SIAM J. Discrete Math.
Vol. 22, No. 2, pp. 747–775

WALKERS ON THE CYCLE AND THE GRID*

J. DÍAZ[†], X. PÉREZ[†], M. J. SERNA[†], and N. C. WORMALD[‡]

Abstract. We present a model of the establishment and maintenance of communication between mobile agents. We assume that the agents move through a fixed environment modeled by a motion graph and are able to communicate if they are within distance at most d of each other. As the agents move randomly, we analyze the evolution in time of the connectivity between a set of w agents, asymptotically for a large number N of vertices, when w also grows large. The particular topologies of the environment we study here are the cycle and the toroidal grid.

Key words. mobile agents, ad hoc network, connectivity, distance graph

AMS subject classifications. 68R10, 60C05

DOI. 10.1137/060659855

1. Introduction.

Consider a setting in which a large number of mobile agents can perform concurrent basic movements: either ahead/behind/left/right, determining a grid pattern, or left/right, describing a line. Each agent can communicate directly with any other agents which are within a given distance d. This enables communications with agents at a further distance using several intermediate agents. At each step in time there is an ad hoc network defined by the dynamic graph whose vertex set consists of the agents, with an edge between any two agents iff they are within distance d of each other.

Various aspects of such networks (connectivity, transport capacity, protocols, etc.) have been studied in the static case, i.e., the case in which the agents are randomly placed but fixed (see, for example, [XK06, H06] and references therein). Also, there has been quite a bit of experimental study dealing with the dynamic situation, i.e., the situation in which the agents are mobile: connections in the network are created and destroyed as the agents move further apart or closer together; see, for example, [AGE02, JBAS03, RMM01]. The paper [GHSZ] also deals with the problem of maintaining connectivity of mobile agents communicating by radio frequency, but from a perspective orthogonal to the one in the present paper. It describes a *kinetic data structure* to maintain the connected components of the union of unit-radius disks moving in the plane.

To the best of our knowledge, ours is the first work in which the dynamic features of such a network are studied theoretically. Moreover, we obtain much sharper results on the static properties than previously obtained for this type of network. In the static case, our results bear a strong resemblance to analogous results for similar problems of random points in continuous spaces (see Holst [H80] and Penrose [P97], [P99]). The

*Received by the editors May 14, 2006; accepted for publication (in revised form) December 20, 2007; published electronically April 18, 2008. An extended abstract of a preliminary version of this paper appeared in STACS [DPSW05].

http://www.siam.org/journals/sidma/22-2/65985.html

[†]Departamento de Llenguatges i Sistemes, Universitat Politecnica de Catalunya, 08034 Barcelona, Spain (diaz@lsi.upc.edu, xperez@lsi.upc.edu, mjserna@lsi.upc.edu). This research was supported by the EU 6th FP under contract 001907 (DELIS) and by the Spanish CICYT project TIC-2001-4917-E. The first author was also supported by the Distinció per a la Recerca de la Generalitat de Catalunya.

[‡]Department of Combinatorics and Optimization, University of Waterloo, Waterloo, ON N2L 3G1 Canada (nickwor@math.uwaterloo.ca). This author was supported by the Canada Research Chairs program and NSERC.

study of the behavior of multiple, simultaneous random walks has its owns merits and could have further applications in computer science.

We propose what we call the *walkers model*, defined as follows. A connected graph $G = (V, E)$ with $|V| = N$ is given, as well as a number w of *walkers* (agents). Also given is a "distance" d. A set W of walkers, with $|W| = w$, is placed randomly and independently on the vertices of G (a vertex may contain more than one walker). Each walker has a range d for communication; that is, two walkers w_1 and w_2 can communicate in one hop if the distance, in G, between the position of the walkers is at most d. Two walkers can communicate if they can reach each other by a sequence of such hops. In addition, each walker takes an independent standard random walk on G, i.e., moves at each time step to a neighboring vertex, each neighbor chosen with equal probability.

The interesting features of the walkers model are encapsulated by the *random graph of walkers*, $\mathcal{W}(G, w, d)$. This random graph is generated using a random assignment $f : W \to V$ of walkers into the vertices of G. The vertices of $\mathcal{W}(G, w, d)$ are the vertices in G that receive at least one walker, two vertices in $\mathcal{W}(G, w, d)$ being joined by an edge iff they are at distance at most d in G. We are interested in the probability of $\mathcal{W}(G, w, d)$ being connected, or in the number of components and their sizes, with certain mild asymptotic restrictions on w and d.

Our primary goal with the walkers model is to characterize the dynamics of the connectivity of the communication network as represented by the random graph process $(\mathcal{W}_t(G, w, d))_{t\geq 0}$. Here $\mathcal{W}_t(G, w, d)$ denotes the graph of walkers at time $t = 0, 1, \ldots$, and, at each time step, each walker simultaneously moves one step to a randomly selected neighbor vertex in G. In order to study this model, we first examine $\mathcal{W}(G, w, d)$, which we call the *static* model. This is a snapshot of the process at one point in time: we are interested in the distribution of the number of components, as well as some other information which helps to answer the dynamic questions. In particular, we are interested in studying the birth and death of components and the sudden connection and disconnection of $\mathcal{W}_t(G, w, d)$. We generally abbreviate $\mathcal{W}(G, w, d)$ to $\mathcal{W}(G)$ when w and d are understood and, similarly, $\mathcal{W}_t(G, w, d)$ to $\mathcal{W}_t(G)$.

We consider a sequence of graphs G with increasing numbers of vertices N for N tending to infinity. The parameters w and d are functions of N. We restrict our attention to the case $w \to \infty$ in order to avoid considering small-case effects. Of course we take $d \geq 1$. We make further restrictions on w and d in order to rule out noninteresting cases, such as values of the parameters in which the network is a.a.s. disconnected or a.a.s. connected. (Throughout this paper, a.a.s. will abbreviate "asymptotically almost surely," which denotes "with probability tending to 1 as $N \to \infty$.") We study the walkers model for two particular sequences of graphs G: the cycle C_N of length N and the $n \times n$ toroidal grid T_N of size $N = n^2$. (In the case of the grid, we use the ℓ^p distance for any $1 \leq p \leq \infty$.) The two cases have an essential difference that prevents a unified treatment: for the interesting values of w and d, disconnectedness of the graph of walkers for the cycle is basically caused by the presence of at least two large "gaps" between the walkers around the cycle, whereas for the grid, it is caused by the presence of one or more isolated walkers.

We now give a description of some of our main results. Throughout this paper, ϱ denotes w/N, which is the expected number of walkers at a vertex.

1.1. The cycle. First we consider the cycle C_N on N vertices. We will show that if $d = \Omega(N)$, then $\mathcal{W}(C_N)$ is a.a.s. connected, so all results here refer to $d = o(N)$.

To study connectedness in this case, we introduce the concept of a *hole*. There is a *hole* between two vertices u and v if u and v each contain at least one walker, but no vertex in the clockwise path from u to v contains a walker. We say that such a hole *follows* u or that u is the *start vertex* of the hole. The number of internal vertices in a hole is its *exact size*. A *k-hole* is a hole whose exact size is at least k.

Let H be the random variable counting the number of d-holes in the walkers model for C_N. Notice that at least two d-holes are needed to disconnect the walkers on C_N. To be precise,

$$\mathcal{W}(C_N) \text{ is connected iff } H \leq 1. \tag{1}$$

We define a new parameter $\mu = N(1-e^{-\varrho})\,e^{-d\varrho}$, which plays a key role in characterizing the connectedness of $\mathcal{W}(C_N)$. It is straightforward to check that

$$\mu \sim \begin{cases} we^{-d\varrho} & \text{if } \varrho \to 0, \\ N(1-e^{-\varrho})\,e^{-d\varrho} & \text{if } \varrho \to c, \\ Ne^{-d\varrho} & \text{if } \varrho \to \infty. \end{cases}$$

Here, and throughout the paper, asymptotics refer to $N \to \infty$ unless otherwise stated.

Regarding the behavior of H, and connectedness, we have the following.

THEOREM 1.1. *The expected number of holes satisfies* $\mathbf{E}[H] \sim N(1-e^{-\varrho})(1-d/N)^w$. *Furthermore,*

(i) *if* $\mu \to 0$*, then* $\mathcal{W}(C_N)$ *is connected a.a.s.,*
(ii) *if* $\mu \to \infty$*, then* $\mathcal{W}(C_N)$ *is disconnected a.a.s., and*
(iii) *if* $\mu = \Theta(1)$*, then H is asymptotically Poisson with mean* μ.

The following corollary gives the asymptotic probability that $\mathcal{W}(C_N)$ is connected. It follows immediately from the theorem in view of (1).

COROLLARY 1.2. $\mathbf{P}(\mathcal{W}(C_N)$ *is connected*$) = e^{-\mu}(1+\mu) + o(1)$.

We now turn to dynamic properties of the random graph of walkers on the cycle. One of our main end results concerns the expected time that the graph of walkers remains disconnected after a point in time at which it becomes disconnected. Define a *disconnected period* to be a maximal sequence of consecutive time steps for which the graph of walkers $\mathcal{W}_t(C_N)$ is disconnected. We define for $t > 0$

$$LD_t = \begin{cases} \max\{k \in \mathbb{N} \mid \mathcal{W}_t(C_N), \dots, \mathcal{W}_{t+k-1}(C_N) \text{ disconnected}\} & \text{if } \mathcal{W}_{t-1}(C_N) \text{ is connected,} \\ 0 & \text{otherwise.} \end{cases}$$

This is the random variable counting the length of the disconnected period starting at time step t provided that it really starts then. Formally we define the *average length* of a disconnected period starting at time t to be

$$LD_{\text{av}} := \mathbf{E}[LD_t \mid LD_t > 0].$$

By symmetry, this is independent of t and so is a function of N, d, and w. The next result finds its size.

THEOREM 1.3. *For the walkers model on the cycle C_N, the average length of a disconnected period of* $(\mathcal{W}_t(C_N))_{t\geq 0}$ *satisfies*

$$LD_{\text{av}} \sim \begin{cases} 2\frac{e^{\mu}-1-\mu}{\mu^2}\varrho^{-1} & \text{if } \varrho \to 0, \\ \frac{e^{\mu}-1-\mu}{1+\mu-(1+\mu+\lambda+\lambda^2)e^{-\lambda}} & \text{if } \varrho \to c, \\ \frac{e^{\mu}}{1+\mu} & \text{if } \varrho \to \infty, \end{cases}$$

where $\lambda = (1 - \frac{3e^{-\varrho} - e^{-\frac{3}{2}\varrho}}{1+e^{-\frac{1}{2}\varrho}})\mu$. *Here* $0 < \lambda < \mu$ *for* $\varrho \to c$.

To convey a feeling of the complexity of the question of how long a network remains disconnected once it becomes so, we introduce the following *train paradox*. A visitor to Barcelona goes every morning to Catalunya station, where there are two metro lines meeting (red and green), though at different levels. He wishes to measure the average length of a train. Since there is plenty to see in Barcelona, each morning he chooses either the red line or the green line at random. He waits for the first train to leave on that line and records its length. He finds after many days that the average length recorded is 9 cars. But he notices that, restricted to the days that the train is already at the platform when he arrives, the average length is only 8 cars.

Could it be that the shorter trains wait longer for him? No, because the trains in Barcelona stop at stations for equal times. Moreover, on any given line they arrive regularly at equally spaced intervals, so the well-known bus paradox does not directly apply.

Which is a better measure—the length of the first train to arrive or the length of trains arriving at a prescribed time? The former, yielding the answer 9, might seem more natural at first. However, the explanation for the differing answers reveals the other to be meaningful, and perhaps even more so. The data above, in both versions of the paradox, arise if the red line has trains of average length 12 arriving every 10 minutes, and the green line has trains of average length 6 arriving at 5 minute intervals. In an extended time period, recording all the lengths of trains on all lines will yield 8 as the average.

Returning to the walkers model, as we shall see in section 3, when N is even there are many different configurations of walkers that cannot arise from a given initial configuration. As a Markov chain, the process is not ergodic; each configuration belongs to a class of mutually reachable configurations. The different classes of states correspond to different metro lines in the train paradox. The quantity LD_{av} in Theorem 1.3 is roughly equivalent to our traveler's measurement of the length of trains restricted to those days that a train is just arriving. However, the train paradox shows that this is not the only reasonable measure of average length. Moreover, the situation is even more complicated, as the train paradox would be if the intertrain time periods were variable on a given line. The average length of the first train to arrive would then be affected by any dependence between the length of a train and the time before the previous train. We wish to study the analogue of the average length of trains on a given metro line—the average length of a period of disconnection given the initial state. If N is divisible by 2 and the initial configuration is conditioned upon, the walkers process is "locked in" to a future in which the (conditional) average length of connected periods may be different from LD_{av}. However, we show that for almost all initial configurations this average is essentially asymptotically equal to LD_{av}. For $T \in \mathbb{Z}^+$, we define the *average disconnection time* of the graph of walkers in $[1, T]$ to be

$$LD_T = \frac{\sum_{t=1}^{T} LD_t}{|\{t \in \{1, \dots, T\} : LD_t > 0\}|}.$$

We show that LD_T converges in probability as $T \to \infty$ to a random variable which may depend on the class of the initial state but nothing else. (Actually, the value is in general different for different classes.) The notation $f \sim g$ a.a.s. used in the following theorem denotes that for all $\epsilon > 0$, a.a.s. $|f/g - 1| < \epsilon$ (see, for example, [W04]).

