Random walks on quasirandom graphs

Ben Barber

Department of Pure Mathematics and Mathematical Statistics
University of Cambridge
United Kingdom

b.a.barber@dpmms.cam.ac.uk

Eoin Long

School of Mathematical Sciences
Queen Mary University of London
United Kingdom

e.p.long@qmul.ac.uk

Submitted: April 7, 2013; Accepted: Nov 17, 2013; Published: Nov 29, 2013
Mathematics Subject Classifications: 05C81

Abstract

Let G be a quasirandom graph on n vertices, and let W be a random walk on G of length αn^2. Must the set of edges traversed by W form a quasirandom graph? This question was asked by Böttcher, Hladký, Piguet and Taraz. Our aim in this paper is to give a positive answer to this question. We also prove a similar result for random embeddings of trees.

1 Introduction

Given a graph G and sets $A, B \subseteq V(G)$, let

$$e_G(A,B) = |\{(a,b) \in A \times B : ab \in E(G)\}|$$

be the number of edges from A to B. A graph G with n vertices and $\rho\binom{n}{2}$ edges is *ϵ-quasirandom* if

$$|e_G(A,B) - \rho|A||B|| < \epsilon|A||B|$$

for all sets A, $B \subseteq V(G)$ with $|A|$, $|B| \geqslant \epsilon n$. Thus a quasirandom graph resembles a random graph with the same density, provided we do not look too closely. Quasirandom graphs were introduced by Thomason [10] and have come to play a central role in probabilistic and extremal graph theory. The reader is referred to the excellent survey article by Krivelevich and Sudakov [5] for further details.

The random graph $G_{n,p}$, in which edges appear independently with probability p, is quasirandom with high probability. More generally, given a quasirandom graph G we can, with high probability, obtain a new quasirandom graph $G_{edge}(p)$ by retaining edges of G with some fixed probability p. (The random graph $G_{n,p}$ can be thought of as the result of applying this process to the complete graph K_n.)

Another natural way to choose a random set of edges from an n-vertex graph G is given by the following process. A random walk W on G is a sequence of vertices $W_0, W_1, \dots, W_l$ where W_0 is chosen from some initial distribution and W_{i+1} is selected uniformly from the neighbours of W_i, with all choices made independently. We will be interested in the case when W traverses some constant fraction of the edges of G, so we take the length l of W to be αn^2 for some constant $\alpha > 0$. To avoid confusion with the random walk W, let $G_{walk}(\alpha)$ denote the random subgraph of G consisting of those edges traversed by W. The main question that we are interested in here is the following: given a quasirandom graph G, is it true (as with $G_{edge}(p)$) that $G_{walk}(\alpha)$ is also quasirandom with high probability?

First note that this is true when this process is applied to the complete graph $G = K_n$. Indeed, $G_{walk}(\alpha)$ is very close to $G_{n,p}$ for some p. This is because the sequence $W_0, W_1, \dots$ is very nearly a sequence of independent random vertices of G ('very nearly' because consecutive terms of the sequence are forced to be distinct). Then $W_0W_1, W_2W_3, \dots$ and $W_1W_2, W_3W_4, \dots$ are very nearly two sequences of independent random edges of G, so $G_{walk}(\alpha)$ is very close to a random subgraph of G.

The following heuristic suggests that $G_{walk}(\alpha)$ should also be quasirandom for a general quasirandom graph G.

(1) The graph G is approximately regular, of degree $\rho(n-1)$, so the equilibrium distribution of W is approximately uniform. The random walk W 'mixes rapidly', so most sequence terms have distributions close to the equilibrium distribution, and W visits each vertex around the same number of times: approximately αn times.

(2) For each vertex v of G, the random walk leaves v αn times, so picks up a random set (chosen with replacement) of αn of the ρn edges at v. Taking the union of these sets of edges gives a random subgraph of G.

While this simple plan seems quite plausible, there are two reasons why it is not easy to implement, both connected to (1) above. The first is related to what exactly it means to say that W 'mixes rapidly'. Since quasirandomness does not say anything about small parts of the graph, G might have small configurations of low degree vertices that can trap the random walk for long periods of time. The second is that, even if we know the distribution of each W_i, these random variables are not independent for different i.

These difficulties can, however, be overcome, and an argument of the above form can be used to show that the subgraph of G spanned by W will be quasirandom with high probability.

Theorem 1. *Given $\alpha, \rho, \eta > 0$ there exists $\epsilon > 0$ such that the following holds. Let G be an n-vertex ϵ-quasirandom graph with $\rho\binom{n}{2}$ edges, and let W be a random walk on G of length αn^2 starting at any vertex W_0 of G with degree in $[(\rho-\epsilon)n, (\rho+\epsilon)n]$. Then, with*

probability $1 - o_\epsilon(1)$*, the graph* $G_{walk}(\alpha)$ *is* η*-quasirandom with* $(1 - e^{-2\alpha/\rho} + o_\epsilon(1))\rho\binom{n}{2}$ *edges.*

Here $o_\epsilon(1)$ means a quantity which is less than $f(\epsilon)$ for n sufficiently large, for some $f(\epsilon)$ tending to zero with ϵ. This dependence on ϵ, and not just n, is necessary to deal with the 'bad small subgraph' problem described above.

If we also have a lower bound on the minimum degree of G then we can say much more: if we start with a graph that is ϵ-quasirandom, then, with high probability, $G_{walk}(\alpha)$ will also be ϵ-quasirandom.

Theorem 2. *Let* $\alpha, \epsilon, \rho > 0$ *and let* $\gamma = C\epsilon^{1/4}$ *for some absolute constant* $C > 0$*. Let* G *be an* n*-vertex* ϵ*-quasirandom graph with* $\rho\binom{n}{2}$ *edges and minimum degree at least* γn*, and let* W *be a random walk on* G *of length* αn^2*. Then, with probability* $1 - o(1)$*, the graph* $G_{walk}(\alpha)$ *is* ϵ*-quasirandom with* $(1 - e^{-2\alpha/\rho} + o(1))\rho\binom{n}{2}$ *edges.*

In Section 2 we define an explicit model for our random walks and show that step (2) above works straightforwardly in this setting. In Section 3 we carry out (1) for the case where we have a bound on the minimum degree of G. This proves Theorem 2, and illustrates why we get the weaker conclusion in Theorem 1. In Section 4 we use a more elaborate argument to perform (1) in the general case, proving Theorem 1.

The problems considered in this paper were suggested by Böttcher, Hladký, Piguet and Taraz [3] after they encountered similar problems in connection with their work on tree packing. Suppose that we are trying to pack many trees into a copy of K_n. One approach is to embed some of the trees randomly. If we succeed in packing a small number of trees, then it would be good to know that the subgraph consisting of unused edges has nice enough properties that we can iterate the argument and therefore pack a much larger number of trees. If H is a subgraph of G, and both graphs are quasirandom, then $G - H$ is also quasirandom. So it would be useful to have a result like Theorem 1, but for random images of trees rather than paths. We consider such a generalisation in Section 5.

