A conjecture concerning a limit of non-Cayley graphs

Reinhard Diestel and Imre Leader

Abstract

Our aim in this note is to present a transitive graph that we conjecture is not quasi-isometric to any Cayley graph. No such graph is currently known. Our graph arises both as an abstract limit in a suitable space of graphs and in a concrete way as a subset of a product of trees.

1. Introduction

Woess [7] asked the following beautiful and natural question: does every transitive graph 'look like' a Cayley graph? More precisely, is every connected locally finite vertex-transitive graph quasi-isometric to some Cayley graph?

Let us recall that graphs G and H are said to be *quasi-isometric* if there exist Lipschitz mappings $\theta\colon V(G) \to V(H)$ and $\phi\colon V(H) \to V(G)$ such that $\theta \circ \phi$ and $\phi \circ \theta$ are bounded. Equivalently, G and H are quasi-isometric if there exists a *quasi-isometry* from G to H, a function $\theta\colon V(G) \to V(H)$ for which there are constants $C, D \geqslant 1$ such that

$$d(\theta x, \theta y) \leq C d(x, y) \text{ for all } x, y \in G,$$

$$d(\theta x, \theta y) \geq \frac{1}{C} d(x, y) \text{ for all } x, y \in G \text{ with } d(x, y) \geq D,$$

$$d(\theta G, y) \leq D \text{ for all } y \in H,$$

where as usual d denotes the graph distance (in G or H) and $d(A, y) = \min\{d(x, y) : x \in A\}$.

Thus quasi-isometry is the natural notion of 'looks the same as, from far away'. Many properties of a graph are preserved under quasi-isometry – for example, the space of ends is preserved. As another example, if G and H are transitive graphs that are quasi-isometric then they have the same type of growth: polynomial or sub-exponential or exponential. See [2] for background on quasi-isometry.

Let us also recall that a *Cayley graph* is a graph arising in the following way. Let G be a group, with a finite generating set S closed under inversion (ie. $a \in S$ implies $a^{-1} \in S$). Then the *(left) Cayley graph* of G with respect to S has vertex-set G, with x joined to y if for some $a \in S$ we have $x = ay$. Note that G acts freely (ie. with no non-identity element having a fixed point) and transitively on this graph. In fact, Cayley graphs are characterised by this property: if G is any locally finite connected graph whose automorphism group $Aut\, G$ has a subgroup that acts transitively and freely on G then G is easily seen to be isomorphic to a Cayley graph of that subgroup. See [3] for more background on Cayley graphs. Let us also mention here that, up to quasi-isometry, the Cayley graph of a (finitely-generated) group does not depend on which generating set one chooses.

Several transitive graphs are known that are not (isomorphic to) Cayley graphs (see [4], [5]), but each of these is quasi-isometric to a Cayley graph. Indeed, the answer to Woess' question is known to be in the affirmative for several classes of graphs, including those of polynomial growth [6].

Our aim in this note is to present a graph that we believe is a counterexample to Woess' question. We construct a sequence of graphs that seem to look less and less like Cayley graphs. It turns out that this sequence has a limit when viewed in a certain natural space of graphs. We give this construction in Section 2.

Fortunately, this limit graph can also be expressed 'concretely', as a certain subset of a product of two trees. We do this in Section 3. We hope that this should make the conjecture that this graph is not quasi-isometric to a Cayley graph more susceptible to proof.

2. A limit of non-Cayley graphs

Our starting point is the following example of Thomassen and Watkins [5] of a non-Cayley graph. Let H be the graph obtained from a T_5 (the infinite 5-regular tree) by replacing each vertex by a $K_{2,3}$ (the complete bipartite graph with vertex classes of size 2 and 3) in the following way. Replace each vertex of T_5 by a disjoint copy of $K_{2,3}$, and then, for each edge uv of the T_5, identify a vertex of the $K_{2,3}$ corresponding to u with a vertex of the $K_{2,3}$ corresponding to v, in such a way that no point in any $K_{2,3}$ is identified more than once, and a vertex in a class of size 2 is always identified with a vertex in a class of size 3 and vice versa (see Figure 1). Then H is certainly transitive (of degree 5); why is it not a Cayley graph?

