Invariant random graphs with iid degrees in a general geography

Johan Jonasson [*†‡]

May 2008

Abstract

Let D be a non-negative integer-valued random variable and let $G = (V, E)$ be an infinite transitive finite-degree graph. Continuing the work of Deijfen and Meester [5] and Deijfen and Jonasson [4], we seek an $\mathbf{Aut}(G)$-invariant random graph model with V as vertex set, iid degrees distributed as D and finite mean connections (i.e. the sum of the edge lengths in the graph metric of G of a given vertex has finite expectation). It is shown that if G has either polynomial growth or *rapid* growth, then such a random graph model exists if and only if $\mathbb{E}[D\,R(D)] < \infty$. Here $R(n)$ is the smallest possible radius of a combinatorial ball containing more than n vertices. With rapid growth we mean that the number of vertices in a ball of radius n is of at least order $\exp(n^c)$ for some $c > 0$. All known transitive graphs have either polynomial or rapid growth. It is believed that no other growth rates are possible.

When G has rapid growth, the result holds also when the degrees form an arbitrary invariant process. A counter-example shows that this is not the case when G grows polynomially. For this case, we provide other, quite sharp, conditions under which the stronger statement does and does not hold respectively.

Our work simplifies and generalizes the results for $G = \mathbb{Z}$ in [4] and proves e.g. that with $G = \mathbb{Z}^d$, there exists an invariant model with finite mean connections if and only if $\mathbb{E}[D^{(d+1)/d}] < \infty$. When G has exponential

*Chalmers University of Technology
†Göteborg University
‡`jonasson@math.chalmers.se`

growth, e.g. when G is a regular tree, the condition becomes $\mathbb{E}[D \log D] < \infty$.

Keywords: Random graphs, degree distribution, automorphism, invariant model, mass-transport principle, unimodular graph, polynomial growth, intermediate growth, exponential growth.

AMS 2000 Subject Classification: 05C80, 60B15, 60G10, 60G50, 60K35.

1 Introduction

In recent years there has been an increasing interest in the use of random graph models as models for different complex structures. For such applications the original Erdős-Rényi model will not do. One reason for this is that the degree distribution of these structures is widely different from what one gets from the Erdős-Rényi model. Therefore it has been a natural step to construct models, where the degrees of different vertices are iid random variables with a given desired distribution F. A handful of models of this type have been proposed by different authors, see [5] and [4] and the references therein.

A second reason for the need for new random graph models is that many of the networks one wants to model, exhibit a notion of geography; the vertices have a well-defined position in space. Therefore it is natural to ask for models which are, in addition to the above, also geographically invariant in some proper sense. This problem was introduced by Deijfen and Meester [5], who constructed an invariant model on $\mathbb{Z}$ (i.e. a random graph model on the vertices of $\mathbb{Z}$ whose edge configuration is invariant under the automorphisms of $\mathbb{Z}$). Their model leads to well-defined graphs, provided that F has finite mean. However, the expected edge lengths, and hence the expected total edge length of a vertex, turn out to be infinite for any F. This leads to the question if one can construct models where this is not the case and, if so, what is a necessary and sufficient condition on F for this to be possible. A first answer came in [4] where an invariant model on $\mathbb{Z}$ with finite expected total edge length of a vertex, was constructed under the condition that F has finite second moment. It is easily seen that finite second moment is also necessary.

The present paper is a natural continuation of [4]; we seek to extend the results from there to other geographies. We will do this in the most general sense possible. Let $G = (V, E)$ be an infinite transitive finite-degree graph with either polynomial or rapid (for definition, see the next section) growth. (In fact, it is believed, but still not confirmed, that no other growth rates are possible.) We will prove the existence of an $\mathbf{Aut}(G)$-invariant model on G, with finite expected total edge length per vertex, when $\mathbb{E}[D\,R(D)] < \infty$. Here $R(n)$, $n = 1, 2, \ldots$, is the radius of the smallest possible combinatorial ball with more than n vertices and D is distributed according to F.

In the polynomial growth case, our model will be based on a discrete version of a "stable marriage of Poisson and Lebesgue" of Hoffman, Holroyd and Peres [7]. For $G = \mathbb{Z}$ this model is similar in spirit to the one in [4], but it turns out to be more amenable to generalization. For $G = \mathbb{Z}^d$, $d \geq 1$, $R(n)$ is of order $n^{1/d}$ so the condition $\mathbb{E}[D\,R(D)] < \infty$ becomes $\mathbb{E}[D^{(d+1)/d}] < \infty$. When G is a regular tree of degree at least 3, we get $\mathbb{E}[D \log D] < \infty$. The latter condition also applies to some more exotic geographies such as the Trofimov graph, the Diestel-Leader graphs (see e.g. [3]) or the lamplighter groups (see e.g. [9]), all of exponential growth. If G is the Grigorchuk group (see [6]) whose growth rate is known to be between $\exp(\sqrt{n})$ and $\exp(n^{0.768})$ (see [1]) we get that $\mathbb{E}[D(\log D)^2] < \infty$ is sufficient and $\mathbb{E}[D(\log D)^{1.302}] < \infty$ is necessary for the existence of an automorphism invariant graph with finite mean connections.

