King's Research Portal

DOI:

[10.1002/rsa.20822](https://doi.org/10.1002/rsa.20822)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Cooper, C., McDowell, A., Radzik, T., Rivera Aburto, N., & Shiraga, T. (2018). Dispersion processes. *RANDOM STRUCTURES AND ALGORITHMS*, 53(4), 561-585. <https://doi.org/10.1002/rsa.20822>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Dispersion processes*

Colin Cooper[†] Andrew McDowell[‡] Tomasz Radzik[§] Nicolás Rivera[¶]
Takeharu Shiraga^{||}

September 7, 2018

Abstract

We study a synchronous process called *dispersion*. Initially M particles are placed at a distinguished *origin vertex* of a graph G . At each time step, at each vertex v occupied by more than one particle at the beginning of this step, each of these particles moves to a neighbour of v chosen independently and uniformly at random. The process ends at the first step when no vertex is occupied by more than one particle.

For the complete graph K_n , for any constant $\delta > 1$, with high probability, if $M \leq n/2(1 - \delta)$, the dispersion process finishes in $O(\log n)$ steps, whereas if $M \geq n/2(1 + \delta)$, the process needs $e^{\Omega(n)}$ steps to complete, if ever. A lazy variant of the process exhibits analogous behaviour but at a higher threshold, thus allowing faster dispersion of more particles.

For paths, trees, grids, hypercubes and Abelian Cayley graphs of large enough size, we give bounds on the time to finish and the maximum distance traveled from the origin as a function of M .

Keywords: Random processes on graphs; Dispersion of particles; Random walk

*This work was supported by EPSRC grant EP/M005038/1, “Randomized algorithms for computer networks”, and Becas CHILE.

[†]Department of Informatics, King’s College London, UK. colin.cooper@kcl.ac.uk

[‡]Department of Informatics, King’s College London, UK. andrew.mcdowell@kcl.ac.uk

[§]Department of Informatics, King’s College London, UK. tomasz.radzik@kcl.ac.uk

[¶]Department of Informatics, King’s College London, UK. nicolas.rivera@kcl.ac.uk

^{||}Department of Information and System Engineering, Chuo University, Japan. shiraga@ise.chuo-u.ac.jp

1 Introduction

A dispersion process can be described as follows. Initially a number of identical particles are located at a single vertex of a graph. The particles move apart in a distributed fashion until no more than one particle occupies any vertex. When this occurs we say the particles are dispersed. We require the behaviour of the particles during dispersion to be identical, and that no communication, prioritization or other symmetry breaking occurs.

The method we use to achieve this, hereafter called dispersion, works as follows. The process is synchronous and proceeds in discrete steps. At any step, whenever two or more particles occupy the same vertex, each particle moves independently to a random neighbour. If only a single particle occupies a vertex, it stays there until another particle arrives. The process ends at the first step when no vertex is occupied by more than one particle.

At each step of the process at which a particle moves, it makes one step of an independent random walk. However, the steps at which a particle moves are correlated with the arrival of other particles. The particles make random walks which stop and start, and up to the step at which a particle finally stops, each time it moved another particle moved with it.

We analyse this process on a number of different graphs, including the complete graph, the star graph, paths, grids, the hypercube, Cayley graphs and regular trees.

Interacting particle systems. Processes in which particles make independent random walks subject to conditions which describe the interactions between the particles have been widely studied in statistical mechanics as models of physical systems, a notable early example of this work being the kinetic theory of gases. In a discrete context of particles moving on a finite or countably infinite graph, three classical models are the contact process, the voter model and its dual the coalescing random walk, and the exclusion process. See, e.g., the monographs of Liggett [9, 10] for the development of these topics.

In an exclusion process, it is an initial condition that each vertex is occupied by at most one particle. The process is asynchronous, and at any step a randomly chosen particle moves to a randomly chosen neighbour provided the destination vertex is unoccupied. If the destination vertex is occupied, the particle stays where it is. The continuous time exclusion process was introduced by Spitzer [11]. The following description of the process is from Liggett [9]. An exclusion process on a graph is an asynchronous process in which particles with independent unit exponential holding times try to move around the graph. An activated particle at vertex x jumps to y with probability $p(x, y)$. If vertex y is already occupied, the particle remains at x and begins a new exponential holding time.

In terms of random processes whose aim is to disperse particles on a network, perhaps the best known is Internal Diffusion Limited Aggregation (IDLA). In this process particles start from an origin vertex one at a time. The next particle does not start until the previous particle has stopped moving. Once activated, a particle moves randomly until it reaches an unoccupied vertex, which it then occupies permanently and does not move any further. A particle which arrives at an occupied vertex continues to walk randomly until it reaches a vacant site at which point it stops.

The IDLA process was introduced by Diaconis and Fulton [4], who considered the integer line and gave the limiting length and position of the segment of the line covered by the particles. The induced subgraph formed by IDLA on any graph is connected, as no particle goes further than the next unoccupied neighbour. Lawler, Bramson and Griffeath [7] subsequently analysed IDLA on

d -dimensional grids. For the two dimensional grid they proved the limiting shape is a disk. The shapes made by the rotor-router analogue of IDLA have also been studied. For two dimensional grids, Levine and Peres [8] prove the limiting rotor-router shape is also a disk.

Limitations of IDLA as a distributed model of particle dispersion include the fact that the process is asynchronous, and that the next particle is only activated after the previous one stops. This implies communication, and the obvious asymmetry in particle behaviour. The first two objections (asynchronous, communication) can be overcome using the following synchronous version of IDLA. If a single particle occupies a vertex at any step it halts permanently, whereas if two or more particles arrive simultaneously, they all move independently to a random neighbour. If particles arrive at an already occupied vertex, the new arrivals move independently to a random neighbour.

However, this process does not respect the requirement for symmetry of behaviour of indistinguishable particles. Particles still behave in an asymmetric fashion, stopping permanently on becoming the unique occupant of a vertex.

In the dispersion process particles act independently and in an indistinguishable manner. All particles at any multiply occupied vertex behave identically in a synchronous fashion. Particles stop and start, and a stopped particle cannot know when it will move again. This stop-start behaviour makes it difficult to relate the walk steps of a particle to the underlying steps of the dispersion process, and thus presents an additional obstacle to analysis.

The dispersion process differs from the type of methods considered for dispersing robots or sensors, in that we do not explicitly require the particles to disperse uniformly throughout the network, but merely to move away from one another and establish a personal space. Neither do we allow communication or collaboration which are usually a feature of robot dispersion. This means that the degree of self-organization is less than required for swarm systems. Random dispersion of swarms is considered by Beal [2, 3]. The particles use Lévy Flights to move a random biased distance d with probability proportional to $1/d$ (within some large finite range).

Dispersion of particles can be used to model some load balancing processes in distributed computing systems, with the vertices of the graph representing homogeneous network resources and the particles representing tasks or users. If a resource v is overloaded with users, then they may decide to move to neighbouring resources for better service. It may be difficult for the users to acquire advance information about the current levels of usage of neighbouring resources, and the users may be unable, or unwilling, to cooperate to resolve overloading. In which case, they might just as well try a random neighbouring resource. This scenario was considered by Ackermann *et al.* [1] in the context of systems where users have *quality-of-service demands*.

Definitions. We analyse the synchronous dispersion process in which M particles are initially placed at a single vertex of a graph, which we call the origin vertex. If two or more particles occupy the same vertex at the end of step t , then all particles at that position move independently to a random neighbour at step $t + 1$. Thus it can be that (by chance) the particles move to the same place and have to move again at step $t + 2$, and so on. The process ends once the particles have all stopped moving. This occurs when all vertices are occupied by either one or no particles. Trivially, for the process to end the graph must have at least as many vertices as there are particles.

We distinguish between the dispersion process and the random walks made by the individual particles. We call the steps $t = 0, 1, 2, \dots$ of the dispersion process, *time steps* or (*dispersion*) *process steps*. We refer to the steps of the random walk made by an individual particles as the *walk steps* of the particle. For a given particle the walk steps are indexed by the number of times the

particle moves, which can be much less than the number of steps made by the dispersion process, and are embedded in the process steps in an irregular fashion. The time steps t of the dispersion process go on forever, but after dispersion the particle locations do not change.

We are interested in typical properties of the dispersion process as a function of the number of particles M , and the size n and structure of the underlying graph G . We define two random variables the dispersion time $\mathbf{T}_{\text{Disp}}$ and the dispersion distance $\mathbf{D}_{\text{Disp}}$. The dispersion time is the number of steps of the process needed to disperse the particles. The dispersion distance is the maximum graph distance of any particle from the origin at dispersion.

Throughout this paper we use the notation ω to denote a quantity which tends slowly to infinity with increasing n or M .

A threshold result for dispersion time on the complete graph K_n . The complete graph K_n exhibits a threshold in dispersion time from $O(\log n)$ when the number of particles M is at most $(1 - \delta)n/2$, to $e^{\Omega(n)}$ when the number of particles M is at least $(1 + \delta)n/2$, where $\delta > 0$ is an arbitrarily small constant. The following theorem is proven in Section 2.

Theorem 1. *For the complete graph K_n and the star S_n the following hold for any constant $\delta > 0$.*

- (i) *If the number of particles M satisfies $M/n \leq (1 - \delta)/2$, then with probability $1 - O(1/n)$, the dispersion process terminates in $\mathbf{T}_{\text{Disp}} = O(\log n)$ steps.*
- (ii) *If the number of particles M satisfies $(1 + \delta)/2 \leq M/n < 1$, then there is a constant $c = c(\delta) > 0$ such that the probability that $\mathbf{T}_{\text{Disp}} \leq e^{cn}$ is less than e^{-cn} .*

In order to disperse more than $n/2$ particles in $O(\log n)$ process steps, we introduce a variant process which we call *lazy dispersion*. For some $0 < p \leq 1$, any particle which occupies a vertex containing other particles, moves with probability p and stays put at its current vertex with probability $1 - p$. A particle which is the sole occupant of a vertex does not move.

The following theorem shows that choosing a value of $p < 1$ allows logarithmic dispersion for a larger value of M . More precisely, the threshold M/n which separates fast and slow dispersion moves upwards from $1/2$ to $1 - p/2$.

Theorem 2. *Given $0 < p \leq 1$, and $\alpha > 0$, both of which may depend on n , the lazy dispersion of M particles on the complete graph K_n behaves in the following way.*

- (i) *If $M = (1 - \frac{p}{2} - \alpha)n$, then with probability $1 - O(1/n)$ the dispersion process terminates in $\mathbf{T}_{\text{Disp}} = O((p\alpha)^{-1} \log n)$ steps.*
- (ii) *If $M = (1 - \frac{p}{2} + \alpha)n$, then there exists a constant $c > 0$ (independent of p and α), such that the probability that $\mathbf{T}_{\text{Disp}} \leq e^{cnp^3\alpha^4}$ is less than $e^{-cnp^3\alpha^4}$.*

Observe that the more detailed Theorem 2 tells us also what happens in Theorem 1 (where $p = 1$) for complete graphs, if $0 < \delta = o(1)$. If $M/n = (1 - \delta)/2$, then $\mathbf{T}_{\text{Disp}} = O(\delta^{-1} \log n)$, and if $M/n = (1 + \delta)/2$, then $\mathbf{T}_{\text{Disp}} = e^{\Omega(n\delta^4)}$ with probability at least $1 - e^{-\Omega(n\delta^4)}$.

