

Coalescing random walks and voting on graphs

Colin Cooper¹, Robert Elsässer², Hirotaka Ono³, and Tomasz Radzik¹

¹ Department of Informatics, King's College London, UK

² Department of Computer Science, University of Paderborn, Germany

³ Department of Economic Engineering, University of Kyushu, Fukuoka, Japan

Abstract. In a *coalescing random walk*, a set of particles make independent random walks on a graph. Whenever one or more particles meet at a vertex, they unite to form a single particle, which then continues the random walk through the graph. Coalescing random walks can be used to achieve consensus in distributed networks, and is the basis of the self-stabilizing mutual exclusion algorithm of Israeli and Jalfon [11].

Let $G = (V, E)$, be an undirected, connected n vertex graph. Let $C(n)$ be the expected time for all particles to coalesce, when initially one particle is located at each vertex of an n vertex graph.

We study the problem of bounding the coalescence time $C(n)$ for general classes of graphs. For d -regular graphs with expansion parameterized by the eigenvalue gap $1 - \lambda_2$, where λ_2 is the second eigenvalue of the transition matrix of the random walk, we establish that $C(n) = O(n/(1 - \lambda_2))$. This result also extends to near regular graphs; where a graph is near regular if the ratio of the maximum and minimum degrees Δ/δ is constant.

Our general result is, that $C(n) = O(n/(\nu(1 - \lambda_2)))$, where $\nu = (\sum d^2(v))/(d^2 n)$, and d is the average node degree. The parameter ν is an indicator of the variability of node degrees: $1 \leq \nu = O(n)$, with $\nu = 1$ for regular graphs. The result holds provided the maximum node degree is $O(m^{1-\epsilon})$.

A system of coalescing particles where initially one particle is located at each vertex, corresponds to a *voter model*. Initially each vertex has a distinct opinion, and at each step each vertex changes its opinion to that of a random neighbour. The voting process can be used for leader election in a distributed context. Let $\mathbf{E}(C_{\mathbf{v}})$ be the expected time for voting to complete, i.e. for a unique opinion to emerge. It is known that $E(C_{\mathbf{v}}) = C(n)$, so our results imply that $E(C_{\mathbf{v}}) = O(n/(\nu(1 - \lambda_2)))$.

We also investigate how the voting time improves when a vertex elicits more than one opinion at each step. In a model which we call *min-voting*, each vertex initially holds an opinion from $\{1, 2, \dots, n\}$. At each step each vertex takes the opinions of two random neighbours and keeps the smaller. We show that for regular graphs with very good expansion properties, voting is completed in $O(\log n)$ time with high probability. This result can be viewed as an example for the “power of two choices” in distributed voting.

1 Introduction

In a *coalescing random walk*, a set of particles make independent random walks in an undirected connected graph. Whenever one or more particles meet at a vertex, then they unite to form a single particle which then continues to make a random walk through the graph.

The expected time for the particles to coalesce to a single particle depends on their initial positions. For a connected graph $G = (V, E)$ with n vertices and m edges, let $C_k(i_1, \dots, i_k)$, $2 \leq k \leq n$, be the coalescence time when there are initially k particles starting from distinct vertices i_1, \dots, i_k . The quantity we study is $C(k) = \max_{i_1, \dots, i_k} \mathbf{E}(C_k(i_1, \dots, i_k))$, the worst case expected coalescence time.

A system of n coalescing particles where initially one particle is located at each vertex, corresponds to another classical problem, the *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.

Let $C_{\mathbf{v}}$ be the number of steps for voting to be completed, i.e. for a unique opinion to emerge. The random variable $C_{\mathbf{v}}$ has the same distribution, and hence the same expected value, as the coalescence time C_n of n coalescing particles, where one particle is initially located at each vertex, (see [2]). Thus $C(n) \equiv E(C_n) = E(C_{\mathbf{v}})$. The expected completion time $\mathbf{E}(C_{\mathbf{v}})$ is also called the *voting time*, *trapping time* or the *consensus time*.

The coalescing random walk is the basis of the self-stabilizing mutual exclusion algorithm of Israeli and Jalfon [11]. Initially each vertex emits a token which makes a random walk on G . On meeting at a vertex, tokens coalesce. Provided the graph is connected, and not bipartite, eventually only one token will remain, and the vertex with the token has exclusive access to some resource. The token makes a random walk on G , so in the long run it will visit all vertices of G in proportion to their stationary distribution.

The results given in this paper. We study the problem of bounding the coalescence time $C(n)$ for general classes of graphs. For d -regular graphs with expansion parameterized by the eigenvalue gap $1 - \lambda_2$, where λ_2 is the second eigenvalue of the transition matrix of the random walk, we establish that $C(n) = O(n/(1 - \lambda_2))$. This bound is actually a consequence of our general result that

$$C(n) = O\left(\frac{n}{\nu(1 - \lambda_2)}\right).$$

where parameter $\nu = (\sum_{v \in V} d^2(v))/(d^2n) \geq 1$ measures the variability of the degree sequence; d is the average degree. This parameter ranges from 1 for regular

graphs to $\Theta(n)$ for a star graph. Our bound on $C(n)$ holds provided the maximum degree Δ satisfies the (weak) condition that $\Delta = O(m^{1-\epsilon})$, where m is the number of edges and ϵ can be an arbitrarily small positive constant.

We also investigate the voting time when each vertex elicits more than one opinion at each step. Our model, which we call *min-voting*, is as follows. Initially each vertex holds a distinct opinion, chosen uniformly at random without replacement from $\{1, 2, \dots, n\}$, and, at each step each vertex takes the opinions of two random neighbours and keeps the smaller of the two (disregarding its own opinion). We establish that for regular graphs with very good expansion properties (in the sense of [7]), voting can be completed in $O(\log n)$ time with high probability. (With high probability, **whp** with probability tending to 1 as $n \rightarrow \infty$).

