

BROADCASTING ON PATHS AND CYCLES

REAZ HUQ AND PAWEŁ PRAŁAT

ABSTRACT. Consider the following broadcasting process run on a connected graph $G = (V, E)$. Suppose that $k \geq 2$ agents start on vertices selected from V uniformly and independently at random. One of the agents has a message that she wants to communicate to the other agents. All agents perform independent random walks on G , with the message being passed when an agent that knows the message meets an agent that does not know the message. The broadcasting time $\xi(G, k)$ is the time it takes to spread the message to all agents. We provide tight bounds for $\xi(P_n, k)$ and $\xi(C_n, k)$ that hold asymptotically almost surely for the whole range of the parameter k .

1. INTRODUCTION

In this paper, we investigate the problem of broadcasting messages between agents that randomly move on a connected graph $G = (V, E)$. The assumption is that $k \geq 2$ agents start the process at random locations on the graph and then perform a random walk along its vertices. One agent, selected in advance, initially possesses some information. If at some point during the process two agents meet at some vertex or pass each other at some edge and only one of them possesses the information, it is passed along to the other agent. The broadcasting time $\xi(G, k)$ is the time it takes to spread the message to all agents. (Formal definition will be provided in Section 2.)

The performance of a random walk in a network is a fundamental process that has found applications in many areas of computer science. Since this paper contains theoretical results, we will focus on prior results of related processes that were investigated rigorously and in a theoretical context. As this is still a very broad topic, we only scratch the surface and focus on multiple random walks performed simultaneously (which has many applications in distributed computing, such as sampling). For more on other directions, we direct the reader to one of the many books on Markov chains; see, for example [18].

Suppose there are $k \geq 2$ particles, each making a simple random walk on a graph G . Even if the particles are oblivious of each other, it is important and non-trivial to estimate the (vertex) cover time, an extensively studied graph parameter that is defined as the expected time required for the process to visit every vertex of G . The first paper [3] on this problem was concerned with the walk starting on the worst case vertices and subsequent papers [10, 22] dealt specifically with starting positions selected randomly from the stationary distribution. Questions become more interesting (and difficult) once we allow particles to interact once they meet. We assume that interaction occurs only when meeting at a vertex, and that the random walks made by

the particles are otherwise independent. There are at least four interesting variants of this process:

- *Predator-Prey*: estimate the expected time-to-extinction of the prey particles under the assumption that k predator and ℓ prey particles walk independently; predators eat prey particles upon meeting at a vertex.
- *Coalescing particles*: estimate the expected time to coalesce to a single particle under the assumption that k particles walk independently and coalesce upon meeting at a vertex.
- *Annihilating particles*: estimate the expected time-to-extinction of all particles under the assumption that $k = 2\ell$ particles walk independently and destroy each other (pairwise) upon meeting at a vertex.
- *Talkative particles*: estimate the expected time to broadcast a message—this is exactly the problem that we are concerned with in this paper.

All of these variants have been studied for random d -regular graphs $\mathcal{G}_{n,d}$ [6]. In particular, if $d \geq 3$ is a fixed constant and $k \leq n^\varepsilon$ for a sufficiently small constant $\varepsilon > 0$, then asymptotically almost surely (see the next section for a definition and a notation used)

$$\xi(\mathcal{G}_{n,d}, k) \sim \frac{2H_{k-1}}{k} \cdot \frac{d-1}{d-2} \cdot n.$$

Moreover, in a recent paper the authors of this paper (along with three co-authors) provided a complete characterization of $\xi(K_n, k)$ for the whole range of the parameter k . Interestingly, $\xi(K_n, k)$ is well concentrated around $2n \ln k/k$ for a wide range of possible values of k , but the behaviour changes when k is very large, namely, when k is linear in n [14]. These are the only theoretical results on the broadcasting time that we are aware of. However, the *frog model*, a well-known and well-studied epidemic model, is somewhat related to our problem [4]. There are a few differences between the two models. For example, in the frog model the number of agents that start on a given vertex is an independent Poisson random variable, some agents do not perform walks, and agents have a given lifespan. Though variations of the frog model have been studied (for example, in [23], where every agent performs a random walk), we have seen none which are direct analogues of the process we study in this paper.

On the other hand, the variant of coalescing particles is very well-studied, mainly because of its surprising connection to the voter model [2, 19]. The state of the process at a given time t is described by a function $\mu_t : V(G) \rightarrow O$, where $V(G)$ is the vertex set of a graph G and O is a given set of possible opinions. Each vertex $v \in V(G)$ “wakes up” at rate 1. When it wakes up at a time $t > 0$, v chooses one of its neighbours w uniformly at random and updates its value $\mu_t(v)$ to the opinion of w ; all other opinions remain the same. A classical duality result (see, for example, [2, 19]) directly relates the state of the process at a given time to a system of coalescing random walks on G moving backwards in time. As we already mentioned, there are many interesting results on the coalescing time. Let us only mention a beautiful conjecture posed by Aldous and Fill in the mid-nineties (Open problem 13, Chapter 14 of [2]). They conjectured an upper bound for the mean coalescent time in terms of the mean hitting time of a single random walk. The conjecture was proved in [20, 21].

Let us now briefly discuss the following well-known and well-studied rumour spreading protocols: *Push* and *Push & Pull*. Suppose that one vertex in a network is aware of a piece of information, the “rumour”, and wants to spread it to all vertices. In each round of the *Push* protocol, every informed vertex contacts a random neighbour and sends the rumour to it (“pushes” the rumour). In *Push & Pull*, uninformed vertices can also contact a random neighbour to get the rumour if the neighbour knows it (“pulls” the rumour).

There is a long sequence of interesting and important papers studying the runtime of *Push* on the complete graph. The first paper considering this protocol is [11] but more precise bounds were provided in [24] and then in [9], in which it was shown that the process can essentially be stochastically bounded (from both sides) by coupon collector-type problems. A very recent paper [7] both determines the limiting distribution and explains why it is a difficult problem: the runtime, scaled appropriately by $(\log_2 n + \ln n)$, has no limiting distribution; instead, it exhibits double-oscillatory behaviour. *Push* has been extensively studied on several other graph classes besides complete graphs.

The *Push & Pull* protocol has an equally long sequence of interesting papers studying it. The synchronous version of the protocol (as described above) was introduced in [8] and popularized in [16]. However, such synchronized models (that is, models in which all vertices take action simultaneously at discrete time steps) are not plausible for many applications, including real-world social networks. As a result, an asynchronous version of the model with a continuous timeline was introduced in [5]. In this variant, each vertex has its own independent clock that rings at the times of a rate 1 Poisson process with the protocol specifying what a vertex has to do when its own clock rings. The first theoretical relationships between the spread times in the two variants was provided in [1].

In this paper, we focus on paths and cycles. We show that the behaviour of $\xi(G, k)$ is similar for both families of graphs and is approximately equal to n^2/k (that is, up to poly-log factors), provided that $k \leq n$. (See Theorem 2.1 for the precise statement.) The paper is structured as follows. In the next section, we formally define our problem, introduce asymptotic notation, and state the main result. The whole of Section 3 is devoted to proving the main result.

