paths of \mathcal{G} . This will be the only topology we will consider on this space.

1.2 Regularity for locally finite graphs

Even though some of the following results are seldom stated in this way, no result of Section 1.2 can be considered to be new. However, I believe it is useful to present these results and their proofs in order to provide context for our work in the non locally finite setup.

1.2.1 Infinite geodesic paths

The power of local finiteness comes from the following observation:

Fact 1.1 *In a locally finite rooted graph, the set consisting of the paths starting at the root is compact.*

Proof. The set \mathfrak{P} of the paths in $\mathcal{G} = (V, E)$ that start at o is a subset of $X := \prod_{k \in \mathbb{Z}} B(o, |k|) \cup \{\emptyset\}$, endowed with the product of discrete topologies. Since \mathcal{G} is locally finite, every $B_{\mathcal{G}}(o, |k|)$ is finite hence compact. By Tychonoff’s Theorem¹, X is compact. For an element x of X , being a path is the conjunction over (k, l, m) in \mathbb{Z}^3 such that $k < l < m$ of the following closed conditions:

- $x(l) = \emptyset \implies (x(k) = \emptyset \text{ or } x(m) = \emptyset)$;
- $\{x(k), x(k+1)\} \cap \{\emptyset\} = \emptyset$
 $\implies \{x(k), x(k+1)\} \in E$.

As a result, the set \mathfrak{P} is closed in the compact space X , hence compact. \square

The following result, which is useful in the locally finite framework, does not require local finiteness. We leave its proof to the reader.

Fact 1.2 *In any rooted graph, the geodesic paths starting at the root form a closed subset of the space of paths.*

Corollary 1.3 *Every locally finite graph of infinite diameter contains an infinite geodesic path.*

Proof. Let \mathcal{G} be a locally finite graph of infinite diameter. Fix o a vertex of \mathcal{G} . By assumption, we can find a sequence (γ_n) of finite geodesic paths that start at o and such that the length of γ_n tends to infinity when n goes to infinity. As \mathcal{G} is locally finite, by Fact 1.1, this sequence (γ_n) has at least one accumulation point in the space of paths: take γ_∞ to be such an accumulation point. By Fact 1.2, γ_∞ is actually geodesic. As γ_∞ is infinite, Corollary 1.3 is established. \square

Under an additional assumption of transitivity, we will prove better: see Corollary 1.6. Once again, we will use some fact that does not require local finiteness (and actually not even transitivity), namely Fact 1.4.

Fact 1.4 *In any graph, the geodesic paths form a closed subset of the space of paths, and this subset is invariant by reparametrisation² and by graph automorphisms.*

By using Fact 1.1, we can prove the following lemma.

Lemma 1.5 *Let \mathfrak{P} denote a set of paths of some transitive graph \mathcal{G} . Assume that the following conditions hold:*

1. \mathcal{G} is locally finite;
2. the automorphism group of \mathcal{G} leaves \mathfrak{P} globally invariant;
3. if $k_0 \in \mathbb{Z}$ and $\kappa \in \mathfrak{P}$, then $k \mapsto \kappa(k + k_0)$ belongs to \mathfrak{P} ;
4. \mathfrak{P} is closed;
5. \mathfrak{P} contains paths of arbitrarily large length.

¹ No axiom of choice is needed here, as the product is countable and the factors are metrisable. The reader may prefer to proceed by diagonal extraction instead of resorting to Tychonoff’s Theorem.

² i.e. by $\gamma \mapsto (\gamma(ak + b))_k$ for any $a \in \{-1, 1\}$ and $b \in \mathbb{Z}$

Then \mathfrak{P} contains a bi-infinite path.

Proof. Let (κ_n) denote a sequence of elements of \mathfrak{P} such that the length of κ_n is at least n . Such a sequence exists by Condition 5. Let o denote a vertex of \mathcal{G} . As \mathcal{G} is transitive, by Conditions 2, 3 and 5, one may assume that for every n , the domain of κ_n contains $-\lfloor n/2 \rfloor$ and $\lfloor n/2 \rfloor$ and $\kappa_n(0) = o$. By Condition 4, Condition 1 and Fact 1.1, the elements of \mathfrak{P} that start at o form a compact set. The sequence (κ_n) thus admits an accumulation point in \mathfrak{P} , which must be bi-infinite. \square

From Fact 1.4 and Lemma 1.5, it results that:

Corollary 1.6 *Every transitive locally finite graph of infinite diameter contains a bi-infinite geodesic path.*

Of course, this corollary is false without the transitivity assumption: consider a one-sided infinite ray.

1.2.2 Local topology is Hausdorff

Local finiteness also plays an important role in the investigation of the so-called local topology. Let us denote by $\mathfrak{G}_{\text{lf}}^\dagger$ the set consisting in the isomorphism classes of locally finite rooted graphs. Notice that $\mathfrak{G}_{\text{lf}}^\dagger$ can indeed be realised as a set, as every locally finite graph is isomorphic to a graph the vertex-set of which is a subset of \mathbb{N} . By abuse of language, we may, whenever this is not harmful, identify a locally finite rooted graph with its isomorphism class.

Given two rooted graphs (\mathcal{G}, o) and (\mathcal{G}', o') , we set $d((\mathcal{G}, o), (\mathcal{G}', o')) := 2^{-K}$, where

$$K := \sup\{n \in \mathbb{N} : B_{\mathcal{G}}(o, n) \simeq B_{\mathcal{G}'}(o', n)\} \in \mathbb{N} \cup \{\infty\}$$

and the isomorphism between the balls is taken relative to their structure of rooted graphs. The following proposition is well-known.

Proposition 1.7 *The map d defines a distance on $\mathfrak{G}_{\text{lf}}^\dagger$.*

Proof. The most interesting point to check is that if $d((\mathcal{G}, o), (\mathcal{G}', o')) = 0$ then (\mathcal{G}, o) and (\mathcal{G}', o') are isomorphic. We will only check this.