Previous work on coalescing walks and voting systems. We next summarize some of what is known about these topics. Cox [6] considered coalescence time of random walks and the consensus time of the voter model for d -dimensional tori. In a variant of the voter model, *the two-party model*, initially there are only two opinions A and B . The two-party model was considered by Donnelly and Welsh [9]. Hassin and Peleg [15] and Nakata *et al.* [14] also consider the two-party model, and discuss its application to agreement problems in distributed systems. Papers [14, 15] focus on analysing the probability that all vertices will eventually adopt the opinion which is initially held by a given group of vertices. The central result is that the probability that opinion A wins is $d(A)/(2m)$, where $d(A)$ is the sum of the degrees of vertices initially holding opinion A , and m is the number of edges in G .

Let $H_{u,v}$ denote the *hitting time* of vertex v starting from vertex u , that is, the random variable which gives the time taken by a random walk starting from vertex u to reach vertex v ; and let $H_{\max} = \max_{u,v} \mathbf{E}(H_{u,v})$. Aldous [1] showed that $C(2) = O(H_{\max})$, which implies that $C(n) = O(H_{\max} \log n)$ (since the number of particles halves in $O(H_{\max})$ steps), and conjectured that $C(n)$ is actually $O(H_{\max})$. Cox's results [6] imply that the conjecture $C(n) = O(H_{\max})$ is true for constant dimension tori and grids. Aldous [1] also states a lower bound for $C(2)$. For graphs, this bound can be simplified to $C(2) = \Omega(m/\Delta)$, where Δ is the maximum degree of a vertex in G ([1] considers more general Markov Chains). For the class of local expanders we study in this paper, this gives $C(2) = \Theta(n)$. However, the bounds $C(2) = \Omega(m/\Delta)$ and $C(2) = O(H_{\max})$ can be far apart. For example, for a star graph (with loops), $C(2) = \Theta(1)$ whereas the bounds give $\Omega(1) \leq C(2) = O(n)$.

Aldous and Fill [2] showed that for regular graphs $C(n) \leq e(\log n + 2)H_{\max}$, for d -regular s -edge connected graphs $C(n) \leq \frac{dn^2}{4s}$, and for complete graphs $C(n) \sim n$ (where $f(n) \sim g(n)$ means that $f(n) = (1 \pm o(1))g(n)$).

Cooper *et al.* [5] showed that the conjecture $C(n) = O(H_{\max})$ is true for the family of random regular graphs. They proved that for r -regular random graphs, $\mathbf{E}(C) = C(n) \sim 2((r-1)/(r-2))n$, **whp**.

As noted above, our voting processes can be viewed as consensus, or aggregation. There is a large amount of research focusing on distributed selection and aggregation in different scenarios and various settings (see e.g. [16, 17] or [3] for a survey). Here, we only mention the result of [8], in which the following process related to min-voting is analysed. At the beginning each vertex of a complete graph has its own opinion. In each step every vertex contacts two other nodes uniformly at random, and changes its opinion to the median of the opinions of these two nodes and its own opinion. It is shown in [8] that this rule converges in time $O(\log n)$ to a single opinion, **whp**. Unfortunately, the results of [8] cannot be directly extended to sparse graphs with good expansion properties.

1.1 Results: Coalescence and Voting

We assume the graphs G we consider are not bipartite, or that if G is bipartite, then we assume the coalescing random walk is lazy, and pauses with probability $1/2$ at each step. Equivalently, for the voting process, we assume that vertices may choose their own opinion with this probability.

We prove the following very general theorem.

Theorem 1. *Let G be a connected graph with n vertices, m edges, average vertex degree d , and maximum degree $\Delta = O(m^{1-\epsilon})$, for an arbitrary constant $\epsilon > 0$. Let $\nu = (\sum_{v \in V} d^2(v))/(d^2n)$. Let $C(n)$ be the expected coalescence time for a system on n particles making a lazy random walk on G , where originally one particle starts at each vertex. Then*

$$C(n) = O\left(\frac{n}{\nu(1-\lambda_2)}\right). \quad (1)$$

Thus by the equivalence between coalescence and voting, the expected time $\mathbf{E}(C_{\mathbf{v}})$ to complete voting on G is

$$\mathbf{E}(C_{\mathbf{v}}) = O\left(\frac{n}{\nu(1-\lambda_2)}\right). \quad (2)$$

The quantity $\nu = (\sum_{v \in V} d^2(v))/(d^2n)$ measures the variability of the degree sequence. It is easy to check that $1 \leq \nu \leq \Delta/d \leq n$. The result (1) holds for many classes of graphs, provided the bound $\Delta = O(m^{1-\epsilon})$, which is not that restrictive.

For *near regular* graphs, when the ratio of the largest to the smallest vertex degree is bounded by a constant, we have $\nu \leq \Delta/d = O(1)$, so the bound (2) becomes

$$C(n) = O\left(\frac{n}{1 - \lambda_2}\right), \quad (3)$$

In particular, If G is an expander in the classic sense that the eigenvalue gap $(1 - \lambda_2)$ is constant, then $C(n) = O(n)$.

As $1 - \lambda_2 \geq 1/2n^2$ for any connected graph (see e.g. Sinclair [18]) our bound shows that coalescence is completed in $O(n^3)$ expected time on any connected n vertex graph, provided the required bound on Δ . Hassin and Peleg [15] showed that the voting (hence also coalescence) is completed in expected $O(n^3 \log n)$ time on any connected graph. Our bound parameterized by the eigenvalue gap can be viewed as refinement of such absolute bounds.

The proof of our results is based on the following theorem bounding the time to first meeting between any of k particles.

Theorem 2. *For $2 \leq k \leq k^* = \log^3 n$ particles starting from general position in G , let M_k be the time to first meeting. Let $\Delta(G) = o(m/\log^6 n)$. Then for $\nu = (\sum d^2(v))/d^2n$,*

$$\mathbf{E}(M_k) = O\left(\frac{1}{1 - \lambda_2} \left(k \log n + \frac{n}{\nu k^2}\right)\right). \quad (4)$$

1.2 Results: Min-Voting on regular expanders

Here we consider d -regular graphs $G = (V, E)$ with good global expansion properties. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be the eigenvalues of the transition matrix of the random walk on G , and let $\lambda = \max\{\lambda_2, |\lambda_n|\}$. We assume that $\lambda \leq c/\sqrt{d}$ for some constant c , where $c \ll \sqrt{d}$ if $d = O(1)$. In terms of the previous theorems, this corresponds to a second eigenvalue gap $1 - \lambda_2(P)$ of the transition matrix of size $1 - c/\sqrt{d} > 0$. We call such a connected graph an *almost Ramanujan* or *expanding* (cf. [7]).