The threshold in Theorems 1 and 2 relies implicitly on the arrival of the process at a fixed point (as a function of M, n) at which the expected number of happy and unhappy particles does not change at the next and subsequent steps.

Results for dispersion time on Cayley graphs. We give results for dispersion on grids, the hypercube and indeed any other symmetric Cayley graph of an Abelian group (Abelian Cayley

graph). In such a graph, the simple random walk transition at any vertex (an element of the underlying group) is determined by sampling uniformly from a symmetric generator set S which defines the graph (symmetric means that if $g \in S$, then also $-g \in S$). Transitions at vertex u are made to $v = u + g$, where “+” denotes the group operation and $g \in S$ is the group element which labels edge (u, v) . The edge (v, u) from v to $u = v + (-g)$ is also present. For example, on the line the transitions are defined by $S = \{-1, +1\}$. Denoting by X_t the position of the random walk at step t , we have $X_{t+1} = X_t \pm 1$, equiprobably. For the two dimensional infinite grid, the transitions are defined by $S = \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$.

Because many Cayley graphs are bipartite, a simple random walk on the graph is periodic and does not have a well defined mixing time T . However, for a finite, connected, regular bipartite graph, the corresponding two-step walk converges to the uniform distribution on the relevant side of the bipartition.

Theorem 3. (i) Let G be an infinite Abelian Cayley graph. Let $R(2t)$ be the expected number of returns to the origin in $2t$ steps by a simple random walk on G . With probability at least $1 - 1/\omega$, a system of M particles disperses on G in $\mathbf{T}_{\text{Disp}} = O(\omega M^2 R(2t))$ process steps.

(ii) Let $G = (V, E)$ be a finite n -vertex Abelian Cayley graph. If G is bipartite, let $n' = n/2$, and if G is non-bipartite, let $n' = n$. Let P be the transition matrix of a simple random walk on G , and let T be such that for all $s \geq T$, or for all even $s \geq T$ in case of bipartite graphs, and for all $u \in V$, $|P^s(u, u) - 1/n'| \leq 1/2n'$. With probability at least $1 - 1/\omega$, a system of $M = o(\sqrt{n/\omega})$ particles disperses on G in $\mathbf{T}_{\text{Disp}} = O(TM^2)$ process steps.

The proof of Theorem 3 part (i) is given in Section 5, and part (ii) in Section 6. The following examples are applications of Theorem 3.

For infinite Abelian Cayley graphs. We give values of $\mathbf{T}_{\text{Disp}}$ for grids of dimension $d \geq 1$. More detail for the 2-dimensional grid is given in Section 5.

The values of $R(2t)$ for the line, 2-dimensional grid, and grids of dimension $d \geq 3$ are $\Theta(\sqrt{t})$, $\Theta(\log t)$ and $\Theta(1)$ respectively. This gives values of $\mathbf{T}_{\text{Disp}} = O(\omega^2 M^4)$ for the line, $\mathbf{T}_{\text{Disp}} = O(\omega M^2 \log M)$ for the 2-dimensional grid and $\mathbf{T}_{\text{Disp}} = O(\omega M^2)$ for grids of dimension $d \geq 3$.

For finite Abelian Cayley graphs. We give values of $\mathbf{T}_{\text{Disp}}$ for the hypercube and n -cycle. A detailed proof for the hypercube is given in Section 6.

The hypercube H_d on $n = 2^d$ vertices consists of vertices labeled as vectors in $\{0, 1\}^d$ and edges uv between vertices u and v whenever the vertex labels differ in a single coordinate.

The value of T for the hypercube is $T = O(\log^2 n)$, and for the n -cycle $T = O(n^2 \log n)$. Thus, provided $M = o(\sqrt{n})$, with probability $1 - o(1)$ for the hypercube $\mathbf{T}_{\text{Disp}} = O(M^2 \log^2 n)$ and for the n -cycle $\mathbf{T}_{\text{Disp}} = O(M^2 n^2 \log n)$.

Results for dispersion on paths and regular trees. We give upper bounds on $\mathbf{T}_{\text{Disp}}$ and $\mathbf{D}_{\text{Disp}}$ for paths and sufficiently large k -regular trees. The proof for trees is given in Sections 3 and for paths in 4.

Theorem 4. For path (or cycle) of length at least $12M \log M$, and M particles initially placed at the central vertex of the path, the following holds w.h.p. for any $\varepsilon > 0$. The dispersion time $\mathbf{T}_{\text{Disp}} = \Omega(M^2)$, and is bounded above by $\mathbf{T}_{\text{Disp}} \leq 4(1 + \varepsilon)M^3 \log M$. The maximum distance $\mathbf{D}_{\text{Disp}}$ of any particle from the origin is bounded by

$$\lfloor M/2 \rfloor \leq \mathbf{D}_{\text{Disp}} \leq 4(1 + \varepsilon)M \log M. \quad (1)$$

We note that recently Frieze and Pegden, [6], improved the upper bound for paths in Theorem 4 to $\Theta(M)$ using an argument which exploits the natural linear ordering of vertices on a path.

Theorem 5. *Let $k \geq 3$, and $\varepsilon > 0$, constant. For a k -regular tree of depth at least $(2 + \varepsilon) \log_{k-1} M$, with M particles initially placed at the central vertex of the tree, w.h.p. the dispersion time $\mathbf{T}_{Disp} = O(M \log_{k-1} M)$. The maximum distance $\mathbf{D}_{Disp}$ from the origin is w.h.p. bounded by*

$$\left(1 + \frac{1}{2(\log_{k-1} k) - 1} - \varepsilon\right) \log_{k-1} M \leq \mathbf{D}_{Disp} \leq \left(\frac{5}{3} + \frac{1}{3 \log_{k-1} k} + 2\varepsilon\right) \log_{k-1} M. \quad (2)$$

To make the nature of the relation between $\mathbf{D}_{Disp}$ and $\log_{k-1} M$ clearer, observe that for any $k \geq 3$ and an appropriately small ε , the coefficient at $\log_{k-1} M$ in the upper bound in (2) is less than 2. For the lower bound in (2), observe that for growing k , $(\log_{k-1} k) - 1 \sim 1/(k \ln k)$, so the coefficient there is at least $2 - 1/k$, for sufficiently large k and small ε . Thus there is a constant k_0 such that for any constant $k \geq k_0$, w.h.p.

$$2 - \frac{1}{k} \leq \frac{\mathbf{D}_{Disp}}{\log_{k-1} M} \leq 2. \quad (3)$$

Proof methodology. It is a condition of the dispersion process that the particles make independent random walks whenever they move on the underlying graph G . To remove any suspicion of correlation between the walks we adopt the following device, and predetermine the movements the particles will take when they move. For each particle, independently predetermine an infinite random walk on G . Whenever the particle is required to move in the dispersion process, it reads the next movement from its own random walk and follows it. In this way the walk is independent of the dispersion process and the movement of any other particle. However the number of steps taken by the walk at a given step of the dispersion process, and the step of the walk at which the particle will next move or stop forever is entirely determined by the underlying dispersion process.

Our methods for bounding the dispersion time $\mathbf{T}_{Disp}$ and dispersion distance $\mathbf{D}_{Disp}$ are based (in various ways) on the underlying random walks made by the particles.

The first method uses the number of meetings between particles, which allows us to find a bound on the total dispersion time of the process that holds with high probability. The transitions of random walks on Abelian Cayley graphs can be expressed in a simple way, e.g., $\{-1, +1\}$ on the line, which allows us to treat the number of meetings of two particles as the number of returns to the origin of a combined walk, by reversing the movements of one of the particles. The number of meetings a particle has (as measured by total returns to the origin by particle pairs, over all particles it meets) is an upper bound on the number of walk steps it has made. If the number of returns to the origin grows too slowly with the assumed number of walk steps of the particle, this leads to a contradiction of the assumption that this many steps could have been made.

The second method, which we use on the k -regular tree, is to take advantage of the branching structure and the fact that there are many distinct vertices a large distance from the origin. In particular, for a particle to first reach a distance d from the origin and continue moving, another particle must also visit the same vertex at the same time.

The first method allows us to bound the total dispersion time $\mathbf{T}_{Disp}$ of the process, which we may be able to use in turn to bound the distance a particle can move. Conversely, the second gives direct bounds on the distance $\mathbf{D}_{Disp}$ any particle can move, from which we may also be able to deduce bounds on the dispersion time.

The hypercube is a graph which satisfies both of these approaches, since it is an Abelian Cayley graph but also has good local tree-like expansion, so it is a good example of how these two methods can be applied. In Section 6, we derive bounds of $\mathbf{T}_{\text{Disp}} = O(M^2 \log^2 n)$ and $\mathbf{T}_{\text{Disp}} = O(M \log^3 n)$ respectively, by using these two methods.

2 Dispersion on the complete graph

A proof of Theorem 1 for the complete graph K_n is given in Section 2.1. To keep the proof tidy, we first analyse the case of K_n with loops, that is, we assume that a particle jumping from a vertex v lands back on v with probability $1/n$. Details for K_n without loops are given afterwards.

The behavior of dispersion on the star graph S_n (one centre vertex and n leaves) is analogous to that on K_n (without loops) in the following way. At alternate steps of the dispersion process on S_n , any active particles congregate at the central vertex, and then jump to a random leaf. Thus, if the origin is the centre vertex (resp. a leaf), then in steps 2, 4, 6, ... (resp. in steps 3, 5, 7, ...), all particles are in the leaves of S_n with the same distribution as the distribution of the particles in steps 2, 3, 4, ... in dispersion on K_n .

A proof of Theorem 2 is given in Section 2.2.

2.1 Proof of Theorem 1

Case of K_n with loops. Let M be the total number of particles. A particle is happy at a given time step if it is the only particle at its current vertex, otherwise it is unhappy. Let $H(t)$ be the number of happy particles at step t , and $U(t)$ the unhappy ones. Thus $H(t) + U(t) = M$. The process ends when $H(t) = M$.

In what follows, we bound the value of $H(t+1)$ given the values of $H(t)$ (and thus $U(t)$) at step t . At each time step t any unhappy particle moves to a random vertex $v \in [n]$. The particles which are happy do not move. Suppose $U(t) > 0$. At the next step there are $X = X(t)$ previously happy particles which became unhappy because (unhappy) particles landed on top of them. Also $Y = Y(t)$ previously unhappy particles became happy by being the only particle to move to one of the $n - H(t)$ unoccupied vertices. This gives $H(t+1) = H(t) - X + Y$. We obtain $\mathbf{E}(X(t) \mid H(t))$, $\mathbf{E}(Y(t) \mid H(t))$ and hence $\mathbf{E}(H(t+1) \mid H(t))$.

To simplify notation we do not explicitly state the conditioning on $H(t)$, and we abbreviate $H(t)$ and $H(t+1)$ to H and H' , respectively, and similarly abbreviate $U(t)$ to U etc.