Suppose there is a subgroup S of $Aut\, H$ that acts freely and transitively on H, and let K be one of the $K_{2,3}$s making up H – say K has vertex classes $\{x_1, x_2\}$ and $\{y_1, y_2, y_3\}$. Any automorphism that sends an element of $\{y_1, y_2, y_3\}$ back into $\{y_1, y_2, y_3\}$ must fix K – indeed, it must map the set $\{x_1, x_2\}$ to itself, as $\{x_1, x_2\}$ is the only pair of two vertices that has 3 common neighbours and has a common neighbour in the set $\{y_1, y_2, y_3\}$. Hence the $\theta \in S$

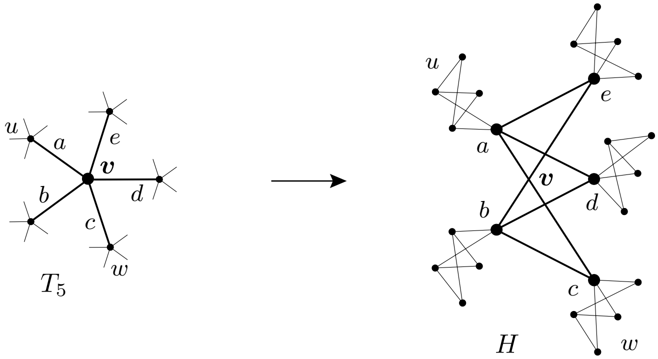

FIGURE 1. Constructing the non-Cayley graph H from T_5

sending y_1 to y_2 must swap x_1 and x_2, as must the $\theta' \in S$ sending y_1 to y_3. But then $\theta'\theta^{-1}$ sends y_2 to y_3 and fixes x_1, a contradiction.

Of course, H is still quasi-isometric to T_5 (which is the Cayley graph of the free group with 5 generators, each of order 2): we just have to map each $K_{2,3}$ back to the vertex of T_5 from which it was expanded. Thus the $K_{2,3}$s are too local to affect quasi-isometry: we would like to introduce something like 'larger $K_{2,3}$s' to have the same effect more globally. The following idea shows that these can indeed be obtained.

Roughly speaking, the reason why H is not Cayley is that the insertion of $K_{2,3}$s has introduced an 'orientation' which all automorphisms must preserve (but cannot all preserve without a fixed point). Indeed, each $K_{2,3}$ has a natural orientation of its edges from the 2-set to the 3-set, and put together they make H into a regular directed graph of in-degree 2 and out-degree 3. Our key observation now is that we can reverse this process of obtaining an orientation from $K_{2,3}$s to one of obtaining $K_{2,3}$s from an orientation. Indeed, if we *start from* a suitable orientation D_0 of T_5, namely, the regular orientation of in-degree 2 and out-degree 3, then our directed version of H (with all its useful 'Cayley-inhibiting' $K_{2,3}$s) is obtained from D_0 by one simple operation, which moreover can be iterated canonically to yield 'larger and larger $K_{2,3}$s' (see Figures 2 and 5): the operation of taking a directed line graph.

Let us do this in more detail. Given a directed graph D, the *line graph* of D is the directed graph D' whose vertices are the arcs uv of D, and in which such a vertex $uv \in V(D')$ sends an arc (of D') to another vertex $v'w' \in V(D')$ if and only if $v = v'$. Note that if D is regular with in-degree a and out-degree b then so is D'. The operation of taking a line graph can thus be iterated on regular directed graphs without increasing their degrees – a fact that will be vital to our whole approach.