Since we will only prove asymptotic results we make a number of simplifying assumptions. We assume that ϵ is sufficiently small compared to the other parameters, and are only interested in statements for n sufficiently large. We omit notation indicating the taking of integer parts, and ignore questions of divisibility when breaking walks into pieces of a given size.

2 The list model

We now define a third model of a random subgraph to act as a staging post between $G_{walk}(\alpha)$ and $G_{edge}(p)$. The subgraph $G_{list}(\nu)$ of G is obtained by selecting $\nu d(v)$ edges at each vertex v of G to be retained, with all choices made independently and with replacement. We give a rather elaborate formal definition in order to introduce some ideas which will be useful later.

For each $v \in V(G)$, let L_v be an infinite list of uniform selections from the neighbourhood of v, with all choices made independently. The entry u on the list L_v corresponds

naturally to the edge uv of G, and we define

$$G_{list}(\nu) = \bigcup_{v \in V(G)} \{uv : u \text{ appears in the first } \nu d(v) \text{ entries of } L_v\}.$$

In this section we will show that $G_{list}(\nu)$ is very close to $G_{edge}(p)$ for some p, in the sense that large subgraphs have similar densities in each model. It will then follow that $G_{list}(\nu)$ is quasirandom with high probability.

We first calculate the expected density of $G_{list}(\nu)$ in G. The reader is encouraged to focus on the case where we have a bound on the minimum degree.

Lemma 3. *Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges, and let $A, B \subseteq V(G)$ with $|A|, |B| \geqslant \epsilon^{0.99}n$. Then*

$$\mathbb{E}\left(e_{G_{list}(\nu)}(A, B)\right) = (1 - e^{-2\nu} + o_\epsilon(1))e_G(A, B).$$

Moreover, if the minimum degree of G is at least γn, then, for all $A, B \subseteq V(G)$,

$$\mathbb{E}\left(e_{G_{list}(\nu)}(A, B)\right) = (1 - e^{-2\nu} + o(1))e_G(A, B).$$

The exact lower bound on the sizes of A and B in the general case is unimportant; any value asymptotically larger than ϵ would work equally well.

Proof. Write $S = \{v \in V(G) : d(v) \geqslant (\rho - \epsilon)n\}$. Then

$$|e_G(V(G), S^c) - \rho n|S^c|| > \epsilon n|S^c|,$$

so, by ϵ-quasirandomness, $|S^c| < \epsilon n$.

The edge uv of G appears in $G_{list}(\nu)$ if and only if u appears in the first $\nu d(v)$ entries of L_v, or v appears in the first $\nu d(u)$ elements of L_u. If $u, v \in S$, then the probability of this occurring is

$$1 - (1 - 1/d(v))^{\nu d(v)}(1 - 1/d(u))^{\nu d(u)} = 1 - e^{-2\nu} + o(1),$$

since $d(v), d(u) \geqslant (\rho - \epsilon)n$. So

$$\begin{aligned}\mathbb{E}\left(e_{G_{list}(\nu)}(A, B)\right) &= (1 - e^{-2\nu} + o(1))e_G(A, B) + O(\epsilon n(|A| + |B|)) \\ &= (1 - e^{-2\nu} + o_\epsilon(1))e_G(A, B),\end{aligned}$$

since $e_G(A, B) \geqslant (\rho - \epsilon)|A||B|$ and $|A|, |B| \geqslant \epsilon^{0.99}n$.

If $d(v) \geqslant \gamma n$ for every $v \in V(G)$, then the probability of being retained is $1-e^{-2\nu}+o(1)$ for every edge of G, so

$$\mathbb{E}\left(e_{G_{list}(\nu)}(A, B)\right) = (1 - e^{-2\nu} + o(1))e_G(A, B). \qquad \square$$

To show that the number of edges retained in any subgraph is close to its expectation we use Talagrand's concentration inequality [9]. In its usual form Talagrand's inequality is asymmetric and bounds a random variable in terms of its median. We use the following symmetric version (see [7, Chapter 20]) that gives concentration of the random variable about its mean.[1]

Theorem 4 (Talagrand's inequality)**.** *Let $\Omega = \prod_{i=1}^{N} \Omega_i$ be a product of probability spaces with the product measure. Let X be a random variable on Ω such that*

(i) $|X(\omega) - X(\omega')| \leqslant c$ whenever ω and ω' differ on only a single coordinate for some constant $c > 0$;

(ii) whenever $X(\omega) \geqslant r$ there is a set $I \subseteq \{1, \dots, N\}$ with $|I| = r$ such that $X(\omega') \geqslant r$ for all $\omega' \in \Omega$ with $\omega'_i = \omega_i$ for all $i \in I$;

and suppose that the median of X is at least $100c^2$. Then, for $0 \leqslant s \leqslant \mathbb{E}(X)$,

$$\mathbb{P}\left(|X - \mathbb{E}(X)| \geqslant s + 90c\sqrt{\mathbb{E}(X)}\right) \leqslant 4e^{-s^2/8c^2\mathbb{E}(X)}.$$

Lemma 5. *Let G be an n-vertex ϵ-quasirandom graph, and fix $\nu > 0$. Then, with probability $1 - o(1)$,*

$$e_{G_{list}(\nu)}(A, B) = (1 - e^{-2\nu} + o_\epsilon(1))e_G(A, B),$$

for all $A, B \subseteq V(G)$ with $|A|, |B| \geqslant \epsilon^{0.99}n$. Moreover, if the minimum degree of G is at least γn, then the same result holds for $|A|, |B| \geqslant \epsilon n$, with $o_\epsilon(1)$ replaced by $o(1)$.

Proof. We apply Theorem 4 to the space $\Omega = \prod_{v \in V(G)} \prod_{i=1}^{\nu d(v)} N(v)$, where each neighbourhood has the uniform probability measure; we can view Ω as the space of choices for the first $\nu d(v)$ entries of each list L_v. For $A, B \subseteq V(G)$ with $|A|, |B| \geqslant \epsilon^{0.99}n$, let $X_{A,B} = e_{G_{list}(\nu)}(A, B)$. It is easy to see that $X_{A,B}$ satisfies the conditions of Talagrand's inequality. Indeed, (i) holds since changing a list entry can change $X_{A,B}$ by at most $c = 2$. Furthermore, (ii) holds since, if $X_{A,B} \geqslant s$, then there are s list entries witnessing this fact. Finally, the technical condition on the median holds for large n, since the probability that $X_{A,B}$ is less than any constant is $o(1)$. Therefore, by Theorem 4, for $180\sqrt{\mathbb{E}(X_{A,B})} \leqslant t \leqslant \mathbb{E}(X_{A,B})$ and large n we have

$$\mathbb{P}(|X_{A,B} - \mathbb{E}(X_{A,B})| \geqslant 2t) \leqslant 4e^{-t^2/32\mathbb{E}(X_{A,B})}.$$

By Lemma 3 we have $\mathbb{E}(X_{A,B}) = (1 - e^{-2\nu} + o_\epsilon(1))e_G(A, B)$. Since $e_G(A, B) \geqslant (\rho - \epsilon)\epsilon^{1.98}n^2$, taking $t = C'\sqrt{n\mathbb{E}(X_{A,B})}$ $(= o(\mathbb{E}(X_{A,B})))$ for large enough $C' > 0$ gives that

$$\mathbb{P}\left(\left|X_{A,B} - (1 - e^{-2\nu} + o_\epsilon(1))e_G(A, B)\right| \geqslant 2t\right) \leqslant 8^{-n}.$$

[1] We note that the proof in [7] is incorrect, as a simplified form of the failure probability valid for a small deviation t is used when t is large. However, when t is large, a different simplification can be used instead at the cost of (a) increasing the constants slightly and (b) imposing a very mild additional condition on the random variable X.

