

k Workers in a Circular Warehouse: A Random Walk on a Circle, without Passing*

Joseph D. Skufca[†]

Abstract. We consider the problem of stochastic flow of multiple particles traveling on a closed loop, with a constraint that particles move without passing. We use a Markov chain description that reduces the problem to a generalized random walk on a hyperplane (with boundaries). By expressing positions via a moving reference frame, the geometry of the no-passing criteria is greatly simplified, with the resultant condition expressible as the coordinate system planes which bound the first orthant. To determine state transition probabilities, we decompose transitions into independent events and construct a digraph representation in which calculating transition probability is reduced to a shortest path determination on the digraph. The resultant decomposition digraph is self-converse, and we exploit that property to establish the necessary symmetries to find the stationary density for the process.

Key words. generalized random walks, Markov chains, self-converse

AMS subject classifications. 60G50, 60J10, 05C99

DOI. 10.1137/S0036144503432011

I. Introduction. In dense traffic flow, the speed of individual vehicles is constrained because the cars are unable to pass one another. In protein-controlled diffusion through a membrane, the narrow passage forces the transported molecules to move single-file. A similar phenomenon develops in the motion of people along multiple serving stations, such as occurs in a buffet line. Motivated by the natural problems that develop when physical objects try to move independently along the same path, we consider a stochastic flow of particles on a closed loop, with the constraint that one particle may not pass another. We limit this paper to the specific problem in which motion occurs over discrete units of time and the particles are limited to moving in some integer number of steps. We expect that our approach may have some applicability to models of microscopic transport, cellular automata models of traffic flow, and some queuing theory problems, as well as providing some general insight into diffusive flows in constrained channels.

Problem Statement. Consider a warehouse with parts stored in n locations (parts bins) arranged in a closed path. k workers move around the warehouse in a continuous loop, filling an order on each trip through the warehouse. In real life, a worker would collect items from those bins which contain the parts required to fill

*Received by the editors July 23, 2003; accepted for publication (in revised form) June 29, 2004; published electronically April 29, 2005.

<http://www.siam.org/journals/sirev/47-2/43201.html>

[†]University of Maryland, College Park, MD, and United States Naval Academy, Annapolis, MD (skufca@maxinter.net).

an order. We make a simplifying assumption and model this selection process as a Bernoulli trial, with probability s that the worker will select a part from any given bin. All the workers move in the same direction as they fill their orders, but without passing each other. (Perhaps they are all driving forklifts, and there is no room to pass; or maybe the workers are actually robotic carts moving on a common track.) To define “time,” we assume that during one time step, each worker moves around the closed loop until one of two conditions is met:

1. The worker has a Bernoulli success at some bin and stops to collect a part.

Or

2. since we assume the workers are not allowed to pass one another, it may happen that one worker has stopped to gather a part, but the trailing worker may have no Bernoulli successes before reaching the bin directly behind the worker who is gathering a part. We call this condition a *blockage*. The trailing worker is required to wait in that adjacent bin for the entire time step and collects no part.

Blockages, therefore, represent reduced output from a worker. Our goal is to evaluate the efficiency of our warehouse by *determining the expected value of the fraction of the time steps that an individual worker is blocked*.

We apply the theory of Markov chains to this problem, which leads to a description of the process as a $(k - 1)$ -dimensional generalized random walk on a compact lattice. We develop the state space description in a moving reference frame, where the constraint of single file motion can be simplified to a boundary condition expressible as a coordinate plane. To determine transition probabilities, we develop an associated digraph that allows us to decompose transitions in a way that relates transition probabilities to shortest paths between vertices of that digraph. We then exploit the self-converse structure of that digraph to establish a stationary density on the state space. The paper is organized as follows: in section 2, we define and enumerate the states; in section 3, we develop an expression for the transition probabilities; in section 4, we establish a stationary distribution for the general problem; and in section 5, we answer the specific efficiency question of the problem statement.

2. Labeling and Counting States. Markov chain theory requires that the probabilities of future states of the system depend only upon the current state of the system, so the initial step in applying these techniques is to identify a state space description that satisfies this requirement. Since all the bins have the same probability for selection, a blockage event does not depend upon the specific bin location of the worker, but rather it is the distance from one worker to the next that is of critical importance. Therefore, we focus on these distances to provide a state space description. Let

$$d = w - z \pmod{n}$$

be the distance from a worker at bin z forward to a worker at bin w (in the direction of motion of the workers). We specify that two different workers may not occupy the same bin, so $d \geq 1$. We let $X = (x_1, \dots, x_k)$, $X \in \mathbb{R}^k$, with x_i the distance from worker i forward to worker $i + 1 \pmod{k}$, and we say that vector X defines a *physical configuration* of the system (or simply *configuration*). This vector description can be viewed as using a moving reference frame (moving with the workers), which replaces the fixed frame of the bin numbers. Each configuration is actually an equivalence class of bin locations for the workers. Figure 1 provides a visual example of this method of configuration identification.