As shown in [4, 7], these graphs satisfy the following expansion properties. If d is large enough, then there exists a constant β such that for any constant ϕ , sufficiently small $\alpha > 0$, and any $A \subset V$ with $|A| \leq n/2$:

1. $\alpha d|A| \leq |E(A, V \setminus A)| \leq \frac{d|A||V \setminus A|}{n} + \lambda d\sqrt{|A| \cdot |V \setminus A|}$;
2. if $|A| \leq \min\{\phi n/d, n/2\}$, then $|N(A)| \geq \alpha d|A|$;

3. the number of vertices in $V \setminus A$ with at least $\alpha d|A|/n$ neighbours in A is at least $|V \setminus A| - \frac{\beta n^2}{d|A|}$.

Here $E(A, V \setminus A)$ represents the set of edges between A and $V \setminus A$ and $N(A)$ is the set of neighbours of A in $V \setminus A$. In the following, we assume that d is large enough, and α , ϕ , and β do not depend on d . It is easy to see that if the definition holds for some α (where ϕ is fixed), then it also holds for values smaller than α . Hence, we assume that α is arbitrarily small and ϕ is arbitrarily large.

For this class of graphs, we consider the following voting model called *min-voting*. Initially each node has a distinct opinion from $\{1, \dots, n\}$. In each step, every node contacts two neighbours independently and uniformly at random, and changes its opinion to the smallest among the opinions of the two nodes. The two neighbours are not necessarily distinct. We obtain the following result for this voting model.

Theorem 3. *Let $G = (V, E)$ be a d -regular almost Ramanujan graph, where d exceeds some large constant value. Applying min-voting, after a certain number of $O(\log n)$ steps all nodes will have the same opinion, **whp**.*

The proof consists of three main phases. In the first phase, we show that a small opinion is either removed from the graph or it resides at $\Omega(\log^2 n)$ nodes after $(\log \log n)^2$ steps. In the second phase, we show that after $O(\log n)$ steps, at least $n/2$ nodes have the same (currently smallest) opinion in the graph. (which at this time is not necessarily 1). Finally, we show that after additional $O(\log n)$ steps, all nodes will adopt this opinion. The full proof of Theorem 3 for min-voting is given in the Appendix.

2 Random Walk Properties

Let $G = (V, E)$ denote a connected graph, $|V| = n$, $|E| = m$, and let $d(v)$ be the degree of a vertex v . A *simple random walk* \mathcal{W}_u , $u \in V$, on graph G is a Markov chain modeled by a particle moving from vertex to vertex according to the following rule. The probability of transition from vertex v to vertex w is equal to $1/d(v)$, if w is a neighbour of v , and 0 otherwise. The walk \mathcal{W}_u starts from vertex u at $t = 0$. Denote by $\mathcal{W}(t)$ the vertex reached at step t ; $\mathcal{W}(0) = u$.

We assume the random walk \mathcal{W}_u on G is ergodic with stationary distribution π , where $\pi_v = d(v)/(2m)$. If this is not the case, e.g. G is bipartite, then the walk can be made ergodic, by making it *lazy*. A random walk is *lazy*, if it moves from

v to one of its neighbours w with probability $1/(2d(v))$, and stays where it is (at vertex v) with probability $1/2$.

Let $P = P(G)$ be the matrix of transition probabilities of the walk and let $P_u^t(v) = \mathbf{Pr}(\mathcal{W}_u(t) = v)$. Let the eigenvalues $\lambda(P)$ of $P(G)$ be $\lambda_1 = 1 \geq \lambda_2 \geq \dots \geq \lambda_n > -1$, as we assume G is not bipartite. Let $\lambda_{\max} = \max(\lambda_2, |\lambda_n|)$. The rate of convergence of the walk is given by

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

see for example, Lovasz [13]. We assume henceforth that $\lambda_{\max} = \lambda_2$. If not, the standard way to ensure that $\lambda_{\max} = \lambda_2$, is to make the chain lazy.

We use the following definition of mixing time T_G , for a graph G . For all vertices u and x in G and any $t \geq T_G$,

$$|P_u^{(t)}(x) - \pi_x| \leq o\left(\frac{1}{n^2}\right). \quad (6)$$

For convenience we assume that $T_G = \Omega(\log n)$, even if this is not necessary.

Let $\mathbf{E}_\pi(H_w)$ denote the expected hitting time of a vertex w from the stationary distribution π . The quantity $\mathbf{E}_\pi(H_w)$ can be expressed as (see e.g. [2], Chapter 2)

$$\mathbf{E}_\pi(H_v) = Z_{vv}/\pi_v, \quad (7)$$

where

$$Z_{vv} = \sum_{t=0}^{\infty} (P_v^{(t)}(v) - \pi_v). \quad (8)$$

Let $A_v(t; u)$ denote the event that \mathcal{W}_u does not visit vertex v in steps $0, \dots, t$. The following lemma gives a bound on the probability of this event in terms of $\mathbf{E}_\pi(H_v)$ and the mixing time of the walk.

