Multiple random walks and interacting particle systems

Colin Cooper* Alan Frieze† Tomasz Radzik*

Draft

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 our 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 $k, n \to \infty$.

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 $k \to \infty$.

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

*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.

opinion to that of a random neighbour. The expected time for a unique opinion to emerge is the expected time for all the particles to coalesce, which is asymptotic to $2\theta_r n$.

Combining results from the predator-prey and multiple random walk models allows us to compare expected detection time in the following cops and robbers 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

A random walk is a simple process in which particles or messages move randomly from vertex to vertex in a graph. Random walks are an established method of graph exploration and connectivity testing with limited memory. If we consider the case where several random walks occur simultaneously, many questions and different types of application arise: In graph exploration, to what extent do the multiple random walks speed up the process? If the walks can interact how effective is communication, such as broadcasting, between the walks? If there are two different types of particles making walks, then we can model predator-prey processes (cops and robbers). In the case where each vertex of the graph initiates a random walk, there are applications in distributed data collection, gossiping and voting.

In this paper, we study properties of multiple random walks on a graph under various assumptions of interaction between the particles. To give detailed results for comparison purposes, we make the analysis for random regular graphs. The technique used is not specific to random graphs, nor to regular graphs. It can be applied to many graphs with at least reasonable edge expansion, and whose local edge structure around vertices has enough symmetry to be describable in a precise sense.

For brevity we restrict our proofs to random r-regular graphs, where $r \geq 3$. Our results also apply to many non-random regular graphs e.g. Lubotsky-Phillips-Sarnak type expanders and, with minor alterations, to many regular graphs where $r \to \infty$ slowly, e.g. the hypercube on $n = 2^r$ vertices. In the case where $r \to \infty$, the parameter θ_r used throughout this paper takes the value 1. To make our analysis, we reduce the multiple random walks to a single random walk on a suitably defined product graph, to which we apply the technique of [6]. The main difficulty was to analyze the structure of the product graph, in particular the pair-wise interaction of the walks. Once established, the reduction approach allows us to address a wide range of problems, some of which we now describe.

Suppose there are $k \geq 1$ particles, each making a simple random walk on a 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.

The paper gives precise results for the following topics on random regular graphs:

1. **Multiple walks**. For k particles walking independently, we establish the cover time $C_G(k)$ of G.
2. **Talkative particles**. For k particles walking independently, which communicate on meeting at a vertex, we give the expected time to broadcast a message.
3. **Predator-Prey**. For k predator and ℓ prey particles walking independently, we give the expected time to extinction of the prey particles, when predators eat prey particles on meeting at a vertex.
4. **Annihilating particles**. For $k = 2\ell$ particles walking independently, which destroy each other (pairwise) on meeting at a vertex, we give the expected time to extinction.
5. **Coalescing particles**. For k particles walking independently, which coalesce on meeting at a vertex, we give the expected time to coalesce to a single particle. In the case where a walk starts at each vertex, we extend the analysis to a distributed model of voting, the **Voter model**.

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 approach, and it is natural to try to speed this up by parallel searching. Similarly, properties of communication between particles moving in a network, such as broadcasting and gossiping, are natural questions. In this context, the predator-prey model could represent interaction between server and client particles, where each client needs to attach to a server. 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. An application of this, is with the predators as cops and the prey as robbers.

Coalescing and annihilating particle systems are part of the classical theory of interacting particles; and our paper makes a new contribution to this area. A system of coalescing particles where initially one particle is located at each vertex, is dual 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 can be shown that the distribution of time taken for a unique opinion to emerge, is the same as the distribution of time for all the particles to coalesce. By establishing the expected coalescence time, we obtain the expected time to complete voting in the voter model.

Most known results for interacting particle systems are for the infinite d-dimensional grid Z^d (see e.g. Liggett [13]). As far as we know, the results presented here are the first which give precise answers for finite graphs, especially for the Voter model (Theorem 8). For an informative discussion on models of interacting particle systems see Chapter 14 of Aldous and Fill [2].

If one step of a random walk corresponds to a vertex forwarding a message to a random neighbour, and vertices combine messages they receive, the coalescing particle system gives the time taken to combine all messages. Another application is to calculate the average value of a vertex based function $f(v),\ v \in V$; for example temperature. To do this each vertex initiates a message, and the messages then perform a coalescing random walk. The voter model allows the distributed nomination of a central vertex, to e.g. relay messages. This can be used to implement the leader election problem in a distributed network.

