

 Open access • Journal Article • DOI:10.1137/080729542

Multiple Random Walks in Random Regular Graphs — [Source link](#)

[Colin Cooper](#), [Alan Frieze](#), [Tomasz Radzik](#)

Institutions: [Carnegie Mellon University](#)

Published on: 01 Nov 2009 - [SIAM Journal on Discrete Mathematics](#) (Society for Industrial and Applied Mathematics)

Topics: [Random regular graph](#), [Random walk](#), [Random graph](#), [Vertex \(graph theory\)](#) and [Cover \(topology\)](#)

Related papers:

- [Many random walks are faster than one](#)
- [Tight bounds for the cover time of multiple random walks](#)
- [Markov Chains and Mixing Times](#)
- [Random Walks on Graphs: A Survey](#)
- [On the coalescence time of reversible random walks](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/multiple-random-walks-in-random-regular-graphs-1h4yr779ga>

Multiple random walks in random regular graphs

Colin Cooper* Alan Frieze† Tomasz Radzik*

June 1, 2009

Abstract

We study properties of multiple random walks on a graph under various assumptions of interaction between the particles. To give precise results, we make the analysis for random regular graphs.

The cover time of a random walk on a random r -regular graph was studied in [6], where it was shown with high probability (**whp**), that for $r \geq 3$ the cover time is asymptotic to $\theta_r n \ln n$, where $\theta_r = (r-1)/(r-2)$.

In this paper we prove the following (**whp**) results, arising from the study of multiple random walks on a random regular graph G . For k independent walks on G , the cover time $C_G(k)$ is asymptotic to C_G/k , where C_G is the cover time of a single walk. For most starting positions, the expected number of steps before any of the walks meet is $\theta_r n / \binom{k}{2}$. If the walks can communicate when meeting at a vertex, we show that, for most starting positions, the expected time for k walks to broadcast a single piece of information to each other is asymptotic to $\frac{2 \ln k}{k} \theta_r n$, as $k, n \rightarrow \infty$.

We also establish properties of walks where there are two types of particles, predator and prey, or where particles interact when they meet at a vertex by coalescing, or by annihilating each other. For example, the expected extinction time of k explosive particles (k even) tends to $(2 \ln 2) \theta_r n$ as $k \rightarrow \infty$.

The case of n coalescing particles, where one particle is initially located at each vertex, corresponds to a voter model defined as follows: Initially each vertex has a distinct opinion, and at each step each vertex changes its opinion to that of a random neighbour. The expected time for a unique opinion to emerge is the same as the expected time for all the particles to coalesce, which is asymptotic to $2\theta_r n$.

*Department of Computer Science, King's College, University of London, London WC2R 2LS, UK (colin.cooper@kcl.ac.uk, tomasz.radzik@kcl.ac.uk). Research supported by Royal Society Grant 2006/R2-IJP.

†Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213. Email: alan@random.math.cmu.edu. Research supported in part by NSF grant CCF0502793.

Combining results from the predator-prey and multiple random walk models allows us to compare expected detection time of all prey in the following scenarios: both the predator and the prey move randomly, the prey moves randomly and the predators stay fixed, the predators move randomly and the prey stays fixed. In all cases, with k predators and ℓ prey the expected detection time is $\theta_r H_\ell n/k$, where H_ℓ is the ℓ -th harmonic number.

1 Introduction

Let $G = (V, E)$ be a connected graph, let $|V| = n$, and $|E| = m$. For $v \in V$ let C_v be the expected time taken for a simple random walk W on G starting at v , to visit every vertex of G . The *vertex cover time* C_G is defined as $C_G = \max_{v \in V} C_v$. The (vertex) cover time of connected graphs has been extensively studied. It is a classic result of Aleliunas, Karp, Lipton, Lovász and Rackoff [3] that $C_G \leq 2m(n-1)$. It was shown by Feige [11], [12], that for any connected graph G , the cover time satisfies $(1 - o(1))n \ln n \leq C_G \leq (1 + o(1))\frac{4}{27}n^3$. The complete graph K_n , is an example of a graph achieving the lower bound. The *lollipop* graph consisting of a path of length $n/3$ joined to a clique of size $2n/3$ has cover time asymptotic to the upper bound.

Let \mathcal{G}_r denote the set of r -regular graphs with vertex set $V = \{1, 2, \dots, n\}$ and the uniform measure. In [6] which studied the cover time of random regular graphs \mathcal{G}_r , the following theorem was proved:

Theorem 1. *Let $r \geq 3$ be constant. Let $\theta_r = \frac{r-1}{r-2}$. Let G be chosen randomly from \mathcal{G}_r , then **whp***

$$C_G \sim \theta_r n \ln n.$$

A few words on notation etc. The results in this paper are always asymptotic in $n = |V|$. The notation $A_n \sim B_n$ means that $\lim_{n \rightarrow \infty} A_n/B_n = 1$, and **whp** (with high probability) means with probability tending to 1 as $n \rightarrow \infty$. For $r \geq 3$, we define a subset of \mathcal{G}_r of *typical* graphs \mathcal{G}'_r such that $|\mathcal{G}'_r| \sim |\mathcal{G}_r|$ and for which our results hold. If G is chosen randomly from \mathcal{G}_r , then results which are true for typical graphs, hold **whp** over our choice of G . In Theorems 5–10 we assume that G is typical, although this is not stated explicitly.

Suppose there are $k \geq 1$ particles, each making a simple random walk on the graph G . Essentially there are two possibilities, either the particles are oblivious of each other or can interact on meeting. *Oblivious particles* act independently of each other, with no interaction on meeting. *Interactive particles* can interact directly in some way on meeting. For example they may exchange information, coalesce, reproduce, destroy each other. We assume that interaction occurs *only when meeting at a vertex*, and that the random walks made by the

particles are otherwise independent. For such models we study various properties of the walks, for example:

Multiple walks. For k particles walking independently, we establish the cover time $C_G(k)$ of G .

Talkative particles. For k particles walking independently, which communicate on meeting at a vertex, we study the expected time to broadcast a message.

Predator-Prey. For k predator and ℓ prey particles walking independently, we study the expected time to extinction of the prey particles, when predators eat prey particles on meeting at a vertex.

Coalescing particles. For k particles walking independently, which coalesce on meeting at a vertex, we study the expected time to coalesce to a single particle.

Annihilating particles. For $k = 2\ell$ particles walking independently, which destroy each other (pairwise) on meeting at a vertex, we study the expected time to extinction of all particles.

The motivation for these models comes from many sources, and we give a brief introduction. A further discussion, with detailed references is given in the appropriate sections below.

Using random walks to test graph connectivity is an established algorithm, and it is obvious to try to speed this up by parallel searching. Similarly, properties of communication, such as broadcasting and gossiping between particles moving in a network, is a natural question. The predator-prey model is an established model of interacting particle systems. Combining results from the predator-prey and multiple random walk models allows us to compare expected detection time for the following scenarios: both the predator and the prey move, the prey moves and the predators stay fixed, the predators move and the prey stays fixed.

Coalescing and annihilating particle systems are part of the classical theory of interacting particles; an established area to which our paper makes a new contribution. A system of coalescing particles where initially one particle is located at each vertex, corresponds to another classical problem, the voter model, which is defined as follows: Initially each vertex has a distinct opinion, and at each step each vertex changes its opinion to that of a random neighbour. It is known that the expected time for a unique opinion to emerge, is the same as the expected time for all the particles to coalesce. By establishing the expected coalescence time, we obtain the expected time for voting to be completed.

Remark. The results in this paper hold for any r -regular graph which has the *typical* properties listed in Section 2. Many r -regular vertex transitive graphs of high girth have these typical properties, for example the Lubotzky, Phillips and Sarnak construction of Ramanujan

graphs [16]. The high girth is not a necessary condition for our proofs, but simplifies the calculations somewhat.

Results: Oblivious particles

Our first result concerns speedup of cover time. Let $T(k, v_1, \dots, v_k)$ be the time taken to visit all vertices using k independent walks starting at vertices v_1, \dots, v_k . Define the k -particle cover time $C_k(G)$ in the natural way as $C_k(G) = \max_{v_1, \dots, v_k} \mathbf{E}(T(k, v_1, \dots, v_k))$ and define the speedup as $S_k = C(G)/C_k(G)$. The speedup can vary considerably depending on the graph structure. This can be seen from the following results, which are easily proved. For the complete graph K_n , the speedup is k ; for P_n , the path of length n the speedup is $\Theta(\ln k)$.

Using k independent random walks to improve s - t connectivity testing was studied by Broder, Karlin, Raghavan and Upfal [5]. They proved that for k random walks starting from (positions sampled from) the stationary distribution, the cover time of an m edge graph is $O((m^2 \ln^3 n)/k^2)$. In the case of r -regular graphs, Aldous and Fill [2] (Chapter 6, Proposition 17) give an upper bound on the cover time of $C_k \leq (25 + o(1))n^2 \ln^2 n/k^2$, which holds for $k \geq 6 \ln n$.

More recently, the value of $C_k(G)$ was studied by Alon, Avin, Koucký, Kozma, Lotker and Tuttle [4] for general classes of graphs. They found that for expanders the speedup was $\Omega(k)$ for $k \leq n$ particles. They also give an example, the barbell graph, (two cliques joined by a long path) for which the speed-up is exponential in k provided $k \geq 20 \ln n$.

In the case of random r -regular graphs, we establish the k -particle cover time. Comparing Theorem 2 with Theorem 1, we see that $C_G(k) \sim C_G/k$, i.e. the speedup is exactly linear, as is the case for the complete graph.

Theorem 2. Multiple particles walking independently.

Let $r \geq 3$ be constant. Let G be chosen randomly from \mathcal{G}_r , then **whp**.

(i) For $k = o(n/\ln^2 n)$ the k -particle cover time $C_G(k)$ satisfies

$$C_G(k) \sim \frac{\theta_r}{k} n \ln n,$$

and this result is independent of the initial positions of the particles.

(ii) For any k , $C_G(k) = O\left(\frac{n}{k} \ln n + \ln n\right)$.

Suppose we distinguish two types of particles, mobile, and fixed; and that mobile particles are predators and the fixed particles are prey (or vice versa). An application of the methods used

in Theorem 2 give the following theorem, concerning e.g. cops and robbers. For comparison with the case where both predator and prey move, we have also included a result of Theorem 7 for the predator-prey model. The moral of the story is that as long as at least one particle type moves, the expected detection time is the same.

Theorem 3. Comparison of search models.