Lemma 1. *Let $T = T_G$ be a mixing time of a random walk \mathcal{W}_u on G satisfying (6). Then*

$$\mathbf{Pr}(A_v(t; u)) \leq e^{-\lfloor t/(T+3\mathbf{E}_\pi(H_v)) \rfloor}.$$

Proof. Let $\rho \equiv P_u^{(T)}$ be the distribution of \mathcal{W}_u on G after T steps. Then (6) and the fact that $\pi_x = \Omega(1/n^2)$ for any connected graph imply

$$\mathbf{E}_\rho(H_v) = (1 + o(1))\mathbf{E}_\pi(H_v). \quad (9)$$

Let $H_v(\rho)$ be the time to hit v starting from ρ , and let $\tau = T + 3\mathbf{E}_\pi(H_v)$. Then, noting that $\mathbf{E}_\rho(H_v) \equiv \mathbf{E}(H_v(\rho))$,

$$\begin{aligned} \mathbf{Pr}(A_v(\tau; u)) &= \mathbf{Pr}(A_v(T; u) \text{ and } H_v(\rho) \geq 3\mathbf{E}_\pi(H_v)) \\ &\leq \mathbf{Pr}(H_v(\rho) \geq 3\mathbf{E}_\pi(H_v)) \leq \mathbf{Pr}(H_v(\rho) \geq e \cdot \mathbf{E}(H_v(\rho))) \leq \frac{1}{e}. \end{aligned}$$

By restarting the process \mathcal{W}_u at $W(0) = u, W(\tau), W(2\tau), \dots, W(\lfloor t/\tau \rfloor \tau)$ we obtain

$$\mathbf{Pr}(A_v(t; u)) \leq e^{-\lfloor t/\tau \rfloor}.$$

□

3 Multiple random walks

In this paper we consider the coalescence of $k \geq 2$ independent random walks on a graph $G = (V_G, E_G)$. To do this we replace the k walks by a single walk as follows.

Let graph $Q = Q_k = (V_Q, E_Q)$ have vertex set $V_Q = V^k$. A vertex \mathbf{v} of Q_k is a k -tuple $\mathbf{v} = (v_1, v_2, \dots, v_k)$ of vertices $v_i \in V_G, i = 1, \dots, k$, with repeats allowed. Two vertices $\mathbf{v}, \mathbf{w} \in V_Q$ are adjacent if $\{v_1, w_1\}, \dots, \{v_k, w_k\}$ are edges of G . There is a direct equivalence between k random walks $\mathcal{W}_{u_i}(t)$ on G with starting positions u_i and a single random walk $\mathcal{W}_{\mathbf{u}}(t)$ on Q_k with starting position $\mathbf{u} = (u_1, u_2, \dots, u_k)$.

For any starting positions $\mathbf{u} = (u_1, \dots, u_k)$ of the walks, let $M_k(\mathbf{u})$ be the time until the first meeting in G . Let $S_k \subset V(Q_k)$, the *diagonal set of vertices*, be defined by

$$S = S_k = \{(v_1, \dots, v_k) : v_i = v_j \text{ some } 1 \leq i < j \leq k\}.$$

If the random walk on Q_k visits this set, two particles occupy the same vertex in the underlying graph G and a (coalescing) meeting occurs.

Because visits to a set by a random walk is not a readily manipulated quantity, the easiest approach is to contract S_k to a single vertex $\gamma = \gamma_k = \gamma(S_k)$, thus replacing Q_k by a graph $\Gamma = \Gamma_k$. On contraction, all edges, including loops, are retained. Thus $d_\Gamma(\gamma) = d_Q(S)$, where d_F denotes vertex degree in graph F , and the degree $d_F(X)$ of a set X is the sum of the degrees of the vertices in X . Moreover Γ and Q have the same total degree, and the degree of any vertex of Γ other than γ is the same as in graph Q . Let π and $\hat{\pi}$ be the stationary distributions of a random walk on Q and Γ , respectively. If $\mathbf{v} \notin S$ then $\hat{\pi}_{\mathbf{v}} = \pi_{\mathbf{v}}$, and $\hat{\pi}_\gamma = \pi_S \equiv \sum_{\mathbf{x} \in S} \pi_{\mathbf{x}}$.

It follows that, if T_Γ is a mixing time satisfying (6) in Γ , then

$$\mathbf{E}(M_k(\mathbf{u})) \leq T_\Gamma + (1 + o(1))\mathbf{E}_{\hat{\pi}}(H_{\gamma_k}), \quad (10)$$

where $\mathbf{E}_{\hat{\pi}}(H_{\gamma_k})$ is the hitting time of γ_k in Γ from stationarity.

Because we have replaced k individual walks on G by a single walk on Q_k , and then on Γ , we need to relate mixing times on T_Q and T_Γ directly to a given mixing time T_G of a single random walk on the underlying graph G . (We will need T_Γ to apply Lemma 1 to graph Γ .)

Lemma 2. *For random walks in graphs G , Q and Γ , there are mixing times*

$$T_G = \frac{\log n}{1 - \lambda_2(G)}, \quad T_Q = O(kT_G), \quad \text{and} \quad T_\Gamma = O(kT_G) \quad (11)$$

such that

$$\max_{u, x \in V_F} |P_u^t(x) - \pi_x| = o(1/n_F^2), \quad \text{for any } t \geq T_F,$$

where F is any of the graphs G , Q or Γ , and $n_F = |V_F|$.

Proof. We get the bound on T_G from (5), since $\pi_x/\pi_u = O(n)$ and $\lambda_2^{1-\lambda_2}$ has a constant $c < 1$ upper bound. To use (5) also to derive bounds on T_Q and T_Γ , we need to know the eigenvalues of Q_k and Γ in terms of the eigenvalues of G . We claim that $\lambda_2(\Gamma) \leq \lambda_2(Q_k)$, and $\lambda_2(Q_k) = \lambda_2(G)$. This follows from established results, as we next explain.

In the jargon of Markov processes, the random walk on Q_k is known as the *tensor product chain*, and its eigenvalues are the k -wise products of the eigenvalues of G . Thus, assuming $\lambda_2(G) = \lambda_{\max}(G)$, it follows that $\lambda_2(Q_k) = \lambda_2(G)$. See [12] page 168 for more details.

In the notation of [2, Ch. 3], the random walk on Γ is the random walk on Q_k with S collapsed to $\gamma(S)$. It is proved in [2, Ch. 3], Corollary 27, that if a subset A of vertices is collapsed to a single vertex, then the second eigenvalue of the transition matrix cannot increase (in that corollary the variable $\tau_2 = 1/(1 - \lambda_2)$). Thus $\lambda_2(Q) \geq \lambda_2(\Gamma)$.

We get the factor k in the bounds (11) on the mixing times because we need $|P_u^T(x) - \pi_x| = o(1/n^{3k})$, as the number of vertices in graphs Q and Γ is $O(n^k)$ and $\pi_x/\pi_u = O(n^{2k})$. \square

For reference, we record the salient facts for G, Q, Γ in Table 1. The bound on π_γ will be established in Lemma 4.