Results: Oblivious particles

A standard measure of efficiency of graph exploration by a single random walk, is the cover time, which is defined as follows: Let $G = (V, E)$ be a connected graph, with $|V| = n$ vertices and $|E| = m$ edges. For a given starting vertex $v \in V$ let C_v be the expected time taken for a simple random walk 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 $C_G \le 2m(n-1)$. It was shown by Feige [9], [10], that for any connected graph G, the cover time satisfies $(1-o(1))n \ln n \le C_G \le (1+o(1))\frac{4}{27}n^3$.

For many classes of graphs the cover time can be found precisely. For random regular graphs, the following result was proved in [6].

Theorem 1 *Let $\mathcal{G}_r$ denote the space of r-regular graphs with vertex set $V = \{1, 2, \ldots, n\}$ and the uniform measure. Let $r \geq 3$ be constant, and let $\theta_r = \frac{r-1}{r-2}$. If G is chosen randomly from $\mathcal{G}_r$, then* **whp**

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

The results given are asymptotic in n, the size of the vertex set. Thus $A_n \sim B_n$ means that $\lim_{n\to\infty} A_n/B_n = 1$, and **whp** (with high probability) means with probability tending to 1 as $n \to \infty$.

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

Improving s-t connectivity testing by using k independent random walks 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 $C_k \leq (25 + o(1))n^2 \ln^2 n/k^2$. This bound 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 result. For comparison with the case where both predator and prey move, we have included the result of Theorem 5 below, 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.

(i) Suppose there are k mobile predator particles walking randomly, and ℓ prey particles fixed at randomly chosen vertices of the graph. 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 the graph. 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 5. 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.

The proofs of the above theorems are given in Section 7 of the appendix.

Results: Interacting particles

Consider a pair of random walks, starting at vertices u and v. Let $M(u,v)$ be the number of steps before the walks first meet at a vertex. Clearly if $u = v$, then $M(u,v) = 0$. We say the walks are in *general position*, if the starting vertices of the walks are not too near. For our definition of general position $(v_1, v_2, ..., v_k)$, we choose a pairwise separation $d(v_i, v_j) \geq \omega = \omega(k,n)$ between particles, where

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

For the results given in this section, we assume that $r \geq 3$ is constant, that G is chosen randomly from $\mathcal{G}_r$, and that the results hold **whp** over our choice of G.

We first consider problems of passing information between particles. We assume that particles can only communicate when they meet at a vertex. 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).

Theorem 4 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 $k \to \infty$, $\mathbf{E}(B_k) \sim \frac{2\theta_r \ln k}{k} n$.

An alternative and less efficient way to pass on a message, is for the the originating agent to tell it directly to all other agents. 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 4, with $\mathbf{E}(D_{1,k-1})$ of Theorem 5 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. One variant of interacting particles is the predator-prey model, in which both types of particles make independent random walks. If a predator encounters prey on a vertex it eats them.

Theorem 5 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 the prey. Then

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

A variant of predator-prey is interacting sticky particles, in which all particles are predatorial, and only one particle survives an encounter.

Theorem 6 Coalescence time: sticky particles.
Let $k \leq n^\epsilon$ for a sufficiently small positive constant ϵ. Let S_k be the time to

coalesce, when there are originally k sticky particles walking randomly, starting from general position. Then,

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

and the expected value for k constant is given by (7).

As a twist on predator-prey, we consider particles which destroy each other (pairwise) on meeting at a vertex.

Theorem 7 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 D_k be the time to extinction, then

$$\lim_{k\to\infty} \mathbf{E}(D_k) \sim 2\theta_r \ln 2n,$$

and the expected value for k constant is given by (8).

The proofs of Theorems 4-7 are given in Section 5.

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_{vm} be the number of steps needed for a unique opinion to emerge in the voter model and let C_{crw} be the number of steps to complete a coalescing random walk when one particle starts at each vertex. By a duality argument these random variables have the same expected value.

Theorem 8 Voter model. whp *for random r-regular graphs,*

$$\mathbf{E}C_{vm} = \mathbf{E}C_{crw} \sim 2\theta_r n.$$

The proof of Theorem 8 for the Voter model is given in Section 10 of the appendix.

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 11 are met, then the probability that a vertex v is unvisited at step $T, ..., t$ 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 has not visited the diagonal (set of vertices $\boldsymbol{v} = (v_1, ..., v_k)$ 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.