Let $k, \ell \leq n^\epsilon$ for a sufficiently small positive constant ϵ . Let G be chosen randomly from \mathcal{G}_r , then **whp**.

- (i) Suppose there are k mobile predator particles walking randomly, and ℓ prey particles fixed at randomly chosen vertices of G . Let $\mathbf{E}(F_{k,\ell}^{(i)})$ be the expected detection time of all prey particles.
- (ii) Suppose there are ℓ mobile prey particles walking randomly, and k predator particles fixed at randomly chosen vertices of G . Let $\mathbf{E}(F_{k,\ell}^{(ii)})$ be the expected detection time of all prey particles.

Let $\mathbf{E}(D_{k,\ell})$ be the expected extinction time of ℓ mobile prey using k mobile predators, as given by Theorem 7. Then **whp**

$$\mathbf{E}(F_{k,\ell}^{(i)}) \sim \mathbf{E}(F_{k,\ell}^{(ii)}) \sim \mathbf{E}(D_{k,\ell}) \sim \frac{\theta_r H_\ell}{k} n,$$

where H_ℓ is the ℓ -th harmonic number.

First meeting between particles

Meetings between a pair of particles walking randomly on a graph have been studied in many contexts. The simplest model, the one we consider here, is to allow both particles to move simultaneously at each step. A meeting occurs when both particles occupy the same vertex at the same step; but not when they pass on an edge.

Consider a pair of random walks, starting at vertices u and v . Let $\widehat{M}(u, v)$ be the number of steps before the walks first meet at a vertex. Clearly if $u = v$, then $\widehat{M}(u, v) = 0$. Let (v_1, v_2, \dots, v_k) be the starting positions of k walks. We say the walks start in *general position*, if there is a pairwise separation $d(v_i, v_j) \geq \omega(k, n)$ between the starting positions, where

$$\omega(k, n) = \Omega(\ln \ln n + \ln k). \tag{1}$$

Theorem 4. Let $k \leq n^\epsilon$, for sufficiently small positive constant ϵ .

For k independent random walks starting in a general position $\mathbf{v} = (v_1, v_2, \dots, v_k)$, let $\widehat{M}(\mathbf{v})$

denote the number of steps before any of the particles meet. Then

$$\mathbf{E}\widehat{M}(\mathbf{v}) \sim \frac{2\theta_r}{k(k-1)}n.$$

Our requirement that the particles should be in general position is not very restrictive. In the stationary distribution, the probability that some pair of particles are at most α apart from each other is $O(k^2r^\alpha/n)$. Thus a separation of $\alpha = \Omega(\max(\ln \ln n, \ln k))$, is both feasible and occurs **whp** in stationarity, provided $k \leq n^\epsilon$, for a sufficiently small positive constant ϵ . The motivation for the value (1) is as follows: Let $T_k = kT$, where T is the single particle mixing time. It is shown in Lemma 20 that with separation $\omega(k, n)$, the probability of some pair of particles meeting during the mixing time T_k is $o(1)$.

By Theorem 4, for a pair of particles starting in general position on a random r -regular graph, the expected first meeting time is $\mathbf{E}\widehat{M}(u, v) \sim \theta_r n$. The quantity $\mathbf{E}\widehat{M}(u, v)$ is closely related to the expected time to reach a random vertex. For a random walk starting at vertex u , let $H(u, v) = \mathbf{E}_u T_v$ be the (expected) hitting time of vertex v . It follows directly from [6] that (**whp**) an r -regular random graph G has the following property. Independently of u , all but $O((\ln n)^{O(1)})$ vertices v have hitting time $H(u, v) \sim \theta_r n$, and this is an upper bound for hitting time of the other vertices. Thus for most pairs of vertices $H(u, v) \sim \theta_r n \sim \mathbf{E}\widehat{M}(u, v)$. The quantity $\theta_r n$, has the interpretation of the hitting time $\sum_j \pi_j \mathbf{E}_i T_j$ of a random target.

Turning briefly to other models of random walks, one possibility governing the movement of the particles, is to have a controller (e.g. demonic, angelic) who allows only one particle to move at each step. The aim of this is to delay or speedup the meeting between the particles. In [18] Tetali and Winkler give general $O(n^3)$ worst case expected bounds for these processes. The demonic controller model was also studied by Coppersmith, Tetali and Winkler, in [10], and by Aldous and Fill [2], as *Cat and Mouse Games* (see Chapters 3, 6).

Random walks with demonic controllers is a model which arises in the context of self-stabilizing networks. Two particles starting at u, v make random walks, but only one particle is allowed to move at each step. The particle which moves at a given step is determined by a demon, whose aim is to delay the meeting as long as possible. In [10] it is shown that $M(u, v)$, the (worst case) expected number of steps before meeting, satisfies

$$H(u, v) \leq M(u, v) \leq H(u, v) + H(v, t) - H(t, v), \quad (2)$$

for some fixed vertex t . The existence of the vertex t , which is called *hidden*, is established in [10], and is defined by the property that $H(t, v) \leq H(v, t)$, for all $v \in V$.

In our study, there is no demon, and both of the particles walk randomly until they meet. As already remarked, for a pair of particles starting at vertices u, v , in general position we find that $\mathbf{E}\widehat{M}(u, v) \sim \theta_r n$.

In fact, for most vertex pairs (u, v) we also have that $M(u, v) \sim \theta_r n$. We briefly outline why this is the case. The results we claim follow from [6], and are summarized in Lemma 15. Say that a vertex v is *tree-like* if there is no cycle in the neighbourhood $N(v, L_0)$, the set of vertices within distance L_0 of v , where $L_0 = \alpha \ln \ln n$ for some absolute constant $\alpha > 0$. Most vertices v of a random regular graph are *tree-like* **whp**. If u is not in $N(v, L_0)$ then $H(u, v) \sim \theta_r n$, and moreover $H(u, v) \leq \theta_r n(1 + o(1))$ for all u, v . Let t be hidden and pick a tree-like v , of distance at least L_0 from both u and t . There are $n - O((\ln n)^A)$ such vertices for some constant $A > 0$. It is immediate from (2) that

$$H(u, v) \leq M(u, v) \leq H(u, v)(1 + o(1)).$$

Thus for most choices of u and v , $M(u, v) \sim \theta_r n$.

Results: Interacting particles

We first consider problems of passing messages between particles. We assume that when particles meet at a vertex they exchange all messages which they know. We refer to such particles as *agents*, to distinguish them from non-communicating particles. If initially one agent has a message it wants to pass to all the others, we refer to this process as *broadcasting* (among the agents). Formally, there are two sets $I(t), U(t)$ of informed and uninformed vertices respectively. Initially $I = \{\rho_1\}$, where ρ_1 is the agent with the message, and $U = \{\rho_2, \dots, \rho_k\}$. If a member of I meets a member ρ of U , then ρ becomes informed and is moved from U to I . The broadcast time is the step at which $U = \emptyset$.

Theorem 5. Broadcast time.

Let $k \leq n^\epsilon$ for a sufficiently small positive constant ϵ . Suppose k agents make random walks starting in general position. Let B_k be the time taken for a given agent to broadcast to all other agents. Then

$$\mathbf{E}(B_k) \sim \frac{2\theta_r}{k} H_{k-1} n,$$

where H_k is the k -th harmonic number. Thus when $k \rightarrow \infty$, $\mathbf{E}(B_k) \sim \frac{2\theta_r \ln k}{k} n$.

If each agent has a message it wants to pass to all other agents, we refer to this process as *gossiping*. When $k \rightarrow \infty$, the proof of Theorem 5 also gives an upper bound for the agents to complete gossiping among themselves.

Corollary 6. Gossip time.

Let $k \rightarrow \infty$, then **whp** gossiping among the agents can be completed in time $O(n(\ln k)^2/k)$.

An alternative and less efficient way to pass on a message, is for the the originating agent to tell it directly to all the others. Compared to this, broadcasting improves the expected time

for everybody to receive the message by a multiplicative factor of $k/2$, for large k . To see this, compare $\mathbf{E}(B_k)$ of Theorem 5, with $\mathbf{E}(D_{1,k-1})$ of Theorem 7 below. Meeting directly with all other agents corresponds to a predator-prey process with one predator (the broadcaster) and $k - 1$ prey.

Our next results are for particles which interact in a far from benign manner. Problems related to such interacting random walks are considered by Aldous and Fill [2]. In particular Chapter 14, *Interacting Particle Systems*, contains a wealth of results on this topic. Among the processes considered in this chapter are coalescing and annihilating particle systems, and their duals: the voter and anti-voter model.

One variant of interacting particles is a predator-prey model, in which both types of particles make independent random walks. If a predator encounters prey on a vertex it eats them. Formally, there is a set of predators W , and a set of prey R . When a member of W meets a member r of R , then r is deleted. The extinction time of the prey particles is the step at which $R = \emptyset$.

Theorem 7. Predator-prey.

Let $k, \ell \leq n^\epsilon$ for a sufficiently small positive constant ϵ . Suppose k predator and ℓ prey particles make random walks, starting in general position. Let $D_{k,\ell}$ be the extinction time of all the prey particles. Then

$$\mathbf{E}(D_{k,\ell}) \sim \frac{\theta_r H_\ell}{k} n.$$

A variant of predator-prey is the coalescing particle system in which particles stick together on meeting. Thus when two (or more) particles meet they are replaced by a single particle, which continues its random walk from the meeting vertex. The coalescence time is the step at which only one particle remains.

Theorem 8. Coalescence time: sticky particles.

Let $k \leq n^\epsilon$ for a sufficiently small positive constant ϵ . Let S_k be the coalescence time, when there are originally k sticky particles walking randomly, starting from general position. Then,

$$\lim_{k \rightarrow \infty} \mathbf{E}(S_k) \sim 2\theta_r n,$$

and the expected value for k constant is given by (36).

Coalescence time on r -regular graphs was studied by Aldous and Fill [2]. They considered an equivalent continuous time model, in which initially one particle starts at each vertex of an n vertex r -regular graph. For this model they prove that expected time to coalesce $\mathbf{E}C$ is $O(rn^2)$ and for the complete graph $\mathbf{E}C \sim n$. For further discussion of this, see the voter model below, and Section 8 Theorem 25 of this paper.

As a last twist on predator-prey, we consider a set P of particles which destroy each other (pairwise) on meeting at a vertex. The extinction time (of the particle system) is the step at which $P = \emptyset$.

Theorem 9. Extinction time: explosive particles.