But there are at most 2^n choices for A and 2^n choices for B. Therefore, with probability at least $1-2^{-n}$, we have that $X_{A,B} = (1-e^{-2\nu}+o_\epsilon(1))e_G(A,B)$, for all pairs (A,B) with $|A|,|B| \geqslant \epsilon^{0.99}n$. The 'moreover' statement is proved identically, using the 'moreover' statement from Lemma 3. □

This is enough to ensure that $G_{list}(\nu)$ is quasirandom with high probability.

Theorem 6. *Let $\nu,\gamma > 0$ and let G be an n-vertex ϵ-quasirandom graph with $\rho\binom{n}{2}$ edges. Then, with probability $1-o(1)$, $G_{list}(\nu)$ is $o_\epsilon(1)$-quasirandom. Moreover, if the minimum degree of G is at least γn, then, with probability $1-o(1)$, $G_{list}(\nu)$ is ϵ-quasirandom.*

Proof. By Lemma 3, with probability $1-o(1)$,

$$|e_{G_{list}(\nu)}(A,B) - (1-e^{-2\nu})e_G(A,B)| = o_\epsilon(1)e_G(A,B),$$

for all A,B with $|A|,|B| \geqslant \epsilon^{0.99}n$. By the definition of quasirandomness, we also have that

$$|e_G(A,B) - \rho|A||B|| < \epsilon|A||B|.$$

So by the triangle inequality,

$$\begin{aligned} |e_{G_{list}(\nu)}(A,B) - (1-e^{-2\nu})\rho|A||B|| &< o_\epsilon(1)e_G(A,B) + (1-e^{-2\nu})\epsilon|A||B| \\ &\leqslant (o_\epsilon(1) + (1-e^{-2\nu})\epsilon)|A||B| \\ &= o_\epsilon(1)|A||B|. \end{aligned}$$

Hence $G_{list}(\nu)$ is δ-quasirandom with $\delta = \max(\epsilon^{0.99}, o_\epsilon(1))$, where the $o_\epsilon(1)$ is taken from the last line. Since $(1-e^{-2\nu})\epsilon < \epsilon$, the 'moreover' statement follows from the 'moreover' statement of Lemma 3. □

Having shown that $G_{list}(\nu)$ is quasirandom with high probability, it suffices to show that $G_{walk}(\alpha)$ is close to $G_{list}(\nu)$ for some ν. The construction of the random walk W requires, at each visit to a vertex v, a choice of a random neighbour of v. We obtain a coupling of $G_{walk}(\alpha)$ and $G_{list}(\nu)$ by, at the jth visit to v, taking this choice to be the jth entry of the list L_v. Then $G_{walk}(\alpha)$ and $G_{list}(\nu)$ both consist of the edges corresponding to some initial segments of the lists L_v, and it is enough to show that we can choose ν such that the lengths of these initial segments are similar: that is, that the number of times the random walk W visits each vertex of G is roughly proportional to its degree.

We give two arguments. The first, appearing in Section 3, applies when we have a good lower bound on the minimum degree of G. The second, appearing in Section 4, applies to a general quasirandom graph G, but necessarily gives a weaker result. We include the argument for the special case where G has large minimum degree for two reasons. First, it proves the stronger result that there is no loss of quasirandomness when we pass from G to $G_{walk}(\alpha)$, which could be useful for some applications. Second, it illustrates why the natural approach cannot work in general, justifying the use of a more technical argument in Section 4.

3 Bounded minimum degree

To begin this section we recall some useful facts. A random walk W on a graph G is a Markov chain with transition matrix P given by

$$P_{uv} = \begin{cases} 1/d(u) & \text{if } uv \in E(G); \\ 0 & \text{if } uv \notin E(G). \end{cases}$$

Thus P is a normalised version of the adjacency matrix A, where each row has been scaled by the degree of the corresponding vertex. The eigenvalues of P are all real; let these be $\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_n$ and write $\lambda = \max(|\lambda_2|, |\lambda_n|)$. The first eigenvalue λ_1 of P is always equal to 1 and has a corresponding eigenvector $\pi = (\pi_v)$ given by $\pi_v = \frac{d(v)}{2e(G)}$. This vector π is called the *stationary distribution* of the walk W. If G is disconnected, then $\lambda_2 = 1$; if G is bipartite, then $\lambda_n = -1$. If G is connected and non-bipartite, then $\lambda < 1$ and the walk W is irreducible and aperiodic. Later in this section, G will be ϵ-quasirandom and have minimum degree at least γn with $\gamma > \epsilon$. It will therefore be both connected (as the minimum degree condition forces every component to have size greater than ϵn, and components have no edges between them) and non-bipartite (else the larger class would have no edges to itself). In this case it is well-known (for example, see [6, Theorem 4.9]) that, for any initial distribution of W_0, the distribution of W_i converges to π as $i \to \infty$ (i.e. $\mathbb{P}(W_i = v) \to \pi_v$ as $i \to \infty$ for each v).

The following standard result, which can be read out of Jerrum and Sinclair [4], gives control on the *rate* of this convergence.

Lemma 7. *For any n-vertex graph G with minimum degree at least γn, and any initial distribution of W_0, we have*

$$\max_{v \in V(G)} |\mathbb{P}(W_i = v) - \pi_v| \leqslant c_\gamma \lambda^i,$$

for some c_γ depending on γ.

If G is a regular ϵ-quasirandom graph then λ is small on the scale of ϵ. (This is because the 'spectral gap' of a quasirandom graph is large [2], and P is a scalar multiple of A when G is regular.) For a general ϵ-quasirandom graph this need not be true: for example, if G contains a small connected component, then $\lambda = 1$ (the 1-eigenspace is spanned by the stationary distributions of each connected component of G). Similarly, λ can be very close to 1 if there is a small set of vertices that is only weakly connected to the rest of the graph. However, a lower bound on the minimum degree of G is enough to recover an upper bound on λ.

Lemma 8. *Let G be an n-vertex ϵ-quasirandom graph with $\rho\binom{n}{2}$ edges and minimum degree at least γn, where $\gamma \geqslant C\epsilon^{1/4}$ for some absolute constant $C > 0$. Then, for n sufficiently large, $\lambda \leqslant 1/2$.*

Before proving Lemma 8, we note the following simple observation about quasirandom graphs which we will use repeatedly.

