Macroscopic Modeling of Stochastic Deployment Policies with Time Delays for Robot Ensembles

T. William Mather and M. Ani Hsieh

Abstract We consider the dynamic assignment and reassignment of a homogeneous robot ensemble to multiple spatially located tasks with deterministic or near-deterministic task execution times. Similar to Halasz et al (2007); Berman et al (2008), we consider the development of agent-level, *i.e.*, *microscopic*, stochastic control policies through the analysis of an appropriate *macroscopic* analytical model that describes the dynamics of the ensemble. Specifically, we present an approach to better approximate the effects of deterministic microscopic time delays at the macroscopic level based on Padé approximants. We present, analyze, and compare the frequency response of our approach to the one presented by Berman et al (2008) using different agent-based simulations.

1 Introduction

We consider the synthesis of agent-level stochastic control policies to allocate a homogeneous ensemble of robots to a collection of tasks through the analysis of an appropriate ensemble model. These tasks may be spatially distributed at v different locales and/or have to be executed following specific precedence constraints. This problem is relevant to various applications including the scheduling of automated transportation systems, surveillance of multiple locations for large scale environmental monitoring, or providing aerial coverage for various ground units. In all these applications, the ensemble must have the ability to distribute themselves among the various tasks/locations and autonomously redistribute to ensure task completion which may be affected by robot failures or changes in the environment.

T. William Mather
Drexel University, Philadelphia, PA 19104, e-mail: twm32@drexel.edu

M. Ani Hsieh
Drexel University, Philadelphia, PA 19104 e-mail: mhsieh1@drexel.edu

This is similar to the task/resource allocation problem where the objective is to determine the optimal assignment of robots to tasks. In the multi-robot domain, existing methods for the solution of these combinatorial optimization problems often reduce to market-based approaches (Lin and Zheng, 2005; Guerrero and Oliver, 2003; Jones et al, 2006) where robots must execute complex bidding schemes to determine the appropriate allocation based on the various perceived costs and utilities. While market-based approaches have gained much success in various multi-robot applications (Dias, 2004; Vail and Veloso, 2003; Gerkey and Mataric, 2002; Jones et al, 2007) and can be further improved when learning is incorporated (Dahl et al, 2006), these methods often scale poorly in terms of team size and number of tasks (Dias et al, 2006; Golfarelli et al, 1997). Furthermore, in applications such as mining and search and rescue, where inter-agent wireless communication may be unreliable or non-existent, it is often difficult to devise reliable strategies to ensure timely communication of the various local costs and utilities required by existing allocation approaches.

Other recent task allocation strategies include the work by Shen and Salemi (2002) where the allocation problem is formulated as a Distributed Constraint Satisfaction Problem. This approach requires the explicit modeling of tasks, their requirements, and robot capabilities which makes implementation for large populations difficult.

More recently, task allocation strategies have been devised by modeling the dynamics of the ensemble distribution across various tasks. This is achieved by developing an appropriate continuous model to describe the ensemble dynamics. The approach by Milutinovic and Lima (2006) maximizes robot occupation at a desired position using a centralized optimal control policy where the ensemble distribution is modeled using a partial differential equation. Other similar efforts include the work by Martinoli et al (2004) and Lerman et al (2006) where the continuous models are obtained by defining individual robot controllers and averaging their performance. These models are then used to describe the performance of the swarm in a collaborative stick-pulling task (Martinoli et al, 2004) and in an adaptive multi-robot foraging task (Lerman et al, 2006). The foraging task in (Lerman et al, 2006) is modeled as a stochastic process that does not involve explicit communication or global knowledge; however, the only way to control robot task reallocation is to modify the task distribution in the environment.

The dynamic distribution of a team of robots across multiple sites without the use of explicit inter-agent wireless communication was achieved by Halasz et al (2007); Hsieh et al (2008) through the use of the ensemble model to synthesize stochastic control policies for individual robots. In these works, the behavior of each robot is represented by a probabilistic finite state machine. Robots transition from one state to another based on a set of predefined transition rates. The allocation and re-allocation of the team is achieved by designing the transition rates to ensure the desired allocation across the various sites. Similar to Martinoli et al (2004) and Lerman et al (2006), these work employs a multi-level representation of the ensemble activity where the macroscopic analytical model is used to determine the set of microscopic, or agent-level, transition rates that results in the desired ensemble

performance. These results where then extended to account for inter-task navigation delays that are often stochastic in nature by Berman et al (2008).

In all these examples, the macroscopic continuous models are obtained by mapping the multi-agent robotic system to an equivalent chemical reaction network and then obtaining a set of rate equations to describe the birth-death process of various chemical species. In this work, we build on the results by Halasz et al (2007); Hsieh et al (2008); Berman et al (2008) and present an approach to incorporate the effects of deterministic agent-level time delays into the macroscopic analytical models. Different from Berman et al (2008), we analyze the effects of the deterministic time delays on the behavior of the macroscopic models in the frequency domain and show how these time delays can be better approximated at the macroscopic level using Padé approximants rather than an equivalent expanded linear system. We show that the Padé approximants have the ability to account for higher-order effects resulting in more accurate prediction of the mean behavior of the ensemble in the frequency domain. This is important since the top-down design paradigm rely on these macroscopic models in the optimization of the microscopic transition rates for the individual robots.

This paper is organized as follows: We formulate the time-delayed assignment problem in Section 2 and present our methodology in Section 3. We summarize the approach described by Berman et al (2008) in Section 4 and compare the two models in simulation in Section 5. We discuss our findings in Section 6 and conclude with some directions for future work in Section 7.

2 Problem Statement

2.1 Definitions

We consider the distribution of N robots among M spatially distributed tasks similar to Berman et al (2008) and model the precedence constraints between the M tasks via a directed graph, $\mathscr{G} = (\mathscr{V}, \mathscr{E})$, where the set of vertices, $\mathscr{V}$, represents tasks $\{1, \ldots, M\}$ and the set of directed edges, $\mathscr{E}$, the set of precedence constraints. We say two nodes $i, j \in \{1, \ldots, M\}$ are *adjacent*, (i, j), if there exists a precedence constraint between task i and j, and we represent this relation by the ordered pair, $(i, j) \in \mathscr{V} \times \mathscr{V}$ if i precedes j with the set $\mathscr{E} = \{(i, j) \in \mathscr{V} \times \mathscr{V}\}$. We begin with the assumption that the graph $\mathscr{G}$ is strongly connected, *i.e.* a path exists for any $i, j \in \mathscr{V}$.