2. FORMULATION OF THE PROBLEM AND THE MAIN RESULT

In this section, we formally define the process we aim to analyze (Subsection 2.1). Though we define it for any connected graph, in this paper we focus on paths and cycles. As our results are asymptotic in nature, we need to introduce the asymptotic notation that is used throughout the entire paper (Subsection 2.2). Finally, we state the main result that combines all ranges for the number of agents involved (Subsection 2.3).

2.1. Problem. Suppose that we are given a connected graph $G = (V, E)$ on $n = |V|$ vertices, and let $k \geq 2$ be any natural number. There are k agents, one of which is green with the rest being white. The process starts at round $t = 0$ with agents located randomly on vertices of G ; that is, each agent starts at any vertex $v \in V$ with probability $1/n$, independently of other agents and independently of her colour. Each

agent synchronously performs an independent random walk, regardless of whether she is green or white. In other words, an agent occupying vertex $v \in V$ moves to any neighbour of v with probability equal to $1/\deg(v)$. A white agent becomes green when she meets a green agent at some *round* $t \geq 0$. In particular, all agents that start at the same vertex as the initial green agent become green at the very beginning. For most graphs it does not make a substantial difference but, for example, for bipartite graphs this definition has a flaw: if G is a bipartite graph with parts X, Y , a given white agent and the green agent will not meet if they do not start on the same partition of the graph. In order to solve this potential issue, we will allow a white agent to become green if they move from u to v at the same round that a green agent moves from v to u . Let $\xi = \xi(G, k)$ be the time it takes for all agents to become green. (Note that ξ is a random variable even when G is a deterministic graph.)

We say that the process is at *phase* ℓ ($1 \leq \ell \leq k$) if there are ℓ green agents (and so $k - \ell$ white agents). The first phase is usually phase 1, unless some white agents start at the same vertex as the green agent but this is rare if k is small. Clearly, the process always moves from a smaller phase to a larger phase. If k is small, then typically it takes some number of rounds for the process to move to another phase though some phases may be skipped. If k is large, then skipping phases is quite common. The process ends at the end of round ξ , when we are about to move to phase k .

2.2. Asymptotic Notation. Our results are asymptotic in nature, that is, we will assume that $n \rightarrow \infty$. Formally, we consider a sequence of graphs $G_n = (V_n, E_n)$ (paths and cycles on n vertices) and $k = k(n)$ may be a function of n that tends to infinity as $n \rightarrow \infty$. We are interested in events that hold *asymptotically almost surely* (*a.a.s.*), that is, events that hold with probability tending to 1 as $n \rightarrow \infty$.

Given two functions $f = f(n)$ and $g = g(n)$, we will write $f(n) = O(g(n))$ if there exists an absolute constant $c \in \mathbb{R}_+$ such that $|f(n)| \leq c|g(n)|$ for all n , $f(n) = \Omega(g(n))$ if $g(n) = O(f(n))$, $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$, and we write $f(n) = o(g(n))$ or $f(n) \ll g(n)$ if $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$. In addition, we write $f(n) \gg g(n)$ if $g(n) = o(f(n))$ and we write $f(n) \sim g(n)$ if $f(n) = (1 + o(1))g(n)$, that is, $\lim_{n \rightarrow \infty} f(n)/g(n) = 1$.

Finally, for any $\ell \in \mathbb{N}$ we will use $[\ell]$ to denote the set of ℓ smallest natural numbers, that is, $[\ell] := \{1, 2, \dots, \ell\}$.

2.3. Main Result. Let us summarize the main results for paths and cycles in one theorem. More detailed and stronger statements can be found in the next section.

Theorem 2.1. *Let $\omega = \omega(n)$ be any function that tends to infinity as $n \rightarrow \infty$. Let $G = P_n$ (a path on n vertices) or $G = C_n$ (a cycle on n vertices). Depending on the parameter $k = k(n)$, the following properties hold a.a.s.:*

(a) *If $k \leq \omega$, then*

$$\frac{n^2}{\omega^3 \ln \omega} \leq \xi(G, k) \leq n^2 \omega.$$

(b) If $\omega < k \leq \omega \ln n$, then

$$\frac{n^2}{\omega k^2 \ln k} \leq \xi(G, k) \leq n^2 \omega.$$

(c) If $\omega \ln n < k \leq n / \ln^2 n$, then

$$\Omega \left(\frac{n^2}{k(\ln k)(\ln n)} \right) = \xi(G, k) = O \left(\frac{n^2 \ln n}{k} \right).$$

(d) If $n / \ln^2 n < k \leq 50n \ln n$, then

$$\Omega(n) = \xi(G, k) = O \left(\frac{n^2 \ln n}{k} \right).$$

(e) If $k = k(n) \geq 50n \ln n$, then $\xi(G, k) = \Theta(n)$.

In any case, if $k = n^{x+o(1)}$ for some $x \in [0, 1]$, then a.a.s.

$$\xi(G, k) = \frac{n^{2+o(1)}}{k} = n^{2-x+o(1)}.$$

Parts (a) and (b) follow from Theorems 3.7 and 3.5. Parts (c) and (d) follow from Theorems 3.10 and 3.9. Finally, part (e) follows from Theorem 3.4. The arguments for paths P_n and cycles C_n are almost the same. Because of the boundary effect, the argument for paths is usually slightly more challenging. In order to avoid reproving theorems for the two classes, we provide a coupling that shows that $\xi(P_n, k) \leq \xi(C_{2(n-1)}, k)$, provided that $k = o(n)$ —see Lemma 3.1. It is a simple but interesting and useful observation as it proves that to establish the asymptotic behaviour for both classes, one only needs to prove upper bounds for cycles and lower bounds for paths. Having said that, we have to admit that we more often than not make exceptions to this rule. For example, the argument for large values of k (Section 3.4) is the same for both classes of graphs so there is no need for coupling which does not apply to this range of values of k anyway. The proofs for the corresponding upper bounds for small values of k (Section 3.6) are substantially different for paths and cycles, so we decided to include both arguments. (Of course, the result for paths is alternatively implied by the coupling.) Finally, since part of the range of medium values of the parameter k (Section 3.7) is not covered by the coupling, we decided to present an argument for paths that is simpler and only mention straightforward adjustments to cycles.

3. PROOFS

This whole section is devoted to proving Theorem 2.1. We investigate the process running on P_n , the path on n vertices, and C_n , the cycle on n vertices. It will be convenient to label vertices of P_n as follows: $V(P_n) = [n] := \{1, 2, \dots, n\}$ and $E(P_n) = \{i(i+1) : i \in [n-1]\}$. Similarly, $V(C_n) = [n] := \{1, 2, \dots, n\}$ and $E(C_n) = \{i(i+1) : i \in [n-1]\} \cup \{1n\}$.

