Navigation-based Optimization of Stochastic Strategies for Allocating a Robot Swarm among Multiple Sites

Spring Berman, Ádám Halász, M. Ani Hsieh and Vijay Kumar

Abstract—We present a decentralized, communication-less approach to the dynamic allocation of a swarm of homogeneous robots to a target distribution among multiple sites. Building on our work in [1], we optimize stochastic control policies for the robots that cause the population to quickly redistribute among the sites while adhering to a limit on inter-site traffic at equilibrium. We propose a way to account for delays due to navigation between sites in our controller synthesis procedure. Control policies that are designed with and without the use of delay statistics are compared for a simulation in which 240 robots distribute themselves among four buildings.

I. INTRODUCTION

We address the problem of quickly and efficiently deploying a swarm of homogeneous robots to occupy multiple locations in a predefined distribution for parallel task execution at each site. The robots must autonomously redistribute themselves among the sites to ensure task completion in the presence of robot failures or changes in the environment. This problem is relevant to applications such as the surveillance of multiple buildings, search-and-rescue, and largescale environmental monitoring.

In the multi-robot domain, methods to optimally allocate robots to tasks or resources often reduce to market-based approaches [2], [3], in which robots execute complex bidding schemes to determine an allocation based on perceived costs and utilities. While these approaches have been successful in various applications, optimality is often sacrificed to reduce the computation and communication requirements, which scale poorly as the number of robots and tasks increase. Another approach is to model the swarm as a partial differential equation and use a centralized optimal control strategy [4].

As the number of robots increases, it becomes less likely that resource-constrained robots will always have the communication and computational capabilities required for centralized control. It therefore makes sense to consider a decentralized approach to allocation that is efficient, scalable in the number of robots and sites, robust to changes in robot population, and uses little to no communication. Some work on task allocation has used the self-organized behavior of

S. Berman and V. Kumar are with the GRASP Laboratory, University of Pennsylvania, 3330 Walnut Street, Philadelphia, PA 19104, USA, {spring, kumar}@grasp.upenn.edu

A. Halász is with the Dept. of Mathematics, West Virginia University, 307G Armstrong Hall, Morgantown, WV 26506, USA, halasz@math.wvu.edu

M. A. Hsieh is with the Dept. of Mechanical Engineering and Mechanics, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104, USA, mhsiehl@drexel.edu

We gratefully acknowledge partial support from NSF grants CSR-CPS 0720801, IIS-0427313, NSF IIP-0742304, and IIS-0413138, ARO grant W911NF-05-1-0219, and ONR grant N00014-07-1-0829.

insect colonies as inspiration for decentralized strategies in which robots switch between simple behaviors in response to local sensing [5]–[7]. We have adopted this distributed paradigm, inspired in particular by the process of house-hunting in ant colonies [8].

In recent work on decentralized control for task allocation, a physical multi-robot system is abstracted to an accurate differential equation model [9], [10]. In contrast to this "bottomup" analysis procedure, our methodology is based on a "top-down" design approach that gives theoretical guarantees on performance. The robots in our scenario redistribute themselves by switching stochastically between pairs of sites at probability rates specific to each transition. Our strategy is to use a continuous approximation of the system to design these rates to achieve our global objective.

We first applied this approach to the design of ant-inspired behaviors that produce a predefined swarm allocation between two sites [11], [12]. We extended our methodology to the problem of redistributing a swarm among many sites [1] and introduced quorum-based control policies [13]. In [1], we noted the idea of including in the analysis the effects of robot navigation between sites when the travel times are not negligible compared to the waiting times for transitions. In this work, we extend our controller design methodology to account for inter-site travel times. We optimize the rates at which robots switch between sites for fast convergence to a desired distribution subject to a constraint on equilibrium traffic between sites. We compare the resulting control policies using a four-site surveillance simulation in which robot travel time distributions have significant variability.

II. PROBLEM STATEMENT

A. Definitions and Assumptions

Consider N robots to be distributed among M sites. We denote the number of robots at site $i \in \{1, ..., M\}$ at time t by $n_i(t)$ and the desired number of robots at site i by n_i^d , a positive integer. The population fraction at site i at time t is $x_i(t) = n_i(t)/N$. Then the system state vector is given by $\mathbf{x}(t) = [x_1(t) \dots x_M(t)]^T$. We define the *target distribution* as a set of desired population fractions to occupy each site, $\mathbf{x}^d = [x_1^d \dots x_M^d]^T$, where $x_i^d = n_i^d/N$. A specification in terms of fractions rather than absolute robot numbers is useful for scaling and for applications in which losses of robots to attrition and breakdown are common.