A moment's thought shows that our directed version of H is indeed the line graph of D_0. So for $i = 1, 2, \ldots$ let D_i be the (directed) line graph of D_{i-1}, and let G_i denote the undirected graph underlying D_i. (Thus, $G_1 = H$.) Since

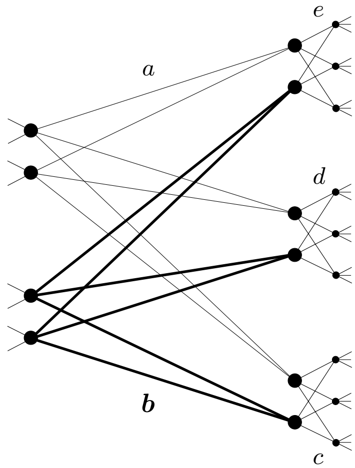

FIGURE 2. The portion of G_2 corresponding to the central $K_{2,3}$ in Figure 1

every D_i is regular with in-degree 2 and out-degree 3, all the G_i are 5-regular; it is therefore not unreasonable to expect that they converge to a graph 'at infinity' in some natural sense, and that this limit graph might not be quasi-isometric to a Cayley graph.

In order to define this limit graph precisely, let us pause to explain the (very simple) space of graphs we are working with. For a fixed positive integer r (which for us will always be 5), let $Q = Q_r$ denote the set of (isomorphism classes of) all connected r-regular transitive graphs. We introduce a metric on Q by setting $d(G, H) = 1/(n + 1)$ if n is the maximum positive integer such that there exists an isomorphism from the ball $B_G(0, n)$ to $B_H(0, n)$ sending 0 to 0. (Here 0 is any particular point of G or H, and $B_G(0, n)$ denotes the set of all points at graph distance at most n from 0.) This is a natural metric to use on Q; see for example [1]. The following easy compactness argument shows that it is indeed a metric.

Proposition 1. *Let $G, H \in Q$ with $d(G, H) = 0$. Then G and H are isomorphic.*

Proof. For each n, we have an isomorphism $\theta_n : B_G(0, n) \to B_H(0, n)$ sending 0 to 0. Now, there are only finitely many choices for an isomorphism from $B_G(0, 1)$ to $B_H(0, 1)$, so among the restrictions $\theta_1|B_G(0, 1), \theta_2|B_G(0, 1), \ldots$ there are infinitely many that agree: say

$$\theta_{i_1}|B_G(0, 1) = \theta_{i_2}|B_G(0, 1) = \ldots = \overline{\theta}_1.$$

Then, among the restrictions $\theta_{i_1}|B_G(0, 2), \theta_{i_2}|B_G(0, 2), \ldots$ there must be infinitely many that agree: say

$$\theta_{j_1}|B_G(0,2) = \theta_{j_2}|B_G(0,2) = \ldots = \overline{\theta}_2.$$

Continuing in this way, we obtain a sequence of isomorphisms $\overline{\theta}_n : B_G(0,n) \to B_H(0,n)$ with the property that for all $m \leq n$ we have $\overline{\theta}_n|B_G(0,m) = \overline{\theta}_m$. It follows that the union $\bigcup_{n\geq 1} \overline{\theta}_n$ is a (well-defined) isomorphism from G to H. □

A very similar argument shows that Q is compact:

Proposition 2. *Every sequence in Q has a convergent subsequence.*

Proof (sketch). Let $G_1, G_2, \ldots$ be any sequence of graphs in Q, each with a chosen point 0. Infinitely many of the G_i must have isomorphic 1-balls $B_{G_i}(0,1)$: say $B_{G_{i_1}}(0,1), B_{G_{i_2}}(0,1), \ldots$ are all isomorphic (with 0 mapping to 0). Among $G_{i_1}, G_{i_2}, \ldots$ we can find infinitely many graphs whose 2-balls are isomorphic (extending the isomorphisms of their 1-balls), and so on.