4 Expected hitting time from stationarity

Our proofs of Theorems 1 and 2 are based on a good upper bound on the expected hitting time of vertex γ by a random walk in Γ which starts from the stationary

Graph F	n_F	π_v – stationary distribution	T_F – mixing time
G	$n_G = n$	$\pi_v = d(v)/2m$	$T_G = O(\log n/(1 - \lambda_2))$
Q_k	$n_Q = n^k$	$\pi_v \leq (\Delta/2m)^k$	$T_Q = O(kT_G)$
Γ_k	$n_\Gamma \leq n^k$	$\pi_\gamma \leq (k^2(2m)^{k-2}n\Delta^2)/(2m)^k$	$T_\Gamma = O(T_Q)$

Table 1. The main parameters of the random walks on graphs G , Q_k and Γ_k .

distribution. We obtain such a bound using (7) and deriving an upper bound on $Z_{\gamma\gamma}$ (Lemma 3) and a lower bound on π_γ (Lemma 4).

Lemma 3. *Let F be a graph with the eigenvalue gap $1 - \lambda_2$, then*

$$Z_{vv} \leq \frac{1}{1 - \lambda_2}. \quad (12)$$

In particular, for any vertex v of Q or Γ , $Z_{vv} \leq 1/(1 - \lambda_2(G))$.

Proof. Let $\lambda_2 = \lambda_2(F)$. Using (5) with $x = u = v$, then

$$|P_v^t(v) - \pi_v| \leq \lambda_2^t,$$

and

$$Z_{vv} = \sum_{t \geq 0} (P_v^t(v) - \pi_v) \leq \sum_{t \geq 0} \lambda_2^t = \frac{1}{1 - \lambda_2}.$$

From the proof of Lemma 2, both $F = Q, \Gamma$ satisfy $(1 - \lambda_2(\Gamma)) \geq 1 - \lambda_2(Q) = 1 - \lambda_2(G)$. \square

Lemma 4. *Let $k \leq \log^3 n$. Let G be a connected n vertex, m edge graph satisfying $\Delta(G) = o(m/\log^6 n)$. Let $\gamma = \gamma_k$ in Γ be the contraction of $S = S_k$ in Q . Then there exists $c_k > 0$ constant such that*

$$\pi_\gamma = \frac{d(\gamma)}{(2m)^k} \geq \frac{c_k k^2 \nu}{n}. \quad (13)$$

Proof. By definition, $d(\gamma) = d(S)$. For $1 \leq x < y \leq k$, define the following subsets of S :

$$S_{(x,y)} = \{(v_1, \dots, v_k) : v_x = v_y\}.$$

We have

$$S = \bigcup_{1 \leq x < y \leq k} S_{(x,y)},$$

$$d(S_{(x,y)}) = (2m)^{k-2} \sum_{v \in V} d^2(v),$$

and for $\{x, y\} \neq \{p, q\}$,

$$d(S_{(x,y)} \cap S_{(p,q)}) = \begin{cases} (2m)^{k-4} \sum_{v,u \in V} d^2(v) d^2(u), & \text{if } \{x, y\} \cap \{p, q\} = \emptyset, \\ (2m)^{k-3} \sum_{v \in V} d^3(v), & \text{if } |\{x, y\} \cap \{p, q\}| = 1. \end{cases}$$

Therefore, from the inclusion-exclusion principle and recalling $\sum_{v \in V} d^2(v) = \nu d^2 n$,

$$\begin{aligned} d(S) &\geq \sum_{\{x,y\}} d(S_{(x,y)}) - \sum_{\{x,y\} \neq \{p,q\}} d(S_{(x,y)} \cap S_{(p,q)}) \\ &\geq \binom{k}{2} (2m)^k \frac{\nu}{n} - 3 \binom{k}{4} (2m)^k \frac{\nu^2}{n^2} - 3 \binom{k}{3} (2m)^k \frac{\Delta \nu}{2mn} \\ &\geq \binom{k}{2} (2m)^k \frac{\nu}{n} \left(1 - \frac{k^2 \nu}{n} - \frac{k \Delta}{2m}\right) = \binom{k}{2} (2m)^k \frac{\nu}{n} (1 - o(1)). \end{aligned} \quad (14)$$

The last equality follows from the assumptions that $\Delta = o(m/\log^6 n)$ and $k \leq \log^3 n$, and the fact that $\nu \leq \Delta n/(2m) = o(n/\log^6 n)$. The bound (13) follows from (14). \square

5 Proof of coalescence time

Proof of Theorem 2

Let M_k be the time of the first meeting among $k \leq k^*$ particles in G , and let $\gamma = \gamma_k$ be the contraction of the diagonal set $S = S_k$. Using (7) for graph Γ and with $v = \gamma$, and Lemmas 3 and 4 we have, that the hitting time H_γ of γ from stationarity has expected value

$$\mathbf{E}_\pi(H_\gamma) \leq \frac{1}{\pi(\gamma)} \frac{1}{1 - \lambda_2} \leq \frac{1}{c_k k^2} \frac{n}{\nu} \frac{1}{1 - \lambda_2}. \quad (15)$$

Since $T_\Gamma = O(kT_G)$, and referring to (10), and Table 1,

$$\mathbf{E}(M_k) \leq O(kT_G) + (1 + o(1)) \mathbf{E}_\pi(H_\gamma) = O\left(\frac{1}{1 - \lambda_2} \left(k \log n + \frac{n}{\nu k^2}\right)\right). \quad (16)$$

Proof of Theorem 1

Let C_k be the time for $k \leq k^* = \log^3 n$ particles to coalesce. Then, using (16), we get

$$\mathbf{E}(C_k) \leq \sum_{s=2}^k \mathbf{E}(M_s) = O\left(\frac{1}{1-\lambda_2} \left(k^2 \log n + \frac{n}{\nu}\right)\right), \quad (17)$$

since $\sum_s (1/s^2) \leq \pi^2/6$ is constant.