Let $(\mathcal{G}, o) = ((V, E), o)$ and $(\mathcal{G}', o') = ((V', E'), o')$ be two rooted graphs such that for every n , $B_{\mathcal{G}}(o, n)$ and $B_{\mathcal{G}'}(o', n)$ are isomorphic as rooted graphs. For every n , one can thus fix some isomorphism φ_n from $B_{\mathcal{G}}(o, n)$ to $B_{\mathcal{G}'}(o', n)$. Each φ_n can be seen as a map from V to $V' \cup \{\emptyset\}$, and more precisely as an element of the space $\prod_{v \in V} (B_{\mathcal{G}'}(o', d(o, v)) \cup \{\emptyset\})$ — which is compact by local finiteness of \mathcal{G}' . Any accumulation point of the sequence (φ_n) is an isomorphism from \mathcal{G} to \mathcal{G}' . \square

This distance defines the so-called **local topology** — see [3, 7, 12]. It is very useful in statistical mechanics: see e.g. [2, 6, 13, 14, 16, 20]. It is also important in the topic of soficity, which is a major notion at the interplay between group theory and ergodic theory [22]. The reader is referred to [8, 9] for some interaction with group theory (via the notion of marked groups

[10, 11, 15]), and to Section 10 in [1] for some graph-theoretic/probabilistic counterpart of soficity.

Say that a graph \mathcal{G} is **weakly transitive** if for any vertices o and o' and every positive integer n , the graphs $B_{\mathcal{G}}(o, n)$ and $B_{\mathcal{G}}(o', n)$ are isomorphic as rooted graphs. This definition applies to graphs that may not be locally finite.

Corollary 1.8 *A locally finite graph is transitive if and only if it is weakly transitive.*

Proof. Every transitive graph is weakly transitive. Let thus \mathcal{G} denote a locally finite weakly transitive graph. Let o and o' denote two vertices of \mathcal{G} . By weak transitivity of \mathcal{G} , we have $d((\mathcal{G}, o), (\mathcal{G}, o')) = 0$. By Proposition 1.7, (\mathcal{G}, o) and (\mathcal{G}, o') are isomorphic, which means that there is an automorphism of \mathcal{G} that maps o to o' . \square

Given a weakly transitive graph, for every k , one can speak — up to isomorphism as a rooted graph — of its ball of radius k . Say that two weakly transitive graphs are **weakly isomorphic** if for every positive integer k , their balls of radius k are isomorphic the one to the other. This definition applies to graphs that may not be locally finite.

The following statement results from Proposition 1.7.

Corollary 1.9 *For locally finite (weakly) transitive graphs, weak isomorphism is equivalent to isomorphism.*

1.3 Beyond the locally finite case

In this section, we will use the word “fact” for results whose essence is classical. Except for facts, relevant proofs will be deferred to subsequent sections.

Let us consider the corresponding situations in the non locally finite world. Before getting to precise statements, we want to explain why such a study is relevant. Apart from the exploration of the limits of regularity properties for their own sake, and our interest for the constructions and questions it leads to, we would like to mention the following motivation: the study of *geometric* properties of some *quasi-transitive* non locally finite graphs plays a significant role across various branches of mathematics. For instance, in the study of surfaces, the *curve graph* of an orientable surface S with genus g and m punctures satisfying $3g + m \geq 5$, and more specifically its hyperbolicity, plays a major role in the study of the mapping class group of S and its Teichmüller space — see [18, 19]. The hyperbolicity of this graph led to the resolution of Thurston’s ending lamination conjecture – see [17]. Likewise, in group theory, the hyperbolicity of the graph of free factors and that of free splittings of the free group with $n \geq 3$ generators \mathbf{F}_n is important in the study of the exterior automorphisms of \mathbf{F}_n .

1.3.1 The study of geodesic paths

Fact 1.1 is *always* false in the non locally finite case.

Fact 1.10 *In a non locally finite rooted graph, the set of the paths that start at the root is never compact.*

Corollary 1.3 is false in the non locally finite case. Indeed, let \mathcal{T} be the graph defined as follows, and depicted on Figure 1. The vertex-set is

$$V := \{(k, n) \in \mathbb{N}^2 : 1 \leq k \leq n\} \cup \{0\}.$$

The edge-set is defined to be

$$\begin{aligned} & \{ \{(k, n), (l, n)\} : 1 \leq k \leq n, 1 \leq l \leq n, |k - l| = 1 \} \\ & \cup \{ \{0, (1, n)\} : n \geq 1 \}. \end{aligned}$$




Fig. 1: Two ways of picturing the graph \mathcal{T} .

Fact 1.11 *The graph \mathcal{T} has infinite diameter but does not admit any geodesic ray.*

Corollary 1.3 is still false in the *transitive* non locally finite case: actually one can even find counterexamples that are Cayley graphs of abelian groups.

Propositions 2.1 and 2.2 .

The groups \mathbb{Z} and $\bigoplus_{n>2} \mathbb{Z}/n\mathbb{Z}$ admit Cayley graphs of infinite diameter that do not contain any geodesic ray.

We introduce in page 47 the generalised diameter of a graph. The generalised diameter of a graph is either an ordinal, or the symbol ∞ . It is a finite ordinal if and only if the usual diameter is finite, and in that case the two notions agree. Moreover, the generalised diameter is ∞ if and only if the considered graph contains a geodesic ray. When this generalised diameter is an infinite ordinal, it captures finely the extension properties of geodesic paths.