THEOREM 1.4. *For the walkers model on* C_N, LD_T *converges in probability as*

$T \to \infty$ (with N fixed) toward a random variable LD, uniquely determined almost everywhere. Furthermore, (as $N \to \infty$) we have $LD \sim LD_{\mathrm{av}}$ a.a.s.

Intuitively, LD is the average length of the disconnected periods appearing in all the trajectories of the random process starting from the initial state.

In section 3.2 we give analogous results for maximal connected periods of $\mathcal{W}(C_N)$ (see Theorems 3.10 and 3.11).

During our examination of the cycle, we derive an incidental result (Corollary 3.7) on the expected time required before a certain infinite set of random walkers sprinkled on the nonnegative integers meets another such set sprinkled on the negative integers.

1.2. The grid. We turn now to the toroidal grid T_N with $N = n^2$ vertices, for which our results apply with any normed ℓ^p distance, for $1 \le p \le \infty$. We will show that if $d = \Omega(n)$, then $\mathcal{W}(T_N)$ is a.a.s. connected, so all results here refer to $d = o(n)$. We define h to be the number of grid points within distance d of any fixed point in T_N. This depends on the particular metric chosen, but in all cases it is easy to see that $h = \Theta(d^2)$.

We call a component of $\mathcal{W}(T_N)$ *simple* if it is formed by only one isolated vertex. Let X be the number of simple components. To examine the connectedness of $\mathcal{W}(T_N)$, we need to redefine the μ used for the cycle. For the following grid results, $\mu = N\left(1 - e^{-\varrho}\right) e^{-h\varrho}$. Hence

$$\mu \sim \begin{cases} we^{-h\varrho} & \text{if } \varrho \to 0, \\ N\left(1 - e^{-\varrho}\right) e^{-h\varrho} & \text{if } \varrho \to c, \\ Ne^{-h\varrho} & \text{if } \varrho \to \infty. \end{cases}$$

THEOREM 1.5.

(i) *For $\mu \to \infty$, $\mathcal{W}(T_N)$ is disconnected a.a.s.*

(ii) *For $\mu = \Theta(1)$, a.a.s. all but one component of $\mathcal{W}(T_N)$ are simple, and the number X of simple components is asymptotically Poisson with expected value μ.*

(iii) *For $\mu \to 0$, $\mathcal{W}(T_N)$ is connected a.a.s.*

The theorem immediately gives the following.

COROLLARY 1.6. $\mathbf{P}\left(\mathcal{W}(T_N) \textit{ is connected}\right) = e^{-\mu} + o(1)$.

For dynamic considerations, we also need a function $b = \Theta(d)$ which is related to the boundary of a ball of radius d in the grid. Like h, this function also depends on the metric. It is defined in section 4.2.

We define LD_{av} and LD_T for disconnected periods as for the cycle case above.

THEOREM 1.7. *For the walkers model on the grid T_N with $d > 1$,*

$$LD_{\mathrm{av}} \sim \begin{cases} \frac{e^{\mu}-1}{\mu b \varrho} & \text{if } d\varrho \to 0, \\ \frac{e^{\mu}-1}{1-e^{-\lambda}} & \text{if } d\varrho \to c, \\ e^{\mu} & \text{if } d\varrho \to \infty, \end{cases}$$

where $\lambda = \left(1 - e^{-b\varrho}\right)\mu$. Here $0 < \lambda < \mu$ for $d\varrho \to c$. Furthermore, LD_T converges in probability for T growing large (N fixed) toward a random variable LD, uniquely determined almost everywhere, where $LD \sim LD_{\mathrm{av}}$ a.a.s.

The case $d = 1$ is excluded for this result because of a technical complication that will be explained in the proofs, though our basic approach could be modified to deal with this case also. We also obtain exactly the same result as in Theorem 1.7 for connected periods as opposed to disconnected periods.

The remainder of the paper gives proofs of these theorems as well as stating and proving related ones. In section 2 we give basic definitions and technical lemmas to be used throughout the paper. In section 3 we deal with the cycle C_N, and in section 4 the toroidal grid T_N. One of the main differences between this case and the cycle is the need for a geometric lemma that may be of independent interest (Lemma 4.3). This bounds from below the number of nonoccupied vertices at distance at most d from the boundary of any connected component in $\mathcal{W}(T_N)$, as a function of a measure of the length of the boundary. The last section contains some discussion and related problems.

2. General definitions and basic results. We begin with some definitions and results which are common for all G. We call a component *simple* if it is formed by only one isolated vertex. Recall that $\varrho = w/N$ is the expected number of walkers at a vertex. For any $v \in V$, define h_v to be the number of vertices in G at distance at most d from v, and define $h = \min_{v\in V} h_v$. We say that a vertex or set of vertices is *empty of walkers (e.o.w.)*, or simply *empty*, if it contains no walkers, and *occupied* otherwise. Note that there must be h empty vertices in G within distance d of a simple component.

By considering the coupon collector's problem, we observe that if $w = N\log N + \omega(N)$, then $\mathcal{W}(G, w, d)$ is trivially a.a.s. connected due to every vertex being occupied. Moreover, for the graphs G which we consider in this paper, if $h \in \Omega(N/\sqrt{w})$, then $\mathcal{W}(G, w, d)$ is a.a.s. connected as well. This last claim will be seen in Observations 3.1 and 4.1. Thus, we consider throughout the paper $w < N\log N + O(N)$ and $h = o(N/\sqrt{w})$. In fact, our proofs will just assume h to be $o(N)$. Note that, for the cycle, $h = 2d$.

We will need to compute the probability of certain configurations of walkers at two consecutive time steps t and $t+1$ in order to record the event that walkers jump to the appropriate place at that step. There is a convenient way to formulate this in terms of occupancy of directed edges. Let us regard G as a directed graph by considering each edge as a symmetric pair of directed edges. For any directed edge e going from vertex u to v, we say that a walker is on e between time steps t and $t+1$ if the walker is on u at time step t and jumps onto v within one step. Note that in this way we can encode dynamic transitions in terms of static configurations of walkers.

There is an alternative formulation in terms of cells. Each vertex is divided into as many cells as its degree, and each cell is associated with one of the directed edges stemming from the vertex. Then, the transition of the system between two consecutive time steps can be described by the placement of the walkers in the cells (see Figure 1). We use this representation particularly in our figures for the sake of simplicity and visual clarity.

Assign *size* 1 to all vertices in G. For a given directed edge stemming from a vertex v with degree δ_v, its *size* will be $1/\delta_v$. Given a set A of vertices or directed edges, we define the *size* of A to be the sum of the sizes of its elements.

The following technical lemma is used in most of our proofs to compute the probability of having a certain configuration of walkers (in terms of vertex or directed edge occupancy) at a given time step. The lemma applies to any graph G. Its proof, using an inclusion-exclusion argument, is fairly straightforward and so is omitted. (See [DPSW06] for details.)

LEMMA 2.1. *Let $A_0, \dots, A_m$ be pairwise disjoint sets of vertices (or directed edges) in G, with sizes $S_0, \dots, S_m$, respectively. Let $N = |V(G)|$. If $\sum_{i=0}^{m} S_i = o(N)$,*

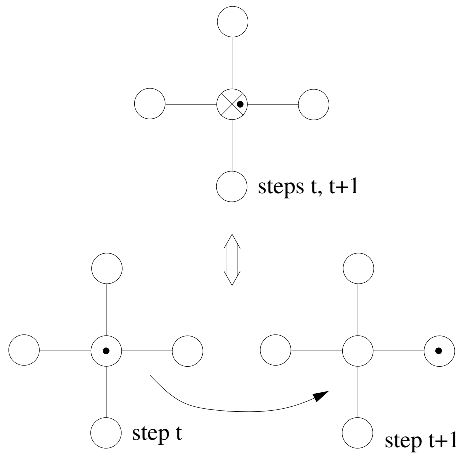

FIG. 1. *The walker jumps to a neighbor according to which cell it is placed on.*

then

$$\mathbf{P}\left(A_0 \text{ e.o.w. } \wedge \bigwedge_{i=1}^{m}(A_i \text{ not e.o.w.})\right) \sim \left(1-\frac{S_0}{N}\right)^w \prod_{i=1}^{m}\left(1-e^{-S_i\varrho}\right).$$

To cover large sizes S, not necessarily $o(N)$, we need the following variation.

LEMMA 2.2. *Let A be a set of vertices in G of size S, and $v_1,\ldots,v_m$ vertices not in A, with $m \geq 1$. Assume $N-S\to\infty$. The probability that no vertex in A is occupied and $v_1,\ldots,v_m$ are all occupied is at most $p_0p^{m-1}\alpha^w(1+o(1))$, where $p_0 = 1-e^{-\varrho/\alpha}$, $\alpha = 1-S/N$, and $p=\min(1,\varrho/\alpha)$. Here the asymptotics are uniform over all m.*

Proof. The probability of the event $\mathcal{E}$ that all walkers avoid A is α^w. The probability that v_1 is occupied conditional upon $\mathcal{E}$ is $1-(1-(N-S)^{-1})^w$, which is asymptotic to p_0. The lemma follows for $\varrho > \alpha$, i.e., $w > N-S$. Otherwise, conditional upon $\mathcal{E}$ and the event that $v_1,\ldots,v_i$ are occupied, the probability that the next is occupied is clearly decreasing with i and is thus at most $w/(N-S) = \varrho/\alpha$. The lemma follows. ∎

3. The cycle.

Here $G = C_N$, the cycle with N vertices.

OBSERVATION 3.1. *Cover C_N with $\lceil \frac{N}{\lceil d/2\rceil}\rceil$ paths of $\lceil d/2\rceil$ vertices. If $d = \Omega(N/\sqrt{w})$, then the probability that some path is e.o.w. is at most*

$$\left\lceil \frac{N}{\lceil d/2\rceil}\right\rceil \left(1-\frac{\lceil d/2\rceil}{N}\right)^w \leq O(\sqrt{w})e^{-\Omega(\sqrt{w})} \to 0.$$

Thus, a.a.s. each of these paths is occupied (by at least one walker), and $\mathcal{W}(G,w,d)$ is connected.

In view of this observation, we assume for the rest of the section that $d = o(N)$. If $d = \Omega(N)$, then $\mathcal{W}(G,w,d)$ is a.a.s. connected.

Recall the definition of a hole from section 1 and that H is the random variable counting the number of d-holes. For this section, μ is defined as for Theorem 1.1.

3.1. Static properties. Here, we study the connectedness of the graph of walkers $\mathcal{W}(G,w,d)$ in the static situation by analyzing the behavior of H. In view of (1), if $\mathbf{E}[H] \to 0$, then $\mathcal{W}(G,w,d)$ is a.a.s. connected. Theorem 1.1 gives the asymptotic distribution of H.

Proof of Theoerem 1.1. For any vertex v in C_N, let H_v be an indicator random variable such that $H_v = 1$ iff there is a d-hole following vertex v. Then,

$$H = \sum_{v\in V} H_v \qquad \text{and} \qquad \mathbf{E}[H] = \sum_{v\in V} \mathbf{P}(H_v = 1). \tag{2}$$

We shall compute the kth factorial moment:

$$\mathbf{E}[H]_k = \sum_{v_1\neq\cdots\neq v_k\in V} \mathbf{P}(H_{v_1} = 1 \wedge \cdots \wedge H_{v_k} = 1). \tag{3}$$

Let $\mathcal{S}$ denote the set of tuples such that each v_i and v_j, $i \neq j$, have distance at least $d+1$ around the cycle. For $(v_1,\dots,v_k) \notin \mathcal{S}$, the probability in (3) is 0 since one of the v_i "lies in" the hole following some v_j, and yet v_i must be occupied. For $(v_1,\dots,v_k) \in \mathcal{S}$, the probability in (3) is easily computed by applying Lemma 2.1:

$$\mathbf{P}(H_{v_1} = 1 \wedge \cdots \wedge H_{v_k} = 1) \sim \left(1 - \frac{kd}{N}\right)^w \left(1 - e^{-\varrho}\right)^k.$$

Since $d = o(N)$, we have $|\mathcal{S}| \sim N^k$, and thus from (3)

$$\mathbf{E}[H]_k \sim \left[N\left(1 - e^{-\varrho}\right) e^{-d\varrho + O(\frac{d^2w}{N^2})}\right]^k. \tag{4}$$

In particular,

$$\begin{aligned}\mathbf{E}[H] &\sim N\left(1 - e^{-\varrho}\right)\left(1 - \frac{d}{N}\right)^w \\ &\sim N\left(1 - e^{-\varrho}\right) e^{-d\varrho + O(\frac{d^2w}{N^2})}.\end{aligned} \tag{5}$$

In the case $\mu \to 0$, we have also $\mathbf{E}[H] \to 0$ since $(1 - d/N)^w \le e^{-d\varrho}$. Then $\mathbf{P}(H = 0) \to 1$, and $\mathcal{W}(G,w,d)$ is connected a.a.s.. In the case that μ is bounded away from 0, taking logarithms,

$$d\varrho = \log N\left(1 - e^{-\varrho}\right) - \log\mu. \tag{6}$$

Considering separately the cases when $\varrho \to 0$ and when $\varrho \to c, \infty$, we get from (6) that $\frac{d^2w}{N^2} = o(1)$, so we can ignore the term $O(\frac{d^2w}{N^2})$ in the expression (4) and obtain

$$\mathbf{E}[H]_k \sim \left[N\left(1 - e^{-\varrho}\right) e^{-d\varrho}\right]^k = \mu^k. \tag{7}$$

Moreover, if μ is bounded, then it follows, from (7) and (for instance) from Theorem 1.23 in [Bol02], that H is asymptotically Poisson. For $\mu \to \infty$, we have that $\mathbf{E}[H]_2 \sim \mu^2$, and so it follows from Chebyshev's inequality that a.a.s. $H > \mu/2$ (and $\mathcal{W}(G,w,d)$ is disconnected). $\square$

3.2. Dynamic properties. Assume that from an initial random placement f of the walkers, at each step, every walker moves from its current position to one of its neighbors, with probability $1/2$ of going either way. This is a standard random walk on the cycle for each walker. To study the connectivity properties of the dynamic graph of walkers we need to introduce some notation.