Fig. 1 A warehouse of $n = 16$ bins. The positions of the $k = 4$ workers are shown in circles. Note that $x_1 + x_2 + x_3 + x_4 = 3 + 6 + 2 + 5 = 16 = n$. The bin locations of the four workers in the fixed frame, $(2, 5, 11, 13)$, constitute a member of the equivalence class denoted by configuration $(3, 6, 2, 5)$.

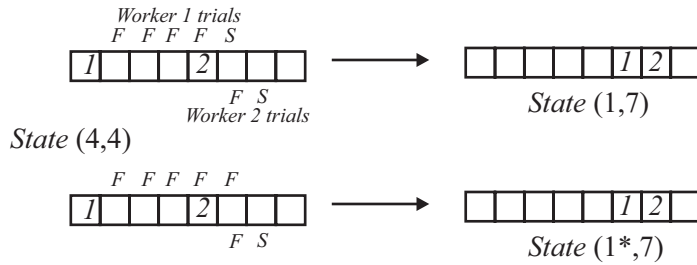


Fig. 2 With $k = 2$ and $n = 8$, we see two possible transitions from the state $(4, 4)$. The Bernoulli trial results are indicated with F for failure and S for success. Both transitions end in the same physical configuration. However, in the second case, worker 1 does not have a Bernoulli success before he is forced to stop because he cannot occupy the same bin as worker 2. In this second case, worker 1 would not collect a part.

Since the workers are restricted to a warehouse of n bins, we have $\sum_{i=1}^k x_i = n$. Counting configurations is simply a way of approaching the problem, *How many ways can you arrange k positive integers that sum to n ?* This well-known problem, generally covered in a first course in combinatorics, yields the formula

$$(1) \quad C(k, n) = \binom{n-1}{k-1}$$

possible configurations. To extend the set of configurations to describe all the states of the system, we must consider blockages. In a *blocked* situation, the blocked worker is physically a distance $d = 1$ from the worker ahead. The physical configuration is the same in both cases. The difference, however, is that at the bin where he is stopped, the blocked worker had a Bernoulli failure instead of a Bernoulli success (and therefore does not collect a part on that time step). To distinguish this condition, we use the symbol 1^* to represent that blocked condition. Figure 2 illustrates this difference.

By defining $|1^*| := 1$, we can formally represent states as vectors:

$$(2) \quad X = (x_1, \dots, x_k), \quad x_i \in \{1^*, 1, 2, \dots, n - k + 1\}, \quad \sum_{i=1}^k |x_i| = n.$$

Geometrically, the blockage condition forms a boundary on the configuration space, where these boundaries can be treated as hyperplanes that are parallel to each of

the coordinate planes.¹ The simplicity of this boundary expression results from our having shifted to the moving coordinate system. In a fixed reference (with states described by bin locations), the geometry of the blockage state becomes significantly more complex as k increases [2].

We enumerate the blocked states as follows: Suppose we want to count the number of states with b blockages. An arbitrary state can be considered an ordered list of k elements, and we seek to count the number of valid ways to fill those k elements. Forming such a sequence can be broken down into two steps: first, place a 1^* for each worker that is blocked; then, fill the other elements with positive integers to satisfy the summation constraint of (2). There are $\binom{k}{b}$ ways to insert the 1^* 's into the stream. The other $k - b$ elements must be filled with positive integers that sum to $n - b$. We use (1) with new arguments $(k - b)$ and $(n - b)$ to count the possibilities for the other integers, which come to $\binom{n-b-1}{k-b-1}$. Applying the multiplicative rule of counting, we find that N , the number of states with b blockages, is given by

$$(3) \quad N(b, k, n) = \binom{k}{b} \binom{n-b-1}{k-b-1}.$$

Example. Let $k = 5$ and $n = 16$. An arbitrary state is $(\square, \square, \square, \square, \square)$. Suppose $b = 2$. We must place two 1^* 's into five possible slots; there are $\binom{5}{2}$ ways to do so, and $(\square, \square, 1^*, \square, 1^*)$ would be a specific choice. We now must fill the other three positions with positive integers which sum to 14, and there are $\binom{13}{2}$ ways to do so. Then $N(2, 5, 16) = \binom{5}{2} \binom{13}{2} = 780$.