Continuing in this way, and choosing suitable partially nested isomorphisms to some fixed reference set X of vertices, we build up a nested sequence of finite graphs whose union G is a graph on X. Then G is connected and r-regular. To show that G is transitive, it is enough to show that for every choice of $x, y \in X$ and every n there is an isomorphism $B_G(x,n) \to B_G(y,n)$ mapping x to y; then the method of the proof of Proposition 1 yields an automorphism of G that takes x to y. But this is immediate: $B_G(x,n)$ and $B_G(y,n)$ are both contained in some ball $B_G(0,m)$; this ball coincides with the ball $B_{G_i}(0,m)$ in each of the graphs G_i of our mth subsequence; and G_i (being transitive) has an automorphism that takes x to y, and therefore also $B_G(x,n)$ to $B_G(y,n)$. Thus, $G \in Q$.

Finally, it is clear that any diagonal subsequence of the subsequences of $G_1, G_2, \ldots$ that we have chosen converges to G, as required. □

We remark in passing that, although it does not seem to help us, it is interesting to note that the set of Cayley graphs is a closed subset of Q: this may be proved by arguments similar to those in the proof of Proposition 2.

Let G be any limit point of the sequence $G_1, G_2, \ldots$. (A little thought shows that this sequence is actually convergent and thus has a unique limit; we shall prove this formally in the next section.) Is G still quasi-isometric to T_5? No, it is not: it will not be difficult to prove (see the next section) that G has only one end, and so cannot be quasi-isometric to T_5.

Of course, it is very hard to think about an abstract limit graph. Luckily, there is a far more down-to-earth description of G, which we give now.

3. An explicit construction

Our starting point here is that the (directed) line graph D_1 of D_0 is precisely the set of all directed paths in D_0 of length 1, with path uv joined to path wx if $v = w$. Similarly D_2, the line graph of the line graph of D_0, can be thought of as the set of all directed paths in D_0 of length 2, with uvw joined to xyz if $v = x$ and $w = y$. And so on:

Proposition 3. *The directed graph D_n is isomorphic to the graph whose vertices are the directed paths of length n in D_0, with an arc from $x_1x_2 \ldots x_{n+1}$ to $y_1y_2 \ldots y_{n+1}$ if $y_i = x_{i+1}$ for all $1 \le i \le n$.*

Proof. Induction on n. □

Let us see what, when n is large, a 'small' neighbourhood (of radius much less than n) of a vertex $v \in G_n$ looks like. Let P be the path in D_0 corresponding to v. Suppose that we wish to move from v to one of its five neighbours v' in G_n: how do we obtain the path P' corresponding to v' from the path P? If the edge $e = vv'$ is directed from v to v' in D_n, then P' is obtained from P by moving the last vertex of P to one of its three out-neighbours in D_0, while all the other vertices of P simply move to their successors along P. Similarly, if e is directed from v' to v, we obtain P' from P by moving the first vertex of P to one of its two in-neighbours in D_0, while all the other vertices of P are forced: they just move to their predecessors on P. See Figure 3.

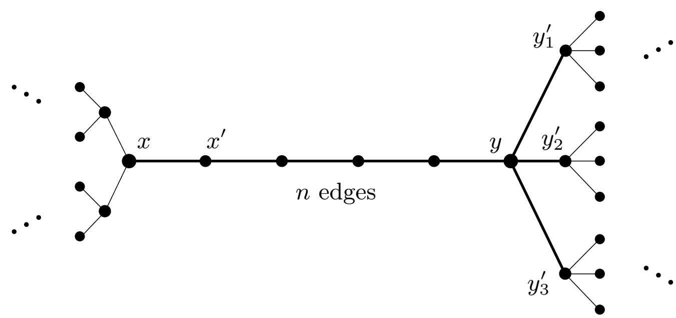

FIGURE 3. A path $x \ldots y \subset D$ corresponding to a vertex $v \in D_n$, and the paths $x' \ldots y_i \subset D$ corresponding to the 3 out-neighbours of v in D_n

So what does the open $n/2$-neighbourhood N of a point $v \in G_n$ look like? If (the path of) v has start vertex x and end vertex y, then the set of the start vertices of the points of N is disjoint from the set of their end vertices: indeed, these sets are contained in the open balls of radius $n/2$ about x and y respectively. So we may view the start and end vertices as behaving 'independently': as long as we stay in the ball of radius $n/2$ about v, the start vertices trace out part of a tree of in-degree 2 and out-degree 1, while the end vertices trace out part of a tree of in-degree 1 and out-degree 3.