The properties of X, Y are as follows. We randomly allocate U balls (the unhappy particles) to n boxes (the vertices), of which H are non-empty. The number X of the H non-empty boxes receiving at least one ball and the number Y of empty boxes receiving exactly one ball have expected values

$$\mathbf{E}X = H \left(1 - \left(1 - \frac{1}{n} \right)^U \right), \quad (4)$$

$$\mathbf{E}Y = U \left(\frac{n-H}{n} \right) \left(1 - \frac{1}{n} \right)^{U-1}. \quad (5)$$

The concentration of $H' = H - X + Y$ follows from considering the Doob martingale $B_i = \mathbf{E}_{Z_{i+1}, \dots, Z_U}[H' \mid Z_1, Z_2, \dots, Z_i]$, where Z_i is the box (vertex) chosen by the i -th unhappy ball. Thus

$B_0 = \mathbf{E}[H']$, $B_U = H'$, $\mathbf{E}B_i = B_{i-1}$, and $|B_i - B_{i-1}| \leq 2$ because a difference in choice of bin by ball i (with all other choices remaining the same) can only alter the value of H' by at most 2. The Azuma-Hoeffding inequality for the concentration of the values of martingales implies

$$\mathbf{Pr}(|H' - \mathbf{E}H'| \geq \lambda) \leq 2 \exp\left(-\frac{\lambda^2}{8U}\right). \quad (6)$$

Let $\Delta H = H(t+1) - H(t) = Y - X$.

$$\mathbf{E}\Delta H = \frac{U}{n}(n-H) \left(1 - \frac{1}{n}\right)^{U-1} - H \left(1 - \left(1 - \frac{1}{n}\right)^U\right) \quad (7)$$

$$= \left(1 - \frac{1}{n}\right)^U \left(U + H - \frac{U(H-1)}{n-1}\right) - H. \quad (8)$$

Case (i): $M \leq (1 - \delta)n/2$.

For $U \geq 0$, $(1 - 1/n)^U \geq (1 - U/n)$. Substituting this into (8), and using $U + H = M$ gives

$$\begin{aligned} \mathbf{E}\Delta H &\geq \left(1 - \frac{U}{n}\right) \left(U + H - \frac{UH}{n}\right) - H \\ &= U \left(1 - \frac{M}{n} - \frac{M-U}{n} \left(1 - \frac{U}{n}\right)\right). \end{aligned}$$

Thus

$$\mathbf{E}(U(t+1) \mid U(t)) \leq \frac{U}{n} \left(M + (M-U) \left(1 - \frac{U}{n}\right)\right) \leq U(t) \frac{2M}{n}.$$

Now use $M \leq (n/2)(1 - \delta)$ and iterate to get

$$\mathbf{E}U(t) \leq U(0) (1 - \delta)^t \leq U(0) \exp(-t\delta).$$

Choosing

$$t = (2/\delta) \log n \quad (9)$$

gives $\mathbf{E}U(t) = O(1/n)$ and thus $\mathbf{Pr}(U(t) \geq 1) = O(1/n)$. Hence $\mathbf{T}_{\text{Disp}} \leq t$ with probability $1 - O(1/n)$.

Case (ii): $M \geq (1 + \delta)n/2$.

Let $H(t) \leq (n/2)(1 + \delta/2)$. We prove below that

$$\mathbf{E}H(t+1) \leq (1 + o(1)) \frac{n}{2} \left(1 + \frac{\delta}{2} \left(1 - \frac{3\delta}{8}\right)\right). \quad (10)$$

By the concentration of $H(t+1)$ (see discussion below (5)),

$$\mathbf{Pr}\left(H(t+1) \geq \frac{n}{2} \left(1 + \frac{\delta}{2}\right)\right) \leq e^{-2cn}, \quad (11)$$

for some constant $c = c(\delta) > 0$. To disperse the particles requires H to equal $M = (n/2)(1 + \delta)$. The Inequality (11) implies, however, that H remains below $(n/2)(1 + \delta/2)$ for e^{cn} steps with probability at least $(1 - e^{-2cn})^{e^{cn}} \geq 1 - e^{-cn}$.

Proof of Equation (10). Let $H(t) = (n/2)(1 + \varepsilon)$, where $-1 \leq \varepsilon \leq \delta/2$. From (8), with $U = (n/2)(\delta - \varepsilon)$, we have that

$$\mathbf{E}H(t+1) \leq (1 + o(1))e^{-\frac{U}{n}} \left(M - \frac{U(M-U)}{n} \right) = (1 + o(1))\frac{n}{2}A(\varepsilon, \delta),$$

say, where

$$\begin{aligned} A(\varepsilon, \delta) &\leq e^{-\frac{1}{2}(\delta-\varepsilon)} \left((1+\delta) - \frac{1}{2}(\delta-\varepsilon)(1+\varepsilon) \right) \\ &= e^{\varepsilon/2-\delta/2} \left(1 + \frac{\delta}{2} + \frac{\varepsilon}{2}(1-\delta) + \frac{\varepsilon^2}{2} \right). \end{aligned}$$

Thus $A(\varepsilon, \delta)$ is monotone increasing in ε and for $\varepsilon \leq \delta/2$,

$$A(\varepsilon, \delta) \leq A(\delta/2, \delta) = e^{-\delta/4} \left(1 + \frac{3\delta}{4} - \frac{\delta^2}{8} \right).$$

For $0 \leq x \leq 1$, $e^{-x} \leq 1 - x + x^2/2$, so that

$$\begin{aligned} A(\delta/2, \delta) &\leq \left(1 - \frac{\delta}{4} + \frac{\delta^2}{32} \right) \left(1 + \frac{3\delta}{4} - \frac{\delta^2}{8} \right) \\ &= 1 + \frac{\delta}{2} - \frac{\delta^2}{32} \left(9 - \frac{7\delta}{4} + \frac{\delta^2}{8} \right). \end{aligned}$$

However $9 - \frac{7\delta}{4} + \frac{\delta^2}{8}$ is monotone decreasing in δ for $0 \leq \delta \leq 1$ and

$$A(\delta/2, \delta) \leq 1 + \frac{\delta}{2} \left(1 - \frac{6\delta}{16} \right).$$

Case of K_n without loops. For $\mathbf{E}X$ in (4), the value of $(1 - 1/n)$ becomes $(1 - 1/(n-1))$. The effect on Y is to slightly increase the value of $\mathbf{E}Y$ in (5) as follows. Let $V(H)$ be the happy vertices, and $u(v)$ the number of unhappy particles at $v \in V$. The upper tail of $u(v)$ is stochastically dominated by $Z \sim \text{Bin}(M, 1/(n-1))$. Using a Chernoff bound that $\mathbf{Pr}(Z \geq \alpha\mu) \leq (e/\alpha)^{\alpha\mu}$

$$\mathbf{Pr}(Z \geq n/\log n) \leq \left(\frac{eM \log n}{n(n-1)} \right)^{n/\log n} = e^{-\Theta(n)}.$$

Thus (w.h.p.)

$$\begin{aligned} \mathbf{E}Y &= \sum_{i \in U} \sum_{v \in [n] - V(H) - \{v_i\}} \frac{1}{n-1} \left(1 - \frac{1}{n-1} \right)^{U-u(v)-1} \\ &= (1 + O(1/\log n)) U \left(1 - \frac{H}{n} \right) \left(1 - \frac{1}{n} \right)^U, \end{aligned}$$

where v_i is the vertex currently occupied by unhappy particle i . The rest of the proof is the same.

2.2 Lazy Dispersion on K_n

We have shown that on the complete graph, the dispersion process disperses the particles in logarithmic time (w.h.p.) if the number of particles is less than half the number of vertices. If the number of particles is more than half the number of vertices, there is a double exponential leap, as it now requires exponential time (w.h.p.) to disperse the particles. We next show that, perhaps counter intuitively, slowing down the particles can allow the process to disperse more quickly. More precisely, we show that if instead of all unhappy particles moving, each unhappy particle moves with some probability $0 < p \leq 1$, then for suitable choices of p , we can disperse many more particles in logarithmic time.

To have some intuition as to why slowing particles down may speed up the process, consider that for small enough p , we can assume at any time step that at most one particle moves with high probability. In this range, we have a process in which particles that are happy stay still and at most one unhappy particle moves. This ensures that any vertex that is occupied will never become unoccupied. This is identical, other than the order that the particles move and that at some time steps, nothing changes, to the IDLA process, which we know completes in polynomial time, even for $M = n$.

As before, we let M be the total number of particles. A particle is happy (at a given step) if it is the only particle at its current vertex, otherwise it is unhappy. Let $H(t)$ be the number of happy particles at step t , and $U(t)$ the unhappy ones. Thus $H(t) + U(t) = M$. The process ends when $H(t) = M$. At each time step t any unhappy particle moves to a random vertex $v \in [n]$ with probability p and stays still with probability $(1 - p)$ independently of any other particle. The happy particles do not move.

Proof of Theorem 2(i)

Recall that we have $M = (1 - \delta)n$ particles moving on K_n (with loops) where $\delta = \frac{p}{2} + \alpha$.

The proof of Theorem 2 part (i) uses different method from the proof of Theorem 1. The distribution of unhappy particles at a given vertex affects the probability of a single unhappy particle becoming happy by remaining at this vertex, while the other particles at the location all leave.

For a vertex v , let O_v be the number of particles at v . Let $\mathcal{V}$ be the set of vertices with unhappy particles, and E be the number of vertices such that $O_v = 0$ (the empty vertices). Let $R = H + |\mathcal{V}|$ be the number of occupied vertices, i.e. the number of vertices with non-zero occupancy. At a given time step, let R_+ be the number of unoccupied vertices receiving at least one particle, and let R_- be the number of (necessarily unhappy) vertices losing all of their current particles.

Using $(1 - x)^k \leq 1 - kx + k^2x^2/2$, if k is a positive integer and $kx \leq 1$, we have,

$$\mathbf{E}R_+ = E \left(1 - \left(1 - \frac{p}{n} \right)^U \right) \geq \frac{EU p}{n} \left(1 - \frac{Up}{2n} \right). \quad (12)$$

$$\begin{aligned} \mathbf{E}R_- &= \sum_{v \in \mathcal{V}} p^{O_v} \left(1 - \frac{1}{n} \right)^{O_v} \left(1 - \frac{p}{n} \right)^{U - O_v} = \left(1 - \frac{p}{n} \right)^U \sum_{v \in \mathcal{V}} \left(\frac{p(n-1)}{n-p} \right)^{O_v} \\ &\leq \left(1 - \frac{Up}{n} + \frac{U^2 p^2}{2n^2} \right) \sum_{v \in \mathcal{V}} p^{O_v} \leq \left(1 - \frac{Up}{n} + \frac{U^2 p^2}{2n^2} \right) \frac{U}{2} p^2 \\ &= \frac{Up^2}{2} \left(1 - \frac{Up}{n} \left(1 - \frac{Up}{2n} \right) \right) \leq \frac{Up^2}{2} \left(1 - \frac{Up}{2n} \right). \end{aligned} \quad (13)$$

The expected change in R satisfies,

$$\begin{aligned}\mathbf{E}\Delta R &= \mathbf{E}[R_+ - R_-] \\ &\geq \frac{EU p}{n} \left(1 - \frac{Up}{2n}\right) - \frac{Up^2}{2} \left(1 - \frac{Up}{2n}\right) \\ &= Up \left(1 - \frac{Up}{2n}\right) \left(\frac{E}{n} - \frac{p}{2}\right).\end{aligned}$$

We consider the case $\delta = p/2 + \alpha$, so $E/n \geq \delta = p/2 + \alpha$. We also have $(Up)/(2n) \leq 1/2$, so

$$\mathbf{E}\Delta R \geq Up\alpha/2. \quad (14)$$

We define $D(t) \equiv M - R(t)$. The change of D in each time step is equal to the negative of the change in R and we have $D = 0$ if and only if $R = M$ and dispersion has occurred. By (14) and using $D(t) \leq U(t)$, we have that

$$\mathbf{E}(D(t+1)|D(t)) \leq D(t) - U(t)p\alpha/2 \leq D(t)(1 - p\alpha/2).$$

Iterating our argument we have,

$$\mathbf{E}D(t) \leq D(0)(1 - p\alpha/2)^t \leq D(0)\exp(-tp\alpha/2).$$

Clearly $D(0) < n$ and so choosing

$$t = \frac{4\log n}{p\alpha} = O((p\alpha)^{-1}\log n), \quad (15)$$

gives $\mathbf{E}D(t) = O(1/n)$. Hence $\mathbf{T}_{\text{Disp}} \leq t$, with probability $1 - O(1/n)$ as required.