The proofs require unique approaches depending on the number of agents involved (parameter $k = k(n)$). Thus, we will deal with each sub-range of k independently. However, before we start, let us state some concentration inequalities that we will use often, and introduce the coupling between the processes run on paths and cycles.

3.1. Chernoff inequality. Throughout the paper, we will be using the following concentration inequality: let $X \in \text{Bin}(n, p)$ be a random variable characterized as a binomial distribution with parameters n and p . Then, a consequence of *Chernoff's bound* (see e.g. [15, Corollary 2.3]) is that

$$\mathbb{P}(|X - \mathbb{E}[X]| \geq \varepsilon \mathbb{E}[X]) \leq 2 \exp\left(-\frac{\varepsilon^2 \mathbb{E}[X]}{3}\right) \quad (1)$$

for $0 < \varepsilon < 3/2$.

3.2. Hoeffding-Azuma inequality. Let X_0, X_1, \dots be an infinite sequence of random variables that is a martingale; that is, for any $a \in \mathbb{N}$ we have $\mathbb{E}[X_a | X_{a-1}] = X_{a-1}$. Suppose that there exist constants $c_a > 0$ such that $|X_a - X_{a-1}| \leq c_a$ for each $a \leq t$. Then, the *Hoeffding-Azuma inequality* implies that for every $b > 0$,

$$\mathbb{P}(\exists i(0 \leq i \leq t) : |X_i - X_0| \geq b) \leq 2 \exp\left(-\frac{b^2}{2 \sum_{a=1}^t c_a^2}\right). \quad (2)$$

3.3. Coupling. We will show now that one may couple the processes run on paths and cycles. This coupling will allow us to translate bounds obtained for one class to another one. The argument applies provided that $k = o(n)$, and the coupling can only be established a.a.s. but it is enough as our main result holds a.a.s. anyway.

Lemma 3.1. *Suppose that $k = o(n)$. The processes on P_n and $C_{2(n-1)}$ can be coupled such that a.a.s.*

$$\xi(P_n, k) \leq \xi(C_{2(n-1)}, k).$$

Proof. To simplify the notation, we are going to label the vertices of P_n and $C_{2(n-1)}$ slightly differently than in the rest of the paper. Vertices of P_n are labelled as follows: $V(P_n) = \{0, 1, \dots, n-1\}$. On the other hand, vertices of $C_{2(n-1)}$ have labels from the set $V(C_{2(n-1)}) = \{-(n-2), -(n-3), \dots, -1, 0, 1, \dots, (n-2), (n-1)\}$. The coupling will identify vertices i and $-i$ on $C_{2(n-1)}$ with a vertex i on P_n ($i \in [n-2]$); vertex 0 and $n-1$ on the cycle will be mapped to 0 and, respectively, $n-1$ on the path—see Figure 1. There is a slight complication with making sure the agents start the process from a uniform distribution on the corresponding set of nodes. Because of that the result holds only a.a.s. and we need an assumption that $k = o(n)$.

Agents start independently and uniformly at random on $C_{2(n-1)}$, as they should. Each agent, independently, becomes *unusual* with probability $1/n$. Unusual agents put their *avatars* on P_n at vertex 0 with probability $1/2$, and on vertex $n-1$ otherwise. Usual (that is, not unusual) agents that start at vertex i on $C_{2(n-1)}$, place their avatars on P_n at vertex $|i|$. It is easy to see that avatars are distributed uniformly at random on P_n ; in particular, the probability that a given agent puts her avatar at vertex 0 on the path is equal to

$$\frac{1}{n} \cdot \frac{1}{2} + \left(1 - \frac{1}{n}\right) \cdot \frac{1}{2(n-1)} = \frac{1}{2n} + \frac{1}{2n} = \frac{1}{n}.$$

Since it is assumed that $k = o(n)$, the expected number of unusual agents is equal to $k/n = o(1)$ and so a.a.s. there is no unusual agent by the first moment method. If there is at least one unusual agent, then we simply stop the coupling and claim no




FIGURE 1. Coupling between $C_{2(n-1)}$ and P_n . Agents walking on the cycle and their avatars walking on the path are in red cycles.

bound for the two random variables. On the other hand, if there is no unusual agent, then agent occupying vertex i on $C_{2(n-1)}$ has her avatar placed on vertex $|i|$ on P_n and we may continue. Agents start walking randomly on the cycle and their avatars will follow them accordingly, that is, if an agent moves from vertex a to vertex b on the cycle, her avatar moves from vertex $|a|$ to vertex $|b|$ on the path. It is clear that avatars also perform independent random walks but on the path instead of the cycle. The two corresponding random walks are coupled but the broadcasting processes are performed independently on both graphs. By coupling, if two agents meet, then their avatars also meet but the converse might not be true—avatars meeting on the path might correspond to agents that occupy different vertices (again, see Figure 1). Hence, when all agents become green, then all avatars must be green too. This establishes the desired relationship between $\xi(P_n, k)$ and $\xi(C_{2(n-1)}, k)$, and the proof of the lemma is finished. \square

3.4. Large k . In this section, we show that a.a.s. $\xi(P_n, k) = \Theta(n)$, provided that k is sufficiently large, that is, $k \geq 50n \ln n$. The coupling (Lemma 3.1) cannot be applied to this range of parameter k but the proof works for cycles as well as paths. This proves part (e) of Theorem 2.1.

Let us start with the following useful observation. Suppose that an agent starts at vertex $i \in [n]$. For a given $t \in \mathbb{N} \cup \{0\}$ and $j \in [n]$, let $P_t(i, j)$ be the probability that the agent occupies vertex j at round t . Then the following holds.

Lemma 3.2. *For any $t \in \mathbb{N} \cup \{0\}$ and $i, j \in [n]$, we have that*

$$P_t(i, j) = P_t(j, i), \quad (3)$$

provided that $i \notin \{1, n\}$ and $j \notin \{1, n\}$. More importantly, we always have that

$$\frac{P_t(j, i)}{2} \leq P_t(i, j) \leq 2P_t(j, i). \quad (4)$$

Proof. The lemma is an instant corollary of the fact that the associated simple random walk is reversible, see [18, Section 1.6], and that the stationary distribution is near

uniform. Indeed, since walking on a graph is reversible,

$$\pi(i)P_t(i, j) = \pi(j)P_t(j, i),$$

where $\pi(x)$ is the stationary distribution. The conclusion follows from the fact that

$$\pi(x) = \frac{\deg(x)}{2|E|} = \begin{cases} \frac{1}{2(n-1)} & \text{if } x \in \{1, n\}, \\ \frac{2}{2(n-1)} & \text{otherwise.} \end{cases}$$

The proof of the lemma is finished. \square

Adjustment to cycles: The lemma holds for cycles. In fact, property (3) holds for all i and j (since $\pi(x)$ is uniform on cycles) and so the weaker property (4) trivially holds.

We will now show that there are plenty of agents on each vertex at any round of the process, provided that it ends in at most n rounds.