The interconnection topology of the sites can be modeled as a directed graph, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} , the set of vertices, represents sites $\{1, \ldots, M\}$ and \mathcal{E} , the set of $N_{\mathcal{E}}$ edges, represents physical one-way routes between sites. Sites *i* and j are *adjacent*, denoted by $i \sim j$, if robots can travel from i to j. We represent this relation by the ordered pair $(i, j) \in \mathcal{V} \times \mathcal{V}$, with the set $\mathcal{E} = \{(i, j) \in \mathcal{V} \times \mathcal{V} \mid i \sim j\}$. More generally, we can define \mathcal{V} as a set of M tasks and \mathcal{E} as the set of possible transitions between tasks; then \mathcal{G} models precedence constraints between the tasks. It is assumed that \mathcal{G} is *strongly connected*, *i.e.* a directed path exists between any pair of distinct vertices. This property facilitates redistribution by allowing robots to travel to any site from any other site; no sites act as sources or sinks.

We consider $\mathbf{x}(t)$ to represent the distribution of the state of a Markov process on \mathcal{G} , for which \mathcal{V} is the state space and \mathcal{E} is the set of possible transitions. Every edge in \mathcal{E} is assigned a constant positive *transition rate*, k_{ij} , which defines the probability per unit time for one robot at site *i* to go to site *j*. It follows that the number of transitions between two adjacent sites in a time interval *t* has a Poisson distribution with parameter $k_{ij}t$. We use *constant* k_{ij} in order to be able to abstract the system to a continuous model (see Section II-B). In general, $k_{ij} \neq k_{ji}$.

We assume that each robot has knowledge of \mathcal{G} , all k_{ij} , and the task to perform at each site, as well as the behaviors necessary to navigate between sites and execute the tasks. We also assume that each robot has a map of the environment and can localize itself and sense neighboring robots.

B. Base Model

Our strategy for redistributing robots among the sites is to program each robot to switch stochastically from site *i* to site *j* with probability $k_{ij}\Delta t$ at each time step Δt [1]. In the limit $N \rightarrow \infty$, the physical system of individual robots can be abstracted to a linear ordinary differential equation (ODE) model according to the theoretical justification provided by [14]. The model quantifies $\dot{x}_i(t)$ as the difference between the total influx and total outflux of robots at site *i*,

$$\dot{x}_i(t) = \sum_{\forall j \mid (j,i) \in \mathcal{E}} k_{ji} x_j(t) - \sum_{\forall j \mid (i,j) \in \mathcal{E}} k_{ij} x_i(t) .$$
(1)

Then the system of equations for all M sites is given by

$$\dot{\mathbf{x}} = \mathbf{K}\mathbf{x} , \qquad (2)$$

where $\mathbf{K} \in \mathbb{R}^{M \times M}$ is a matrix defined as

$$\mathbf{K}_{ij} = \begin{cases} k_{ji} & \text{if } i \neq j , \quad (j,i) \in \mathcal{E} ,\\ 0 & \text{if } i \neq j , \quad (j,i) \notin \mathcal{E} , \\ -\sum_{(i,l)\in \mathcal{E}} k_{il} & \text{if } i = j . \end{cases}$$
(3)

Since the number of robots is conserved, the population fractions satisfy the equation

$$\mathbf{1}^T \mathbf{x} = 1 \ . \tag{4}$$

We will refer to equation (2) subject to (4) as the *switching model*, since it describes a system in which robots switch instantaneously between sites. The following theorem was proved in [1].

Theorem 1: If the graph \mathcal{G} is strongly connected, then the switching model has a unique, stable equilibrium $\bar{\mathbf{x}}$.

This equilibrium can be calculated as [15]:

$$\bar{x}_i = K_{ii} / \sum_{j=1}^M K_{jj} , \quad i = 1, ..., M$$
, (5)

where K_{ij} is the cofactor of **K** obtained by deleting row *i* and column *j*.

Theorem 1 implies that we can achieve the target distribution $\mathbf{x}^{\mathbf{d}}$ from any initial distribution by specifying that $\bar{\mathbf{x}} \equiv \mathbf{x}^{\mathbf{d}}$ through the following constraint on **K**,

$$\mathbf{K}\mathbf{x}^{\mathbf{d}} = 0 \ . \tag{6}$$

When the k_{ij} are chosen to satisfy (6), robots that use the k_{ij} as stochastic transition rules will collectively occupy the sites in distribution \mathbf{x}^{d} at steady state, assuming instant switching.