Proposition 9. *Let G be an n-vertex ϵ-quasirandom graph with $\rho\binom{n}{2}$ edges, and let X be a set of vertices with $|X| \geqslant \epsilon n$. Let $Y = \{v \in V(G) : |e_G(v, X) - \rho|X|| \geqslant \epsilon|X|\}$. Then $|Y| < 2\epsilon n$.*

Proof. We have $Y = Y^+ \cup Y^-$ where

$$Y^+ = \{v \in V(G) : e_G(v, X) \geqslant \rho|X| + \epsilon|X|\},$$
$$Y^- = \{v \in V(G) : e_G(v, X) \leqslant \rho|X| - \epsilon|X|\}.$$

Clearly

$$\left|e_G(X, Y^+) - \rho\,|X|\,\left|Y^+\right|\right| \geqslant \epsilon\,|X|\,\left|Y^+\right|,$$
$$\left|e_G(X, Y^-) - \rho\,|X|\,\left|Y^-\right|\right| \geqslant \epsilon\,|X|\,\left|Y^-\right|.$$

But then, since G is ϵ-quasirandom and $|X| \geqslant \epsilon n$, we must have $|Y^+|, |Y^-| < \epsilon n$. □

In particular, taking $X = V(G)$ there are at least $(1 - 2\epsilon)n$ vertices v of G with $|d(v) - \rho n| \leqslant \epsilon n$. We will call such vertices *balanced*.

Proof of Lemma 8. The proof follows a well-known argument (see for example [2]). We first estimate the number of labelled copies of C_4 in G, and then evaluate the trace of P^4 in two different ways. Note that the implicit constants in our use of $O(\cdot)$ notation here are absolute (i.e. independent of ϵ, ρ, γ and λ).

The number of labelled copies of C_4 in G is

$$C_4(G) = 2 \sum_{u \in V(G)} \sum_{v \in V(G)\setminus\{u\}} \binom{|N(u) \cap N(v)|}{2}.$$

If u is balanced, then, by Proposition 9, the number of v with $|N(u) \cap N(v)|$ far from $\rho^2 n$ is small, so

$$\begin{aligned} C_4(G) &= 2 \cdot (1 + O(\epsilon))n \cdot (1 + O(\epsilon))n \cdot \binom{(\rho + O(\epsilon))^2 n}{2} + O(\epsilon)n^2\binom{n}{2} \\ &= (\rho + O(\epsilon))^4 n^4 + O(\epsilon)n^4 \\ &= \big(1 + O\big(\epsilon/\rho^4\big)\big)\,\rho^4 n^4, \end{aligned}$$

where the main term here accounts for balanced vertices u and v with close to $\rho^2 n$ common neighbours, and the error term bounds the contribution to the sum from each other pair by $\binom{n}{2}$.

Now the trace of P^4 is a weighted sum of the closed walks of length 4 in G, where the weight of the closed walk $uvwx$ is $1/(d(u)d(v)d(w)d(x))$. Thus

$$\begin{aligned} \sum_{v \in V(G)} (P^4)_{vv} &= \frac{(1 + O(\epsilon/\rho^4))\rho^4 n^4}{((\rho + O(\epsilon))n)^4} + \frac{O(\epsilon)n^4}{(\gamma n)^4} + \frac{O(n^3)}{(\gamma n)^4} \\ &= 1 + O\big(\epsilon/\rho^4\big) + O\big(\epsilon/\gamma^4\big) + O\big(1/(\gamma^4 n)\big), \end{aligned}$$

where the main term counts the contribution from 4-cycles containing only balanced vertices and the error terms account for the contributions from 4-cycles with at least one unbalanced vertex and from closed walks of length 4 which are not 4-cycles respectively. (The lower bound on the minimum degree of G gives an upper bound of $1/(\gamma n)^4$ for the weight of any one walk.) But we also have

$$\sum_{v\in V(G)} (P^4)_{vv} = \sum_{i=1}^{n} \lambda_i^4 = 1 + \sum_{i=2}^{n} \lambda_i^4,$$

from which it follows that

$$\lambda^4 \leqslant \sum_{i=2}^{n} \lambda_i^4 = O\big(\epsilon/\rho^4\big) + O\big(\epsilon/\gamma^4\big) + O\big(1/(\gamma^4 n)\big) \leqslant 1/16,$$

for $\rho \geqslant \gamma \geqslant C\epsilon^{1/4}$ and n sufficiently large. □

For the next lemma we will need to approximate one probability measure by another on the same space. Given a finite probability space Ω, the *total variation distance* between two probability measures μ_1 and μ_2 is defined by

$$d_{TV}(\mu_1, \mu_2) = \frac{1}{2} \sum_{\omega\in\Omega} |\mu_1(\omega) - \mu_2(\omega)|.$$

This is the amount of probability mass that would have to be moved to turn one distribution into the other.

Combining Lemma 7 with Lemma 8, it is easy to see that the total variation distance between W_t and a vertex sampled from the stationary distribution is small when t is moderately large. In fact, we get much more.

Let $L = (\log n)^2$, and let $K = \alpha n^2/L$. Given $i < L$, let $W^{(i)}$ denote the subsequence of W obtained by starting from W_i and taking L steps at a time: that is, $W^{(i)} = (W_1^{(i)}, \dots, W_K^{(i)})$ where $W_j^{(i)} = W_{i+(j-1)L}$ for all $j \leqslant K$. For each $v \in V(G)$, let $X_v^{(i)}$ be the random variable which counts the number of times $W^{(i)}$ visits v. Our next lemma shows that with high probability $X_v^{(i)}$ is close to its mean.

Lemma 10. *Let G be a graph satisfying the conditions of Lemma 8 and let $v \in V(G)$. Then we have*

$$\mathbb{P}\left(\big|X_v^{(i)} - K\pi_v\big| \geqslant \sqrt{\frac{8\log n}{K\pi_v}} K\pi_v\right) = O\big(n^{-3}\big).$$

Proof. Let $\mu = \pi^K$ be the K-fold product measure of π on $V(G)^K$; that is, $\mu(w) = \prod_{i=1}^{K} \pi_{w_i}$ for $w \in V(G)^K$. By Lemma 7 (absorbing the constant c_γ into the $O(\cdot)$ notation) and Lemma 8, we have

$$\begin{aligned}
\mathbb{P}\left(W^{(i)} = w\right) &= \mathbb{P}\big(W_1^{(i)} = w_1\big)\mathbb{P}\big(W_2^{(i)} = w_2 | W_1^{(i)} = w_1\big) \cdots \mathbb{P}\big(W_K^{(i)} = w_K | W_{K-1}^{(i)} = w_{K-1}\big) \\
&= \left(\pi_{w_1} + O\left(2^{-(\log n)^2}\right)\right)\left(\pi_{w_2} + O\left(2^{-(\log n)^2}\right)\right) \cdots \left(\pi_{w_K} + O\left(2^{-(\log n)^2}\right)\right) \\
&= \big(\pi_{w_1} + O\big(n^{-6}\big)\big)\big(\pi_{w_2} + O\big(n^{-6}\big)\big) \cdots \big(\pi_{w_K} + O\big(n^{-6}\big)\big) \\
&= \big(1 + O\big(n^{-3}\big)\big)\,\mu(w),
\end{aligned}$$

since $\frac{\gamma}{\rho n} \leqslant \pi_v \leqslant \frac{1}{\rho n}$ for all v and $K = O(n^2)$. Summing over all w gives that

$$d_{TV}(\mathbb{P}, \mu) = O\big(n^{-3}\big),$$

where $\mathbb{P}$ is the measure on $V(G)^K$ induced by $W^{(i)}$. Now let

$$A = \left\{ w \in V(G)^K : |X_v^{(i)}(w) - K\pi_v| \geqslant \sqrt{8 \log n K \pi_v} \right\}.$$