A *configuration* or *state* is an arrangement of the w walkers on the vertices of C_N. Consider the graph of configurations, where the vertices are the N^w different configurations. Each configuration can be represented by a vector $\mathbf{a} = (a_1, \ldots, a_w) \in (\mathbb{Z}_N)^w$, where a_i indicates the label of the vertex being occupied by walker i. Given a configuration $\mathbf{a}$, there exists an edge between $\mathbf{a}$ and all configurations $(a_1 \pm 1, \ldots, a_w \pm 1)$. Thus, any configuration has 2^w neighbors, and the relationship of being neighbors is symmetric. The dynamic process can be viewed as a random walk on the graph of configurations, in particular, a Markov chain $\mathcal{M}(N, w, d)$. We denote by $\mathcal{M}_t$ the state of the process at time step t.

For N even, given any two configurations $\mathbf{a}$ and $\mathbf{b}$, we say that they have the *same parity* if for all i and j, $a_i - a_j \equiv b_i - b_j \pmod 2$. There are 2^{w-1} different parities. Note that the initial parity stays invariant during the dynamic process. The following lemma is straightforward, and the proof is left to the reader.

LEMMA 3.2. *Let* $\mathbf{a}$ *and* $\mathbf{b}$ *be any two configurations, and let* $h_{\mathbf{a},\mathbf{b}}$ *denote the hitting time from* $\mathbf{a}$ *to* $\mathbf{b}$. *If* N *is odd, then* $\mathbf{b}$ *is reachable from* $\mathbf{a}$ *and* $h_{\mathbf{a},\mathbf{b}}$ *is finite for any* $\mathbf{a}$ *and* $\mathbf{b}$. *If* N *is even, then* $\mathbf{b}$ *is reachable from* $\mathbf{a}$ *provided that* $\mathbf{a}$ *and* $\mathbf{b}$ *have the same parity, and in this case* $h_{\mathbf{a},\mathbf{b}}$ *is finite.*

Then, if N is odd, $\mathcal{M}(N, w, d)$ consists of a single closed class of states, so it is irreducible and positive recurrent. It is trivial to verify aperiodicity, and hence we have that the chain is ergodic. However, if N is even, there are 2^{w-1} closed classes of states, where each class consists of all configurations with the same parity. We denote by $\mathcal{A}$ the set of classes. Let $A \in \mathcal{A}$ be any class of states, and let $\mathbf{a} \in A$ be a configuration. For this particular configuration, we can partition the set of walkers $W = W_1 \cup W_2$ so that the ones in W_1 lie in odd positions of the cycle and the ones in W_2 lie in even positions. Let A_1 be the set of all states which lead to this same partition, and A_2 the set of those which lead to the complementary one. Clearly, $A = A_1 \cup A_2$. Those states in A_1 are only reachable by an even number of steps from $\mathbf{a}$, and those in A_2 by an odd number of steps. Hence, if we restrict the Markov chain to any class of states, it is irreducible and positive recurrent but 2-periodic and hence not ergodic.

OBSERVATION 3.3. *Note that for any fixed* t, *the distribution of* $\mathcal{W}_t(C_N)$ *is just that of* $\mathcal{W}(C_N)$ *in the static case. That is, the initial uniform distribution stays invariant, even though when* N *is even the chain is not ergodic and there is no limit distribution. Hence, by Theorem* 1.1, *if* $\mu \to 0$ *or* ∞, *then for any fixed* t, $\mathcal{W}_t(C_N)$ *is a.a.s. connected or a.a.s. disconnected, respectively.*

In view of this observation, we assume that $\mu = \Theta(1)$ for the remainder of the section, since we wish to study only the nontrivial dynamic situations. One thing this implies, in view of the definition of μ, is that

$$d\varrho \to \infty. \tag{8}$$

We define $H(t)$ to be the random variable that counts the number of d-holes at time step t. Then from section 3.1, $H(t)$ is asymptotically Poisson with expectation $\mu = \Theta(1)$.

For the dynamic properties of $\mathcal{W}_t(C_N)$, we are interested in the probability that a new d-hole appears at a given time step. Moreover, we require knowledge of this probability conditional upon the number of d-holes already existing.

If there is a d-hole from u to v at time step t and all walkers at u and v move in the same direction on the next step, a new d-hole may appear following one of the neighbors of u (provided no new walkers move in to destroy this). These two d-holes, though different, are related, and the presence of the first makes the second more likely to occur. Similarly, the exact size of a d-hole following u may change in one step, making it technically a different d-hole but, again, related. In all these cases, the start vertex of the d-hole "moves" by at most 1; a d-hole at time step $t+1$ which does not follow u or a neighbor of u is not related to a d-hole following u at time step t. We need some definitions to make this loose description rigorous. Define a *d-hole line* to be a maximal sequence of pairs $(h_1, t_1), \dots, (h_l, t_l)$, where h_i is a d-hole existing at time step t_i for $1 \le i \le l$, and such that $t_i = t_{i-1} + 1$ and the start vertex of h_i is adjacent to, or equal to, the start vertex of h_{i-1} for $2 \le i \le l$. Fix two consecutive time steps t and $t+1$. If $t_1 = t+1$, we say that the line is *born* between t and $t+1$; if $t_l = t$, the line *dies* between t and $t+1$; and if $t = t_i$, $i \in \{1, \dots, l-1\}$, we say that the line *survives* during the interval $[t, t+1]$. Note that the time-reversal of the process has a d-hole line born at vertex u between $t+1$ and t iff the d-hole line dies at u between t and $t+1$.

We define random variables $S(t)$, $B(t)$, and $D(t)$ to be the number of d-hole lines surviving, being born, and dying between t and $t+1$. We have $D(t) + S(t) = H(t)$ and $B(t) + S(t) = H(t+1)$.

The proof of the following result is similar to that of Theorem 1.1 but rather more complicated, so we give only an abbreviated proof here. A full proof may be found in [DPSW06] or [P07]. Here "asymptotically independent Poisson" means that the joint distribution tends to that of independent Poisson variables.

PROPOSITION 3.4. *For any fixed t, the random variables $S(t)$, $B(t)$, and $D(t)$ are asymptotically independent Poisson, with the expectations*

$$\mathbf{E}[S(t)] \sim \begin{cases} \mu & \text{if } \varrho \to 0, \\ \mu - \lambda & \text{if } \varrho \to c, \\ 3\mu e^{-\varrho} & \text{if } \varrho \to \infty, \end{cases} \qquad \text{and} \qquad \mathbf{E}[B(t)] = \mathbf{E}[D(t)] \sim \begin{cases} \frac{1}{2}\mu\varrho & \text{if } \varrho \to 0, \\ \lambda & \text{if } \varrho \to c, \\ \mu & \text{if } \varrho \to \infty, \end{cases}$$

where λ is defined in Theorem 1.3.

Proof. We will estimate factorial moments of the appropriate variables. In C_N, let right denote "clockwise" and left "counterclockwise." Moreover, for a vertex $v \in V$ and $i \ge 0$, let $v + i$ (respectively, $v - i$) denote the vertex i positions to the right (respectively, left) from v. All probabilities and events in this proof will involve two consecutive time steps t and $t+1$. We can describe these events from a static point of view, as explained in section 2.

We consider births in detail. There are four ways that a d-hole line can be born at v between time steps t and $t+1$ according to the following descriptions, as shown in Figure 2:

b1. At time step t, there is a hole between $v+1$ and $v+d$ of exact size $d-2$. Then all walkers at v and $v+1$ move left and all walkers at $v+d$ and $v+d+1$ move right.

b2. At time step t, there is a hole between $v+1$ and $v+d+1$ of exact size $d-1$. Then all walkers at v and $v+1$ move left and all walkers at $v+d+1$ and $v+d+2$ move right.

b3. At time step t there is a hole between $v+1$ and $v+d+1$ of exact size $d-1$, and $v+d+2$ is occupied. Then all walkers at v and $v+1$ move left, all walkers at $v+d+1$ move right, and at least one walker at $v+d+2$ moves left.

b4. At time step t there is a hole between v and $v+d$ of exact size $d-1$, and $v-1$ is occupied. Then all walkers at v move left, all walkers at $v+d$ and $v+d+1$ move right, and at least one walker at $v-1$ moves right.

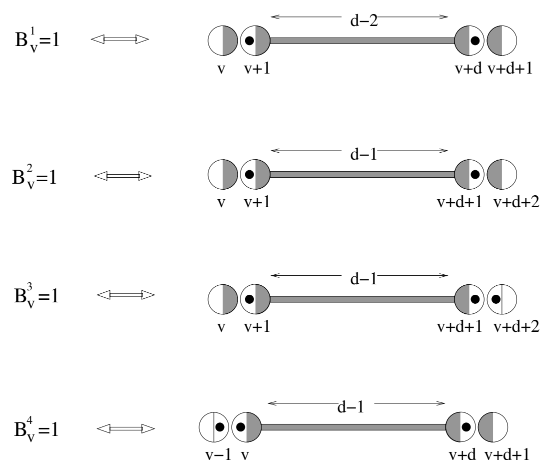

FIG. 2. *Birth of a d-hole line at vertex v.*

The expected number of d-hole lines born at a given time is N times the sum of the probabilities of these four events. Applying Lemma 2.1, together with the fact that $(1-d/N)^w \sim e^{-d\varrho}$ (Theorem 1.1), we get (writing the four terms separately)

$$\frac{1}{N}\mathbf{E}\left[B(t)\right] \sim \left(1-e^{-\varrho/2}\right)^2 e^{-d\varrho} + \left(1-e^{-\varrho/2}\right)^2 e^{-(d+1)\varrho}$$
$$+ \left(1-e^{-\varrho/2}\right)^3 e^{-(d+1/2)\varrho} + \left(1-e^{-\varrho/2}\right)^3 e^{-(d+1/2)\varrho}.$$

This leads to the expression for $\mathbf{E}\left[B(t)\right]$ stated in the theorem. The other expected values are obtained in a similar way (with three types of survival event and four types of death event). Computing the joint factorial moments of random variables counting births, deaths, and survivals gives the claim about Poisson distribution, similar to the proof of Theorem 1.1. We omit the remaining details. ∎

As $H(t) = S(t) + D(t)$ and $H(t+1) = S(t) + B(t)$, the following is immediate.

COROLLARY 3.5. *$H(t)$ and $B(t)$ are asymptotically independent, and so are $D(t)$ and $H(t+1)$.*

It is natural to define the *lifespan* of a d-hole line as the number of time steps for which the line is alive. For any vertex v and time step t, the random variable $L_{v,t}$ is the lifespan of the d-hole line born at vertex v between time steps t and $t+1$. If no such birth takes place, $L_{v,t}$ is defined to be 0. Note that the random variables $L_{v,t}$ are identically distributed for all v and t.

In view of (8), it is easy to see that a state can be reached in which there are no holes. One way to do this is to force the walkers to move to positions in which they are almost equally spaced around the cycle. Then, by Lemma 3.2, for any initial state, the process will reach some such state within finite expected time. Thus the expected lifespan of any given d-hole line, given the configuration of walkers at its birth, is bounded (and the bound is simply a function of N, d, and w).

Considering the train paradox discussed in section 1, we define the *average lifespan* of d-hole lines to be the expected time that a d-hole line will survive once born. Formally,

$$L_{\mathrm{av}} := \mathbf{E}\left[L_{v,t} \mid L_{v,t} > 0\right].$$

By symmetry, this is independent of v and t and so is a function of N, d, and w. The next result finds its size.

THEOREM 3.6. *The average lifespan of d-hole lines satisfies*

$$L_{\mathrm{av}} = \frac{\mathbf{E}\left[H(0)\right]}{\mathbf{E}\left[B(0)\right]} \sim \begin{cases} 2\varrho^{-1} & \text{if } \varrho \to 0, \\ \frac{\mu}{\lambda} & \text{if } \varrho \to c, \\ 1 & \text{if } \varrho \to \infty, \end{cases}$$

where λ is defined as in Theorem 1.3.