It will also follow, for graphs of rapid growth, that there exists a model, that gives finite expected edge length per vertex under the same condition, as soon as the degrees form *any* automorphism invariant process. On the other hand, we demonstrate that the same condition does not apply when G has polynomial growth. In this case we provide other, quite sharp, conditions for when a model of the desired type does and does not exist.

For a longer and fuller introduction to the subject we refer to [4] and the references therein.

In the next section, the necessary concepts and tools are introduced. In the third, and final, section, we state and prove the main result (i.e. we construct an analyze the promised models.)

2 Preliminaries

Let $G = (V, E)$ be a graph. We will write $\mathbf{dist}_G$ for the graph metric on G, i.e. for two vertices, u and v, $\mathbf{dist}_G(u, v)$ is the minimum number of edges of a path between u and v.

An *automorphism* on G is a bijective map $g : V \to V$ such that $\{gu, gv\} \in E$ if and only if $\{u, v\} \in E$. We put $\mathbf{Aut}(G)$ for the group of automorphisms on

G. The graph G is said to be *transitive* if, for every $u, v \in V$, there exists a $g \in \mathbf{Aut}(G)$ such that $gu = v$.

Assume from now on that G is infinite, transitive and of finite degree. Fix an arbitrary vertex, o, henceforth to be thought of as the "origin". Let, for $n = 0, 1, 2, \ldots$, $\mathcal{V}(n)$ be the number of vertices in the combinatorial ball, $B[o, n] := \{v \in V : \mathbf{dist}_G(v, o) \leq n\}$, of radius n at the origin, and let $\mathcal{A}(n) := \mathcal{V}(n) - \mathcal{V}(n-1)$ with $\mathcal{A}(0) = 1$. Put $R(x) := \min\{n : \mathcal{V}(n) > x\}$, $x \in [0, \infty)$ and let $\mathcal{R} : [0, \infty) \to [0, \infty)$ be the concave hull of R (i.e. the smallest concave function greater than or equal to R). Note that since $\mathcal{V}(n)$ increases with n at at least linear rate, $|\mathcal{R}(x) - R(x)| < 1$ and $\mathcal{R}(n) = R(n)$ for every n in the range of $\mathcal{V}$. Thus finiteness of $\mathbb{E}[D\,R(D)]$ is equivalent to finiteness of $\mathbb{E}[D\mathcal{R}(D)]$. It should also be kept in my mind that $\mathcal{R}$ is the inverse of a suitable extension of $\mathcal{V}$. (The reason for introducing the function $\mathcal{R}$ is that later on we will need this concavification of R in order to apply Jensen's inequality.)

We will make use of a few, among graph theorists well known, facts about the possible behavior of $\mathcal{V}(n)$; for background information, see [8] and the references there. The function $\mathcal{V}(n)$ grows either at rate n^k for some integer $k \geq 1$, at rate at least e^{cn} for some constant $c > 0$ or at a rate such that for any $k < \infty$ and any $c > 0$, $n^k < \mathcal{V}(n) < e^{cn}$ for all sufficiently large n. These possible growth rates are referred to as polynomial growth, exponential growth and intermediate growth respectively. The existence of transitive graphs with intermediate growth is highly non-trivial. The first example was found by Grigorchuk [6] in the early 1980's; his example is nowadays known as the Grigorchuk group (or Grigorchuk's first group). It is today widely believed that there exists $c > 0$ such that no transitive graph of intermediate growth can have a growth rate lower than $e^{c\sqrt{n}}$.

We will say that G has *rapid* growth if, for some $c > 0$, $\mathcal{V}(n) > \exp(n^c)$ for all sufficiently large n. Hence, as the matter stands today, the common belief is that any G with super-polynomial growth, must grow rapidly.

For $v \in V$, write $S_v := \{g \in \mathbf{Aut}(G) : gv = v\}$, the *stabilizer* of v. We say that G is *unimodular* if, for every $u, v \in V$,

$$|S_v u| = |S_u v|.$$

Unimodularity is a very mild condition, satisfied by any graph you are likely to meet in practice. E.g. any Cayley graph of a finitely generated group is unimodular, see e.g. [3]. The simplest example of a non-unimodular transitive graph is the so called Trofimov graph, which one gets by identifying a direction in a binary tree and then drawing an edge from each vertex to its grandparent with respect to this direction, see [3] again. A more substantial class of examples, also to be found in [3], is the Diestel-Leader graphs.

Next we introduce the *mass-transport principle*. Let $X \in \{0,1\}^V$ be an $\mathbf{Aut}(G)$-invariant random process and let $m : V \times V \times \{0,1\}^V \to [0,\infty)$ be such that $m(u,v,x) = m(gu,gv,gx)$ for all $u,v \in V$, $x \in \{0,1\}^V$ and $g \in \mathbf{Aut}(G)$. The function m is thus a diagonally invariant function and we think of $m(u,v,x)$ as the amount of mass transported from the vertex u to the vertex v when $X = x$.