Lemma 3.3. *Consider the process on a path P_n with $k = k(n) \geq 50n \ln n$ agents. Then, a.a.s. the following holds: for any $t \in [n] \cup \{0\}$ and any $j \in [n]$, the number of agents occupying vertex j at round t is at least $12 \ln n$.*

Proof. Fix any $t \in [n] \cup \{0\}$ and any $j \in [n]$, and let us concentrate on a given agent A . Let $B(i)$ be the event that agent A starts at vertex i , and let $C(j)$ be the event that agent A occupies vertex j at time t . It follows that

$$\mathbb{P}(C(j)) = \sum_{i \in [n]} \mathbb{P}(C(j) \wedge B(i)) = \sum_{i \in [n]} \mathbb{P}(C(j) \mid B(i)) \cdot \mathbb{P}(B(i)).$$

Since agent A starts on a vertex selected uniformly at random from V , $\mathbb{P}(B(i)) = 1/n$. After noticing that $\mathbb{P}(C(j) \mid B(i))$ is exactly $P_t(i, j)$, we get from Lemma 3.2 that

$$\mathbb{P}(C(j)) = \frac{1}{n} \sum_{i \in [n]} P_t(i, j) \geq \frac{1}{2n} \sum_{i \in [n]} P_t(j, i) = \frac{1}{2n},$$

as, trivially, $\sum_{i \in [n]} P_t(j, i) = 1$.

Since $k \geq 50n \ln n$ agents select their starting points independently and perform independent random walks afterwards, the number of agents occupying vertex j at round t can be stochastically lower bounded by the random variable $X \sim \text{Bin}(50n \ln n, 1/(2n))$. Note that $\mathbb{E}[X] = 25 \ln n$ and so it follows from the Chernoff inequality (1) applied with $\varepsilon = 1/2$ that

$$\begin{aligned} \mathbb{P}(X \leq 12 \ln n) &\leq \mathbb{P}\left(|X - \mathbb{E}[X]| \geq \frac{1}{2} \mathbb{E}[X]\right) \\ &\leq 2 \exp\left(-\frac{1}{12} \mathbb{E}[X]\right) = 2n^{-25/12} = o(n^{-2}). \end{aligned}$$

Since there are $n+1$ choices for t and n choices for j , the desired property fails for some pair of t and n with probability at most $n(n+1) \cdot o(n^{-2}) = o(1)$ and so the desired property holds a.a.s. and the proof is finished. \square

Adjustment to cycles: Since Lemma 3.2 holds also for cycles, the exact same proof of the above lemma extends to cycles.

Now, we are ready to show both an upper and a lower bound for $\xi(P_n, k)$ for $k = k(n) \geq 50n \ln n$.

Theorem 3.4. *For any $k = k(n) \geq 50n \ln n$, a.a.s. $\lfloor n/2 \rfloor \leq \xi(P_n, k) \leq n - 1$.*

Proof. By Lemma 3.3, since we aim for a result that holds a.a.s., we may assume that for any round $t \in [n] \cup \{0\}$ and any vertex $j \in [n]$, the number of agents occupying vertex j at round t is at least $12 \ln n$. We will say that a vertex j is green if it is occupied by green agents; otherwise, it is white, that is, it is occupied by white agents. In particular, this means that in the very first round ($t = 0$) there is precisely one green vertex.

Suppose that at the end of some round $t \in [n - 2] \cup \{0\}$, some vertex j is green whereas a neighbouring vertex i ($i \in \{j - 1, j + 1\}$) is white. Since there are at least $12 \ln n$ green agents occupying j at the end of round t , the probability that i stays white in round $t + 1$ is at most

$$(1/2)^{12 \ln n} = \exp(-12(\ln 2) \ln n) \leq n^{-8} = o(n^{-1}).$$

After applying this argument $n - 1$ times, we get that a.a.s. at round $t \in [n - 1] \cup \{0\}$ vertices at distance at most t from the initial green vertex are green. In particular, a.a.s. all vertices become green in at most $n - 1$ rounds, and so the desired upper bound holds. In fact, conditioning on the event that the initial green vertex is vertex $j \in [n]$, we get that a.a.s. $\xi(P_n, k) = \max\{j - 1, n - j\}$. Since $\max\{j - 1, n - j\} \geq \lfloor n/2 \rfloor$, the desired lower bound holds too, and the proof is finished. \square

Adjustment to cycles: The same argument works for cycles. Due to the symmetry, a.a.s. $\xi(C_n, k) = \lfloor n/2 \rfloor$.

3.5. Walking on Integers. Let us take a short break from our problem and briefly discuss a closely related and classical problem: walking on integers. The *simple random walk on \mathbb{Z}* starts with $X_0 = 0$ and in each round $t \in \mathbb{N}$, $X_t = X_{t-1} - 1$ with probability $1/2$; otherwise, $X_t = X_{t-1} + 1$. Alternatively, the *lazy simple random walk on \mathbb{Z}* starts with $X_0 = 0$ and in each round $t \in \mathbb{N}$, $X_t = X_{t-1} - 1$ with probability $1/4$, $X_t = X_{t-1} + 1$ with probability $1/4$, and $X_t = X_{t-1}$ otherwise.

It is easy to see that the sequence X_0, X_1, \dots is a martingale. In particular, the Hoeffding-Azuma inequality (2) can be applied to show that for small values of t , X_t has to be relatively close to zero a.a.s. On the other hand, if t is large, then a.a.s. X_t moves away from the origin. We will need this well-known observation to establish some of our bounds. We provide the proof for completeness but for more details we direct the reader to, for example, [18] or any other book on random walks.

Let us first concentrate on the simple random walk. Observe that X_t and t are of the same parity, that is, $t - X_t$ is even. Provided that $t - a$ is even, there are $\binom{t}{\frac{t-a}{2}}$ walks of length t from 0 to a . Combining the two observations together we get that for any

$$-t \leq a \leq t$$

$$\mathbb{P}(X_t = a) = \begin{cases} \binom{t}{\frac{t-a}{2}} 2^{-t} & \text{if } t-a \text{ is even,} \\ 0 & \text{otherwise.} \end{cases}$$

It follows that for any $-t \leq a \leq t$, we have

$$\mathbb{P}(X_t = a) \leq \binom{t}{\lceil t/2 \rceil} 2^{-t} = \frac{t!}{\lceil t/2 \rceil! \lfloor t/2 \rfloor!} 2^{-t} \sim \sqrt{\frac{2}{\pi t}},$$

where the asymptotic bound follows from *Stirling's formula* ($t! \sim \sqrt{2\pi t}(t/e)^t$). Similarly, for the lazy simple random walk, we get that for any $-t \leq a \leq t$

$$\mathbb{P}(X_t = a) \leq \sum_{s=0}^t \mathbb{P}(\text{Bin}(t, 1/2) = s) \cdot \binom{s}{\lceil s/2 \rceil} 2^{-s}.$$