Proof. For any vertex v and time step t, we have

$$\begin{aligned} N\mathbf{E}\left[L_{v,t}\right] - \mathbf{E}\left[H(0)\right] &= \sum_{v \in V} \mathbf{E}\left[L_{v,t}\right] - \mathbf{E}\left[H(0)\right] \\ &= \frac{1}{T}\left(\sum_{v \in V}\sum_{t=0}^{T-1} \mathbf{E}\left[L_{v,t}\right] - \sum_{t=0}^{T-1} \mathbf{E}\left[H(t)\right]\right) \quad \text{for any } T \in \mathbb{Z}^{+} \\ &= \frac{1}{T}\mathbf{E}\left[\sum_{v \in V}\sum_{t=0}^{T-1} L_{v,t} - \sum_{t=0}^{T-1} H(t)\right] \\ &= \lim_{T \to \infty} \frac{1}{T}\mathbf{E}\left[\sum_{v \in V}\sum_{t=0}^{T-1} L_{v,t} - \sum_{t=0}^{T-1} H(t)\right] \\ &= 0, \end{aligned}$$

since $H(t)$ is uniformly bounded (the bound depending on N) and so is the expected lifespan of any line. Thus, the portions of lifespans omitted in $\sum_{t=0}^{T-1} H(t)$ are bounded in expectation.

Now the rest follows easily:

$$\begin{aligned} \mathbf{E}\left[L_{v,t} \mid L_{v,t} > 0\right] &= \frac{\mathbf{E}\left[L_{v,t}\right]}{\mathbf{P}\left(L_{v,t} > 0\right)} \\ &= \frac{\mathbf{E}\left[H(0)\right]}{N\mathbf{P}\left(L_{v,t} > 0\right)} \\ &= \frac{\mathbf{E}\left[H(0)\right]}{\sum_{v} \mathbf{P}\left(L_{v,t} > 0\right)} \\ &= \frac{\mathbf{E}\left[H(0)\right]}{\mathbf{E}\left[B(0)\right]}. \end{aligned}$$

The final formula for this expression comes immediately from Proposition 3.4, noting that $H(t) = S(t) + D(t) \sim S(t)$ for $\varrho \to 0$ and $D(t)$ dominates for $\varrho \to \infty$. ∎

The previous theorem has the following consequence of a rather different nature. Although this is entirely incidental to our main goals, it seems interesting enough to mention. It is not clear how to establish this result without reference to our result on the cycle.

COROLLARY 3.7. *For constant* $0 < p < 1$*, randomly place walkers on the integers, each independently with probability* p*, but with probability* 1 *for integers* 0 *and* 1*. Let them all take the usual independent random walks. Then the expected time* $g(p)$ *before one of the initially nonpositive walkers meets or passes one of the initially positive walkers satisfies* $g(p) \sim 2/p$ *as* $p \to 0$*.*

Proof. We call the model introduced in the statement of the theorem the p-model. For fixed p, let $m = \lceil p^{-1} \rceil$. For fixed $i > 0$, consider the intervals $I_{-i} = [-im, -(i-1)m - 1]$ and $I_i = [(i-1)m + 1, im]$. Let H_i denote the event that in the first i^2m^2 steps, some walker initially in I_i meets some walker initially in I_{-i}. Then $\mathbf{P}(H_i) > c > 0$ for some constant c independent of i (and p). This is because the probability that at least one site in I_i is initially occupied is at least $1 - (1 - p)^m$, which is bounded below, and because after i^2m^2 steps a walker's position is binomially distributed, subject to parity and relative to its starting point, with variance i^2m^2, not forgetting that the walkers move independently.

Call a walker *positive* if it is initially positive and *negative* otherwise. For fixed M, the events $H_1, H_2, \dots, H_M$ are mutually independent, so the probability that none of them occurs is at most $(1 - c)^M$. Thus the probability that none of the negative walkers meets or passes one of the positive walkers in the first Mi^2m^2 steps is at most $(1 - c)^M$, i.e., exponentially small in M. It follows that $g(p)$ is finite and that a.a.s. the number of steps before the first positive walker passes the first negative walker is bounded in probability.

Still for fixed p, define a sequence of integers $d = d(N)$ satisfying $d = p^{-1} \log N + O(1)$, where the $O(1)$ term ensures that such an integer sequence can be found. Define $\tilde{g}(p)$ analogously to $g(p)$ but restricted to the walkers initially occupying the set $\mathbb{Z}(d) = \{k : |k| \leq d\}$. By the observation at the end of the previous paragraph, $\tilde{g}(p) \sim g(p)$.

Next consider our standard model of w walkers on a cycle of length $N + d$, where $w = \lfloor pN \rfloor$. Let $g_1(p)$ denote the expected lifespan of a d-hole line conditional upon it being born of size d. This is clearly independent of the time and place of birth, but let us fix the birth at a vertex v between time steps t and $t+1$. All possible configurations that are feasible at time steps t and $t + 1$ and such that this birth occurs are equally likely. At time step t there is a walker at v and one at $v + d + 1 \pmod N$, with none in between. For any such configuration of walkers at time step $t + 1$, we can map the part of the configuration near v (to its "left") and near $v + d + 1$ (to its "right") onto the integers, taking $[v - d, v]$ onto the integers $[-d, 0]$ and $[v + d + 1, v + 2d]$ onto $[1, d]$. Let V_d denote the set of sites mapped in this way to $\mathbb{Z}(d)$, and let f denote the mapping.

For the rest of the proof we omit some technical details. For any particular configuration in V_d, the probability that it arises is close to the probability of its image under f occurring in the p-model. In fact, it tends toward the same probability as $N \to \infty$. Similarly, the probability that two walkers occupy the same position within $\mathbb{Z}(d)$ or V_d will tend to 0 as $N \to \infty$ (in both cases). Hence $g_1(p) \to g(p)$ as $N \to \infty$.

Now let $g_2(p)$ denote the expected lifespan of a d-hole line conditional upon it being born of size $d + 1$. The same argument as above gives $g_2(p) \to g(p)$ as $N \to \infty$.

It follows that for the walkers on the cycle with these parameters, $L_{\text{av}} \to g(p)$ as $N \to \infty$. Letting $p \to 0$ slowly and applying Theorem 3.6 now gives the theorem. ∎

It is also interesting to ask how the average lifespan relates to the initial configuration of the walkers (see the train paradox in section 1). As noted earlier, if N is odd, then $\mathcal{M}(N, w, d)$ is ergodic and parity is immaterial. However, if N is even and the initial configuration is conditioned upon, the (conditional) average lifespan of d-hole lines may be different from $L_{\rm av}$. Define $T \in \mathbb{Z}^+$,

$$L_T = \frac{\sum_{t=0}^{T-1} \sum_{v \in V} L_{v,t}}{|\{(v,t) : L_{v,t} > 0\}|},$$

where the denominator runs over all pairs $(v,t) \in V \times \{0, \ldots, T-1\}$. If the denominator is zero (or the numerator is infinite, which happens with probability 0), the value is immaterial and may be defined as 0. Note that L_T is a function of a given trajectory of the process.

We next aim for an analogue of Theorem 1.4. Define

$$L = \frac{\mathbf{E}\left[H(0)|\Lambda\right]}{\mathbf{E}\left[B(0)|\Lambda\right]},$$

where Λ is the random variable which is the closed class in which the initial state $\mathcal{M}_0$ lies.

THEOREM 3.8. *For the walkers model on C_N, L_T converges in probability as $T \to \infty$ (with N fixed) toward a random variable L. Furthermore, (as $N \to \infty$) we have $L \sim L_{\rm av}$ a.a.s.*

Proof. The proof proceeds by a renewal theory type of argument, showing a sharp concentration of the number of times that the associated Markov chain revisits a given state over a given long time period.

Let us define the truncated average lifespan of d-hole lines in $[0, T-1]$ to be

$$\overline{L_T} = \frac{\sum_{t=0}^{T-1} H(t)}{H(0) + \sum_{t=0}^{T-2} B(t)}$$

(defined by convention to be 0 if the denominator is 0). As we prove below, this is an approximation of L_T.

We first deal with the case that N is even. To prove the result we need to take into account the class of states containing the initial state, since different starting configurations of walkers may lead to different expected numbers of holes and births.

Let Λ be the random variable accounting for the closed class of states where the initial state $\mathcal{M}_0$ lies. Let A be any closed class of states. We condition on the event $(\Lambda = A)$, which we call A with some abuse of notation. By Lemma 3.2, the hitting time between any two states is finite.

We consider the Doob martingale $\Sigma_0, \ldots, \Sigma_T$ defined by

$$\Sigma_i = \mathbf{E}\left[\sum_{t=0}^{T-1} H(t) \,\middle|\, A, \mathcal{M}_0, \ldots \mathcal{M}_{i-1}\right], \quad i = 0, \ldots, T.$$

We have $\Sigma_0 = T\mathbf{E}\left[H(0)|A\right]$ and $\Sigma_T = \sum_{t=1}^{T} H(t)$. (We recall that this last expression is regarded in the probability space induced by restricting to the event A.)

From the fact that the expected hitting time between any two states is finite, we deduce that the differences $|\Sigma_{i+1} - \Sigma_i|$ are bounded above by a constant independent of T. Hence as an immediate consequence of Azuma's inequality we get that, conditional upon the event A,

$$\lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} H(t) = \mathbf{E}\left[H(0)|A\right] \quad \text{in probability} \tag{9}$$

and by a similar argument

$$\lim_{T\to\infty} \frac{1}{T}\left(H(0) + \sum_{t=0}^{T-2} B(t)\right) = \mathbf{E}\left[B(0)|A\right] \quad \text{in probability.} \tag{10}$$

Then, by taking the ratio of (9) and (10), we get that

$$\lim_{T\to\infty} \overline{L_T} = L \quad \text{in probability.} \tag{11}$$

The case that N is odd is easier. There is just one closed class of states, and $\mathcal{M}(N, w, d)$ is ergodic. By the martingale argument as above, we get (noting Theorem 3.6)

$$\lim_{T\to\infty} \overline{L_T} = \frac{\mathbf{E}\left[H(0)\right]}{\mathbf{E}\left[B(0)\right]} = L_{\text{av}} \quad \text{in probability.} \tag{12}$$

Moreover, we obtain

$$\lim_{T\to\infty} L_T - \overline{L_T} = 0 \quad \text{in probability} \tag{13}$$

since $H(t)$ and $B(t)$ are uniformly bounded (the bound depending on N) and so is the expected lifespan of any line. Thus the portions of lifespans omitted in $\overline{L_T}$ are bounded in expectation, which is insignificant since the denominator grows linearly with T, as shown in (10).

In order to finish the proof, it suffices to show that as $N \to \infty$

$$L \sim L_{\text{av}} \quad \text{a.a.s.}$$

We note that the quantity $\mathbf{E}\left[H(0)|A\right]/\mathbf{E}\left[B(0)|A\right]$ may vary depending on the particular closed class of states A. Let us study this in more detail. Let $\mathbf{a} \in A$ be a configuration. As in section 3.2, let us partition the set of walkers W into W_1 and W_2 according to the parity of their positions in the cycle. Let us define the *imbalance* of the configuration as $\Delta(\mathbf{a}) = |w_1 - w_2|$, where $w_i = |W_i|$. It makes sense to define $\Delta(A) = \Delta(\mathbf{a})$ since it does not depend on the choice of a.

We can compute the expectations of S, B, and D conditional upon $(\mathcal{M}_0 \in A)$ by proceeding the same way as in Proposition 3.4. The only difference is that w_1 walkers must go to $\frac{N}{2}$ of the vertices (say, those with odd position) and w_2 must go to the other $\frac{N}{2}$. We omit the details here since they are fairly tedious and completely analogous to the previous computations. We note that these expectations do not depend on the particular partition (W_1, W_2) but only on the imbalance $\Delta(A)$. In all cases we get

$$\frac{\mathbf{E}\left[H(0)|A\right]}{\mathbf{E}\left[B(0)|A\right]} = \Theta\left(\frac{\mathbf{E}\left[H(0)\right]}{\mathbf{E}\left[B(0)\right]}\right). \tag{14}$$

In fact, for $\frac{\Delta(A)}{N} = O(1)$ we have $\mathbf{E}\left[H(0)|A\right] = \Theta\left(\mathbf{E}\left[H(0)\right]\right)$ and $\mathbf{E}\left[B(0)|A\right] = \Theta\left(\mathbf{E}\left[B(0)\right]\right)$. For $\frac{\Delta(A)}{N} \to \infty$, these two statements are no longer true, but the extra factors in the numerator and denominator of (14) cancel out to within a factor of $\Theta(1)$.