By Chernoff's inequality (see [1, A.1.11 and A.1.13]),

$$\mu(A) \leqslant 2e^{-(4+o(1))\log n} = O\big(n^{-3}\big).$$

Since $\mathbb{P}(A) \leqslant \mu(A) + d_{TV}(\mathbb{P}, \mu)$, the result follows. □

Now let $X_v = \sum_{i=0}^{L-1} X_v^{(i)}$ be the number of visits W makes to vertex v. Observing that $LK\pi_v = \alpha n^2 \pi_v = (1 + \frac{1}{n-1})\frac{\alpha}{\rho} d(v)$, we obtain the following corollary by summing over i and v.

Corollary 11. *Let $\alpha, \epsilon, \rho, \gamma > 0$ with $\rho, \gamma \geqslant C\epsilon^{1/4}$ for some absolute constant $C > 0$. Let G be an n-vertex ϵ-quasirandom graph with minimum degree at least γn, and let W be a random walk on G of length αn^2. Then*

$$\mathbb{P}\left(\left| X_v - \frac{\alpha}{\rho} d(v) \right| \geqslant \sqrt{\frac{8 \log n}{K\pi_v}} \frac{\alpha}{\rho} d(v) \text{ for some } v \right) = O\big(n^{-1}\big). \qquad \square$$

Since $K\pi_v = \Omega(n/(\log n)^2)$, with high probability the number of visits W makes to each $v \in V(G)$ is $\left(\frac{\alpha}{\rho} + o(1)\right) d(v)$. We can now complete the proof of Theorem 2.

Proof of Theorem 2. By Corollary 11, we have that, with probability $1 - o(1)$,

$$G_{list}(\alpha/\rho - o(1)) \subseteq G_{walk}(\alpha) \subseteq G_{list}(\alpha/\rho + o(1)).$$

From the proof of Theorem 6, we have that, with probability $1 - o(1)$,

$$|e_{G_{walk}(\alpha)}(A, B) - (1 - e^{-2\frac{\alpha}{\rho}})\rho|A||B|| < (1 - e^{-2\frac{\alpha}{\rho}} + o(1))\epsilon|A||B|,$$

for all $A, B \subseteq V(G)$ with $|A|, |B| \geqslant \epsilon n$. Since $1 - e^{-2\frac{\alpha}{\rho}} < 1$, $G_{walk}(\alpha)$ is ϵ-quasirandom with probability $1 - o(1)$. □

4 General case

We now move to the case of a general ϵ-quasirandom graph G with edge density ρ. Such G must always contain a connected component of order at least $(1-\epsilon)n$ (as otherwise we can find two sets of size at least ϵn with no edges between them), so by restricting our walk to this component we can assume that G is connected.

The extra difficulty in the general case is that there might be small sets of vertices that are only weakly connected to the rest of the graph in which the random walk can get stuck. For example, let G be a graph consisting of a small clique of order $\epsilon^2 n/2$ joined to a large clique of order $(1-\epsilon^2/2)n$ by a single edge. Then G is ϵ-quasirandom, but it is not even true that the number of edges in $G_{walk}(\alpha)$ is concentrated near some value. Indeed, if we start our random walk in the large clique then with positive probability (depending on ϵ but not on n) W will lie entirely within the large clique, but there is also a positive probability (depending on ϵ but not on n) that W will cross to the small clique in the first ϵn^2 steps and remain there. So for general quasirandom graphs we cannot hope for as strong a result as Theorem 2, and our assertions about high probability will necessarily depend on ϵ as well as n. In this section we use 'with high probability' to mean 'with probability $1-o_\epsilon(1)$', with $o_\epsilon(1)$ small (depending on ϵ) for large n as defined in Section 1.

Our task in this section is to find a weaker replacement for Corollary 11 in Section 3. Instead of saying that the random walk visits every vertex v around $\frac{\alpha}{\rho}d(v)$ times, we ask instead that the random walk visits *most* vertices of G around $\frac{\alpha}{\rho}d(v)$ times. Recall that we call a vertex v *balanced* if $|d(v)-\rho n| \leqslant \epsilon n$. We will show that, if W is a random walk of length αn^2 on G with W_0 balanced, then, with high probability, W hits most vertices of G about the right number of times. The results in Section 2 can then be used to prove Theorem 1 in the same way that Theorem 2 was deduced from Corollary 11.

Our first lemma gives a lower bound on the probability that a given step of a random walk W is in a set $S \subseteq V(G)$. Write $\mathbf{1}_X$ for the indicator function of a set X and $\mathbf{1}_v$ for the indicator function of the set $\{v\}$. Note that if the initial distribution for W_0 is π then $\mathbb{P}(W_i \in S) = \sum_{v\in S}\pi_v = \pi\cdot\mathbf{1}_S$ for any set $S \subseteq V(G)$ when $i \geqslant 0$. The next result shows that this is still almost true if W starts from a balanced vertex, S is large and $i \geqslant 2$.

Lemma 12. *Let G be a connected n-vertex ϵ-quasirandom graph with $\rho\binom{n}{2}$ edges, and let v be a balanced vertex. Let $S \subseteq V(G)$ with $|S| \geqslant \epsilon n$. Then, for a random walk W starting at v, we have*

$$\mathbb{P}(W_i \in S) \geqslant \pi\cdot\mathbf{1}_S - 8\sqrt{\epsilon}/\rho \geqslant |S|/n - 9\sqrt{\epsilon}/\rho,$$

for $i \geqslant 2$ and n sufficiently large.