We define τ_i as the amount of time it takes a robot to execute task i and assume once a robot commits to a task, it must complete the task. To model the robots switching from one task to another, we assign every edge in $\mathscr{E}$ constant *transition rates*, $k_{ij} > 0$, where k_{ij} defines the transition probability per unit time for any robot previously executing task i to arrive at task j. In general, the transition rate from i to j does not equal the transition rate form j to i, $k_{ij} \neq k_{ji}$. The individual robot controller is modeled using a finite state machine where each controller state is

associated with a specific task controller. Robots transition from one controller state (or task) to another based on the time required for each task, τ_i, and the transition rates, k_{ij}. In this work, we assume that every robot has the set of task controllers, complete knowledge of the $\mathscr{G}$, as well as all the transition rates k_{ij}. The individual robot controller and the associated ensemble model are shown in Fig. 1.

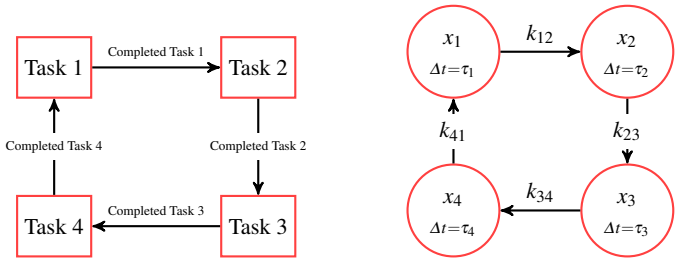

(a) Robot Controller (b) Ensemble Model

Fig. 1 (a) The robot controller. The robot changes controller states dependent on guard conditions. (b) The ensemble model. Here x_i denotes the population fraction executing task i. τ_i denotes the time to execute task i. On average, robots transition from finishing task i to starting task j with the stochastic transition rates k_{ij} .

We denote the actual number of robots executing task $i \in \{1,\ldots,M\}$ at time t by $n_i(t)$ and the desired number of robots for task i by $\bar{n}_i$ which is specified by the user. We define the population fraction executing each task at time t as $x_i(t)$ where $x_i(t) = n_i(t)/N$. Then the system state vector is given by $\mathbf{x} = [x_1,\ldots,x_M]^T$. For some initial distribution of the robots given by $\{n_i\}(t_0)$, $i = 1,...,M$, we denote the desired distribution of the ensemble as a set of population fractions for each task given by

$$\bar{x}_i = \frac{\bar{n}_i}{N} \, .$$

The specification in terms of fractions rather than absolute robot numbers is chosen such that a team size invariant formulation can be achieved. Such a framework is practical for scaling purposes as well as in situations where losses of robots to attrition and breakdown are common. Our objective is to deploy a team of robots to achieve the desired distribution among the various tasks, given by the $\bar{\mathbf{x}}$, starting from an initial distribution, $\mathbf{x}(t_0)$, with no inter-agent wireless communication.

In this work, we leverage on classical frequency-domain analysis of linear systems. Thus, given a real-valued function $f(t)$ for all real numbers $t \geq 0$, we denote the Laplace transform of $f(t)$ as $F(s) = \mathscr{L}[f(t)]$. Similarly, we denote the inverse Laplace transform of a function $F(s)$ as $f(t) = \mathscr{L}^{-1}[F(s)]$.

2.2 Linear and Time-Delayed Models

We consider the linear and delayed differential models first presented by Halasz et al (2007) and later refined by Berman et al (2008). We briefly summarize the two models in this section.

In the absence of the task execution times τ_i's, the individual robot behaviors can be described by a continuous-time Markov process in the limit of large N. As such, the time evolution of the population fraction executing task i is given by

$$\frac{dx_i(t)}{dt} = \sum_{(j,i)\in\mathscr{E}} k_{ji}x_j(t) - \sum_{(i,j)\in\mathscr{E}} k_{ij}x_i(t), \tag{1}$$

with the rate of change of the population fractions across all M tasks given by

$$\frac{d\mathbf{x}}{dt} = \mathbf{K}\mathbf{x} \tag{2}$$

where $\mathbf{K}_{ij} = k_{ji}$ for $i \neq j$ and $\mathbf{K}_{ii} = -\sum_{(i,j)\in\mathscr{E}} k_{ij}$. We note that the columns of $\mathbf{K}$ sum to 0 and since the number of agents is conserved, the system is subject to the conservation constraint

$$\sum_{i=1}^{M} x_i(t) = 1 . \tag{3}$$

The system (2) describes the average rates of change of the population fractions executing each task and are referred as the reaction rate equations (RREs). From this model, one can design the steady-state distribution of the team across the various tasks, *i.e.*, $\bar{\mathbf{x}}$, by appropriately selecting the individual transition rates (Halasz et al, 2007). The choice of $\bar{\mathbf{x}}$ can be selected to satisfy some minimal average task completion rate for the ensemble. Halasz et al (2007); Hsieh et al (2008) showed that for strongly connected graphs, the system (1) is always stable regardless of the choice of $\mathbf{K}$. As such, given $\bar{\mathbf{x}}$ and the corresponding $\mathbf{K}$, the ensemble will automatically distribute itself accordingly without explicit inter-agent wireless communication.

The linear model given by (1) assumes that robots transition from one task to another instantaneously. However, in practice, once a robot commits to task, there is a delay between finishing its current task and beginning its next one. The linear model can be extended to take into consideration the delayed transitions by converting (1) into a delayed differential equation.

$$\frac{dx_i(t)}{dt} = \sum_{(j,i)\in\mathscr{E}} k_{ji}x_j(t-\tau_j) - \sum_{(i,j)\in\mathscr{E}} k_{ij}x_i(t), \tag{4}$$

For the system with near deterministic task-times, τ_i for all $i = 1,\ldots,M$, we claim this is a more natural description of the robotic task allocation problem. When the tasks are spatially distributed within the workspace, these transition rates can repre-

sent the uncertainty in switching between tasks such as navigation delays between tasks due to concerns such as traffic congestion and collision avoidance maneuvers.

Both (1) and (4) result in robots switching between states at equilibrium. This is because these models force a trade-off between maximizing the transition rates for fast equilibration and achieving long-term efficiency during equilibrium.