However, not all imbalances are equally likely. In fact, for any $\epsilon > 0$, we have

$$\mathbf{P}\left(\Delta(\mathcal{M}_0) \geq w^{\frac{1+\epsilon}{2}}\right) = \mathbf{P}\left(\left|w_1(\mathcal{M}_0) - \frac{w}{2}\right| \geq \frac{1}{2} w^{\frac{1+\epsilon}{2}}\right) \leq \frac{w}{w^{1+\epsilon}} = o(1). \tag{15}$$

Moreover, for a (typical) class A such that $\Delta(A) < w^{\frac{1+\epsilon}{2}}$ and by the method explained above, we get

$$\mathbf{E}\left[H(0)|A\right] \over \mathbf{E}\left[B(0)|A\right] \sim \frac{\mathbf{E}\left[H(0)\right]}{\mathbf{E}\left[B(0)\right]}. \tag{16}$$

From this last fact, together with (14) and (15), the theorem follows. ∎

We turn now to connectivity issues, for which we use (1). The next lemma gives us the probability that there is one component at time step t but at least two at time step $t+1$.

LEMMA 3.9. *The probability that $\mathcal{W}_t(C_N)$ is connected and that $\mathcal{W}_{t+1}(C_N)$ is disconnected is given by*

$$\mathbf{P}\Big(H(t+1) \geq 2 \wedge H(t) < 2\Big) \sim \begin{cases} \frac{1}{2}\mu^2 e^{-\mu}\varrho & \text{if } \varrho \to 0, \\ e^{-\mu}\left(1+\mu-(1+\mu+\lambda+\lambda^2)e^{-\lambda}\right) & \text{if } \varrho \to c, \\ (1+\mu)e^{-\mu}(1-(1+\mu)e^{-\mu}) & \text{if } \varrho \to \infty, \end{cases}$$

where λ is defined as in Theorem 1.3.

Proof. We can split the probability on the left-hand side according to the events $H(t) = 0$ and $H(t) = 1$. Noting that

$$\mathbf{P}\Big(H(t+1) \geq 2 \wedge H(t) = 0\Big) = \mathbf{P}\Big(H(t) = 0 \wedge B(t) \geq 2\Big),$$

and

$$\begin{aligned} &\mathbf{P}\Big(H(t+1) \geq 2 \wedge H(t) = 1\Big) = \mathbf{P}\Big(S(t) + B(t) \geq 2 \wedge S(t) + D(t) = 1\Big) \\ &= \mathbf{P}\Big(S(t) = 1 \wedge B(t) \geq 1 \wedge D(t) = 0\Big) + \mathbf{P}\Big(S(t) = 0 \wedge B(t) \geq 2 \wedge D(t) = 1\Big), \end{aligned}$$

the result follows from Proposition 3.4 and Corollary 3.5 provided $\varrho \not\to 0$. If $\varrho \to 0$, we need, for instance, to estimate the vanishingly small quantity $\mathbf{P}(S(t) = 1 \wedge B(t) \geq 1 \wedge D(t) = 0)$ asymptotically, but the moment calculations, as for the proof of Proposition 3.4, suffice for this, and we omit the routine details. ∎

Recall the definition of a disconnected period from section 1. By Lemma 3.2, from a disconnected state the graph will always reach some connected one (for example, one in which all walkers occupy one of two adjacent sites) within finite expected time. Thus the expected length of any disconnected period is bounded (this bound depending on N).

Proof of Theorem 1.3. Arguing as in the proof of Theorem 3.6, and obtaining the asymptotic expressions for the numerator and the denominator from Theorem 1.1 and Lemma 3.9, we obtain

$$LD_{\mathrm{av}} = \frac{\mathbf{P}\left(H(1) \geq 2\right)}{\mathbf{P}\left(H(0) < 2 \wedge H(1) \geq 2\right)} \sim \begin{cases} 2\frac{e^{\mu}-1-\mu}{\mu^2}\varrho^{-1} & \text{if } \varrho \to 0, \\ \frac{e^{\mu}-1-\mu}{1+\mu-(1+\mu+\lambda+\lambda^2)e^{-\lambda}} & \text{if } \varrho \to c, \\ \frac{e^{\mu}}{1+\mu} & \text{if } \varrho \to \infty. \end{cases} \tag{17}$$ ∎

Proof of Theorem 1.4. Let the random variable $\Phi(t) = 1_{[H(t)\geq 2]}$ be the indicator of the event that $\mathcal{W}_t(C_N)$ is disconnected, and let $\Psi(t) = 1_{[H(t-1)<2 \wedge H(t)\geq 2]}$ be the indicator of the event that a disconnected period starts at time step t.

We define the truncated average disconnection time of the graph of walkers in $[1, T]$ as

$$\overline{LD_T} = \frac{\sum_{t=1}^{T} \Phi(t)}{\sum_{t=1}^{T} \Psi(t)}.$$

Proceeding as in the proof of Theorem 3.8, we get that LD_T converges in probability for large T toward a random variable LD such that $LD \sim LD_{\mathrm{av}}$. ∎

From the definitions in section 1 relating to disconnected periods, by changing disconnected to connected, we define *connected periods* LC_t, LC_{av}, and LC_T. Using

$$LC_{\mathrm{av}} = \frac{\mathbf{P}\left(H(1) < 2\right)}{\mathbf{P}\left(H(0) \geq 2 \wedge H(1) < 2\right)}$$

and an analogue of Lemma 3.9, we obtain the following.

THEOREM 3.10. *The average length of a connected period satisfies*

$$LC_{\mathrm{av}} \sim \begin{cases} 2\frac{1+\mu}{\mu^2}\varrho^{-1} & \text{if } \varrho \to 0, \\ \frac{1+\mu}{1+\mu-(1+\mu+\lambda+\lambda^2)e^{-\lambda}} & \text{if } \varrho \to c, \\ \frac{e^{\mu}}{e^{\mu}-(1+\mu)} & \text{if } \varrho \to \infty, \end{cases}$$

where λ is defined as in Theorem 1.3.

The result for connection times analogous to Theorem 1.4 is the following. The proof is almost identical and so is omitted.

THEOREM 3.11. *LC_T converges in probability as $T \to \infty$ (N fixed) toward a random variable LC, where $LC \sim LC_{\mathrm{av}}$ a.a.s.*

4. The grid. In this section, we study the walkers model for $G = T_N$, the toroidal grid with $N = n^2$ vertices. We can refer to vertices by using coordinates in $\mathbb{Z}_n \times \mathbb{Z}_n$. For the grid we encounter significant new obstacles as compared to the cycle; see, for instance, the Geometric Lemma below.

For any $p \in [1, \infty]$ and any two vertices u and v in T_N, we define the distance $\mathrm{dist}_{\ell^p}(u, v)$ as the minimal ℓ^p distance between any two points u' and v' in the square grid such that the coordinates of u' are congruent to those of u taken modulo n, and similarly for v and v'. Let us fix any such metric for our study of the connectedness of $\mathcal{W}(G, w, d)$ and write $\mathrm{dist}(u, v)$ for short. We use dist to refer to this measure of distance, distinguishing it from Euclidean distance. Note that

$$\frac{1}{2} \cdot \mathrm{dist}_{\ell^1}(u, v) \leq \mathrm{dist}(u, v) \leq \mathrm{dist}_{\ell^1}(u, v). \tag{18}$$

The number of vertices at distance at most d from any given vertex is $h = \Theta(d^2)$. The exact expression of h depends on the choice of the metric. Some examples are found in Table 1.

OBSERVATION 4.1. *Assume that $d < 2n$ (otherwise, the graph of walkers is always complete). For each $i, j < 4n/d$, let v_{ij} denote the vertex in T_N with coordinates $(\lfloor id/4 \rfloor, \lfloor jd/4 \rfloor)$. Let S_{ij} denote the set of grid points closer to v_{ij} than any of the other $v_{i'j'}$. Then there are $\Theta(N/d^2)$ disjoint sets S_{ij} each containing $\Theta(d^2)$ points. The probability that at least one of these S_{ij} is empty of walkers is at most*

$$\Theta(N/d^2)(1 - \Theta(d^2/N))^w = O(\sqrt{w})e^{-\Omega(\sqrt{w})},$$

TABLE 1
Number of vertices at distance at most d from a given vertex.

Metric	h
ℓ^1	$h = 2d(d+1)$
ℓ^2	$h \sim \pi d^2$ if $d \to \infty$
ℓ^∞	$h = 4d(d+1)$

which goes to 0 *if* $d^2 = \Omega(N/\sqrt{w})$*. Thus, a.a.s. each of these pieces is occupied by at least one walker, and* $\mathcal{W}(G, w, d)$ *is connected.*

In view of the observation, we assume for the rest of the section that $h = o(N)$, i.e., $d = o(n)$. If $d = \Omega(n)$, then $\mathcal{W}(G, w, d)$ is a.a.s. connected.

We wish to study the connection and disconnection of $\mathcal{W}(G, w, d)$ in a similar way to the cycle. For the grid, the notion of hole does not help, and we deal directly with components. Recall from the introduction that a simple component is one with just one vertex. These play a major role, and we shall prove that, for the interesting values of the parameters, a.a.s. there exist only simple components aside from one giant one.

Let $\mathcal{C}$ be any given component. The *edges* of $\mathcal{C}$ are the straight edges joining occupied vertices in $\mathcal{C}$ of distance at most d. The associated *forbidden region* $A_{\mathcal{C}}$ is the set of vertices not in $\mathcal{C}$ but at distance at most d from some vertex in $\mathcal{C}$ (i.e., those vertices which must be free of walkers for $\mathcal{C}$ to exist as a component). The *exterior* $\mathcal{E}_{\mathcal{C}}$ of $\mathcal{C}$ is the set containing all those vertices not in $\mathcal{C} \cup A_{\mathcal{C}}$. We partition $\mathcal{E}_{\mathcal{C}}$ into *external regions* as follows: two vertices belong to the same external region when they can be joined by a continuous arc not intersecting any edge of $\mathcal{C}$. Figure 3 shows a component with different external regions.

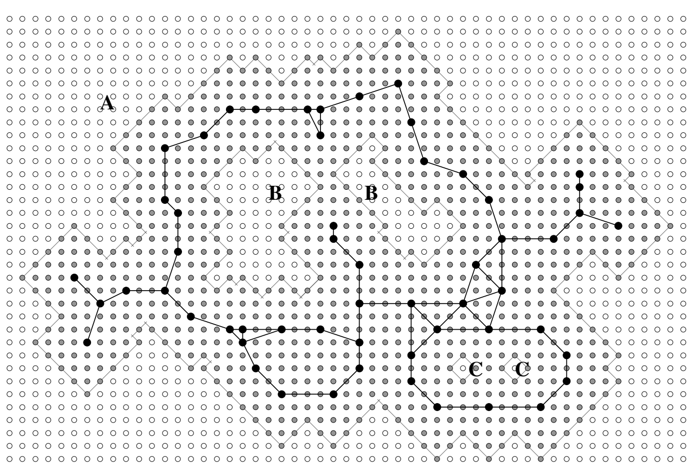

FIG. 3. *Component (black), empty area (gray), external regions A, B, and C (white).*

Recall that, in the terminology of planar maps, the *bounding cycle* of a face is a walk around the boundary of the face. Given an external region ${\mathcal{E}_{\mathcal{C}}}^i$, let $\mathcal{C}'$ be any connected subgraph of $\mathcal{C}$ that has no edges crossing and such that no vertices of $\mathcal{C}$ are contained in the face F of $\mathcal{C}'$ which contains ${\mathcal{E}_{\mathcal{C}}}^i$. Such subgraphs always exist: for instance, take the spanning tree of $\mathcal{C}$ whose length (sum of lengths of edges) in terms of dist has been minimized and, subject to this, has the shortest Euclidean length. We refer to the bounding cycle of this face F as a *boundary walk* β in $\mathcal{C}$ with respect to ${\mathcal{E}_{\mathcal{C}}}^i$. Such a walk is *maximal* if the face F does not properly contain a face of some other subgraph of $\mathcal{C}$ of the same type. In Figure 4, $\{1, 5, 3, 7, 9, 11, 13, 15, 13, 12, 14, 12, 10, 8, 6, 2, 4, 1\}$ is a nonmaximal boundary walk, and $\{1, 5, 3, 7, 9, 11, 13, 15, 13, 14, 12, 10, 8, 6, 2, 4, 1\}$ is a maximal one. A maximal walk always exists, because any nonmaximal walk can be diverted around any face that prevents it from being maximal.

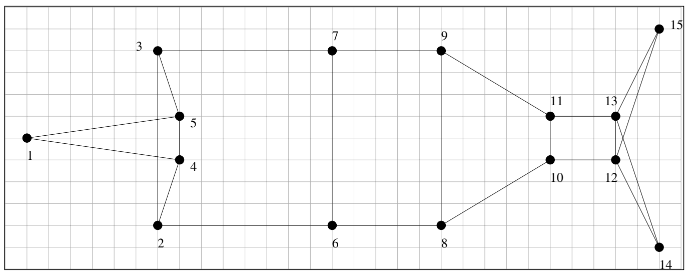

FIG. 4. *The boundary walk.*

We call a (directed) closed walk in $\mathcal{C}$ *regular* if, for each edge entering a vertex v, the next edge in the walk is the next edge in the clockwise direction around v.

For $i < n$, let us call a *v-band* of width i to any subset of T_N defined by $\{a, \dots, a+i-1\} \times \mathbb{Z}_n$. Similary, we define an *h-band* of height j. Define a *rectangle* of width i and height j to be the intersection of a v-band of width i and an h-band of height j. We can compare vertices in a rectangle according to their coordinates and use statements such as v_1 *is more left than* v_2 or v_3 *is an uppermost vertex* in the rectangle.

