A Simple Network of Nodes Moving on the Circle

D. Cheliotis* I. Kontoyiannis† M. Loulakis‡ S. Toumpis§

August 2, 2018

Abstract

Two simple Markov processes are examined, one in discrete and one in continuous time, arising from idealized versions of a transmission protocol for mobile, delay-tolerant networks. We consider two independent walkers moving with constant speed on either the discrete or continuous circle, and changing directions at independent geometric (respectively, exponential) times. One of the walkers carries a message that wishes to travel as far and as fast as possible in the clockwise direction. The message stays with its current carrier unless the two walkers meet, the carrier is moving counter-clockwise, and the other walker is moving clockwise. In that case, the message jumps to the other walker. The long-term average clockwise speed of the message is computed. An explicit expression is derived via the solution of an associated boundary value problem in terms of the generator of the underlying Markov process. The average transmission cost is also similarly computed, measured as the long-term number of jumps the message makes per unit time. The tradeoff between speed and cost is examined, as a function of the underlying problem parameters.

Keywords: Markov process, generator, ergodicity, stochastic networks, network routing

*Department of Mathematics, Panepistimiopolis, Athens 15784, Greece. Email: dcheliotis@math.uoa.gr.

†Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK; and Department of Informatics, Athens University of Economics and Business, Patission 76, Athens 10434, Greece. Email: i.kontoyiannis@eng.cam.ac.uk. I.K. was supported in part by RC-AUEB and by EU and Greek national funds through the Operational program ELF of the NSRF through the program Thales-Investing in Knowledge Society.

‡School of Applied Mathematical and Physical Sciences, National Technical University of Athens, Polytechnioulopis Zographou, Athens 15780, Greece; and Institute of Applied and Computational Mathematics, Foundation for Research and Technology - Hellas, Greece. Email: loulakis@math.ntua.gr.

§Department of Informatics, Athens University of Economics and Business, Patission 76, Athens 10434, Greece. Email: toumpis@aueb.gr.

[0]Preliminary versions of some of the present results appeared in the conference papers [9, 10].

1 Introduction

Consider a network that consists of many mobile nodes moving around randomly in some large area. Suppose that each node moves with constant speed, changing its direction of travel at random times, and that one of the nodes carries a message that she wants to transmit to a far away destination in some specific, fixed direction. The message stays with its current carrier until the first time she comes within a certain distance from some other node moving in a better direction, i.e., in a direction closer to that of the intended recipient. In that case, she transmits her message to the other node, and the new carrier then proceeds in the same fashion. What is the long-term average speed with which the message travels towards its destination, as a function of, say, the nodes' individual speeds and their density? How often, on the average, does the message get transmitted form one node to another?

Networks of this type, where messages propagate via a combination of physical transport (moving with their carrier) and wireless transmissions (being sent from one node to another) belong to the wide class of delay-tolerant networks (DTNs) [29]. Examples of DTNs arising in applications include space [2], vehicular [4], sensor [26], and pocket-switched networks [17]. In earlier work by some of the authors [8, 7, 19], the questions of the previous paragraph were considered under very general assumptions on the movement of the nodes and on the protocol under which the message gets transmitted between nodes. In that line of work, as in much of the related earlier work in this area, e.g., [16, 12, 18, 27], the complexity of the models involved prohibits the derivation of exact, explicit answers. For that reason, typical results are in the form of asymptotics, approximations, performance bounds, or estimates based on simulation experiments.

In this work we examine two variants of a simple, idealized model, where it is possible to derive explicit, closed-form expressions for the performance metrics of interest. We first consider a collection of $m \geq 2$ nodes moving independently on a discrete circle consisting of $N \geq 3$ locations, in discrete time. Each node maintains their current direction of travel for a geometrically distributed amount of time with parameter $\epsilon \in (0, 1)$, and one of the nodes carries a message intended to travel as far as possible in the clockwise direction. The message stays with its current carrier unless, while moving counter-clockwise, it finds itself in the same location as a different node moving clockwise. In that case the message gets transmitted to the other node, and the same process is repeated.

For the case of $m = 2$ nodes, in Section 3.1, Theorem 3.1, we show that the long-term average clockwise speed $s = s(N, \epsilon)$ of the message is $s = \frac{1-\epsilon}{2(1+\epsilon(N-2))}$. The proof, given in Section 4, involves the construction of a martingale that solves an associated (discrete) boundary value problem. Similar techniques allow us to compute the average transmission cost $c = c(N, \epsilon)$, measured as the long-term average number of message transmissions per unit time. In Theorem 3.3 we show that $c(N, \epsilon) = \epsilon \times s(N, \epsilon)$, for all N and ϵ. Therefore, the message travels a clockwise distance of $1/\epsilon$ units between successive jumps (on the average), regardless of N. The tradeoff between speed and cost for different values of the parameters N and ϵ is also discussed in Section 3.1.

Section 3.2 contains continuous-time analogs of Theorems 3.1 and 3.3. Here we consider $m \geq 2$ nodes moving with constant speed $v > 0$ on a continuous circle of circumference $N > 0$, changing directions at independent exponential times with rate $r > 0$. The corresponding expressions for the long-term average speed $s(N, v, r)$ and cost $c(N, v, r)$ are established in

Theorems 3.4 and 3.6, respectively. Again, it turns out that the speed and cost satisfy a simple scale-free relationship: $c(N, v, r) = (r/v) \times s(N, v, r)$. In other words, the message travels a (clockwise) distance of v/r units between successive jumps, on the average.

The proofs of Theorems 3.4 and 3.6, given in Section 4, involve shorter and somewhat cleaner arguments than their discrete-time counterparts. In the continuous-time case, it is more straightforward to construct appropriate solutions to the relevant boundary value problems, which are stated in terms of the infinitesimal generator of the underlying Markov process. What is somewhat cumbersome, is the proof that this Markov process is exponentially ergodic, uniformly in its initial state. The relevant ergodic properties are stated in Proposition 2.1 and Theorem 2.2, both proved in the Appendix.

Although perhaps the most restrictive of our assumptions is that nodes are assumed to move along the circumference of a circle, we note that there has been much recent interest in one-dimensional models of DTNs, particularly in connection with the important class of vehicular networks (VANETS); see [4, 5, 30] and the references therein. Finally, a somewhat less closely related but quite extensively studied problem, in terms of a Markov chain describing the movement of a finite collection of nodes on a circle, is the *k-server* problem introduced in [21]; see, e.g., [11] or [6] for more recent developments.

2 Models and Problem Statement

2.1 Random walk on the discrete circle

Let $S := \{0, 1, \ldots, N-1\} = \mathbb{Z}/N\mathbb{Z}$ denote the discrete N-circle, for a fixed odd $N \geq 3$. We place $m \geq 2$ independent random walkers on S, located at $X_t = (X_t(1), X_t(2), \ldots, X_t(m))$ at time $t = 0, 1, 2, \ldots$, and with each walker j we associate a random direction $D_t(j)$ at time t, where $D_t(j)$ is either $= +1$ (clockwise motion) or $D_t(j) = -1$ (counter-clockwise motion). The initial positions X_0 and directions D_0 are arbitrary. The Markov chain $\{(X_t, D_t)\ ;\ t \geq 0\}$ evolves on the state space $S^m \times \{-1, +1\}^m$ as follows.

Let $\{Z_t = (Z_t(1), Z_2(2), \ldots, Z_t(m))\}$ be a sequence of independent Bernoulli random variables with parameter $\epsilon \in (0, 1)$. Given the current state (X_t, D_t), each walker j takes a step in the direction given by $D_t(j)$,

$$X_{t+1}(j) = X_t(j) + D_t(j) \pmod N, \quad t \geq 0,\ j = 1, 2, \ldots, m,$$

and then decides to either continue moving in the same direction with probability $(1 - \epsilon)$, or to switch to the opposite direction, with probability ϵ:

$$D_{t+1}(j) = (1 - Z_t(j))D_t(j) - Z_t(j)D_t(j), \quad t \geq 0,\ j = 1, 2, \ldots, m.$$

We also define an index process $\{I_t\}$ evolving on $\{1, 2, \ldots, m\}$, with I_0 chosen arbitrarily and I_t trying to track walkers that move clockwise: Given $(X_t, D_t, I_t = i)$, let (X_{t+1}, D_{t+1}) be defined as above. If $D_{t+1}(i) = -1$ and there is at least one more walker, j, say, at the same location, $X_{t+1}(i) = X_{t+1}(j)$, but its direction $D_{t+1}(j) = +1$, then $I_{t+1} = j$ (or a uniformly chosen such j if there are multiple candidates). In all other cases, $I_{t+1} = I_t = i$.

It is easy to see from the above construction that $\mathbf{\Phi} = \{\Phi_t = (X_t, D_t, I_t)\ ;\ t \geq 0\}$ is an irreducible and aperiodic chain on the state space Σ consisting of all configurations of the

form,

$$(x(1), x(2), \ldots, x(m), d(1), d(2), \ldots, d(m), i) \in S^m \times \{+1, -1\}^m \times \{1, 2, \ldots, m\},$$

except those where $d(i) = -1$ and there is a $j \neq i$ such that $x(j) = x(i)$ and $d(j) = +1$. Moreover, under the unique invariant distribution π of $\mathbf{\Phi}$, the distribution of (X_t, D_t) is uniform: The positions $X_t(i)$ are independent of each other and uniformly distributed on S, and the directions $D_t(i)$ are independent of the positions X_t and each $D_0(i) = \pm 1$ with probability $1/2$, independently of the others.

We are primarily interested in the following three quantities, as functions of N, m and ϵ: (i) *Direction*: What is the limiting distribution of the direction $D_t(I_t)$ of the message at time t? (ii) *Speed*: What is the long-term average speed of the message? (iii) *Cost*: What is the long-term average number of jumps per unit time? We are also interested in the relationship between the speed and cost: Do higher speeds always imply an increase in cost? Or is there a range of parameter values that improve the speed and cost simultaneously?