To count the entire state space, we simply sum over all possible values for b . (We note that the maximum number of blockages is $b = k - 1$, since one worker must be stopped in order to block the other workers.) The total number of states would be given by

$$(4) \quad M(k, n) = \sum_{b=0}^{k-1} N(b, k, n) = \sum_{b=0}^{k-1} \binom{k}{b} \binom{n-b-1}{k-b-1}.$$

3. Transition Probabilities. Given the state description of section 2, we number those states from 1 to M . To establish the Markov property, we must identify the transition probabilities

$$(5) \quad p_{ij} = \mathbb{P}[\text{System in state } X_j \text{ at time } \tau + 1 \mid \text{System in state } X_i \text{ at time } \tau].$$

(For our general problem, there is no obvious scheme for a preferred numbering of the states. Therefore, we develop transition probabilities as p_{XY} , the probability of transitioning from state X to state Y , and identify $p_{ij} \equiv p_{X_i X_j}$.) First, we will consider only the transitions between unblocked states, and then we extend that formulation to describe transitions from and to the blocked states.

3.1. Transition Probabilities for Unblocked States. Recall that s is the probability of Bernoulli success, and let $q = 1 - s$. We decompose a transition from unblocked

¹Alternatively, we could have used 0 instead of 1^* to represent a blockage, which makes it easier to interpret the blockage states as the coordinate hyperplanes $x_j = 0$ forming the boundary of the physical configurations which lie strictly in the first orthant. We find the $|1^*| = 1$ symbology more useful in providing an algebraic and physical description of the state.

state X to unblocked state Y into a sequence of independent events:

	Each worker has Q Bernoulli failures before one of the workers has a success.	All workers have one success.	After one worker has stopped, the other workers have additional failures to move the system from configuration X to configuration Y .
Event:	$\underbrace{\hspace{15em}}$	$\underbrace{\hspace{15em}}$	$\underbrace{\hspace{15em}}$
Probability:	q^{kQ}	s^k	$f(X, Y)$

We compute the transition probability by taking the product of the independent events and summing over all Q to give

$$(6) \quad p_{XY} = \sum_{Q=0}^{\infty} q^{kQ} s^k f(X, Y) = s^k \frac{1}{1 - q^k} f(X, Y),$$

where $f(X, Y)$ remains to be specified. We call $f(X, Y)$ the *arrangement probability* and note that it is a conditional probability (conditioned upon one of the workers having stopped). Our visualization is that all workers simultaneously move to an adjacent bin and conduct a trial. If all workers had failures, they would all index by one bin and repeat. Until the first success, the physical configuration does not change. However, when one of the workers has a success, that worker stops, and the other workers keep moving. The change in configuration is the result of workers moving, i.e., Bernoulli failures. Therefore, $f(X, Y) = q^l$, where l is the total number of bins that the workers must pass to take the system from configuration X to configuration Y , given that one of the workers passes 0 bins. We use a geometric and graph-theoretic approach to resolve the problem of counting these Bernoulli failures.

Background Geometry to Address Arrangement Probability. Let X be an unblocked state: a k -tuple of positive integers which sum to n . Then X lies on the hyperplane in \mathbb{R}^k given by $\sum_{i=1}^k x_i = n$, with orthogonal vector $u = (1, \dots, 1) \in \mathbb{R}^k$. Specifically, X is a lattice point in the first orthant that lies on that hyperplane. Define vector $V^r = (v_1, \dots, v_k)$ by

$$(7) \quad v_i = \begin{cases} n - k + 1, & i = r, \\ 1, & i \neq r. \end{cases}$$

Geometrically, each V^r specifies a vertex of the polytope containing all the configurations. Figure 3 illustrates this geometry for the case $k = 3, n = 7$: Figure 3(a) shows how the hyperplane and lattice points are situated in \mathbb{R}^3 , and Figure 3(b) shows the view from a position orthogonal to the plane, illustrating the structure on the $(k - 1 = 2)$ -dimensional subspace. Vertices V^i are labeled on both figures.