C. Time-Delayed Model

In reality, the influx of robots to site j from site i is delayed by the time taken to travel between the sites, τ_{ij} . If we assume a constant delay τ_{ij} for each edge (i, j), this effect can be included by rewriting equation (1) as a delay differential equation (DDE):

$$\dot{x}_i(t) = \sum_{\forall j \mid (j,i) \in \mathcal{E}} k_{ji} x_j(t - \tau_{ji}) - \sum_{\forall j \mid (i,j) \in \mathcal{E}} k_{ij} x_i(t) .$$
(7)

Modeling the time delays has the effect that $\mathbf{1}^T \mathbf{x}(t) < 1$ for t > 0, since some robots are traveling between sites. Let $n_{ij}(t)$ be the number of robots traveling from site *i* to site *j* at time *t* and $y_{ij}(t) = n_{ij}(t)/N$. Then the conservation equation for this system is:

$$\sum_{i=1}^{M} x_i(t) + \sum_{i=1}^{M} \sum_{\forall j \mid (i,j) \in \mathcal{E}} y_{ij}(t) = 1 .$$
 (8)

III. ANALYSIS

In application, robot travel times between sites can be highly variable due to changes in navigation patterns caused by collision avoidance, crowding, and errors in localization. Hence, model (7) can be made more realistic by defining the delays τ_{ij} as random variables, T_{ij} . A reasonable form for the probability density of the T_{ij} can be estimated from an analogous scenario in which vehicles deliver items along roads to various sites. Vehicle travel times in this system have been modeled as following an Erlang distribution to capture the properties that the times have minimum possible values, a small probability of being large due to accidents, breakdowns, and low energy, and their distributions tend to be skewed toward larger values [16]. We assume that each T_{ij} follows this distribution with parameters ω_{ij} , a positive integer, and θ_{ij} , a positive real number:

$$g(t;\omega_{ij},\theta_{ij}) = \frac{\theta_{ij}^{\omega_{ij}}t^{\omega_{ij}-1}}{(\omega_{ij}-1)!}e^{-\theta_{ij}t} .$$
(9)

In practice, the parameters are estimated by fitting empirical travel time data to density (9).

Under this assumption, the DDE model (7) can be transformed into an equivalent ODE model of the form (1), which



Fig. 1. A labeled edge (i, j) = (1, 2) that consists of (a) the physical sites, corresponding to model (2), and (b) both physical and virtual sites (for $\omega_{12} = 2$), corresponding to model (10).

allows us to optimize the rates k_{ij} using the method we develop for this type of model. We use the fact that T_{ij} has the same distribution as the sum of ω_{ij} independent random variables, $T_1, ..., T_{\omega_{ij}}$, with a common distribution $f(t; \theta_{ij}) = \theta_{ij} e^{-\theta_{ij}t}$ [17]. Each of the variables represents a portion of the travel time between sites i and j. To model these portions of the journey, we define a directed path composed of a sequence of virtual sites, $u = 1, ..., \omega_{ij}$, between the physical sites i and j. Assume that robots transition instantaneously from virtual site u to u+1, which is site j when $u = \omega_{ij}$, at a constant probability per unit time, θ_{ij} . It follows that $f(t; \theta_{ij})$ is the distribution of the time that a robot spends at virtual site u, and so we can define T_u as this site occupancy time.

The expected value of T_u is $E(T_u) = \theta_{ij}^{-1}$. Using the property $E(T_{ij}) = \sum_{u=1}^{\omega_{ij}} E(T_u)$, we see that $\theta_{ij} = \omega_{ij}/E(T_{ij})$. Then the variance of T_{ij} is $Var(T_{ij}) = E(T_{ij})^2/\omega_{ij}$. Note that $Var(T_{ij}) \to 0$ as $\omega_{ij} \to \infty$; in this case the system is described by model (7), assuming that each $\tau_{ij} = E(T_{ij})$.

The population fraction at virtual site u along edge (i, j) will be denoted by $y_{ij}^{(u)}$. Then $\sum_{u=1}^{\omega_{ij}} y_{ij}^{(u)}$ represents y_{ij} , the fraction of robots traveling from site i to j. Fig. 1 illustrates how an edge from model (1) is expanded with two virtual states $y_{ij}^{(u)}$. The dynamics of the population fractions at all physical and virtual sites in the expanded system can be written as a set of linear ODE's, as in Section II-B:

$$\dot{x}_{i}(t) = \sum_{j|(j,i)\in\mathcal{E}} \theta_{ji} y_{ji}^{(\omega_{ji})}(t) - \sum_{j|(i,j)\in\mathcal{E}} k_{ij} x_{i}(t) ,
\dot{y}_{ij}^{(1)}(t) = k_{ij} x_{i}(t) - \theta_{ij} y_{ij}^{(1)}(t) ,
\dot{y}_{ij}^{(m)}(t) = \theta_{ij} \left(y_{ij}^{(m-1)}(t) - y_{ij}^{(m)}(t) \right) ,
m = 2, ..., \omega_{ij} ,$$
(10)