Additionally, both (1) and (4), are macroscopic models of the ensemble activity. The desired distribution of the ensemble, $\bar{\mathbf{x}}$, across the various M tasks can be achieved through the selection of the individual transition rates, k_{ij}, (Halasz et al, 2007). Hsieh et al (2008) used (1) to determine the set of individual robot transition rates that result in fast convergence and minimum steady-state transitions. Berman et al (2008) employed an expanded linear model similar to (1) to approximate the model given by (4) to account for stochastic time delays that may arise when robots navigate between locations.

In this work, we use the formulation given by (4) and apply Padé approximants for the deterministic time delays in the frequency domain. We show how such an approach more accurately, in terms of the frequency response, and efficiently, in terms of the polynomial degree, models the steady-state oscillations observed in the underlying multi-agent robotic system when delays are deterministic or near deterministic.

3 Methodology

We begin by applying the Laplace Transform to the system given by (4) to obtain

$$sX_i = -\sum_{(i,j)\in\mathscr{E}} k_{ij}X_i + \sum_{(j,i)\in\mathscr{E}} k_{ji}e^{-s\tau_j}X_j \text{ for } i = 1,\ldots,M. \tag{5}$$

In general, the Laplace transform converts a differential equation in the time domain into an algebraic equation in the frequency domain where the resulting equations are purely sums of polynomials of s, thus simplifying the analysis. However, the time delay introduces an exponential term which makes the rate expression transcendental. To retain the algebraic structure, a common approach is to apply a Padé approximation of the exponential term in the frequency domain resulting from the time delay (Silva et al, 2004).

A Padé approximation is obtained by approximating the function as a rational function which is a fraction of two polynomials. In general, a Padé approximant of order (q,r) for some function $f(s)$ is the rational function of the form

$$R(s) = \frac{\alpha_0 + \alpha_1 s + \alpha_2 s^2 + \ldots + \alpha_q s^q}{\beta_0 + \beta_1 s + \beta_2 s^2 + \ldots + \beta_r s^r} \tag{6}$$

that has a Maclaurin expansion that agrees with $f(s)$ up to its $f^{(q+r)}(s)$. In the case of modeling the complex exponential time delay, the Padé approximation has the simpler form,

$$R(s) = \frac{1-\alpha_1 s + \ldots + (-1)^q \alpha_q s^q}{1+\alpha_1 s + \ldots + \alpha_q s^q} = \prod_{i=1}^{q} \frac{1-p_i s}{1+p_i s} \tag{7}$$

Note that the poles and zeros of the (q,q) Padé approximation are equal with opposite sign. This maintains the unity gain element of the systems.

Applying the above approximation to (5), we obtain

$$X_i = \frac{1}{s+\sum_{(i,j)\in\mathscr{E}} k_{ij}} \sum_{(i,j)\in\mathscr{E}} \prod_{l=1}^{q} \frac{1-p_l s}{1+p_l s} X_j.$$

To demonstrate the effects of such an approximation in the time domain, consider a 1^{st} order Padé approximation and $\tau_i = \tau$ for all $i = 1,\ldots,M$. Taking the inverse Laplace transform of the 1^{st} order Padé approximated system, we obtain

$$\ddot{x}_i + (\frac{2}{\tau} + \sum_{(j,i)\in\mathscr{E}} k_{ij})\dot{x}_i + \frac{2}{\tau} \sum_{(j,i)\in\mathscr{E}} k_{ij} x_i = \frac{2}{\tau} \sum_{(i,j)\in\mathscr{E}} k_{ji} x_j - \sum_{(j,i)\in\mathscr{E}} k_{ji} \dot{x}_j \tag{8}$$

for all $i = 1,\ldots,M$. To express the above equations in state-space form, let $z_i = x_i$ and $z_{2i} = \dot{x}_i$ for all $i = 1,\ldots,M$. This results in the following system of equations

$$\begin{aligned} \dot{z}_{2i-1} &= z_{2i}, \\ \dot{z}_{2i} &= -(k_i + \frac{2}{\tau})\, z_{2i} - \frac{2}{\tau}\, z_{2i-1} + \sum_{(j,i)\in\mathscr{E}} \frac{2}{\tau} k_{ji}\, z_{2j-1} - \sum_{(j,i)\in\mathscr{E}} k_{ji}\, z_{2j}. \end{aligned}$$

The above equations is akin to a second-order version of (1) and can be rewritten in matrix form as

$$\dot{z} = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \frac{2}{\tau}\mathbf{K} & \mathbf{F} - \frac{2}{\tau}\mathbf{I} \end{bmatrix} z$$

where $[F]_{ij} = -k_{ji}$ for $i \neq j$ and $[F]_{ii} = -\sum_j k_{ij}$. Unlike, $\mathbf{K}$, the columns of $\mathbf{F}$ do not sum to 0.

Our main concern is to extend the RREs to provide a more accurate description of the steady-state oscillations experienced by the underlying robotic system. While it is possible to formulate the ensemble assignment problem in terms of the Chemical Master Equation (CME), the RRE framework provides a more computationally efficient approach for two reasons. First, the dimension of the system state space in the RRE formulation is equal to the number of tasks as compared to the CME formulation where the dimension depends on both the number of robots and tasks. Second, approximations for deterministic time delays quickly become intractable within the CME formulation due to the high number of additional states required to adequately model the delays. We discuss this in the context the of model proposed by Berman et al (2008) in the following section. Finally, while the CME formulation

describes both the expected value and the variance of the population density function over time, the steady-state oscillation of the underlying system can be difficult to discern. The ability to predict the frequency of oscillation can inform the design of agent-level adaptation strategies especially in situations where the oscillations degrade the overall performance of the ensemble.

4 The Multi-Pole Model

Berman et al (2008) considered the dynamic distribution of N robots to M physically distinct locales. In this framework, the time delays in (4) were assumed to arise from the time required to navigate from a site (or task) to another. Rather than model the navigation time as a deterministic variable, Berman *et al.* observed that such navigation delays can be more accurately modeled as random variables rather than deterministic delays. As such, the model proposed by Berman et al (2008) only considered the effects of stochastic agent-level time delays. We briefly summarize the model in this section and refer to it as the Multi-Pole model.

We begin with the assumption that every transition time, τ_i, is drawn from an Erlang distribution with mean and variance given by τ_i and $Var(\tau_i)$. The Erlang distribution is the sum of k independent identically exponentially distributed random variables. The differential delayed system given by (4) can be transformed into an equivalent linear ordinary differential equation of the form (1) by representing each edge in $\mathscr{G}$ as a directed path composed of a finite set of "dummy states", represented by extra vertices. The transitions between states are governed by exponential wait times and therefore the transition time through a set of dummy transitions becomes an Erlang distributed waiting time. This approach is similar to the *linear chain trick* described by MacDonald (1978).