Let $k \leq n^\epsilon$ for a sufficiently small positive constant ϵ . Suppose there are $k = 2\ell$ explosive particles walking randomly, starting in general position, and that particles destroy each other pairwise on meeting at a vertex. Let A_k be the time to extinction of all particles, then

$$\lim_{k \rightarrow \infty} \mathbf{E}(A_k) \sim 2\theta_r \ln 2n$$

and the expected value for k constant is given by (37).

Finally we consider the *voter model*. In this model, each vertex initially has a distinct opinion. At each time step, each vertex i contacts a random neighbour j , and changes its opinion to the opinion held by j . The number of opinions is non-increasing at each step. Let C be the number of steps needed for a unique opinion to emerge. It is shown in Chapter 14 of [2] that the random variable C for the voter model has the same distribution (and hence the same expected value) as the coalescence time in the case of n coalescing particles, where one particle is initially located at each vertex. Theorem 10 answers, in part, some open problems posed in Chapter 14 of [2].

Theorem 10. Voter model.

Let \mathbf{EC} be the expected coalescence time of n simple random walks when one particle starts on each vertex. Then, for typical r -regular vertex transitive graphs and **whp** for random r -regular graphs,

$$\mathbf{EC} \sim 2\theta_r n.$$

Structure of the paper

Methodology. For oblivious particles, we use the techniques and results of [6] and [8] to establish the probability that a vertex is unvisited by any of the walks at a given time t . Let T be a suitably large mixing time. Provided the graph is typical (Section 2) and the technical conditions of Lemma 13 are met, then the probability that a vertex v is unvisited at step T, \dots, t tends to $(1 - \pi_v/R_v)^t$. Here R_v is the number of returns to v during T by a walk starting at v . This value is a property of the structure of the graph around vertex v . For most vertices of a typical graph $R_v \sim \theta_r$, which explains the origin of this quantity.

In [6] a technique, vertex contraction, was used to estimate the probability that the random walk had not visited a given set of vertices. For interacting particles, we use this technique to derive the probability that a walk on a suitably defined product graph H_k has not visited

the diagonal (set of vertices $\mathbf{v} = (v_1, \dots, v_k)$ with repeated vertex entries v_i) at a given time t . Basically we contract the diagonal to a single vertex, γ , and analyse the walk in the contracted graph Γ .

Layout of the paper. In Section 2 we define the properties of r -regular graphs we require in order to make our analysis. These properties, which we call *typical* occur **whp** in random r -regular graphs. In Section 3 we establish some mixing properties of the product graph H_k , and state a lemma (Lemma 13) which gives the probability distribution of first visit time to a vertex. In Section 4 we establish the cover time for k particles walking independently. In Section 5 we establish R_γ the expected number of returns to the contracted vertex γ during the mixing time of the interacting particle system. The main task is a careful study of the interaction between pairs of particles. In Section 6 we prove the conditions of Lemma 13 apply to interacting particle systems. In Section 7 we complete the proofs of Theorems 4-9. Finally in Section 8 we give a more detailed summary of results for the voter model in [2] Chapter 14, and prove Theorem 10.

2 Typical r -regular graphs

We say an r -regular graph G is *typical* if it satisfies the properties **P1-P4** listed below. We first give some definitions. Let

$$L_1 = \lfloor \epsilon_1 \log_r n \rfloor, \quad (3)$$

where $\epsilon_1 > 0$ is a sufficiently small constant. A cycle C is *small* if $|C| \leq L_1$. A vertex v is *tree-like* if there is no cycle in the subgraph $G[v, L_1]$ induced by the set of vertices within distance L_1 of v .

P1. G is connected, and not bipartite.

P2. The second eigenvalue of the adjacency matrix of G is at most $2\sqrt{r-1} + \epsilon$, where $\epsilon > 0$ is an arbitrarily small constant.

P3. There are at most $n^{2\epsilon_1}$ vertices on small cycles.

P4. No pair of cycles C_1, C_2 with $|C_1|, |C_2| \leq 100L_1$ are within distance $100L_1$ of each other.

Note that P3 implies that at most n^{ϵ_C} vertices of a typical r -regular graph are not tree-like, where

$$n^{\epsilon_C} = O(r^{L_1} n^{2\epsilon_1}) = O(n^{3\epsilon_1}). \quad (4)$$

Theorem 11. Let $\mathcal{G}'_r \subseteq \mathcal{G}_r$ be the set of typical r -regular graphs. Then $|\mathcal{G}'_r| \sim |\mathcal{G}_r|$.

That P2 holds **whp** is a very difficult result of Friedman [14]. The other properties are straightforward to establish, see e.g. [6].

3 Estimating first visit probabilities

3.1 Convergence of the random walk

Let G be a connected graph with n vertices and m edges. For random walk \mathcal{W}_u starting at a vertex u of G , let $\mathcal{W}_u(t)$ be the vertex reached at step t . Let $P = P(G)$ be the matrix of transition probabilities of the walk and let $P_u^t(v) = \Pr(\mathcal{W}_u(t) = v)$. Assuming G is not bipartite, the random walk \mathcal{W}_u on G is ergodic with stationary distribution π . Here $\pi(v) = d(v)/(2m)$, where $d(v)$ the degree of vertex v . We often write $\pi(v)$ as π_v .

Let the eigenvalues of $P(G)$ be $\lambda_0 = 1 \geq \lambda_1 \geq \dots \geq \lambda_{n-1} > -1$, as P1 assumes G is not bipartite. Let $\lambda_{\max} = \max(\lambda_1, |\lambda_{n-1}|)$. The rate of convergence of the walk is given by

$$|P_u^t(x) - \pi_x| \leq (\pi_x/\pi_u)^{1/2} \lambda_{\max}^t. \quad (5)$$

For a proof of this, see for example, Lovasz [15].

The standard way to ensure that $\lambda_{\max} = \lambda_1$, is to make the chain *lazy* i.e. the walk only moves to a neighbour with probability $1/2$. Otherwise it stays where it is. If we do this until every vertex has been covered, then this will double the cover time. It is simplest therefore to assume that we keep the chain lazy for T_G steps. At this point it is mixed, and we can stop being lazy. All of our walks will be assumed to be lazy until the mixing time.

In this paper we consider the joint convergence of $k \geq 1$ independent random walks on a graph $G = (V_G, E_G)$. Let $H_k = (V_H, E_H)$ have vertex set $V_H = V^k$ and edge set $E_H = E^k$. The vertices \mathbf{v} of H_k consist of k -tuples $\mathbf{v} = (v_1, v_2, \dots, v_k)$ of vertices $v_i \in V(G), i = 1, \dots, k$, with repeats allowed. Two vertices \mathbf{v}, \mathbf{w} are adjacent if $\{v_1, w_1\}, \dots, \{v_k, w_k\}$ are edges of G . The graph H_k replaces k random walks $\mathcal{W}_{u_i}(t)$ on G with current positions \mathbf{v}_i and starting positions u_i by a single walk $\mathcal{W}_{\mathbf{u}}(t)$.

If $S \subseteq V_H$, then $\Gamma(S)$ is obtained from $H = H_k$ by contracting S to a single vertex $\gamma(S)$. All edges, including loops are retained. Thus $d_\Gamma(\gamma) = d_H(S)$, where d_F denotes vertex degree in graph F . Moreover Γ and H have the same total degree $(nr)^k$, and the degree of any vertex of Γ , except γ , is r^k .

Lemma 12. *Let G be typical. Let $F = G, H, \Gamma$. Let $k \geq 1$ be fixed. Let S be such that $d_H(S) \leq k^2 n^{k-1} r^k$. Let $\mathcal{W}_{u,F}$ be a random walk starting at $u \in V_F$. Let T_F be such that, for*

graph $F = (V_F, E_F)$, and $t \geq T_F$,

$$\max_{u, x \in V_F} |P_u^t(x) - \pi_x| \leq \frac{1}{n^3}.$$

Then for $k \leq n$,

$$T_G = O(\ln n), \quad T_H = O(\ln n) \quad \text{and} \quad T_\Gamma = O(k \ln n).$$

Proof

Case (i): Single random walk.

We choose $\epsilon = 0.1$ in P2 so that for $r \geq 3$ we have

$$\lambda_{\max} \leq 0.977. \quad (6)$$

Let

$$T_G = \frac{5 \ln n}{-\ln \lambda_{\max}}. \quad (7)$$

Using (5) we see that for $t \geq T_G$,

$$\max_{u, x \in V} |P_{u,G}^t(x) - \pi_x| \leq \max_u \sum_{x \in V} |P_{u,G}^t(x) - \pi_x| \leq n^{-3}. \quad (8)$$

Case (ii): k independent random walks.

Let $\mathcal{W}_{\mathbf{u}, H}^t$ be the corresponding random walk in H . As the k associated walks in G are independent, we have $P_{\mathbf{u}}^t(\mathbf{x}) = P_{u_1}^t(x_1) \dots P_{u_k}^t(x_k)$ and $\pi(\mathbf{x}) = \pi(x_1) \dots \pi(x_k)$. At step t , the total variation distance $\Delta_{\mathbf{u}}(t, H)$ of the walk is

$$\Delta_{\mathbf{u}}(t, H) = \frac{1}{2} \sum_{\mathbf{x} \in V_H} |P_{\mathbf{u}}^t(\mathbf{x}) - \pi(\mathbf{x})|.$$

To simplify notation let $P_i = P_{u_i}^t(x_i)$, where $\mathbf{u} = (u_1, \dots, u_k)$, and let $\pi_i = \pi(x_i)$. Then

$$\begin{aligned} |P_{\mathbf{u}}^t(\mathbf{x}) - \pi(\mathbf{x})| &\leq |P_1 \dots P_k - P_1 \dots P_{k-1} \pi_k| + |P_1 \dots P_{k-1} \pi_k - P_1 \dots P_{k-2} \pi_{k-1} \pi_k| + \dots \\ &\quad + |P_1 \dots P_\ell \pi_{\ell+1} \dots \pi_k - P_1 \dots P_{\ell-1} \pi_\ell \dots \pi_k| + \dots + |P_1 \pi_2 \dots \pi_k - \pi_1 \dots \pi_k|. \end{aligned}$$

It follows that

$$\Delta_{\mathbf{u}}(t, H) \leq \frac{k}{2} \max_{i=1, \dots, k} \left(\sum_{x \in V(G)} |P_{u_i}^t(x) - \pi(x)| \right) \leq k \Delta(t, G),$$

where $\Delta(t, G) = \max_{u \in V(G)} \Delta_u(t, G)$. If we choose

$$T_H = \frac{\ln k + 5 \ln n}{-\ln \lambda_{\max}},$$

then $\Delta(t, G) \leq 1/(kn^3)$ (see (8)) and so $\Delta(t, H) \leq 1/n^3$.