where i = 1, ..., M and $(i, j) \in \mathcal{E}$. Let **y** be the vector of $y_{ij}^{(u)}$, $u = 1, ..., \omega_{ij}$, $(i, j) \in \mathcal{E}$. The system state vector is then $\mathbf{z} = [\mathbf{x} \ \mathbf{y}]^T$. We can interpret each component of z as the population fraction at a site $i \in \{1, ..., M'\}$, where M' is the sum of all physical and virtual sites. The interconnection topology of these sites can be modeled as a directed graph, $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$, where $\mathcal{V}' =$ $\{1, ..., M'\}$ and $\mathcal{E}' = \{(i, j) \in \mathcal{V}' \times \mathcal{V}' \mid i \sim j\}$. Since \mathcal{G} is strongly connected, so is \mathcal{G}' . Then the ODE model (10) can be written in the form of model (2):

$$\dot{\mathbf{z}} = \hat{\mathbf{K}}\mathbf{z} , \qquad (11)$$

where $\hat{\mathbf{K}} \in \mathbb{R}^{M' \times M'}$ has structure (3) with entries \hat{k}_{ij} (in place of k_{ij}) defined by the corresponding coefficients in model (10). The conservation equation (8) can be written as

$$\mathbf{1}^T \mathbf{z} = 1 \ . \tag{12}$$

We will refer to system (11) subject to (12) as the chain model, since it incorporates a chain of virtual sites between each pair of physical sites.

At equilibrium, the incoming and outgoing flux at each virtual site along the path from site i to j is $k_{ij}\bar{x}_i$, yielding the following equilibrium values of $y_{ij}^{(u)}$, $u = 1, ..., \omega_{ij}$:

$$\bar{y}_{ij}^{(u)} = k_{ij}\bar{x}_i/\theta_{ij}$$
 (13)

Substituting $\bar{y}_{ij} = \sum_{u=1}^{\omega_{ij}} \bar{y}_{ij}^{(u)}$ into equation (8) gives the conservation equation for this system at equilibrium:

$$\sum_{i=1}^{M} \bar{x}_i \left(1 + \sum_{j \mid (i,j) \in \mathcal{E}} k_{ij} \omega_{ij} / \theta_{ij} \right) = 1$$
(14)

The equilibrium values \bar{x}_i can be shown to be [15]:

$$\bar{x}_i = K_{ii} / \sum_{p=1}^{M} (1 + \sum_{j \mid (p,j) \in \mathcal{E}} k_{pj} \omega_{pj} / \theta_{pj}) K_{pp} , \quad i = 1, ..., M .$$
(15)

Comparing the equilibrium values (15) of the chain model with the values (5) of the corresponding switching model, it is evident that the ratio of \bar{x}_i between any two sites is the same in both models. However, since $k_{pj}\omega_{pj}/\theta_{pj} > 0$, the \bar{x}_i of the chain model are *lower* than those of the switching model. The following theorem shows that the equilibrium distribution will be achieved from any initial state.

Theorem 2: If \mathcal{G} is strongly connected, then the chain model has a unique, stable equilibrium given by (13), (15).

Proof: Since the system can be represented in the same form as model (2) subject to (4). Theorem 1 can be applied to show that there is a unique, stable equilibrium.

IV. METHODOLOGY

We consider the problem of computing the rates k_{ij} that cause a swarm of robots, modeled as system (2) or (11), to redeploy from an initial distribution to a target distribution. The redistribution can be made arbitrarily fast by choosing high k_{ii} , since the rates of convergence of systems (2) and (11) are governed by the real parts of the eigenvalues of K and $\mathbf{\hat{K}}$, respectively, which are positive homogenous functions of the k_{ij} [18]. However, as shown by equation (13), raising k_{ij} increases the equilibrium fraction of travelers on the route corresponding to edge (i, j). This extraneous traffic between sites at equilibrium expends power and can lead to backups due to congestion. Thus, when choosing the k_{ij} , we are faced with a tradeoff between rapid equilibration and long-term system efficiency, i.e. few idle trips between sites once the target distribution is achieved. In light of this tradeoff, we define our objective as the design of an optimal transition *rate matrix* \mathbf{K}^* or $\hat{\mathbf{K}}^*$ that maximizes the convergence rate of the system to the target distribution while not exceeding a limit on the inter-site traffic at equilibrium.