2 Typical r-regular graphs

We say an r-regular graph G is *typical* if it has the properties **P1-P4** listed below: Let $\epsilon_1 > 0$ be a sufficiently small constant. Let a cycle C be *small* if $|C| \leq L_1$, where

$$L_1 = \lfloor \epsilon_1 \log_r n \rfloor. \tag{2}$$

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.

The results of this paper are valid for any typical r-regular graph G, and indeed most r-regular graphs have this property.

Theorem 9 *Let $\mathcal{G}'_r \subseteq \mathcal{G}_r$ be the set of typical r-regular graphs. Then $|\mathcal{G}'| \sim |\mathcal{G}_r|$.*

P2 is a deep result of Friedman [12]. The other properties are easy to check. Note that P3 implies that most vertices of a typical r-regular graph are tree-like.

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 $P_u^{(t)}(v) = \mathbf{Pr}(\mathcal{W}_u(t) = v)$. 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 $\lambda_0 = 1 \geq \lambda_1 \geq \cdots \geq \lambda_{n-1} \geq -1$, and let $\lambda_{\max} = \max(\lambda_1, |\lambda_{n-1}|)$. The rate of convergence of the walk is given by

$$|P_u^{(t)}(x) - \pi_x| \leq (\pi_x/\pi_u)^{1/2}\lambda_{\max}^t. \tag{3}$$

For a proof of this, see for example, Lovasz [14].

In this paper we consider the joint convergence of k independent random walks on a graph $G = (V_G, E_G)$. It is convenient to use the following notation. Let $H_k = (V_H, E_H)$ have vertex set $V_H = V^k$ and edge set $E_H = E^k$. If $S \subseteq V_H$, then $\Gamma(S)$ is obtained from H 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.

Let $k \geq 1$ be fixed, and let $H = H_k$. For $F = G, H, \Gamma$ let $\mathcal{W}_{u,F}$ be a random walk starting at $u \in V_F$. Thus $\mathcal{W}_{u,G}$ is a single random walk, and $\mathcal{W}_{u,H}$ corresponds to k independent walks in G.

Lemma 10 *Let G be typical. Let $F = G, H, \Gamma$. Let S be such that $d_H(S) \leq k^2 n^{k-1} r^k$. Let T_F be such that, for graph $F = (V_F, E_F)$, and $t \geq T_F$, the walk $\mathcal{W}_{u,F}$ satisfies*

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

for any $u \in V_F$. Then for $k \leq n$,

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

The proof, which is straightforward but technical is given in Section 6 of the appendix.

3.2 First visit time lemma: Single vertex v

Considering a 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 $t = 0, 1, \ldots$. Let

$$R_T(z) = \sum_{j=0}^{T-1} r_j z^j, \tag{4}$$

generate returns during steps $t = 0, 1, ..., T_1$. Our definition of return includes $r_0 = 1$.

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].

Lemma 11 *Let T be a mixing time such that*

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

Let $R_T(z)$ be given by (4)*, let $R_v = R_T(1)$, and let*

$$p_v = \frac{\pi_v}{R_v(1 + O(T\pi_v))}. \tag{5}$$

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* $T\pi_v = \Omega(n^{-2})$.

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

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

4 Interacting particles: Applying the first visit time lemma

Recall the definition of H_k consisting of k copies of G, and let $S = \{(v_1, ..., v_k) :$ at least two v_i are the same$\}$. The particles making random walks are at the components of the vector corresponding to the vertex in question. Thus S is the set of particle positions in which at least two particles coincide at a given step. As before, let $\gamma(S)$ be the contraction of S to a single vertex, and let $\Gamma(S)$ be H_k with S contracted.

In order to usefully apply Lemma 11, and estimate the first visit probability of γ (and hence S) we need to establish three things.
(i) The value of R_γ, the expected number of returns to the diagonal S of H_k for k particles, and of $\pi(\gamma)$ the stationary distribution of γ in Γ.

(ii) The conditions of Lemma 11 hold with respect to the vertex γ of the graph Γ. This is straightforward and given in Section 9 of the appendix.
(iii) The probability that any particles meet during the mixing time T_Γ.

These points are formally summarized in Lemmas 12-13 below, whose proofs are given in the appendix.