(Variable s in the above formula controls the number of rounds the walk actually moves.) Chernoff's bound (1) applied with $\varepsilon = 1/t^{1/3}$ implies that

$$\begin{aligned} \mathbb{P}(X_t = a) &\leq o(1/t) + \sum_{s=t/2-t^{2/3}}^{t/2+t^{2/3}} \mathbb{P}(\text{Bin}(t, 1/2) = s) \cdot \binom{s}{\lceil s/2 \rceil} 2^{-s} \\ &\leq o(1/t) + (1 + o(1)) \sum_{s=t/2-t^{2/3}}^{t/2+t^{2/3}} \mathbb{P}(\text{Bin}(t, 1/2) = s) \cdot \sqrt{\frac{2}{\pi s}} \\ &\leq o(1/t) + (1 + o(1)) \sqrt{\frac{2}{\pi(t/2)}} \sum_{s=t/2-t^{2/3}}^{t/2+t^{2/3}} \mathbb{P}(\text{Bin}(t, 1/2) = s) \sim \sqrt{\frac{4}{\pi t}}. \end{aligned}$$

Hence, regardless of whether we deal with lazy random walks or not, for t large enough and any $a \geq 1$ we have

$$\mathbb{P}(|X_t| < a) \leq \frac{4a}{\sqrt{t}}. \quad (5)$$

We will also need the following result on the hitting time defined as follows:

$$\tau_a = \min\{t \geq 0 : X_t = a\},$$

that is, τ_a is the first time the walk hits a . Using the *reflection principle*, one can show that

$$\mathbb{P}(\tau_a > t) = \mathbb{P}(-a < X_t \leq a) \leq \frac{4a}{\sqrt{t}}. \quad (6)$$

(See Lemma 2.21 in [18] that applies to both lazy and non-lazy simple random walks.)

3.6. Small k . Let $\omega = \omega(n)$ be any function that tends to infinity as $n \rightarrow \infty$. In this section, we show that a.a.s. $\xi(P_n, k) \geq n^2/(\omega k^2 \ln k)$ and $\xi(C_n, k) \leq n^2\omega$, provided that $k \leq \omega \ln n$. These two bounds, together with the coupling (Lemma 3.1), prove parts (a) and (b) of Theorem 2.1.

We start by proving an upper bound for $\xi(C_n, k)$. It is a strong bound for small values of k but a weak one for large values of k . Recall that, in particular, $\xi(C_n, k) = O(n)$

for $k \geq 50n \ln n$. However, since it holds for all values of k , we state it here in full generality.

Theorem 3.5. *Let $\omega = \omega(n)$ be any function that tends to infinity as $n \rightarrow \infty$. For any $k = k(n) \geq 2$, a.a.s. $\xi(C_n, k) \leq n^2\omega$.*

Proof. We will couple our process with the random walk on integers we discussed above in the most natural way. If an agent starts at vertex $i \in [n]$ on the cycle, her avatar starts at integer $X_0 = i$ on \mathbb{Z} . If X_t increases, then the agent occupying vertex $i < n$ moves to $i + 1$ and she moves to 0 if she occupies vertex n . Similarly, if X_t decreases, then the agent occupying vertex $i > 1$ moves to $i - 1$ and she moves to n if she occupies vertex 1.

Concentrate on the initial green agent and an arbitrary white agent that are at distance d from each other. Our goal is to control random variable Y_t , the “distance” between the corresponding avatars walking on \mathbb{Z} . We initiate the auxiliary process with $Y_0 = d$ and for each $t \in \mathbb{N}$, $Y_t = Y_{t-1} + 2$ with probability $1/4$, $Y_t = Y_{t-1} - 2$ with probability $1/4$, and $Y_t = Y_{t-1}$ otherwise. (Hence, effectively, it is a lazy random walk.) The distance between the two avatars is $|Y_t|$.

Consider the first $t = n^2\omega$ rounds. Our bound (5) applied with $a = 2n + 1$ implies that a.a.s. $|Y_t| \geq 2n + 1$. But this implies that the two agents met at some point (when $|Y_t| = n$ or $|Y_t| = n + 1$) and then met again (when $|Y_t| = 2n$ or $|Y_t| = 2n + 1$), after making in the meantime everyone else green. \square

Adjustment to paths: This is the only situation when the argument for cycles *cannot* be easily adjusted to deal with paths. We provide an independent, direct argument if one does not want to use the coupling between the two families of graphs.

Theorem 3.6. *Let $\omega = \omega(n)$ be any function that tends to infinity as $n \rightarrow \infty$. For any $k = k(n) \geq 2$, a.a.s. $\xi(P_n, k) \leq n^2\omega$.*

Proof. We will estimate the number of rounds needed for the initial green agent to travel to one of the endpoints of the path and then to walk to the other endpoint. We will show that a.a.s. it happens in at most $n^2\omega$ rounds. This will finish the proof as it guarantees that all other agents have to meet her at some point and so all of them eventually become green.

We will couple our process with the random walk on integers we discussed above in the most natural way. If X_t increases and the agent occupies vertex $i < n$, she moves to $i + 1$. Similarly, if X_t decreases and the agent occupies vertex $i > 1$, she moves to $i - 1$. However, if she occupies one of the endpoints of the path (vertex 1 or vertex n), her move is deterministic as she is forced to stay on the path, regardless of what the random walk does.

Consider the first $t = n^2\omega/2$ rounds. Our bound (5) applied with $a = n$ implies that a.a.s. $|X_t| \geq n$. This implies that the agent must bump into one of the endpoints (say, vertex 1) during this time period (say, at time $T \leq t$). Let us now concentrate on the next $t = n^2\omega/2$ rounds following time T and let us restart the coupled random walk by fixing $X_T = 1$. Using (6) applied with $a = n$ we conclude that a.a.s. the random walk hits integer n during that period of time, and so the agent has to visit vertex n as well. This concludes the proof. \square

Let us now turn our attention to a lower bound for $\xi(P_n, k)$.

Theorem 3.7. *Let $\omega = \omega(n)$ be any function that tends to infinity as $n \rightarrow \infty$. For any $k = k(n) \leq \omega \ln n$, a.a.s. $\xi(P_n, k) \geq n^2/(\omega k^2 \ln k)$.*

Proof. Note that with probability $1 - O(1/\omega^{1/3}) \sim 1$, the initial green agent starts the process at distance at least $n/\omega^{1/3} = o(n)$ from both endpoints. Similarly, with probability $1 - O(1/(k\omega^{1/3}))$, a given white agent starts the process at distance at least $n/(k\omega^{1/3}) = o(n/k)$ from the green agent. Hence, all white agents are at distance at least $n/(k\omega^{1/3})$ from the green agent with probability

$$(1 - O(1/(k\omega^{1/3})))^{k-1} = 1 - O(k/(k\omega^{1/3})) \sim 1.$$

Since we aim for a conclusion that holds a.a.s., we may assume that this property is satisfied at the end of round 0.