In this section, we outline our methodology of determining \mathbf{K}^* and $\hat{\mathbf{K}}^*$ for a simulated scenario in which a swarm of robots surveys the perimeters of several buildings while reallocating to a target distribution among the buildings.

A. Surveillance simulation

1) Robot motion control: Each robot is represented as a planar agent governed by a kinematic model. A robot that is monitoring a building circulates around the perimeter by aligning its velocity vector with the straight lines that comprise the perimeter; this motion can also be achieved with feedback controllers of the form given in [19]. The robot slows down if a robot in front of it enters its sensing range, which results in an approximately uniform distribution of robots around the perimeter.

To implement inter-site navigation, we first performed a convex cell decomposition of the free space. This resulted in a discrete roadmap on which shortest-path computations between cells can be obtained using any standard graph search algorithm. Each edge $(i, j) \in \mathcal{E}$ is defined as a sequence of cells to be traversed by robots traveling from a distinct exit point on the perimeter of building *i* to an entry point on the perimeter of building j. We used Dijkstra's algorithm to compute the sequence with the shortest cumulative distance between cell centroids, starting from the cell adjacent to the exit at i and ending at the cell adjacent to the entrance at j. The robots are provided a priori with the sequence of cells corresponding to each edge. Navigation between cells is achieved by composing local potential functions such that the resulting control policy ensures arrival at the last cell in the sequence [20]. We combine these navigation controllers with ones derived from repulsive potential functions to achieve inter-robot collision avoidance [21]. At each time step, the robots compute the feedback controller to move from one cell to the next based on their current position and the positions of robots within their sensing ranges.

2) Site-to-site transitions: Gillespie's Direct Method [14] was used to simulate a sequence of robot site transition events and their initiation times using the rates k_{ij} from the switching model or chain model. Each event is identified with the commitment of an individual robot to travel to another site. A transition from building *i* to *j* is assigned to a random robot on the perimeter of *i*. This robot continues to track the perimeter until it reaches the exit for edge (i, j), at which point it begins navigating to the entrance on building *j*. For more details on this stochastic simulation, see [11] and [13].

The travel time τ_{ij} is measured as the the sum of τ_{ij}^a , the time for a robot to reach the exit on building *i* from the position at which it commits to the transition, and τ_{ij}^b , the travel time from the exit to building *j*'s entrance. Because the robots at *i* are uniformly distributed around the perimeter and are randomly selected for transitions, τ_{ij}^a has a uniform distribution. The distribution of τ_{ij}^b is affected by the congestion on the roads and at the target sites, which determines the amount of time spent avoiding collisions.

B. Computation of \mathbf{K}^* and $\hat{\mathbf{K}}^*$

We computed \mathbf{K}^* for the switching model and $\hat{\mathbf{K}}^*$ for two versions of the chain model. To determine the parameters

of a *full chain model* that would most accurately emulate the travel time distributions of the surveillance simulation, we collected a set of τ_{ij} from the simulation for each edge (i, j), plotted a histogram of the τ_{ij} , and then fit an Erlang distribution (9) to the histogram to obtain ω_{ij} and θ_{ij} . The $\hat{\mathbf{K}}^*$ for this model is called $\hat{\mathbf{K}}^*_{full}$. We also computed a $\hat{\mathbf{K}}^*$, called $\hat{\mathbf{K}}^*_{one}$, for a *one-site chain model* in which each $\omega_{ij} =$ 1 and each θ_{ij} is $1/E(T_{ij}) = \theta_{ij}/\omega_{ij}$ from the full chain model. In this case, the Erlang distribution reduces to an exponential distribution with the same mean value.

We measure the degree of convergence to $\mathbf{x}^{\mathbf{d}}$ in terms of the *fraction of misplaced robots*, defined as the 2-norm

$$\Delta(\mathbf{x}) = ||\mathbf{x} - \mathbf{x}^{\mathbf{d}}||_2 . \tag{16}$$

We say that one system converges faster than another if it takes less time for $\Delta(\mathbf{x})$ to decrease to a small fraction, such as 0.1, of its initial value. We compute the transition rate matrix that directly minimizes this convergence time using Metropolis optimization [22] with the k_{ij} as the optimization variables. This method was chosen for its simplicity and the fact that it provides reasonable improvements in convergence time with moderate computing resources.