Lemma 12 *For typical graphs and k particles, the expected number of returns to γ in T_Γ steps is*

$$R_{\gamma(S)} = \theta_r + O\left(\frac{k^2}{n^{\Omega(1)}}\right). \tag{6}$$

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

Lemma 13 *If $k \leq n^\epsilon$ then the conditions of Lemma 11 hold with respect to the vertex γ of a typical graph Γ.*

From (5) with $v = \gamma$, and Lemma 12 we have

$$p_\gamma = \frac{\pi_\gamma}{\theta_r(1 + O(n^{-\Omega(1)}))}.$$

It follows from Lemma 19 of the appendix, that the value of π_γ corresponding to a meeting among k particles is $\pi_\gamma = (1 + o(1))\binom{k}{2}/n$, and for a meeting between a given set of s particles and another set of k particles is $\pi_\gamma = (1 + o(1))sk/n$. Applying this to Corollary 11 we have the following theorem.

Theorem 14 *Let $\boldsymbol{A}_k(t)$ be the event that a first meeting among the k particles after the mixing time T_Γ, occurs after step t. Let $p_k = \frac{\theta_r}{\binom{k}{2}n}(1 + O(n^{-\Omega(1)}))$. Then*

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

Let $\boldsymbol{B}_{s,k}(t)$ 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 $q_{sk} = \frac{\theta_r}{skn}(1 + O(n^{-\Omega(1)}))$. Then

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

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 15 *For typical graphs G and $k \leq n^\epsilon$ particles,*
(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 is less than L. Then $p_L = O(k^2 r^L / n)$.
(ii) Suppose the particles start walking on G with minimum separation at least $\alpha(\max\{\ln\ln n, \ln k\})$. Then

$$\mathbf{Pr}(\textit{Some pair of particles meet during } T_\Gamma) = o(1).$$

From Lemma 15, we see that **whp** particles starting from general position do not meet during the mixing time T_Γ. When some set of particles do coincide after the mixing time, the remaining particles are in general position **whp**.

Corollary 16 *Let M_k (resp. $M_{s,k}$) be the time at which a first meeting of the particles occurs, then $\mathbf{E}(M_k) = (1 + o(1))/p_k$ (resp. $\mathbf{E}(M_{s,k}) = (1 + o(1))/q_{s,k}$).*

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

5 Results for interacting particles

After an encounter, we allow the remaining particles time $T = T_G$ to re-mix . In any of Theorem 4-7 the total number of particle interactions k^2. Recall that $T_\Gamma = O(kT)$. From Lemma 15, the event that some particles meet during one of these kT_Γ mixing times has probability $O(k^3 T / n^{\Omega(1)}) = o(1)$ (by assumption).

The proof of Theorem 4-7 will now follow from Lemma 15 and Corollary 16.

5.1 Broadcasting, Predator-Prey: Theorems 4, 5

Recall that $D_{k,\ell}$ is the extinction time of the ℓ prey using k predators. Thus

$$\begin{aligned}\mathbf{E}(D_{k,\ell}) &= O(k\ell T) + \sum_{s=1}^{\ell} \mathbf{E}(M_{s,k}) \\ &\sim n\theta_r \sum_{s=1}^{\ell} \frac{1}{sk} = \frac{2n\theta_r}{k} H_\ell,\end{aligned}$$

where H_ℓ is the ℓ-th harmonic number. Similarly, the time B_k, for a given agent to broadcast to all other agents is $\sum_{s=1,\dots,k-1} M_{s,k-s}$, and thus

$$\mathbf{E}(B_k) = O(k\ell T) + \sum_{s=1}^{k-1} \frac{(1+o(1))}{s(k-s)} \sim \frac{2n\theta_r}{k} H_{k-1}.$$

5.2 Expected time to coalescence: Theorem 6

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

$$\mathbf{E}(S_k) = O(kT) + \sum_{s=1}^{k} \frac{(1+o(1))}{p_s} \sim n\theta_r \sum_{s=1}^{k} \frac{2}{s(s-1)}. \tag{7}$$

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

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

5.3 Expected time to extinction: explosive particles: Theorem 7

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

$$\mathbf{E}(D_k) = O(kT) + \sum_{s=1}^{\ell} \frac{(1+o(1))}{p_{2s}} \sim n\theta_r \sum_{s=1}^{\ell} \frac{2}{2s(2s-1)}. \tag{8}$$

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

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

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 then one, to appear.

[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 (2008).