The approach is shown pictorially in Fig. 2(b) for the two state toy example shown in Fig. 2(a). In general, the mean and variance of the Erlang distribution is a function of both the number of "dummy states" associated with each edge in $\mathscr{E}$ and the associated transition rates, λ_{ij}, given by

$$E(\tau)_{Erl} = \frac{D_{ij}}{\lambda_{ij}}, \qquad Var(\tau)_{Erl} = \frac{D_{ij}}{\lambda_{ij}^2}.$$

The parameters can be adjusted for each edge to control the variance and mean of the associates transition time, $D_{ij} = \tau_i^2/Var(\tau_i)$ and $\lambda_{ij} = D_{ij}/\tau_i$ respectively. For the toy example shown in Fig. 2(a), the resulting expanded linear model is given by

$$\begin{aligned}
\dot{x}_i &= \sum_{(i,j)\in\mathscr{E}} \lambda_{ji} y_{ji}^{(D_{ji})} - \sum_{(i,j)\in\mathscr{E}} k_{ij} x_i, \\
\dot{y}_{ij}^{(1)} &= k_{ij} x_i - \lambda_{ij} y_{ij}^{(1)}, \\
\dot{y}_{ij}^{(m)} &= \lambda_{ij} \left(y_{ij}^{(m-1)} - y_{ij}^{(m)} \right),
\end{aligned} \tag{9}$$

for $m = 2, \ldots, D_{ij}$ and $y_{ij}^{(m)}$ denotes the fraction of the population in dummy "state" m associated with the edge $(i, j) \in \mathscr{E}$. This results in the linear system given by $[\dot{\mathbf{x}}^T \dot{\mathbf{y}}^T]^T = \mathbf{K}_{equiv}[\mathbf{x}^T \dot{\mathbf{y}}^T]^T$ similar to (1).

Applying the Laplace transform to the typical rate equations shown in Fig. 2(a) result in

$$\begin{array}{ccc}
\dot{y}_{12}^{(1)} = \lambda_{12} x_1 - \lambda_{12} y_{12}^{(1)} & \xrightarrow{\mathscr{L}} & sY_{12}^{(1)} = \lambda_{12} x_1 - \lambda_{12} Y_{12}^{(1)} \\
\vdots & & \vdots \\
\dot{y}_{12}^{(i)} = \lambda_{12} y_{12}^{(i-1)} - \lambda_{12} y_{12}^{(i)} & \xrightarrow{\mathscr{L}} & sY_{12}^{(1)} = \lambda_{12} Y_{12}^{(i-1)} - \lambda_{12} Y_{12}^{(i)} \\
\vdots & & \vdots \\
\dot{x}_2 = \lambda_{ij} (y_{12}^{(m)}) & \xrightarrow{\mathscr{L}} & sX_2 = \lambda_{ij} (y_{12}^{(m)}) \\
& \vdots & \\
& sX_2 = \left(\frac{\lambda_{12}}{s+\lambda_{12}} \right)^m X_1 &
\end{array}$$

We note that the addition of extra "dummy states" is equivalent to adding more real poles to the system in the frequency domain. By adding more poles, the approximation begins to look like another typical approximation of the exponential,

$$e^{-\tau s} = \lim_{n\to\infty} \left(1 + \tfrac{s\tau}{n}\right)^{-n} \sim \left(\frac{1}{1 + \frac{s}{m\lambda_{12}}} \right)^m . \tag{10}$$

The location of these poles are dependent on the desired mean and variance of the transition time. While the additional poles result in a sparse matrix $\mathbf{K}_{equiv}$, the amount of poles needed to approximate τ_i as $Var(\tau_i) \to 0$ is very large because of the slow convergence rate of $Var(\tau_{ij}) \sim 1/D_{ij}$. This is analogous to the slow convergence rate of equation (10). To give an example, a transition that with half the variance and the same mean requires a doubling of the amount of poles, or equivalently, dummy sites. The approximation of the delays as a product of poles can attenuate the high frequencies in the system resulting in poor modeling of the transient response and possibly leading to instability.

In practice, a time delay in the input signal does not affect the gain of the frequency response. Rather, the time delay should simply retard the phase of the response signal. In the frequency domain, the time delay is modeled as an exponential variable as shown in Section 3. This is equivalent to an exponential delay in the phase response of the output signal. As the frequency increases, the system delays the output signal by more and more periods. The effects of the deterministic micro-

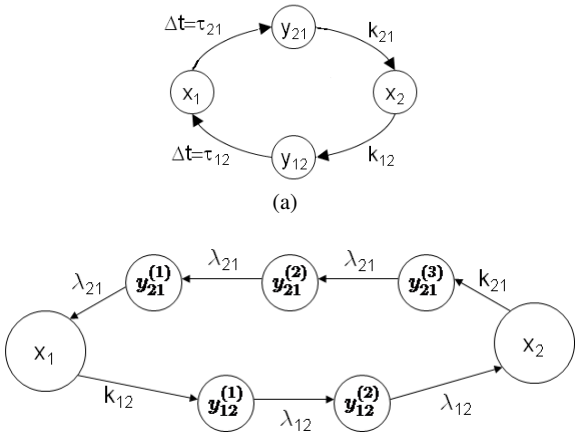

Fig. 2 A toy example with two states or tasks. (a) Differential delayed model. (b) Expanded linear model or Multi-Pole model. Here the $y_{ij}^{(k)}$ denotes the "dummy states".

scopic time delays can potentially be better approximated at the macroscopic level by employing a Padé approximant for the exponential term in the frequency domain since it limits the effects on the gain of the response. This is because as the Padé approximant, *i.e.*, equation (7), grows in order, it adds matching pole-zero pairs such that each extra pole-zero increases the phase shift by π while effectively canceling the loss in gain.

To illustrate this, consider the simple input-output system given by $x(t) = u(t - \tau)$. Fig. 3 shows the step responses for the pure delayed system, a Multi-Pole approximated system, and a Padé approximated system. We note that the frequency content of the delayed response depends on the order of the approximation, *i.e.*, number of poles added to represent each delay similar to (9) or the order of the Padé approximant. For this example, we employed a 64^{th} order Multi-Pole delay and a 4^{th} order Padé approximant. The orders of approximation were selected to display qualitatively similar responses in the rise time step dynamics. While the Multi-Pole response may look like a better approximation of delayed step response, the number of poles required to match the Padé approximated step response grows exponentially compared to the Padé order.