2.2 Continuous motion on the circle

Let $S := \mathbb{R}/N\mathbb{Z}$ denote the one-dimensional circle of circumference $N > 0$, where N is not necessarily an integer. We place $m \geq 2$ independent random walkers $X_t = (X_t(1), \ldots, X_t(m))$ on S, and with each walker j we associate a random direction $D_t(j)$ at time t, where $D_t(j)$ is either $= +1$ (clockwise motion) or $D_t(j) = -1$ (counter-clockwise motion). The initial positions X_0 and directions D_0 are arbitrary. The continuous-time Markov process $\mathbf{W} = \{(X_t, D_t) \; ; \; t \geq 0\}$ evolves on the state space $S^m \times \{-1, +1\}^m$ as follows. The jth walker continues moving at constant speed v in its present direction, $D_t(j) = d$, say, for an exponentially distributed amount of time with mean $1/r$, for some $r > 0$; during that time its direction remains constant, and afterwards it switches to $-d$. The process continues in the same fashion, by choosing a new, independent exponential time for the jth walker, and with the different walkers moving independently of one another.

We assume that the transitions between directions are such that the sample paths of the process $\mathbf{W} = \{(X_t, D_t) \; ; \; t \geq 0\}$ are right continuous, and observe that $\mathbf{W}$ is strong Markov and, therefore, a Borel right process [28, 13]. And since $S^m \times \{-1, +1\}^m$ is compact, $\mathbf{W}$ is also non-explosive [24]. The following simple proposition is proved in the Appendix.

Proposition 2.1. (i) *The Markov process $\mathbf{W} = \{(X_t, D_t) \; ; \; t \geq 0\}$ is ψ-irreducible and aperiodic on $S^m \times \{-1, +1\}^m$, with respect to $\psi := \mathcal{L}^m \times \kappa^m$, where $\mathcal{L}$ denotes the Lebesgue measure on S and κ the counting measure on $\{+1, -1\}$.*

(ii) *The process $\mathbf{W}$ is positive Harris recurrent.*

(iii) *The uniform distribution is the unique invariant probability measure of $\mathbf{W}$.*

We also define an index process $\{I_t\}$ evolving on $\{1, 2, \ldots, m\}$, with I_0 chosen arbitrarily and I_t trying to track walkers that move clockwise. Specifically, I_t stays constant most of the time, and its value only changes when $X_t(I_t) = X_t(j)$ for some $j \neq I_t$, and the direction $D_t(j)$ of the jth walker at the time is $+1$ while $D_t(I_t) = -1$. In that case, the value of I_t switches to j (or to a uniformly chosen such j if there are multiple candidates) and remains there at least until the first time walker j encounters a different walker.

Next we show that the Markov process $\mathbf{\Phi} = \{\Phi_t = (X_t, D_t, I_t) \; ; \; t \geq 0\}$ is uniformly ergodic on the state space Σ, which consists of all elements $\phi \in S^m \times \{-1, +1\}^m \times \{1, 2, \ldots, m\}$,

$$\phi = (x(1), x(2), \ldots, x(m), d(1), d(2), \ldots, d(m), i),$$

where we identify pairs of states $s = (x, d, i)$ and $s' = (x', d', i')$ of the following form: The message is with a different walker in each state, i.e., $i \neq i'$, all positions and directions are identical, $x = x'$ and $d = d'$, the ith and i'th walkers are in the same position $x(i) = x(i')$, and the two walkers move in opposite directions, i.e., $d(i) = -d(i')$.

As with $\mathbf{W}$, we assume that the transitions between directions and between successive values of the process $\{I_t\}$ are such that the sample paths of $\mathbf{\Phi}$ are right-continuous, so that $\mathbf{\Phi}$ is a non-explosive, Borel right process [28, 13]. Its ergodicity properties are summarized in Theorem 2.2, proved in the Appendix. Here, and throughout the paper, for an arbitrary measure μ and function g we write $\mu(g)$ for $\int g d\mu$, whenever the integral exists.

Theorem 2.2. *(i)* $\mathbf{\Phi} = \{\Phi_t = (X_t, D_t, I_t) \; ; \; t \geq 0\}$ *is ψ-irreducible and aperiodic with respect to $\psi := \mathcal{L}^m \times \kappa^m \times \kappa_m$, where, as before, $\mathcal{L}$ and κ denote the Lebesgue and counting measures on S and $\{+1, -1\}$, respectively, and κ_m denotes the counting measure on $\{1, 2, \ldots, m\}$.*

(ii) $\mathbf{\Phi}$ *is uniformly ergodic, with a unique invariant probability measure π.*

(iii) $\mathbf{\Phi}$ *converges to equilibrium uniformly exponentially fast: There are constants $C < \infty, \rho > 0$ such that,*

$$|P_\phi(\Phi_t \in A) - \pi(A)| = |\Pr(\Phi_t \in A | \Phi_0 = \phi) - \pi(A)| \leq Ce^{-\rho t},$$

for all $\phi \in \Sigma$, all measurable $A \subset \Sigma$, and all $t > 0$.

(iv) The following ergodic theorem holds for $\mathbf{\Phi}$: For any bounded (measurable) function $f : \Sigma \to \mathbb{R}$ and any initial state $\Phi_0 = \phi \in \Sigma$,

$$\lim_{t \to \infty} \frac{1}{t} \int_0^t f(\Phi_s) ds = \pi(f), \quad a.s.$$

Finally we note that the dynamics of $\mathbf{\Phi}$ can be described by its infinitesimal generator L. Let A_L denote the collection of all functions continuous functions $f : \Sigma \to \mathbb{R}$, such that f is continuously differentiable in $x(j)$ for each $1 \leq j \leq m$. This is a dense subset of $C(\Sigma)$, and the infinitesimal generator L of $\mathbf{\Phi}$ acts on each $f \in A_L$ as,

$$Lf(x, d, i) = \sum_{j=1}^{m} \left\{ vd(j) \frac{\partial f}{\partial x(j)}(x, d, i) + r \left[f(x, \sigma^j d, i) - f(x, d, i) \right] \right\}, \tag{1}$$

where, for any m-tuple of directions $d \in \{-1, +1\}^m$, $\sigma^j d$ is the same as d but with its jth coordinate having the opposite sign from that of d, $1 \leq j \leq m$. The first term in the sum on the right-hand side above corresponds to the motion of the jth walker at constant velocity $vd(j)$, while the second one corresponds to its change of direction at rate r. It is easy to see that L

defined on A_L is a closed operator and there is $\lambda > 0$ small enough so that $(I-\lambda L)(A_L) = C(S)$, cf. [14, Proposition 1.3.5], so that A_L is the domain of L.

Once again, we are interested in the following three quantities, as functions of m, N, v and r: (i) The limiting distribution of the direction $D_t(I_t)$ of the message at time t; (ii) The long-term average speed of the message; (iii) The long-term average number of jumps per unit time. Also, we wish to examine the nature of the tradeoff between the speed and cost.

3 Results: Speed and Cost with *m=2* walkers

Here we state and discuss our main results for both the discrete and the continuous case. The proofs are given in Section 4. We adopt the following standard notation: For the probabilities of events depending on an underlying Markov process $\{\Phi_t\}$ we write P_ϕ for the measure describing the distribution of the process conditional on $\{\Phi_0 = \phi\}$, and P_μ when $\Phi_0 \sim \mu$ for some probability measure μ. Similarly, E_ϕ and E_μ denote the corresponding expectation operators.

3.1 The discrete circle

Consider the problem of $m = 2$ walkers on the N-circle, changing directions with rate ϵ, as described in Section 2.1.

Theorem 3.1 (**Message speed**). *In the case of $m = 2$ walkers, for any initial state, the long-term average speed of the message is:*

$$s := \pi(D_1(I_1)) = \lim_{n\to\infty} \frac{1}{n} \sum_{t=0}^{n-1} D_t(I_t) = \frac{1-\epsilon}{2(1+\epsilon(N-2))}, \quad a.s.$$

Note that the speed $s = s(N, \epsilon)$ is always less than or equal to $1/2$, and it is decreasing in both N and ϵ; see Figure 1.

In the boundary case $\epsilon = 0$, the speed $s(N, \epsilon)$ is either -1 or 1, depending on the initial directions of the two walkers. Therefore, $s(N, \epsilon)$ is discontinuous at $\epsilon = 0$, since $s(N, \epsilon) \uparrow 1/2$ as $\epsilon \downarrow 0$, for any N. Figure 2 shows the results of two simulation experiments, illustrating the convergence of the speed of the message to the corresponding value $s(N, \epsilon)$ computed in Theorem 3.1.

Theorem 3.1 answers question (ii) of Section 2.1. The answer to question (i) is a simple consequence of the theorem, given in Corollary 3.2 below.

Corollary 3.2 (**Message direction**). *In the case of $m = 2$ walkers, for any initial state $\Phi_0 = \phi \in \Sigma$, the steady state probability that the message moves in the clockwise direction is:*

$$P_\pi(D_1(I_1) = +1) = \lim_{t\to\infty} P_\phi(D_t(I_t) = +1) = \frac{s+1}{2} = \frac{3+\epsilon(2N-5)}{4(1+\epsilon(N-2))}.$$

Next we examine the asymptotic cost of message transmissions. Theorem 3.3 describes the long-term average number of jumps of the message, $c = c(N, \epsilon)$ per unit time.

