Multiple Random Walks and Interacting Particle Systems

Colin Cooper[1,⋆], Alan Frieze[2,⋆⋆], and Tomasz Radzik[1,⋆⋆⋆]

[1] Department of Computer Science, King's College London, London WC2R 2LS, UK
{colin.cooper,tomasz.radzik}@kcl.ac.uk
[2] Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA
alan@random.math.cmu.edu

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. For k independent walks on a random regular graph 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 $2\theta_r n(\ln k)/k$, 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 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.

⋆ Research supported in part by Royal Society grant 2006/R2-IJP and EPSRC grant EP/D00232X/1.
⋆⋆ Research supported in part by NSF grant CCF0502793.
⋆⋆⋆ Research supported in part by Royal Society grant 2006/R2-IJP.

S. Albers et al. (Eds.): ICALP 2009, Part II, LNCS 5556, pp. 399–410, 2009.

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, $r \geq 3$. 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 becomes 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 is 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 G:

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, 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.
4. **Annihilating particles.** For $k = 2\ell$ particles walking independently, which destroy each other (pairwise) on meeting, we give the expected time to extinction.

5. **Coalescing particles.** For k particles walking independently, which coalesce on meeting, 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. The formal definitions of the random variables above can be largely found in [2], Chapter 14.

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 [14]). 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 $C_G \leq 2m(n-1)$. It was shown by Feige [10], [11], that for any connected graph G, the cover time satisfies $(1 - o(1))n \ln n \leq C_G \leq (1 + o(1))\frac{4}{27}n^3$.

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. The paper gives an example, the barbell graph, (two cliques joined by a long path) for which the speed-up is exponential in k provided $k \geq 20 \ln n$.

The paper [4] found that for expanders the speedup was $\Omega(k)$ for $k \leq n$ particles. The class of r-regular graphs we consider are expanders. For these graphs, comparing Theorem 2 with Theorem 1, we see that $C_G(k) \sim C_G/k$, i.e. the asymptotic speedup is *exactly linear*.

Theorem 2 Multiple particles walking independently
Let $r \geq 3$ be constant. Let G be chosen randomly from $\mathcal{G}_r$, then **whp** *and independently of the initial positions of the particles:*

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

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

(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 $k,\ \ell \le n^{\epsilon}$ *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**, *where* H_ℓ *is the* ℓ*-th harmonic number,*

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

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 $d(v_i, v_j) \ge \omega = \omega(k,n)$ between particles, where

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

For the results given in this section, we assume that $r \ge 3$ 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 $k \le n^{\epsilon}$ *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 originating agent to tell it directly to all other agents (by meeting directly with 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,

$$\mathbf{E}(S_k) \ \sim\ 2\theta_r n(k-1)/k,$$

so $\mathbf{E}(S_k) \sim 2\theta_r n,$ *if* $k \to \infty$.

As a twist on predator-prey, we consider "explosive" particles which destroy each other (pairwise) on meeting at a vertex (that is, if two meet, then both are destroyed, but if, say, five meet, then two pairs are destroyed and one particle survives).

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

$$\mathbf{E}(D_k) \ \sim\ 2\theta_r n(H_{2\ell} - H_\ell),$$

so $\mathbf{E}(D_k) \sim 2\theta_r(\ln 2)n,$ *if* $k \to \infty$.

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,*

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

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 2 are met, then the probability that a vertex v is unvisited at steps $T, ..., t$ tends to $(1 - \pi_v/R_v)^t$. Here π is the stationary distribution and R_v is the number of returns to v during T by a walk starting at v. The value R_v 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 analyze the walk in the contracted graph Γ.

Proof of Theorems. Because of space restrictions, we only give results and ideas of proofs in this extended abstract. Full proofs of the theorems of this paper are in [9].

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 [13]. 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 [15].

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 1. *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 $k \leq n$,

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

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 $t = 0, 1, \ldots, T_1$. 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 2. *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 2, 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 the value of $\pi(\gamma)$, the stationary distribution of γ in Γ.
(ii) The conditions of Lemma 2 hold with respect to the vertex γ of the graph Γ.
(iii) The probability that any particles meet during the mixing time T_Γ.