Proof of Theorem 2(ii)

Recall that now we have $M = (1 - \delta)n$ particles moving on K_n (with loops) where $\delta = \frac{p}{2} - \alpha$.

The proof of Theorem 2 part (ii) is similar to that of Theorem 1. We measure the expected change in happy and unhappy particles. However the calculations are more involved.

Suppose $U(t) > 0$. At the next step $t+1$ there are X previously happy particles which become unhappy because unhappy particles land on top of them. Also Y previously unhappy particles become happy by being the only unhappy particle which (i) moves to an already unoccupied vertex (and no one else moves there), or (ii) moves to a vertex which contained several unhappy particles which all left, or (iii) does not move from its current location (or moves to the same location following the loop edge) while all other particles at that vertex left. This gives $H(t+1) = H(t) - X + Y$.

For each unhappy particle P , let O_P be the number of other particles at the same vertex. For a vertex v , let O_v be the total number of particles at v . Let $\mathcal{U}$ be the set of unhappy particles, $\mathcal{V}$ be the set of vertices with unhappy particles, and E be the number of vertices such that $O_v = 0$. We therefore have $n = H + E + |\mathcal{V}|$. The discussion above implies the following formulas for the expected values of X and Y . The three main sums on the right-hand side of (17) correspond to

the cases (i)-(iii) above. Inequality (17) would become equality, if we excluded from the inner sum the vertex occupied by the particle P .

$$\mathbf{E}X = H \left(1 - \left(1 - \frac{p}{n}\right)^U\right) = H - \frac{H(n-p)}{n} \left(1 - \frac{p}{n}\right)^{U-1}. \quad (16)$$

$$\begin{aligned} \mathbf{E}Y &\leq \sum_{P \in \mathcal{U}} \left(E \frac{p}{n} \left(1 - \frac{p}{n}\right)^{U-1} \right) + \sum_{P \in \mathcal{U}} \left(\sum_{v \in \mathcal{V}} \left(\frac{p}{n} p^{O_v} \left(1 - \frac{1}{n}\right)^{O_v} \left(1 - \frac{p}{n}\right)^{U-O_v-1} \right) \right) \\ &\quad + \sum_{P \in \mathcal{U}} \left(\left(1 - p + \frac{p}{n}\right) p^{O_P} \left(1 - \frac{1}{n}\right)^{O_P} \left(1 - \frac{p}{n}\right)^{U-O_P-1} \right) \\ &= \frac{UEp}{n} \left(1 - \frac{p}{n}\right)^{U-1} + \frac{Up}{n} \sum_{v \in \mathcal{V}} \left(p^{O_v} \left(1 - \frac{1}{n}\right)^{O_v} \left(1 - \frac{p}{n}\right)^{U-O_v-1} \right) \\ &\quad + \sum_{v \in \mathcal{V}} O_v \left(\left(1 - p + \frac{p}{n}\right) p^{O_v-1} \left(1 - \frac{1}{n}\right)^{O_v-1} \left(1 - \frac{p}{n}\right)^{U-O_v} \right) \\ &= \left(1 - \frac{p}{n}\right)^{U-1} \times \\ &\quad \left(\frac{UEp}{n} + \sum_{v \in \mathcal{V}} \left(p \left(\frac{n}{n-p} \right) \left(1 - \frac{1}{n}\right) \right)^{O_v} \left(\frac{Up}{n} + O_v \frac{(1-p+p/n)}{p} \frac{n}{n-1} \left(1 - \frac{p}{n}\right) \right) \right) \\ &= \left(1 - \frac{p}{n}\right)^{U-1} \left(\frac{UEp}{n} + \sum_{v \in \mathcal{V}} \left(\frac{p(n-1)}{n-p} \right)^{O_v} \left(\frac{Up}{n} + O_v \frac{(1-p)(n-p)}{p} \frac{1}{n-1} \right) \right) + O(1) \\ &\leq \left(1 - \frac{p}{n}\right)^{U-1} \left(\frac{UEp}{n} + \left(\frac{p(n-1)}{n-p} \right)^2 \left(\frac{U^2 p}{2n} + U \frac{(1-p)(n-p)}{p} \frac{1}{n-1} \right) \right) + O(1). \end{aligned} \quad (17)$$

The last inequality above holds because $p(n-1)/(n-p) \leq 1$, $O_v \geq 2$, $\mathcal{V} \leq U/2$ and $\sum_{v \in \mathcal{V}} O_v = U$. As before, the concentration of $H' := H(t+1)$ will follow from a martingale argument as each of the U particles chooses whether and where to move independently. A difference in whether or not to move and choice of vertex by one particle can only alter the final value of H' by at most 2, so Inequality (6) applies.

Using (16) and (18), we have the following bound on the expected value of H' , given $H := H(t)$.

$$\begin{aligned} \mathbf{E}[H' | H] &= H - \mathbf{E}X + \mathbf{E}Y \\ &\leq \left(1 - \frac{p}{n}\right)^{U-1} \left(\frac{EUp + H(n-p)}{n} + \left(\frac{p(n-1)}{n-p} \right)^2 \left(\frac{U^2 p}{2n} + U \frac{(1-p)(n-p)}{p} \frac{1}{n-1} \right) \right) + O(1) \\ &\leq \left(1 - \frac{p}{n}\right)^{U-1} \left(\frac{EUp + H(n-p)}{n} + \frac{U^2 p^3}{2n} + Up(1-p) \right) + O(1). \end{aligned} \quad (19)$$

We write the current value of H as

$$H = M - \varepsilon n = (1 - \delta - \varepsilon) n, \quad (20)$$

where $0 < \varepsilon \leq 1 - \delta$. Thus

$$U = \varepsilon n \quad \text{and} \quad E \leq (\delta + \varepsilon) n. \quad (21)$$

We use (20) and (21) in (19).

$$\begin{aligned}
\mathbf{E}[H' | H] &\leq n \left(1 - \frac{(U-1)p}{n} + \frac{U^2 p^2}{2n^2} \right) \left(\frac{EUp + H(n-p)}{n^2} + \frac{U^2 p^3}{2n^2} + \frac{Up(1-p)}{n} \right) + O(1) \\
&\leq n \left(1 - \varepsilon p + \frac{p}{n} + \frac{\varepsilon^2 p^2}{2} \right) \left((\delta + \varepsilon) \varepsilon p + (1 - \delta - \varepsilon) (1 - p/n) + \frac{\varepsilon^2 p^3}{2} + \varepsilon p(1-p) \right) + O(1) \\
&= n \left(1 + \frac{p}{n} - \varepsilon p \left(1 - \frac{\varepsilon p}{2} \right) \right) \left((1 + \delta + \varepsilon) \varepsilon p + 1 - \delta - \varepsilon + O\left(\frac{p}{n}\right) + \frac{\varepsilon^2 p^3}{2} - \varepsilon p^2 \right) + O(1) \\
&\leq n \left(1 - \varepsilon p \left(1 - \frac{\varepsilon p}{2} \right) \right) \left(1 - \delta - \varepsilon + \varepsilon p \left(1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p \right) \right) + O(1) \\
&= n(1 - \delta - \varepsilon \\
&\quad + \varepsilon p \left(1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p - \left(1 - \frac{\varepsilon p}{2} \right) \left(1 - \delta - \varepsilon + \varepsilon p \left(1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p \right) \right) \right) + O(1) \\
&= n(1 - \delta - \varepsilon + \varepsilon p A) + O(1)
\end{aligned} \tag{22}$$

We now bound A :

$$\begin{aligned}
A &= 1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p - \left(1 - \frac{\varepsilon p}{2} \right) \left(1 - \delta - \varepsilon + \varepsilon p \left(1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p \right) \right) \\
&= 1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p - 1 + \delta + \varepsilon - \varepsilon p \left(1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p \right) \\
&\quad + \frac{\varepsilon p}{2} \left(1 - \delta - \varepsilon + \varepsilon p \left(1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p \right) \right) \\
&= 2\delta + 2\varepsilon - p + \varepsilon p \left(\frac{p}{2} + \frac{1}{2} - \frac{\delta}{2} - \frac{\varepsilon}{2} + \frac{\varepsilon p}{2} \left(1 + \delta + \varepsilon + \frac{\varepsilon p^2}{2} - p \right) - 1 - \delta - \varepsilon - \frac{\varepsilon p^2}{2} + p \right) \\
&\leq 2\varepsilon - 2\alpha + \varepsilon p \left(\frac{p}{2} - \frac{1}{2} - \frac{3}{2}\delta - \frac{3}{2}\varepsilon + \varepsilon p - \frac{\varepsilon p^2}{2} + p \right)
\end{aligned} \tag{23}$$

$$\begin{aligned}
&= 2\varepsilon - 2\alpha + \varepsilon p \left(p - \left(\frac{1}{2} - \frac{p}{2} \right) - \frac{3}{2}\delta - \varepsilon \left(\frac{3}{2} - p \right) - \frac{\varepsilon p^2}{2} \right) \\
&\leq 2\varepsilon - 2\alpha + \varepsilon p.
\end{aligned} \tag{24}$$

For Inequality (23), use $\delta = (p/2) - \alpha$ and observe that $\varepsilon \leq 1 - \delta$ implies $1 + \delta + \varepsilon + \varepsilon p^2/2 - p \leq 2 - p(1 - (1 - \delta)p/2) \leq 2$. Combining (22) and (24), we get

$$\begin{aligned}
\mathbf{E}[H' | H] &\leq n(1 - \delta - \varepsilon + \varepsilon p A) + O(1) \\
&\leq n(1 - \delta - \varepsilon + 2\varepsilon p(\varepsilon(1 + p) - \alpha)) + O(1).
\end{aligned} \tag{25}$$

For $\varepsilon \leq \varepsilon_0 := \alpha/(3(1 + p))$, we have $2(\varepsilon(1 + p) - \alpha) \leq -4\alpha/3 \leq -\alpha$ and so,

$$\mathbf{E}[H' | H] \leq n(1 - \delta - \varepsilon - \varepsilon p \alpha) + O(1).$$