We say that a component $\mathcal{C}$ with at least 2 vertices is a *rectangular component* (*r-component*) if all of its vertices and edges and its forbidden region are contained in a rectangle as defined above. In particular, this implies that $\mathcal{C}$ contains no nonseparating cycle of the torus. Otherwise, $\mathcal{C}$ is an *nr-component*. For a given r-component $\mathcal{C}$, we define its *origin* as the leftmost of the lowermost vertices of $\mathcal{C}$. The *outside region* of an r-component is the only external region of the component having vertices outside any containing rectangle. We say that an nr-component is *type* 1 if is not the subgraph of some graph of walkers containing more than one nr-component and *type* 2 otherwise. Note that this property is simply determined by the subgraph of the toroidal grid remaining when the component and the vertices of distance at most d from it are deleted. By definition, $\mathcal{W}(G, w, d)$ has at most one nr-component of type 1.

Let X, Y, and Z be the number of simple components, r-components, and nr-components, respectively, and note that $Z = Z_1 + Z_2$, where Z_i is the number of nr-components of type i. Then $\mathbf{E}[Z] = \mathbf{P}(Z_1 = 1) + \mathbf{E}[Z_2]$.

4.1. Static properties. In this section, we study the connectedness of $\mathcal{W}(G, w, d)$ in the static situation for the case $G = T_N$; in particular, we analyze the behavior of X, Y, and Z. For section 4, the value of μ is as defined for Theorem 1.5, rather than the value for the cycle.

PROPOSITION 4.2. *The expected number of simple components of* $\mathcal{W}(G, w, d)$ *for* $G = T_N$ *satisfies* $\mathbf{E}[X] \sim N(1 - e^{-\varrho})\left(1 - \frac{h}{N}\right)^w$. *Furthermore,*

(i) *if* $\mu \to 0$, *then* $\mathbf{E}[X] \to 0$ *and there are no simple components a.a.s.,*

(ii) *if $\mu \to \infty$, then there exist simple components a.a.s. (and $\mathcal{W}(G, w, d)$ is disconnected), and*

(iii) *if $\mu = \Theta(1)$, then X is asymptotically Poisson with mean μ.*

Proof. We repeat the proof of Theorem 1.1 in the present context. To compute $\mathbf{E}[X]_k$, we focus on the set $\mathcal{S} = \{(v_1, \dots, v_k) | \mathrm{dist}(v_i, v_j) > 2d \text{ if } i \neq j\}$, with size $|\mathcal{S}| \sim N^k$. Applying Lemma 2.1, we obtain

$$\mathbf{E}[X]_k \sim \left[N \left(1 - e^{-\varrho}\right) e^{-h\varrho + O(\frac{h^2 w}{N^2})} \right]^k . \tag{19}$$

Comparing with (4), the rest of the proof is as for Theorem 1.1. □

From part (ii) of the proposition, if $h\varrho = O(1)$, then $\mu \to \infty$ and $\mathcal{W}(G, w, d)$ is disconnected a.a.s. In view of this, we may restrict our attention to the condition $h\varrho \to \infty$ in the study of r-components and nr-components.

Given a boundary walk $\beta = (v_1, \dots, v_k)$, we define $\mathrm{length}(\beta) = \sum_{1 \leq i < k} \mathrm{dist}(v_i, v_{i+1})$ as the sum of the distances (using the chosen metric) between consecutive vertices in β. We shall write $\mathrm{length}_{\ell^p}(\beta)$ when we want specify that we are measuring distances in ℓ^p. Similarly, we define $\mathrm{length}_v(\beta)$ (the vertical length) as the sum of the differences between y coordinates of consecutive vertices along the cycle and $\mathrm{length}_h(\beta)$ (the horizontal length) using x coordinates in the same way.

The next lemma relates the size of the forbidden region outside a boundary cycle of a component to the length of the cycle and will play a key role in proving the main results.

LEMMA 4.3 (Geometric Lemma). *Let $\mathcal{C}$ be a component in T_N with β one of its maximal boundary walks and $l = \mathrm{length}(\beta)$ its length. Assume that $\mathcal{C}$ has at least two occupied sites. Then the size of the forbidden region A_β outside β is bounded below by $|A_\beta| \geq dl/J$ for some sufficiently large constant J. Moreover, if $\mathcal{C}$ is rectangular, and β is a maximal boundary walk with respect to the outside region, we have $|A_\beta| \geq h + dl/J$.*

Proof. The main idea in this proof is to consider intervals in the grid "sticking outward" from the boundary walk and show that at least some small proportion of these intervals is on average free of walkers.

For convenience we take $J = 10^{10}$, though probably without large modifications the proof method will yield the result for $J = 1000$. (We do not attempt to optimize the constant J, since probably the theorem holds for J close to 2 when d is large, while for $d = 1$ in all but the case $p = 1$ we need $J \geq 4$. Some experiments suggest that $J = 4$ suffices even in this case.) Observe that in the 3×3 subgrid centered on the endpoint of an edge of β, there must be at least one vertex in the forbidden region outside β; otherwise the boundary walk can be rerouted to contradict its maximality. Each point of the forbidden region can be counted in this way at most 8 times. Thus the forbidden region has size at least $l/(8d)$. This implies the first statement in the lemma provided that $8d^2 \leq 10^{10}$.

Throughout this proof, all distances referred to are measured using dist. For the second statement in the lemma, we begin with the fact that for a rectangular component, there are "caps" of empty region of size at least $h/2 - d$ on the top and bottom of the component, being the set of grid points within distance d, but above the leftmost point of $\mathcal{C}$ of the greatest vertical coordinate (or below the leftmost point of the least vertical coordinate). Without loss of generality, we assume that the component has vertical height greater than 1 (otherwise, we may interchange "vertical" and "horizontal"). Then there are also intervals of empty region of length

d projecting outward from the left- and rightmost vertices of the top level of $\mathcal{C}$. These caps and intervals account for an empty region of size h but use up some of the forbidden region counted in relation to ends of edges of β in the argument above. The number of endvertices of edges of β that are involved in the region of size h described is at most $4d+2$, and there is still at least one endvertex of an edge of β not involved (the rightmost one in the bottom row of $\mathcal{C}$, say—we cannot claim two vertices here because it might be the same as the leftmost one), which means the above argument is valid if the bound on the size of the forbidden region is reduced by a factor of $4d+3$. That is, the size of the forbidden region is at least $h + l/8d(4d+3)$. So the second statement in the lemma holds provided that $8d^2(4d+3) \leq 10^{10}$. Hence we may now assume that $d > 400$.

Choose $k = \lceil \frac{d}{100} \rceil$. Then $k \leq (d+99)/100 \leq (d/100)(1+99/d) < d/80$. Assume without loss of generality that for the boundary walk β we have $\text{length}_v(\beta) \geq \text{length}(\beta)/2$. We place intervals of length k (note that all the ℓ_p metrics measure these intervals the same) horizontally along all grid lines from a maximal boundary cycle toward the outside. Those starting from a vertex of β point toward the outside according to the previous edge of β. (We assume β is oriented in some direction.) We delete any intervals that touch the boundary cycle in two or more places. Then each remaining interval will touch k vertices of the forbidden region outside β, and each such vertex will be touched by at most two different intervals. (Two intervals coming from opposite directions may touch the same vertex.)

We need to bound the number of intervals which were deleted. Call an edge of $\mathcal{C}$ *short* if the distance between its endpoints is less than $d/4$ and *long* otherwise. Suppose that an interval (that is to be deleted) touches a short edge at a point E and a point F on another edge of β, with the part between E and F in the exterior of β. Suppose that F has distance at most $d - d/4 - k$ from both endvertices of its edge. Then by the triangle inequality each endvertex (say, P and Q) of that edge is of distance at most d from each endvertex (R and T) of E's edge. Thus the quadrilateral $PQRT$ (or triangle, if two points coincide) has diameter at most d, and thus the walk β can be changed to make a smaller face "outside" (meaning the side which was minimized). This can clearly be done also if other parts of β enter this quadrilateral, contradicting the maximality of β. Thus, F has distance at least $d - d/4 - k > 2d/3$ from one of the endvertices of its edge. So F's edge is long, or, if F lies on more than one edge, they are all long.

We call a *middle* interval any interval originating from a long edge from point at least 1/8 of the length of the edge from each end. Suppose that such a middle interval originating at a point E (called a *middle point*), on edge e of β, is deleted. Then it hits some other edge f at a point F of horizontal distance at most k from E. If f has an endvertex of distance less than $d/8 - k \geq d/8 - d/80 \geq d/9$ from F, we get a contradiction as above. Since e is long, its endvertices have distance at least $d/32$ from E. Since e and f are straight, either point on e of distance $rd/32$ from E, for any $r > 0$, has horizontal distance at most $(r+1)k$ from f or the extension of f. (Furthermore, note for later that any two points on e and f at the same vertical coordinate have distance at most $32k$ apart.) It follows that the endvertices P and R of e and f above the line EF have distance less than d apart, as do the ends Q and T below (or on) the line EF. If either $P = R$ or $Q = T$, we now obtain a contradiction as before by rerouting β. A similar contradiction arises if either PT or QR has length at most d. So, picturing e and f with the line EF as making a near-perfect but very thin "H," the distances along the H from P to T and from Q to R are both greater than d. Thus the sum of the lengths of e and f is at least $2d-2k$ (since the edge EF is included twice in these distances), and so each of e and f has length at least $d-2k$, and the lengths of the arms of e and f above EF differ by at most k (as for the arms below EF). It follows that both PR and QT have length at most $34k$ (recalling the observation above about horizontal distance), which is at most $d/2$. Without loss of generality, PR is not an edge of β. This again contradicts the choice of β, either by shortening it at PR or, if another edge of β crosses PR, by joining P or R to an end of such an edge. We conclude that no middle intervals are deleted. Moreover, this shows that the point F in the previous paragraph cannot be a middle point. So every interval starting at a short edge that is deleted first hits a nonmiddle point of a long edge. Such a point can be hit by only one interval from a short edge unless it is a vertex of β, in which case it can be hit by two intervals. Thus if i_1 is the number of intervals starting at short edges that are deleted, j_1 is the number of long edges of β, and j_2 is the number of nonmiddle intervals originating at nonends of long intervals,

$$i_1 \leq 4j_1 + j_2.$$

Moreover, since no middle intervals are deleted, the number i_2 of intervals originating on long edges that are deleted similarly satisfies

$$i_2 \leq 4j_1 + j_2.$$

If j_3 is the total number of intervals before any deletions occur, we clearly have

$$j_2 \leq j_3/4$$

and also since long edges have length at least $d/4$,

$$j_3 \geq \text{length}_v(\beta) \geq \text{length}(\beta)/2 \geq j_1 d/8. \tag{20}$$

Combining these gives

$$i_1 + i_2 \leq 8j_1 + 2j_2 \leq j_3 \left(\frac{64}{d} + \frac{1}{2}\right) \leq 3j_3/4,$$

which shows that at least $j_3/4$ of all the intervals are not deleted. These intervals each cover $k \geq d/100$ vertices of the forbidden region, at most two covering any one such vertex, so the first part of the lemma follows for $d > 400$, using (20).

For the second claim of the lemma when $d > 400$, we may again add the caps of size $h - 2d$ but also two extra intervals of length d at the sides: assuming that β is oriented in the clockwise direction, the intervals of length k projecting from the leftmost vertex in the bottom level of $\mathcal{C}$, and from the rightmost vertex in the top level, are not used. It is here that extra intervals of the empty region of length d may be found. ∎

The next lemma will be used to show that nr-components of type 2 are rare.

Lemma 4.4. *Let $\mathcal{C}$ be an nr-component of type 2. Then $\mathcal{C}$ has a maximal boundary walk β with length$(\beta) \geq n - o(n)$.*

Proof. Let us divide T_N into $\lfloor \frac{n}{d+1} \rfloor$ v-bands $c_1, \dots, c_{\lfloor \frac{n}{d+1} \rfloor}$ of width $\geq d+1$, and similarly into $\lfloor \frac{n}{d+1} \rfloor$ h-bands $r_1, \dots, r_{\lfloor \frac{n}{d+1} \rfloor}$ of width $\geq d+1$.

Let $\mathcal{C}$ be an nr-component. Since $\mathcal{C}$ is connected, the v-bands (or h-bands) not containing vertices of $\mathcal{C}$ must be consecutive. If there were at least 2 consecutive v-bands and at least 2 consecutive h-bands without vertices of $\mathcal{C}$, then $\mathcal{C}$ would be rectangular since $\mathcal{C}$ and $A_{\mathcal{C}}$ would be embeddable in the complements of the v-bands and h-bands. Hence, at most one v-band and some consecutive h-bands (or at most one h-band and some consecutive columns) may be without vertices of $\mathcal{C}$.

We assume now that $\mathcal{C}$ coexists with another nr-component $\mathcal{C}'$.