Theorem 2.3 *There are Cayley graphs of abelian groups of every generalised diameter.*

As shown by Propositions 2.1 and 2.2, Corollary 1.6 is false without the assumption of local finiteness. However, one may ask the following question:

Question 1.14 *Does every transitive graph that contains a geodesic ray contain a bi-infinite geodesic path?*

Even though I do not know the answer to this question, we can understand the situation for Cayley graphs of abelian groups (Proposition 2.4 and Proposition 2.5): if the assumption of local finiteness is replaced by that of being a Cayley graph of an abelian group, then Lemma 1.5 does not hold, but its Corollary 1.6 remains true.

Proposition 2.4 *Let \mathcal{G} be a Cayley graph of an abelian group. Assume that \mathcal{G} contains a geodesic ray. Then, \mathcal{G} contains a bi-infinite geodesic path.*

Proposition 2.5 *Lemma 1.5 does not hold if Condition 1 is removed, even if \mathcal{G} is assumed to be a Cayley graph of an abelian group and \mathfrak{B} is assumed to be invariant under $\kappa \mapsto (k \mapsto \kappa(-k))$.*

1.3.2 Weak isomorphism and transitivity

Proposition 1.7 fails to hold for two reasons in the non locally finite case. First, the space of isomorphism classes of rooted graphs cannot be realised as a set. Second, two rooted graphs (\mathcal{G}, o) and (\mathcal{G}', o') may satisfy $B_{\mathcal{G}}(o, n) \simeq B_{\mathcal{G}'}(o', n)$ in the rooted sense for every $n \in \mathbb{N}$, but still not be isomorphic. Consider the graph \mathcal{T} introduced before Fact 1.11, rooted at 0, and the same construction where n is taken in $\mathbb{N} \cup \{\infty\}$ and k in \mathbb{N} . “See” Figure 2.

One can also find such counterexamples that are Cayley graphs of abelian groups, so that Corollary 1.9 does not hold beyond the locally finite case.

Proposition 3.1 *There are two Cayley graphs of abelian groups that are weakly isomorphic but not isomorphic.*

Corollary 1.8 does not hold without local finiteness.

Proposition 3.2 *There exists a graph that is weakly transitive and weakly isomorphic to a transitive graph, but that is not transitive.*

Here is a natural question:

Question 1.19 *Is there a weakly transitive graph which is not weakly isomorphic to any transitive graph?*

I currently do not know how to solve this question, but I expect such a graph to exist.

Remark 1.20 Such a graph would have the flavour of a Penrose tiling (a tiling of the plane made of tiles that cannot tile the plane periodically). More precisely, instead of tiles, we would have a system of balls, i.e. the data for each radius $k \in \mathbb{N}$ of a rooted graph which is thought of as the prescribed ball of radius k up to isomorphism. As is the case for the tiles of a Penrose tiling, this system of balls could be put together but not in a periodic way. In this new setting, “putting




Fig. 2: Two ways of picturing the same rooted graph.

together” does not mean “tiling the plane” but “there is a graph such that, at every vertex, the balls of every radius are isomorphic to the prescribed ones”. Likewise, “putting together in a periodic way” now means that no way of putting together the balls of the system yields a transitive graph.

Question 1.19 naturally leads to the following problem, which I cannot solve either.

Problem 1.21 *Given $(m, n) \in \mathbb{N}^2$ such that $m \leq n$, determine whether there is a rooted graph (\mathcal{B}, o) such that the following two properties hold:*

- *there is a graph $\mathcal{G} = (V, E)$ such that, for every $v \in V$, the rooted graphs $(B_{\mathcal{G}}(v, n), v)$ and (\mathcal{B}, o) are isomorphic*
- *and there is no transitive graph $\mathcal{G} = (V, E)$ such that, for every $v \in V$, the rooted graphs $(B_{\mathcal{G}}(v, m), v)$ and $(B_{\mathcal{B}}(o, m), o)$ are isomorphic.*

I cannot answer this problem either in the finite or locally finite settings.

2 The study of geodesic paths

2.1 Cayley graphs of infinite diameter with no geodesic ray

Here is an easy example of a transitive graph of infinite diameter that does not contain any geodesic ray.

Proposition 2.1 *Let G denote the group $\bigoplus_{n \geq 2} \mathbb{Z}/n\mathbb{Z}$ and S be the generating subset $\{g \in G : \forall n, g_n \in \{-1, 0, 1\}\}$. The Cayley graph of G relative to S has infinite diameter but does not contain any geodesic ray.*

Proof. Let G and S be as in the statement of the proposition, and let \mathcal{G} denote the corresponding Cayley graph. The graph \mathcal{G} has infinite diameter, as $g \mapsto g_n$ defines a 1-Lipschitz map onto $\text{Cay}(\mathbb{Z}/n\mathbb{Z}, \{1\})$ and the diameter of $\text{Cay}(\mathbb{Z}/n\mathbb{Z}, \{1\})$ tends to infinity as n goes to infinity.

Now, let $\kappa : \mathbb{N} \rightarrow G$ be a path in \mathcal{G} and let us prove that κ is not geodesic. Without loss of generality, we may assume that $\kappa(0) = 0$. There is some $N \geq 2$ such that $\forall n \geq N, \kappa(1)_n = 0$. Fix such an N and notice that $d(\kappa(0), \kappa(N)) \leq N - 1$. Indeed, for every integer $n \geq 2$, either $n < N$ and the diameter of $\text{Cay}(\mathbb{Z}/n\mathbb{Z}, \{1\})$ is at most $N - 1$, or $n \geq N$ and $\kappa(1)_n = 0$ implies that $\kappa(N)_n$ lies in the ball of centre 0 and radius $N - 1$ in $\text{Cay}(\mathbb{Z}/n\mathbb{Z}, \{1\})$. \square

We also present the following example for its flavour of additive number theory.