We define Ω to be the convex hull of $\{V^r : r = 1, \dots, k\}$. For $X, Y \in \Omega$, we define $\Delta_{XY} = (d_1, \dots, d_k) := Y - X$. Since $\sum x_i = \sum y_i = n$, a necessary condition for Δ is

$$(8) \quad \sum_{i=1}^k d_i = 0.$$

Geometrically, each Δ_{XY} is a lattice point on the hyperplane orthogonal to vector $u = (1, \dots, 1)$ that passes through the origin. We call this hyperplane \mathcal{P} . We seek




Fig. 3 Visualization of states for $k = 3, n = 7$.

to describe $f(X, Y)$ in terms of some function $g(\Delta_{XY})$, with $g(\cdot)$ independent of X and Y . The justification for this approach is that if no blockages occur, the workers are not affected by each other. The rearrangement probability can be completely characterized by the *change* in spacing between workers, independent of the initial and final configurations. If such a function g exists, then the process can be treated as a generalized random walk on \mathbb{R}^{k-1} (the hyperplane dimension). The physical configurations simply characterize which Δ vectors are feasible, because we require $X + \Delta = Y \in \Omega$. Since X, Y , and Ω are contained within a bounded set, our description becomes a random walk with boundary.

Canonical Representation of Δ . Suppose X describes a configuration. If worker $i + 1$ has a Bernoulli failure and moves one bin, then the distance from the worker behind has increased by 1, while the distance to the worker ahead has decreased by 1. The new configuration \tilde{X} would be given by

$$\begin{cases} \tilde{x}_i = x_i + 1, \\ \tilde{x}_{i+1} = x_{i+1} - 1, \\ \tilde{x}_j = x_j, & j \notin \{i, i + 1\}. \end{cases}$$

Motivated by this effect of a single Bernoulli failure, we define vectors $\delta_i = (\delta_{i1}, \dots, \delta_{ik})$, where

$$\delta_{ij} = \begin{cases} 1, & j = i, \\ -1, & j = i + 1 \pmod{n}, \\ 0, & \text{otherwise.} \end{cases}$$

(Example: $\delta_3 = (0, 0, 1, -1, 0, \dots, 0)$.) We note that $\{\delta_i\}_{i=1}^k \subset \mathcal{P}$ and $\text{span}(\{\delta_i\}_{i=1}^k) = \mathcal{P}$. These spanning vectors have a natural representation as edges in a digraph which connect the configurations that result from a single Bernoulli failure. We call this representation the *rearrangement digraph*. Figure 4 shows the rearrangement digraph for $k = 3, n = 7$.

Any point $\Delta \in \mathcal{P}$ is a linear combination of the δ_i 's, so there exists vector α with integer components such that

$$(9) \quad \Delta = \sum_{i=1}^k \alpha_i \delta_i.$$




Fig. 4 A digraph representation of rearrangements for $k = 3$ and $n = 7$ using δ vectors. Each edge represents the effect of a single Bernoulli failure by one worker.

Since \mathcal{P} is $(k - 1)$ -dimensional, the representation is not unique. However, we want that decomposition to satisfy some additional constraints: As previously stated, rearrangement probability is conditioned upon one of the workers having no failures. Therefore, $\alpha_j = 0$ for some $j \in \{1, \dots, k\}$. There are at most k representations that satisfy this requirement. We also need the representation to be a path in the rearrangement digraph, which implies $\alpha_i \geq 0$ for all i . The following paragraphs describe the algorithm that identifies this canonical representation.

Construction of Canonical Representation. Define matrix $A_{k \times k} = [\delta_1, \dots, \delta_k]$, where the δ_i 's form columns of A . Solutions to (9) are equivalent to solutions of $A\alpha = \Delta$. We note that $\text{rank } A = k - 1$, so A is singular. It is easy to verify that the eigenvector for eigenvalue $\lambda = 0$ is $u = (1, \dots, 1)$. The one-parameter family of solutions to (9) is given by

$$(10) \quad \alpha = \alpha^0 + tu,$$

where α^0 is any solution to (9) and t is an integer.

The geometric interpretation of α is that each α describes a path associated with the displacement Δ , and each α_i describes how many steps we take in the δ_i direction when traversing that path. Since we want the decomposition to represent a valid rearrangement of the workers, we seek vector β such that

$$(11) \quad \begin{cases} A\beta = \Delta, \\ \beta_i \geq 0 & \forall i \in \{1, \dots, k\}, \text{ and} \\ \beta_j = 0 & \text{for at least one } j \in \{1, \dots, k\}. \end{cases}$$

To construct this vector β , we proceed as follows: Define $\gamma := (\gamma_1, \dots, \gamma_k)$, where

$$(12) \quad \gamma_j = \sum_{i=1}^j d_i.$$

From (8), we see that $\gamma_k = 0$, simplifying the computation which verifies that

$$(13) \quad A\gamma = \begin{bmatrix} \gamma_1 - \gamma_k \\ -\gamma_1 + \gamma_2 \\ \vdots \\ -\gamma_{j-1} + \gamma_j \\ \vdots \end{bmatrix} = \begin{bmatrix} d_1 + 0 \\ -d_1 + (d_1 + d_2) \\ \vdots \\ -\sum_{i=1}^{j-1} d_i + \sum_{i=1}^j d_i \\ \vdots \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \\ \vdots \\ d_j \\ \vdots \end{bmatrix} = \Delta.$$