Now from (6), assuming additionally that $\varepsilon = \omega(1/(p\alpha n))$, we have

$$\Pr(H' \leq n(1 - \delta - \varepsilon)) \geq 1 - e^{-\Omega(n\varepsilon p^2 \alpha^2)}. \tag{26}$$

In the range $\varepsilon_0 \leq \varepsilon \leq 1 - \delta$, the bound in (25) does not seem strong enough to separate H' from $n(1 - \delta)$, but the following bound on $U' := U(t + 1)$, which holds for any $0 \leq \varepsilon \leq 1 - \delta$, will help.

$$\mathbf{E}[U' | U] \geq U \left(1 - \left(1 - \frac{p}{n} \right)^{U-1} \right) \quad (27)$$

$$\begin{aligned} &\geq U \left(1 - \left(1 - \frac{(U-1)p}{n} + \frac{U^2 p^2}{2n^2} \right) \right) = \frac{U^2 p}{n} \left(1 - \left(\frac{1}{U} + \frac{Up}{2n} \right) \right) \\ &\geq \frac{U^2 p}{n} \left(\frac{1}{2} - \frac{p}{n} \right) \geq \frac{U^2 p}{3n} = \frac{n\varepsilon^2 p}{3}. \end{aligned} \quad (28)$$

Inequality (27) holds because the probability that a given unhappy particle P remains unhappy is at least the probability that at least one other unhappy particle decides to move to the vertex chosen by P . The first inequality on line (28) holds because the maximum value of $1/U + Up/(2n)$ for $U \in [2, n]$ is $1/2 + p/n$ (achieved at $U = 2$).

The bound (28) implies

$$\mathbf{E}[H' | H] \leq n \left(1 - \delta - \frac{\varepsilon^2 p}{3} \right), \quad (29)$$

so from (6),

$$\Pr \left(H' \leq n \left(1 - \delta - \frac{\varepsilon^2 p}{4} \right) \right) \geq 1 - e^{-\Omega(n\varepsilon^3 p^2)}. \quad (30)$$

We use bounds (26) and (30) to conclude the proof. Assume that $\varepsilon \geq \varepsilon_0^2 p/4$, that is, assume that $H \leq n(1 - \delta - \varepsilon_0^2 p/4)$. Then for some constant c ,

$$\Pr \left(H' \leq n \left(1 - \delta - \frac{\varepsilon_0^2 p}{4} \right) \right) \geq 1 - e^{-2cnp^3 \alpha^4}.$$

The above inequality follows from (26) for the case when $\varepsilon_0^2 p/4 \leq \varepsilon \leq \varepsilon_0$, and from (30) for the case $\varepsilon_0 \leq \varepsilon \leq 1 - \delta$. Thus the probability that H reaches $n(1 - \delta) = M$, and dispersion completes, within $e^{cnp^3 \alpha^4}$ steps is at most $e^{-cnp^3 \alpha^4}$.

3 Dispersion on k -regular trees

For convenience we restate Theorem 5.

Theorem. Let $k \geq 3$, and $\varepsilon > 0$, constant. For a k -regular tree of depth at least $(2 + \varepsilon) \log_{k-1} M$, with M particles initially placed at the central vertex of the tree, w.h.p. the dispersion time $\mathbf{T}_{Disp} = O(M \log_{k-1} M)$. The maximum distance $\mathbf{D}_{Disp}$ from the origin is w.h.p. bounded by

$$\left(1 + \frac{1}{2(\log_{k-1} k) - 1} - \varepsilon \right) \log_{k-1} M \leq \mathbf{D}_{Disp} \leq \left(\frac{5}{3} + \frac{1}{3 \log_{k-1} k} + 2\varepsilon \right) \log_{k-1} M. \quad (31)$$

We require the following lemma.

Lemma 6. For $\ell \geq 2$, $\varepsilon > 0$ and M independent random walks starting from the same origin vertex of the infinite k -regular tree, the probability that there is a vertex at depth $d \geq \frac{\ell + \varepsilon}{\ell - 1} \log_{k-1} M$ visited by ℓ or more walks is $O(M^{-\varepsilon})$.

Proof. What is the probability that a given random walk visits a given vertex v at depth d from the origin vertex? Let X_t be the distance of this walk from v at time t , so $X_0 = d$. With probability $p = \frac{k-1}{k}$, the particle will move away from v and $X_{t+1} = X_t + 1$. Conversely, with probability $q = 1 - p = \frac{1}{k}$, the walk will move towards v and $X_{t+1} = X_t - 1$. The properties of random walks with bias are given in Feller [5]. By equation (3.6) of Chapter XIV, the probability of reaching v (ultimate ruin) starting from distance d is

$$\Pr(X_t = 0, \text{ for some } t) = \left(\frac{q}{p}\right)^d = \left(\frac{1}{k-1}\right)^d. \quad (32)$$

By the union bound over the groups of ℓ walks and over the $k(k-1)^{d-1}$ vertices at depth d , the probability that there are ℓ or more walks visit the same vertex at depth d is at most

$$\binom{M}{\ell} k(k-1)^{d-1} \left(\frac{1}{k-1}\right)^{d\ell} \leq M^\ell \frac{k}{k-1} \left(\frac{1}{k-1}\right)^{(\ell-1)d} < 2M^{-\varepsilon}$$

□

An immediate consequence of Lemma 6 is that in the dispersion process on the infinite k -regular tree, with probability at least $1 - O(M^{-\varepsilon})$ no particle reaches depth $d = (2 + \varepsilon) \log_{k-1} M$ from the origin, because with probability at least $1 - O(M^{-\varepsilon})$ no two random walks of the particles visit the same vertex at depth $d - 1$. This means that while Theorem 5 refers to a k -regular tree of a finite depth at least $(2 + \varepsilon) \log_{k-1} M$, we can assume that we have the infinite tree. Another consequence of Lemma 6 is that during the dispersion process, with probability at least $1 - O(M^{-\varepsilon})$, no vertex at depth $d \geq \frac{\ell + \varepsilon}{\ell - 1} \log_{k-1} M$ is ever visited by ℓ or more particles.

Proof of lower bound in (31). The number of vertices at depth d or less from the origin is equal to $(k(k-1)^d - 2)/(k-2)$. For $d = (1 - \varepsilon/2) \log_{k-1} M$, this is equal to $m_d = (kM^{1-\varepsilon/2} - 2)/(k-2) = o(M)$, so at dispersion, $M - m_d = (1 - o(1))M$ particles must occupy a vertex at distance greater than d . For this to happen, the particle must be at distance d together with at least one other particle at some process step. Let t_0 be the earliest process step at which this occurs for any particle during dispersion. At step t_0 , pair up the particles on each vertex with multiple occupancy at distance d , ignoring the leftover particle if the vertex occupancy is odd. Continue such pairing (of particles at distance d) in all subsequent process steps. It is entirely possible that one or both particles in the current pair occurred as part of a previous pair.

Let $\mathcal{B}_i$ be the event that the i -th pair exists. With probability 1, $(1 - o(1))M$ particles leave at some step a vertex at distance d , so $\Pr(\mathcal{B}_i) = 1$, for each $1 \leq i \leq M/3$. Let B be the event that a pair of particles which start from the same non-root vertex of the tree move together (without separating) directly forward away from the origin for $d' = A \log_{k-1} M$ steps. In our dispersion process, for each $1 \leq i \leq M/3$, let B'_i be the event that the i -th pair exists and it moves together directly forward away from the origin for d' steps (reaching the distance $d + d'$ from the origin). We can view the events B'_i as events $\mathcal{B}_i \cap B_i$, where events B_i are independent copies of B . As mentioned above, the same particle may participate in more than one pair, but this is irrelevant for the independence of B_i 's. Each return to distance d is a renewal event for the random walk of the particle.

The probability of the event B is given by

$$\Pr(B) = \left(\frac{k-1}{k} \times \frac{1}{k}\right)^{d'} = \left(\frac{k-1}{k^2}\right)^{(A \log_{k-1} M)}$$

$$\begin{aligned}
&= (k-1)^{(1-2\log_{k-1} k)A \log_{k-1} M} \\
&= M^{-((2\log_{k-1} k)-1)A} \\
&= M^{-1+\varepsilon/2},
\end{aligned} \tag{33}$$

where the last equality is by taking $A = \frac{1-\varepsilon/2}{2(\log_{k-1} k)-1}$. Thus the expected number of events B_i is equal to $M^{\varepsilon/2}/3$, so by a Chernoff bound,

$$\Pr(\text{fewer than } M^{\varepsilon/2}/4 \text{ events } B_i \text{ occur}) \leq e^{-\Theta(M^{\varepsilon/2})}. \tag{34}$$

The probability that fewer than $M^{\varepsilon/2}/4$ events B'_i occur, that is, the probability that fewer than $M^{\varepsilon/2}/4$ of the pairs go straight to distance $d+d'$, is equal to the probability that fewer than $M^{\varepsilon/2}/4$ events $\mathcal{B}_i \cap B_i$ occur, which is at most

$$\Pr(\neg \mathcal{B}_{M/3} \text{ or fewer than } M^{\varepsilon/2}/4 \text{ events } B_i \text{ occur}) \leq e^{-\Theta(M^{\varepsilon/2})},$$

where the inequality follows from $\Pr(\neg \mathcal{B}_{M/3}) = 0$ and (34).

Let $\mathcal{A}$ be the event that at the completion of dispersion no particle is at distance $d+d'$ or greater from the origin. Either there have been fewer than $M^{\varepsilon/2}/4$ particle pairs arriving at distance $d+d'$ from the origin, or there are $M^{\varepsilon/2}/2$ independent random walks starting from distance $d+d'$ which return to distance $d+d'-1$. By (32), the probability a random walk ever returns to distance $d+d'-1$ from $d+d'$ is $1/(k-1)$. Even if some of these walks are made by the same particle, each time the walk arrives at $d+d'$ from d , a subsequent return to $d+d'-1$ is independent of the previous trajectory of this walk.