Case (iii): Random walk in Γ .

Let $\lambda_H = \lambda_{\max}(H)$, and let

$$\tau(\epsilon, H) = \min \{t : \Delta(t, H) \leq \epsilon \text{ for all } t' \geq t\},$$

then it is a result of [1] (see also [17]) that

$$\tau(\epsilon, H) \geq \frac{1}{2} \frac{\lambda_H}{1 - \lambda_H} \ln \frac{1}{2\epsilon}.$$

Let $\lambda_G = \lambda_{\max} \leq 0.977$ from (6). On the assumption that $k \leq n$ and using $\epsilon = n^{-3}$ and $\tau(\epsilon, H) \leq T_H$, we find that

$$\lambda_H \leq \frac{992}{1000}.$$

For a simple random walk on a graph G , the conductance Φ is given by

$$\Phi(G) = \min_{\substack{X \subset V \\ d(X) \leq m(G)}} \frac{e(X : \bar{X})}{d(X)},$$

where $m(G) = |E(G)|$ is the number of edges of G , $d(X)$ is the degree of set X , and $e(X : \bar{X})$ is the number of edges between X and $V \setminus X$. The second eigenvalue λ_1 of a reversible Markov chain satisfies

$$1 - 2\Phi \leq \lambda_1 \leq 1 - \frac{\Phi^2}{2}. \quad (9)$$

Using a lazy walk, we can assume that $\lambda_{\max} = \lambda_1$, we find that

$$\Phi(H) \geq 1/250. \quad (10)$$

The quantity we need is $\Phi(\Gamma)$, where Γ is the contraction of H . From the construction of Γ it follows that $\Phi(\Gamma) \geq \Phi(H)$; every set of vertices in V_Γ corresponds to a set in V_H , and edges are preserved on contraction. Thus for any starting position \mathbf{u} of a walk $\mathcal{W}_{\mathbf{u}}(\Gamma)$ we have, from (5) and (9) that, provided $t \geq T_\Gamma = 10^5 k \ln n$,

$$|P_{\mathbf{u}}^t(\mathbf{x}) - \pi(\mathbf{x})| \leq \left(\frac{d(\gamma)}{r^k} \right)^{1/2} e^{-t\Phi^2/2} \leq \frac{1}{n^3},$$

where $d(\gamma) \leq k^2 n^{k-1} r^k$. □

3.2 Generating function formulation

We will use the approach of [6] (Section 2), and [8] (Section 2). Let $d(t) = \max_{u,x \in V} |P_u^t(x) - \pi_x|$, and let T be such that, for $t \geq T$

$$\max_{u,x \in V} |P_u^t(x) - \pi_x| \leq n^{-3}. \quad (11)$$

It follows from e.g. Aldous and Fill [2] that $d(s+t) \leq 2d(s)d(t)$ and so for $\ell \geq 1$,

$$\max_{u,x \in V} |P_u^{(\ell T)}(x) - \pi_x| \leq \frac{2^\ell}{n^{3\ell}}. \quad (12)$$

Fix two vertices u, v . Let $h_t = \mathbf{Pr}(\mathcal{W}_u(t) = v)$ be the probability that the walk \mathcal{W}_u visits v at step t . Let

$$H(z) = \sum_{t=T}^{\infty} h_t z^t \quad (13)$$

generate h_t for $t \geq T$. This changes the definition of $H(z)$ from that used in [6], [7] where we included the coefficients h_0, h_1, \dots, h_{T-1} in the definition of $H(z)$ and gave rise to technical problems.

Next, considering the walk \mathcal{W}_v , starting at v , let $r_t = \mathbf{Pr}(\mathcal{W}_v(t) = v)$ be the probability that this walk returns to v at step $t = 0, 1, \dots$. Let

$$R(z) = \sum_{t=0}^{\infty} r_t z^t$$

generate r_t . Our definition of return involves $r_0 = 1$.

For $t \geq T$ let $f_t = f_t(u \rightarrow v)$ be the probability that the first visit of the walk \mathcal{W}_u to v in the period $[T, T+1, \dots]$ occurs at step t . Let

$$F(z) = \sum_{t=T}^{\infty} f_t z^t$$

generate f_t . Then we have

$$H(z) = F(z)R(z). \quad (14)$$

Finally, for $R(z)$ let

$$R_T(z) = \sum_{j=0}^{T-1} r_j z^j. \quad (15)$$

We remark that (14) is also valid for visits by \mathcal{W}_u to a set S of vertices. By contracting the set S of vertices of G to a single vertex $\gamma(S)$, we obtain a graph Γ and an equivalent relationship $\tilde{H}(z) = \tilde{F}(z)\tilde{R}(z)$. We next prove that if $t > T$ then the first visit probabilities $f_t(u \rightarrow S, G)$ to S in G , and $f_t(u \rightarrow \gamma(S), \Gamma)$ to $\gamma(S)$ in Γ , are asymptotically equal.

For a walk \mathcal{W}_u in G starting at u , and with t -step transition probabilities $P_u^t(v)$ to vertex v ,

$$f_t(u \rightarrow S, G) = \sum_{w \notin S} P_u^T(w) \phi_w(t - T, S, G),$$

where $\phi_w(\tau, S, G)$ is the probability that a first visit from w to S occurs at step $\tau > 0$.

For the graphs G and Γ , and $w \notin S$, the value of $\phi_w(\tau, S, G)$ equals $\phi_w(\tau, \gamma(S), \Gamma)$ as the degrees and neighbourhood structure of $G \setminus S$ and $\Gamma \setminus \{\gamma\}$ are identical. Thus provided we choose T to be a sufficiently large mixing time in both G and Γ we have that $P_u^T(w, G) \sim P_u^T(w, \Gamma) \sim \frac{1}{n}$ and thus

$$f_t(u \rightarrow S, G) = (1 + o(1))f_t(u \rightarrow \gamma(S), \Gamma).$$

3.3 First visit time lemma: Single vertex v

The following lemma should be viewed in the context that G is an n vertex graph which is part of a sequence of graphs with n growing to infinity. For a proof see [8], Section 2, Lemma 6 and Corollary 7.

Lemma 13. *Let T be a mixing time such that (11) holds. Let $R_T(z)$ be given by (15), let $R_v = R_T(1)$, and let*

$$p_v = \frac{\pi_v}{R_v(1 + O(T\pi_v))}. \quad (16)$$

Suppose the following conditions hold.

- (a) *For some constant $0 < \theta < 1$, we have $\min_{|z| \leq 1 + \lambda} |R_T(z)| \geq \theta$, where $\lambda = \frac{1}{KT}$ for some sufficiently large constant K .*
- (b) *$T^2\pi_v = o(1)$ and $T\pi_v = \Omega(n^{-2})$.*

Then for all $t \geq T$,

$$f_t(u \rightarrow v) = (1 + O(T\pi_v)) \frac{p_v}{(1 + p_v)^{t+1}} + o(e^{-\lambda t/2}). \quad (17)$$

Let v be a (possibly contracted) vertex, and for $t \geq T$ let $\mathbf{A}_t(v)$ be the event that \mathcal{W}_u does not visit v during steps $T, T+1, \dots, t$. Then

$$\Pr(\mathbf{A}_t(v)) = \sum_{s>t} f_s(u \rightarrow v),$$

and we have the following corollary.

Corollary 14.

$$\Pr(\mathbf{A}_t(v)) = \frac{(1 + O(T\pi_v))}{(1 + (1 + O(T\pi_v))\pi_v/R_v)^t} + o(e^{-t/KT}).$$

4 Oblivious particles

4.1 Cover time for k particles walking independently

Recall from Section 2 that the set of typical graphs \mathcal{G}'_r satisfies $|\mathcal{G}'_r| = (1 - o(1))|\mathcal{G}_r|$. Let

$$p = \frac{1}{n\theta_r} \left(1 + O\left(\frac{k \ln n}{n}\right)\right). \quad (18)$$

Lemma 15. *Let $T = T_{H_k}$ be a mixing time given by Lemma 12. Let $K > 0$ constant. The following properties hold for $G \in \mathcal{G}'_r$.*

(i) *Let $\mathbf{A}_t(v)$ be the event that a walk starting at a fixed vertex x does not visit v during steps T, \dots, t . If v is tree-like then*

$$\Pr(\mathbf{A}_t(v)) = (1 - p)^{(t+o(t))} + O(Te^{-t/(2KT)}), \quad (19)$$

Moreover for all $v \in V$,

$$\Pr(\mathbf{A}_t(v)) \leq (1 - p/2)^{(t+o(t))} + O(Te^{-t/(2KT)}). \quad (20)$$

(ii) *Let $\mathbf{A}_t(u, v)$ be the event that a walk starting at a fixed vertex x does not visit u or v during steps T, \dots, t . There exists a set $S \subseteq V$ of size $n^{1-o(1)}$ such that for all $u, v \in S$*

$$\Pr(\mathbf{A}_t(u, v)) = (1 + o(1))\Pr(\mathbf{A}_t(u))\Pr(\mathbf{A}_t(v)) + O(Te^{-t/(2KT)}).$$

(iii) *Let F_v be the time before a first visit to a tree-like vertex v after T , then $\mathbf{E}F_v \sim 1/p$.*

Proof of Lemma 15 and Theorem 2.

Lemma 15 is obtained from Corollary 14 with the help of [6]. For convenience, we use the definition of *nice* graphs in the notation of [6]. Nice graphs are only tree-like up to $\sigma = \lceil \ln \ln \ln n \rceil$, and have $O(r^{5\sigma})$ non tree-like vertices. Lemma 8 of [6] gives θ_r for tree-like vertices, and pruning arguments can show that $R_v \leq 2\theta_r$ for non-treelike vertices. Lemma 9 of [6] confirms the conditions of Lemma 13, and then Corollary 14 of this paper gives $\mathbf{A}_t(v)$ in Lemma 15. The simplification made in later papers (e.g. [8]) that we only consider first visits after T means we can ignore Lemma 10 of [6] (i.e. $c_{u,v} = (1 + o(1))$).

Let $\mathbf{A}_{k,t}(v)$ be the event that no walk visits vertex v in steps T, \dots, t . As the walks are independent, and $T = T_H$ is a mixing time for all k particles, we have that

$$\Pr(\mathbf{A}_{k,t}(v)) = [\Pr(\mathbf{A}_t(v))]^k, \tag{21}$$

and similarly for $\Pr(\mathbf{A}_{k,t}(u, v))$. The proof of Theorem 2 is a straightforward and simplified adaptation of the proof of Theorem 1 as given in Section 5 of [6].