Trivially, at the end of the whole process (that is, when all agents become green), at least one agent (either the one that was initially green or one of the white ones) has to move at least $n/(2k\omega^{1/3})$ away from her initial position; otherwise, no white agent turns green. We will show that this is highly unlikely after only $n^2/(\omega k^2 \ln k)$ rounds. Applying the Hoeffding-Azuma inequality (2) with $b = n/(2k\omega^{1/3})$, $c_a = 1$, and $t = n^2/(\omega k^2 \ln k)$ implies that a given agent moves that far with probability at most

$$2 \exp\left(-\frac{b^2}{2t}\right) = 2 \exp\left(-\frac{\omega^{1/3}}{8} \ln k\right) = o(1/k).$$

Hence, the probability that at least one agent moves far is $O(k) \cdot o(1/k) = o(1)$, and the proof is finished. \square

Adjustment to cycles: The argument is easily adjusted for cycles. In fact, it is slightly simpler as one does not need to pay attention to the two endpoints of the path.

3.7. Medium k . Let $\omega = \omega(n)$ be any function that tends to infinity as $n \rightarrow \infty$. In order to prove upper bounds in parts (c) and (d) of Theorem 2.1, we need to concentrate on $k = k(n) \gg \ln n$ and $k = k(n) < 50n \ln n$.

Let us start with the following definition. We partition the set of vertices of the path P_n into $b = b(n) := \lfloor k/(500 \ln n) \rfloor$ blocks. (Note that $b \gg 1$ and $b < n/10$.) Each block consists of either $\lfloor n/b \rfloor \geq \lfloor 500n \ln n/k \rfloor \geq 10$ or $\lceil n/b \rceil$ vertices.

We will first adjust the proof of Lemma 3.3 to show that there are plenty of agents on each block at any round of the process, provided that it ends in at most n^2 rounds. The adjustment is easy and straightforward but we provide the proof for completeness. Moreover, since Lemma 3.2 holds for both paths and cycles, the lemma below holds for both families of graphs too.

Lemma 3.8. *Consider the process on a path P_n with k agents such that $\ln n \ll k = k(n) < 50n \ln n$. Then, a.a.s. the following holds: for any $t \in [n^2] \cup \{0\}$ and any $j \in [b]$, the number of agents occupying block j at round t is at least $100 \ln n$ and at most $1800 \ln n$.*

Proof. The number of agents occupying block j at round t can be stochastically lower bounded by random variable $X \sim \text{Bin}(k, (450n \ln n/k)/(2n))$; recall that each block has length at least $\lfloor 500n \ln n/k \rfloor \geq (9/10)(500n \ln n/k) = 450n \ln n/k$ and, by Lemma 3.2, each vertex is occupied by a given agent with probability at least $1/(2n)$. Note that $\mathbb{E}[X] = 225 \ln n$ and so it follows from the Chernoff inequality (1) applied with $\varepsilon = 1/2$ that

$$\begin{aligned} \mathbb{P}(X \leq 100 \ln n) &\leq \mathbb{P}\left(|X - \mathbb{E}[X]| \geq \frac{1}{2} \mathbb{E}[X]\right) \\ &\leq 2 \exp\left(-\frac{1}{12} \mathbb{E}[X]\right) \leq 2n^{-225/12+o(1)} = o(n^{-3}). \end{aligned}$$

Since there are $n^2 + 1$ choices for t and $b = \lfloor k/(500 \ln n) \rfloor < n/10$ choices for j , the desired property fails for some pair of t and j with probability at most

$$(n^2 + 1)(n/10) \cdot o(n^{-3}) = o(1).$$

It follows that the desired lower bound for the number of agents holds a.a.s.

Similarly, the number of agents occupying block j at round t can be stochastically upper bounded by random variable $Y \sim \text{Bin}(k, (600n \ln n/k)(2/n))$; recall that each block has length at most $\lceil n/b \rceil \leq \lceil 501n \ln n/k \rceil \leq (11/10)(501n \ln n/k) \leq 600n \ln n/k$ and, by Lemma 3.2, each vertex is occupied by a given agent with probability at most $2/n$. We get that $\mathbb{E}[Y] = 1200 \ln n$ and $Y \geq (3/2)\mathbb{E}[Y] = 1800 \ln n$ with probability $o(n^{-3})$. The desired upper bound holds a.a.s. too and the proof is finished. \square

Adjustment to cycles: As mentioned above, the above lemma holds for cycles and the proof is exactly the same since Lemma 3.2 holds for both paths and cycles.

We are now ready to prove an upper bound. Note that the coupling (Lemma 3.1) cannot be applied to the whole range of parameter k . Since the arguments used to deal with cycles and paths are the same but some technicalities are slightly more involved for paths, we decided to present an argument for paths instead of using the coupling.

Theorem 3.9. *Let $\omega = \omega(n)$ be any function that tends to infinity as $n \rightarrow \infty$. For any k such that $\omega \ln n \leq k = k(n) < 50n \ln n$, a.a.s. $\xi(P_n, k) = O(n^2 \ln n/k)$.*

Proof. Since the argument is quite involved, let us first provide a high level overview of the proof. First, we will show that the initial green agent quickly meets a white agent. Then, we will track the distance between them and show that at some point they are far apart from each other so that they are separated by at least one block. As a result, not only are these two agents green but, in particular, all agents present on that internal block are green. The final step is to show that the sequence of blocks consisting of only green agents keeps expanding, eventually reaching both endpoints of the path.

Without loss of generality, we may assume that $\omega \leq \ln \ln n$. Suppose that the initial green agent starts the process at vertex $j \in [n]$. Clearly, a.a.s. the initial green agent is at distance at least $n/\omega^{1/5}$ from both endpoints of the path, that is, $j > n/\omega^{1/5}$ and $i < n - n/\omega^{1/5}$. The probability that no white agent starts the process at distance at

most $d := \omega n/k$ from the initial green agent is at most

$$\left(1 - \frac{d}{n}\right)^{k-1} \leq \exp\left(-\frac{\omega(k-1)}{k}\right) \leq \exp(-\omega/2) = o(1).$$

Since we aim for the statement that holds a.a.s., we may assume that initially the green agent is at distance at least $n/\omega^{1/5}$ from both endpoints and one of the white agents starts at distance at most $d = \omega n/k = o(n/\omega^{1/5})$ from her.

Let us focus on the initial green agent and a white agent that is initially the closest to her (if there are multiple white agents with this property, pick one of them arbitrarily). Provided that they did not yet meet, our process can be coupled with the lazy simple random walk on \mathbb{Z} starting at $X_0 = \lceil d/2 \rceil$ in such a way that the distance at round t between the two agent is at most $2X_t$. It follows from (6), applied with $a = \lceil d/2 \rceil$ and $t = \omega d^2 = \omega^3 n^2/k^2 \leq n^2/k = o(n^2 \ln n/k)$, that a.a.s. they meet at some round $T \leq t$. On the other hand, the Hoeffding-Azuma inequality (2), applied with $t = \omega d^2$ and $b = \sqrt{\omega t} = \omega d = \omega^2 n/k \leq n/(3\omega^{1/5})$, implies that a.a.s. both agents are at distance at most $n/(3\omega^{1/5})$ from their initial positions when they meet and so both of them are still at a distance of at least $n/(2\omega^{1/5})$ from both endpoints of the path. As before, since we aim for the statement that holds a.a.s., we may assume that this property holds.