Let \mathbf{x}^0 be the initial robot distribution among the sites. We quantify the *traffic* associated with edge (i, j) as $k_{ij}x_i$, the fraction of robots per unit time that are exiting site i to travel along the edge. For the switching model, the objective is to find a K with structure (3) that minimizes the time for $\Delta(\mathbf{x})$ to converge to $0.1\Delta(\mathbf{x^0})$ subject to constraint (6) and a limit c on the total traffic between sites at equilibrium, $\sum_{(i,j)\in\mathcal{E}} k_{ij} x_i^d$. At each iteration, the k_{ij} are perturbed by random amounts such that the resulting K satisfies (6). Then to apply the traffic constraint, the k_{ij} are multiplied by $c / \sum_{(i,j) \in \mathcal{E}} k_{ij} x_i^d$, which maximizes the traffic capacity. Since the model is a linear system, the convergence time to $0.1\Delta(\mathbf{x^0})$ can be easily calculated. The resulting K is decomposed into its normalized eigenvectors and eigenvalues, system (2) is mapped onto the space spanned by the normalized eigenvectors, and a transformation is applied to compute $\mathbf{x}(t)$ using the matrix exponential of the diagonal matrix of eigenvalues multiplied by time. Since the system is stable by Theorem 1, $\Delta(\mathbf{x})$ always decreases monotonically with time, so a Newton scheme can be used to calculate the exact time when $\Delta(\mathbf{x})/\Delta(\mathbf{x^0}) = 0.1$.

We use the same procedure to compute $\hat{\mathbf{K}}^*$ for the chain models, with \mathbf{z} in place of \mathbf{x} , $\mathbf{z}^{\mathbf{0}} = [\mathbf{x}^{\mathbf{0}^T} \mathbf{0}]^T$, and the target distribution $\mathbf{z}^{\mathbf{d}}$ defined as the null space of $\hat{\mathbf{K}}$ at each iteration. The θ_{ij} are fixed and the k_{ij} are constrained such that the portion of $\mathbf{z}^{\mathbf{d}}$ associated with the physical sites is a fraction of $\mathbf{x}^{\mathbf{d}}$ (the remainder represents the travelers). The same traffic constraint is applied so that the switching and chain models have the same equilibrium traveler fraction, which is necessary to have a basis for comparing the system convergence rates due to the tradeoff between these properties that we discussed. Note that although this constraint is formulated in terms of the x_i^d , in practice the total traffic at equilibrium is actually $\sum_{(i,j)\in \mathcal{E}} k_{ij} dx_i^d = dc$, where d is the population fraction at the physical sites (i.e., not in transit).



Fig. 2. Cell decomposition of the free space used for navigation. The surveyed buildings are highlighted and numbered.



Fig. 3. Snapshots of a simulation using \mathbf{K}^* . The maroon (dark) robots are not engaged in a transition; the orange (light) robots have committed to travel or are in the process of traveling. Robots navigate between sites at 1.3 m/s, which is attainable by some mobile robots that are suited to surveillance tasks, such as PatrolBot[®] and Seekur[®]. The perimeter surveillance speed is 4.5 times slower. We set c = 0.06 robots/s.

V. RESULTS

To investigate the utility of the chain model in optimizing the k_{ij} , we simulated a surveillance task as described in Section IV-A with k_{ij} from the matrices \mathbf{K}^* , $\hat{\mathbf{K}}^*_{one}$, and $\hat{\mathbf{K}}^*_{full}$ computed according to Section IV-B. The swarm consists of 240 robots, and the four buildings to be monitored are located on the section of the University of Pennsylvania campus shown in Fig. 2. We used a graph \mathcal{G} for these four sites with the edges given in Table 1. The robots are initially split equally between sites 3 and 4, and they are required to redistribute to occupy all sites in equal fractions. Fig. 3 illustrates this redistribution for one trial.

The travel time data that we used to determine the chain model parameters were a set of $750 - 850 \tau_{ij}$ per edge collected from a simulation using **K**^{*}. Fig. 4 shows a sample fitting of an Erlang distribution to τ_{ij} data for one edge. Table 1 lists $E(T_{ij})$ (the average τ_{ij}) and ω_{ij} for each edge.

Fig. 5(b),(d) show that starting at ~ 15000 sec, each average traveler fraction over 40 runs oscillates close to the average equilibrium value of the \mathbf{K}^* runs. Thus, \mathbf{K}^* , $\hat{\mathbf{K}}^*_{one}$,



Fig. 4. Histogram of the travel times from site 1 to site 4 (758 data points) and the approximate Erlang distribution.

and $\hat{\mathbf{K}}_{full}^*$ yield approximately the same equilibrium intersite traffic. Fig. 5(a),(b) show that \mathbf{K}^* and $\hat{\mathbf{K}}_{one}^*$ produce very similar average $\Delta(\mathbf{x})$, average traveler fraction, and associated standard deviations. The same can be said of the results for \mathbf{K}^* and $\hat{\mathbf{K}}_{full}^*$ in Fig. 5(c),(d), since the relatively high standard deviations indicate that any disparities may not be significant. One disparity is the slightly lower average traveler fraction for $\hat{\mathbf{K}}_{full}^*$ than for \mathbf{K}^* during the transient phase, which implies that the $\hat{\mathbf{K}}_{full}^*$ runs achieve the target distribution in about the same amount of time as the \mathbf{K}^* runs using less energy to travel between sites.