Now we consider q particles, where $\log^3 n < q \leq n$. We prove that **whp** there cannot be a set of $k = \log^3 n$ particles which has not had a meeting by time t^* , where

$$t^* = k^2 (T_\Gamma + 3\mathbf{E}_\pi(H_{\gamma_k})) = O\left(\frac{1}{1-\lambda_2} \left(k^3 \log n + \frac{n}{\nu}\right)\right),$$

Let $\mathcal{P}(k) = \mathcal{P}(k, \mathbf{v})$ be the set of particles starting from vertices $\mathbf{v} = (v_1, \dots, v_k)$. Either there has been a meeting during the mixing time T_Γ , or if not, Lemma 1 applies for any $t \geq T_\Gamma$. The probability that the particles do not meet by time t is the same as the probability that the random walk in Γ_k starting from \mathbf{v} does not visit γ_k by time t . Therefore, applying Lemma 1 to graph Γ_k , we get

$$\bar{\rho}_k = \mathbf{Pr}(\text{No meeting among } \mathcal{P}(k) \text{ particles has occurred by } t^*) \leq e^{-k^2} = e^{-\log^6 n}.$$

Hence

$$\mathbf{Pr}(\exists \text{ a set } \mathcal{P}(k) \text{ having no meeting by } t^*) \leq \binom{n}{k} \bar{\rho}_k = O\left(e^{-\log^5 n + \log^4 n}\right).$$

Thus by t^* fewer than k particles remain, and an upper bound on the expected time for all particles to coalesce is

$$t^* + \mathbf{E}(C_k) + O\left(n^{-(\log^4 n)/2} t^*\right) = O\left(\frac{1}{1-\lambda_2} \left(k^3 \log n + \frac{n}{\nu}\right)\right). \quad (18)$$

The second term, $\mathbf{E}(C_k)$, (see (17)) is the expected coalescence time of the (at most) k particles remaining after t^* . The last term is the expected time to coalesce, restarting the process at t^* , as many times as needed, under the assumption that some set of $k = \log^3 n$ particles had not met at that time. It can be shown that if $\Delta = O(m^{1-\epsilon})$, then the n/ν term dominates the right-hand side of (18)

References

1. D. Aldous. Meeting times for independent Markov chains. *Stochastic Processes and their Applications*, 38 (1991) 185-193.
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. J. Aspnes. Randomized protocols for asynchronous consensus. *Distributed Computing*, 16 (2003) 165-176.
4. F.R.K. Chung: *Spectral Graph Theory*. American Mathematical Society, 1997.
5. C. Cooper, A. M. Frieze, T. Radzik. Multiple Random Walks in Random Regular Graphs. *SIAM J. Discrete Math.*, 23(4) (2009), 1738-1761.
6. J. T. Cox. Coalescing random walks and voter model consensus times on the torus in \mathbb{Z}^d . *The Annals of Probability*, 17(4) (1989), 1333-1366.
7. B. Doerr, T. Friedrich, T. Sauerwald: Quasirandom Rumor Spreading: Expanders, Push vs. Pull, and Robustness. *Proc. of ICALP'09*, pp. 366-377, 2009.
8. B. Doerr, L.A. Goldberg, L. Minder, T. Sauerwald, C. Scheideler: Stabilizing Consensus with the Power of Two Choices. *Manuscript*, 2010, full version available at www.upb.de/cs/scheideler.
9. P. Donnelly and D. Welsh. Finite particle systems and infection models. *Math. Proc. Camb. Phil. Soc.* 94 (1983), 167-182.
10. S. Hoory, N. Linial, A. Wigderson: Expander Graphs and their Applications. *Bulletin of the American Mathematical Society*, 43 (2006), 439-561.
11. A. Israeli and M. Jalfon. Token management schemes and random walks yield self stabilizing mutual exclusion. *Proc. of PODC'90*, pp. 119-131, 1990.
12. D. Levin, Y. Peres, E. Wilmer. Markov chains and Mixing Times. AMS 2009.
13. L. Lovász. Random walks on graphs: a survey. Bolyai Society Mathematical Studies, *Combinatorics, Paul Erdős is Eighty* (Vol. 2), (1993) 1-46.
14. T. Nakata, H. Imahayashi, M. Yamashita. Probabilistic local majority voting for the agreement problem on finite graphs. *Proc. of COCOON'99*, pp. 330-338, 1999.
15. Y. Hassin and D. Peleg. Distributed probabilistic polling and applications to proportionate agreement. *Information & Computation* 171 (2002), 248-268.
16. D. Kempe, A. Dobra, J. Gehrke. Gossip-based computation of aggregate information. *Proc. of FOCS'03*, pp. 482-491, 2003.
17. F. Kuhn, T. Locher, R. Wattenhofer. Tight bounds for distributed selection. *Proc. of SPAA'07*, pp. 145-153, 2007.
18. A. Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing* 1 (1992) 351-370.

Appendix: Proof of Theorem 3 for min-voting

Lemma 5. *Let V_q be the set of nodes with a certain opinion $q \in O_{\log^2 n}$ at a given time step t . Furthermore, let $W_q = \cup_{i=1}^q V_q$. If $\rho \leq |W_q| \leq \phi n/d$ for a certain constant ρ and opinion q , then with probability $1 - \exp(-\Theta(|W_q|))$ the size of the set W_q increases by a constant factor $c' > 1$ in a step.*

Proof. First, we show that with probability $1 - \exp(-\Theta(|W_q|))$, the nodes of W_q are contacted by at least $(2 - \alpha^3)|W_q|$ edges. We consider a vertex exposure Martingale sequence $(X_i)_{1 \leq i \leq |W_q|}$, representing the number of edges contacting the nodes $v_1, \dots, v_{|W_q|} \in V_q$. We assume that if more than $1/\alpha^2$ edges contact a node, then $1/\alpha^2$ of them are selected uniformly at random, which are allowed to keep their contacts. The rest is released, and the corresponding nodes may contact other nodes in the remaining steps of the exposure process. Since the released edges may choose their contacts outside W_q in the subsequent steps of the process, we only make the number of edges contacting W_q smaller. This implies that the Martingale satisfies the $1/\alpha^2$ -Lipschitz condition, and the Azuma-Hoeffding bound implies that

$$\Pr(|X_{|W_q|} - X_1| \geq \alpha^3 |W_q|) \leq 2 \exp\left(-\frac{\alpha^6 |W_q|^2}{2|W_q|/\alpha^4}\right) \leq \exp(-\Theta(|W_q|)). \quad (19)$$