This motivates the following explicit definition of a graph G^*, which will turn out to be the unique limit of our sequence $G_1, G_2, \ldots$. Let E be a 3-regular tree, oriented to have in-degree 2 and out-degree 1, and let F be the oriented 4-regular tree of in-degree 1 and out-degree 3. Fix a point $0 \in E$ and a point $0 \in F$. Let the *rank* $r(x)$ of a point $x \in E$ be the signed distance from 0 to x (so if the unique undirected path from 0 to x in E has s forward edges and t backward edges then $r(x) = s - t$), and define $r(y)$ in the same way for $y \in F$. Now define the directed graph D^* as follows. The vertex set of D^* is the set $\{(x, y) \in E \times F : r(x) = r(y)\}$, and D^* has an arc from (x, y) to (x', y') whenever $xx' \in E$ and $yy' \in F$ (Figure 4). Finally, let G^* be the undirected version of D^*.

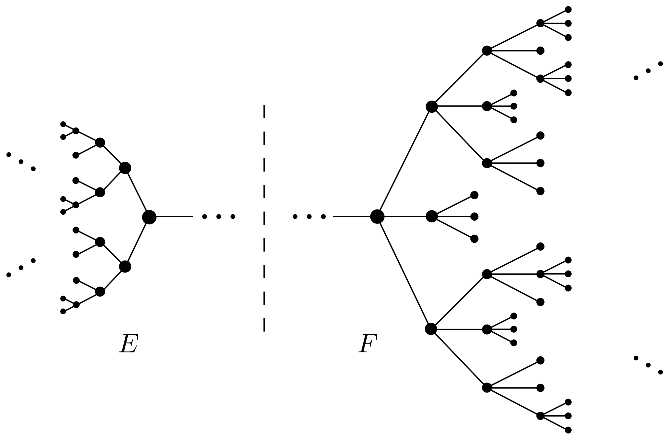

FIGURE 4. All directions are from left to right

Let us verify that G^* is indeed the unique limit of the sequence $G_1, G_2, \ldots$:

Proposition 4. *The sequence (G_n) converges to G^*.*

Proof. The directed graphs D_n and D^* have isomorphic $n/2$-neighbourhoods, so $d(G_n, G^*) \leq \frac{2}{n+2}$. □

We remark that it is now possible to define precisely what we mean by 'large $K_{2,3}$s' in the graph G^*. Given a vertex (x, y) of G^*, we have $r(x) = r(y)$ by definition of G^* and call this number the *rank* of (x, y), denoted again by $r(x, y)$. Given an integer $k > 0$, we call each of the (isomorphic) components of the subgraph of G^* spanned by the vertices of rank between 0 and k a $K_{2,3}$ *of order* k. It is not difficult (if a little tedious) to write down a formal partition of the vertex set of such a $K_{2,3}$ of order k into five classes, together with an adjacency rule between these classes based on adjacencies in G^*, so that the resulting graph is indeed a $K_{2,3}$. Instead, we offer a picture of a $K_{2,3}$ of order 4, shown in Figure 5.

Perhaps the most tangible evidence that we have for our conjecture that G^* is not quasi-isometric to a Cayley graph is that it is certainly not quasi-isometric to the obvious candidate of such a Cayley graph, the graph T_5:

Proposition 5. *G^* has only one end.*

Proof. We show that the deletion of any finite set S of vertices from G^* leaves only one infinite component. Let r be the smallest and s the largest rank of a vertex in S, and let S' be the set of all vertices that can be reached from S by a path whose vertices all have rank between r and s. Clearly S' is finite, so it suffices to show that $G^* - S'$ is connected.