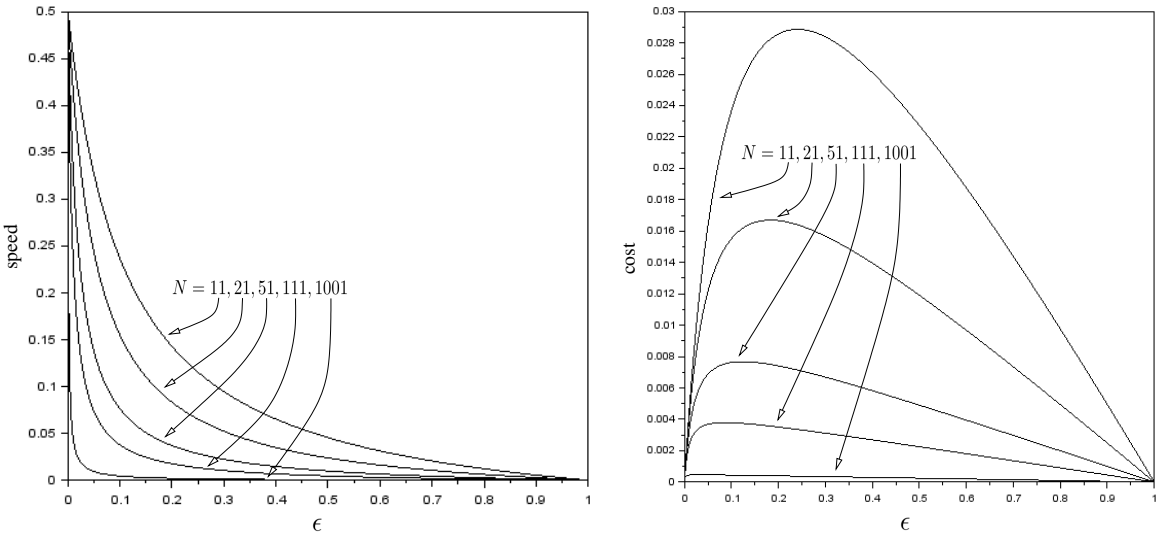

Figure 1: Plots of the asymptotic speed $s = s(N, \epsilon)$ of the message (left) and of the asymptotic cost $c = c(N, \epsilon)$ (right), as functions of ϵ, for different values of N.

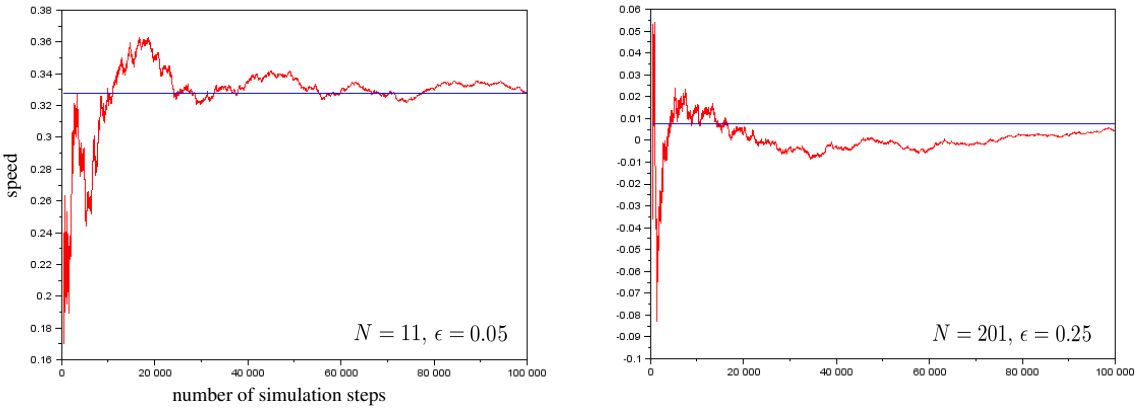

Figure 2: Simulation results for the speed of the message during $T = 10^5$ steps, in two independent realizations of the chain with different parameter values. In each case, the horizontal line is the limiting value of the speed s predicted by Theorem 3.1.

Theorem 3.3 (**Transmission cost**). *In the case of $m = 2$ walkers, for any initial state, the long-term average cost of message transmissions is:*

$$c := P_\pi(I_2 \neq I_1) = \lim_{n\to\infty} \frac{1}{n} \sum_{t=0}^{n-1} \mathbb{I}_{\{I_{t+1} \neq I_t\}} = \frac{\epsilon(1-\epsilon)}{2[1+\epsilon(N-2)]}, \quad a.s.$$

We observe that the cost $c = c(N, \epsilon)$ is decreasing in N, and for each fixed N it is a concave function of ϵ; see Figure 1. Also, unlike the speed $s = s(N, \epsilon)$, the cost $c = c(N, \epsilon)$ is continuous and equal to zero at $\epsilon = 0$.

Speed vs. cost. It is interesting to observe the following simple, scale-free relationship between the asymptotic speed and cost: $c(N,\epsilon) = \epsilon \times s(N,\epsilon)$, for all N and ϵ. Therefore, on the average, the message travels a (clockwise) distance of $1/\epsilon$ units between successive jumps, regardless of the value of N.

In terms of the speed/cost tradeoff, note that for each N there is an ϵ^* below which the speed increases *and* the cost decreases as $\epsilon \downarrow 0$. This suggests that, if such a protocol were to be implemented in practice, it is the relatively smaller values of ϵ that would be most effective in the long run.

3.2 The continuous circle

Now we turn to the problem of $m = 2$ walkers on the continuous circle of circumference N, moving with constant speed v and changing directions at rate r. Theorem 3.4 gives the natural continuous analog of the discrete-time result in Theorem 3.1.

Theorem 3.4 (Message speed). *In the case of $m = 2$ walkers, for any initial state, the long-term average speed of the message is:*

$$s := v\pi(D_1(I_1)) = \lim_{t\to\infty} \frac{1}{t} \int_0^t v D_s(I_s) ds = \frac{v^2}{2v + rN}, \quad a.s.$$

Note that the speed $s = s(N, v, r)$ is always no greater than $v/2$, as in the discrete case. Also observe that, as would be expected, $s = s(N, v, r)$ is decreasing in the circumference length N and increasing in the walker speed v. Moreover, s is also decreasing in the reversal rate r.

Theorem 3.4 answers question (ii) of Section 2.2. The answer to question (i), given below, is an immediate consequence of Theorem 3.4.

Corollary 3.5 (Message direction). *In the case of $m = 2$ walkers, for any initial state $\Phi_0 = \phi \in \Sigma$, the steady state probability that the message moves in the clockwise direction is:*

$$P_\pi(D_1(I_1) = +1) = \lim_{t\to\infty} P_\phi(D_t(I_t) = +1) = \frac{s + v}{2v} = \frac{3v + rN}{2(2v + rN)}.$$

In our final result we determine the long-term average number of jumps $c = c(N, v, r)$ per unit time.

Theorem 3.6 (Transmission cost). *In the case of $m = 2$ walkers, for each time $t > 0$ let M_t denote the (random) number of times the message jumps from one walker to the other up to time t. Then, for any initial state, the long-term average cost of message transmissions is:*

$$c := \lim_{t\to\infty} \frac{M_t}{t} = \frac{rv}{2v + rN}, \quad a.s.$$

Observe that the cost $c = c(N, v, r)$ is naturally increasing in v and decreasing in N. But, unlike in the discrete case, $c = c(N, v, r)$ is monotonically increasing in r. Again we also observe that there is a simple, scale-free relationship between the asymptotic speed and cost, $c(N, v, r) = (r/v) \times s(N, v, r)$: In the long-run, the message travels a (clockwise) distance of v/r units between successive jumps.

A comparison of the results of Theorems 3.4 and 3.6 with their discrete-time analogs is perhaps informative. Consider a large discrete circle of size $N \gg 1$ and a small rate of direction updates $\epsilon \approx 0$. Then, noting that the circumference N_c of the continuous circle corresponds to $2N$ in the discrete case (since, because N is odd, that is the number of steps required for two walkers starting in the same location and moving in opposite directions to meet again), we have the following scaling limit. Taking the speed $v = 1$ in the continuous case, and the rate ϵ in the discrete case to be such that $2N\epsilon = N_c r$, passing to the continuous limit $N \to \infty$ we obtain,

$$s(N, \epsilon) \to \frac{1}{2 + rN_c} = s(N_c, v, r) \quad \text{and} \quad \frac{r}{\epsilon} c(N, \epsilon) \to \frac{r}{2 + rN_c} = c(N_c, v, r).$$

Finally we note that the parameters of the problem define the dimensionless quantity $\alpha = rN/v > 0$, and dimensional analysis alone (meaning, speeding up time by a constant factor, or dilating space by a constant factor) shows that $s = vf(\alpha)$ and $c = rg(\alpha)$, for suitable functions f, g. In this light, our results can be interpreted as showing that $f(\alpha) = g(\alpha) = 1/(2 + \alpha)$, for all $\alpha > 0$.

4 Proofs

4.1 The discrete circle

Proof of Theorem 3.1. First, consider the reduced chain,

$$\mathbf{\Psi} = \{\Psi_t = (Y_t = X_t(1) - X_t(2), D_t(1), D_t(2), I_t) \; ; \; t \geq 0\},$$

where the differences $Y_t = X_t(1) - X_t(2)$ are taken modulo N. Clearly $\mathbf{\Psi}$ is irreducible and aperiodic on the corresponding reduced state space Σ_ψ consisting of all configurations, of the form,

$$(y, d, d', i) \in S \times \{+1, -1\}^2 \times \{1, 2\},$$

except $(0, +1, -1, 2)$ and $(0, -1, +1, 1)$. Let π_ψ denote the unique invariant measure of $\mathbf{\Psi}$. The limit in the theorem exists a.s. by ergodicity; in order to compute its actual value, we define the following regeneration time,

$$T = \inf\{t \geq 1 \; ; \; Y_t = 0 \text{ and } D_t(1) \neq D_t(2)\}, \tag{2}$$

and we consider two special states of $\mathbf{\Psi}$: $\psi_1 = (0, +1, -1, 1)$ and $\psi_2 = (0, -1, +1, 2)$. Let ν denote the probability measure on Σ_ψ given by,

$$\nu = \frac{1}{2}\delta_{\psi_1} + \frac{1}{2}\delta_{\psi_2}. \tag{3}$$

Then T is indeed a regeneration time for ν in the sense that, with $\Psi_0 \sim \nu$, we also have $\Psi_T \sim \nu$. We will use the following general version of Kac's formula; cf. [1, Corollary 2.24].