Theorem 2.1 (Mass-transport principle) *We have that*

$$\mathbb{E}\sum_{v\in V} m(u,v,X) = \mathbb{E}\sum_{v\in V} m(v,u,X)\frac{|S_v u|}{|S_u v|}.$$

In particular, if G is unimodular, then the expected amount of mass transported out of a vertex equals the expected amount of mass transported into it, i.e. for any vertex u,

$$\mathbb{E}\sum_{v\in V} m(u,v,X) = \mathbb{E}\sum_{v\in V} m(v,u,X).$$

When G is a Cayley graph, the mass-transport principle is almost immediate. For a proof of the general case, see [3]. We can now prove the following important lemma, which generalizes Proposition 4.2 of [4]. Since the result may be of independent interest, we give it, with very little extra effort, in a slightly more general form than necessary. (Only the case $k = 1$ will be needed.) Let again $X \in \{0,1\}^V$ be an automorphism invariant random process. Call a vertex, v, such that $X(v) = 1$, a *site* of X and put $N_v(k)$ for the distance from v to its k'th nearest site, other than v itself, $k = 1,2,\ldots$. (Let v make a uniform random ordering of all vertices at the same distance from v for all distances. In this way the k'th nearest site is well defined.) Let $p := P(X(v) = 1)$.

Lemma 2.1 *There exists a constant $K < \infty$, depending on X only via p, such that*

$$\mathbb{E}[N_v(k)|X(v) = 1] < K\mathcal{R}\Big(\frac{k}{p}\Big),$$

Proof. For all sites v, let $C_v = C_v(1)$, the 1-*cell* of v, be the set of vertices that are closer to v than to any other site. (I.e. C_v is v's "discrete Voronoi-cell".) In general, define $C_v(k)$, the k-*cell* of v, $k = 2,3,\ldots$, as the set of vertices that are closer to v than to any other vertex, when the $k-1$ nearest sites of v are ignored. Fix k. Let

$$m(u,v,x) = \mathbf{1}_{\{u \in C_v(k)\}}(x),$$

i.e. let every vertex u send unit mass to every site whose k-cell contains u. Then, since no vertex will send any mass to any vertex that is not among its k nearest,

$$\sum_{u \in V} m(v,u,X) \leq k.$$

Now assume first that G is unimodular. We have that

$$\mathbb{E}\sum_{u \in V} m(u,v,X) = \mathbb{E}\Big[|C_v(k)|\mathbf{1}_{\{X(v)=1\}}\Big] = p\mathbb{E}\Big[|C_v(k)|\Big|X(v)=1\Big].$$

Hence the mass-transport principle yields

$$\mathbb{E}\Big[|C(v)|\Big|X(v)=1\Big] < \frac{k}{p}.$$

However, by definition of $\mathcal{R}$, $N_v(k)$ and $C_v(k)$, we have that $N_v(k) < 3\mathcal{R}(|C_v(k)|)$. Therefore Jensen's inequality implies that

$$\mathbb{E}[N_v(k)|X(v)=1] < 3\mathcal{R}\Big(\mathbb{E}\Big[|C_v(k)|\Big|X(v)=1\Big]\Big) < 3\mathcal{R}\Big(\frac{k}{p}\Big)$$

as desired.

In the non-unimodular case, let

$$a := \max\Big\{\frac{|S_u o|}{|S_o u|} : \mathbf{dist}_G(u,o) = 1\Big\}.$$

Then $a > 1$. If $N_v(k) \geq r$, then, by transitivity, $C_v(k)$ must contain a path $v_0 = v, v_1, x_2, \ldots, v_{\lfloor r/2 \rfloor}$ where $|S_{v_{i+1}}v_i|/|S_{v_i}v_{i+1}| = a$ for every i. Hence,

$$\sum_{v \in V} m(v,u,X)\frac{|S_v u|}{|S_u v|} > \mathbf{1}_{\{X(v)=1\}} \sum_{i=0}^{\lfloor N_v(k)/2 \rfloor} a^i > \mathbf{1}_{\{X(v)=1\}} e^{bN_v(k)}$$

for $b > \log(a)/3$. The mass-transport principle now gives

$$\mathbb{E}[e^{bN_v(k)}\Big|X(v)=1] < \frac{k}{p}.$$

Taking logarithms and using Jensen's inequality gives

$$\mathbb{E}[N_v(k)|X(v)=1] < \frac{1}{b}\log(\frac{k}{p}) \leq \frac{\log(d-1)}{b}\mathcal{R}\Big(\frac{k}{p}\Big).$$

(In fact, $\mathcal{R}$ is indeed logarithmic; any non-unimodular graph is nonamenable and hence grows exponentially, see [3].) Taking $K = 3\log(d-1)/\log a$ finishes the proof. □

For the next lemma, let again X be an $\mathbf{Aut}(G)$-invariant $\{0,1\}^V$-valued random process and $p := P(X(v) = 1)$. Recall that a *matching* of a set of elements is a graph having these elements as vertices, for which the degree of every vertex is 1. In a *partial matching* every vertex has degree 0 or 1.