The probability of the event $\mathcal{A}$ is at most $e^{-\Theta(M^{\varepsilon/2})} + (k-1)^{-M^{\varepsilon/2}/2} = e^{-\Theta(M^{\varepsilon/2})}$. Thus with probability at least $1 - e^{-\Theta(M^{\varepsilon/2})}$, when the dispersion is completed, at least one particle is away from the origin at distance at least

$$d'' = d + d' = (1 - \varepsilon/2 + \frac{1 - \varepsilon/2}{2(\log_{k-1} k) - 1}) \log_{k-1} M \geq (1 - \varepsilon + \frac{1}{2(\log_{k-1} k) - 1}) \log_{k-1} M. \tag{35}$$

Proof of upper bound in (31). We define

$$\begin{aligned}
d' &= \frac{3+\varepsilon}{2} \log_{k-1} M, \\
d &= \left(\frac{5}{3} + \varepsilon\right) \log_{k-1} M > d', \\
\ell &= \left(\frac{1+\varepsilon}{3 \log_{k-1} k}\right) \log_{k-1} M = \left(\frac{1+\varepsilon}{3}\right) \log_k M.
\end{aligned}$$

We first outline the proof. From Lemma 6, with probability at least $1 - O(M^{-\varepsilon})$, no three particles go through the same vertex at depth d' . Thus beyond depth d' , particles move in fixed pairs (that is, particles do not change partners while moving beyond depth d'). We show that w.h.p. only $O(M^{1/3-\varepsilon})$ particles ever reach depth d . With probability at least $1 - O(M^{-\varepsilon})$, the particles which reach depth d and then move on, do so in disjoint pairs and settle in their final location at the first step when they separate from their partner. Thus with probability at least $1 - O(M^{-\varepsilon})$, the movement of particles from depth d consists of $O(M^{1/3-\varepsilon})$ disjoint pairs of particles moving

together until they separate. The probability that a given pair of particles which reside at the current step at the same vertex will move together for ℓ subsequent steps is at most $O(M^{-1/3-\varepsilon})$. Thus the probability that any of the $O(M^{1/3-\varepsilon})$ pairs starting from depth d move together for ℓ steps is at most $O(M^{-2\varepsilon})$. Putting all these together, we conclude that with probability at least $1 - O(M^{-\varepsilon})$, no particle reaches depth $d + \ell$, which is less than the upper bound in (31). We proceed with formal arguments.

With probability 1, the pre-determined walks of all particles eventually reach depth d , and each of these walks has a corresponding vertex that it first visits at this depth. In the dispersion process, for a particle P_1 to reach for the first time depth d , say at vertex v , and continue moving, at least one other particle P_2 must also have this vertex in its path. If vertex v is not the first at depth d in the path of particle P_2 , then to reach this point P_2 must have passed earlier through some vertex w at d . This required another particle $P_3 \neq P_1$ to have been at w . This can happen only if all three (distinct) particles P_1 , P_2 and P_3 have passed through the same vertex at depth d' (the probability of this is $O(M^{-\varepsilon})$) or particles P_2 and P_3 moved together without separating from vertex w until they returned to depth d' . If we can demonstrate that the second event (two particles returning from depth d to depth d') happens also only with probability $O(M^{-\varepsilon})$, then we only need consider pairs of particles that share a vertex at their first visit to depth d .

The label of the vertex first visited at depth d by the pre-determined random walk of a particle is distributed independently and uniformly at random for all particles. Let X be the number of vertices at depth d where the walks of at least two particles first reach this depth. Then X corresponds to the number of multiply occupied bins, in a balls-in-bins problem with M balls (particles) going into $k(k-1)^d = \Theta\left(M^{\frac{5}{3}+\varepsilon}\right)$ bins (vertices at depth d).

Examining the particles in a fixed order, there are at most $M-1$ vertices already claimed by previous particles, and so the probability of a collision is $O\left(M/M^{\frac{5}{3}+\varepsilon}\right) = O\left(M^{-\left(\frac{2}{3}+\varepsilon\right)}\right)$. This bound holds independently of the outcome of previous trials. Therefore we can bound the number of multiply occupied bins X by the sum of independent random indicator variables, each with expectation $O\left(M^{-\left(\frac{2}{3}+\varepsilon\right)}\right)$. Thus

$$\mathbf{E}(X) = M \times O\left(M^{-\left(\frac{2}{3}+\varepsilon\right)}\right) = O\left(M^{\frac{1}{3}-\varepsilon}\right),$$

and using Hoeffding's inequality, w.h.p. $X = O\left(M^{\frac{1}{3}-\varepsilon}\right)$.

Even assuming that each of these collisions occurs during dispersion, for a given pair of particles to return from depth d to depth d' , they would have to move together a distance of at least $d - d' \geq \frac{1}{6} \log_{k-1} M$ towards the root vertex without separating. The movements of the particles are independent random walks and so by (32), the probability that a particle visits this ancestor vertex at any point in its subsequent infinite walk is at most

$$(k-1)^{-\frac{1}{6} \log_{k-1} M} = M^{-\frac{1}{6}}.$$

Therefore, the probability that both particles visit this ancestor is at most $M^{-\frac{1}{3}}$, and by taking a union bound over the $O(M^{\frac{1}{3}-\varepsilon})$ pairs, we see that with probability at least $1 - O(M^{-\varepsilon})$ no pair can return to their ancestor at depth d' , and no other particle visits this ancestor.

Next we show that with probability at least $1 - O(M^{-\varepsilon})$ no pair of particles at depth d will move together for more than ℓ further steps before separating. The probability that two particles move a distance ℓ without separating is $k^{-\ell} = M^{-(1+\varepsilon)/3}$. Taking a union bound, the probability that any of the pairs move together for at least ℓ steps is at most $O(M^{1/3-\varepsilon}) \cdot M^{-(1+\varepsilon)/3} = O(M^{-\varepsilon})$.

Thus with probability at least $1 - O(M^{-\varepsilon})$, any pair of particles that leave a vertex at depth d move together fewer than ℓ further steps. Therefore, at the end of the dispersion process, with probability at least $1 - O(M^{-\varepsilon})$ no particle will have reached further than to depth

$$d + \ell = \left(\frac{5}{3} + \varepsilon\right) \log_{k-1} M + \left(\frac{1 + \varepsilon}{3 \log_{k-1} k}\right) \log_{k-1} M < \left(\frac{5}{3} + \frac{1}{3 \log_{k-1} k} + 2\varepsilon\right) \log_{k-1} M.$$

Upper bound on $\mathbf{T}_{\text{Disp}}$. The upper bound $O(M \log_{k-1} M)$ on the dispersion time $\mathbf{T}_{\text{Disp}}$ follows from the upper bound $O(\log_{k-1} M)$ on the dispersion distance $\mathbf{D}_{\text{Disp}}$ given in (31). The proof of the upper bound in (31) applies not only to the maximum final distance $\mathbf{D}_{\text{Disp}}$ but also to the maximum distance $\tilde{\mathbf{D}}_{\text{Disp}}$ over all steps of the process. Observe that for some small ε , the right-hand side in (31) is less than $2 \log_{k-1} M$.

Let $\ell \leq \mathbf{T}_{\text{Disp}}$ denote a step of the dispersion process, and let $t(\ell)$ be the *total number of random walk steps* made by all particles during ℓ process steps. Let $\mathbf{S}(t(\ell))$ be the sum of the distances of all particles from the origin at ℓ .

Consider a particle making a biased random walk on the integer line starting from 0, which at each step moves right or left with probabilities $(k-1)/k$ and $1/k$, respectively. Let S_t be the position of this particle at step t . We have $\mathbf{E}(S_t) = ((k-2)/k)t$, and by the Azuma-Hoeffding inequality, $\Pr(S_t \leq t/4) \leq e^{-\Theta(t)}$.

We couple $\mathbf{S}(t)$ and S_t so that $\mathbf{S}(t) \geq S_t$. If a particle in the dispersion process moves from the origin, then $\mathbf{S}(t+1) = \mathbf{S}(t) + 1$ but $S_{t+1} \in \{S_t - 1, S_t + 1\}$. Otherwise, when a particle in the dispersion process moves from a vertex other than the origin, both $\mathbf{S}(t)$ and S_t change by the same value $+1$ or -1 .

As long as $\ell \leq \mathbf{T}_{\text{Disp}}$, at least 2 particles move at each step, so $t(\ell) \geq 2\ell$. Let $d(\ell)$ be the maximum displacement from the origin at ℓ , where $d(\ell) \leq \tilde{\mathbf{D}}_{\text{Disp}}$. Also $d(\ell) \geq \mathbf{S}(t(\ell))/M \geq S_{t(\ell)}/M$, and w.h.p. $S_t \geq t/4$. Putting this all together,

$$\tilde{\mathbf{D}}_{\text{Disp}} \geq d(\ell) \geq t(\ell)/4M \geq \ell/2M.$$

Assuming the w.h.p. event that $\tilde{\mathbf{D}}_{\text{Disp}} \leq 2 \log_{k-1} M$, it follows that $\mathbf{T}_{\text{Disp}} \leq 4M \log_{k-1} M$.

4 Dispersion on paths: Proof of Theorem 4

To analyse this case we consider that G is the integer path, with $V(G) = \mathbb{Z}$ and $(i, j) \in E(G)$ if and only if $|i - j| = 1$. Initially all M particles are placed at the origin 0. We will use the following standard result about random walks on the line.

Lemma 7. *For a simple random walk on the integer path, let $R(2T, r)$ be the probability of at least r returns to the origin in $2T$ steps. For $\alpha > 0$, we have*

$$R(2T, \alpha\sqrt{2T}) = O\left(\frac{1}{\alpha} e^{-\alpha^2/2}\right). \quad (36)$$

For completeness, a proof of Lemma 7 is given in the Appendix.

The movement of particle i at time step t takes a value $X_i(t) \in \{-1, 0, 1\}$, with $X_i(t) = 0$ if the particle doesn't move. For any particle the next non-zero movement is uniformly distributed in $\{-1, 1\}$ and independent of the choice of particle, or of the action of any other particles.

If we consider only the walk steps when a given particle moves, and ignore the process time steps at which the particle does not move, the particle makes a random walk on the line. The particle moves only when its random walk intersects that of another particle. When two particles meet at a vertex, reversing the walk of the second particle and taking the union of these two walks, gives a walk which has returned to the origin.

We consider the walk steps of particles 1 and 2, and build a sequence $Y(t) = Y(t, \{1, 2\})$ as follows. Let $W \in \{-1, 1\}$ denote the movement of a particle at a given walk step of that particle. Note that if W is uniformly and independently distributed on $\{-1, 1\}$, then so is $-W$. The entries of $Y(t)$ are the movements $(W_1), -(W_2)$, if any, made by the two particles up to the end of time step t in time step order; the movements of the second particle reversed. If both particles move at a given time step, the order is the movement of particle 1 followed by that of particle 2. The entries for particle 2 are the negative of the step direction W_2 . Thus $Y(0) = (W_1(0), -W_2(0))$ as both move from the origin. Thus, $Y(t) = Y(t-1)$ if neither move, $Y(t) = (Y(t-1), W_1(t))$ if particle 1 moves but not particle 2, $Y(t) = (Y(t-1), -W_2(t))$ if particle 2 moves but not particle 1, $Y(t) = (Y(t-1), W_1(t), -W_2(t))$ if both move.

Let $s_i(t)$ be the number of walk steps taken by particle i at the end of time step t . The length of $Y(t)$ is $s = s_1(t) + s_2(t)$. As $t \rightarrow \infty$, either $Y(t)$ is infinite (dispersion never stops) or has a finite length T . If so, let $Y = (Y_1, Y_2, \dots, Y_T)$ be the entries of $Y(t)$. If T is finite, extend Y for $i > T$ by setting $Y_i \in \{-1, +1\}$, chosen independently with probability $1/2$, and let $Y = (Y_i : i \geq 1)$.

Note that without knowledge of the value of T , each Y_i is distributed uniformly and independently at random and so Y lists the transitions of an infinite random walk on the line. The positions of this walk are $(0, Y^{(1)}, \dots, Y^{(s)}, \dots)$, where $Y^{(s)}$ is the partial sum $Y^{(s)} = \sum_{i=1}^s Y_i$. Importantly, the partial sum $Y^{(s)}$ is zero whenever the two particles meet, up to the time step when they moved s walk steps in total. Note that the partial sums of Y can be zero more often than the particles meet at a vertex. When the particles move simultaneously, two entries Y_i, Y_{i+1} are made in $Y(t)$, and the partial sum up to i can be zero.