Lemma 4.1. *For any function $f : \Sigma_\psi \to \mathbb{R}$ and any regeneration time T for ν:*

$$E_\nu \left[\sum_{t=0}^{T-1} f(\Psi_t) \right] = E_\nu(T) \pi_\psi(f).$$

To apply Lemma 4.1, we first compute $E_\nu(T)$:

Lemma 4.2. $E_\nu(T) = 2N$.

PROOF. Consider the (further restricted) chain $\boldsymbol{\Upsilon} = \{\Upsilon_t = (Y_t, D_t(1), D_t(2)) \; ; \; t \geq 0\}$ on the state space $S \times \{+1, -1\}^2$, and note that its unique invariant measure ρ is uniform. Write, $D = \{(0, +1, -1), (0, -1, +1)\}$, let ρ_D denote the measure ρ conditioned on D, and let,

$$T_D^+ = \inf\{t \geq 1 \; ; \; \Upsilon_t \in D\},$$

so that, in fact, $T_D^+ = T$. Then, by Kac's formula [1], we have,

$$E_\nu(T) = E_{\rho_D}(T_D^+) = \frac{1}{\rho(D)} = \frac{4N}{|D|} = 2N,$$

as claimed. □

The central step in the proof of the theorem is an application of Lemma 4.1 with $f(\Psi_t) = D_t(I_t)$, which, combined with Lemma 4.2 gives us that $s = \pi(D_1(I_1)) = \pi_\psi(D_1(I_1))$ can be expressed as,

$$s = \frac{1}{2N} E_\nu \left[\sum_{t=0}^{T-1} D_t(I_t) \right] = \frac{1}{4N} E_{\psi_1} \left[\sum_{t=0}^{T-1} D_t(1) \right] + \frac{1}{4N} E_{\psi_2} \left[\sum_{t=0}^{T-1} D_t(2) \right] = \frac{1}{2N} E_{\psi_1} \left[\sum_{t=0}^{T-1} D_t(1) \right],$$

where the sums above (and in what follows) correspond to addition over $\mathbb{Z}$ (as opposed to modulo N addition over S), and where the second equality follows from the fact that, by the definition of T, the message is with walker 1 up to time $T-1$. Therefore, writing $X_{t+1}^*(j) = X_t^*(j) + D_t(j)$, for $j = 1, 2$, $t \geq 1$, we have,

$$\begin{aligned} s &= \frac{1}{2N} E_{\psi_1}(X_T^*(1)) \\ &= \frac{1}{2N} E_{\psi_1} \left(\frac{X_T^*(1) - X_T^*(2)}{2} \right) + \frac{1}{2N} E_{\psi_1} \left(\frac{X_T^*(1) + X_T^*(2)}{2} \right) \\ &= \frac{1}{2N} E_{\psi_1} \left(\frac{X_T^*(1) - X_T^*(2)}{2} \right), \end{aligned}$$

where we noted that $E_{\psi_1}(X_T^*(1) + X_T^*(2))$ is zero by symmetry, since the two walkers start off in opposite directions.

Now write, $A = \{(0, +1, +1), (0, +1, -1), (0, -1, +1), (0, -1, -1)\}$, and let T_A^+ denote the first time when the two walkers meet,

$$T_A^+ = \inf\{t \geq 1 \; ; \; \Upsilon_t \in A\} = \inf\{t \geq 1 \; ; \; Y_t = 0\},$$

so that T_A^+ can be expressed in terms of either $\boldsymbol{\Psi}$ or $\boldsymbol{\Upsilon}$. We observe that, at time T_A^+, either the two walkers decide to go in opposite directions, in which case $T_A^+ = T$, or they continue moving together until they choose opposite directions, in which case the difference of their locations $X_t^*(i)$ stays constant; therefore,

$$s = \frac{1}{2N} E_{\psi_1} \left(\frac{X_{T_A^+}^*(1) - X_{T_A^+}^*(2)}{2} \right).$$

Since the last expectation above is conditioned on the two walkers starting from the same position, in opposite directions, and with the first one moving in the positive (clockwise) direction, there are exactly two possible scenarios for their first meeting time T_A^+: In the first scenario, at time $t = T_A^+ - 1$ walker 1 is two steps "ahead" in the clockwise direction of walker 2 (as they are, e.g., at time $t = 1$). In this case, we will necessarily have $X^*_{T_A^+}(1) - X^*_{T_A^+}(2) = 0$. We call this event C. In the second scenario, the relative positions of the two walkers at time $t = T_A^+ - 1$ will be reversed, which necessarily means that the first walker travelled a whole circle "around" the second one before they met, so that (since N is odd) on C^c, we must have $X^*_{T_A^+}(1) - X^*_{T_A^+}(2) = 2N$. Therefore,

$$s = \frac{1}{2N}\left[\frac{0}{2}\cdot P_{\psi_1}(C) + \frac{2N}{2}\cdot P_{\psi_1}(C^c)\right] = \frac{1}{2}P_{\psi_1}(C^c). \tag{4}$$

Finally we compute the probability of the event C:

Lemma 4.3. $P_{\psi_1}(C^c) = \frac{1-\epsilon}{1+\epsilon(N-2)}$.

Proof. Here we consider the chain $\boldsymbol{\Upsilon}^* = \{\Upsilon_t^* = (Y_t^*, D_t(1), D_t(2)) \ ; \ t \geq 0\}$ on $\Sigma^* = \mathbb{Z} \times \{+1, -1\}^2$, where $Y_t^* = X_t^*(1) - X_t^*(2)$. Note that, for the state $u_1 := (0, +1, -1)$, the initial condition $\Upsilon_0^* = u_1$ corresponds to $\Psi_0 = \psi_1$.

We will only need to examine the evolution of $\boldsymbol{\Upsilon}^*$ until time $t = T_A^+$, which, since N is odd, can equivalently be expressed as,

$$T_A^+ = \min\{t \geq 1 \ ; \ Y_t^* = 0 \pmod{2N}\},$$

and the same argument as in the last paragraph before the statement of the lemma shows that, given $\Upsilon_0^* = u_1$, the only two possible values of $Y^*_{T_A^+}$ are 0 and $2N$, on C and on C^c, respectively. Therefore, letting,

$$\begin{aligned} T_R &= \min\{t \geq 1 \ ; \ Y_t^* = 0\}, \\ \text{and } T_L &= \min\{t \geq 1 \ ; \ Y_t^* = 2N\}, \end{aligned}$$

we have that $T_A^+ = \min\{T_L, T_R\}$ and $P_{\psi_1}(C^c) = P_{\psi_1}(T_L < T_R)$; cf. Figure 3.

In fact, for this computation it will suffice to consider the *trace* of $\boldsymbol{\Upsilon}^*$ on the set,

$$\Sigma^t := \{0, 2, 4, \ldots, 2N\} \times \{(+1, -1), (-1, +1)\} \subset \Sigma^*;$$

cf. [25, Example 1.4.4.]. The evolution of this Markov chain is fairly simple and its transition probabilities are easy to compute; e.g., the probability of the transition from $(0, +1, -1)$ to $(2, +1, -1)$ is equal to,

$$(1-\epsilon)^2 + \epsilon(1-\epsilon)\frac{1}{2} + \epsilon(1-\epsilon)\frac{1}{2} = 1 - \epsilon.$$

The first term above corresponds to the case when the two walkers both maintain their original directions after their first step; the second term corresponds to the case when only the first walker changes direction, after which they keep moving at a distance two apart, until one of them changes direction again and they either reach the state $(2, +1, -1)$ or the state

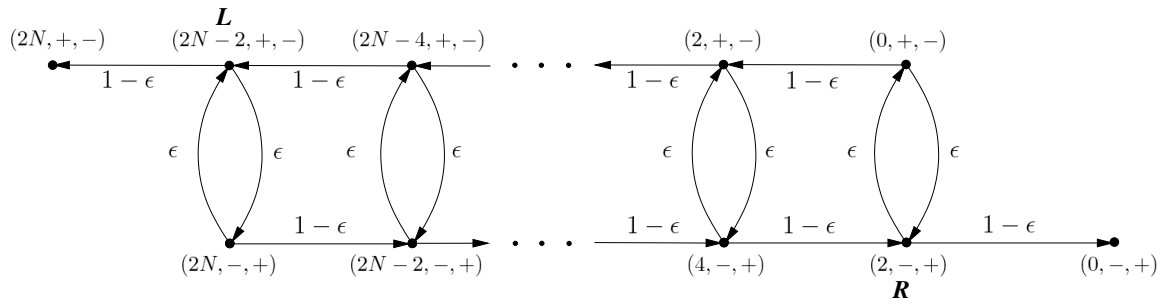

Figure 3: Evolution of the trace of the chain $\boldsymbol{\Upsilon}^*$ on $\{0, 2, 4, \ldots, 2N\} \times \{(+1, -1), (-1, +1)\}$.

$(2, -1, +1)$, each having probability $1/2$ by symmetry; and the third term corresponds to the case when only the second walker changes direction after their first step, and its value is the same as the second term again by symmetry. The remaining transition probabilities can be similarly computed; see Figure 3.