Case 1. Let us suppose first that $\mathcal{C}$ has no vertices in more than one v-band or h-band (v-band without loss of generality). Let c_1, c_2 be two consecutive v-bands not containing vertices of $\mathcal{C}$. Hence all h-bands, excepting at most one, contain vertices in $\mathcal{C}$. For each such h-band r_i, choose a vertex v_i in $\mathcal{C} \cap r_i$. We can also find some vertex w_i in $(c_1 \cup c_2) \cap r_i$ such that w_i is at distance $\geq d+1$ from any vertex in $\mathcal{C}$. By this construction, all w_i belong to the same external region of the component. Let β be any maximal boundary walk of $\mathcal{C}$ with respect to this external region. Then, the straight line joining v_i and w_i must intersect an edge of β, part of the edge contained in r_i. Hence, β crosses all h-bands except at most 3 and length$(\beta) \geq n - 4d - 3$.

Case 2. On the other hand, let us suppose that $\mathcal{C}$ has vertices in all v-bands and h-bands except for at most one of each. Without loss of generality, the other component, $\mathcal{C}'$, has vertices in all v-bands except for at most one. Thus, there are at least $\lfloor \frac{n}{d+1} \rfloor - 2$ v-bands with some vertices of both components. For each such v-band c_i, let us take vertices $v_i \in \mathcal{C} \cap c_i$, $w_i \in \mathcal{C}' \cap c_i$ and join them by a straight line. Notice that all the w_i belong to the same external region of $\mathcal{C}$, and let β be any maximal boundary walk with respect to this region. Then the line joining v_i and w_i must intersect an edge of β, part of the edge contained in c_i. Hence, β crosses all v-bands except at most 4 and length$(\beta) \geq n - 5d - 4$.

In any case, $\mathcal{C}$ has a maximal boundary walk β with length$(\beta) \geq n - o(n)$. ∎

The next technical result shows that simple components are predominant a.a.s. in T_N. The proof uses the Geometric Lemma.

LEMMA 4.5. *If $h\varrho \to \infty$, then $\mathbf{E}[Y] = o(\mathbf{E}[X])$ and $\mathbf{E}[Z_2] = o(\mathbf{E}[X])$.*

Proof. Let us first bound the expected number $\mathbf{E}[Y]$ of rectangular components with more than one vertex. The amount of area that is e.o.w., as guaranteed by the Geometric Lemma, is large enough to make such components rare.

Notice from (18) that the ℓ^1-length of the edges of any boundary walk of a component are integers between 1 and $2d$.

Let $\mathcal{B}$ be the set of walks in T_N which are (for some configuration of the walkers) a maximal boundary walk of some r-component with respect to its outside region. For each $\beta \in \mathcal{B}$, choose a rooted spanning tree $T(\beta)$ of the graph induced by the edges of β. Note that given any such tree T of m vertices, we may recover β by joining certain pairs of vertices (with no edges crossing). The edges added are just diagonals added to a face of degree $2m-2$. For each vertex $v \in V$, natural $m \geq 2$, and tuple $\boldsymbol{l} = (l_1, \ldots, l_{m-1})$ of naturals $1 \leq l_i \leq 2d$, let $\mathcal{B}_{v,m,\boldsymbol{l}}$ be the set of all $\beta \in \mathcal{B}$ such that $T(\beta)$ has m vertices, is rooted at v, and has edges of ℓ^1-lengths $l_1, \ldots, l_{m-1}$. The number of such trees is at most $\prod_{j=1}^{m-1} 16 l_j$, where a factor 4^m comes from the number of rooted plane trees, and each factor $4l_j$ is the number of vertices of ℓ^1 distance l_j from a given vertex. Therefore,

$$|\mathcal{B}_{v,m,\boldsymbol{l}}| \leq \prod_{j=1}^{m-1} C l_j, \tag{21}$$

where C is constant, and also clearly

$$\mathbf{E}[Y] \leq \sum_{\substack{v \in V \\ m \geq 2 \\ 1 \leq l_1, \ldots, l_{m-1} \leq 2d}} \sum_{\beta \in \mathcal{B}_{v,m,l_1,\ldots,l_{m-1}}} \mathbf{P}(Y_\beta = 1), \tag{22}$$

where Y_β indicates the event of having some r-component with β one of its maximal boundary walks with respect to its outside region.

By Lemma 4.3, the size $|A_\beta|$ of the forbidden region outside β is an integer bounded below by $h + dl/J$, where $l = \text{length}(\beta)$. For technical purposes we consider a subset of A_β of size $h + \lceil dl/(2J+1) \rceil$, representing a region free of walkers. By Lemma 2.2, and noting that $\alpha > 1/2$ and hence $p < 2\varrho$, we obtain an upper bound for the probability of this emptiness occurring and the m vertices in β being occupied. Since this is necessary for the event Y_β to occur, we have

$$\begin{aligned} \mathbf{P}(Y_\beta = 1) &= O(1 - e^{-\varrho/\alpha})(2\varrho)^{m-1} \left(1 - \frac{S}{N}\right)^w \\ &= O(1 - e^{-\varrho})(2\varrho)^{m-1} \left(1 - \frac{h}{N}\right)^w e^{-\lceil dl/(2J+1) \rceil \varrho} \end{aligned} \tag{23}$$

since $\alpha \leq 1$. Furthermore, let $l' = l_1 + \cdots + l_{m-1}$. Then since the spanning tree has length no more than the length of β, $\text{length}_{\ell^1}(\beta) > l'$. By (18), we have $l \geq l'/2$, and hence we get

$$\mathbf{P}(Y_\beta = 1) = O(1 - e^{-\varrho})(2\varrho)^{m-1} \left(1 - \frac{h}{N}\right)^w e^{-dl'\varrho/J'}, \tag{24}$$

where $J' = 2(2J+1)$.

From (22), (21), and (24), we get

$$\mathbf{E}[Y] = O(1) \sum_{\substack{m \geq 2 \\ 1 \leq l_1, \ldots, l_{m-1} \leq 2d}} N \left(\prod_{j=1}^{m-1} C l_j \right) \left(1 - e^{-\varrho}\right) (2\varrho)^{m-1} \left(1 - \frac{h}{N}\right)^w e^{-dl'\varrho/J'}.$$

Therefore, using Proposition 4.2 for the asymptotic value of $\mathbf{E}[X]$,

$$\begin{aligned} \mathbf{E}[Y]/\mathbf{E}[X] &= O(1) \sum_{\substack{m \geq 2 \\ 1 \leq l_1, \ldots, l_{m-1} \leq 2d}} \left(\prod_{j=1}^{m-1} C l_j \right) (2\varrho)^{m-1} e^{-dl'\varrho/J'} \\ &= O(1) \sum_{m \geq 2} \left(\frac{C'}{d} \sum_{k=1}^{2d} k d\varrho e^{-kd\varrho/J'} \right)^{m-1}. \end{aligned} \tag{25}$$

In the case where $d\varrho \to \infty$, we have $\varrho e^{-c' d\varrho} = o(\max(1, \varrho)) = o(1)$ and hence $\mathbf{E}[Y] = o(\mathbf{E}[X])$. In the case where $d\varrho = O(1)$, we use $\sum_{k \geq 1} k c^{-\epsilon k d\varrho} = O((d\varrho)^{-2})$.

In the case where $d\varrho = O(1)$, we use $\sum_{k \geq 1} k c^{-kd\varrho} = O((d\varrho)^{-2})$ for $c < 1$, and (25) gives

$$\mathbf{E}[Y]/\mathbf{E}[X] = O(1) \sum_{m \geq 2} \left(\frac{C''}{d^2 \varrho} \right)^{m-1} = o(1) \qquad \text{as } d^2 \varrho \to \infty.$$

To prove $\mathbf{E}[Z_2] = o(\mathbf{E}[X])$, from Lemma 4.4, each component counted by Z_2 has some maximal boundary walk β with $\text{length}(\beta) \geq n - o(n)$. If we apply Lemma 4.3 to this β, we have $|A_\beta| \geq ld/J$, where $l = \text{length}(\beta)$. Using $d = o(n)$ (since $h = o(N)$), we have $|A_\beta| \geq h + ld/2J$ for large N, and we then proceed similarly as for Y. ∎

Proof of Theorem 1.5. From Proposition 4.2, if $\mu \to \infty$, then $\mathcal{W}(G, w, d)$ is disconnected a.a.s. In the other two cases, $\mu = O(1)$ and we must have $h\varrho \to \infty$. In this case we can apply Lemma 4.5 and get

$$\mathbf{P}(Y > 0) \leq \mathbf{E}[Y] = o(\mathbf{E}[X]) = o(1), \qquad \mathbf{P}(Z_2 > 0) \leq \mathbf{E}[Z_2] = o(\mathbf{E}[X]) = o(1).$$

Thus, a.a.s. we have only simple components and at most one nr-component of type 1. The rest of the theorem follows from the asymptotic distribution of X given in Proposition 4.2. ∎

4.2. Dynamic properties. According to the model, from an initial random placement f of the walkers, at each step, every walker moves from its current position to one of its neighbors, with a probability of 1/4 of going either way. This is a standard random walk on the grid for each walker. We wish to study the connectivity properties of $\mathcal{W}_t(T_N)$. The analysis of the dynamic case is quite similar to that of the cycle, so we state the major results and point to the differing details in the proofs.

We define *states* (or *configurations*) and the graph of configurations in an analogous way to the cycle (see section 3.2). In this case, there are $N^w = n^{2w}$ different configurations of walkers, each one represented by a vector $\mathbf{a} = (a_1, \dots, a_w) \in (\mathbb{Z}_n \times \mathbb{Z}_n)^w$, where $a_i = (a_{ix}, a_{iy})$ indicates the label of the vertex being occupied by walker i. Given a configuration $\mathbf{a} = (a_1, \dots, a_w)$, there exists an edge between $\mathbf{a}$ and all configurations $\mathbf{b} = (b_1, \dots, b_w)$, such that $\text{dist}(a_i, b_i) = 1$. Thus, any configuration has 4^w neighbors, and the relationship of being neighbors is symmetric. As in the case of the cycle, the dynamic process can be seen as a random walk on the graph of configurations and thus as a Markov chain $\mathcal{M}(N, w, d)$.

For N even, given any two configurations $\mathbf{a}$ and $\mathbf{b}$, we say that they have the *same parity* if, for all i and j, $(a_{i,x} - a_{j,x}) + (a_{i,y} - a_{j,y}) \equiv (b_{i,x} - b_{j,x}) + (b_{i,y} - b_{j,y}) \pmod 2$. With this definition of parity, Lemma 3.2 and its consequences also apply to the grid. Then, if N is odd, $\mathcal{M}(N, w, d)$ is ergodic, and if N is even, there are 2^{w-1} closed classes of states, where each class consists of all configurations with the same parity. The Markov chain restricted to any of these classes of states is irreducible and positive recurrent, but also 2-periodic, so we do not have ergodicity.

Observation 4.6. *Using the same argument as in Observation* 3.3*, for any fixed t, we can consider $\mathcal{W}_t(T_N)$ as a static $\mathcal{W}(T_N)$.*

In view of this observation and Theorem 1.5, we assume $\mu = \Theta(1)$ for the remainder of the section. This covers the nontrivial dynamic situations where $\mathcal{W}(G, w, d)$ is neither a.a.s. disconnected nor a.a.s. connected. Furthermore, in this case we need only to focus on the study of simple components. We also assume for the present subsection that $d \geq 2$. The case $d = 1$ is excluded for technical reasons.

We define $X = X(t)$ to be the random variable that counts the number of simple components at time step t. Given our assumptions about μ, for t in any fixed bounded time interval, $X(t)$ is asymptotically Poisson with expectation $\mu = \Theta(1)$, as studied in Proposition 4.2.

In analogy with d-hole lines in section 3.2, we define a *simple component line* to be a maximal sequence of pairs $(v_1, t_1), \dots, (v_l, t_l)$, where v_i is a simple component existing at time step t_i for $1 \leq i \leq l$, and such that $t_i = t_{i-1} + 1$ and the vertex v_i is adjacent to v_{i-1}, for $2 \leq i \leq l$. Birth, death, and survival of lines and the random variables $B(t)$, $D(t)$, and $S(t)$ are defined analogously to the cycle case. For the case that $d = 1$, the concept of simple component lines is not quite adequate to describe what is going on unless $\varrho \to 0$, since otherwise it is reasonably likely that a vertex is occupied by several walkers that in one step jump to several different neighbors and create several component lines simultaneously. For this reason we exclude $d = 1$ at this point, though it is very easy to obtain most of the results claimed if $\varrho \to 0$. Presumably a suitable analysis of the simultaneous creation or destruction of several simple component lines gives results even for $\varrho \not\to 0$, but we have not done this.

We need one more definition referred to in section 1. Let v and v' be any two adjacent vertices in the grid. Let b be the number of directed edges whose origin is a vertex at distance strictly greater than d from v and whose destination is at distance at most d from v'. Note that this quantity b does not depend on the particular v and v' but just on parameter d. We have that $b = \Theta(d)$, but the exact expression of this b depends on the particular chosen metrics. Some examples are found in Table 2.

TABLE 2
Parameter b.

Metrics	b
ℓ^1	$b = 2d + 1/2$
ℓ^p $(p < \infty)$	$b \sim \left(1/\sqrt[p]{2} + 3/2\right) d$, if $d \to \infty$
ℓ^∞	$b = 3d + 1$

We next have a result analogous to Proposition 3.4.