Proposition 2.2 *Let S denote the (generating) subset of \mathbb{Z} consisting of the elements that can be written as $\sum_{n \geq 1} \varepsilon_n n!$ for some finitely supported element ε of $\{-1, 0, 1\}^{\mathbb{N}^*}$. The Cayley graph of \mathbb{Z} relative to S has infinite diameter but does not contain any geodesic ray.*

Proof. Let S be as in the proposition above and let \mathcal{G} denote the Cayley graph of \mathbb{Z} relative to S . First, let us prove that \mathcal{G} has infinite diameter. Let $N \geq 2$. Let π denote the canonical projection $\mathbb{Z} \rightarrow \mathbb{Z}/N!\mathbb{Z}$. The map π is 1-Lipschitz from \mathcal{G} onto $\text{Cay}(\mathbb{Z}/N!\mathbb{Z}, \pi(S))$. Besides, the identity map of $\mathbb{Z}/N!\mathbb{Z}$ is $(1 + 2 + \dots + (N - 1)!)$ -Lipschitz, hence $e(N - 1)!$ -Lipschitz, from $\text{Cay}(\mathbb{Z}/N!\mathbb{Z}, \pi(S))$ to $\text{Cay}(\mathbb{Z}/N!\mathbb{Z}, \{1\})$. As a result, \mathcal{G} has diameter at least $\frac{N!}{2} \times \frac{1}{e(N-1)!} = \frac{N}{2e}$. As N is arbitrary, the diameter of \mathcal{G} must be infinite.

Now, let us assume for contradiction that \mathcal{G} admits a geodesic ray $\gamma : \mathbb{N} \rightarrow \mathbb{Z}$. For every $k \in \mathbb{N}^*$, fix $\varepsilon^{(k)}$ such that $\Delta(k) := \gamma(k) - \gamma(k - 1) = \sum_{n \geq 1} \varepsilon_n^{(k)} n!$. First, let us show that without loss of generality, we may change γ — and the $\varepsilon^{(k)}$'s — so that for every n in the support of $\varepsilon^{(1)}$ and every $k \in \mathbb{N}^*$, we have $\varepsilon_n^{(k)} = \varepsilon_n^{(1)}$. If there is some n in the support of $\varepsilon^{(1)}$ such that there is some $k_n \in \mathbb{N}^*$ satisfying $\varepsilon_n^{(k_n)} \neq \varepsilon_n^{(1)}$, then we may change γ to have the same initial point, and the same $\Delta(k)$'s except for $\Delta_{\text{new}}(1) := \Delta_{\text{old}}(1) - \varepsilon_n^{(1)} n!$ and

$$\Delta_{\text{new}}(k_n) := \Delta_{\text{old}}(k_n) + \varepsilon_n^{(1)} n!$$

The new γ is still a geodesic path as it is a 1-Lipschitz map from \mathbb{N} to \mathcal{G} that agrees with the geodesic path γ_{old} on $\{0\} \cup \{k_n, k_n + 1, \dots\}$. As a result, if γ does not satisfy the desired property, then we may change $(\gamma, (\varepsilon^{(k)}))$ so that γ is still a geodesic ray but with an $\varepsilon^{(1)}$ of strictly smaller support. By Fermat's infinite descent, this process must stop at some suitable γ . Henceforth, we assume that we work with such a γ . We may further assume that $\Delta(1) > 0$, as $x \mapsto -x$ is a graph automorphism of \mathcal{G} .

Let M denote the maximum of the support of $\varepsilon^{(2)}$ — which is non-empty as γ is geodesic. Let $K > \max(M, \Delta(1))$. As γ has been suitably modified, one may define a 1-Lipschitz map

$$\lambda : \left\{ 0, \dots, \frac{K!}{\Delta(1)} \right\} \rightarrow \text{Cay}(\mathbb{Z}, S)$$

by setting $\lambda(0) = \gamma(0)$ and $\forall k \in \{1, \dots, K!/\Delta(1)\}$,

$$\lambda(k) - \lambda(k-1) = \Delta(k) - \Delta(1) + \mathbf{1}_{k=2}K!$$

As $\lambda(1) = \gamma(0)$ and $\lambda(K!/\Delta(1)) = \gamma(K!/\Delta(1))$, the fact that λ is 1-Lipschitz contradicts the fact that γ is geodesic, which ends the proof. \square

2.2 The generalised diameter

In order to capture the extension properties of geodesic paths, let us define the generalised diameter of a graph and the generalised radius of a rooted graph.

Let \mathcal{G} be a graph. We denote by $\text{Geod}(\mathcal{G})$ the set of the geodesic paths of \mathcal{G} that are well-defined at 0 but not at -1 . If γ and κ are two elements of $\text{Geod}(\mathcal{G})$, we say that γ **extends** κ if the domain of κ — some interval of \mathbb{Z} — is a subset of that of γ and γ and κ coincide on the domain of κ . We define a partial order on $\text{Geod}(\mathcal{G})$ by setting

$$\gamma \leq \kappa \iff \gamma \text{ extends } \kappa.$$

A *minimal element* of $\text{Geod}(\mathcal{G})$ is precisely a geodesic path which cannot be geodesically extended (while keeping the domain inside \mathbb{N}). For example, every geodesic ray of \mathcal{G} is a minimal element of $\text{Geod}(\mathcal{G})$.

The **generalised diameter** of \mathcal{G} is the smallest ordinal η such that there is a way to assign to each element γ of $\text{Geod}(\mathcal{G})$ an ordinal $\xi(\gamma)$ such that:

- for every $\gamma \in \text{Geod}(\mathcal{G})$, $\xi(\gamma) \leq \eta$,
- and for every γ and κ in $\text{Geod}(\mathcal{G})$, if $\gamma < \kappa$, then $\xi(\gamma) < \xi(\kappa)$.

If some η satisfies these conditions, there is a smallest one satisfying them: in that case, the generalised diameter³ is said to be **transfinite**. If there is no such η , we

³ which is well-defined and equal to the Krull ordinal of $(\text{Geod}(\mathcal{G}), \leq)$; see Section 2.6 in [4]

define the generalised diameter to be ∞ , and say that it is **truly infinite**, or **intransfinite**: it is larger than any ordinal number.