We note that the output value of the Padé approximation can cause the modeled population fraction to reach negative values, whereas the actual robotic population fraction cannot. This is a consequence of modeling the time delay with a finite order rational polynomial and can be addressed via saturation functions to limit the output variables to the appropriate range, $x(t) \in [0, 1]$.

To further illustrate the differences between the Padé and the Multi-Pole systems, consider their frequency responses shown in Fig. 4. While both systems have the same break frequency, the 4^{th} order Padé response closely resembles the frequency response of the actual time-delayed system. We note that location of the

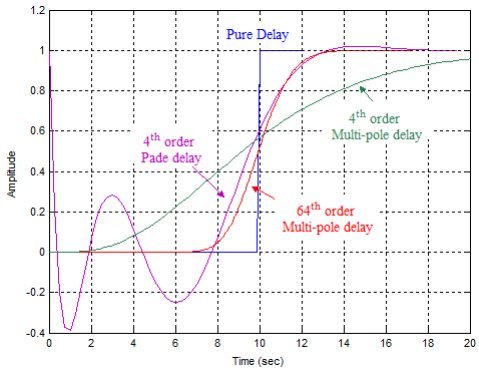

Fig. 3 Step responses for the system $x(t) = u(t-\tau)$ where $\tau = 10$ seconds. The response of the Multi-Pole model, the Padé model, and of the actual system is shown.

break frequency is a function of the transition rates. Depending on the application, the periodicity of various cycles in $\mathscr{G}$, or *loop frequencies*, can be designed by appropriately selecting the transition rates.

In this work, we employ a 4^{th} order Padé approximations for $e^{-s\tau_j}$. We evaluate our Padé approximated model and compare it to the Multi-Pole model in the following section using different agent-based simulations.

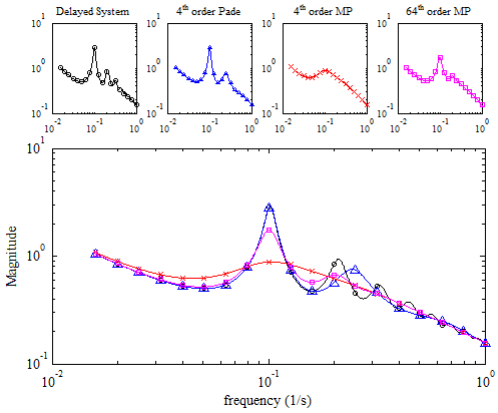

Fig. 4 Magnitude portion of the frequency response for the 4^{th} and 64^{th} order Padé response and the 4^{th} order Multi-Pole response compared to the system response given by (4). The transition rates and delays are $k_{12} = 1\ sec^{-1}$, $k_{21} = 1\ sec^{-1}$, $\tau_1 = 1\ sec$, $\tau_2 = 7\ sec$. These rates result in a loop time of 10 seconds.

5 Simulation Results

For ease of comparison, we consider the allocation problem presented by Hsieh et al (2008); Berman et al (2008). The objective of this work is to determine whether the Padé approximated *macroscopic continuous* (macro-continuous) model can provide an accurate description of the average ensemble behavior for the underlying microscopic system subject to deterministic time delays. These high level representations provide a lower dimensional description of the underlying multi-agent robotic system and can be more efficient when used to optimize the various system parameters to improve the performance of the robot ensemble (Hsieh et al, 2008; Berman et al, 2009). As such, the accuracy of the macro-continuous models will not only determine our ability to predict ensemble behavior but also affect our ability to optimize ensemble performance.

Similar to existing work, we employ a multi-level simulation methodology that represents three separate levels of abstraction of our deployment strategies. At the top-most level are the macro-continuous reaction rate equations (RREs) that model the evolution of the mean robot populations executing different tasks. At the lowest level, is the full model of the multi-agent robotic system, what we refer to as the *microscopic discrete* (micro-discrete) model. At this level, we conduct a full physics simulation of N robots where each robot is programmed to execute the assigned tasks. We also perform simulations at the middle level, or *macroscopic discrete* level (macro-discrete), using a simulation procedure that is mathematically equivalent to an agent-based simulation (Gillespie, 1976, 2007). One can view the macro-discrete representation as a stochastic formulation of the macro-continuous models. We summarize our simulation methodologies and present our results in the following sections.

5.1 Simulation Methodology

We consider the deployment of an ensemble of N robots moving in the plane to M distinct locations/sites. In this work, we consider the case where time delays are deterministic and arise because each robot must spend τ_i seconds executing the task at site i as compared to stochastic delays due to navigation between sites as described by Berman et al (2008). At the macro-continuous level, simulations are performed by numerically integrating the continuous-time differential equations.

5.1.1 Micro-Discrete Simulation

For our micro-discrete simulations, we developed an agent-based simulation for an ensemble of 10 SRV-1 robots in USARSim (USARSim, 2007). USARSim is a high fidelity simulator based on the Karma physics engine in the Unreal Tournament game. The SRV-1s are differential-drive robots equipped with an embedded pro-

cessor, color camera, and 802.11 wireless capability. The team is tasked to surveil two buildings each located in a different location. Each building is represented by a rectangular block in the workspace. We assume the robots have a map of the environment and navigate from building to building via simple potential functions. Initially, the robots are assigned to either building. Once the robot reaches its targeted building, it circles the building in a clockwise direction. The task execution time, τ_i, was chosen to reflect the amount of time it takes the robot to circle the building twice. After accomplishing its task, the robot moves onto the next building to surveil. Robots navigate from one building to another using a potential field controller. Once they have completed the surveillance task at each site. There is variability in the amount of time it takes a robot to travel between the two sites depending on the amount of traffic which may affect the number of collision avoidance maneuvers each robot must execute. Collision avoidance among the robots is achieved through a combination of gyroscopic forces and potential functions (Chang et al, 2003; Hsieh et al, 2007). For these experiments, the transition rates were chosen so the team would evenly distribute themselves among the two buildings. The transition rates and initial conditions for the simulation are summarized in Table 1.