Upper bound on cover time. Let $t_k^* = (\theta_r n/k) \ln n$. Choose $t_0 = (1 + \epsilon)t_k^*$ where $\epsilon \rightarrow 0$ sufficiently slowly. We apply (the first part of) the upper bound proof given in Section 5.1 of [6]. For a walk starting at u , let U_s be the number of unvisited vertices at step s . As in (5.4) of [6], the cover time C_u , starting from u , is bounded by

$$C_u \leq t + \sum_{s \geq t} \mathbf{E}U_s = t + \sum_{v \in V} \sum_{s \geq t} \Pr(\mathbf{A}_{k,s}(v)).$$

Using (21) and (19), the event that tree-like vertex v is unvisited at t is at most $e^{-kp(t+o(t))}$. For the $O(r^{5\sigma}) = o(\ln n)$ non tree-like vertices v , using (20), the probability that v is unvisited is at most $e^{-kp(t+o(t))/2}$. It follows that $C_G(k) \leq t_0 = (1 + o(1))t_k^*$.

Lower bound on cover time. Let $t_1 = (1 - \epsilon)t_k^*$ where $\epsilon \rightarrow 0$ sufficiently slowly. Substituting t_1 into the lower bound proof in Section 5.2 of [6], we find that there is a set of vertices S as given by Lemma 15 above, which **whp** are not all covered at time t_1 . The conclusion is that $C_G(k) \geq t_1$. \square

4.2 Comparison of search methods: Proof of Theorem 3

Suppose particles are of two types, either fixed at a vertex, or mobile. Assume there are k mobile particles, and ℓ fixed particles. When a mobile particle reaches a vertex concealing a fixed particle (an encounter), one particle destroys the other, so exactly one of the parameters ℓ, k decreases.

Lemma 16. *Suppose the ℓ fixed particles are positioned at randomly chosen vertices of the graph G . Then **whp** no mobile particle will visit any fixed particle during any of the (at most) $j = \max(k, \ell)$ mixing times T occurring during the search process.*

Proof Choose a definition of general position $\omega = \beta(\ln \ln n + \ln j)$ corresponding to (1), where $\beta \ln r \geq 6$. The assumption that the fixed particles are at randomly chosen vertices, corresponds **whp** to the following statements.

S1 : The fixed particles are in general position relative to each other, and to the starting positions of the mobile particles.

S2 : The fixed particles are at tree-like vertices of the graph G .

Given that S1 and S2 hold, the probability that any mobile particle reaches any fixed particle in any of j mixing times T , $T = O(k \ln n)$, is $O((j^3 T)/((r-1)^\omega)) = O(k^4 \ln n / (k \ln n)^{\beta \ln r})$. This follows by comparison with a biased random walk on a line of length ω (see (22) in the next section).

There are $n^{\epsilon c}$ non tree-like vertices (see Section 2), and $O(\ell r^\omega)$ vertices within distance ω of all fixed particles. Suppose an encounter occurs t steps ($t \geq T$) after the previous encounter (or the start). The probability that any other mobile particle is within distance ω of any fixed particle at that step is $O(r^\omega j^2/n)$. Thus the probability of some mobile particle (not involved in an encounter at that step) being within distance ω of any fixed particle, after any of the (at most j) encounters is $O((j^3 r^\omega)/n) = o(1)$. \square

We can now apply Lemma 15 to complete the proof of Theorem 3. Let A be the set of remaining fixed particles, and B the set of remaining mobile particles, and let $|A| = a$, $|B| = b$. Conditional on S2, the value of $p = p_A$ in (18) is $p \sim a/(n\theta_r)$.

The probability that none of the B particles visits A during T, \dots, t is

$$[\Pr(\mathbf{A}_t(A))]^b = (1-p)^{-b(t+o(t))} + O(bT e^{-t/(2KT)}).$$

Let $F_{A,B}$ be the first visit time, it follows that the expected visit time

$$\mathbf{E}(F_{A,B}) \sim \frac{1}{1-(1-p)^b} \sim \frac{\theta_r n}{ab}.$$

Thus choosing (i) $a = 1, \dots, \ell$ and $b = k$ or (ii) $a = k$ and $b = 1, \dots, \ell$ we have

$$\mathbf{E}(F_{k,\ell}^{(i)}) \sim \mathbf{E}(F_{k,\ell}^{(ii)}) \sim \frac{\theta_r n}{k} \sum_{i=1}^{\ell} \frac{1}{i} = \frac{\theta_r H_\ell}{k} n.$$

5 Probability two or more particles meet at a given step

When two particles meet, we say they are coincident at a vertex. The main task of this section (Lemma 17) is to estimate the probability that two coincident particles meet up again during

the mixing time. We prove in Lemma 18 that provided k is not too large (e.g. $k = o(\sqrt{n})$) the probability of more than a single pairwise coincidence occurring at any step is small.

We note the following result (see e.g. [13]), for a random walk on the line $= \{0, \dots, a\}$ with absorbing states $\{0, a\}$, and transition probabilities q, p, s for moves left, right and looping respectively. Starting at vertex z , the probability of absorption at the origin 0 is

$$\rho(z, a) = \frac{(q/p)^z - (q/p)^a}{1 - (q/p)^a} \leq \left(\frac{q}{p}\right)^z, \quad (22)$$

provided $q \leq p$. Similarly, for a walk starting at z on the half line $\{0, 1, \dots, \}$, with absorbing states $\{0, \infty\}$, the probability of absorption at the origin is

$$\rho(z) = (q/p)^z. \quad (23)$$

We first consider the case of a meeting between two of the particles.

Lemma 17. *Let G be a typical r -regular graph, and let v be a vertex of G , tree-like to depth $L_1 = \lfloor \epsilon_1 \log_r n \rfloor$. Suppose that at time zero, two independent random walks $(\mathcal{W}_1, \mathcal{W}_2)$ start from v . Let (x, y) denote the position of the particles at any step. Let $S = \{(u, u) : u \in V\}$. Let f_T be the probability of a first return to S within $T = T_\Gamma$ steps given that the walks leave v by different edges at time zero. Then*

$$f_T = \frac{1}{(r-1)^2} + O(n^{-\Omega(1)}).$$

Proof

We write $f_T = g_T + h_T$ where g_T is the probability of a first return to S up to time L_1 . Assume the walks leave v by distinct edges at time 0, let x_t, y_t denote the positions of the particles after t steps and let $Y_t = \text{dist}(x_t, y_t)$.

To estimate g_T we extend $N(v, L_1)$ to an infinite r -regular tree \mathcal{T} rooted at v . Let X_t be the distance between the equivalent pair of particles walking in \mathcal{T} . Thus provided $t \leq L_1$, we have that $Y_t = X_t$, and $g_T = \mathbf{Pr}(\exists t \in [1, L_1] : X_t = 0)$. The values of X_t are as follows: Initially $X_1 = 2$. If $X_t = 0$, then $\mathbf{Pr}(X_{t+1} = 0) = 1/r$, $\mathbf{Pr}(X_{t+1} = 2) = (r-1)/r$. If $X_t > 0$, then

$$X_t = \begin{cases} X_{t-1} - 2 & \text{with probability } q = \frac{1}{r^2} \\ X_{t-1} & \text{with probability } s = \frac{2(r-1)}{r^2} \\ X_{t-1} + 2 & \text{with probability } p = \left(\frac{r-1}{r}\right)^2. \end{cases} \quad (24)$$

Finally let Z_t be a walk on the even numbers $\{0, \pm 2, \pm 4, \dots\}$ of the infinite line, with $Z_1 = 2$, and with transition probabilities p, q, s . By coupling Z_t and X_t , we have inductively that $X_t \geq Z_t$. Note that

$$\mathbf{E}(Z_t - Z_{t-1}) = 2 - \frac{4}{r}. \quad (25)$$

Now let g_∞ denote the probability that $\{\exists t \geq 1 : X_t = 0\}$, i.e. the particles meet in \mathcal{T} . Equation (23) implies that

$$g_\infty = \frac{1}{(r-1)^2}. \quad (26)$$

Furthermore,

$$g_\infty = g_T + g'_T \quad (27)$$

where g'_T is the probability that $\{\exists t > L_1 : X_t = 0\}$, i.e. the particles meet after L_1 steps. If the particles meet, given a starting separation of $L_1/2$, then one of them must have traveled a distance at least $L_1/4$ towards the other. Using (22), we see that

$$\begin{aligned} g'_T &\leq \Pr(X_{L_1} \leq L_1/2) + \left(\frac{1}{(r-1)^2}\right)^{L_1/4} \\ &\leq O(n^{-\Omega(1)}) + \left(\frac{1}{(r-1)^2}\right)^{L_1/4}. \end{aligned} \quad (28)$$

As $\Pr(Z_{L_1} \leq L_1/2) \geq \Pr(X_{L_1} \leq L_1/2)$, the bound $O(n^{-\Omega(1)})$ follows from (25) and the Hoeffding inequality for the sums of bounded random variables Z_{L_1} . We use $\Omega(1)$ throughout this proof, instead of providing explicit constants, but remark that, as the right hand side of (25) is at least $2/3$ for any r , we can insert absolute constants independent of r .

It follows from (26), (27) and (28) that we can write

$$g_T = \frac{1}{(r-1)^2} + O(n^{-\Omega(1)}). \quad (29)$$

Recall that h_T is the probability of a first return to S after time L_1 . We next prove that $h_T = O(n^{-\Omega(1)})$. Let $h_T = h'_T + h''_T$ where

$$h'_T = \Pr(\text{The walks meet during steps } \{L_1, \dots, T\} \text{ and } Y_{L_1} > L_1/2), \quad (30)$$

and we know from (28) that

$$h''_T \leq \Pr(Y_{L_1} \leq L_1/2) = O(n^{-\Omega(1)}). \quad (31)$$

Let $\sigma \leq T$ be the step at which the particles first meet again, and let s be the *last* step before σ at which the distance between the particles is $L_1/2$ or more. Let $x = x_s$, $y = y_s$ denote the positions of the particles at time s .

Let $N = N(x, L_1)$, be the neighbourhood of depth at most L_1 centered at x . It follows from property P4 that there are most two paths xPy , $xP'y$ between x and y in N , both of length at least $L_1/2$. Let ρ_1, ρ_2 denote the particles at x, y respectively.