In general, γ will not satisfy the second and third elements of (11). To find the desired solution, we define $\gamma_{\min} = \min_i \gamma_i$ and assign $\beta_i = \gamma_i - \gamma_{\min}$. Clearly, all of the components of β are nonnegative, with $\beta_j = 0$ if γ achieves its minimum at j . In vector form, $\beta = \gamma - \gamma_{\min}u$. Since $A\gamma = \Delta$, we use (10) to see that β must be the solution to (11).

Example. Let $k = 5$ and $n = 12$. Suppose $X = (2, 1, 4, 4, 1)$ and $Y = (3, 1, 2, 2, 4)$. Then $\Delta = Y - X = (1, 0, -2, -2, 3)$. We first find a solution to $A\gamma = \Delta$ by using (12) to find $\gamma = (1, 1, -1, -3, 0)$. Since $\gamma_{\min} = -3$, we find $\beta = (4, 4, 2, 0, 3)$.

β_i gives the number of Bernoulli failures required of worker $i + 1$ to achieve the rearrangement specified by Δ , and $\|\beta\| = \sum \beta_i$ gives the total number of failures to achieve rearrangement.² Since any other rearrangement $\alpha = \beta + tu$ that describes Δ would result in either negative components of α (for $t < 0$) or all components positive (for $t > 0$), both of which would fail to satisfy (11), the representation is unique. Additionally, any other path representation would require $t > 0$ and $\|\alpha\| > \|\beta\|$. Therefore, $\|\beta\|$ describes the shortest path length. We define $\phi(\Delta) = \|\beta\|$, where β is the unique vector satisfying (11). We can now state that the rearrangement probability is given by

$$f(X, Y) = g(\Delta_{XY}) = q^{[\phi(\Delta_{XY})]}.$$

To simplify notation, we let $l_{ij} = \phi(\Delta_{X_i X_j})$. Then

$$(14) \quad p_{ij} = s^k \frac{1}{1 - q^k} q^{l_{ij}}$$

gives the transition probability *from an unblocked state to an unblocked state*.

3.2. Transition Probabilities for Blocked States. Transition probabilities *from* blocked states are trivial: since the system does not care about past Bernoulli trials, transition probabilities depend only upon the initial configuration, not whether that configuration was the result of a blockage. So if we define a subtraction operation of 1^* by $m - 1^* := m - 1$, and calculate Δ using that arithmetic, transition probabilities *from* blocked states are the same as the transition probability from the unblocked state of the same physical configuration.

To determine transition probabilities *to* blocked states, we note the following: when worker j is *blocked*, his final action is the result of a Bernoulli failure, with no Bernoulli success. So if we compare the two transitions from some state Y ,

$$(15) \quad \begin{aligned} T_1 : & Y \rightarrow (x_1, \dots, x_{j-1}, 1, x_{j+1}, \dots, x_k), \\ T_2 : & Y \rightarrow (x_1, \dots, x_{j-1}, 1^*, x_{j+1}, \dots, x_k), \end{aligned}$$

we can express one transition probability in terms of the other. Specifically,

$$(16) \quad P[T_2] = P[T_1] \cdot \frac{q}{s},$$

indicating that regardless of the specific collection of motions of the workers, transition T_2 requires one fewer success and one additional failure when compared to the T_1 transition. We apply (16) recursively to (14) (once for each blockage) to yield

$$(17) \quad p_{ij} = r^{b_j} \frac{s^k}{1 - q^k} q^{l_{ij}},$$

²In this paper, the only vector norm used is the 1-norm. For ease of notation, we drop the subscript notation for vector norms.

where $r = q/s$ and b_j is the number of blockages in state X_j . This formula reduces to (14) when $b_j = 0$, and therefore it applies to all transitions.

4. Stationary Distribution. A direct approach for finding stationary distributions is to develop a *state transition matrix*, $P = [p_{ij}]$. Then the unit left eigenvector of P associated with eigenvalue $\lambda = 1$ gives a stationary distribution of the states [3]. However, finding such eigenvectors for a general problem is difficult unless P is very well structured. Although we have a formula for p_{ij} , we can find no “good” ordering of the states, where we interpret “good” as giving enough structure to the state transition matrix to directly determine the eigenvector for the general problem of k workers and n bins.