Lemma 2.2 *If G has at least quadratic growth, then there exists an $\mathbf{Aut}(G)$-invariant matching of the sites of X such that*

$$\mathbb{E}[\mathbf{dist}_G(o,w_o)\mathbf{1}_{\{X(o)=1\}}] < 10Kp\mathcal{R}\Big(\frac{1}{p}\Big),$$

where w_o is the vertex matched with o and K is the constant from Lemma 2.1.

Proof. Since $\mathcal{V}(n)$ grows at least quadratically, $\mathcal{R}(x)$ grows at most at rate $\sqrt{x}$.

Put $X_1 = X$ and $p_1 = p$. Construct a random directed graph on the sites of X_1 by drawing a directed edge from each site to its nearest other site. In order to break ties, let W_v, $v \in V$, be iid uniform[0,1] random variables, to be thought of as *weights*. Then, if the nearest site is not unique, pick the one of the candidates with the largest weight. Call a site whose in-degree, in the resulting digraph, is positive a *stem site*. It is clear that a.s. the digraph will consist of finite components, whose underlying graphs are trees and whose stem sites are connected to each other by a directed path ended with two sites pointing to each other.

By Lemma 2.1, the expected $\mathbf{dist}_G$-length given that o is a stem site of X_1, of a uniformly chosen directed edge pointing to o, is bounded by $K\mathcal{R}(1/p_1)$. (This is so, since choosing a typical directed edge amounts to choosing at random an edge pointing to a typical stem site.) Since the edge pointing away from o is at most as long as the shortest one pointing to it, the bound also applies if we choose among all the directed edges incident to o.

Now we partially match the sites in the following way. Fix a given component and consider the underlying tree. If the tree consists of only two sites, then match these to each other. If not, then for each of the stem sites, match its neighbors to each other at random, as far as possible (in the sense that if the number of neighbors is odd, then a randomly chosen one is left out.) Then at least $2/3$ of the sites have been connected to some other site. The length of a connection is bounded by the sum of two uniformly chosen directed edges incident to a stem sites. Thus the expected length of a connection of a given site on the event that it exists, is bounded by $2Kp_1\mathcal{R}(1/p_1)$. Now, what we produced may not be a partial matching of the given component, since some paths of more than two sites may have formed. If so, then remove randomly every second connection of such a path. Then the same bound obviously still applies and still at least $4/9$ of the sites are matched. Letting B_1 be the event that o is a site and gets matched by the this procedure, we have shown that

$$\mathbb{E}[\mathbf{dist}(o, w_o)\mathbf{1}_{B_1}] < 2Kp_1\mathcal{R}\Big(\frac{1}{p_1}\Big).$$

Next, define a new invariant process, $X_2(v)$, $v \in V$, by letting the sites of X_2 be the sites of X_1 that were not matched. Let $p_2 := P(X_2(v) = 1) \leq 4p_1/9$. Now repeat the matching procedure above with X_2 replacing X_1. The same arguments tell us that the probability that a site of X_2 gets matched is also at least $4/9$ and

$$\mathbb{E}[\mathbf{dist}(o, w_o)\mathbf{1}_{B_2}] < 2Kp_2\mathcal{R}\Big(\frac{1}{p_2}\Big),$$

where B_2 is the event that o is a site of X_2 and gets matched when the matching procedure is applied to X_2.

Repeat this recursively, by letting each new process X_n consist of the sites of the previous process X_{n-1} that do not get matched when the matching procedure is applied to X_{n-1}. Then a.s. every site of $X = X_1$ will get matched eventually, and we get

$$\mathbb{E}[\mathbf{dist}_G(o, w_o)\mathbf{1}_{\{X(o)=1\}}] = \sum_{i=1}^{\infty} \mathbb{E}[\mathbf{dist}(o, w_o)\mathbf{1}_{B_i}]$$
$$< 2K\sum_{i=1}^{\infty} p_i\mathcal{R}\Big(\frac{1}{p_i}\Big) = 2Kp_1\mathcal{R}\Big(\frac{1}{p_1}\Big)\Big(1 + \sum_{i=2}^{\infty}\frac{p_i\mathcal{R}(1/p_i)}{p_1\mathcal{R}(1/p_1)}\Big).$$

Since $p_i \leq (5/9)^{i-1}p_1$ and $\mathcal{R}(x)$ grows at most at rate $\sqrt{x}$,

$$1 + \sum_{i=2}^{\infty}\frac{p_i\mathcal{R}(1/p_i)}{p_1\mathcal{R}(1/p_1)} \leq \sum_{i=0}^{\infty}\Big(\frac{\sqrt{5}}{3}\Big)^i < 5.$$

Inserting into the above formula completes the proof. □

3 Main results and proofs

Fix an arbitrary infinite transitive finite-degree graph G. Consider an $\mathbf{Aut}(G)$-invariant random graph model on the vertices of G, where the degree process, $\{D_v\}_{v\in V}$, is distributed according to an automorphism invariant measure on $\mathbb{Z}_+^V$. Let $D = D_o$ be the degree of the origin, distributed according to the given distribution function F. We will write $f(n) := 1 - F(n-1) = P(D \geq n)$. Put T for the sum of the edge lengths for the edges having the origin as one of its end vertices. Of course, we measure edge lengths with the $\mathbf{dist}_G$-distance between its end vertices.