State	Initial Condition	Transition Rate	τ Delay	State Description
x_1	0	-	$\tau_1 = 50$	Robots surveilling building 1
x_2	0	-	$\tau_2 = 50$	Robots surveilling building 2
y_1	5	$k_{1,2} = \frac{1}{17.3}$	-	Robots transitioning from 1 to 2
y_2	5	$k_{2,1} = \frac{1}{18.2}$	-	Robots transitioning from 1 to 2

Table 1 System parameter values for the micro-discrete simulations.

5.1.2 Macro-discrete Simulation

In the absence of time delays, an exact stochastic simulation algorithm can be used to simulate the trajectory of the state variables following Gillespies Direct Method (Gillespie, 1976). This is achieved by first converting the continuous population model into a set of "reactions" that describe individual robot behavior transitions. In a manner analogous to a system of chemical reactions, each transition occurs at a rate derived from the transition rates. In our simulations, the behaviors are represented by tasks/sites and the transition rates are chosen to achieve a uniform distribution across the tasks. The next behavior transition is selected from a uniform distribution over the relevant propensities, and the time interval until the next occurrence, δt_{k+1}, is computed from an exponential distribution with the total propensity as its parameter. The time is advanced by δt_{k+1} and the transition is simulated by decreasing the number of robots in one site and incrementing the number in the second. For a detailed description on how to compute the transition propensities using these rates, we refer the reader to the work by Berman et al (2006).

To simulate the Padé approximated macro-discrete system given by (8), we used the method presented by Bratsun et al (2005) which extends the Direct Method to work with an inherent time delay. Since, the mean and variance of an exponential wait time are both given by the inverse of the arrival rate, λ, long wait times will invariably result in large variances. As a consequence, Bratsun et al (2005) propose the delay times to be drawn from a shifted exponential distribution with expected value given by $\lambda^{-1} + \tau$. This ensures that the wait times are still exponentially distributed while providing the ability to affect the variance of the distribution without significantly altering the structure Gillespie's Direct Method.

5.2 Results

We begin by comparing the frequency responses of the Padé and Multi-Pole macro-continuous models with the micro-discrete simulations. We employed 4^{th} order Padé and Multi-Pole approximations.

We ran 54 micro-discrete simulations in USARSim where each robot was tasked to survey each building for 50 seconds before switching to the other building. To obtain the frequency response of the micro-discrete simulations, we logged the population fractions at each site over time and applied the Fast Fourier Transform (FFT) to these variables for each run. The FFT results were then averaged for 54 runs. The result is shown in Fig. 5. The micro-discrete system exhibits a maximum gain at approximately 7.5 mHz. While both the Padé and Multi-Pole macro-continuous models exhibit peaks at approximately the same frequency, the Padé model shows larger gain.

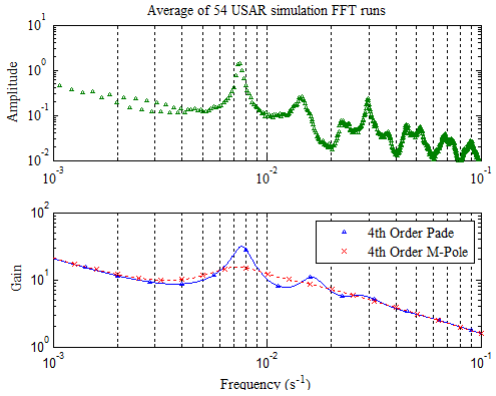

Fig. 5 Top: Average of the FFT of the population fraction at building 2 obtained from 54 micro-discrete simulations. Bottom: Magnitude portion of the Bode plots relating to the number of Robots at building 2. for the 4^{th} order Padé, 4^{th} order Multi-Pole, and 16^{th} order Multi-Pole macro-continuous systems.

In general, it is difficult to directly compare the macroscopic results with the microscopic results. This is because the FFT of the system states only considers the outputs of the system. The magnitude portion of the Bode plots, on the other hand, gives the response of the ratio of the output to input of the system for all frequencies. In other words, the macroscopic frequency response is based on a unity gain input at all frequencies. The difference between the two plots is dependent on the form of the noise input to the system and are related by the shape of the frequency spectrum of the noise input to the system. Given the same transition rates, *i.e.*, **K** in (1), our microscopic simulations result in the same peak frequency as predicted by the macro-continuous models.

The difference in frequency gain between the Padé and Multi-Pole approximations are significant, with the 4^{th} order Pade system closely matching the delayed system dynamics within 0.1% while the 4^{th} order Multi-Pole systems gain was off by 51%. The effects of this discrepancy are significant. The settling times for the delayed system and the Pade system are 547.3 and 545.8 seconds respectively, where as the settling time for the Multi-Pole system is 164 seconds. A long settling time implies that oscillations generated by the noise in the system will remain in the system. For example, in situations when robots break down, the system will experience what looks like a noise impulse and as a result the populations at the affected sites will begin to oscillate. In the agent based simulation, where collision avoidance and traffic can lead to the clustering of robots, this spurious oscillation can be sustained.

Fig. 6 shows the number of robots at Building 2 for a sample micro-discrete simulation. The number of robots for each site was initialized at the desired steady-state levels. As expected, the initial variations in population sizes were random, however, the population size eventually began to oscillate. The oscillations were partly due to the changes in traffic patterns en-route to the different sites which affects the amount of collision maneuvers individual robots must execute. When traffic is high, it becomes more difficult for the team to quickly re-establish equilibrium.

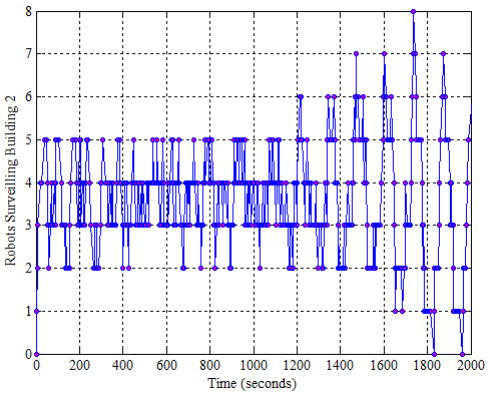

Fig. 6 Number of robots surveilling Building 2 over time. The number of robots at Building 2 was initialized at the desired steady-state population.

Fig. 7 shows a series of snapshots for a typical microscopic simulation. Robots are either conducting surveillance of a building, represented by the blocks on either side, or switching from one building to another. Fig. (a) shows the initial positions of the robots in the workspace, Fig. (f) shows 3 robots transitioning from building 2 to building 1 and Fig. (h) shows a typical equilibrium state.