PROPOSITION 4.7. *Let $d > 1$. For t in any fixed bounded time interval, the random variables $S(t)$, $B(t)$, and $D(t)$ are asymptotically independent Poisson, with the expectations*

$$\mathbf{E}\left[S(t)\right] \sim \begin{cases} \mu & \text{if } d\varrho \to 0, \\ \mu - \lambda & \text{if } d\varrho \to c, \\ 4\frac{1-e^{-\varrho/4}}{1-e^{-\varrho}} e^{-(b+3/4)\varrho}\mu & \text{if } d\varrho \to \infty, \end{cases}$$

$$\mathbf{E}\left[B(t)\right] = \mathbf{E}\left[D(t)\right] \sim \begin{cases} b\varrho\mu & \text{if } d\varrho \to 0, \\ \lambda & \text{if } d\varrho \to c, \\ \mu & \text{if } d\varrho \to \infty, \end{cases}$$

where $\lambda = \left(1 - e^{-b\varrho}\right)\mu$ as in Theorem 1.7.

Proof. For a vertex $v \in V$ with coordinates (x, y), let v_N, v_S, v_E, v_W denote the vertices with coordinates $(x, y+1)$, $(x, y-1)$, $(x+1, y)$, $(x-1, y)$. The proof is so similar to that of Proposition 3.4 that we discuss in detail only the ways that a simple component line can be born at v between time steps t and $t+1$. We classify these events in 4 main classes, as shown in Figure 5.

- b1. At time step t, just one vertex of v_N, v_S, v_E, v_W is occupied and belongs to some bigger component. Then, all walkers there jump to v and stop communicating with other walkers.
- b2. At time step t, just two vertices of v_N, v_S, v_E, v_W are occupied. Then, all walkers there jump to v and do not communicate with other walkers.
- b3. At time step t, just three vertices of v_N, v_S, v_E, v_W are occupied. Then, all walkers there jump to v and do not communicate with other walkers.

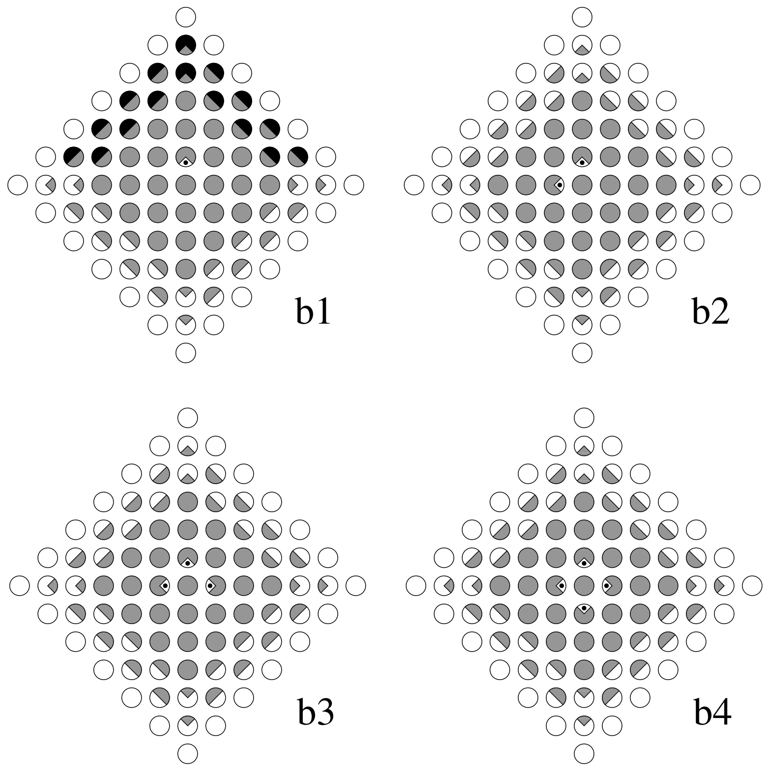

FIG. 5. *Birth of a simple component line at vertex v.*

b4. At time step t, vertices v_N, v_S, v_E, v_W are occupied. Then, all walkers there jump to v and do not communicate with other walkers.

The first three classes have several subtypes (four for b1, six for b2, and four for b3) which correspond to choosing a subset of the four possible compass directions.

The rest of the proof follows that of Proposition 3.4. (See [DPSW06] for the full details.) ∎

From Proposition 4.7, we can easily derive important consequences analogous to those of the cycle, always under the assumption stated after Observation 4.6. The first one gives us the probability that $\mathcal{W}_t(T_N)$ is connected but $\mathcal{W}_{t+1}(T_N)$ is disconnected. Note that this is a.a.s. equivalent to having no simple components at time step t and just one at $t+1$. The proof is immediate by the same argument as for Lemma 3.9.

LEMMA 4.8. *Let $d > 1$. The probability that $\mathcal{W}_t(T_N)$ is connected and that $\mathcal{W}_{t+1}(T_N)$ is disconnected is given by*

$$\mathbf{P}\left(X(t+1) \geq 1 \wedge X(t) = 0\right) \sim \begin{cases} \mu e^{-\mu} b\varrho & \text{if } d\varrho \to 0, \\ e^{-\mu}(1-e^{-\lambda}) & \text{if } d\varrho \to c, \\ e^{-\mu}(1-e^{-\mu}) & \text{if } d\varrho \to \infty, \end{cases}$$

where λ is defined as in Theorem 1.7.

In a similar way to the cycle case, we define the *lifespan* of a simple component line as the number of time steps for which the line is alive. For any vertex v and time t, the random variable $L_{v,t}$ counts the lifespan of the simple component line born at vertex v between time steps t and $t+1$. If this birth does not take place, $L_{v,t}$ is defined to be 0.

We finish by considering lifespans and connected periods. Define L_{av} (the average lifespan of simple component lines), L_T, and L as in section 3.2 but in terms of the new definition of $L_{v,t}$. Similarly, define LC_{av} to be the average length of the period from time t for which the graph of walkers on the grid is connected. We state the following theorem without proof because it is not difficult to adapt the argument used for Theorems 1.3, 1.4, 3.6, and 3.8. The proof is entirely straightforward in the case that L_{av} is bounded (i.e., $d\varrho \not\to 0$). If L_{av} is unbounded, then there are two modifications to the argument. First, one needs to do the moment calculations more carefully (this is analogous to the situation at the end of the proof of Proposition 3.4; there is less independence when d is large, so this takes a bit of care). Second, for connectivity issues the observation that a complex component exists at any one time with probability tending to 0 does not show that we can ignore them, since the lifespan is no longer bounded. However, we may easily adapt the analysis to consider the creation or destruction of components (not just simple ones) in any given step to obtain the required result. We omit these details; they can be found in [P07].

THEOREM 4.9. *For the walkers model on T_N with $d > 1$,*

$$L_{\mathrm{av}} \sim \begin{cases} \frac{1}{b\varrho} & \text{if } d\varrho \to 0, \\ \frac{\mu}{\lambda} & \text{if } d\varrho \to c, \\ 1 & \text{if } d\varrho \to \infty, \end{cases} \qquad \text{and} \qquad LC_{\mathrm{av}} \sim \begin{cases} \frac{1}{\mu b\varrho} & \text{if } d\varrho \to 0, \\ \frac{1}{1-e^{-\lambda}} & \text{if } d\varrho \to c, \\ \frac{1}{1-e^{-\mu}} & \text{if } d\varrho \to \infty, \end{cases}$$

where λ is defined as in Theorem 1.7*. Furthermore, L_T (LC_T) converges in probability for T growing large (N fixed) toward a random variable L (LC, respectively), uniquely determined almost everywhere, where $L \sim L_{\mathrm{av}}$ and $LC \sim LC_{\mathrm{av}}$ a.a.s.*

The proof of Theorem 1.7 is a simple adjustment of the proof of Theorem 4.9 (as in the cycle case).

5. Conclusions and open problems. In this work we have characterized connectivity issues of a very large set of moving agents which move through a prescribed real or virtual graph. We believe it is the first time that these kinds of characterizations have been obtained, and it could open an interesting line of research. We gave characterizations for the cycle and the grid. The results obtained for the grid could easily be extended to the grid with diagonals. Also, an approach similar to ours should work with the k-dimensional toroidal grid, but a suitable substitute for the Geometric Lemma needs to be found.

Our results are based on results for the static problem, and the most difficult was to show that, for the crucial values of w and d, there is a.a.s. just one big component apart from simple ones. There is a corresponding argument in Penrose [P97], [P99] for the related problem in which w walkers are placed at random in a unit box in continuous Euclidean space (or the toroidal counterpart), and walkers of distance at most r can communicate directly. Our setting has one more parameter, but it is clear that the behavior for our problem should be similar to the continuous one with $r = d/n$, where our toroidal grid is $n \times n$. However, proving this equivalence would not be entirely straightforward in the case that d is small (constant or near-constant), since then the discrete ball of radius d is not well approximated by the continuous one.

In our model we use a fixed number w of walkers. One could alternatively place walkers randomly so that each vertex is occupied independently with probability p.

For example, they could be Poisson-distributed at each vertex (as suggested by an anonymous referee) with parameter λ such that $1 - e^{-\lambda} = p$. In the static case this alternative would bear the same relation to our model as the relation between the random graph models $G(n, m)$ and $G(n, p)$, and as for that case, one would expect similar properties when p is approximately w/N (or, more precisely, $1 - e^{-w/N}$, to capture the case that w is close to or greater than N). This would simplify some of our analysis (e.g., the proof of Lemma 2.1). However, it would be difficult to deduce all the results for our model in such a way. For one thing, some of the properties we study are not convex in the required sense (see [JLR00]). There are also other obstacles to using models with independent occupancy probabilities. For instance, if $w(s_0/N)^2 \neq o(1)$, then $(1 - (S_0/N))^w$, as in Lemma 2.1, is not asymptotic to $e^{-S_0\varrho}$.

Further planned work is the extension of the results to random geometric graphs, which will provide a model for omnidirectional radio communication in the setting described in [GHSZ]. Another interesting extension is the hypercube, as the number of neighbors of a vertex is then not constant. A further project is to study the connectivity of walkers when the underlying topology has obstacles which can interfere with communication.

REFERENCES

[AGE02] P. K. Agarwal, L. Guibas, H. Edelsbrunner, J. Erickson, M. Isard, S. Har-Peled, J. Hershberger, C. Jensen, L. Kavraki, P. Koehl, M. Lin, D. Manocha, D. Metaxas, B. Mirtich, and D. Mount, *Algorithmic issues in modeling motion*, ACM Computing Surveys, 34 (2002), pp. 550–572.

[Bol02] B. Bollobás, *Random Graphs*, 2nd ed., Academic Press, London, 2002.

[DPSW05] J. Díaz, X. Pérez, M. Serna, and N. C. Wormald, *Connectivity for wireless agents moving on a cycle or grid*, in Proceedings of STACS 2005, Lecture Notes in Comput. Sci. 3404, Springer-Verlag, Berlin, 2005, pp. 353–364.

[DPSW06] J. Díaz, X. Pérez, M. Serna, and N. C. Wormald, *Walkers on the Cycle and the Grid*, Research Report CORR 2006-09, Faculty of Mathematics, University of Waterloo, Waterloo, Canada, 2006.

[GHSZ] L. Guibas, J. Hershberger, S. Suri, and Li Zhang, *Kinetic Connectivity for Unit Disks*, Discrete Comput. Geom., 25 (2001), pp. 591–610.

[H06] R. Hekmat, *Ad-hoc Networks: Fundamental Properties and Network Topologies*, Springer-Verlag, New York, 2006.

[H80] L. Holst, *On multiple covering of a circle with random arcs*, J. Appl. Probab., 16 (1980), pp. 284–290.

[JLR00] S. Janson, T. Łuczak, and A. Ruciński, *Random Graphs*, Wiley, New York, 2000.

[JBAS03] A. Jardosh, E. Belding-Royer, K. Almeroth, and S. Suri, *Towards realisitic mobility models for mobile ad-hoc networks*, in Proceedings of ACM Mobicom, San Diego, CA, 2003, pp. 217–229.

[P97] M. D. Penrose, *The longest edge of the random minimal spanning tree*, Ann. Appl. Probab., 7 (1997), pp. 340–361.

[P99] M. D. Penrose, *On k-connectivity for a geometric random graph*, Random Structures Algorithms, 15 (1999), pp. 145–164.

[P07] X. Pérez, *Aspects of Random Graphs*, Doctoral Dissertation, Universitat Politècnica de Catalunya, Barcelona, Spain, 2007.

[RMM01] E. M. Royer, P. M. Melliar-Smith, and L. E. Moser, *An analysis of the optimum node density for ad hoc mobile networks*, in Proceedings of the IEEE International Conference on Communications, IEEE Communications Society, Piscataway, NJ, 2001, pp. 857–861.

[W04] N. C. Wormald, *Random graphs and asymptotics*, in Handbook of Graph Theory, J. L. Gross and J. Yellen, eds., CRC, Boca Raton, FL, 2004, pp. 817–836.

[XK06] F. Xue and P. R. Kumar, *Scaling Laws for Ad Hoc Wireless Networks: An Information Theoretic Approach*, NOW Publishers, Delft, The Netherlands, 2006.