Proof. We first show that the random walk is quite well mixed after only two steps. Let A be the set of neighbours of v with degree at most $(\rho+\epsilon)n$ and B be the set of vertices with at least $(\rho-\epsilon)|A|$ neighbours in A; thus A and B are the 'well-behaved' first and second neighbourhoods of v. By ϵ-quasirandomness, $|A| \geqslant d(v)-\epsilon n \geqslant (\rho-2\epsilon)n$ and $|B| \geqslant (1-\epsilon)n$. We have

$$\mathbf{1}_v P = \frac{1}{d(v)}\mathbf{1}_{N(v)} \geqslant \frac{1}{(\rho+\epsilon)n}\mathbf{1}_A,$$

where the inequality holds in each coordinate. For $x \in B$,

$$(\mathbf{1}_A P)_x = \sum_{\substack{y\in A\\ xy\in E(G)}} \frac{1}{d(y)} \geqslant \frac{(\rho-\epsilon)(\rho-2\epsilon)n}{(\rho+\epsilon)n} \geqslant \rho(1-4\epsilon/\rho),$$

where the first inequality holds since each $y \in A$ has degree at most $(\rho+\epsilon)n$, x has at least $(\rho-\epsilon)|A|$ neighbours in A and $|A| \geqslant (\rho-2\epsilon)n$. Since the entries of P are non-negative we can compose these inequalities to obtain

$$\mathbf{1}_v P^2 \geqslant \frac{(1-5\epsilon/\rho)}{n}\mathbf{1}_B.$$

Let $\mathbf{b} = \frac{(1-5\epsilon/\rho)}{n}\mathbf{1}_B$. Since $\pi_x = \frac{d(x)}{2\rho\binom{n}{2}}$, if x is a balanced vertex then $\frac{(1-\epsilon/\rho)}{n-1} \leqslant \pi_x \leqslant \frac{(1+\epsilon/\rho)}{n-1}$; otherwise we have the weaker bound $\pi_x \leqslant \frac{1}{\rho n}$. Since at most $2\epsilon n$ vertices are unbalanced and at most ϵn vertices are not in B,

$$\|\mathbf{b}-\pi\|_2 \leqslant \left(n\left(\frac{7\epsilon}{\rho n}\right)^2 + 3\epsilon n\left(\frac{2}{\rho n}\right)^2\right)^{1/2} \leqslant \left(\frac{64\epsilon}{\rho^2 n}\right)^{1/2}.$$

Then, for $i \geqslant 2$,

$$\begin{aligned}\mathbb{P}(W_i \in S) &= \mathbf{1}_v P^i \mathbf{1}_S \\ &= \mathbf{1}_v P^2 \cdot P^{i-2}\mathbf{1}_S \\ &\geqslant \mathbf{b}P^{i-2}\mathbf{1}_S \\ &= \pi P^{i-2}\mathbf{1}_S + (\mathbf{b}-\pi)P^{i-2}\mathbf{1}_S.\end{aligned}$$

By Cauchy-Schwarz, and the fact that the eigenvalues of P are at most 1 in absolute value,

$$\|(\mathbf{b}-\pi)P^{i-2}\mathbf{1}_S\|_2 \leqslant \|\mathbf{b}-\pi\|_2 \|\mathbf{1}_S\|_2 \leqslant \left(\frac{64\epsilon|S|}{\rho^2 n}\right)^{1/2} \leqslant 8\sqrt{\epsilon}/\rho,$$

and so

$$\mathbb{P}(W_i \in S) \geqslant \pi \cdot \mathbf{1}_S - 8\sqrt{\epsilon}/\rho,$$

proving the first inequality. Since at least $|S| - 2\epsilon n$ elements of S are balanced,

$$\pi \cdot \mathbf{1}_S = \sum_{x\in S}\frac{d(x)}{2\rho\binom{n}{2}} \geqslant \frac{(|S|-2\epsilon n)(\rho-\epsilon)}{\rho n} \geqslant |S|/n - 2\epsilon - \epsilon/\rho \geqslant |S|/n - \sqrt{\epsilon}/\rho,$$

which proves the second inequality. □

We now consider the following variant of the list model for constructing a random walk. Fix some small length L and let $K = \alpha n^2/L$. By a *block rooted at* v we mean a random walk of length L starting at v. For each vertex v, let Λ_v be an infinite list of blocks rooted at v. We construct a random walk of length αn^2 as follows. Choose W_0 from the given initial distribution, and, at each stage $s = 1, \ldots, K$, let $W_{(s-1)L}\cdots W_{sL}$ be the first unused block rooted at $W_{(s-1)L}$. At the end of the construction we have examined K blocks in total from the top of the n lists. Let M be the set of blocks examined (equivalently, the *multiset* of roots of blocks used).

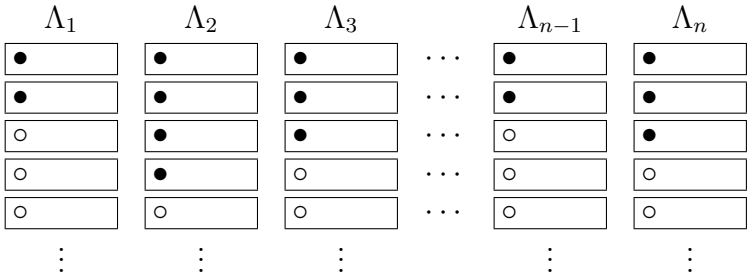


Figure 1: The construction examines K blocks from the top of the lists Λ_v, but we cannot tell in advance which blocks these will be.

This construction generalises the simple list model (which corresponds to the case $L = 1$), and we again hope to exploit the independence of blocks by applying standard concentration inequalities. There are two main obstacles. One is that we do not know anything about the distribution of a block rooted at a vertex v which is not balanced. We therefore first show that most of the root vertices are balanced. The second obstacle is that we do not know in advance which set of blocks we will examine. We handle this by approaching the problem from the other direction: for a given multiset M, what is the probability that the corresponding blocks do not contain an even distribution of the vertices? This turns out to be small enough that summing over all possible M gives the bound we require.

Lemma 13. *Let G be a connected n-vertex ϵ-quasirandom graph with $\rho\binom{n}{2}$ edges, and let W be a random walk of length αn^2 starting at a balanced vertex of G. Let $\delta = 3\sqrt[4]{\epsilon}/\sqrt{\rho}$ and suppose that n is sufficiently large. Then with probability at least $1 - 3\delta$ there exists a set $B \subseteq V(G)$, with $|B| \geqslant (1-\delta)n$, such that each vertex in B is hit at least $(1-4\delta)\alpha n$ times by W.*

Proof. Take $L = \omega_n$ for any $\omega_n \ll n/\log n$ which tends to infinity as $n \to \infty$, and let $K = \alpha n^2/L$. Construct a random walk W as described above and let $x_1, \ldots, x_K$ be the roots of the K blocks used. We first show that with high probability many of the vertices $\{x_1, \ldots, x_K\}$ are balanced.

Let U be the number of x_i that are unbalanced. The start vertex x_1 is always balanced by assumption. By Lemma 12, for $i \geqslant 2$,

$$\mathbb{P}\left(x_i \text{ is unbalanced}\right) \leqslant 1 - \left((1-2\epsilon) - \delta^2\right) \leqslant 2\delta^2,$$

since there are at least $(1-2\epsilon)n$ balanced vertices and $\delta^2 > 2\epsilon$. By Markov's inequality,

$$\mathbb{P}\left(U \geqslant \delta K\right) \leqslant \frac{\mathbb{E}\left(U\right)}{\delta K} \leqslant \frac{2\delta^2 K}{\delta K} = 2\delta,$$

so with probability at least $1 - 2\delta$, at least $(1-\delta)K$ blocks are used whose starting vertex is balanced. We now show that, with high probability, for any such multiset of

balanced vertices M, the corresponding blocks contain most balanced vertices about the right number of times.