Often, finding general solutions requires some insight (in other words, a good guess). Some of the observations that lead to a “good guess” for this problem are the following: (1) since the probabilities are the same for each worker, we expect significant symmetry in the problem; (2) the Δ vector description of transitions yields a random walk, and we see no reason for the random walk to prefer any particular unblocked state; and (3) blocked states are closely related to the associated physical configuration, acting like boundaries on the state space, and that relationship is governed by the ratio $r = q/s$ and the number of blockages. So motivated, we postulate an initial distribution and show that it is stationary.

Postulated Density and Outline of Proof of Stationarity. Let $\{X_i\}_{i=1}^M$ be the collection of all the states and b_i the number of blockages in state X_i . Define

$$(18) \quad \omega_i = r^{b_i}.$$

We will treat these as weighting values for the density in each state and let

$$(19) \quad \nu = \frac{1}{\|\omega\|} \omega$$

be the initial distribution of states. We want to show that ν is stationary or, equivalently, that $\nu_j = \sum_{i=1}^M \nu_i p_{ij}$ for all j . One way to describe this balance is to say that at each time step, the density that transitions *from* state X_j to all states is exactly equal to the sum of all the densities that transition *into* X_j . To establish the required balance, we will prove the following logical argument:

1. Regardless of the density at time τ , at time $\tau + 1$, the density of a state X with b blockages is r^b times the weight of the unblocked state which has the same physical configuration as X .
2. Each unblocked state can be placed in a one-to-one correspondence with another unblocked state (which we call the *co-state*), which has the property that when the density is ν , the density that transitions *into* the state is the same as transitions *from* the co-state. Since ν assigns the same density to all unblocked states, the implication is that ν is stationary on the unblocked states.
3. Since the densities of the blocked states are completely determined by item 1 (which is satisfied by ν), stationarity of the unblocked states immediately implies stationarity of the blocked states.

The first element of the proof is straightforward.

LEMMA 1 (item (1)). *Let W^b be a state with b blockages and let W be the state with the same physical arrangement as W^b , but with no blockages. If the total density that is transitioned into state W is α , then the total density that transitions into state W^b will be $r^b \alpha$.*

Proof. From repeated application of (16), we know that for every state X ,

$$(20) \quad p_{XW^b} = r^b p_{XW}.$$

Suppose the density of state X_i at time τ is ξ_i . Let

$$\alpha = \sum_i \xi_i p_{X_i W}$$

be the total density transitioning into W . Then the density transitioning into W^b is

$$(21) \quad \sum_i \xi_i p_{X_i W^b} = \sum_i \xi_i r^b p_{X_i W} = r^b \alpha. \quad \square$$

Item 3 of the proof follows directly from items 1 and 2, so the crux of the argument is rigorously proving that the unblocked states are stationary. Although an algebraic argument can be shown, it is essentially a tedious symbolic manipulation that fails to reveal the underlying structure. Instead, we will develop a proof that exploits (and illuminates) the symmetry that is visible in the graph-theoretic description (the rearrangement digraph) previously introduced.

Motivating, Defining, and Analyzing Co-states. We want to show that $\nu_j(\tau + 1) = \nu_j(\tau)$ —that the density of state j is the same at time $\tau + 1$ as at time τ . Direct computation is intractable. Instead, we show that for each X_j , we can match it with another indexed state, X_{j^*} , such that $\nu_j(\tau + 1) = \nu_{j^*}(\tau)$. By assumption on ν , we know that if both states are unblocked, then $\nu_j(\tau) = \nu_{j^*}(\tau)$, and we achieve our goal (for unblocked states). An alternate description is that we seek to find a state X_{j^*} so that the density that the transitions out of state X_{j^*} is the same as that transitioning into state X_j . Algebraically, we want

$$(22) \quad \sum_{i=1}^M \nu_{j^*} p_{j^* i} = \sum_{i=1}^M \nu_i p_{ij}.$$

We first consider a slightly different constraint from (22). Suppose we seek j^* such that

$$(23) \quad \sum_{i=1}^C \nu_{j^*} p_{j^* i} = \sum_{i=1}^C \nu_i p_{ij},$$

where the summation is over only unblocked states. Since $\nu_i = \nu_j$ for all unblocked states, we can use (14) to simplify to

$$(24) \quad \sum_{i=1}^C q[\phi(\Delta_{X_{j^*} X_i})] = \sum_{i=1}^C q[\phi(\Delta_{X_i X_j})].$$

Since q is arbitrary, (24) requires the equality of two polynomial expressions in q , which implies that both the left and right sides must have the same collection of exponents. In the digraph representation, each exponent ($\phi(\Delta)$) is determined from the shortest path length between vertices. So (24) will hold if the collection of paths going into X_j is the same as the collection going out of X_{j^*} . When such a relationship holds, we call X_{j^*} the co-state of X_j . Figure 5 illustrates such a pairing for our example problem, where $X_j = (2, 1, 4)$, with co-state $(4, 1, 2)$.