Finally, for every state $u \in \{0, 2, 4, \ldots, 2N\} \times \{(+1, -1), (-1, +1)\}$ we define $h(u) = P_u(T_L < T_R)$, so that $h(u_1) = P_{\psi_1}(C^c)$. Writing L and R for the states $(2N - 2, +1, -1)$ and $(2, -1, +1)$, respectively, we have $h(L) = 1$, $h(R) = 0$, and in fact it is easy to see that the one-step conditional expectation of h given any state $u \neq (2N, +1, -1)$ or $(0, -1, +1)$, is equal to $h(u)$. This relationship can be expressed as a simple recursion: Letting $f(k) = h(2k, +1, -1)$ and $g(k) = h(2k + 2, -1, +1)$, we have,

$$\begin{cases} f(k) = (1 - \epsilon) f(k + 1) + \epsilon g(k), & \text{for } 0 \leq k \leq N - 1, \\ g(k + 1) = (1 - \epsilon) g(k) + \epsilon f(k + 1), & \text{for } -1 \leq k \leq N - 2, \\ g(0) = 0 \text{ and } f(N - 1) = 1. \end{cases}$$

Adding the first two equations above shows that $f(k) - g(k)$ is a constant, say A, independent of k, and substituting this in the recursion for g gives $g(k) = A\epsilon k/(1 - \epsilon)$. Similarly solving for f we obtain, $f(k) = A + A\epsilon k/(1 - \epsilon)$, and from the boundary values we can solve for A to get, $A = (1 - \epsilon)/(1 + \epsilon(N - 2))$. Therefore,

$$P_{\psi_1}(C^c) = h(u_1) = f(0) = A = (1 - \epsilon)/(1 + \epsilon(N - 2)),$$

as claimed. □

Combining (4) with the result of Lemma 4.3 completes the proof of the theorem. □

Proof of Theorem 3.3. Recall the ergodic chain $\boldsymbol{\Psi}$ defined in the beginning of the proof of Theorem 3.1. Write Σ_ψ for its state space, π_ψ for its unique invariant measure, and let $\boldsymbol{P}$ denote its transition kernel, $\boldsymbol{P}(\psi, \psi') = \Pr(\Psi_{t+1} = \psi' | \Psi_t = \psi)$, $\psi, \psi' \in \Sigma_\psi$. Now consider the bivariate chain $\tilde{\boldsymbol{\Psi}} = \{\tilde{\Psi}_t = (\Psi_t, \Psi_{t+1}) \,;\, t \geq 0\}$. Then $\tilde{\boldsymbol{\Psi}}$ is also ergodic, with unique invariant measure,

$$\tilde{\pi}(\psi, \psi') = \pi_\psi(\psi) \boldsymbol{P}(\psi, \psi'),$$

for every state (ψ, ψ') of $\tilde{\boldsymbol{\Psi}}$. Therefore, the limit in the statement indeed exists a.s., and it equals,

$$c := P_\pi(I_2 \neq I_1) = \tilde{\pi}\Big(\mathbb{I}_{\{I_2 \neq I_1\}}\Big) = \tilde{\pi}(B),$$

where B consists of the following 8 states,

$$\begin{aligned} B = \Big\{ &\Big((0,+1,+1,1),(0,-1,+1,2)\Big), \quad \Big((0,-1,-1,1),(0,-1,+1,2)\Big), \\ &\Big((0,+1,+1,2),(0,+1,-1,1)\Big), \quad \Big((0,-1,-1,2),(0,+1,-1,1)\Big), \\ &\Big((2,-1,+1,1),(0,-1,+1,2)\Big), \quad \Big((2,-1,+1,2),(0,+1,-1,1)\Big), \\ &\Big((-2,+1,-1,1),(0,-1,+1,2)\Big), \quad \Big((-2,+1,-1,2),(0,+1,-1,1)\Big)\Big\}, \end{aligned}$$

and where, with a slight abuse of notation, the negative values of the Y_t variables above are again interpreted modulo N.

In order to compute the actual value of $c = \tilde{\pi}(B)$, we first observe that,

$$\begin{aligned} \tilde{\pi}(B) &= \Big(\pi_\psi(0,+1,+1,1) + \pi_\psi(0,-1,-1,1) + \pi_\psi(0,+1,+1,2) + \pi_\psi(0,-1,-1,2)\Big)\epsilon(1-\epsilon) \\ &\quad + \Big(\pi_\psi(2,-1,+1,2) + \pi_\psi(-2,+1,-1,1)\Big)\epsilon^2 \\ &\quad + \Big(\pi_\psi(2,-1,+1,1) + \pi_\psi(-2,+1,-1,2)\Big)(1-\epsilon)^2 \\ &= \frac{1}{2N}\epsilon(1-\epsilon) + \Big(\pi_\psi(-2,+1,-1,1) - \pi_\psi(2,-1,+1,1) + \frac{1}{4N}\Big)\epsilon^2 \\ &\quad + \Big(\pi_\psi(2,-1,+1,1) - \pi_\psi(-2,+1,-1,1) + \frac{1}{4N}\Big)(1-\epsilon)^2, \end{aligned}$$

where we used the fact that the invariant distribution of $(X_t(1), X_t(2), D_t(1), D_t(2))$ is uniform, which implies that $\pi_\psi(y,d,d',1) + \pi_\psi(y,d,d',2) = 1/(4N)$, for any $y \in S$ and $d, d' \in \{+1,-1\}$. Simplifying,

$$\tilde{\pi}(B) = \frac{1}{4N} + (1-2\epsilon)[\pi_\psi(2,-1,+1,1) - \pi_\psi(-2,+1,-1,1)]. \tag{5}$$

To compute the difference of the two probabilities in (5), recall the definition of the regeneration time T and the measure ν in (2) and (3), respectively. By Lemma 4.1, for any state $\psi \in \Sigma_\psi$ of the form $\psi = (y,d,d',1)$, we have,

$$E_\nu(T)\pi_\psi(\psi) = E_\nu\left[\sum_{t=0}^{T-1} \mathbb{I}_{\{\Psi_t=\psi\}}\right] = \frac{1}{2}E_{\psi_1}\left[\sum_{t=0}^{T-1} \mathbb{I}_{\{\Psi_t=\psi\}}\right] = \frac{1}{2}E_{u_1}\left[\sum_{t=0}^{T-1} \mathbb{I}_{\{\Upsilon_t^*=(y,d,d')\}}\right],$$

where the chain $\boldsymbol{\Upsilon}^*$ was defined in the proof of Lemma 4.3 and $u_1 = (0,+1,-1)$ as before. Therefore, substituting this twice in (5) and recalling the discussion of the evolution of $\boldsymbol{\Upsilon}^*$ until time T from the proof of Lemma 4.3, we have,

$$c = \tilde{\pi}(B) = \frac{1}{4N} + \frac{(1-2\epsilon)}{4N}E_{u_1}\left[\sum_{t=0}^{T-1}\Big(\mathbb{I}_{\{\Upsilon_t^*=(2,-1,+1)\}} - \mathbb{I}_{\{\Upsilon_t^*=(2N-2,+1,-1)\}}\Big)\right], \tag{6}$$

where we also used the result of Lemma 4.2. By the definition of T, the value of the sum inside the last expectation above is either $0-1$ or $1-0$, and the corresponding probabilities can be found by looking at the trace of the chain $\boldsymbol{\Upsilon}^*$ on the set Σ^t, as defined in the proof

of Lemma 4.3. Indeed, referring to Figure 3, and in the notation of the proof of Lemma 4.3, the case $1-0$ has probability $1-P_{u_1}(T_L < T_R)$, whereas the case $0-1$ has probability $P_{u_1}(T_L < T_R)$. So (6) becomes,

$$c = \tilde{\pi}(B) = \frac{1}{4N} + \frac{(1-2\epsilon)}{4N}[1 - 2P_{u_1}(T_L < T_R)],$$

and now substituting the result of Lemma 4.3, $P_{u_1}(T_L < T_R) = \frac{1-\epsilon}{1+\epsilon(N-2)}$, and simplifying, yields precisely the claimed result. □

4.2 The continuous circle

Proof of Theorem 3.4. We begin with some simple notation. Let,

$$\begin{aligned} D &= \{(z,z) : z \in S\} \subset S^2 \\ F &= \{(z_1,z_2,d_1,d_2,i) \in \Sigma : z_1 = z_2,\ d_1 d_2 = -1\} \subset \Sigma, \end{aligned}$$

and note that, by the definition of Σ, for any $(z_1,z_2,d_1,d_2,i) \in F$ we can always take, without loss of generality, $d_i = +1$.

The limit in the statement of the theorem exists a.s. by Theorem 2.2; in order to compute its value we first define the stopping time,

$$T = \inf\big\{t > 0 :\ \Phi_t \in F\big\}. \tag{7}$$

Let ν denote the uniform probability measure on F. Then, T is a regeneration time for ν, in the sense that, if $\Phi_0 \sim \nu$, then $\Phi_T \sim \nu$, as well. We go on to compute $\pi(D_1(I_1))$ using the following natural continuous-time generalization of Kac's formula, proved in the Appendix.

Lemma 4.4. *If T is a regeneration stopping time for ν with $E_\nu(T) < \infty$, then for any bounded measurable function $f : \Sigma \to \mathbb{R}$ we have,*

$$E_\nu\left(\int_0^T f(\Phi_s)ds\right) = E_\nu(T)\pi(f).$$

First, we compute the expectation of the regeneration time T, conditional on the initial state $\Phi_0 = \phi$ being in F.

Lemma 4.5. *For any initial state $\phi \in F$, we have: $E_\phi(T) = \frac{N}{v}$.*

Proof. For any state $\phi = (x,d,i) = (x_1,x_2,d_1,d_2,i) \in \Sigma$, write,

$$\delta(x) := x_1 - x_2 \ (\mathrm{mod}\, N) \in [0,N).$$

We can compute $E_\phi(T)$ using simple tools from the potential theory of Markov processes. To that end, we will construct a function $H : \Sigma \to \mathbb{R}$ that formally satisfies $LH(\phi) = -1$, for all $\phi \in \Sigma \setminus F$. Note that such a function would not be in the domain of the generator, as can be easily seen by integrating both sides with respect to the invariant measure. Nevertheless, if we start the process from $\Phi_0 = \phi \notin F$, then it can be checked that $\{H(\Phi_t) + t\}$ is a martingale up to time T.