So let M be a fixed multiset of $(1-\delta)K$ balanced vertices and let $W^{(1)}, W^{(2)}, \ldots, W^{((1-\delta)K)}$ be the corresponding blocks. Note that, as M is fixed, the blocks $W^{(i)}$ are independent. Let $S \subseteq V(G)$ with $|S| \geqslant \delta n$. By Lemma 12, for every $1 \leqslant i \leqslant (1-\delta)K$ and every $j \geqslant 2$ we have $\mathbb{P}\left(W_j^{(i)} \in S\right) \geqslant \delta - \delta^2$. Let X_{ij} be the indicator of the event $W_j^{(i)} \in S$, let $X_j = \sum_{i=1}^{K} X_{ij}$ and let $X_{M,S} = \sum_{j=1}^{L} X_j$. For fixed j the X_{ij} are independent, so by Chernoff's inequality (see [1, Appendix A]),

$$\mathbb{P}\left(X_j < (\delta - 2\delta^2)|M|\right) \leqslant e^{-2\delta^4|M|}.$$

Hence

$$\begin{aligned}\mathbb{P}\left(X_{M,S} < (\delta - 4\delta^2)\alpha n^2\right) &\leqslant \mathbb{P}\left(X_{M,S} < (\delta - 3\delta^2)(1-\delta)KL\right)\\ &\leqslant \mathbb{P}\left(X_j < (\delta - 2\delta^2)|M| \text{ for some } 2 \leqslant j \leqslant L\right)\\ &\leqslant Le^{-2\delta^4|M|},\end{aligned}$$

where the second inequality holds for large n because the contribution from X_1 is negligible as $L \to \infty$.

Let $A = \{x \in V(G) : x \text{ is visited at least } (1-4\delta)n \text{ times by } W\}$. If $|A| < (1-\delta)n$, then either δK of the x_i are unbalanced, or there is set S of δn vertices and a multiset M of $(1-\delta)K$ balanced vertices of $V(G)$ for which $X_{M,S} < (\delta - 4\delta^2)\alpha n^2$. Hence

$$\begin{aligned}\mathbb{P}\left(|A| < (1-\delta)n\right) &\leqslant \mathbb{P}\left(U \geqslant \delta K\right) + \sum_M \sum_S Le^{-2\delta^4|M|}\\ &\leqslant 2\delta + \binom{K+n-1}{n-1}\binom{n}{\delta n} Le^{-2\delta^4(1-\delta)K}\\ &\leqslant 2\delta + O(K)^n \cdot 2^n \cdot L \cdot e^{-2\delta^4(1-\delta)K}\\ &\leqslant 2\delta + \exp\bigl(O(n\log n) + O(n) + O(\log n) - 2\delta^4(1-\delta)K\bigr)\\ &\leqslant 3\delta,\end{aligned}$$

for n sufficiently large, since $K \gg n\log n$. □

We now have everything we need to complete the proof of Theorem 1.

Proof of Theorem 1. We will show that, with probability $1 - o_\epsilon(1)$, the graph $G_{walk}(\alpha)$ is close to $G_{list}(\alpha/\rho)$. It then follows from Theorem 6 that $G_{walk}(\alpha)$ is $o_\epsilon(1)$-quasirandom with probability $1 - o_\epsilon(1)$.

Since there are at most $2\epsilon n < \delta$ unbalanced vertices in G, by Lemma 13, with probability at least $1 - 3\delta$, there is a set B of $(1-2\delta)n$ balanced vertices such that every $v \in B$ is hit at least $(1-4\delta)\alpha n \geqslant (1-5\delta)\frac{\alpha}{\rho}d(v)$ times by W. This accounts for $(1-2\delta)n \cdot (1-4\delta)\alpha n \geqslant (1-7\delta)\alpha n^2$ of the list entries examined, so $G_{walk}(\alpha)$ differs from $G_{list}(\alpha/\rho)$ by at most $14\delta\alpha n^2$ edges. Since δ tends to 0 with ϵ, the result follows. □

5 Trees

A *homomorphism* from a graph H to a graph G is an edge-preserving map $\phi : V(H) \to V(G)$. A random walk can be viewed as a random homomorphism of a path; a natural generalisation is to consider a random homomorphism of some other tree T (sometimes called a *tree-indexed random walk*). Just as we traversed a path in one direction, our trees will be rooted and we think of them as directed 'downwards', away from the root. In this section we will explore to what extent the methods of Section 4 can be applied in this more general setting.

We generate a random homomorphism as follows. Enumerate the vertices of T as $v_0, v_1, \dots, v_k$ where, for each j, $T[v_0, \dots, v_j]$ is a connected subtree of T containing the root v_0. First choose $\phi(v_0)$ from a given initial distribution. Then, at each stage $j > 0$, let u be the parent of v_j in T and choose $\phi(v_j)$ uniformly at random from the neighbours of $\phi(u)$. All choices are made independently, and we can think of these choices as being taken from the lists L_v as before.

Suppose now that G is an ϵ-quasirandom graph on n vertices. Let ϕ be a random homomorphism of a tree T of size αn^2 to G, and let $G(T)$ be the subgraph of G consisting of the edges in the image of ϕ. Is $G(T)$ quasirandom with high probability? It is easy to see that in general the answer is no. For example, let $G = K_n$ and let T be an $n/2$-ary tree of depth 2 (here $\alpha = 1/4 + o(1)$). Then with high probability $\phi(T)$ contains a constant fraction of the edges of G. But all of these edges are incident on the neighbourhood of the root, which has only $(1 - e^{-1/2} + o(1))n$ vertices with high probability, so, with high probability, $G(T)$ is not quasirandom.

We seek conditions on T such that we can apply the approach taken in Section 4 with minimal changes. The condition we give here imposes an upper bound on the maximum degree of T.

We need an analogue of the second model for the construction of a random walk. Instead of breaking our path into many short paths, we break our tree into many small edge-disjoint subtrees.

Lemma 14. *Let T be a rooted tree with N edges and let $L \leqslant N$. Then T can be written as an edge-disjoint union of rooted trees $R_1, \dots, R_K$, each of size between L and $3L$.*

Proof. Let v be a vertex of T furthest from the root such that v has at least L descendants. Then each branch of T lying below v has at most L edges, so some union of these branches has size between L and $2L$; let this be R_1. We obtain $R_2, \dots, R_K$ similarly until there are less than L edges of T remaining, which we add to R_K. □

Write $\mathcal{R} = \{R_1, \dots, R_K\}$ for the corresponding set of abstract rooted trees, up to isomorphism. In an abuse of notation we use R_i to refer to both the specific subtree of T and its isomorphism type.