Having now restated the issue as a graph question, we introduce the following definitions from graph theory [1]:




Fig. 5 Using $k = 3$ and $n = 7$, rearrangements into state $(2,1,4)$ and from its co-state $(4,1,2)$ are decomposed by δ vectors, to illustrate an equivalent set of path lengths.

- A digraph D_1 is *isomorphic* to digraph D_2 if there is a one-to-one mapping Φ from the vertex set of D_1 to the vertex set of D_2 such that (u, v) is an edge of D_1 if and only if $(\Phi u, \Phi v)$ is an edge of D_2 .
- The *converse* of graph D (written \overleftarrow{D}) is the digraph with the same vertex set as D , such that (u, v) is an edge of D if and only if (v, u) is an edge of \overleftarrow{D} . (The converse is formed by reversing the direction of each edge.)
- A digraph D is *self-converse* if \overleftarrow{D} is isomorphic to D .

LEMMA 2. Let G be the digraph formed by taking physical configurations (unblocked states) as vertices and edges defined by the rearrangement resulting from a single Bernoulli failure δ_i . Then G is self-converse.

To prove the lemma, we demonstrate the self-converse structure by providing the necessary isomorphism.

Proof. For vertex $X = (x_1, \dots, x_k)$, define map Φ (and notation $(\cdot)^*$) by

$$(25) \quad X^* := \Phi(X) \equiv (x_k, \dots, x_1).$$

Φ is clearly a one-to-one mapping of the vertices of G onto itself. Assume (X, Y) is an edge of G (and, by definition, (Y, X) is an edge in \overleftarrow{G}). Then $X + \delta_i = Y$ for some i . Since vector addition in \mathbb{R}^k is defined componentwise, reversing the order of components (by applying Φ) does not alter the quality, so we have

$$(26) \quad X^* + \delta_i^* = Y^*.$$

By referring to the definition of δ_i , one easily sees that $\delta_i^* = -\delta_{k-i}$. (Example: If $k = 6$, then $\delta_2 = (0, 1, -1, 0, 0, 0)$ and $\delta_2^* = (0, 0, 0, -1, 1, 0) = -\delta_4$.) Substituting into (26) and rearranging, we have

$$Y^* + \delta_{k-i} = X^*,$$

which implies that there is an edge in G from Y^* to X^* . So (Y, X) being an edge of \overleftarrow{G} implies that $(\Phi Y, \Phi X)$ is an edge of G . Since each step above is reversible, Φ is the necessary isomorphism between G and \overleftarrow{G} . \square

Figure 6 illustrates the self-converse property of our example case.

We remark that the k -fold symmetry of the graphs leads to k different isomorphisms that prove the self-converse property. However, we will consider (25) as defining co-states. The graphical structure *from* X^* exactly matches the structure *into* X , which ensures that (24) is satisfied and establishes the balance among the unblocked states. Moreover, the self-converse structure tells us that for any two unblocked states X and Y ,

$$p_{XY} = p_{Y^*X^*}.$$




Fig. 6 Using $k = 3$ and $n = 7$, the first graph is G and the second is its converse, formed by reversing the direction of all the edges. The third graph is simply a redrawing of the converse by reflecting the image over the hashed line. The reflection preserves the direction of adjacency between vertices, so the graph is preserved. The new redrawing more clearly shows that G is self-converse.

However to prove that the unblocked states are stationary, we must demonstrate an additional balance between the density transitioning into X_j from blocked states with the density transitioning from X_j^* to blocked states.

Let W^b be a state with b blockages and W the unblocked state with the same physical configuration as W^b . We recall that transition probabilities from W^b are the same as the transition probabilities from W (because the system does not have “memory” of how it got to some configuration). If the current density is ν , then the density of W^b is $r^b/\|\omega\|$ (by the definition of ν). The density that transitions from blocked state W^b to any unblocked state X is

$$(27) \quad p_{WX} \frac{r^b}{\|\omega\|}.$$

This is exactly the same as the density which transitions from X^* to $\Phi(W^b)$, computed as follows: the density of X^* is $1/\|\omega\|$, so the density transition from X^* to W^* is $\frac{1}{\|\omega\|} p_{X^*W^*}$. We apply (20) b times to find that the density transitioned is

$$(28) \quad \frac{1}{\|\omega\|} r^b p_{X^*W^*} = \frac{r^b}{\|\omega\|} p_{WX}.$$

This additional balance of transitioning density allows us to conclude that under assumed density ν , the unblocked state density is stationary. To complete the proof, we invoke Lemma 1 to claim that the blocked states are stationary as well. Additionally, since $p_{ij} > 0$ for all transition probabilities, it is known that the stationary density is unique [3].