We continue the process for an additional $t = \sqrt{\omega}(n \ln n/k)^2 \leq n^2/\omega^{3/2}$ rounds measured from round T when they met. Similarly as before, provided that they did not reach the end of the path, the process can be coupled with the lazy simple random walk on \mathbb{Z} starting at $X_0 = 0$ in such a way that the distance at round t between the two agent is at least $2X_t$. It follows from (6), applied with $a = \lceil n/b \rceil = \Theta(n \ln n/k)$ and $t = \sqrt{\omega}(n \ln n/k)^2$, that a.a.s. some agent reaches the end of the path or the two agents are at distance more than $2a$ at some point of that period of time. On the other hand, after applying the Hoeffding-Azuma inequality (2), with $t = \sqrt{\omega}(n \ln n/k)^2$ and $b = \sqrt{\omega t} = \omega^{3/4}(n \ln n/k) \leq n/\omega^{1/4} = o(n/\omega^{1/5})$, we get that a.a.s. neither agent reaches the end of the path during that period of time. We conclude that a.a.s. the two agents are at distance more than $2a$ at some point, that is, at some point they occupy two different blocks of the path that are separated by at least one block. A trivial but important property is that all agents occupying that internal block are green.

Suppose that at some point of the process, all agents occupying block i , $2 \leq i \leq b$, are green but some agent occupying block $i-1$ is white. Note that each block has length at most $\lceil n/b \rceil \leq (11/10)(n/b) \leq 600n \ln n/k$. Let $t = (9600n \ln n/k)^2$. Let us concentrate on any agent occupying block i . It follows from equation (5), applied with $a = \sqrt{t}/8 = 1200n \ln n/k$ and the above t , that with probability at least $1/2$ the agent either reaches the endpoint of the path (at some point during the following t rounds) or is at distance at least a from the original place after t rounds. Hence, by symmetry, with probability at least $1/4$, this agent either at some point reaches the endpoint of the path or occupies block $j \leq i-1$ after t rounds. In both scenarios all agents occupying block $i-1$ become green. By Lemma 3.8, there are at least $100 \ln n$ agents occupying block i . Hence the probability that no agent does the job is at most $(3/4)^{100 \ln n} = o(n^{-1})$. By symmetry, the same argument can be applied when all agents

occupying block i , $1 \leq i \leq b - 1$, are green but some agent occupying block $i + 1$ is white.

Since there are $b \leq n$ blocks, by the union bound, we get that a.a.s. after any period of t rounds, the number of blocks with all agents being green increases and so the process is done after at most

$$\begin{aligned} o(n^2 \ln n/k) + t \cdot b &= o(n^2 \ln n/k) + (9600n \ln n/k)^2 \cdot (k/(500 \ln n)) \\ &= O(n^2 \ln n/k) \end{aligned}$$

rounds and so the proof is finished. \square

Adjustment to cycles: As mentioned earlier, the same argument works for cycles and the proof is simpler as one does not need to pay attention to the two endpoints of the path. The same comment applies to the next theorems below.

Let us now move to a lower bound.

Theorem 3.10. *Let $\omega = \omega(n)$ be any function that tends to infinity as $n \rightarrow \infty$. For any k such that $\omega \ln n \leq k = k(n) \leq n/\ln^2 n$, a.a.s.*

$$\xi(P_n, k) = \Omega\left(\frac{n^2}{k(\ln k)(\ln n)}\right) = \Omega\left(\frac{n^2}{k \ln^2 n}\right).$$

Moreover, for any k such that $n/\ln^2 n < k = k(n) < 50n \ln n$, a.a.s.

$$\xi(P_n, k) > (1/2 + o(1))n.$$

Proof. Assume first that $\omega \ln n \leq k = k(n) \leq n/\ln^2 n$. Recall that there are $b = b(n) = \lfloor k/(500 \ln n) \rfloor$ blocks (because of our assumption that $b \gg 1$ and $b \ll n/\ln^3 n$), each of length $(500 + o(1))n \ln n/k$. We assign to each block a label from $[b]$; the first block contains vertex 1 and the last block contains vertex n . By symmetry, we may assume that the initial green agent starts at block $i_0 \geq b/2$. We will say that a green agent is *leading* if she occupies vertex j and no other green agent occupies vertex $\ell < j$. Note that leading agents may (and often do) change during the process, and there could be more than one leading agent at a given round. We will concentrate on leading agents and investigate times t_i when a leading agent leaves block $i + 1$ and enters block i for the first time. We will show that the following property holds a.a.s.: for all values of $i \in [i_0 - 2]$,

$$t_i - t_{i+1} \geq T = T(n) := \frac{n^2}{384 k^2 \ln k}. \quad (7)$$

Since $k = k(n) \leq n/(\sqrt{\ln n \omega})$, we get that $T \gg 1$ (this is the reason we had to introduce an upper bound for k ; in fact, we assumed that $k \leq n/\ln^2 n$ as for larger values of k we will be able to prove a stronger bound anyway). This will yield a lower bound for $\xi(P_n, k)$ as it proves that it takes at least $(i_0 - 1)T \geq (b/2 - 1)T = \Omega(n^2/(k \ln k \ln n))$ steps for the leading agent to reach the first block. Reaching the first block is needed as, by Lemma 3.8, a.a.s. there are agents in that block that by the definition of the leader are still white.

Let us fix any $i \in [i_0 - 2]$ and investigate the situation at time t_{i+1} when a leading agent enters block $i + 1$. It follows from Lemma 3.8 that there are at most $1800 \ln n$

agents occupying block $i+1$. Since each block has length $(500+o(1))n \ln n/k$, there must be a gap between two agents occupying that block that is of length at least $g := n/(4k)$. Trivially, between time t_{i+1} and t_i at least one agent has to move at least $g/2 = n/(8k)$ from her position at time t_{i+1} ; otherwise, no agent crosses the middle vertex of the gap and so no green agent enters block i . The probability that at least one agent crosses the gap during T rounds is, by Hoeffding-Azuma inequality (2) applied with $b = g/2$ and $t = T$, at most

$$k \cdot 2 \exp\left(-\frac{(g/2)^2}{2T}\right) = 2k \exp\left(-\frac{n^2/(64k^2)}{2n^2/(384k^2 \ln k)}\right) = 2/k^2.$$

Property (7) holds for a given i with probability $1 - o(1/k)$. Since $i_0 \leq b \leq k$, by the union bound, it holds a.a.s. for all $i \in [i_0 - 2]$. The desired lower bound holds for this range of k .