Theorem 8. *For all $\varepsilon > 0$, w.h.p. the dispersion process on the integers with M particles, will terminate in $4(1+\varepsilon)M^3 \log M$ time steps, with no particle at a distance greater than $4(1+\varepsilon)M \log M$ from their original position.*

Proof. At each time step $t \geq 0$ let $Z_i(t)$ be the walk length, i.e. the total number of walk steps made by particle i up to the end of time step t . Let K be a large constant, and let F_{ij} be the event that

$$F_{ij} = \{\text{Particles } i, j \text{ meet more than } \alpha(2S)^{1/2} \text{ times while their walk lengths } Z_i, Z_j \leq S + K\},$$

and let $F = \bigcup F_{ij}$, where the union is over all unordered pairs $\{i, j\}$, $i \neq j$. Using the sequence $Y(t, \{i, j\})$ for the particles i, j (as described above for particles 1,2) we can upper bound the number of meetings between the particles, by the number of returns to the origin of the random walk $Y(t)$.

Let $\alpha^2 = 4(1+\varepsilon) \log M$, and $S = 2\alpha^2 M^2$. Let $T = S + K$, $\beta = \alpha\sqrt{S/(S+K)}$. By inequality (39) of Lemma 7, the probability of at least $\beta\sqrt{2(S+K)} = \alpha\sqrt{2S}$ returns in $2(S+K)$ steps, satisfies,

$$R\left(2(S+K), \beta\sqrt{2(S+K)}\right) = O\left(M^{-(2+\varepsilon)}\right).$$

As returns are monotone non-decreasing with the number of walk steps, if $Z_i(t) + Z_j(t) < 2(S+K)$, extend $Y(t)$ to a walk of length $2(S+K)$ and include any extra returns. Thus

$$\Pr(F) \leq M^2 O(M^{-(2+\varepsilon)}) = O(M^{-\varepsilon}).$$

Suppose that at time step $t' > SM/2$ the dispersion process has not stopped. Then there exists a particle which has taken more than S walk steps. Pick the first process time step $t \leq t'$ at which any particle i makes $S + 1$ walk steps, choosing the particle with the lowest label I if there is any choice.

Recall that $Z_i(t)$ is the number of walk steps made by particle i at time step t . For a given particle i , let $R_{ij}(t)$ be the number of walk steps at which particles i and j occupy the same vertex. Each such pair-wise meeting causes both particles to make one step of a random walk, and so $\sum_{j \neq i} R_{ij}(t)$ counts every walk step of i at least once. Thus $Z_i(t) \leq \sum_{j \neq i} R_{ij}(t)$. Assume the event F^c holds. At time step t , no particle has made more than $S + 1$ walk steps, so for all i ,

$$Z_i(t) \leq \sum_{j \neq i} R_{ij}(t) \leq M\alpha(2S)^{1/2}.$$

In particular

$$Z_I(t) = S + 1 \leq M\alpha(2S)^{1/2},$$

and thus

$$(S + 1)^2 \leq 2\alpha^2 M^2 S = S^2.$$

This implies that $(S + 1) \leq S$, contradicting the existence of a first t where some particle exceeds S steps. Thus with probability at least $1 - O(M^{-\varepsilon})$ (that is, at least with the probability of the event F^c), no particle takes more than S steps and we have $\mathbf{T}_{\text{Disp}} \leq MS/2 = 4(1 + \varepsilon)M^3 \log M$.

We next prove that $\mathbf{D}_{\text{Disp}} = O(M \log M)$. By using a Chernoff bound for the sum of s independent uniform $\{-1, 1\}$ random variables, we see that the probability a random walk reaches a distance greater than a in s walk steps is less than $2e^{-\frac{a^2}{2s}}$. If any particle reaches a distance greater than $4(1 + \varepsilon)M \log M$ from the origin, then the event F occurs or one of M random walks of lengths $S = 8(1 + \varepsilon)M^2 \log M$ reaches that distance. By the above, the probability of the latter occurring for one given random walk is less than

$$2e^{-\frac{(4(1+\varepsilon)M \log M)^2}{16(1+\varepsilon)M^2 \log M}} = 2e^{-(1+\varepsilon) \log M} = 2M^{-(1+\varepsilon)}.$$

Taking the union bound over all particles, with high probability no particle reaches a distance of $4(1 + \varepsilon)M \log M$. $\square$

To conclude the proof of Theorem 4, we consider a finite-length path or cycle. The bound $4(1 + \varepsilon)M \log M$ in Theorem 8 applies to the maximum distance any particle will be at from the origin at the end of the process. It is possible that a particle may reach a further distance and return before the process terminates. Using the same argument as in the end of the proof of Theorem 8 but taking a union bound over all particles *and* over the steps $1 \leq s \leq S = 8(1 + \varepsilon)M^2 \log M$ of each particle's walk, we can show that w.h.p. at no point in the process, could any particle reach a distance of $4\sqrt{2}(1 + \varepsilon)M \log M$. Therefore w.h.p. this process will disperse in the same manner as on the infinite line on any finite path or cycle of size larger than $8\sqrt{2}(1 + \varepsilon)M \log M$, which is less than $12M \log M$ for $\varepsilon = 0.2$.

5 Dispersion on grids and other infinite Abelian Cayley graphs

In this section we prove Theorem 3, part (i) and show its implication for the dispersion in the 2-dimensional grid.

We use the same argument as for the line, linking times that two particles meet in a grid or other Abelian Cayley graph with the number of returns to the origin of a single combined random walk.

In an Abelian Cayley graph, $X_i(t) \in S \cup \{\mathbf{0}\}$, where S is the symmetric set of generators which define the graph, i.e., if $a \in S$ then $-a \in S$, and $a + (-a) = \mathbf{0}$. Here $\mathbf{0}$ is the additive identity $a + \mathbf{0} = a$ for all elements a of the group. We assume the particles start at $\mathbf{0}$, which we call the origin. If $A_i \in S$ is the generator sampled at step i of the walk, the position W_t of the walk at step t is $W_0 = \mathbf{0}$, and for $t \geq 1$, $W_t = \sum_{i=1}^t A_i$.

If we consider only the steps where a given particle moves (walk steps), and ignore the time-steps in which the particle does not move, the particle makes a random walk on the graph. The particle moves only when its random walk intersects that of another particle. When two particles meet at a vertex, reversing the walk of the second particle and taking the union of these two walks, gives a walk which has returned to the origin. The following Lemma is a proof of Theorem 3, part (i), and is used in the proof of Theorem 3, part (ii).

Lemma 9. *Let $\omega = \omega(M) \rightarrow \infty$. Let G be an infinite Abelian Cayley graph, and let t be such that $t \geq \omega \binom{M}{2} R(2t)$, where $R(2t)$ is the expected number of returns to the origin in $2t$ steps by a simple random walk on G . Then with probability at least $1 - 1/\omega$, a system of M particles disperses on G in t process steps*

Proof. The first part of the proof is similar to that for the infinite line, see above Theorem 8 of Section 4, and we give only an outline.

Let $W \in S$ denote the movement of a particle at a given walk step. Note that if W is uniformly and independently distributed on S , then so is $-W$. The sequence $Y(t)$ records the transitions $W_1(t), -W_2(t)$ of particles 1 and 2 up to the end of time step t in time step order, the walk of the second particle being negated. Thus if $W_2(t) = a$, for some $a \in S$, the entry in $Y(y)$ is $-W_2(t) = -a$. As before, if both particles move at time step t , the order in $Y(t)$ is the movement of particle 1 followed by that of particle 2.

Let $s_i(t)$ be the number of walk steps taken by particle i at the end of time step t . The length of $Y(t)$ is $s = s_1(t) + s_2(t)$. As $t \rightarrow \infty$, either Y is infinite or has finite length T . If T is finite, extend Y for $i > T$ by setting $Y_i \in S$, each entry chosen independently with probability $1/|S|$. Without knowledge of the value of T , each Y_i is distributed uniformly and independently at random and so Y gives the transitions of an infinite random walk on the graph. The partial sum $Y^{(s)} = \sum_{i=1}^s Y_i$ gives the walk position, and is equal to $\mathbf{0}$ (the combined walk visits the origin) at least once for each time step when the original two particles intersected up to the time step when they have moved s walk steps in total.

Let $r_D(s)$ be the probability $Y^{(s)} = \mathbf{0}$ and let $r(s)$ be the probability of a return to the origin at step s of a simple random walk on the graph. Then

$$\sum_{s=0}^t r_D(s) \leq \sum_{s=0}^{2t} r(s) = R(2t),$$

where $R(2t)$ is the expected number of returns to the origin of a random walk during $2t$ steps.

For a system of M particles dispersing from the origin, the above discussion bounds the expected number of meetings of a given pair of particles in t steps by $R(2t)$. Let $Z(t)$ be the number of pairwise meetings between M particles in t process steps. As any pair of particles makes at most

$2t$ walk steps in t process steps,

$$E(Z(t)) \leq \binom{M}{2} R(2t)$$

Given $\omega \rightarrow \infty$, take any t satisfying

$$\binom{M}{2} R(2t) \leq t/\omega. \quad (37)$$

By Markov's inequality we have $P(Z(t) \geq \omega \mathbf{E}(Z(t))) \leq 1/\omega$. Then with probability $1 - O(1/\omega)$

$$Z(t) \leq \omega \mathbf{E}(Z(t)) \leq \omega \binom{M}{2} R(2t) < t.$$

At least one pair of particles moves at each time step during dispersion, so the total number of meetings $Z(t')$ at or before t' must be at least t' . If not, the process has stopped before step t . $\square$

Upper bound on $\mathbf{T}_{\text{Disp}}$ for the two dimensional grid. *With probability $1 - O(1/\omega)$, a system of M particles disperses on the 2-dimensional grid in $2\omega M^2 \log M$ process steps.*

In the case of the 2-dimensional grid, the movement of particle i at time step t takes a value $X_i(t) \in \{(-1, 0), (1, 0), (0, 0), (0, -1), (0, 1)\}$, with $X_i(t) = (0, 0)$ (the additive identity) if the particle does not move. For any particle the next non-zero movement is uniformly distributed in the set of generators $S = \{(-1, 0), (1, 0), (0, -1), (0, 1)\}$ and is independent of the choice of particle, or of the action of any other particles.

The value of $R(2t)$ is (see e.g. [5] page 328)

$$R(2t) = \sum_{s=0}^t \binom{2s}{s}^2 \frac{1}{4^{2s}} \leq \log t + c,$$

for some constant c . By inspection we see that $t = 2\omega M^2 \log M$ satisfies (37). As at least 2 particles move at any step, $\mathbf{T}_{\text{Disp}} = O(\omega M^2 \log M)$.

6 Finite Abelian Cayley graphs: The Hypercube

In this section we present two different bounds for dispersion on the hypercube using two different methods. The first approach, given in the proof of Lemma 10 below, can be easily generalised to any finite Abelian Cayley graph to give a proof of Theorem 3 (ii). The second approach, Lemma 11, uses the structure of the hypercube.