5. Calculating the Expected Value for Blockages. Finding the stationary density for a Markov process enables us to answer myriad questions regarding the asymptotic behavior of the system. In particular, we are now ready to tackle the question of the original problem statement. Specifically, we seek to develop a formula for $B(k, n, r) =$ the expected value of the fraction of the time steps for which a given worker is blocked. Since the problem is symmetric in each worker, without loss of generality, we analyze worker 1. ν_i gives the expected value of the fraction of the time spent in state i , so

$$(29) \quad B(k, n, r) = \sum_{i \in \mathcal{A}} \nu_i,$$

where \mathcal{A} is the index of all the states when worker 1 is blocked. We note that

$$(30) \quad \nu_i = \frac{\omega_i}{\|\omega\|} = \frac{r^{b_i}}{\|\omega\|},$$

where b_i is the number of blockages in state X_i . We first calculate $\|\omega\|$ by applying (3):

$$(31) \quad \|\omega\| = \sum_{b=0}^{k-1} N(b, k, n) r^b = \sum_{b=0}^{k-1} \binom{k}{b} \binom{n-b-1}{k-b-1} r^b.$$

We now consider the set \mathcal{A} : When worker 1 is blocked, the state is given by

$$(32) \quad X = (1^*, x_2, \dots, x_k).$$

So we may formally define

$$(33) \quad \mathcal{A} = \left\{ j \left| X_j = (1^*, x_2, \dots, x_k), \sum_{i=2}^k |x_i| = n - 1 \right. \right\}.$$

We note that for any state $X_i, i \in \mathcal{A}$, the total number of blockages is exactly one more than the number of blockages in the second through k th components. We can use our previously determined counting formulas, but with these updated arguments, computing $N(b, k - 1, n - 1)$, where b may range from 0 to $k - 2$. We then find that the frequency of blockage of worker 1 is given by

$$(34) \quad B(k, n, r) = \sum_{i \in \mathcal{A}} \nu_i = \sum_{i \in \mathcal{A}} \frac{r^{b_i}}{\|\omega\|} = \frac{\sum_{b=0}^{k-2} \binom{k-1}{b} \binom{n-b-2}{k-b-2} r^{b+1}}{\sum_{b=0}^{k-1} \binom{k}{b} \binom{n-b-1}{k-b-1} r^b}.$$

(Note: The exponent r^{b+1} accounts for the formulation of one 1^* in the first component and b 1^* 's among components x_2, \dots, x_k .)

6. Conclusion. In this paper, we considered the problem of single-file motion of multiple agents on a closed path. By using a moving reference frame of distance between workers (instead of actual bin location), we were able to simplify the geometry of the constraint. In the $(k - 1)$ -dimensional state space, the system reduces to a random walk with the coordinate planes as boundaries. Distinct from the usual state transition graph, we proposed a decomposition digraph for transitions that allowed calculation of transition probabilities. The self-converse structure of that digraph allowed us to identify the stationary density of the process. We expect that the geometric approach may simplify similar boundary descriptions of other single-file diffusion problems. Additionally, we suspect that similar symmetries may be exploitable by decomposition of these processes into fundamental motion digraphs.

The original motivation for this problem was an actual operations analysis question involving zoning arrangements to achieve maximum worker efficiency in warehousing systems. Our simplified expression of that problem (as stated in this paper) implies that workers may move arbitrarily far around the warehouse in one time step before achieving a Bernoulli success. We call this situation the “infinite travel speed” problem. As a more realistic model, one would consider an additional constraint that

a worker could have at most f failures during a time step, since there must be some maximum speed of travel in a real system. The closed form description of the resultant stationary density remains an open problem.

REFERENCES

- [1] M. BEHZAD, G. CHARTRAND, AND L. LESNIAK-FOSTER, *Graphs and Digraphs*, Prindle, Weber & Schmidt, Boston, MA, 1979.
- [2] D. BEN-AVRAHAM, B. M. JOHNSON, C. A. MONACO, P. L. KRAPIVSKY, AND S. REDNER, *Ordering of random walks: The leader and the laggard*, *J. Phys. A: Math. Gen.*, 36 (2003), pp. 1789–1799.
- [3] G. GRIMMETT AND D. STIRZAKER, *Probability and Random Processes*, Oxford University Press, New York, 2001.