It is not hard to find an explicit solution to $LH(\phi) = -1$ for all $\phi \notin F$. Indeed, let, for $\phi \notin F$,

$$H(\phi) := \left(\frac{N - 2\delta(x)}{4v}\right)\big(d_1 - d_2\big) + \frac{1 + d_1 d_2}{4r} + \frac{r\,\delta(x)\big(N - \delta(x)\big)}{2v^2},$$

and for $\phi \in F$ as,

$$H(\phi) := -\frac{N}{2v}.$$

Recalling the form of the generator L from (1), it is straightforward to verify that H satisfies, $LH(\phi) = -1$, for all $\phi \notin F$, and that H is discontinuous across F, in that,

$$\lim_{\delta(x)\downarrow 0} H(x,d,i) - \lim_{\delta(x)\uparrow N} H(x,d,i) = \frac{N}{2v}\big(d_1 - d_2\big).$$

For $\phi \notin F$, an application of the optional stopping theorem for the martingale $\{H(\Phi_t) + t\}$ gives,

$$E_\phi(T) = H(\phi) + \frac{N}{2v}. \tag{8}$$

On the other hand, for $\phi \in F$, the Markov property gives,

$$E_\phi(T) = E_\phi\big[T;\ T \le t\big] + E_\phi\big[E_{\Phi_t}(T);\ T > t\big],$$

for all $t > 0$. Since, as $t \to 0$, we have $P_\phi(T \le t) \to 0$ and $H(\Phi_t) \to N/(2v)$, P_ϕ-a.s., for all $\phi \in F$, letting $t \to 0$ in the preceding equation and recalling (8) completes the proof. □

Now let $d^* := (+1, -1)$, $\phi^* := (0, 0, d^*, 1) \in F$, and note that, for all $x \in D$,

$$E_{(x,d^*,1)}\left[\int_0^T D_t(I_t)\,dt\right] = E_{(x,-d^*,2)}\left[\int_0^T D_t(I_t)\,dt\right] = E_{\phi^*}\left[\int_0^T D_t(1)\,dt\right].$$

Combining this, with an application of Lemma 4.4 with $f(\Phi_t) = D_t(I_t)$, and with Lemma 4.5, gives,

$$s = \frac{v^2}{N} E_\nu\left[\int_0^T D_t(I_t)\,dt\right] = \frac{v^2}{N} E_{\phi^*}\left[\int_0^T D_t(1)\,dt\right].$$

Therefore, writing $X_t^*(i)$ for the total clockwise displacement distance travelled by walker $i = 1, 2$, up to time $t \ge 0$, so that $X_t^*(i) = v\int_0^t D_t(i)\,dt$, we have,

$$\begin{aligned} s &= \frac{v}{N} E_{\phi^*}\big[X_T^*(1)\big] \\ &= \frac{v}{2N} E_{\phi^*}\big[X_T^*(1) + X_T^*(2)\big] + \frac{v}{2N} E_{\phi^*}\big[X_T^*(1) - X_T^*(2)\big] \\ &= \frac{v}{2N} E_{\phi^*}\big[X_T^*(1) - X_T^*(2)\big], \end{aligned}$$

where we used the fact that, since the two walkers start in opposite directions, we have $E_{\phi^*}[X_T^*(1) + X_T^*(2)] = 0$ by symmetry.

There are exactly two scenarios for the first meeting of the two walkers starting from F with $D_0(I_0) = 1$: In the first one, they meet with directions that are opposite to the ones they

started with. In this case, we necessarily have $X_T^*(I_0) - X_T^*(I_0') = 0$, where I_0' denotes the complementary index to I_0, i.e., $I_0' = 3 - I_0$. We call this event C. In the second scenario, corresponding to event C^c, the directions of the walkers are the same as the ones they started with, which necessarily means that $X_T^*(I_0) - X_T^*(I_0') = N$. Therefore, the speed s is:

$$s = \frac{v}{2} P_{\phi^*}(C^c).$$

The proof of the theorem is completed by an application of Lemma 4.6 below. □

Lemma 4.6. *For any initial state $\phi \in F$ of the form $\phi = (x, 1, -1, 1)$, we have:*

$$P_\phi(C^c) = \frac{2v}{2v + rN}.$$

Proof. Consider the function $V : \Sigma \setminus F \to \mathbb{R}$, defined, for any state $\phi = (x, d_1, d_2, i) \notin F$ as,

$$V(\phi) := \frac{r\delta(x) + v[1 + (d_1 - d_2)/2]}{rN + 2v},$$

where $\delta(x) := x_1 - x_2 \pmod N$, as before. It is straightforward to verify that $LV(\phi) = 0$, for every $\phi \notin F$. An application of the optional stopping theorem for the martingale $\{V(\Phi_t)\}$ gives,

$$P_\phi(C^c) = V(\phi),$$

for all $\phi = (x, 1, -1, 1)$ with $x \notin D$. The proof is concluded by repeating the same argument as in the end of the proof of Lemma 4.5, with the random variable $\mathbb{I}_{C^c}$ in place of T, and using the Markov property at time $t > 0$ and letting $t \to 0$. □

Proof of Theorem 3.6. Using the same notation as in the proof of Theorem 3.4, we define, for each $n \geq 0$, the time T_n as the time of the nth return of $\mathbf{\Phi}$ to F, i.e., $T_0 := 0$ and inductively, for all $n \geq 1$,

$$T_{n+1} := \inf\{t > T_n : \Phi_t \in F\}.$$

The number of excursions around F up to time $t \geq 0$ is then,

$$N_t = \max\{n \geq 0 : \ T_n \leq t\}.$$

We also define the independent Bernoulli random variables $\{J_n; n \geq 1\}$, that take the value 1 exactly when the event C occurs during the nth excursion around F. Lemma 4.6 then implies that, for all $n \geq 2$,

$$P_\phi(J_n = 1) = \frac{rN}{2v + rN}.$$

In this notation, the total transmission cost M_t up to time $t \geq 0$ is given by,

$$M_t = \sum_{n=1}^{N_t} J_n.$$

Since, by Lemma 4.5, $N_t/t \to v/N$, a.s., as $t \to \infty$, we have,

$$c := \lim_{t\to\infty} \frac{M_t}{t} = \frac{v}{N} \times \frac{rN}{2v + rN} = \frac{rv}{2v + rN}, \quad \text{a.s.},$$

as claimed. □

A Appendix

A.1 Proof of Proposition 2.1

Parts (i) and (ii) are immediate consequences of the more general result in Theorem 2.2, established below. For (iii) it is easiest to work with the infinitesimal generator L_0 of $\mathbf{W}$. Let A_{L_0} denote the collection of all continuous functions $f : S^m \times \{-1,+1\}^m \to \mathbb{R}$, such that f is continuously differentiable in each $x(j)$, $1 \le j \le m$. Also, for any m-tuple of directions $d \in \{-1,+1\}^m$, let $\sigma^j d$ be the same as d but with its jth coordinate having the opposite sign from that of d, $1 \le j \le m$. Then the action of L_0 on any $f \in A_{L_0}$ is,

$$L_0 f(x,d) = \sum_{j=1}^{m} \left\{ v d(j) \frac{\partial f}{\partial x(j)}(x,d) + r\left[f(x,\sigma^j d) - f(x,d) \right] \right\}.$$

Arguing as for the operator L at the end of Section 2.2 we get that the domain of L_0 is A_{L_0}. It is now a simple computation to show that, if π_0 denotes the uniform distribution on $S^m \times \{-1,+1\}^m$, then $\int L_0 f d\pi_0 = 0$, for any $f \in A_{L_0}$, so that π_0 is indeed invariant [20, Theorem 3.37]. The uniqueness of π_0 follows from Harris recurrence [15]. □

A.2 Proof of Theorem 2.2

For the sake of simplicity, we assume that all three parameters, N, v and r, are equal to 1. A cursory examination of the proof below should immediately reveal that the general case only involves notational modifications.

Given $\epsilon \in (0, 1/(3m))$ arbitrary, let $S_\epsilon \subset S^m$ denote the set,

$$S_\epsilon := S^m \setminus \{x : d(x(j), x(k)) < 3\epsilon, \text{ for some } j \neq k\},$$

where $d(\cdot,\cdot)$ denotes the usual distance on S, so that $d(x(j), x(k))$ is equal to the length of the shortest arc connecting points $x(j)$ and $x(k)$. Note that, since $\epsilon < 1/(3m)$, S_ϵ has nonempty interior. We also define the set $\Sigma_\epsilon := S_\epsilon \times \{-1,1\}^m \times \{1,2,\ldots,m\}$, and the measures $\mathcal{L}_\epsilon$ such that $d\mathcal{L}_\epsilon / d\mathcal{L}^m = \mathbb{I}_{S_\epsilon}$ on S^m, and $\mu := \mathcal{L}_\epsilon \times \kappa^m \times \kappa_m$ on Σ. Note that $\mathcal{L}_\epsilon$ and μ are supported on S_ϵ and Σ_ϵ, respectively.

The main step in the proof is the following Doeblin-like domination condition:

Proposition A.1. *Let $\epsilon \in (0, 1/(3m))$ arbitrary. Then, for every $t \ge t_0 := 2 + 6\epsilon$, every measurable $A \subset \Sigma$, and every initial state $\phi \in \Sigma$, we have,*

$$P_\phi(\Phi_t \in A) \ge c\mu(A), \tag{9}$$

with $c = (\epsilon^2 e^{-1-(3/2)t_0}/4)^m$.