If (\mathcal{G}, o) is a rooted graph, we set $\text{Geod}(\mathcal{G}, o)$ to be the set of the elements of $\text{Geod}(\mathcal{G})$ that start at o . By using $\text{Geod}(\mathcal{G}, o)$ instead of $\text{Geod}(\mathcal{G})$, one defines the **generalised radius** of (\mathcal{G}, o) .

Here are a few important remarks.

- If a graph \mathcal{G} has a vertex o such that the generalised radius of (\mathcal{G}, o) is ∞ , then all its vertices satisfy this property. See Lemma 2.4 of [21], which does not rely on any form of local finiteness, only connectedness.
- The generalised diameter of a graph $\mathcal{G} = (V, E)$ is the supremum over $o \in V$ of the generalised radius of (\mathcal{G}, o) . In particular, if \mathcal{G} is transitive, then for every vertex o of \mathcal{G} , the generalised diameter of \mathcal{G} is equal to the generalised radius of (\mathcal{G}, o) .
- The generalised diameter of a graph is finite if and only if its diameter is finite, and in that case the two notions agree.
- The generalised diameter of a graph is intransfinite if and only if this graph contains a geodesic ray.

Example. The generalised radius and diameter allow us to capture finely the extension properties of the geodesic paths. For instance, a rooted graph (\mathcal{G}, o) has a generalised radius larger than or equal to $\omega 2 + 5$ if and only if there is a geodesic path γ starting at o and of length 5 that can be extended to geodesic paths of arbitrary length that themselves can be extended to geodesic paths of arbitrary length, i.e.

$$\exists \gamma \in \text{Geod}_5(\mathcal{G}, o), \forall n \geq 5, \exists \kappa \in \text{Geod}_n(\mathcal{G}, o), \kappa \leq \gamma \ \& \ \forall m \geq n, \exists \lambda \in \text{Geod}_m(\mathcal{G}, o), \lambda \leq \kappa$$

where $\text{Geod}_n(\mathcal{G}, o)$ denotes the set of the elements of $\text{Geod}(\mathcal{G}, o)$ of length n and the order \leq has been introduced after the proof of Proposition 2.2.

Let us introduce some useful notation. Let (\mathcal{G}, o) be a rooted graph. Define by transfinite induction, for every ordinal η , the set

$$D_\eta := \{ \gamma \in \text{Geod}(\mathcal{G}, o) : \forall \kappa \in \text{Geod}(\mathcal{G}, o), \kappa \leq \gamma \implies \exists \xi < \eta, \kappa \in D_\xi \}.$$

In words, one can say that we remove from $\text{Geod}(\mathcal{G}, o)$ its minimal elements and put them in D_0 ; once this erasure is performed, we remove the minimal elements of the remaining set, and put them, together with the previously erased elements, in D_1 ; etc. The generalised radius of (\mathcal{G}, o) is the smallest ordinal η such that $D_\eta = \text{Geod}(\mathcal{G}, o)$ — which is set to be ∞ if no such η exists.

If A denotes a set of ordinals, $\sup^+ A$ denotes $\sup\{\eta + 1; \eta \in A\}$. It is the unique ordinal η_0 such that for every ordinal η , the following equivalence holds:

$$\eta \geq \eta_0 \iff \forall \xi \in A, \eta > \xi.$$

Besides, we set $\sup^+ A \cup \{\infty\} := \infty$.

Let (\mathcal{G}, o) be a rooted graph. Given two vertices v and w , write $v \rightarrow w$ as an abbreviation for “ $\{v, w\}$ is an edge and $d(o, w) = d(o, v) + 1$ ”. Given a vertex v , any two geodesic paths from o to v belong exactly to the same D_η 's. As for any v there is a geodesic path from o to v , one can define the **label** ℓ_v of v in (\mathcal{G}, o) to be the least η such that γ belongs to D_η — where γ is a geodesic path from o to v and the definition does not depend on the choice of γ . If no such ordinal exists, set ℓ_v to be ∞ . This labeling is the unique assignment L of an ordinal or ∞ to each vertex satisfying the following equation: $\forall v, L_v = \sup^+ \{L_w : v \rightarrow w\}$.

The generalised radius of a rooted graph (\mathcal{G}, o) is equal to ℓ_o . Recall that if \mathcal{G} is transitive, then the generalised diameter of \mathcal{G} is the generalised radius of (\mathcal{G}, o) .

Theorem 2.3 *For every ordinal η , there is a Cayley graph of an abelian group the diameter of which is transfinite and equal to η .*

Proof. For every ordinal η , let us define some abelian group G_η and a generating subset S_η of G_η . We do so by transfinite induction, by using the following rule:

- if η can be written as $\xi + 1$, then set $G_\eta := G_\xi \oplus \mathbb{Z}/2\mathbb{Z}$ and $S_\eta := (S_\xi \times \{0\}) \cup \{(0, 1)\}$;
- otherwise, set $G_\eta := \bigoplus_{\xi < \eta} G_\xi$ and $S_\eta := \{g \in G_\eta : \forall \xi < \eta, g_\xi \in S_\xi\}$.

Let us prove by transfinite induction that for every η , the diameter of the Cayley graph \mathcal{G}_η of (G_η, S_η) is transfinite and equal to η . Let η be an ordinal number such that for every $\xi < \eta$, the diameter of \mathcal{G}_ξ is transfinite and equal to ξ . Let us show that the diameter of \mathcal{G}_η is transfinite and equal to η . Denote by $g \mapsto \ell_g^\eta$ the label-map of $(\mathcal{G}_\eta, 0)$.