Suppose first there is a single path xPy . Consider the particle ρ_1 . Either ρ_1 moves at least $L_1/4$ down xPy , at some step $s < t \leq \sigma$, or if not, then ρ_2 must do; as we now show. Suppose first that both particles stay within N until they meet, then ρ_2 must move at least $L_1/4$ along xPy to meet ρ_1 . Suppose next that $s < t \leq \sigma$ is the first step at which the boundary of N is visited by either particle, and suppose this particle is ρ_1 (or both). As the distance between the particles is at most $L_1/2$, then ρ_2 has moved at least $L_1/4$ along xPy by that step.

Suppose next there are two paths $xPy, xP'y$, where $|P'| \geq |P| \geq L_1/2$. Let $xPy = xQaDbRy$ and let $xP'y = xQaD'bRy$, where $C = aDbD'a$ is a cycle, and possibly $x = a$ or $b = y$ (or both). Say a path section S is *long*, if $|S| \geq L_1/6$. Then either Q or R , or both of D, D' are long. Arguing as before, we see that wherever the particles meet, at least one particle must travel at least halfway down a long section, a distance of at least $L_1/12$.

Thus there is (at least) one of at most four fixed sub-paths of length $L_1/12$ in G which one of the particles has to traverse, an event of probability $O((1/(r-1))^{L_1/12})$, (see (22)). As there are at most T ways of choosing s , and T starting times for traversing the sub-path, an upper bound of $O(T^2(r-1)^{-L_1/12}) = O((k \ln n)^2/n^{\epsilon_1/24})$ follows. \square

At any step t of a walk $\mathcal{W}_{\mathbf{u}}(t)$ on H_k , the corresponding positions of the k particles is given by $(\mathcal{W}_{u_1}(t), \dots, \mathcal{W}_{u_k}(t))$. Let $S = \{(v_1, \dots, v_k) : \text{at least two } v_i \in V(G) \text{ are the same}\}$. We refer to S as the *diagonal* of H_k . Recall the definition $\gamma(S)$, the contraction of S , in which we replace S in H_k by a single vertex $\gamma(S)$ in Γ . The number of loop edges of $\gamma(S)$ is determined by the edge structure induced by S , i.e. $H_k[S]$. Once the walk \mathcal{W}_γ moves away from γ down a non-loop edge, the structure of Γ is exactly that of H_k , and we can maintain this description of the walk up to a first return to γ . Sometimes, if it is intuitively helpful, we will describe moves in and out of S in terms of the corresponding walk moves in G .

Using Lemma 17, we can calculate the expected number of returns to the diagonal S of H_k for k particles.

Lemma 18. *Let k be the number of particles walking on the underlying graph G . Let \mathcal{W}_γ be a random walk in Γ starting at $\gamma = \gamma(S)$. Let f^* denote the probability that \mathcal{W}_γ makes a first return to γ within T_Γ steps. Then*

$$f^* = \frac{1}{r-1} + O\left(\frac{k^2}{n^{\Omega(1)}}\right). \quad (32)$$

If $k \leq n^\epsilon$ for a small constant ϵ , then $f^* \sim \frac{1}{r-1}$.

Proof Every vertex in S has degree r^k in H_k , and the size of S is at most $\binom{k}{2}n^{k-1}$. On the other hand, there are at least $N_2(k)$ vertices of S with exactly two replicates, where

$$N_2(k) \geq \binom{k}{2}n^{k-1} - \binom{k}{3}n^{k-2} - \binom{k}{2}\binom{k-2}{2}n^{k-2} = \binom{k}{2}n^{k-1}\left(1 - O\left(\frac{k^2}{n}\right)\right).$$

Thus the total degree of $\gamma(S)$ is

$$d(\gamma) = \binom{k}{2} n^{k-1} r^k \left(1 - O\left(\frac{k^2}{n}\right) \right). \quad (33)$$

Similarly the loop degree of $\gamma(S)$ is

$$d_\ell(\gamma) = \binom{k}{2} n^{k-1} r^{k-1} \left(1 + O\left(\frac{k^2}{n}\right) \right).$$

The main case is where two coincident particles move to the same neighbour in G . Formally, there is an edge joining vertices $\mathbf{v} \in S$ and $\mathbf{w} \in N(\mathbf{v}) \cap S$ with $v_i = v_j$ and $w_i = w_j$. This gives loop degree $\binom{k}{2} n^{k-1} r^{k-1}$. The other case, contributing $O(k^2 n^{k-2} r^k)$ to the loop degree, is when a different pair of particles coincide at the next step, or a coincident particle moves to the same vertex as some other particle.

Returns to $\gamma(S)$ can be of several types. The simplest type is a loop return (type O), discussed above. For non-loop edges, we distinguish four cases. For simplicity we describe the returns in terms of the underlying walks on G . In the first case (type A), there were initially exactly two particles coincident at a tree-like vertex of G , which meet up again at some vertex. In the second case (type B), the coincident particles do not necessarily meet up again, but instead some other particles which were not initially coincident meet up. In the third case (type C), three or more particles coincide, either initially or finally. In the fourth case (type D), the coincident particles are initially at a non-tree-like vertex of G . Thus we can write f^* , the probability of a first return to γ within T_Γ steps, as

$$f^* = f_O + f_A + O(f_B + f_C + f_D).$$

For type A returns, let $d'(\gamma)$ count the non-loop edges of γ corresponding to the case where the pair of coincident particles are at a tree-like vertex v of G , but move apart at the next step. We do not place any restriction on the position or movement of the remaining $k - 2$ particles. Any coincidences among the other particles causing a loop or earlier return to γ are corrections of type B or type C.

As G is typical, the number of non tree-like vertices is n^{ϵ_C} (see (4)). In this case, there can be at most $O(k^2 n^{\epsilon_C} n^{k-2})$ coincidences at non tree-like vertices, and so

$$d'(\gamma) = d(\gamma)(1 - O(k^2 n^{\epsilon_C}/n)) - d_\ell(\gamma).$$

Thus

$$f_A = \frac{d'(\gamma)}{d(\gamma)} f_{T_\Gamma}. \quad (34)$$

The value of $f_{T_\Gamma} = 1/(r - 1)^2 + O(n^{-\Omega(1)})$ is given by Lemma 17, and $f_O = \frac{d_\ell(\gamma)}{d(\gamma)}$. Thus

$$f_O + f_A = \frac{1}{r - 1} + O(n^{-\Omega(1)}).$$

We can estimate f_B as follows: Of the vertices of S , at most $\nu_1 = \binom{k}{2}^2 r^{L_1} n^{k-2}$, have another pair of entries within the same neighbourhood of depth L_1 and at most $\nu_2 = \binom{k}{2} k r^{L_1} n^{k-2}$, have an entry within the neighbourhood of a coincident pair. For particles distance at least L_1 apart, the probability they coincide in T_Γ steps is $O(n^{-\Omega(1)})$, by the analysis of Lemma 17. Thus

$$f_B = O\left(\frac{(\nu_1 + \nu_2)r^k}{d(\gamma)} + \frac{k^2}{n^{\Omega(1)}}\right) = O\left(\frac{k^2}{n^{\Omega(1)}}\right).$$

Finally,

$$\begin{aligned} f_C &= O\left(\frac{k^3 n^{k-2} r^k}{d(\gamma)} + \frac{k^3 r^{2L_1} n^{k-3} r^k}{d(\gamma)} + \frac{k^2}{n^{\Omega(1)}}\right) = O\left(\frac{k^2}{n^{\Omega(1)}}\right) \\ f_D &\leq \frac{n^{\epsilon C} r_1^L k^2 n^{k-2}}{d(\gamma)} = O\left(\frac{k^2}{n^{\Omega(1)}}\right). \end{aligned}$$

The expression for f_C arises as follows: The term $O\left(\frac{k^3 n^{k-2} r^k}{d(\gamma)}\right)$ is the probability 3 or more particles are coincident initially. The term $O\left(\frac{k^3 r^{2L_1} n^{k-3} r^k}{d(\gamma)}\right)$ is the probability that two particles are initially in the L_1 -neighbourhood of a third. The term $O\left(\frac{k^2}{n^{\Omega(1)}}\right)$ is the probability that at least two particles, initially at distant greater than L_1 meet within time T . \square

The next lemma follows from Lemma 18 and a reworking of some parts of the proof of Lemma 17.

Lemma 19. *For typical graphs and k particles, the expected number of returns to γ in T_Γ steps is*

$$R_{\gamma(S)} = \theta_r + O((k \ln n)^{-\Omega(1)}). \quad (35)$$

If $k \leq n^\epsilon$ for a small constant ϵ , then $R_{\gamma(S)} \sim \theta_r$. \square

Proof We first upper bound $R_{\gamma(S)}$ as follows. Let $R^{(+)} = \frac{1}{1-f^*}$ where f^* is given by (32). Then $R_{\gamma(S)} \leq R^{(+)}$, and $R^{(+)} = \theta_r + O(k^2 n^{-\Omega(1)})$.

We next establish a lower bound for $R_{\gamma(S)}$. In the proof of Lemma 17, we defined $f_T = g_T + h_T$, where g_T was the probability of a type A first return to S up to time L_1 . Define a type $A(\omega)$ event as a type A event where the first return occurs up to time ω , where

$$\omega = \min(\alpha(\ln \ln n + \ln k), L_1),$$

and α is a sufficiently large constant. Redefine g_T as the probability of a type $A(\omega)$ first return. Reworking the proof of Lemma 17, we see that provided $\alpha(\ln \ln n + \ln k) < L_1$, the $O(n^{-\Omega(1)})$ error terms become $O((k \ln n)^{-\Omega(1)})$, where the $\Omega(1)$ term is arbitrarily large, but no other changes occur. Thus

$$g_T = \frac{1}{(r-1)^2} + O((k \ln n)^{-\Omega(1)}) + O(k^2 n^{-\Omega(1)}).$$

Let $g^* = f_O + g_T(d'(\gamma)/d(\gamma))$ be the first return probability for type O and type A(ω) events. Let $\nu = T/(2\omega)$, where $T = Ak \ln n$. Let $R^{(-)} = \sum_{i=0}^{\nu} (g^*)^i$, then

$$R^{(-)} = \frac{1}{1 - g^*} - \sum_{i > \nu} (g^*)^i = \theta_r - O\left(\frac{k^2}{(k \ln n)^{\Omega(1)}}\right) - O\left(\frac{k^2}{n^{\Omega(1)}}\right) - O\left(\frac{(1 + o(1))}{(r - 1)^{T/2\omega}}\right).$$

As $\nu\omega < T$, all excursions counted by $R^{(-)}$ will have been completed by step T , which establishes $R^{(-)}$ as a lower bound for $R_{\gamma(S)}$. \square

By an *occupied vertex*, we mean a vertex visited by at least one particle at that time step. The next lemma concerns what happens during the first mixing time, when the particles start from general position, and also the separation of the occupied vertices when a meeting occurs.