Now we show that with probability $1 - \exp(-\Theta(|W_q|))$ there are at least $\alpha^2 |W_q|$ nodes, which contact exactly one node in W_q . Since $|W_q| \leq \phi n/d$, we have $|N(W_q)| \geq \alpha d |W_q|$. Then, the pigeonhole principle implies that at least $\alpha d |W_q|/2$ nodes in $N(W_q)$ have at most $2/\alpha$ neighbours in W_q . An arbitrary such node u contacts exactly one neighbour in W_q with probability at least

$$2 \frac{|E(u, W_q)|}{d} \left(1 - \frac{|E(u, W_q)|}{d}\right) \geq \frac{2}{d} - \frac{2}{d^2}.$$

Since there are at least $\alpha d |W_q|/2$ such nodes, and each of these acts independently, we use Chernoff bounds to conclude that at least $\alpha^2 |W_q|$ nodes will contact exactly one node in W_q , with probability at least $1 - \exp(-\Theta(|W_q|))$, whenever ρ is large enough. Then, with probability $1 - \exp(-\Theta(|W_q|))$ the nodes of W_q are contacted in total by at least

$$\left(\frac{2 - \alpha^3 - \alpha^2}{2} + \alpha^2\right) |W_q| \geq \frac{2 + \alpha^2 - \alpha^3}{2} |W_q|$$

nodes, which is larger than $|W_q|$ by at least a constant factor $c' > 1$.

Lemma 6. *After step $(\log \log n)^2$, for an arbitrary q we have $|W_q| = 0$ or $|W_q| \geq \log^2 n$, with probability $1 - 1/\log^{\omega(1)} n$. Furthermore, **whp** there will be at least one opinion $q \in O_{\log^2 n}$, which is contained in at least $\log^2 n$ nodes.*

Proof. For simplicity, we assume that $d = o(n/|W_q|)$, i.e., $|N(W_q)| \geq \alpha d |W_q|$. If $|W_q| \geq \rho$, then it follows from Lemma 5 that with probability $1 - \exp(-\Theta(|W_q|))$ the number of nodes in W_q increases by some factor $c' > 1$ in a step. Now, we consider the case when $|W_q| < \rho$. We show that with constant probability the number of nodes having some opinion at most q increases by at least 1 in a step. Since we assumed that d is large enough, and due to the fact that ρ does not depend on d , we may assume $\rho \ll d$. We know that each node has d neighbours, and a node of W_q can have at most $|W_q| - 1$ neighbours in W_q . Then, a node of $N(W_q)$ contacts a neighbour in W_q with probability at least $1 - (1 - 1/d)^2 = 2/d - 1/d^2$. Since $|N(W_q)| \geq d - \rho$, simple Chernoff bounds imply that whenever $d > 2\rho + 1$, there will be $\rho + 1$ nodes in $N(W_q)$ contacting a node with some opinion $1 \dots q$, with constant probability.

Now we describe the process by a simple random walk. Let $P' = (V', E')$ with $V' = \{u_\rho, u_{\rho+1}, \dots, u_{\log^2 n}\}$ be a path of length $O(\log^2 n)$, in which node u_i represents the case when $|W_q| = i$. Furthermore, node $u_{\log^2 n}$ represents the case when *at least* $\log^2 n$ nodes have some opinion $1 \dots q$. We also define path $P'' = (V'', E'')$ with $V'' = \{w_0, \dots, w_\rho\}$ of length $\rho + 1$, whose node w_i corresponds to the case $|W_q| = i$. Note that state ρ is contained in both paths. On path P' there is a transition from each node u_i to the set of nodes $u_{c'i + \Omega(1)}$ with probability $1 - \exp(-\Theta(i))$. On path P'' , there is a constant transition from each node w_i with $i > 0$ to the set of nodes $w_{i + \Omega(1)}$. Furthermore, the random walk is stopped as soon as node $u_{\log^2 n}$ is reached. Therefore, a random walk on the combined path $P' \cup P''$ satisfies the conditions of Lemma A.4 from [8], and we may conclude that the number of nodes with some opinion $1 \dots q$ will exceed $\log^2 n$ within $(\log \log n)^2$ steps, with probability $1 - \log^{-\omega(1)} n$. According to Lemma 5, once some W_q reaches $\log^2 n$, it will always increase within the next $(\log \log n)^2$ steps, with probability $1 - n^{-\omega(1)}$.

Concerning the second statement of the lemma, we observe that $|W_{\log^2 n}| = \log^2 n$ at step $(\log \log n)^2$, with probability $1 - \log^{-\omega(1)} n$. Thus, there is a q with $|W_q| = |V_q| \geq \log^2 n$ at step $(\log \log n)^2$, **whp**.

In the following, we assume that q is chosen such that $W_q = V_q$ at step $(\log \log n)^2$. That is, q is the smallest opinion, which is contained in the graph after $(\log \log n)^2$ steps. Then, it follows from the previous lemma that $|V_q| \geq \log^2 n$ at that time step, **whp**. Now we show the following lemma.

Lemma 7. *Assume that at some time step the number of nodes with the (currently) smallest opinion q in the graph is at least $\log^2 n$. Then, after additional $O(\log n)$ steps the number of nodes with this opinion is at least $n/2$, with probability $1 - n^{-\omega(1)}$.*

Proof. As long as $|V_q| \leq \phi n/d$, the statement of the lemma follows from Lemma 5. Therefore, we only consider the case $\phi n/d \leq |V_q| \leq n/2$ here. First, we assume that there are $n/40$ nodes in $V \setminus V_q$, which have more than $\alpha d|V_q|/n$, but less than $d - \alpha d|V_q|/n$ neighbours in V_q . Then, using similar Martingale arguments as in Lemma 5, Equation (19) implies that there are at least $(2 - \alpha/42)|V_q|$ edges contacting nodes in V_q . Furthermore, one of the $n/40$ nodes above contacts set V_q with exactly one edge with probability $2 \frac{\alpha d|V_q|}{nd} \left(1 - \frac{\alpha d|V_q|}{nd}\right)$. Since there are at least $n/40$ such nodes, there will be in expectation $2 \frac{\alpha|V_q|}{40} \left(1 - \frac{\alpha|V_q|}{40}\right)$ nodes contacting V_q with exactly one edge. Applying simple Chernoff bounds, at least $2\alpha|V_q|/41$ nodes will contact V_q with exactly one edge, with probability $1 - n^{-\Omega(1)}$. Putting everything together, we obtain that within one step, the number of nodes with opinion q increases by a constant factor larger than 1, with probability $1 - n^{-\Omega(1)}$.