Let $R := R(D) - 1$. We have

$$T \geq \sum_{v\in B[o,R]\setminus\{o\}} \mathbf{dist}_G(o,v) = \sum_{n=1}^{R} n\mathcal{A}(n) = R\mathcal{V}(R) - \sum_{n=1}^{R-1} \mathcal{V}(n)$$

$$\geq \frac{1}{2}R\mathcal{V}(R) \geq \frac{1}{2}D(R(D)-1)$$

where the first inequality follows from the (super)linear growth of $\mathcal{V}$. Hence if $\mathbb{E}[D\,R(D)]$ is infinite, then so is $\mathbb{E}T$. This proves the only if-parts of the following results.

Theorem 3.1 *Let G be an infinite transitive finite-degree graph of either polynomial or rapid growth. Let F be a probability distribution supported on the nonnegative integers and let D be a random variable with distribution F. There exists an $\mathbf{Aut}(G)$-invariant simple random graph model on G, with iid degrees distributed according to F and $\mathbb{E}T < \infty$, if and only if $\mathbb{E}[D\mathcal{R}(D)] < \infty$.*

A more general, partly stronger and partly different result is:

Theorem 3.2 *If the degree process $\{D_v\}_{v\in V}$ forms an arbitrary automorphism invariant process, where the individual D_v's are distributed according to F, then the only if-part of Theorem 3.1 still holds.*

The if-part holds if G grows rapidly, but fails if G has polynomial growth. In the polynomial growth case with $\mathcal{V}(n) = \Theta(n^d)$, $d \in \{1,2,\ldots\}$, the following holds.

(i) If $d = 1$, then for any $k < \infty$, there are invariant degree processes for which $\mathbb{E}[D^k] < \infty$, and for which $\mathbb{E}T = \infty$ for any invariant random graph on V.

(ii) If $d = 2$ and $\alpha > 0$, then

$$\mathbb{E}[D^{(d+1)/(d-1)+\alpha}] < \infty$$

is sufficient to guarantee the existence of an invariant random graph model with $\mathbb{E}T < \infty$. On the other hand, there are examples of invariant degree processes with $\mathbb{E}[D^{(d+1)/(d-1)-\alpha}] < \infty$ for which T necessarily has infinite mean for any invariant random graph model on V.

The rest of the paper will be concerned with proving the if-parts of Theorem 3.1 and Theorem 3.2.

3.1 Graphs of rapid growth

We will repeatedly use Lemma 2.2. Assume that the D_v's form an arbitrary $\mathbf{Aut}(G)$-invariant process. Define invariant $\{0,1\}^V$-valued processes $X_1, X_2, \ldots$ by letting $X_n(v) = \mathbf{1}_{\{D_v \geq n\}}$. Note that $P(X_n(v) = 1) = f(n)$. First, match the sites of X_1 according to a matching satisfying Lemma 2.2. Next, color the sites of X_2 with a uniformly chosen 2-coloring such that whenever two sites of X_2 were matched in the matching of X_1, they get different colors. Let Y_2^1 and Y_2^2 be the invariant processes consisting of the sites of X_2 with the two different colors respectively. Now match the sites of each these two processes among themselves according to a matching satisfying Lemma 2.2. Note that $P(Y_2^k(v) = 1) = f(2)/2$.

Keep on doing this recursively; after having defined and matched the processes $Y_k^1, \ldots, Y_k^k$, $k = 1, \ldots, n-1$, color the sites of X_n with a uniformly chosen n-coloring, such that any two sites matched earlier get different colors. Let $Y_n^1, \ldots, Y_n^n$ consist of the sites of X_n with the n different colors respectively, and match the sites of each of these among themselves according to a matching satisfying Lemma 2.2. This leads to a graph with the desired degrees and the coloring procedure makes sure it becomes simple. Note that $P(Y_n^k(v) = 1) = f(n)/n$.

If $D_v \geq n$, then put w_v^n for the vertex v gets matched to in the matching of X_n, and if $D_v < n$, then put $w_v^n = v$. Then, by Lemma 2.2,

$$\mathbb{E}T = \sum_{n=1}^{\infty} \mathbb{E}[\mathbf{dist}(o, w_o^n)] < 10K \sum_{n=1}^{\infty} f(n)\mathcal{R}\Big(\frac{n}{f(n)}\Big). \qquad (1)$$

So far, we have not used that G grows rapidly. This means that (1) can (and will) be used also when G grows polynomially. However, since G does grow rapidly here, there exists some $c < \infty$ such that $\mathcal{R}(n) \le (\log n)^c$. Now either $1/f(n)$ grows at super-polynomial rate, in which case the right hand side sum of (1) trivially converges, or $1/f(n)$ grows polynomially, in which case $\mathcal{R}(n/f(n))$ and $\mathcal{R}(n)$ only differ by a constant and convergence follows from the hypothesis that $\mathbb{E}[D\mathcal{R}(D)] < \infty$.