First, assume that η can be written as $\xi + 1$. Let us prove that for every $(g, x) \in G_\eta = G_\xi \oplus \mathbb{Z}/2\mathbb{Z}$, one has $\ell_{(g,x)}^\eta = \ell_g^\xi + \mathbb{1}_{x=0}$. Notice that this implies that $\ell_{(0,0)}^\eta = \ell_0^\xi + 1 = \xi + 1 = \eta$, i.e. that the diameter of \mathcal{G}_η is transfinite and equal to η . To prove the claim, notice that for every g and h in G_ξ ,

1. $(g, 0) \rightarrow (h, 0) \iff (g, 1) \rightarrow (h, 1) \iff g \rightarrow h$;
2. $(g, 0) \rightarrow (h, 1) \iff g = h$;
3. $(g, 1) \not\rightarrow (h, 0)$.

For $(g, x) \in G_\eta$, let $\ell_{(g,x)} := \ell_g^\xi + \mathbb{1}_{x=0}$. What we need to prove is that for every $(g, x) \in G_\eta$, one has $\ell_{(g,x)} = \sup^+ \{\ell_{(h,y)} : (g, x) \rightarrow (h, y)\}$. Let $(g, x) \in G_\eta$. If $x = 1$, then by 1 and 3, one has

$$\begin{aligned} \ell_{(g,x)} &= \ell_g^\xi = \sup^+ \{\ell_h^\xi : g \rightarrow h\} \\ &= \sup^+ \{\ell_{(h,y)} : (g, x) \rightarrow (h, y)\}. \end{aligned}$$

If $x = 0$, then by 2, one has

$$\begin{aligned} \ell_{(g,x)} &= \ell_g^\xi + 1 = \sup^+ \{\ell_{(g,1)}\} \\ &\leq \sup^+ \{\ell_{(h,y)} : (g, x) \rightarrow (h, y)\}. \end{aligned}$$

Still assuming $x = 0$, it remains to show that

$$\ell_{(g,x)} \geq \sup^+ \{\ell_{(h,y)} : (g, x) \rightarrow (h, y)\}.$$

In other words, let (h, y) be such that $(g, 0) \rightarrow (h, y)$ and let us show that $\ell_{(g,0)} > \ell_{(h,y)}$. If $y = 0$, then by 1, one has $g \rightarrow h$, so that $\ell_{(g,0)} = \ell_g^\xi + 1 > \ell_h^\xi + 1 = \ell_{(h,y)}$. If $y = 1$, then by 2, one has $g = h$, so that $\ell_{(g,0)} = \ell_g^\xi + 1 > \ell_g^\xi = \ell_{(h,y)}$.

Now, assume that η is a limit ordinal. Let us prove that for every $g \in G_\eta = \bigoplus_{\xi < \eta} G_\xi$, one has

$$\ell_g^\eta = \sup \{\ell_{g_\xi}^\xi : \xi \in [0, \eta) \text{ and } d(0, g_\xi) = d(0, g)\}.$$

Notice that this implies that $\ell_0^\eta = \sup \{\ell_0^\xi : \xi < \eta\} = \sup \{\xi : \xi < \eta\} = \eta$, i.e. that the diameter of \mathcal{G}_η is transfinite and equal to η . For $g \in G_\eta$, let $\ell_g := \sup \{\ell_{g_\xi}^\xi : \xi \in [0, \eta) \text{ and } d(0, g_\xi) = d(0, g)\}$. What we need to prove is that for every $g \in G_\eta$, one has $\ell_g = \sup^+ \{\ell_h : g \rightarrow h\}$. Let $g \in G_\eta$. First, let us show that $\ell_g \geq \sup^+ \{\ell_h : g \rightarrow h\}$, i.e. let $h \in G_\eta$ be such that $g \rightarrow h$ and let us prove that $\ell_g > \ell_h$. As $g \rightarrow h$, the element h cannot be 0. Thus, by definition of (G_η, S_η) , the set $E_h := \{\xi \in [0, \eta) \text{ and } d(0, h_\xi) = d(0, h)\}$ is finite. As a result, it is sufficient to establish that for every $\xi \in E_h$, one has $\ell_g > \ell_{h_\xi}^\xi$. Let $\xi \in E_h$. As $d(0, h) = d(0, h_\xi)$ and $g \rightarrow h$, by definition of (G_η, S_η) , one has $d(0, g) = d(0, g_\xi) = d(0, h) - 1$. Thus, one has $g_\xi \rightarrow h_\xi$ and $\ell_g \geq \ell_{g_\xi}^\xi > \ell_{h_\xi}^\xi$. As a result, the inequality $\ell_g \geq \sup^+ \{\ell_h : g \rightarrow h\}$ holds.

Let us now establish that $\ell_g \leq \sup^+ \{\ell_h : g \rightarrow h\}$. Let thus $\xi \in [0, \eta)$ be such that $d(0, g_\xi) = d(0, g)$ and let us prove that $\ell_{g_\xi}^\xi \leq \sup^+ \{\ell_h : g \rightarrow h\}$. Recall that $\ell_{g_\xi}^\xi = \sup^+ \{\ell_{\tilde{h}}^\xi : g_\xi \rightarrow \tilde{h} \in G_\xi\}$. By definition of ℓ_h and as $g_\xi \rightarrow \tilde{h} \implies \exists h \in G_\eta, h_\xi = \tilde{h} \ \& \ g \rightarrow h$, one has $\sup^+ \{\ell_{\tilde{h}}^\xi : g_\xi \rightarrow \tilde{h} \in G_\xi\} \leq \sup^+ \{\ell_h : g \rightarrow h\}$. As a result, the inequality $\ell_{g_\xi}^\xi \leq \sup^+ \{\ell_h : g \rightarrow h\}$ holds, and the proof is complete. \square

2.2.1 From infinite rays to bi-infinite paths

Proposition 2.4 *Let \mathcal{G} be a Cayley graph of an abelian group. Assume that \mathcal{G} contains a geodesic ray. Then, \mathcal{G} contains a bi-infinite geodesic path.*

Proof. Let γ denote a geodesic ray. For $n \in \mathbb{N}$, set $s_n := \gamma(n+1) - \gamma(n)$. Define $\kappa : \mathbb{Z} \rightarrow V$ as follows:

- for every $n \geq 0$, set $\kappa(n) := \sum_{0 \leq k < n} s_{2k}$;
- for every $n < 0$, set $\kappa(n) := \sum_{0 \leq k < -n} s_{2k+1}$.