The first method is similar to that used in the previous section for infinite Abelian Cayley graphs. The key difference is the difficulty in bounding the number of meetings of particles in a finite graph. To work around this, we first allow the dispersion process to approach mixing time, when the probability of being at any given vertex is close to uniform, assuming a trivial bound on the number of meetings in this initial period. Once we have reached this mixing time, we can use this uniformity to derive a bound on two particles being at the same vertex.

Lemma 10. *Let H_d be the hypercube on $n = 2^d$ vertices. Then with probability $1 - O(1/\omega)$, a system of $M \leq \sqrt{n}/\omega$ particles disperses on the hypercube H_d in $O(M^2 \log^2 n)$ process steps.*

Proof. The hypercube H_d on $n = 2^d$ vertices consists of vertices labeled as vectors in $\{0, 1\}^d$ and edges uv between vertices u and v whenever the vertex labels differ in a single coordinate (Hamming distance one). Let e_j be the vector whose entries are zero except at the j -th coordinate whose entry is one. A transition (u, v) of a random walk on H_d can be modeled by sampling $j \in \{1, \dots, d\}$ u.a.r. and setting $v = (u + e_j) \bmod 2$. We refer to e_j as the transition vector.

Mimicking the argument for the grid given previously, we consider the walk steps of particles 1 and 2, and build a sequence $Y(t) = Y(t, \{1, 2\})$ consisting of the transition vectors of the walks of particle 1 and particle 2 at each step in that order. Because $e_j + e_j = 0 \bmod 2$ we do not need to multiply the second transition vector by (-1) as in the case of the grid.

Let T be such that for all even $s \geq T$, $|P^s(u, u) - 2/n| \leq 1/n$, for all $u \in V$. Then $T = O(\log^2 n)$, and for all $u \in V$ and all even $s \geq T$, $P^s(u, u) \leq 3/n$. The expected number of meetings between a pair of particles in $T + t$ process steps is at most the expected number of returns to the origin of the walk $Y(t)$ within $2T + 2t$ steps, which is at most $2T + 6t/n$. Similarly as in the argument in the proof of Lemma 9, with probability at least $1 - O(1/\omega)$ the total number $Z(T + t)$ of pairwise meetings in $T + t$ process steps is at most

$$\binom{M}{2} 2T + \omega \binom{M}{2} \frac{6t}{n},$$

which is less than $T + t$, if $M \leq \sqrt{n}/\omega$ and $t = T \lceil M^2/(1 - 3\omega M^2/n) \rceil \leq 2TM^2$. Thus with probability at least $1 - O(1/\omega)$, all particles stop moving after fewer than $T + t = O(M^2 \log^2 n)$ process steps. $\square$

We can improve (for most values of M and n) on this upper bound for dispersion on the hypercube by using the technique we used in analysing regular trees.

Lemma 11. *Let H_d be the hypercube on $n = 2^d$ vertices and let $\omega \gg \sqrt[4]{d}$. Then with probability at least $1 - O(1/\omega)$, a system of $M \leq \sqrt{n}/\omega$ particles disperses on the hypercube H_d in $O(M \log^3 n)$ process steps, and no particle will lie further than $d/2$ from the origin.*

Proof. Note that if some particle reaches a distance greater than c from the origin, it must have first reached distance c and been joined there by another particle. If this second particle had previously already reached distance c before visiting the common vertex where these particles meet, then it must have met some other particle at the place where it first reached distance c . Iterating, we see that at some point, there must exist a pair of particles who first reached distance c at the same vertex.

There are $\binom{d}{c}$ vertices at distance c from the origin. The event that a given one of these vertices is the first a particle visits in its predetermined walk is uniformly distributed over all vertices at distance c . Therefore, provided the number of particles $M \ll \sqrt{\binom{d}{c}}$, the probability some pair of particles shares a common vertex at their first visit to distance c tends to 0.

If $c = d/2$, then $\binom{d}{c} = \binom{d}{d/2} \geq \frac{2^{d-1}}{\sqrt{d/2}}$. Since $n = 2^d$ and $M \leq \sqrt{n}/\omega$ it follows that $M = 2^{d/2}/\omega$. The probability that two particles share a common vertex as their first visit to distance $d/2$, is equal to $\binom{d}{c}^{-1} \leq \frac{\sqrt{d/2}}{2^{d-1}}$. Taking a union bound over all possible pairs, the probability of some pair of particles sharing a common vertex as their first visit to distance c is at most

$$\binom{M}{2} \frac{\sqrt{d/2}}{2^{d-1}} \leq \frac{M^2 \sqrt{d/2}}{2^d} \leq \frac{\sqrt{d/2}}{\omega^2}. \quad (38)$$

Choosing a value for ω such that $\omega \gg \sqrt[4]{d}$, the right hand side of (38) tends to zero. Thus, with high probability no particle travels further than $d/2$ from the origin in the dispersion process. This will allow us to upper bound the dispersion time of the process.

Consider the walk of a given particle. We note that if a particle is at distance less than $d/2$ from the origin, then since each vertex has degree d , it will move away from the origin with probability at least $1/2$, increasing its distance. Therefore the probability that this walk reaches distance $d/2$ in some number of steps is at least as high as that for a simple random walk on the line. The expected time it takes the simple walk on the line to reach distance $d/2$ is $(d/2)^2$ (by a simple martingale optional stopping time argument) and so the probability that the walk takes more than $\omega(d/2)^2$ steps to reach a distance of $d/2$ is less than ω^{-1} . In particular, the probability that it takes more than $d^2/2$ steps is less than $1/2$ and so with probability at least $1/2$ the walk will have reached $d/2$ by time $d^2/2$. If not, then we restart the analysis, treating the walk from this time onwards as a new random walk. Although we may not be at the origin, this only reduces the expected time to reach distance $d/2$ and so again with probability at least $1/2$, independently of the previous round, we will reach distance $d/2$ in the next $d^2/2$ steps. We iterate this process $d/2$ times, taking at most $d^3/4$ steps in total, and so the probability the walk has not finished after all these rounds, is at most,

$$(1 - 1/2)^{d/2} = 2^{-d/2} = n^{-1/2} \leq (\omega M)^{-1}.$$

Taking a union bound over the M particles, we see that with probability $1 - \omega^{-1}$, no particle reaches a distance of more than $d/2$, and thus stop moving after taking at most $d^3/4$ steps.

Since at any time step, at least two particles must be moving or the process has ended, we have that after $(M/2)(d^3/4)$ time steps, each particle must have moved at least $d^3/4$ walk steps and hence the process will end. This tells us that with probability at least $1 - \omega^{-1}$, the entire process terminates within this many time steps, so $\mathbf{T}_{\text{Disp}} = O(M \log^3 n)$. $\square$

7 Acknowledgements

We thank Tony Johansson and Fiona Skerman for spotting a mistake in an earlier version of the paper, and for offering suggestions on how to correct it.

References

- [1] H. Ackermann, S. Fischer, M. Hoefer, and M. Schöngens. Distributed algorithms for QoS load balancing. *Distributed computing*, 23(5-6), (2011) 321-330.
- [2] J. Beal. Superdiffusive dispersion and mixing of swarms with reactive Lévy walks. 7th International Conference on Self-Adaptive and Self-Organizing Systems (SASO), IEEE 2013.
- [3] J. Beal. Superdiffusive dispersion and mixing of swarms. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 2015.
- [4] P. Diaconis and W. Fulton. A growth model, a game, an algebra, Lagrange inversion, and characteristic classes, *Rend. Sem. Mat. Univ. Pol. Torino* 49 (1991), 95-119.

- [5] W. Feller. An Introduction to Probability Theory and its Applications, Volume I, 2nd Edition, Wiley, 1968.
- [6] A. M. Frieze and W. Pegden. A note on dispersing particles on a line. <https://arxiv.org/abs/1711.02192> (2017).
- [7] G. F. Lawler, M. Bramson, and D. Griffeath. Internal diffusion limited aggregation, Ann. Probab. 20 (1992), 2117-2140.
- [8] L. Levine and Y. Peres. The rotor-router shape is spherical. The Mathematical Intelligencer, Volume 27, (2005) 9-11.
- [9] T. Liggett. Stochastic Interacting Systems: Contact, Voter and Exclusion Processes. Springer-Verlag (1999).
- [10] T. Liggett. Interacting particle systems, an introduction. Lectures given at the School and Conference on Probability Theory, Trieste (2002).
- [11] F. Spitzer. Interaction of Markov Processes. Advances in Mathematics 5, (1970), 246–290

Appendix

Proof of Lemma 7

Lemma 12. *For a simple random walk on the integer path, let $R(2T, r)$ be the probability of at least r returns to the origin in $2T$ steps. For $\alpha > 0$, we have*

$$R(2T, \alpha\sqrt{2T}) = O\left(\frac{1}{\alpha}e^{-\alpha^2/2}\right). \quad (39)$$

Proof. Let $z_{2T}^{(r)}$ be the probability of exactly r returns to zero in $2T$ steps, then for $T \geq 1$, we have from Feller, Theorem 1, Section 3.6 in [5]

$$F_r = z_{2T}^{(r)} = \frac{1}{2^{2T-r}} \binom{2T-r}{T}.$$

We need to calculate

$$R(2T, r) = \sum_{s \geq r} F_s,$$

namely, the probability of at least r returns in time $2T$, for the case when $r = \alpha\sqrt{2T}$. If $r = T$ then $F_T = 1/2^T$ (very small).

For $s \geq r \geq 1$

$$F_{s+1}/F_s = \frac{2T-2s}{2T-s} \leq \frac{2T-2r}{2T-r} = F_{r+1}/F_r = \beta.$$

So

$$R(2T, r) \leq F_r(1 + \beta + \beta^2 + \dots) = F_r \frac{1}{1-\beta} = F_r \frac{2T-r}{r}.$$

Now, provided $r < T$

$$\begin{aligned} F_r &= (1 + (1/T) + (1/(T-r))) \frac{1}{\sqrt{2\pi}} \sqrt{\frac{2T-r}{T(T-r)}} \frac{2^r}{2^{2T}} \frac{(2T-r)^{2T-r}}{T^T (T-r)^{T-r}} \\ &= \Theta(1) \sqrt{\frac{2T-r}{T(T-r)}} \frac{(1-r/2T)^{2T-r}}{(1-r/T)^{T-r}}. \end{aligned}$$

But

$$\begin{aligned} \frac{(1-r/2T)^{2T-r}}{(1-r/T)^{T-r}} &= \exp((2T-r) \log(1-r/2T) - (T-r) \log(1-r/T)) \\ &= \exp\left(-\frac{r^2}{4T} - \frac{r^3}{8T^2} - \dots\right) \\ &\leq \exp\left(-\frac{r^2}{4T}\right). \end{aligned}$$

Thus

$$R(2T, r) \leq \Theta(1) \sqrt{\frac{2T-r}{T(T-r)}} \frac{2T-r}{r} \exp\left(-\frac{r^2}{4T}\right).$$

Put $r = \alpha\sqrt{2T}$ to obtain (assuming $r < T$)

$$R(2T, \alpha\sqrt{2T}) = O\left(\frac{1}{\alpha} e^{-\alpha^2/2}\right),$$

as required. □