[9] U. Feige. A tight upper bound for the cover time of random walks on graphs, *Random Structures and Algorithms*, 6 (1995) 51-54.

[10] U. Feige. A tight lower bound for the cover time of random walks on graphs, *Random Structures and Algorithms*, 6 (1995) 433-438.

[11] W. Feller, *An Introduction to Probability Theory, Volume I,* (Second edition) Wiley (1960).

[12] J. Friedman. A proof of Alon's second eignevalue conjecture, Memoirs of the A.M.S., to appear.

[13] T. M. Liggett. *Interacting Particle Systems.* Springer-Verlag (1985).

[14] L. Lovász. Random walks on graphs: a survey. Bolyai Society Mathematical Studies, Combinatorics, Paul Erdös is Eighty (Volume 2) Keszthely (Hungary), (1993) 1-46.

[15] A. Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing* 1 (1992) 351-370.

Appendix

6 Proof of Lemma 10: Mixing time

Case (i): Single random walk.
We choose $\epsilon = 0.1$ in P1 so that for $r \geq 3$ we have

$$\lambda_{\max} \leq 0.977. \tag{9}$$

Let

$$T_G = \frac{3 \ln n}{-\ln \lambda_{\max}}. \tag{10}$$

Using (3) we see that for $t \geq T_G$,

$$\max_{x \in V} |P^{(t)}_{u,G}(x) - \pi_x| \leq n^{-3}.$$

Case (ii): k independent random walks.

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

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

To simplify notation let $P_i = P^t_{u_i}(x_i)$, where $\boldsymbol{u} = (u_1, ..., u_k)$, and let $\pi_i = \pi(x_i)$. Then

$$\begin{aligned}|P^t_{\boldsymbol{u}}(\boldsymbol{x}) - \pi(\boldsymbol{x})| \leq & |P_1...P_k - P_1...P_{k-1}\pi_k| + |P_1...P_{k-1}\pi_k - P_1...P_{k-2}\pi_{k-1}\pi_k| + \cdots \\ & + |P_1...P_\ell \pi_{\ell+1}...\pi_k - P_1...P_{\ell-1}\pi_\ell...\pi_k| + \cdots + |P_1\pi_2...\pi_k - \pi_1...\pi_k|.\end{aligned}$$

It follows that

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

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

$$T_H = \frac{\ln k + 3 \ln n}{-\ln \lambda_{\max}},$$

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

Case (iii): Random walk in Γ.
Let $\lambda_H = \lambda_{\max}(H)$, and let

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

then it is a result of [1] (see also [15]) 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 (9). 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 with $m(G)$ edges, the conductance Φ is given by

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

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

$$1 - 2\Phi \leq \lambda_1 \leq 1 - \frac{\Phi^2}{2}. \tag{11}$$

On the assumption that $\lambda_{\max} = \lambda_1$, we find that

$$\Phi(H) \geq 1/250. \tag{12}$$

The standard way to ensure that $\lambda_{\max} = \lambda_1$, is to make the chain *lazy*, so that 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_H 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.

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

starting position $\boldsymbol{u}$ of a walk $\mathcal{W}_{\boldsymbol{u}}(\Gamma)$ we have, from (3) and (11) that, provided $t \geq T_\Gamma = 10^5 k \ln n$,

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

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

7 Oblivious particles

7.1 Cover time for k particles walking independently

The proof of the following lemma follows directly from the one given in [6] with the simplification made in later papers (e.g. [8]), that we only consider first visits after T. It is obtained from Corollary 11 by substituting $R_v \sim \theta_r$ for tree-like vertices.

Lemma 17 *Let $\mathcal{G}_r$ denote the set of r-regular random graphs. Let T_H be a mixing time given by Lemma 10. Let*

$$p = \frac{1}{n\theta_r}. \tag{13}$$

There exists a subset $\mathcal{H}$ of $\mathcal{G}_r$ of size $(1-o(1))|\mathcal{G}_r|$ such that the following properties hold for $G \in \mathcal{H}$.

(i) Let $\boldsymbol{A}_t(v)$ be the event that a walk starting at a fixed vertex x does not visit v during steps $T, ..., t$. For all $v \in V$

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

where $K > 0$ is constant. Moreover, if v is tree-like then

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

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

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

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

$$\mathbf{Pr}(\boldsymbol{A}_{k,t}(v)) = \mathbf{Pr}(\boldsymbol{A}_{1,t}(v))^k = e^{-kp(t+o(t))}$$

and similarly for $\mathbf{Pr}(\boldsymbol{A}_{k,t}(u, v))$. The proof of Theorem 2 is a straightforward adaptation of the proof of Theorem 1 as given in Lemma 17.