Let vertices $(x_1, y_1), (x_2, y_2) \in G^* - S'$ be given, and let us show that we can move a token vertex (x, y) from (x_1, y_1) to (x_2, y_2) in G^* without hitting S'. We may assume that $s < r(x_1, y_1) \leqslant r(x_2, y_2)$: the proof for $r(x_1, y_1) \leqslant r(x_2, y_2) < r$ is analogous, and any vertex of rank between r and s can be joined to a vertex of rank $> s$ by any path of increasing rank (which avoids S' by definition of S').

Starting with $(x, y) = (x_1, y_1)$, we first move (x, y) towards the right in Figure 4 (formally: with increasing rank, and thus avoiding S') until x lies on a left (i.e. backward oriented) ray R in E that avoids S'_E, the set of first components of the vertices in S'. We now move (x, y) to the left, keeping x on R, until y lies to the left of y_2 in F. We then move (x, y) right again until $y = y_2$; since x stays on R during this move, this keeps us outside S' until we are back at points of rank $> s$. We now move on towards the right until x lies to the right of x_2 in E, and back again until $(x, y) = (x_2, y_2)$. □

How might one show that G^* is not quasi-isometric to a Cayley graph? The first hope, of course, would be to imitate our proof of why H is not a Cayley graph, using a sufficiently large $K_{2,3}$ instead of the actual $K_{2,3}$s in H. However, we have been unable to make this approach work and are not sure that it can work: although it is straightforward to translate the canonical group action on a hypothetical Cayley graph quasi-isometric to G^* to similar 'quasi-automorphisms' of G^*, the fuzziness introduced seems to blur the difference between the sizes of the two vertex classes even of large $K_{2,3}$s (which are 2^n and 3^n, respectively), a difference central to the 'non-Cayley' proof for H.

As a more global approach we might try to show that every quasi-automorphism of G^* preserves the natural orientation of all sufficiently large $K_{2,3}$s, mapping their left sets (their vertices of minimal rank) to the left of the images of their right sets (their vertices of maximal rank). Then any Cayley graph quasi-isometric to G^* would have two 'directions' invariant under all its automorphisms (not just under its own group action), and in which it grows at different speeds: 2^n 'to the left' and 3^n 'to the right'. Can this happen in a