Before giving the proof we make two simple observations. First, we will actually prove that, under the assumptions of the proposition,

$$P_\phi(\Phi_t \in A) \ge c' e^{-(3/2)tm} \mu(A), \tag{10}$$

with $c' = (\epsilon^2 e^{-1}/4)^m$. Then (9) follows for each $t \geq t_0$ by the Markov property:

$$P_\phi(\Phi_t \in A) = E_\phi\big[P_{\Phi_{t-t_0}}\Phi_{t_0} \in A\big] \geq c\mu(A).$$

Second, it suffices to establish (10) for events A of the form,

$$A = \prod_{j=1}^{m} B_j \times \prod_{j=1}^{m} \{d(j)\} \times \{1\}, \tag{11}$$

with each $d(j) \in \{-1, 1\}$, and each $B_j \subset S$. The proof is based on the following construction.

For any two points $x, x' \in S$, we say that a walker *travels 'clockwise' from x to x' in time t*, when the walker takes the following steps. Let $x'' = \min\{y \in (x, \infty) : y = x' \pmod 1\}$. Then the walker, starting at x, travels counter-clockwise for time $z = (t + x - x'')/2$, then switches directions and travels clockwise for time $t - z$, ending up at point x''. Similarly, we say that a walker *travels 'counter-clockwise' from x to x' in time t*, when the reverse of the above process occurs, so that, now, $x'' = \max\{y \in (-\infty, x) : y = x' \pmod 1\}$.

Let A be as in (11). A simple strategy for the motion of the m walkers that guarantees $\Phi_t \in A$ is the following: Walker 1 travels clockwise from point $x(1)$ to some point in B_1 in time t, while all other walkers travel counter-clockwise from their initial positions $x(k)$ to some point in the corresponding B_k in time t. During the final segment of her travel, walker 1 will encounter each of the other walkers at least once, and she will be moving clockwise while all of the other walkers will be moving counter-clockwise. Therefore, the message will certainly be with walker 1 at time t.

Proof of Proposition A.1. Let $\epsilon \in (0, 1/(3m))$ and $\phi = (x_0, d_0, i_0) \in \Sigma$ arbitrary, and without loss of generality take $A \subset \Sigma$ to be of the form (11).

For any position $x \in S$, any pair of directions $d, d' \in \{-1, +1\}$, any time $t > 0$, and any measurable $C \subset S$, we will define the event $M^+(x, d, C, d', t, \epsilon)$ that, roughly speaking, will describe a scenario in which a walker starts at x with direction d, and ends up in C with direction d' after time t. To make this precise, let e_0, e_1, e_2 and e_3 be independent exponential random variables with mean equal to 1, and imagine a walker starting at time zero in position x with direction d, taking the following steps:

- $+e_0, -e_1, +e_2$, if $d = d' = 1$ (i.e., the walker first moves clockwise for time e_0, then counter-clockwise for time e_1, and then clockwise again for time e_2);
- $+e_0, -e_1, +e_2, -e_3$, if $d = 1$, $d' = -1$;
- $-e_1, +e_2$, if $d = -1$, $d' = 1$;
- $-e_1, +e_2, -e_3$, if $d = d' = -1$.

Now we can compute the location of the walker at time t. Write $E_0 = e_0 \mathbb{I}_{\{d=1\}}$ and consider two cases:

If $d' = 1$, and assuming that,

$$\frac{t-1}{2} < e_1 < \frac{t}{2}, \quad e_2 > t, \quad \text{and} \quad E_0 \leq \epsilon, \tag{12}$$

then at time $2e_1$ the walker has covered the distances E_0, $-e_1$, and $e_1 - E_0$, and therefore is at point x and is moving clockwise. Moreover, since $e_2 > t > 2e_1$, after the remaining $t - 2e_1$ time units, the walker ends up in position $x + t - 2e_1$.

If $d' = -1$, and assuming that,

$$E_0 + e_1 + e_2 < t < E_0 + e_1 + e_2 + e_3, \tag{13}$$

then at time $E_0+e_1+e_2$ the walker is located at $x+E_0-e_1+e_2$ and moving counter-clockwise. Therefore, at time t, i.e., after travelling counter-clockwise for an additional $t-(E_0+e_1+e_2) < e_3$ time units, the walker's final position is $x+E_0-e_1+e_2-(t-E_0-e_1-e_2) = x-t+2(E_0+e_2)$.

To summarize, under assumptions (12) and (13), the position (in $\mathbb{R}$, not necessarily in S), of the walker at time t is,

$$U := \begin{cases} x + t - 2e_1, & \text{if } d' = 1, \\ x - t + 2(E_0 + e_2), & \text{if } d' = -1. \end{cases} \tag{14}$$

The last ingredient we need for the formal definition of M^+ is the following. Given x and C, we write,

$$C_x := (C \cap (x, 1]) \cup (C \cap [0, x] + 1) \subset (x, x + 1],$$

so that C_x is the same as C, but the points before x have been pushed one unit later. Now we can formally define:

$$M^+(x, d, C, d', t, \epsilon) := \begin{cases} \{E_0 \le \epsilon, e_2 > t, U \in C_x\}, & \text{if } d' = 1, \\ \{E_0 \le \epsilon, e_3 > \epsilon, E_0 + e_1 + e_2 \in (t - \epsilon, t), U \in C_x\}, & \text{if } d' = -1. \end{cases}$$

Note that both conditions (12) and (13) above are satisfied on $M^+(x, d, C, d', t, \epsilon)$. We also define the "symmetric" event $M^-(x, d, C, d', t, \epsilon)$ as the event in which the reflection of the walker follows the path described by $M^+(-x, -d, -C, -d', t, \epsilon)$. Finally, let,

$$G := \left\{ \begin{array}{l} \text{The path of walker 1 up to time } t \text{ is in } M^+(x_0(1), d(1), B_1, d'(1), t, \epsilon), \\ \text{and the path of each walker } j \neq 1 \text{ up to } t \text{ is in } M^-(x_0(j), d(j), B_k, d'(j), t, \epsilon) \end{array} \right\}.$$

Now we claim that, for all $t > t_0 := 2 + 6\epsilon$:

$$G \subset \{\Phi_t \in A\}. \tag{15}$$

To see that (15) holds, again we consider two cases. If $d' = 1$, then time t is reached during step e_2, because $E_0 < \epsilon, e_1 = (t+x-U)/2 < t/2$ and $e_2 > t$, so that $E_0+e_1 < t < E_0+e_1+e_2$. And if $d' = -1$, then time t is reached during step e_3. Therefore, a walker following $M^+(x, d, C, d', t, \epsilon)$ travels clockwise at least in the time interval $[t - (1 + 2\epsilon), t - \epsilon]$, which has length $1 + \epsilon$. To see this, note that, if $d' = 1$, then we have $e_1 = (t - (U - x))/2 > \epsilon \ge E_0$, and the last part of the trajectory (i.e., the part of step e_2 until time t) takes time,

$$e_1 - E_0 + U - x = (t + U - x)/2 - E_0 \ge t/2 - \epsilon \ge 1 + 2\epsilon,$$

while if $d' = -1$, then,

$$e_2 = (U - x + t)/2 - E_0 \ge t/2 - \epsilon \ge 1 + 2\epsilon,$$

and also $E_0 + e_1 + e_2 \in (t-\epsilon, t)$. Consequently, in the time interval $[t-(1+2\epsilon), t-\epsilon]$, walker 1 meets every other walker and thus gets the message. In the remaining time interval, $[t-\epsilon, t]$, the walkers do not meet again because each is at distance at most ϵ from their final positions, and the distance between any two of these final positions is at least 3ϵ.

Now, (15) implies that,

$$\begin{aligned} P_\phi(\Phi_t \in A) &\geq P_\phi(G) \\ &= P_\phi\big(M^+(x_0(1), d(1), B_1, d'(1), t, \epsilon)\big) \prod_{j=2}^{m} P_\phi\big(M^-(x_0(j), d(j), B_k, d'(j), t, \epsilon)\big), \end{aligned}$$

and the bound in (10) and hence the result of the proposition follow from the bound in the lemma immediately below. □

Lemma A.2. *In the notation and under the assumptions of the above proof, for any position $x \in S$, and pair of directions $d, d' \in \{+1, -1\}$, any measurable $C \subset S$, and any initial state $\phi \in \Sigma$, we have,*

$$P_\phi\big(M^+(x, d, C, d', t, \epsilon)\big) \geq c_\epsilon \mathcal{L}_\epsilon(C),$$

with $c_\epsilon = \frac{1}{4} e^{-1} \epsilon^2 e^{-3t/2}$. By symmetry, the same bound holds for M^- in place of M^+.

Proof of Lemma A.2. When $d' = 1$, the random variable U defined in (14) has density $f_U(s) = (1/2) e^{-(x+t)/2} e^{s/2} \mathbb{I}_{\{s<x+t\}}$, so that,

$$P_\phi(M^+(x, d, C, d', t, \epsilon)) = (1 - e^{-\epsilon} \mathbb{I}_{\{d=1\}}) e^{-t} \frac{1}{2} e^{-(x+t)/2} \int_{C_x} e^{s/2}\, ds \geq \frac{1}{4} \epsilon e^{-1/2} e^{-3t/2} \mathcal{L}_\epsilon(C),$$

where we used the elementary inequality, $1 - e^{-\epsilon} \geq \epsilon/(1+\epsilon) \geq \epsilon/2$, $\epsilon \in [0, 1]$.