Upper bound on cover time. Let $t_k^* = (\theta_r n/k) \ln n$, then for suitably large t, the event $\boldsymbol{A}_{k,t}(v)$ that vertex v is unvisited in $T, ..., t$ is at most $e^{-kp(t+o(t))}$. Choosing $t_0 = (1 + \epsilon)t^*$ where $\epsilon \to 0$ sufficiently slowly, and substituting this value into the upper bound proof given in Section 5.1 of [6], we find that $C_G(k) \leq t_0 = (1 + o(1))t_k^*$.

Lower bound on cover time. Choosing $t_1 = (1 - \epsilon)t_k^*$ where $\epsilon \to 0$ sufficiently slowly, and substituting this value into the lower bound proof given in Section 5.2 of [6], we find that there is a set of vertices S, given by Lemma 17 above which **whp** are not all covered at time t_1. The conclusion is that $C_G(k) \geq t_1$.

This completes the proof of Theorem 2. □

7.2 Comparison of search methods: Proof of Theorem 3

In this model, particles are of two types, either fixed at a vertex, or mobile. Recall the definition of general position from (1). 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 the starting positions of the mobile particles.
S2: The fixed particles are at tree-like vertices of the graph G.

Conditional on S1, it follows from Corollary 15 that **whp** no mobile particle will visit any fixed particle during any of the (at most) $\max(k, \ell)$ mixing times T occurring during the search process.

Let A be the current set of fixed particles, and B the current set of remaining moving particles, and let $|A| = a, \ |B| = b$. Conditional on S2, the value of $p = p_A$ in (13) is $p \sim a/(n\theta_r)$.

The probability that none of the B particles visits A during $T, .., t$ is

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

Let $F_{A,B}$ be the first visit time, then 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) $a = 1, ..., \ell$ and $b = k$ or (ii) $a = k$ and $b = 1, ..., \ell$ we have

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

8 Proof of Lemma 12 and Lemma 15: Returns to a vertex during the mixing time

The main task of this proof is to estimate the probability two or more particles meet at a given step.

We note the following result (see e.g. [11]), for a random walk on the line $= \{0, ..., a\}$ 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

$$\rho(z, a) = \frac{(q/p)^z - (q/p)^a}{1 - (q/p)^a} \leq \left(\frac{q}{p}\right)^z, \tag{14}$$

provided $q \leq p$.

For a walk starting at z on the line $\{0, 1, ..., \}$, with absorbing states $\{0, \infty\}$, the probability of absorption at the origin is $\rho(z) = (q/p)^z$.

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

Lemma 18 *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 $t \le L$, 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

$$X_t = \begin{cases} X_{t-1} - 2 & \text{with probability } q = \frac{1}{r^2} \\ X_{t-1} & \text{with probability } s = \frac{2(r-1)}{r^2} \\ X_{t-1} + 2 & \text{with probability } p = \left(\frac{r-1}{r}\right)^2. \end{cases} \tag{15}$$

Finally let Z_t be a walk on the even numbers $\{0, \pm 2, \pm 4, ...\}$ 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 $X_t \ge Z_t$. Note that

$$\mathbf{E}(Z_t - Z_{t-1}) = 2 - \frac{4}{r}. \tag{16}$$

Now let g_∞ denote the probability that $\{\exists t \ge 1 : X_t = 0\}$, i.e. the particles meet in $\mathcal{T}$. Equation (14) implies that

$$g_\infty = \frac{1}{(r-1)^2}. \tag{17}$$

Furthermore,

$$g_\infty = g_T + g_T' \tag{18}$$

where g_T' is the probability that $\{\exists t > L_1 : X_t = 0\}$, i.e. the particles meet after L_1 steps. But, using (14) once again, we see that

$$\begin{aligned} g_T' &\le \mathbf{Pr}(X_{L_1} \le L_1/2) + \left(\frac{1}{(r-1)^2}\right)^{L_1/4} \\ &\le O(n^{-\Omega(1)}) + \left(\frac{1}{(r-1)^2}\right)^{L_1/4}. \end{aligned} \tag{19}$$

As $\mathbf{Pr}(Z_{L_1} \le L_1/2) \ge \mathbf{Pr}(X_{L_1} \le L_1/2)$, the bound $O(n^{-\Omega(1)})$ follows from (16) 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 (16) is at least $2/3$ for any r, we can insert absolute constants independent of r.