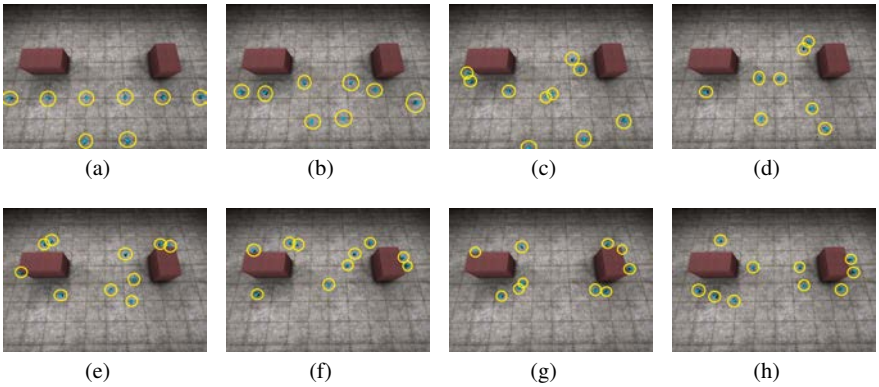

Fig. 7 Snapshots from a typical microscopic simulation. Building 1 and 2 are the right and left structures respectively.

To show that the steady-state oscillations predicted by the time-delayed models are independent from team size, we ran macro-discrete simulations using the modified Direct Method (Bratsun et al, 2005) for ensemble sizes of 10, 200, and 2000 robots using the same transition rates and task execution times. As mentioned previously, the macro-discrete simulation is mathematically equivalent to an agent-based simulation but is more computationally efficient when the ensemble size is large. Fig. 8 shows the FFT of the output of these simulations with the frequency response of the 4^{th} order Padé macro-continuous model. Regardless of the team size, the system exhibits the same peak frequency as predicted by the Padé macro-continuous model. More interestingly is that the shape of the FFT is consistent across the various team sizes.

6 Discussion

Within this framework, the intrinsic noise in the population variables is amplified by the loop gains associated with the transition rates and task delay times. When the system has deterministic or near deterministic time delays, the population values can have high magnitude frequency components. Given the time delays, it is possible to determine these loop frequencies. While such spurious effects do not produce instability, the associated oscillations can adversely effect the performance of the system

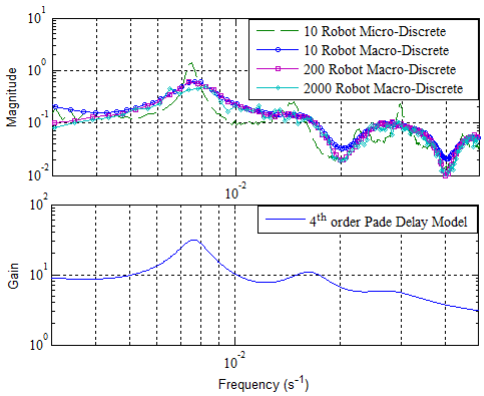

Fig. 8 FFT of different macro-discrete simulations for different team sizes and the frequency response of the 4^{th} order Padé macro-continuous model. The top plot is the FFT results of a macro discrete simulation with different number of robots. The output is the frequency content of the number of robots at building 2 normalized by the total number of robots.

by increasing the amount of time a robot spends switching between tasks. This is because the spurious frequency components are constantly being renewed based on the noise content of the underlying system. The implication here is that these oscillations can be sustained for much longer time scales than what the non-delayed macro-continuous simulations can predict. In the presence of other noise sources, like the variability introduced by collision avoidance or other forms of robot-robot or robot-environment interactions, these oscillation can be sustained.

The presence of these spurious frequency components in the surveillance application creates undesired long term behaviors. The physical meaning of the observed oscillations in the time domain is that the robots are clustering together, *i.e.*, traveling in packs. Consider the extreme case of a single robot, where the robot always travels in a "pack". In this case, the frequency content will be large because there is no traffic and all states will be unit amplitude square waves. For the single robot there is no adverse effect of this frequency peaking. However, for the team of robots traveling together, collision avoidance becomes a significant concern since the local traffic is always high. If the congestion does not dissipate, then such local traffic concerns will tend the follow the robots as the move from one task to another. This leads to degraded performance as the average transit time between sites will increase due to these traffic concerns due to an amplification of the low number intrinsic noise of the system by the frequency dynamics. If we only considered the steady-state behavior of the system, one would expect the majority of the N robots to be executing their surveillance tasks with only a small fraction of them traveling between sites. However, the presence of the unwanted frequency components can result in a significant imbalance between robots at sites and those traveling between them. The Padé approximated macro-continuous model has the ability to predict the spurious frequency component that is present in the agent-based simulation.

The removal of these spurious frequency components at the macroscopic level can be achieved by incorporating a notch filter of the form

$$H(s) = \frac{s^2 + 2\zeta_1\omega_N - \omega_N^2}{s^2 + 2\zeta_2\omega_N - \omega_N^2}$$

with ω_N and the ratio of ζ_1/ζ_2 chosen such that the location and magnitude of the notch are properly located at the spurious frequency.

While applying control on the macro-continuous model is straight-forward, it is not clear how such a controller can be mapped to specific modifications of the individual robot controllers. In fact, inspection of the closed-loop time domain equations suggest in order to realize the macroscopic notch controller, the individual robots must have the ability to estimate the higher order derivatives of the population fractions at both their current sites as well as their adjacent sites. While such a strategy makes intuitive sense, the straight forward implementation of such a controller can be prohibitively costly in terms of the amount of information that needs to be communicated. Furthermore, it is not clear whether such a controller can be achieved by the ensemble in a distributed fashion or whether there are other distributed strategies that will result in the same macroscopic effect. This is a topic of on-going research for our group.

From our results, we have shown that the Padé approximated macro-continuous model is better suited for modeling the effects of deterministic time delays on ensemble performance. As described in Section 4, the Multi-Pole formulation can also provide an adequate approximation as long as the number of dummy states required is reasonable. For transitions times where the magnitude of the variance is on the same order as the mean, the Multi-Pole formulation will model the frequency dynamics accurately. However, if we consider an example where the variance is 1 sec^2 and the mean is 10 sec, from Section 4, we note that this requires 100 dummy transitions. From our results, it seems these near deterministic transitions can be more efficiently modeled by concatenating a single pole dummy transition with the desired variance with a Padé approximation to model the discrete time. This hybrid approach should give similar results to the Multi-Pole model with a significantly lower order and is a direction for further investigation.