When $d' = -1$, we consider two cases. If $d = -1$, then the joint density of $(U, V) := (x - t + 2e_2, e_1 + e_2)$ is,

$$f_{U,V}(u, v) = \frac{1}{2} e^{-v} \mathbb{I}_{\{u>x-t\}} \mathbb{I}_{\{2v-u>t-x\}},$$

and thus,

$$\begin{aligned} P_\phi(M^+(x, d, C, d', t, \epsilon)) &= \Pr(U \in C_x, V \in (t-\epsilon, t)) \\ &= \frac{1}{2} \int_{C_x} \int_{t-\epsilon}^{t} e^{-v}\, dv\, du \\ &\geq \frac{1}{2} \epsilon e^{-1} e^{-t} \mathcal{L}_\epsilon(C). \end{aligned}$$

Finally, when $d = 1$, the joint density of $(U, V, W) := (x - t + 2(e_0 + e_2), e_0 + e_1 + e_2, e_0)$ is,

$$f_{U,V,W}(u, v, w) = \frac{1}{2} e^{-v} \mathbb{I}_{\{2v-u>t-x\}} \mathbb{I}_{\{0<w<(u+t-x)/2\}},$$

and thus,

$$P_\phi(M^+(x, d, C, d', t, \epsilon)) = \frac{1}{2} \int_0^\epsilon \int_{C_x} \int_{t-\epsilon}^{t} e^{-v}\, dv\, du\, dw \geq \frac{1}{2} \epsilon^2 e^{-t} \mathcal{L}_\epsilon(C).$$

Combining the three bounds derived gives the required result. □

We are now in a position to establish the four claims of Theorem 2.2. If $\psi(A) > 0$ for some $A \subset \Sigma$, then $\mu(A) = \psi(A \cap \Sigma_\epsilon) > 0$ for all $\epsilon > 0$ small enough, so Proposition A.1 implies that $\mathbf{\Phi}$ is ψ-irreducible [23]. Moreover, it implies that the state space Σ itself is small [23], hence it is petite, so that $\mathbf{\Phi}$ is aperiodic [13, Eq. (10)]. This establishes (i).

Since Σ is small, the drift condition $(\mathcal{D}_T)$ of [13] holds with Lyapunov function $V \equiv 1$, and [13, Theorem 5.2] implies that $\mathbf{\Phi}$ is uniformly (exponentially) ergodic. This implies (iii) by the definition of V-uniform ergodicity [13, Eq. (11)]. In particular, $\mathbf{\Phi}$ is positive Harris recurrent, hence it has a unique invariant probability measure [15], giving (ii). Finally, the ergodic theorem of part (iv) is a standard consequence of positive recurrence of strong Markov processes, see, e.g., [22, Theorem 5.1] or [3, Proposition 3.7]. □

A.3 Proof of Lemma 4.4

If $P_\nu(T = 0) = 1$, the claim is trivial. Suppose now $0 < E_\nu(T) < \infty$, and consider the following occupation measure ζ on Σ: For any measurable A,

$$\zeta(A) = E_\nu \left(\int_0^T \mathbb{I}_A(\Phi_s) ds \right).$$

Observe that ζ integrates measurable functions on Σ as,

$$\int f \, d\zeta = E_\zeta \left(\int_0^T f(\Phi_s) ds \right).$$

Let now $g \in \mathcal{D}(L)$. By an application of the optional stopping theorem for the martingale $\{g(\Phi_t) - \int_0^t Lg(\Phi_s) ds\}$ at the stopping time T, we get,

$$E_\nu\big[g(\Phi_T)\big] - E_\nu \left(\int_0^T Lg(\Phi_s) ds \right) = E_\nu\big[g(\Phi_0)\big].$$

As T is a regeneration time for ν, we have that, $E_\nu\big[g(\Phi_T)\big] = E_\nu\big[g(\Phi_0)\big]$, and hence,

$$\int Lg \, d\zeta = E_\nu \left(\int_0^T Lg(\Phi_s) ds \right) = 0.$$

Since this holds for all $g \in \mathcal{D}(L)$, the normalized occupation measure $\bar{\zeta} := \zeta / \zeta(\Sigma) = \zeta / E_\nu(T)$ is an invariant probability measure under the dynamics of $\mathbf{\Phi}$, therefore, by uniqueness, $\bar{\zeta} = \pi$, as required. □

References

[1] D. Aldous and J.A. Fill. Reversible markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014, available online at `www.stat.berkeley.edu/~aldous/RWG/book.html`.

[2] G. Araniti, N. Bezirgiannidis, E. Birrane, I. Bisio, S. Burleigh, C. Caini, M. Feldmann, M. Marchese, J. Segui, and K. Suzuki. Contact graph routing in DTN space networks: Overview, enhancements and performance. *IEEE Commun. Mag.*, 53(3):38–46, 2015.

[3] S. Asmussen. *Applied probability and queues*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons, Ltd., Chichester, 1987.

[4] E. Baccelli, P. Jacquet, B. Mans, and G. Rodolakis. Highway vehicular delay tolerant networks: Information propagation speed properties. *IEEE Trans. Inform. Theory*, 58(3):1743–1756, March 2012.

[5] E. Baccelli, P. Jacquet, B. Mans, and G. Rodolakis. Multi-lane vehicle-to-vehicle networks with time-varying radio ranges: Information propagation speed properties. In *2013 IEEE International Symposium on Information Theory (ISIT)*, Istanbul, Turkey, July 2013.

[6] N. Bansal, N. Buchbinder, A. Madry, and J. Naor. A polylogarithmic-competitive algorithm for the k-server problem. *J. ACM*, 62(5):40:1–40:49, November 2015.

[7] R. Cavallari, S. Toumpis, and R. Verdone. Analysis of hybrid geographic/delay-tolerant routing protocols for wireless mobile networks. In *IEEE International Conference on Computer Communications (INFOCOM)*, Honolulu, HI, April 2018.

[8] R. Cavallari, R. Verdone, and S. Toumpis. Cost/speed analysis of mobile wireless DTNs under random waypoint mobility. In *2016 International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, Tempe, AZ, May 2016.

[9] D. Cheliotis, I. Kontoyiannis, M. Loulakis, and S. Toumpis. Exact speed and transmission cost in a simple one-dimensional wireless delay-tolerant network. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 476–480, Aachen, Germany, June 2017.

[10] D. Cheliotis, I. Kontoyiannis, M. Loulakis, and S. Toumpis. Analysis of a one-dimensional continuous delay-tolerant network model. In *2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, Kalamata, Greece, June 2018.

[11] B. Csaba and S. Lodha. A randomized on-line algorithm for the k-server problem on a line. *Random Structures & Algorithms*, 29(1):82–104, 2006.

[12] S.N. Diggavi, M. Grossglauser, and D.N.C. Tse. Even one-dimensional mobility increases the capacity of wireless networks. *IEEE Trans. Inform. Theory*, 51(11):3947–3954, 2005.

[13] D. Down, S.P. Meyn, and R.L. Tweedie. Exponential and uniform ergodicity of Markov processes. *Ann. Probab.*, 23(4):1671–1691, October 1995.

[14] S.N. Ethier and T.G. Kurtz. *Markov processes : Characterization and convergence*. John Wiley & Sons, New York, 1986.

[15] R.K. Getoor. Transience and recurrence of Markov processes. In *Séminaire de Probabilités XIV 1978/79*, pages 397–409. Springer, 1980.

[16] M. Grossglauser and D.N.C. Tse. Mobility increases the capacity of ad-hoc wireless networks. In *IEEE International Conference on Computer Communications (INFOCOM)*, volume 3, pages 1360–1369, Anchorage, AL, April 2001.

[17] P. Hui, J. Crowcroft, and E. Yoneki. Bubble rap: Social-based forwarding in delay-tolerant networks. *IEEE Trans. Mobile Comput.*, 10(11):1576–1589, 2011.

[18] P. Jacquet, B. Mans, and G. Rodolakis. Information propagation speed in mobile and delay tolerant networks. *IEEE Trans. Inform. Theory*, 56:5001–5015, October 2010.

[19] I. Kontoyiannis, S. Toumpis, R. Cavallari, and R. Verdone. On calculating packet speed and cost in a mobile wireless network model. In *2018 IEEE International Symposium on Information Theory (ISIT)*, Vail, CO, June 2018.

[20] T.M. Liggett. *Continuous time Markov processes. An introduction.* Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2010.

[21] M. Manasse, L. McGeoch, and D. Sleator. Competitive algorithms for on-line problems. In *Twentieth annual ACM Symposium on Theory of Computing*, pages 322–333, 1988.

[22] G. Maruyama and H. Tanaka. Ergodic property of N-dimensional recurrent Markov processes. *Memoirs of the Faculty of Science, Kyushu University. Series A, Mathematics*, 13(2):157–172, 1959.

[23] S.P. Meyn and R.L. Tweedie. Stability of Markovian processes II: Continuous time processes and sampled chains. *Adv. Appl. Prob.*, 25(3):487–517, September 1993.

[24] S.P. Meyn and R.L. Tweedie. Stability of Markovian processes III: Foster-Lyapunov criteria for continuous time processes. *Adv. Appl. Prob.*, 25(3):518–548, September 1993.

[25] J.R. Norris. *Markov chains.* Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, UK, 1998.

[26] P.N. Pathirana, N. Bulusu, A.V. Savkin, and S. Jha. Node localization using mobile robots in delay-tolerant sensor networks. *IEEE Trans. Mobile Comput.*, 4(3):285–296, 2005.

[27] Y. Peres, A. Sinclair, P. Sousi, and A. Stauffer. Mobile geometric graphs: Detection, coverage, and percolation. *Probab. Theory Relat. Fields*, 156(1):273–305, 2013.

[28] M. Sharpe. *General theory of Markov processes.* Pure and Applied Mathematics. Academic Press, Inc., Boston, MA, 1988.

[29] A. Vasilakos, Y. Zhang, and T.V. Spyropoulos, editors. *Delay Tolerant Networks: Protocols and applications.* CRC Press, 2011.

[30] M. Zarei, A.M. Rahmani, and H. Samimi. Connectivity analysis for dynamic movement of vehicular ad hoc networks. *Wireless Networks*, 23(3):843–858, 2017.