It follows from (17), (18) and (19) that we can write

$$g_T = \frac{1}{(r-1)^2} + O(n^{-\Omega(1)}). \tag{20}$$

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

$$h'_T = \mathbf{Pr}\left(\text{The walks meet during steps } \{L_1, ..., T\} \text{ and } Y_{L_1} > L_1/2\right),$$

and we know from (19) that

$$h''_T \leq \mathbf{Pr}(Y_{L_1} \leq L_1/2) = O(n^{-\Omega(1)}). \tag{21}$$

Let $\sigma \leq T$ be the step at which the particles first meet again, and let s be the *last* step less than 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 ρ_1, ρ_2 denote the particles at x, y respectively.

Let $N = N(x, L_1)$, 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.

Suppose 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 t is the first step at which the boundary of N 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, $xP'y$ not internally disjoint, and $|P'| \geq |P| \geq L_1/2$. Let w be the mid-point of P so that $xPy = xRwSy$ (resp. w' of P' etc). Then (e.g.) each half-path $xRw, xR'w'$ has a section of length at least $L_1/12$ in common with or disjoint from the other. Repeating the argument above, at least one of these sections must be walked by one of the particles.

Thus there is 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 (14)). As there are at most T ways of choosing s, and T starting times for traversing the sub-path, an upper bound of $O(T^2(r-1)^{-L_1/12})$ follows. □

Using Lemma 18, we can calculate the expected number of returns to the diagonal S of H_k for k particles. Recall the definitions of $\Gamma(S)$ and $\gamma(S)$, the contraction of S.

Lemma 19 *Let k be the number of particles walking on the underlying graph G. Let $\mathcal{W}_\gamma$ be a random walk in Γ starting at γ. Let f^* denote the probability that $\mathcal{W}_\gamma$*

makes a first return to γ within T_Γ steps. Then

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

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

Proof Let $S = \{(v_1, ..., v_k) : \text{at least two } v_i \text{ are the same}\}$. The particles are at the components of the vector corresponding to the vertex in question. Every vertex in S has degree r^k in H_k, and the size of 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

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

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

$$d(\gamma) = \binom{k}{2}n^{k-1}r^k\left(1 - O\left(\frac{k^2}{n}\right)\right). \tag{22}$$

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 correction being for where a different pair of particles coincide at the next step.

A single walk in Γ is related to k independent walks in the underlying graph G. The point of difference being the moves into and out of $\gamma(S)$. First returns to $\gamma(S)$ can be of several types. The simplest type is a loop return (type O), for example two particles move to the same neighbour. If this does not occur, we distinguish four cases. 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 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

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

where $f_O=\frac{d_\ell(\gamma)}{d(\gamma)}$. For type A returns,

$$f_A=\frac{d'(\gamma)}{d(\gamma)}f_{T_\Gamma} \tag{23}$$

Here $d'(\gamma)$ counts the non-loop edges of γ corresponding to tree-like vertices of G, and f_{T_Γ} is from Lemma 18. Thus $d'(\gamma)=(d(\gamma)-d_\ell(\gamma))(1-O(k^2n^{\epsilon_0}/n)-O(k^2/n))$. This follows from P3, as G is typical, and the value of $N_2(k)$ above. 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 $\nu_1=\binom{k}{2}^2r^Ln^{k-2}$, have another pair of entries within the same neighbourhood of depth L and at most $\nu_2=\binom{k}{2}kr^Ln^{k-2}$, have an entry within the neighbourhood of a coincident pair. For particles distance at least L apart, the probability they coincide in T_Γ steps is $O(n^{-\Omega(1)})$, by the analysis of Lemma 18. 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,

$$\begin{aligned}
f_C &= O\left(\frac{k^3n^{k-2}r^k}{d(\gamma)}+\frac{k^3r^{2L}n^{k-2}r^k}{d(\gamma)}+\frac{k^2}{n^{\Omega(1)}}\right)=O\left(\frac{k^2}{n^{\Omega(1)}}\right)\\
f_D &\le \frac{n^{2\epsilon_0}r^Lk^2n^{k-2}}{d(\gamma)}=O\left(\frac{k^2}{n^{\Omega(1)}}\right), \qquad \text{see P3.}
\end{aligned}$$