Lemma 20. *Let G be typical and $k \leq n^\epsilon$.*

(i) *Suppose two (or more) particles meet at time $t > T_\Gamma$. Let p_ℓ be the probability that the minimum separation between some pair of occupied vertices of G is less than ℓ . Then $p_\ell = O(k^2 r^\ell / n)$.*

(ii) *Suppose the particles start walking on G with minimum separation*

$$\ell = \alpha(\ln \ln n + \ln k).$$

Then

$$\Pr(\text{Some pair of particles meet during } T_\Gamma) = O((k \ln n)^{-\Omega(1)}).$$

Proof (i) At step $t > T_\Gamma$ the position of the (at most) k particles, including the coincident pair, is close to the stationary distribution. There are $O(r^\ell)$ vertices within distance ℓ of a given vertex of G , so the probability some pair are within distance ℓ at any coincidence is $O(k^3 r^\ell / n)$.

(ii) Assume $\alpha\epsilon < \epsilon_1$, so that $\ell < L_1$. We return to the analysis of Lemma 17 from (30) onwards, but use a separation between the particles of ℓ instead of $L_1/2$, in the subsequent argument. The probability a given pair of particles meet during T , is now $O(T^2/(r - 1)^{\ell/6})$. The probability some pair of particles meet during any of the k mixing times $T_\Gamma = O(k \ln n)$ is

$$O\left(\frac{k^3 (T_\Gamma)^2}{(r - 1)^{\ell/6}}\right) = O\left(\frac{k^5 (\ln n)^2}{(k \ln n)^{(\alpha \ln(r-1))/6}}\right) = O((k \ln n)^{-\Omega(1)}),$$

where the $\Omega(1)$ term can be made arbitrarily large by choosing α large and small enough $\epsilon > 0$. \square

6 Conditions of the first visit time lemma

We next check that the conditions of Lemma 13 hold with respect to the vertex γ of the graph Γ . Thus in this section, $T = T_\Gamma$, and $v = \gamma$. The conditions are:

- (a) $\min_{|z| \leq 1 + \lambda} |R_T(z)| \geq \theta$, for some constant $\theta > 0$, and $\lambda = 1/KT$ for suitably large K .
- (b) $T^2\pi_v = o(1)$ and $T\pi_v = \Omega(n^{-2})$.

Condition (a) follows from Lemma 21 below. Condition (b) is easily disposed of. Recall from Lemma 12, that $T = O(k \ln n)$. From Lemma 18, $\pi_\gamma = d(\gamma)/(nr)^k$ where $1 \leq d(\gamma) \leq k^2 n^{k-1} r^k$. Thus $T^2\pi_\gamma = o(1)$ provided e.g. $k \leq n^{1/5}$.

Lemma 21. *For $|z| \leq 1 + \lambda$, there exists a constant $\theta > 0$ such that $|R_T(z)| \geq \theta$.*

Proof Let $r_t = r_{t,A} + r_{t,B}$ where $r_{t,A}$ is the probability of a loop return ($t = 1$), or type A return to γ at time $t \leq L_1$. Thus $R_T(z) = R_A(z) + R_B(z)$ where $R_A(z) = \sum_{t=0}^{L_1} r_{t,A} z^t$.

The arguments in the proof of Lemma 17 show that

$$\phi = \sum_{t=L_1+1}^{T_\Gamma} r_t = O(n^{-\Omega(1)}),$$

and thus, for $|z| \leq 1 + \lambda$,

$$|R_B(z)| \leq (1 + \lambda)^{T_\Gamma} [\phi + T_\Gamma O(f_B + f_C + f_D)] = O(n^{-\Omega(1)}) + O\left(\frac{k^3 \ln n}{n^{\Omega(1)}}\right) = o(1).$$

As $|R_T(z)| \geq |R_A(z)| - |R_B(z)|$, we have $|R_T(z)| \geq |R_A(z)| - o(1)$.

As in Lemma 17, let Y_t be the distance between the particles during the first L_1 steps. For $1 \leq t \leq L_1$ let

$$b_t = \mathbf{Pr}(Y_t = 0, Y_1, \dots, Y_{t-1} > 0, \max_{i=1, \dots, t} Y_i < L_1),$$

be the probability that the walks first meet again at step t . Let $B(z) = \sum_{t=1}^{L_1} b_t z^t$, and let $A(z) = \sum_{t=0}^{\infty} a_t z^t = 1/(1 - B(z))$. As $A(z) = \sum_{j=0}^{\infty} (B(z))^j$, then $A(z)$ generates return probabilities for walks where all returns occurred within L_1 steps or less. It follows that, for $t \leq L_1$, $r_{t,A} = a_t$. Thus

$$\begin{aligned} R_A(z) &= \sum_{t=0}^{L_1} r_{t,A} z^t = \sum_{t=0}^{L_1} a_t z^t = \sum_{t=0}^{\infty} a_t z^t - \sum_{t=L_1+1}^{\infty} a_t z^t = \\ &= \sum_{t=0}^{\infty} a_t z^t - O(((1 + \lambda)\zeta)^t) = \sum_{t=0}^{\infty} a_t z^t - o(1), \end{aligned}$$

where the last step but one is explained as follows. By coupling Y_t and Z_t as in Lemma 17, and again applying the Hoeffding Lemma to Z_t , there is an absolute constant $0 < \zeta < 1$, such that $a_t = O(\zeta^t)$.

From Lemma 18, $B(1) = f_O + f_A - o(1) = \frac{1}{r-1} + o(1)$ and so for $|z| \leq 1 + \lambda$

$$B(|z|) \leq B(1 + \lambda) \leq B(1)(1 + \lambda)^T \leq \frac{e^{1/K}}{r-1} + o(1) < e^{1/K}.$$

So,

$$|R_A(z)| + o(1) = \left| \frac{1}{1 - B(z)} \right| \geq \frac{1}{1 + B(|z|)} \geq \frac{1}{1 + e^{1/K}}.$$

□

7 Results for interacting particles

Using Lemma 19 and (33) in Corollary 14 we obtain the following result.

Theorem 22. *Let $\mathbf{A}_k(t)$ be the event that a first meeting among the k particles after the mixing time T_Γ , occurs after step t . Let $p_k = \frac{\binom{k}{2}}{\theta_{r,n}}(1 + o(1))$. Then*

$$\Pr(\mathbf{A}_k(t)) = (1 + o(1))(1 - p_k)^t + O(T_\Gamma e^{-t/2KT_\Gamma}).$$

Let $\mathbf{B}_{s,k}(t)$ be the event that a first meeting between a given set of s particles and another set of k particles after the mixing time T_Γ , occurs after step t . Let $q_{sk} = \frac{sk}{\theta_{r,n}}(1 + o(1))$. Then

$$\Pr(\mathbf{B}_{s,k}(t)) = (1 + o(1))(1 - q_{sk})^t + O(T_\Gamma e^{-t/2KT_\Gamma}).$$

Let M_k (resp. $M_{s,k}$) be the time at which a first meeting of the particles occurs.

Corollary 23. $\mathbf{E}(M_k) = (1 + o(1))/p_k$, $\mathbf{E}(M_{s,k}) = (1 + o(1))/q_k$.

The proof follows from $\mathbf{E}(M_k) = \sum_{t \geq T} \Pr(\mathbf{A}_k(t))$ and $p_k T_\Gamma = o(1)$ etc. □

The proofs of Theorem 4, and Theorem 7 now follow from Lemma 20 and Corollary 23.

7.1 Expected broadcast time: Theorem 5

We allow the particles which met, time $T = T_G$ to re-mix after an encounter. This happens $k - 1$ times. Recall that $T_\Gamma = O(kT)$. From Lemma 20, the event that some particles meet during one of these mixing times has probability $o(1)$.

Assuming this does not happen, the expected time $\mathbf{E}(B_k)$, for a given agent to broadcast to all other agents is,

$$\begin{aligned} \mathbf{E}(B_k) &= O(kT) + \sum_{s=1}^{k-1} \frac{(1 + o(1))}{q_{s(k-s)}} \\ &\sim n\theta_r \sum_{s=1}^{k-1} \frac{1}{s(k-s)}, \\ &= \frac{2n\theta_r}{k} H_{k-1}, \end{aligned}$$

where H_k is the k -th harmonic number. On the assumption that k is large,

$$\mathbf{E}(B_k) \sim \frac{2\theta_r \ln k}{k} n.$$

Let $\omega = A \ln k$. Corollary 24 (and hence Corollary 6) follows from putting $t = (1 + o(1))\omega n\theta_r / (s(k-s))$ in Theorem 22, at step s of the broadcasting process.

Corollary 24. *Let $\omega = A \ln k$.*

$$\Pr(B_k > \omega \mathbf{E}B_k) = O(k^{-A+1}).$$

□

7.2 Expected time to coalescence: Theorem 8

Let S_k be the time for all the particles to coalesce, when there are originally k sticky particles walking in the graph. Then,

$$\begin{aligned} \mathbf{E}(S_k) &= O(kT) + \sum_{s=2}^k \frac{(1 + o(1))}{p_s} \\ &\sim n\theta_r \sum_{s=2}^k \frac{2}{s(s-1)}. \end{aligned} \tag{36}$$

Noting that $\sum_{s=2}^{\infty} 1/(s(s-1)) = 1$ we see that for large k ,

$$\mathbf{E}(S_k) \sim 2\theta_r n.$$

□

7.3 Expected time to extinction; explosive particles: Theorem 9

Let A_k be the time to extinction, when there are originally $k = 2\ell$ explosive particles walking in the graph. Then

$$\begin{aligned} \mathbf{E}(A_k) &= O(kT) + \sum_{s=1}^{\ell} \frac{(1 + o(1))}{p_{2s}} \\ &\sim n\theta_r \sum_{s=1}^{\ell} \frac{2}{2s(2s-1)}. \end{aligned} \tag{37}$$

Noting that $\sum_{s=1}^{\infty} 1/(2s(2s-1)) = \ln 2$ we see that for large k ,

$$\mathbf{E}(A_k) \sim 2\theta_r \ln 2 n.$$

□

8 The voter model: Theorem 10.

We briefly summarize Section 3 of Chapter 14 of [2], *Coalescing random walks and the voter model*. This will introduce the topic, and establish certain results we require.