3.2 Graphs of polynomial growth

First we give an example of an invariant degree process on $G = \mathbb{Z}$ for which $\mathbb{E}[D^k] < \infty$ for a given $k < \infty$ and for which it is not possible to have $\mathbb{E}T$ finite. Let X be a random variable with support on the positive odd integers, such that $\mathbb{E}[X^c] = \infty$ for every $c > 0$. Define the degree process $\{D_v\}_{v \in V}$ in the following way. First make a realization of X. Then, given $X = x$, make an invariant partition of $\mathbb{Z}$ into intervals of length x. For each interval, give the center vertex, which is well defined since x is odd, degree $x^{1/k}$ (for simplicity assume that this number is an integer) and the other vertices degree 0. We have

$$\mathbb{E}[D^k|X] = \frac{1}{X}(X^{1/k})^k = 1$$

and hence $\mathbb{E}D = 1 < \infty$. However, given $X = x$ and $D > 0$, we must necessarily have

$$T \ge 2x \sum_{n=1}^{\lfloor x^{1/k}/2 \rfloor} n > \frac{1}{5} x^{1+2/k}.$$

Thus

$$\mathbb{E}T \ge \frac{1}{5} P(D > 0) \mathbb{E}[X^{1+2/k} | D > 0].$$

However

$$P(X = n | D > 0) = \frac{1}{P(D > 0)} n^{-1} P(X = n).$$

Hence

$$\mathbb{E}T \ge \frac{1}{5} \sum_{n=1}^{\infty} n^{2/k} P(X = n) = \mathbb{E}[X^{2/k}] = \infty.$$

This proves part (i) of Theorem 3.2.

Now we modify the given example to $G = \mathbb{Z}^d$ in order to prove the second part of (ii). Let X have the same distribution as for $d = 1$. Make a realization of

X and, given $X = x$, make in invariant partition of $\mathbb{Z}^d$ into d-dimensional cubes with side length x. For each cube, give the center vertex degree $x^{d(d-1)/(d+1)+\beta}$ for a $\beta > 0$ less than, say, $\alpha/10$. Then

$$\mathbb{E}[D^{(d+1)/(d-1)-\alpha}|X] = \frac{1}{X^d} X^{d-d(d-1)\alpha/(d+1)+(d+1)\beta/(d-1)-\alpha\beta} \leq 1$$

so that

$$\mathbb{E}[D^{(d+1)/(d-1)-\alpha}] \leq 1.$$

However, given $X = x$ and $D > 0$, we must have

$$T \geq x \sum_{n=1}^{\mathcal{R}(x^{d(d-1)/(d+1)+\beta})} n\mathcal{A}(n) \geq Cx^{d+\gamma}$$

for constants $C, \gamma > 0$. Hence

$$\mathbb{E}T \geq CP(D > 0)\mathbb{E}[X^{d+\gamma}|D > 0].$$

Since

$$P(X = n|D > 0) = \frac{1}{P(D > 0)} n^{-d} P(X = n)$$

we get

$$\mathbb{E}T \geq C \sum_{n=1}^{\infty} n^{\gamma} P(X = n) = C\mathbb{E}[X^{\gamma}] = \infty.$$

Next we prove the first part of Theorem 3.2(ii). Assume thus, that G has growth of order n^d, $d \geq 2$. Then $\mathcal{R}(x)$ is of order $x^{1/d}$. Assume that

$$\mathbb{E}[D^{(d+1)/(d-1)+\alpha}] < \infty.$$

This is equivalent to

$$\sum_{n=1}^{\infty} n^{2/(d-1)+\alpha} f(n) < \infty.$$

By (1), finiteness of $\mathbb{E}T$ follows if it can be shown that $\sum_{n=1}^{\infty} f(n)\mathcal{R}(n/f(n))$ is finite, i.e. if

$$\sum_{n=1}^{\infty} n^{1/d} f(n)^{(d-1)/d} < \infty.$$

However, by assumption and Hölder'ss inequality,

$$\sum_{n=1}^{\infty} n^{1/d} f(n)^{(d-1)/d} = \sum_{n=1}^{\infty} n^{-(1+(d-1)\alpha)/d} (n^{2/(d-1)+\alpha} f(n))^{(d-1)/d}$$

$$\leq \Big(\sum_{n=1}^{\infty} n^{-(1+(d-1)\alpha)}\Big)^{1/d} \Big(\sum_{n=1}^{\infty} n^{2/(d-1)+\alpha} f(n)\Big)^{(d-1)/d} < \infty.$$

This finishes the proof of Theorem 3.2.

Remark. One can easily show that if one settles for a random multigraph model without loops instead of a simple graph model, then instead $\mathbb{E}[D^{d/(d-1)+\alpha}] < \infty$ suffices for $\mathbb{E}T < \infty$ being possible. This is done by skipping the coloring part of the proof of Lemma 2.2. One can also easily modify the above arguments to show that $\mathbb{E}[D^{d/(d-1)-\alpha}] < \infty$ is necessary.