Table 1. Row 1:	(i, j), row 2:	$E(T_{ij})$) (sec),	row 3:	ω_{ij}
-----------------	-------	-----------	-------------	----------	--------	---------------

(1,2)	(1,3)	(1,4)	(2,3)	(2,4)	(3,4)	(4,1)
7	757	738	556	1507	1628	1228	1072
	14	15	9	6	5	7	6

VI. DISCUSSION AND FUTURE WORK

In summary, we have extended our framework in [1] for redistributing a swarm of robots among multiple sites, or more generally tasks, to include the effects of travel times between sites. We emulate realistic travel time distributions by augmenting a linear ODE model of the swarm with virtual sites that represent the progress of traveling robots. We design the transition rates between sites in the ODE model for fast convergence to a target distribution and longterm efficiency. These rates define stochastic switching rules for individual robots, whose collective behavior follows the continuous model prediction. In this way, we synthesize decentralized robot controllers that can be computed a priori from a set of $N_{\mathcal{E}}$ rates, where $N_{\mathcal{E}}$ depends on the site graph and not on the population size. The controllers require no communication and have guarantees on performance.

The predictive value of the chain model depends on how well the travel time distributions are characterized. Combinations of Erlang distributions can be used to approximate more complicated (e.g. multimodal) distributions.

The similarity among the results for the surveillance simulations run with rates from the switching and chain models indicates that for the purpose of controller synthesis, the switching model is a sufficiently accurate representation



Fig. 5. (a),(c) Fraction of misplaced robots $\Delta(\mathbf{x})$ and (b),(d) fraction of travelers vs. time for simulations using \mathbf{K}^* , $\hat{\mathbf{K}}_{one}^*$, and $\hat{\mathbf{K}}_{full}^*$. Thick lines are averages over 40 simulation runs; thin lines mark the standard deviations. The horizontal dashed lines mark the mean equilibrium traveler fraction, 0.237, measured from the \mathbf{K}^* runs.

of our system. Hence, we can simply optimize the matrix \mathbf{K} and do not have to incur the greater computational expense that is needed to optimize the larger matrix $\hat{\mathbf{K}}$. In ongoing work on the switching model, we compare our Metropolis \mathbf{K} optimization method with others that maximize functions of the eigenvalues of \mathbf{K} , which govern the model's rate of convergence [23]. A possible extension of our work is to design a time-dependent matrix $\mathbf{K}(t)$ that causes the swarm to redistribute according to a trajectory of desired configurations, $\mathbf{x}^{\mathbf{d}}(t)$. Also, we can devise an extension of our quorum-based strategy [13] in which robots stop moving between sites (and hence expending energy) once they detect that a site is close enough to the target occupancy.

One avenue of future work is to investigate whether using the chain model to optimize the rates improves performance under different conditions. In our simulation, the average τ_{ij} for the edges are within a factor of 3 of each other. The $\hat{\mathbf{K}}^*$ for a scenario with larger differences between average τ_{ij} should assign much higher k_{ij} to edges with low τ_{ij} than to edges with high τ_{ij} since this would speed up convergence; the computation of \mathbf{K}^* does not account for the τ_{ij} and so would be expected to produce slower convergence. It may also be fruitful to study scenarios in which, for each edge (i, j), the ratio of the average τ_{ij} to k_{ij}^{-1} , the average waiting time at site *i*, is higher than in our simulations, in which the highest ratio (over all sets of k_{ij}) is 0.31. Another aspect to consider is the dependence of travel times on the robot population, which may lead to a generalized definition of "traffic capacity" that is not well described by a linear model.