Note that κ is a path. For every $n \in \mathbb{N}$, one has

$$\kappa(n) - \kappa(-n) = \sum_{0 \leq k \leq 2n-1} s_k = \gamma(2n) - \gamma(0).$$

As γ is geodesic, the path κ is also geodesic. \square

For further use, say that a path κ is “zigzag-free” if for every k such that $\kappa(k)$ and $\kappa(k+2)$ are both defined, there is a unique path of length 2 connecting $\kappa(k)$ and $\kappa(k+2)$.

Proposition 2.5 *Lemma 1.5 does not hold if Condition 1 is removed, even if \mathcal{G} is assumed to be a Cayley graph of an abelian group and \mathfrak{P} is assumed to be invariant under $\kappa \mapsto (k \mapsto \kappa(-k))$.*

Proof. Let \mathcal{G} denote the Cayley graph of $G := \bigoplus_{n \geq 2} \mathbb{Z}/n\mathbb{Z}$

relative to the set consisting of the elements of G of which all entries are zero but precisely one which is equal to 1. Label each edge by the index of the coordinate where its extremities disagree. This labelling is preserved by every automorphism of \mathcal{G} , as for any $n \geq 3$, an edge is labelled n if and only if there is a “zigzag-free” path of length $n-1$ connecting its extremities — and an edge has label 2 if and only if its label is not larger than or equal to 3. Denote the label of an edge e by N_e .

Let \mathfrak{P} denote the set of the paths κ such that, for every k such that $\kappa(k-1)$ and $\kappa(k+1)$ are both defined, one has $|N_{\{\kappa(k-1), \kappa(k)\}} - N_{\{\kappa(k), \kappa(k+1)\}}| = 1$ and $N_{\{\kappa(k-1), \kappa(k)\}} \neq N_{\{\kappa(k), \kappa(k+1)\}}$. Since the labelling is preserved by every automorphism, condition 2 of Lemma 1.5 is satisfied. By definition, \mathfrak{P} is invariant under $\kappa \mapsto (k \mapsto \kappa(\pm k + k_0))$. Condition 4 holds because \mathfrak{P} is defined by a conjunction of conditions each of which involves finitely many (here four) coordinates. The set \mathfrak{P} contains the infinite path $\mathbb{N} \rightarrow G$ mapping k to $(\mathbb{1}_{n \leq k+1})_n$, so that Condition 5 also holds. However, \mathfrak{P} cannot contain a bi-infinite path. Indeed, let $\kappa \in \mathfrak{P}$ be bi-infinite. The map $k \mapsto N_{\{\kappa(k), \kappa(k+1)\}}$ is well-defined from \mathbb{Z} to \mathbb{N} and of the form $k \mapsto ak + b$ for some $a \in \{-1, 1\}$ and some $b \in \mathbb{N}$: such a map is contradictory, which ends the proof. \square

3 Weak isomorphy and transitivity

Proposition 3.1 *Let $G := \bigoplus_{n \geq 2} \mathbb{Z}/n\mathbb{Z}$, and let S denote the (generating) subset of G consisting of the elements of which all entries are zero but precisely one which is equal to 1. Let G' denote $G \oplus \mathbb{Z}$, and let S' denote $(S \times \{0\}) \cup \{(0, 1)\}$, which is a generating subset of G' . Then the Cayley graphs of (G, S) and (G', S') are weakly isomorphic but are not isomorphic.*

Proof. Let (G, S) and (G', S') be as in the statement of the proposition. Let \mathcal{G} denote $\text{Cay}(G, S)$ and \mathcal{G}' denote $\text{Cay}(G', S')$.

First, let us prove that \mathcal{G} and \mathcal{G}' are weakly isomorphic. Let G_0 denote $\bigoplus_{n \geq 2} \mathbb{Z}$, and let S_0 be the (generating) subset of G consisting of the elements of which

all entries are zero but precisely one which is equal to 1. Let $H := \bigoplus_{n \geq 2} n\mathbb{Z} \subset G_0$. One can realise (G, S) as $(G_0/H, \overline{S_0})$. Let $N \geq 0$, and let $H_N := \bigoplus_{n \geq 2} k_{N,n}\mathbb{Z} \subset G_0$, where $k_{N,n}$ is equal to n if $n < 2N+2$, 0 if $n = 2N+2$ and $n-1$ otherwise. One can realise (G', S') as $(G_0/H_N, \overline{S_0})$. As H_N and H have the same intersection with the ball of radius $2N+1$ and centre 0 in $\text{Cay}(G_0, S_0)$, the transitive graphs \mathcal{G} and \mathcal{G}' have isomorphic balls of radius N . Since N is arbitrary, \mathcal{G} and \mathcal{G}' are weakly isomorphic.

Now, let us show that \mathcal{G} and \mathcal{G}' are not isomorphic. Recall that we say that a bi-infinite geodesic path γ is “zigzag-free” if for every $n \in \mathbb{Z}$, there is a unique geodesic path between $\gamma(n)$ and $\gamma(n+2)$. Since \mathcal{G} does not admit a “zigzag-free” bi-infinite geodesic path while \mathcal{G}' does, these graphs cannot be isomorphic. \square

Proposition 3.2 *There is a graph that is weakly transitive and weakly isomorphic to a transitive graph, but not transitive.*

Proof. We use the notation of the proof of Proposition 3.1. Let V denote the set of non-empty finite words over the alphabet $G \cup G'$ such that the following two conditions hold:

1. the first letter and only the first letter belongs to G ;
2. only the first and the last letter of the word are allowed to be equal to 0_G or $0_{G'}$.