It is convenient to number the R_i such that $R_1 \cup \dots \cup R_j$ is a subtree of T containing the root for each j. We can then describe the second model for the construction of a random homomorphism as follows. For each $v \in V(G)$ and $R \in \mathcal{R}$, let $\Lambda_{v,R}$ be a list of independent random homomorphisms from R to G that map the root of R to v. Choose

a vertex v_1 from the given distribution for the image of the root of T and identify $\phi(R_1)$ with the first entry from Λ_{v_1,R_1}. (If R_1 has a non-trivial automorphism group then there is a choice of identification of R_1 with the reference copy in $\mathcal{R}$. The choice is unimportant provided the same choice is made every time.) Then at each stage j we have already determined the image v_j of the root of R_j, and we identify $\phi(R_j)$ with the first unused element of Λ_{v_j,R_j}.

Now let T be a rooted tree with αn^2 edges. As before we want to show that T 'visits' most vertices of G about the right number of times. We need to be careful here about what counts as a 'visit': what we want to count is the number of times an edge leaves a vertex, as that is the number of entries of the corresponding list that will be examined. So we say $\phi(T)$ *visits* $x \in V(G)$ whenever uv is an edge of T with u the parent of v and $\phi(u) = x$; the number of visits $\phi(T)$ makes to x is the number of edges uv for which this occurs.

There are three places where the argument in the proof of Lemma 13 needs modification or additional details need to be checked.

(i) In the path case the edges (or vertices) of the blocks had a natural order and the blocks were all the same size. In the tree case we are free to choose a labelling of the edges in each block, but the blocks might still have different sizes: when we look at the $2L$th edge from each block, are there enough blocks with $2L$ edges that Chernoff's inequality will give good concentration?

(ii) In the path case the set of list entries examined was parameterised by multisets of vertices of G. In the tree case the set of list entries examined is instead parameterised by multisets of pairs (v, R) with $v \in V(G)$ and $R \in \mathcal{R}$. So the factor $\binom{K+n-1}{n-1}$ in the final sum needs to be replaced by $\binom{K+n|\mathcal{R}|-1}{n|\mathcal{R}|-1}$, and we must restrict the size of $\mathcal{R}$ to prevent this becoming too large.

(iii) In the path case we had to ignore the first two vertices of each block as we needed to take two steps before we had good information about the distribution over vertices. This was safe because the ignored vertices were only a $o(1)$ fraction of the total number of vertices. In the tree case we must ignore the edges whose start point is the root of the block or is a child of the root. We need to ensure that the number of ignored edges is at most a small fraction of the total number of edges.

Problem (i) is avoided by throwing away the small number of edges that receive a label shared by few other edges. If we throw away all edges that receive a label which is used less than $\epsilon n^2/L^2$ times then the total number of edges thrown away is less than $3\epsilon n^2/L$ as there are at most $3L$ edges in each block.

Problem (ii) is avoided by taking L small: $L = \frac{\log n}{2 \log 3}$ suffices. Indeed, since the number of rooted trees on L vertices is $O((2.9955\ldots)^L)$ (see [8]) and $\frac{\alpha n^2}{3L} \leqslant K \leqslant \frac{\alpha n^2}{L}$, we have in this case that $n|\mathcal{R}| \ll n^{3/2} \ll K$, and

$$\binom{K + n|\mathcal{R}| - 1}{n|\mathcal{R}| - 1} \ll K^{n|\mathcal{R}|} \ll \exp\left(O(n^{3/2} \log n)\right),$$

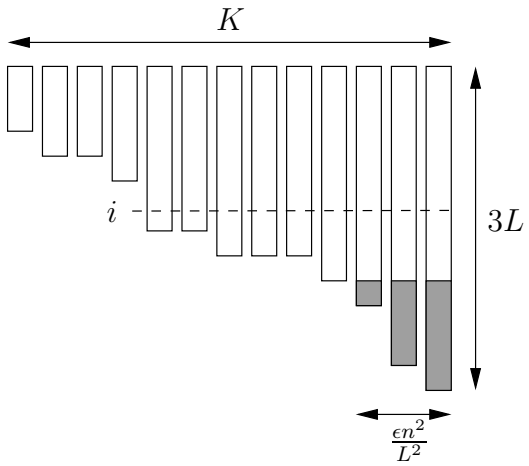


Figure 2: Deleting a $o(1)$ fraction of the edges ensures that the remaining labels i are each used in a large number of blocks.

which is small enough that it will not overpower the e^{-cK}-type decay.

Problem (iii) is avoided by having Δ^2, the square of the maximum degree of T small (depending on the desired level of quasirandomness) compared to L: so Δ can be as large as a small multiple of $\sqrt{\log n}$.

With these modifications to our earlier argument we obtain the following result.

Theorem 15. *Given $\alpha, \rho, \eta > 0$ there exists $\epsilon, c > 0$ such that the following holds. Let G be an n-vertex ϵ-quasirandom graph with $\rho\binom{n}{2}$ edges, let T be a rooted tree of size αn^2 with maximum degree $\Delta \leqslant c\sqrt{\log n}$ and let ϕ be a random homomorphism from T to G such that the image of the root is balanced. Then, with probability $1 - o_\epsilon(1)$, the subgraph $G(T)$ of G consisting of the edges of $\phi(T)$ is η-quasirandom with $(1 - e^{-2\alpha/\rho} + o_\epsilon(1))\rho\binom{n}{2}$ edges.*

It would be interesting to know how large $\Delta(T)$ can be taken in Theorem 15. By the example at the start of this section we must have $\Delta(T)$ small compared to n. Is this already enough?

Acknowledgements

We would like to thank the organisers of the *Probabilistic methods in graph theory* workshop at the University of Birmingham where we heard about this problem. We would also like to thank Jan Hladký for some helpful discussions and for suggesting the extension of our results on paths to more general trees.

References

[1] N. Alon and J. Spencer. *The probabilistic method*. John Wiley & Sons Inc., third edition, 2008.

[2] F. R. K. Chung, R. L. Graham and R. M. Wilson. Quasi-random graphs, *Combinatorica*, 9(4):345-362, 1989.

[3] J. Hladký. Problem session, Probabilistic methods in graph theory, University of Birmingham, 2012.

[4] M. Jerrum and A. Sinclair. Approximate counting, uniform generation and rapidly mixing Markov chains. *Inform. and Comput.*, 82(1):93-133, 1989.

[5] M. Krivelevich and B. Sudakov. Pseudo-random graphs. In *More sets, graphs and numbers*, volume 15 of *Bolyai Soc. Math. Stud.*: pages 199-262. Springer, Berlin, 2006.

[6] D. A. Levin, Y. Peres and E.L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. With a chapter by J.G. Propp and D.B. Wilson.

[7] M. Molloy and B. Reed. *Graph colouring and the probabilistic method*. Springer-Verlag, 2002.

[8] R. Otter. The number of trees. *Ann. of Math. (2)*, 49:583-599, 1948.

[9] M. Talagrand. Concentration of measure and isoperimetric inequalities in product spaces. *Inst. Hautes Études Sci. Publ. Math.*, (81):73-205, 1995.

[10] A. Thomason. Pseudorandom graphs. In *Random graphs '85 (Poznań, 1985)*, volume 144 of *North-Holland Math. Stud.*, p.307-331. North-Holland, Amsterdam, 1987.