REFERENCES

- Á. Halász, M. A. Hsieh, S. Berman, and V. Kumar, "Dynamic redistribution of a swarm of robots among multiple sites," in *Proc. Int'l. Conf. on Intelligent Robots and Syst. (IROS'07)*, 2007, pp. 2320– 2325.
- [2] M. B. Dias, R. M. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: a survey and analysis," *Proc. of the IEEE*, vol. 94, no. 7, pp. 1257 – 1270, 2006.
- [3] E. G. Jones, M. B. Dias, and A. Stentz, "Learning-enhanced marketbased task allocation for oversubscribed domains," in *Proc. Int'l. Conf.* on Intelligent Robots and Syst. (IROS'07), 2007, pp. 2308–2313.
- [4] D. Milutinovic and P. Lima, "Modeling and optimal centralized control of a large-size robotic population," *IEEE Trans. on Robotics*, vol. 22, no. 6, pp. 1280–1285, 2006.
- [5] T. H. Labella, M. Dorigo, and J.-L. Deneubourg, "Division of labor in a group of robots inspired by ants' foraging behavior," ACM Trans. Auton. Adapt. Syst., vol. 1, no. 1, pp. 4–25, 2006.
- [6] M. J. B. Krieger, J.-B. Billeter, and L. Keller, "Ant–like task allocation and recruitment in cooperative robots," *Nature*, vol. 406, pp. 992–995, 2000.
- [7] W. Agassounon and A. Martinoli, "Efficiency and robustness of threshold-based distributed allocation algorithms in multi-agent systems," in *Proc. First Int'l. Joint Conf. on Auton. Agents and Multi-Agent Syst. (AAMAS'02)*, 2002, pp. 1090–1097.
- [8] N. Franks, S. C. Pratt, E. B. Mallon, N. F. Britton, and D. T. Sumpter, "Information flow, opinion polling and collective intelligence in househunting social insects," *Phil. Trans. R. Soc. Lond. B*, vol. 357, no. 1427, pp. 1567–1583, 2002.
- [9] W. Agassounon, A. Martinoli, and K. Easton, "Macroscopic modeling of aggregation experiments using embodied agents in teams of constant and time-varying sizes," *Auton. Robots*, vol. 17, no. 2-3, pp. 163–192, 2004.
- [10] K. Lerman, C. V. Jones, A. Galstyan, and M. J. Mataric, "Analysis of dynamic task allocation in multi-robot systems," *Int. J. of Robotics Research*, vol. 25, no. 3, pp. 225–241, 2006.
- [11] S. Berman, Á. Halász, V. Kumar, and S. Pratt, "Algorithms for the analysis and synthesis of a bio-inspired swarm robotic system," in 9th Int'l. Conf. on the Simul. of Adapt. Behav. (SAB'06), Swarm Robotics Workshop, vol. LNCS 4433, 2007, pp. 56–70.
- [12] —, "Bio-inspired group behaviors for the deployment of a swarm of robots to multiple destinations," in *Proc. Int'l. Conf. on Robotics* and Automation (ICRA'07), 2007, pp. 2318–2323.
- [13] M. A. Hsieh, Á. Halász, S. Berman, and V. Kumar, "Biologically inspired redistribution of a swarm of robots among multiple sites," to appear a Special Issue of the Swarm Intelligence Journal, 2008.
- [14] D. Gillespie, "Stochastic simulation of chemical kinetics," Annu. Rev. Phys. Chem., vol. 58, pp. 35–55, 2007.
- [15] H. Othmer, Analysis of Complex Reaction Networks, Lecture Notes. Minneapolis, MN: Sch. of Mathematics, Univ. Minnesota, Dec. 2003.
- [16] R. Russell and T. Urban, "Vehicle routing with soft time windows and Erlang travel times," vol. 59, no. 9, pp. 1220–1228, 2008, J. Operational Research Society.
- [17] B. Harris, *Theory of Probability*. Reading, MA: Addison-Wesley, 1966.
- [18] J. Sun, S. Boyd, L. Xiao, and P. Diaconis, "The fastest mixing Markov process on a graph and a connection to the maximum variance unfolding problem," *SIAM Review*, vol. 48, no. 4, pp. 681–699, 2006.
- [19] M. A. Hsieh, S. Loizou, and V. Kumar, "Stabilization of multiple robots on stable orbits via local sensing," in *Proc. Int'l Conf. on Robotics and Automation (ICRA'07)*, Apr. 2007, pp. 2312–2317.
- [20] D. C. Conner, A. Rizzi, and H. Choset, "Composition of local potential functions for global robot control and navigation," in *Proc. Int'l. Conf.* on Intelligent Robots and Syst. (IROS'03), vol. 4, 2003, pp. 3546–3551.
- [21] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Flocking in fixed and switching networks," in *IEEE Trans. on Automatic Control*, vol. 52, no. 5, 2007, pp. 863–868.
- [22] D. P. Landau and K. Binder, A guide to Monte-Carlo simulations in statistical physics. Cambridge University Press, 2000.
- [23] S. Berman, Á. Halász, M. A. Hsieh, and V. Kumar, "Optimization of stochastic strategies for redistributing a robot swarm among multiple sites," under review for *IEEE Trans. on Robotics*.