The expression for f_C arises as follows: The term $O\left(\frac{k^3n^{k-2}r^k}{d(\gamma)}\right)$ is the probability 3 or more particles coincident initially. The term $O\left(\frac{k^3r^{2L}n^{k-2}r^k}{d(\gamma)}\right)$ is the probability that two particles are initially in the L-neighbourhood of a third vertex. 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 meet within time T. □

9 Interacting particles: Proof that the conditions of Lemma 11 hold

We next check that the conditions of Lemma 11 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) $\min_{|z|\le 1+\lambda}|R_T(z)| \ge \theta$, 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 20 below. Condition (b) is easily disposed of. Recall from Lemma 10, that $T = O(k\ln n)$. From Lemma 19, $\pi_\gamma = d(\gamma)/(nr)^k$ where $1 \le d(\gamma)/n^{k-1}r^k \le k^2$. Thus $T^2\pi_\gamma = o(1)$ provided $k \le n^{1/5}$.

Lemma 20 *For $|z| \le 1+\lambda$, there exists a constant $\theta > 0$ such that $|R_T(z)| \ge \theta$.*

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 at time $t \le L$. Thus $R_T(z) = R_A(z) + R_B(z)$ where $R_A(z) = \sum_{t=0}^{L} r_{t,A}z^t$.

The arguments in the proof of Lemma 18 show that

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

and thus

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

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

As in Lemma 18, let Y_t be the distance between the particles during the first L steps. For $1 \le t \le L$ let

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

be the probability that the walks first meet at step t. Let $B(z) = \sum_{t=1}^{L} b_t z^t$, and let $A(z) = \sum_{t=0}^{\infty} a_t z^t = 1/(1 - B(z))$. Thus

$$R_A(z) = \sum_{t=0}^{L} r_{t,A}z^t = \sum_{t=0}^{L} a_t z^t = \sum_{t=0}^{\infty} a_t z^t - \sum_{t=L+1}^{\infty} a_t z^t =$$

$$\sum_{t=0}^{\infty} a_t z^t - O(((1+\lambda)\zeta)^t) = \sum_{t=0}^{\infty} a_t z^t - o(1),$$

as we now explain. By coupling Y_t and Z_t as in Lemma 18, 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 19, $B(1) = f_A + f_O = \frac{1}{r-1} + o(1)$ 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,

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

□

10 The voter model: Theorem 8

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.

We next summarize properties of the voter model (see e.g. [2], Chapter 14.3, *Coalescing random walks and the voter model*). The total number of existing opinions can only decrease with time, and at some random step $t = C_{\boldsymbol{vm}}$ there will be only one consensus opinion, k say. The reverse process, using the same edges and working backwards from t constitutes a *coalescing random walk process* of length $C_{\boldsymbol{crw}}$, where k is the vertex at which the particles finally coalesce.

In particular the event for the voter model: By time t_0 everyone's opinion is the opinion initially held by person k, is exactly the same as the event 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 $\mathbf{Pr}(C_{\boldsymbol{vm}} \leq t_0) = \mathbf{Pr}(C_{\boldsymbol{crw}} \leq t_0)$, these two times C have the same distribution and hence the same expected value.

For completeness we briefly explain this 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 $s = 0, 1, ..., t$ each vertex i chooses a random neighbour $j(i) \in N(i)$. This defines a digraph with levels $0, 1, ..., t$, each with vertex set V, and the edges generated at step $s = 0, 1, ..., t-1$ between the levels. Initially at level 0, particle i is at vertex i, $i = 1, ..., n$. 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.

10.1 Voter Model: Proof of Theorem 8

Lower bound. Pick any $k \to \infty$ 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 6 $\mathbf{E}(S_k) \sim 2\theta_r n$, as $k \to \infty$.

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

Let $\mathcal{P}(k, \boldsymbol{v})$ be the set of particles starting from vertices $\boldsymbol{v} = (v_1, ..., v_k)$, not necessarily in general position. Either there has been a meeting during the mixing time T_Γ, or if not, the results of Theorem 14 apply for all $t \geq T_\Gamma$.

Suppose no meeting has occurred. Let

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

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

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

Hence

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

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 $\mathbf{E}C \leq rn^2/(4s)$ from Proposition 9, Chapter 14, Section 3.1 of [2], restarting the process at t^* with however many particles remain at that time.