7 Future Work

We have presented a methodology to more accurately predict the effects of deterministic agent-level time delays on mobile robot ensemble dynamics. Different from existing approaches, we employ a Padé approximation for the time delays in our system to retain frequency components that might have been smoothed out in a purely kinematic macroscopic continuous formulation of the ensemble dynamics. We showed through our multi-level simulations that our approach provides a macro-

scopic description of the ensemble that retains the relevant frequency characteristics of the microscopic system.

An immediate direction for future work is to investigate other approximation techniques that can further enhance these chemical reaction network derived models to provide a low parameter description of the physical mobile robot systems. Of particular interest are techniques that enable us to incorporate more dynamics of the underlying robotic systems. For near deterministic or small variance time delays, one possibility is to employ a combination of Padé and Multi-Pole methods to better capture the microscopic effects. We would also like to extend our results to nonlinear chemical reaction network based systems. Though the Padé results in the Laplace frequency domain are linear, the method for constructing a delay out of derivatives of the input and output will still be applicable. This is of particular interest since nonlinear chemical reaction networks provide a framework to develop macroscopic models that describe the dynamics of heterogeneous and interacting mobile robot ensembles. Lastly, as alluded to in our discussion, we are interested in investigating various distributed realizations of feedback control strategies obtained through the analysis of macroscopic models.

References

Berman S, Halasz A, Kumar V, Pratt S (2006) Algorithms for the analysis and synthesis of a bio-inspired swarm robotic system. In: 9^{th} Int. Conf. on the Simulation of Adaptive Behavior (SAB'06), Swarm Robotics Workshop, Rome, Italy

Berman S, Halasz A, Hsieh MA, Kumar V (2008) Navigation-based optimization of stochastic deployment strategies for a robot swarm to multiple sites. In: Proc. of the 47th IEEE Conference on Decision and Control, Cancun, Mexico

Berman S, Halasz A, Hsieh MA, Kumar V (2009) Optimized stochastic policies for task allocation in swarms of robots. IEEE Transactions on Robotics 25(4)

Bratsun D, Volfson D, Tsimring L, Hasty J (2005) Delay-induced stochastic oscillations in gene regulation. PNAS 62:14,593–8

Chang DE, Shadden S, Marsden JE, Olfati-Saber R (2003) Collision avoidance for multiple agent systems. In: Proc. IEEE Conference on Decision and Control, Maui, Hawaii

Dahl TS, Matarič MJ, Sukhatme GS (2006) A machine learning method for improving task allocation in distributed multi-robot transportation. In: Braha D, Minai A, Bar-Yam Y (eds) Understanding Complex Systems: Science Meets Technology, Springer, Berlin, Germany, pp 307–337

Dias MB (2004) Traderbots: A new paradigm for robust and efficient multirobot coordination in dynamic environments. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA

Dias MB, Zlot RM, Kalra N, Stentz AT (2006) Market-based multirobot coordination: a survey and analysis. Proceedings of the IEEE 94(7):1257–1270

Gerkey BP, Mataric MJ (2002) Sold!: Auction methods for multi-robot control. IEEE Transactions on Robotics & Automation 18(5):758–768

Gillespie D (1976) A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. Journal of Computational Physics 22(4):403–434

Gillespie D (2007) Stochastic simulation of chemical kinetics. Annual Review of Physical Chemistry 58:35–55

Golfarelli M, Maio D, Rizzi S (1997) Multi-agent path planning based on task-swap negotiation. In: Proceedings of the 16th UK Planning and Scheduling SIG Workshop, PlanSIG, Durham, England

Guerrero J, Oliver G (2003) Multi-robot task allocation strategies using auction-like mechanisms. In: Artificial Research and Development, Frontiers in Artificial Intelligence and Applications, IOS Press, vol 100, pp 111–122

Halasz A, Hsieh MA, Berman S, Kumar V (2007) Dynamic redistribution of a swarm of robots among multiple sites. In: Proceedings of the Conference on Intelligent Robot Systems (IROS'07), San Diego, CA, pp 2320–2325

Hsieh MA, Loizou S, Kumar V (2007) Stabilization of multiple robots on stable orbits via local sensing. In: Proceedings of the International Conference on Robotics and Automation (ICRA) 2007, Rome, Italy

Hsieh MA, Halasz A, Berman S, Kumar V (2008) Biologically inspired redistribution of a swarm of robots among multiple sites. Swarm Intelligence

Jones EG, Browning B, Dias MB, Argall B, Veloso M, Stentz AT (2006) Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks. In: Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA'06), IEEE, Los Alamitos, CA, USA, pp 570–575

Jones EG, Dias MB, Stentz A (2007) Learning-enhanced market-based task allocation for oversubscribed domains. In: Proceedings of the Conference on Intelligent Robot Systems (IROS'07), IEEE, Los Alamitos, CA, pp 2308–2313

Lerman K, Jones C, Galstyan A, Mataric MJ (2006) Analysis of dynamic task allocation in multi-robot systems. International Journal of Robotics Research

Lin L, Zheng Z (2005) Combinatorial bids based multi-robot task allocation method. In: Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA'05), IEEE, Los Alamitos, CA, pp 1145–1150

MacDonald N (1978) Time-lags in biological models, Lecture Notes in Biomathematics, vol 27. Springer, Berlin

Martinoli A, Easton K, Agassounon W (2004) Modeling of swarm robotic systems: a case study in collaborative distributed manipulation. International Journal of Robotics Research: Special Issue on Experimental Robotics 23(4-5):415–436

Milutinovic D, Lima P (2006) Modeling and optimal centralized control of a large-size robotic population. IEEE Transactions on Robotics 22(6):1280–1285

Shen WM, Salemi B (2002) Towards distributed and dynamic task reallocation. In: Gini M, Shen WM, Torras C, Yuasa H (eds) Intelligent Autonomous Systems 7, IOS Press International, Marina del Rey, CA, pp 570–575

Silva GJ, Datta A, Bhattacharyya SP (2004) PID Controllers for Time Delay Systems. Birkhäuser Boston

USARSim (2007) Unified system for automation and robot simulation. http://usarsim.sourceforge.net

Vail D, Veloso M (2003) Multi-robot dynamic role assignment and coordination through shared potential fields. In: Schultz A, Parker L, Schneider F (eds) Multi-Robot Systems, Kluwer, pp 87–98