In the *voter model* each vertex initially has a different opinion (vertex i has opinion i say). As time passes, opinions change according to the following rule. At each step, each vertex i contacts a random neighbour j , and changes (if necessary) its opinion to the current opinion of vertex j .

The total number of existing opinions can only decrease with time, and at some random step $t = C_{vm}$ a consensus opinion, k say, emerges. The reverse process, using the same edges and working backwards from t constitutes a *coalescing random walk process* of length C_{crw} , where k is the vertex at which the particles finally coalesce at step t . This symmetry, which we establish formally below, means that the following events are equivalent.

In the voter model: By time t_0 everyone's opinion is the opinion initially held by person k .

For the coalescing random walk process: All particles have coalesced by time t_0 , and the cluster is at k at time t_0 . Thus $\Pr(C_{\mathbf{vm}} \leq t_0) = \Pr(C_{\mathbf{crw}} \leq t_0)$, and these two times, C , have the same distribution and hence the same expected value.

We briefly explain the equivalence between these processes under time reversal. Suppose the coalescence process completes at the start of step t with the particle on vertex k . We can generalize the coalescence process as follows. At each step $s = 0, 1, \dots, t$ each vertex i chooses a random neighbour $j(i) \in N(i)$. This defines a digraph with levels $0, 1, \dots, t$, each with vertex set V , and the edges generated at step $s = 0, 1, \dots, t-1$ between the levels. Initially at level 0, particle i is at vertex i , $i = 1, \dots, n$. The particles move down the edges chosen at each step. If there is no particle at the vertex (the vertex is empty) then nothing moves. Picking any vertex v at level t , we can trace back any paths directed towards that vertex. All such paths originate at a vertex of in-degree 0. There are exactly n paths, with terminal vertex k at level t , which trace back to a vertex at level 0 containing its particle. All other paths trace back to an empty vertex. This may include some paths terminating at vertex k at level t . Finally, there are no components with a terminal vertex rooted at a level $s < t$, as each vertex at this level has out-degree one. Thus all paths in the digraph can be found by tracing back from level t .

We now start at level t and make the first step (step 0) of the voter process, by considering the edges (i, j) directed from level $t-1$ to t , and then the next step (step 1) from $t-2$ to $t-1$ etc. Vertex i takes the opinion of vertex j at the level below. Opinion v becomes extinct if all paths tracing back from v at level t have in-degree zero before level 0. At level 0, everybody has opinion k and the voting process is complete.

Among the results given in [2] about these processes are the following:

Theorem 25. [2] *Let G be r -regular.*

- (a) *If G is s -edge connected then $\mathbf{EC} \leq \frac{rn^2}{4s}$,*
- (b) $\mathbf{EC} \leq e(\ln n + 2) \max_{i,j} \mathbf{E}_i T_j$,
- (c) *For the complete graph K_n , $\mathbf{EC} \sim n$.*

Section 14.3 of [2] also gives two open problems:

- Prove that for a sequence of vertex transitive graphs with $\tau_2/\tau_0 \rightarrow 0$, we have $\mathbf{EC} \sim \tau_0$.
- Prove there exists an absolute constant K such that on any graph $\mathbf{EC} \leq K \max_{v,w} \mathbf{E}_v T_w$.

Here $\tau_2 = 1/(1 - \lambda_2)$ is the relaxation time, and $\tau_0 = \sum_j \pi_j \mathbf{E}_k T_j$ is an average hitting time (random target lemma).

Theorem 10 answers these questions affirmatively for the restricted class of typical r -regular graphs. For this class, $\tau_0 \sim \theta_r n$, $\tau_2 = \Theta(1)$ and $\max_{v,w} \mathbf{E}_v T_w \sim \theta_r n$. The Ramanujan graphs of [16] are vertex transitive and typical.

Note. The process considered in [2] is about twice as fast as the one we consider here. In their model a collision occurs if a particle moves down an edge and the vertex at other end of the edge is already occupied by a particle. Thus if all the vertices were occupied, the expected number of collisions in unit time is $(rn/2) \times 2 \times (1/r) = n$. In the discrete model, both particles have to move from distinct neighbours of a given vertex, to the vertex itself, to collide. This gives $n \times \binom{r}{2} \times (1/r^2) = n(1 - 1/r)/2$ expected collisions. For large r , the results of [2] should be multiplied by 2 for comparison with our model. For example, Theorem 25(c) gives $\mathbf{E}C \sim n$ for the complete graph K_n , and Theorem 10 gives $\mathbf{E}C \sim 2n$ as $r \rightarrow \infty$.

8.1 Proof of Theorem 10

Lower bound. Pick any $k \rightarrow \infty$ particles in general position from among the n particles, initially located one at each vertex. Let S_k be the time for these k particles to coalesce. By Theorem 8, $\mathbf{E}(S_k) \sim 2\theta_r n$, as $k \rightarrow \infty$.

Upper bound. Consider all k -sets of particles, where $k = \ln^3 n$. We prove **whp** there is no k -set which has not had a meeting by time $t^* = n/\ln n$, and thus by t^* at most k particles remain.

Let $\mathcal{P}(k, \mathbf{v})$ be the set of particles starting from vertices $\mathbf{v} = (v_1, \dots, v_k)$, not necessarily in general position. Either there has been a meeting during the mixing time T_Γ , or if not, the results of Theorem 22 apply for all $t \geq T_\Gamma$.

Suppose no meeting has occurred. Let

$$\begin{aligned} \bar{\rho}_k &= \mathbf{Pr}(\text{No meeting among the } \mathcal{P}(k) \text{ has occurred by } t^*) \\ &\leq (1 + o(1))(1 - p_k)^{t^*} + O(T_\Gamma e^{-t^*/(2KT_\Gamma)}) \end{aligned}$$

where $p_k \sim \binom{k}{2}/(\theta_r n)$ and $T_\Gamma = O(k \ln n)$. Thus

$$\bar{\rho}_k = O\left(e^{-(1+o(1))\frac{\ln^5 n}{2\theta_r}} + \ln^4 n e^{-O(\frac{n}{\ln^5 n})}\right).$$

Hence

$$\begin{aligned} \Pr(\exists \text{ a } k\text{-set } \mathcal{P}(k) \text{ having no meeting by } t^*) &\leq \binom{n}{k} \bar{\rho}_k \\ &= O\left(e^{-1/(3\theta_r)\ln^5 n + \ln^4 n}\right) \\ &= O\left(n^{-(\ln^4 n)/7}\right). \end{aligned}$$

Thus an upper bound on the expected time for all particles to coalesce is at most

$$t^* + (1 + o(1))(2\theta_r n) + O\left(n^{-(\ln^4 n)/7} n^3\right) \sim 2\theta_r n.$$

The second term, $2\theta_r n$, is an upper bound for the coalescence time of the (at most) k particles remaining after t^* . The $O(n^3)$ in the last term is $\mathbf{EC} \leq rn^2/(4s)$ from Theorem 25(a), restarting the process at t^* under the assumption that no meeting among some set of $k = \ln^3 n$ particles has occurred by t^* .

9 Conclusions

We have extended the results of [6] to deal with multiple random walks. In particular we have shown once again the usefulness of Lemma 13 in the context of random walks on expander graphs of high girth. Further applications of this lemma will be found in [9] where we allow more extensive interaction between particles and the underlying graph.

Acknowledgement

We thank the anonymous referees for their insightful comments and suggestions about this paper.

References

- [1] D. Aldous. Some inequalities for reversible Markov chains. *J. London Math. Soc.* (2) 25, (1982) 564-576.
- [2] D. Aldous and J. Fill. *Reversible Markov Chains and Random Walks on Graphs*, <http://stat-www.berkeley.edu/pub/users/aldous/RWG/book.html>.

- [3] R. Aleliunas, R.M. Karp, R.J. Lipton, L. Lovász and C. Rackoff. Random Walks, Universal Traversal Sequences, and the Complexity of Maze Problems. *Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science* (1979) 218-223.
- [4] N. Alon, C. Avin, M. Kouchý, G. Kozma, Z. Lotker, M. Tuttle. Many random walks are faster than one. *Proceedings of the twentieth annual symposium on parallelism in algorithms and architectures*, (2008) 119–128.
- [5] A. Broder, A. Karlin, A. Raghavan and E. Upfal. Trading space for time in undirected $s - t$ connectivity. *Proc. ACM Symp. Theory of Computing* (1989) 543-549.
- [6] C. Cooper and A. M. Frieze. The cover time of random regular graphs, *SIAM Journal on Discrete Mathematics*, 18, 728-740 (2005).
- [7] C. Cooper and A. M. Frieze. The cover time of the preferential attachment graph. *Journal of Combinatorial Theory Series B*, 97(2), 269-290 (2007).
- [8] C. Cooper and A. M. Frieze. The cover time of the giant component of a random graph. *Random Structures and Algorithms* 32, 401-439 J. Wiley 2008.
- [9] C. Cooper, A. M. Frieze and T. Radzik, Competition and predation in multiple random walks. (in preparation)
- [10] D. Coppersmith, P. Tetali and P. Winkler. Collisions among random walks on a graph. *SIAM Journal on Discrete Mathematics*, 6, 363-374 (1993).
- [11] U. Feige. A tight upper bound for the cover time of random walks on graphs, *Random Structures and Algorithms*, 6 (1995) 51-54.
- [12] U. Feige. A tight lower bound for the cover time of random walks on graphs, *Random Structures and Algorithms*, 6 (1995) 433-438.
- [13] W. Feller, *An Introduction to Probability Theory, Volume I*, (Second edition) Wiley (1960).
- [14] J. Friedman. A proof of Alon's second eigenvalue conjecture. *Proceedings of the thirty-fifth annual ACM symposium on theory of computing* (2003) 720–724.
- [15] L. Lovász. Random walks on graphs: a survey. Bolyai Society Mathematical Studies, *Combinatorics, Paul Erdős is Eighty* (Volume 2) Keszthely (Hungary), (1993) 1-46.
- [16] A. Lubotzky, R. Phillips, P. Sarnak. Ramanujan graphs. *Combinatorica* 8, (1988) 261-277.
- [17] A. Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing* 1 (1992) 351-370.
- [18] P. Tetali and P. Winkler. Simultaneous reversible Markov chains. *Combinatorics, Paul Erdős Is Eighty* (Vol. 1), Keszthely, Hungary (1993), 433-451.