This is the vertex-set. Two vertices are declared to be adjacent if and only if one of the following (incompatible) conditions holds:

- (i) one word can be written as the other one followed by $0_{G'}$;
- (ii) one word can be written as the other one with the last letter modified by adding to it an element of S or S' .

This “graph” indeed is non-empty and connected. As \mathcal{G} and \mathcal{G}' are weakly transitive, the graph \mathcal{H} we have built is weakly transitive, and weakly isomorphic to the free product of \mathcal{G} (resp. \mathcal{G}') with the graph $K_2 = (\{0, 1\}, \{\{0, 1\}\})$ — which is defined in the same way but by removing Condition 1 and taking all occurrences of G and G' to be G (resp. G'). In particular, \mathcal{H} is weakly transitive and weakly isomorphic to a transitive graph.

Let us now show that \mathcal{H} is not transitive. Notice that removing an edge of \mathcal{H} leaves its endpoints in different connected components if and only if it exists due to Condition (i). As a result, a graph automorphism of \mathcal{H} must map every connected component of \mathcal{H}' onto a connected component of \mathcal{H} , where \mathcal{H}' stands for the “graph” with vertex-set V and where only the edges due to Condition (ii) are kept. As the connected component of $[0_G]$ is isomorphic to \mathcal{G} and that of $[0_G; 0_{G'}]$ is isomorphic to \mathcal{G}' , and as \mathcal{G} and \mathcal{G}' are not isomorphic, no automorphism of \mathcal{H} maps $[0_G]$ to $[0_G; 0_{G'}]$. Thus, \mathcal{H} is not transitive. \square

Acknowledgements. I would like to thank the Weizmann Institute of Science and my postdoctoral hosts — Itai Benjamini and Gady Kozma — for the excellent working conditions they have provided me. I am grateful to my postdoctoral hosts and Emmanuel Jacob for interesting discussions. I finally wish to thank the anonymous referee and the editorial team for useful comments.

References

- [1] David Aldous and Russell Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12(54):1454–1508, 2007.
- [2] Omer Angel and Oded Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.
- [3] László Babai. Vertex-transitive graphs and vertex-transitive maps. *Journal of graph theory*, 15(6):587–627, 1991.
- [4] Hyman Bass. Descending chains and the Krull ordinal of commutative Noetherian rings. *Journal of Pure and Applied Algebra*, 1(4):347–360, 1971.
- [5] Itai Benjamini. *Coarse Geometry and Randomness*, volume 2100 of *Lecture Notes in Mathematics*. Springer, 2013.
- [6] Itai Benjamini, Asaf Nachmias, and Yuval Peres. Is the critical percolation probability local? *Probab. Theory Related Fields*, 149(1-2):261–269, 2011.
- [7] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6(23):1–13, 2001.
- [8] Christophe Champetier and Vincent Guirardel. Limit groups as limits of free groups. *Israel Journal of Mathematics*, 146(1):1–75, 2005.
- [9] Yves de Cornulier, Luc Guyot, and Wolfgang Pitsch. On the isolated points in the space of groups. *Journal of Algebra*, 307(1):254–277, 2007.
- [10] Claude Chabauty. Limite d’ensembles et géométrie des nombres. *Bulletin de la Société Mathématique de France*, 78:143–151, 1950.
- [11] Christophe Champetier. L’espace des groupes de type fini. *Topology*, 39(4):657–680, 2000.
- [12] Reinhard Diestel and Imre Leader. A conjecture concerning a limit of non-Cayley graphs. *Journal of Algebraic Combinatorics*, 14(1):17–25, 2001.
- [13] Geoffrey Grimmett and Zhongyang Li. Locality of connective constants. ArXiv:1412.0150.
- [14] Geoffrey Grimmett and Zhongyang Li. Connective constants and height functions for Cayley graphs *Transactions of the AMS*, 369(8):5961–5980, 2017.
- [15] Rostislav Grigorchuk. Degrees of growth of finitely generated groups, and the theory of invariant means. *Mathematics of the USSR-Izvestiya*, 25(2):259, 1985.
- [16] Sébastien Martineau. The set of connective constants of Cayley graphs contains a Cantor space. *Electron. Comm. Probab.*, 22(12):1–4, 2017.
- [17] Yair Minsky. End invariants and the classification of hyperbolic 3-manifolds. *Current developments in mathematics*, 2002:111–141, 2002.
- [18] Howard Masur and Yair Minsky. Geometry of the complex of curves I: Hyperbolicity. *Inventiones mathematicae*, 138(1):103–149, 1999.
- [19] Howard Masur and Yair Minsky. Geometry of the complex of curves II: Hierarchical structure. *Geometric and Functional Analysis*, 10(4):902–974, 2000.
- [20] Sébastien Martineau and Vincent Tassion. Locality of percolation for abelian Cayley graphs. *Ann. Probab.*, 45(2): 1247–1277, 2017.
- [21] Matthew Tointon and Ariel Yadin. Horofunctions on graphs of linear growth. *C. R. Math. Acad. Sci. Paris*, 354(12):1151–1154, 2016.
- [22] Benjamin Weiss. Sofic groups and dynamical systems. *Sankhyā: The Indian Journal of Statistics, Series A*, 62(3):350–359, 2000.

SÉBASTIEN MARTINEAU

Laboratoire de Mathématiques d’Orsay
Université Paris-Sud,
France

E-mail address: sebastien.martineau@u-psud.fr