The argument for $n/\ln^2 n < k = k(n) < 50n \ln n$ is straightforward. As before, by symmetry we may assume that the initial green agent starts at block $i_0 \geq b/2$, that is, she is at distance at least $d = d(n) := n/2 - (500 + o(1))n \ln n/k \sim n/2$ from the first block. By Lemma 3.8, we may assume that the first block is always occupied by some agents. Trivially (and deterministically), it takes at least $d(n)$ steps for a leading agent to reach that block which yields the desired lower bound. The proof is finished. \square

In fact, in order to get a slightly stronger lower bound, one may use the fact that agents occupying blocks that are further away from the gap have to move more than $g/2$. Agents that are far can be dealt with easily. Hence, the union bound can be taken over $\Theta(\ln n)$ agents occupying close blocks instead of k agents. This would improve the bound by a multiplicative factor of $\Theta(\ln k / \ln \ln n)$. However, since such a lower bound does not match the upper bound we proved above, we stayed with an easier proof of a slightly weaker bound.

4. CLOSING THE GAP AND THE MEET-EXCHANGE PROCESS

Many of the bounds in this paper are proved in a “local” fashion by proving concentration bounds for the “moving parts”, whether these parts are trajectories of walks or numbers of agents in blocks. Once these bounds are established, one essentially needs to consider the worst case bound that holds with desired probability and treats the process as deterministic. This is a classic and natural approach. However, for this problem it will never be enough to establish tight bounds.

In order to obtain tight bounds, one needs to use a “global” approach. Such approach was successfully applied in [17] and then in [13] to analyze similar models. It is possible that it could be used again to get tight bounds for the broadcasting time for cycles and paths (possibly grids too). We leave it as an open problem.

Finally, let us mention the *Meet-Exchange* process that was recently introduced and studied in [12]. This process is closely related to our process but there are a few differences. In Meet-Exchange, agents are placed independently on the vertices of a graph according to the stationary distribution π instead of selecting starting points uniformly at random. A message is left on one of the vertices and needs to be picked up the agents before they start passing it to each other. Moreover, agents perform a

lazy random walk to avoid a problem of agents never meeting if the graph is bipartite (in the broadcasting time we avoid this issue by passing a message if agents go through the same edge but in the opposite directions—see below for more details). Despite these differences, it is quite possible that the bounds proved for the broadcasting time can also be proved for the Meet-Exchange process. Indeed, the stationary distribution is uniform on the cycle and almost uniform on the path. The message is picked by the agents quickly. In [22] it was shown that it takes $O((n/k)^2 \ln^2 k)$ rounds in expectation and so it is negligible unless k is very small. Addressing the fact that agents in the Meet-Exchange perform a lazy random walk should also be possible but disregarding passing a message by crossing agents seems to be the most challenging task. Having said that, since the considered graphs are strongly recurrent, once agents meet they typically do so a few times. With more work one should be able to overcome these technicalities. We also leave it as an open problem.

We would like to thank anonymous reviewers for pointing these papers and the “global” approach to us and for many other valuable comments that substantially improved the quality of this paper.

REFERENCES

- [1] H. Acan, A. Collevocchio, A. Mehrabian, and N. Wormald, On the push&pull protocol for rumour spreading, *SIAM J. Discrete Math.* 31 (2017), 647–668.
- [2] D. Aldous, J.A. Fill, *Reversible Markov Chains and Random Walks on Graphs*, 2002, unfinished monograph, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>
- [3] N. Alon, C. Avin, M. Koucký, G. Kozma, Z. Lotker, and M.R. Tuttle, Many random walks are faster than one, *Combinatorics, Probability and Computing* 20(4) (2011), 481–502.
- [4] I. Benjamini, L.R. Fontes, J. Hermon, and F.P. Machado, On an epidemic model on finite graphs. *Annals of Applied Probability* 30(1) (2020), 208–258.
- [5] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.
- [6] C. Cooper, A. Frieze, T. Radzik, Multiple random walks in random regular graphs, *SIAM J. Discrete Math.* **23(4)** (2009/10), 1738–1761.
- [7] R. Daknama, K. Panagiotou, S. Reisser, Asymptotics for Push on the Complete Graph, In 2020 Proceedings of the 14th Latin American Theoretical Informatics Symposium (LATIN 2020), Lecture Notes in Computer Science 12118 (2020), 451–463.
- [8] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In Proc. 6th Symp. Principles of Distributed Computing (PODC), pages 1–12, 1987.
- [9] B. Doerr and M. Künnemann. Tight analysis of randomized rumor spreading in complete graphs. In 2014 Proceedings of the Eleventh Workshop on Analytic Algorithmics and Combinatorics (ANALCO), pages 82–91. SIAM, 2014.
- [10] K. Efremenko and O. Reingold, How well do random walks parallelize? In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques* (pp. 476–489). Springer, Berlin, Heidelberg, 2009.
- [11] A.M. Frieze and G.R. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10(1):57–77, 1985.
- [12] G. Giakkoupis, F. Mallmann-Trenn, and H. Saribekyan. How to spread a rumor: Call your neighbors or take a walk? In P. Robinson and F. Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 24–33. ACM, 2019.

- [13] G. Giakkoupis, H. Saribekyan, and T. Sauerwald. Spread of Information and Diseases via Random Walks in Sparse Graphs. In DISC 2020 – 34rd International Symposium on Distributed Computing, pages 1–42, Freiburg, Germany, Oct. 2020.
- [14] R. Huq, B. Kamiński, A. Mashatan, P. Prałat, and P. Szufel, On Broadcasting Time in the Model of Travelling Agents, *Discrete Applied Mathematics* 291 (2021), 246–263.
- [15] S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*, Wiley, New York, 2000.
- [16] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized Rumor Spreading. In Proc. 41st Symp. Foundations of Computer Science (FOCS), pages 565–574, 2000.
- [17] H. Kesten and V. Sidoravicius. The spread of a rumor or infection in a moving population. *Ann. Probab.*, 33(6):2402–2462, 2005.
- [18] D.A. Levin, Y. Peres, *Markov Chains and Mixing Times* (Second Edition), AMS, 2017, 447 pp.
- [19] T.M. Liggett. *Interacting Particle Systems*, volume 276 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, 1985.
- [20] R. Oliveira, On the coalescence time of reversible random walks, *Transactions of the American Mathematical Society* 364, no. 4 (2012): 2109–2128.
- [21] R. Oliveira, Mean field conditions for coalescing random walks, *The Annals of Probability* 41, no. 5 (2013): 3420–3461.
- [22] N. Rivera, T. Sauerwald, and J. Sylvester, Multiple Random Walks on Graphs: Mixing Few to Cover Many. 48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, LIPIcs 198 (2021), 107:1–107:16
- [23] Serguei Yu. Popov, Frogs and some other interacting random walks models. *Discrete Random Walks, DRW'03*, 2003, Paris, France. pp. 277–288.
- [24] B. Pittel. On Spreading a Rumor. *SIAM J. Appl. Math.*, 47(1):213–223, 1987.

RYERSON UNIVERSITY, TORONTO, CANADA

Email address: `reaz.huq@ryerson.ca`, `pralat@ryerson.ca`