It remains to prove Theorem 3.1 for G's of polynomial growth. Observe that the polynomial growth entails that there is a constant $C < \infty$ such that for all n, $1/C \leq \mathcal{V}(2n)/\mathcal{V}(n) \leq C$. Fix the integer t so large that $\mathbb{E}H_v < 4^{-d}/C$, where $H_v = H_v^t := D_v \mathbf{1}_{\{D_v > t\}}$. Say that a vertex is *high* if $H_v > 0$ and *low* if not. We will first connect each high vertex v to D_v low vertices, in such a way that no low vertex gets connected to more than one high vertex, and so that that the mean total edge length per vertex is finite. After this has been done, then the "remaining" degrees, i.e. 0 if v is high, $D_v - 1$ if $D_v \geq 1$, v is low and was connected to a high vertex, and D_v otherwise, are bounded by t and form an invariant process. Thus an application of the arguments of Section 3.1 finishes the proof.

The procedure for connecting the high vertices to low vertices is the following stepwise algorithm, which is a discrete version of the "stable marriage of Poisson and Lebesgue" of Hoffman, Holroyd and Peres [7]. It will be obvious from its definition that the algorithm is automorphism invariant. For convenience we start by "disturbing" the positions of the vertices a little, i.e. regarding $v \in V$ as having a position at distance M_v from its actual position, along a uniformly chosen incident edge. The M_v's are iid random variables in $[0, 0.1]$. This is simply in order to make sure that every vertex has a unique nearest vertex, a unique second nearest vertex etc.

Now, in Step 1 of the algorithm, every high vertex claims its H_v nearest low neighbors. Then every low vertex that has been claimed by at least one high vertex, is connected to the nearest one of the high vertices that has claimed it. After this each high vertex will be connected to, say, $H_v(1)$ low vertices, where $H_v(1) \leq H_v$ with equality iff v was connected to every vertex it claimed.

In Step 2, every high vertex v for which $H_v(1) < H_v$, claims its $H_v - H_v(1)$ nearest low neighbors which are not yet connected to some high vertex. Then every now claimed low vertex is connected to the nearest one of the high vertices claiming it. After this every high vertex is connected to, say, $H_v(2)$ low vertices, where $H_v(2) \leq H_v$.

In Step 3, every high vertex v for which $H_v(2) < H_v$, claims its $H_v - H_v(2)$ nearest low neighbors which are not yet connected to some high vertex etc.

This goes on recursively. We must now show that this algorithm a.s. leads to a well defined edge configuration. For simplicity, we will assume that $P(D = 0) = 0$, only noting that removing this condition is straightforward.

On the event that v is high, put R_v for the distance between v and the far-most vertex v gets connected to. Put $R = R_0$ and $H = H_0$. Since no vertex outside $B[o, 2n]$ can influence whether or not a high origin gets connected to a given low vertex in $B[o, n]$, we have

$$P(R \geq n) \leq P\Big(\sum_{v \in B[o,2n]} H_v \geq \frac{1}{2}\mathcal{V}(n)\Big) \leq P\Big(\frac{1}{\mathcal{V}(2n)} \sum_{k=1}^{\mathcal{V}(2n)} X_k \geq \frac{1}{C2^{d+1}}\Big),$$

where $X_1, X_2, \ldots$ are iid and distributed as H. (The $1/2$-factor in the second expression is needed since high vertices are connected only to low vertices. However if the sum of the H_v's in a given region is less than half the number of vertices there, then also less than half the vertices are high.)

Hence

$$\mathbb{E}R \leq \sum_{n=1}^{\infty} P\Big(\frac{1}{\mathcal{V}(2n)} \sum_{k=1}^{\mathcal{V}(2n)} X_k \geq \frac{1}{C2^{d+1}}\Big).$$

The expression on the right hand side is finite if and only if

$$\sum_{n=1}^{\infty} n^{-(d-1)/d} P\Big(\frac{1}{n} \sum_{k=1}^{n} X_k \geq \frac{1}{C2^{d+1}}\Big)$$

is finite; this follows from an integral approximation and a suitable substitution. Now recall that $\mathbb{E}X_k = 4^{-d}/C < 2^{-(d+1)}/C$. Hence, by a standard result of Baum and Katz [2] on convergence rates in the law of large numbers, finiteness of the last expression is equivalent to finiteness of $\mathbb{E}[X_k^{(d+1)/d}] = \mathbb{E}[H^{(d+1)/d}]$. This establishes that our model is well defined. Since, obviously, $T \leq HR$, it also almost proves the main result that $\mathbb{E}T < \infty$. The problem is that H and R are not independent. However, given that HR is large, R turns out to essentially only depend on H_v, $v \neq o$. We now formalize this. We have

$$P(R \geq n|H) \leq P\Big(\sum_{v \in B[o,2n]} H_v \geq \frac{1}{2}\mathcal{V}(n)\Big|H\Big)$$

$$\leq P\Big(H + \sum_{k=1}^{\mathcal{V}(2n)} X_k \geq \frac{1}{2}\mathcal{V}(n)\Big|H\Big),$$

where the X_k's are iid, distributed as H and independent of H.