These points are formally summarized in Lemmas 3-4 below.

Lemma 3. *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 4. *If* $k \leq n^\epsilon$ *then the conditions of Lemma 2 hold with respect to the vertex* γ *of a typical graph* Γ.

From (5) with $v = \gamma$, and Lemma 3 we have

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

It follows from [9] 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 Lemma 2 we have the following theorem.

Theorem 10. *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 $p_k = \frac{\binom{k}{2}}{\theta_r n}(1 + O(n^{-\Omega(1)}))$. 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 $q_{sk} = \frac{sk}{\theta_r n}(1 + O(n^{-\Omega(1)}))$. 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 5. *For typical graphs G and $k \le n^\epsilon$ 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, for a sufficiently large constant α,

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

From Lemma 5, 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 1. *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 $\mathbf{E}(M_k) = \sum_{t\ge T} \mathbf{Pr}(\boldsymbol{A}_k(t))$ and $p_k T_\Gamma = o(1)$. □

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 5, 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 5 and Corollary 1.

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

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

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 $\sum_{s=1}^{k-1} M_{s,k-s}$, and thus

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

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,

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

We see that for $k \to \infty$, $\mathbf{E}(S_k) \sim 2\theta_r n$.

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

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

Noting that $\lim_{\ell\to\infty}(H_{2\ell} - H_\ell) = \ln 2$, we have $\mathbf{E}(D_k) \sim 2\theta_r(\ln 2)\, n$, for $k \to \infty$.

References

1. Aldous, D.: Some inequalities for reversible Markov chains. J. London Math. Soc. 25(2), 564–576 (1982)
2. Aldous, D., Fill, J.: Reversible Markov Chains and Random Walks on Graphs. Monograph in preparation, `http://stat-www.berkeley.edu/pub/users/aldous/RWG/book.html`
3. Aleliunas, R., Karp, R.M., Lipton, R.J., Lovász, L., Rackoff, C.: Random Walks, Universal Traversal Sequences, and the Complexity of Maze Problems. In: Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science, pp. 218–223 (1979)
4. Alon, N., Avin, C., Kouchý, M., Kozma, G., Lotker, Z., Tuttle, M.: Many random walks are faster then one. In: Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures, pp. 119–128 (2008)
5. Broder, A., Karlin, A., Raghavan, A., Upfal, E.: Trading space for time in undirected s-t connectivity. In: Proceedings of the 21st Annual ACM Symposium on Theory of Computing, pp. 543–549 (1989)
6. Cooper, C., Frieze, A.M.: The cover time of random regular graphs. SIAM Journal on Discrete Mathematics 18, 728–740 (2005)

7. Cooper, C., Frieze, A.M.: The cover time of the preferential attachment graph. Journal of Combinatorial Theory Series B 97(2), 269–290 (2007)
8. Cooper, C., Frieze, A.M.: The cover time of the giant component of a random graph. Random Structures and Algorithms 32, 401–439 (2008)
9. Cooper, C., Frieze, A.M., Radzik, T.: Multiple random walks in random regular graphs (2008), http://www.math.cmu.edu/~af1p/Texfiles/Multiple.pdf
10. Feige, U.: A tight upper bound for the cover time of random walks on graphs. Random Structures and Algorithms 6, 51–54 (1995)
11. Feige, U.: A tight lower bound for the cover time of random walks on graphs. Random Structures and Algorithms 6, 433–438 (1995)
12. Feller, W.: An Introduction to Probability Theory, 2nd edn., vol. I. Wiley, Chichester (1960)
13. Friedman, J.: A proof of Alon's second eignevalue conjecture. Memoirs of the A.M.S. (to appear)
14. Liggett, T.M.: Interacting Particle Systems. Springer, Heidelberg (1985)
15. Lovász, L.: Random walks on graphs: a survey. In: Bolyai Society Mathematical Studies, Combinatorics, Paul Erdös is Eighty, Keszthely, Hungary, vol. 2, pp. 1–46 (1993)
16. Sinclair, A.: Improved bounds for mixing rates of Markov chains and multicommodity flow. Combinatorics, Probability and Computing 1, 351–370 (1992)