Now we show that there must exist $n/40$ nodes in $V \setminus V_q$ which have more than $\alpha d|V_q|/n$, but less than $d - \alpha d|V_q|/n$ neighbours in V_q . We know that

$$|E(V_q, V \setminus V_q)| \leq \frac{d|V_q||V \setminus V_q|}{n} + \lambda d \sqrt{|V_q||V \setminus V_q|}.$$

In the rest of the proof, we assume for simplicity that $\lambda = o(1)$ (a more sophisticated analysis leads to the same result if $\lambda \ll 1/\sqrt{d}$, if $d = O(1)$). Assume that there are less than $n/20$ nodes in $V \setminus V_q$ with less than $d - \alpha d|V_q|/n$ neighbours in V_q . Then, the number of edges between V_q and $V \setminus V_q$ is at least

$$\begin{aligned} \left(|V \setminus V_q| - \frac{n}{20}\right) \left(d - \alpha d \frac{|V_q|}{n}\right) &\geq \left(|V \setminus V_q| - \frac{n}{20}\right) d \left(1 - \frac{\alpha}{2}\right) \\ &\geq \frac{9|V \setminus V_q|}{10} d \left(1 - \frac{\alpha}{2}\right), \end{aligned} \tag{20}$$

where $\alpha/2$ can be arbitrarily small. On the other side

$$|E(V_q, V \setminus V_q)| \leq \frac{d|V_q||V \setminus V_q|}{n} (1 + o(1)) \leq \frac{|V \setminus V_q|}{2} d (1 + o(1))$$

which is much smaller than the value in Equation (20) leading to contradiction.

On the other hand, the number of nodes in $V \setminus V_q$ which have less than $\alpha d|V_q|/n$ neighbours in V_q is at most

$$\frac{\beta n^2}{d|V_q|} \leq \frac{\beta n}{\phi} < \frac{n}{40}$$

if $\beta/\phi < 1/40$. This implies that at least $n/40$ nodes in $V \setminus V_q$ have more than $\alpha d|V_q|/n$, but less than $d - \alpha d|V_q|/n$ neighbours in V_q .

Now we consider the case $|V_q| > n/2$.

Lemma 8. *Assume that at some time step the number of nodes with the (currently) smallest opinion q in the graph is at least $n/2$. Then, after additional $O(\log n)$ steps all nodes will have opinion q , **whp**.*

Proof. The proof basically follows the arguments given in the previous lemmas. Let \overline{V}_q represent the set of nodes which do not have opinion q . We show that within one step, \overline{V}_q decreases by a constant factor, with probability $1 - n^{-\Omega(1)}$. Again, we consider two cases. First, we assume that $|\overline{V}_q| \geq \phi n/d$. Then, using the same arguments as in Lemmas 7 and 5, there are at most $(2 + \alpha/42)|\overline{V}_q|$ edges contacting nodes in \overline{V}_q . On the other hand, Lemma 7 also implies that at least $n/40$ nodes in $V \setminus \overline{V}_q$ have more than $\alpha d|\overline{V}_q|/n$, but less than $d - \alpha d|\overline{V}_q|/n$ neighbours in \overline{V}_q . These statements combined with the arguments of Lemma 5 imply that at least $2\alpha|\overline{V}_q|/41$ nodes will contact \overline{V}_q with exactly one edge, with probability $1 - n^{-\Omega(1)}$. Then, the total number of nodes contacting \overline{V}_q with both edges is at most

$$\left(2 + \frac{\alpha}{42} - \frac{2\alpha}{41}\right) \frac{|\overline{V}_q|}{2},$$

with probability $1 - n^{-\Omega(1)}$. Thus, with probability $1 - n^{-\Omega(1)}$ the number of nodes, which do not have opinion q , decreases by a constant fraction.

Now we consider the case $|\overline{V}_q| < \phi n/d$. Using the same Martingale approach as in Lemma 5, Equation (19) implies that with probability $1 - \exp(-\Theta(|\overline{V}_q|))$ there are at most $(2 + \alpha^3)|\overline{V}_q|$ edges contacting nodes in \overline{V}_q . On the other hand, with probability $1 - \exp(-\Theta(|\overline{V}_q|))$ there are at least $\alpha^2|\overline{V}_q|$ nodes in V_q , which contact \overline{V}_q with exactly one edge. Thus, the total number of nodes which contact \overline{V}_q with both edges is at most $(2 + \alpha^3 - \alpha^2)|\overline{V}_q|/2$, with probability $1 - \exp(-\Theta(|\overline{V}_q|))$. This implies that after $\log n$ steps, all but at most $\log^2 n$ nodes will have the same opinion q , with probability $1 - n^{-\Omega(1)}$.

For the case $|\overline{V}_q| \leq \log^2 n$, we use the random walk approach described in Lemma 5. However, in this case, the nodes u_i and w_i in the paths P' and P'' ,

respectively, represent the cases when i nodes do not have opinion q . Then, a node u_i makes a transition to some node $u_{i/c' - \Omega(1)}$ with probability $1 - \exp(-\Theta(i))$, where $c' > 1$. Furthermore, a node w_i makes a transition to a node $w_{i - \Omega(1)}$ with constant probability. According to the arguments above and to Lemma 5, the random walk will reach node u_0 within $(\log \log n)^2$ steps, with probability $1 - \log^{-\omega(1)} n$. Thus, the lemma follows.