The last expression is of course bounded above by 1 and when $n > 4\mathcal{R}(H)$, it is bounded by

$$P\Big(\mathcal{V}(\frac{n}{4}) + \sum_{k=1}^{\mathcal{V}(2n)} X_k \geq \frac{1}{2}\mathcal{V}(n)\Big) \leq P\Big(\sum_{k=1}^{\mathcal{V}(2n)} X_k \geq \frac{1}{4}\mathcal{V}(n)\Big).$$

We get

$$\mathbb{E}[R|H] = \sum_{n=1}^{\infty} P(R \geq n|H) \leq 4\mathcal{R}(H) + \sum_{n=4\mathcal{R}(H)+1}^{\infty} P\Big(\sum_{k=1}^{\mathcal{V}(2n)} X_k \geq \frac{1}{4}\mathcal{V}(n)\Big)$$

$$\leq 4\mathcal{R}(H) + \sum_{n=1}^{\infty} P\Big(\frac{1}{\mathcal{V}(2n)}\sum_{k=1}^{\mathcal{V}(2n)} X_k \geq \frac{1}{C2^{d+2}}\Big).$$

The sum on the right hand side is independent of H and finite by the same Baum-Katz argument as above; denote it S. Hence

$$\mathbb{E}T \leq \mathbb{E}[HR] = \mathbb{E}[H\,\mathbb{E}[R|H]] = \mathbb{E}[H(4\mathcal{R}(H) + S)]$$

$$\leq 4\mathbb{E}[H\mathcal{R}(H)] + S\,\mathbb{E}H < \infty.$$

We are done.

Concluding remarks

- **Other properties.** It would be interesting to find out what other properties our random graph models have, apart from finiteness of $\mathbb{E}T$. For example, could the volume growth of the random graph, H say, on the geography G, be fundamentally different from that of G? Of course, if G grows super-polynomially, then H will in most cases grow at a higher rate, but when $\mathcal{V}_G(n)$ grows like n^d, then the situation is not so clear. Letting $\mu := \mathbb{E}T$, it follows from Markov's inequality that for any number $C \geq 1$ and any n,

$$P\Big(\mathcal{V}_H(n) \geq (C\mu)^d \mathcal{V}_G(n)\Big) \leq \frac{1}{C}.$$

However, it could still for example be the case that $\limsup_n \mathcal{V}_H(n)/\mathcal{V}_G(n)$ is infinite a.s.

Another interesting question is how simple random walk on H behaves. For G of exponential growth, could one have some unexpected behavior of SRW on H? For example, could a heavy tail in the distribution of D cause the drift to be very low or even 0?

- **Finite-graph versions.** The models in this paper rely heavily on the fact that G is infinite. If G instead was finite, then of course the degrees of H could not be completely independent, since that would with possible probability lead to an impossible degree sequence with respect to the number of vertices of G. (E.g. the sum of the degrees could be odd.)

  It is not clear what a finite-graph version would mean. One suggestion is that if G has n vertices, one asks for a model which with probability $1-o(1)$ for iid even degrees produces a random graph, for which $\mathbb{E}T$ is bounded in n. Then of course one needs to let G grow in some natural way. One way to do this is to let G be infinite as before and then let $G_n = B[o, n]$. Let H be a random graph defined on G as before, and let H_n be the induced subgraph of H on the vertices of G_n. Then some of the vertices of H_n will have fewer edges than prescribed. Hopefully this could then be compensated before by connecting these vertices to each other or to other vertices at the cost of deleting other edges. In case G is amenable (i.e. has zero isoperimetric constant) it is easily seen by a Markov inequality argument, that this can indeed be done in such a way that the desired properties are satisfied and so that the edges of o are in the resulting random graph with probability $1 - o(1)$ exactly as in H. In the nonamenable case, however, the problem seems to be more involved.

  Another, less general, way to let G grow naturally, is to let G be a d-dimensional torus. In this case, it is likely, but not obvious, that a modification of our model on the d-dimensional lattice could lead to an invariant model of the desired kind.

References

1. L. Bartholdi, The growth of Grigorchuk's torsion group, *Internat. Math. Res. Notices* no. 20 (1998), 1049-1054.

2. E. Baum and M. Katz, Convergence rates in the law of large numbers, *Trans. Amer. Math. Soc.* **120** (1965), 108-123.

3. I. Benjamini, R. Lyons, Y. Peres and O. Schramm, Group-invariant percolation on graphs, *Geom. Funct. Analysis* **9** (1999), 29-66.

4. M. Deijfen and J. Jonasson, Stationary random graphs on $\mathbb{Z}$ with prescribed iid degrees and finite mean connections, *Electronic Comm. Probab.* **11** (2006), 336-346.

5. M. Deijfen and R. Meester, Generating stationary random graphs on $\mathbb{Z}$ with prescribed iid degrees, *Adv. Appl. Probab.* **38** (2006), 287-298.

6. R. Grigorchuk, On the Milnor problem of Group growth, *Soviet Math. Dokl.* **28** (1983), 23-26.

7. C. Hoffman, A. Holroyd and Y. Peres, A stable marriage of Poisson and Lebesgue, *Ann. Probab.* **34** (2006), 1241-1272.

8. W. Imrich and N. Seifter, A survey on graphs with polynomial growth, *Discrete Math.* **95** (1991), 101-117.

9. V. Kaimanovich and A. Vershik, Random walks on discrete groups: boundary and entropy, *Ann. Probab.* **11** (1983), 457-490.