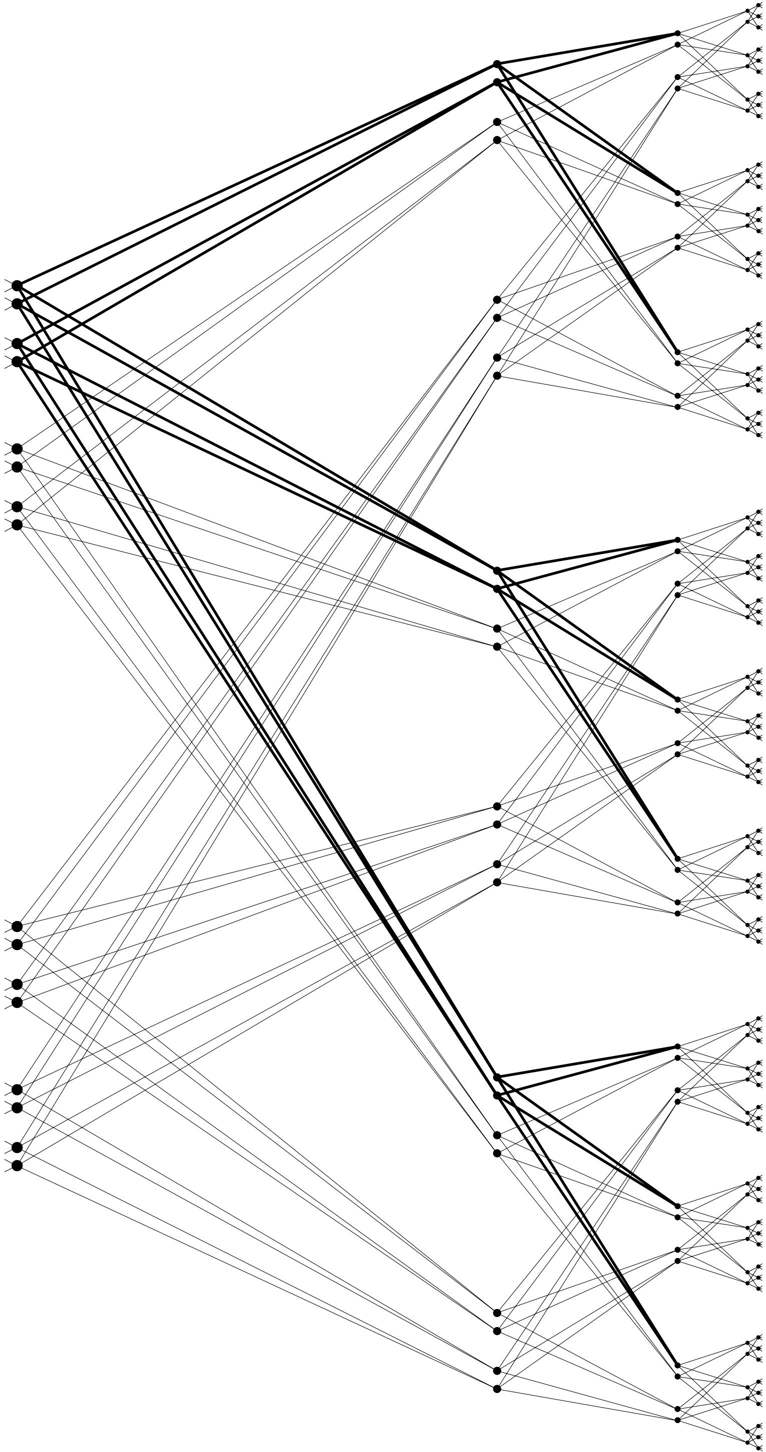

FIGURE 5. A $K_{2,3}$ of order 4 in G^*, and a (bold) $K_{2,3}$ of order 2

Cayley graph? (Recall that the overall growth speed of a graph is not preserved under quasi-isometries: for example, the trees T_3 and T_4 are quasi-isometric.)

Acknowledgement

We are grateful to Wolfgang Woess for many interesting conversations.

References

[1] Babai, L., Vertex-transitive graphs and vertex-transitive maps, *J. Graph Theory* **15** (1991), 587-627.

[2] Gromov, M., Asymptotic invariants of infinite groups, in *Geometric Group Theory*, Volume 2 (Niblo, G.A. and Roller, M.A., eds.), Cambridge University Press, 1993.

[3] Magnus, W., Karrass, A. and Solitar, D., *Combinatorial Group Theory*, Dover, 1976.

[4] Soardi, P.M. and Woess, W., Amenability, unimodularity, and the spectral radius of random walks on infinite graphs, *Math. Z.* **205** (1990), 471-486.

[5] Thomassen, C. and Watkins, M.E., Infinite vertex-transitive, edge-transitive, non-1-transitive graphs, *Proc. Amer. Math. Soc.* **105** (1989), 258-261.

[6] Trofimov, V.I., Graphs with polynomial growth, Math. USSR-Sb. **51** (1985), 405-417.

[7] Woess, W., Topological groups and infinite graphs, *Discrete Math.* **94** (1991), 1-12.

Department of Mathematics, University of Hamburg
Bundesstr. 55, 20146 Hamburg, Germany

and

Department of Mathematics, University College London,
Gower St., London WC1E 6BT, England