Robustness in the Presence of Task Differentiation in Robot Ensembles

M. Ani Hsieh and T. William Mather

Abstract. In the last fifteen years, much interest has been focused on the deployment of large teams of autonomous robots for applications such as environmental monitoring, surveillance and reconnaissance, and automated parts inspection for manufacturing. The objective is to leverage the team's inherent redundancy to simultaneously cover wide regions and achieve massive parallelization in task execution while remaining robust to individual failures. Despite recent successes, significant challenges remain, in part, due to the difficulties associated with managing and coordinating the various redundancies that exist in a large team of homogeneous agents. In this chapter, we present an ensemble approach towards the design of distributed control and communication strategies for the dynamic allocation of a team of robots to a set of tasks. This approach uses a class of stochastic hybrid systems to model the robot team dynamics as a continuous-time Markov jump process. The main advantage is a lower-dimensional representation of the team dynamics that is amenable to system-level analysis of the team's performance in the presence of task differentiation. We show how such analysis can be further used to design and optimize individual robot control policies through simulations and experimental validation.

1 Introduction

In the last fifteen years, much interest has been focused on the deployment of large teams of autonomous robots for applications such as environmental monitoring, surveillance and reconnaissance, and automated parts inspection for manufacturing. The objective is to leverage the team's inherent redundancy to simultaneously cover wide regions and achieve massive parallelization in task execution while remaining robust to individual failures. Despite recent successes, significant challenges remain, in part, due to the difficulties associated with managing and coordinating

M. Ani Hsieh · T. William Mather
Drexel University, 3141 Chestnut St, Philadelphia, PA 19104, USA
e-mail: {mhsieh1,twm32}@drexel.edu

D. Milutinović & J. Rosen (Eds.): Redundancy in Robot Manipulators, LNEE 57, pp. 93–108.
DOI: 10.1007/978-3-642-33971-4_6

the various redundancies that exist in a large team of homogeneous agents. In this chapter, we present an ensemble approach towards the design of distributed control and communication strategies for the dynamic allocation of a team of robots to a set of tasks. This approach uses a class of stochastic hybrid systems to model the robot team dynamics as a continuous-time Markov jump process. The result is a lower-dimensional representation of the team dynamics that is amenable to system-level analysis of the team's performance in the presence of task differentiation.

In our work, we consider the allocation of a team of robots to a collection of tasks distributed within a workspace. Applications include automated warehouses where inventory from the loading docks must be transported and stowed in their designated locations or deploying robots to cover different regions with their sensors for monitoring or surveillance purposes. In these examples, the team must have the ability to autonomously move from one task location to another, distribute themselves accordingly among the various locales to ensure task completion or to provide appropriate sensor coverage, all the while remaining robust to changes in the environment or individual failures.

The assignment of robots to a collection of spatially distributed tasks is a variant of the multi-task (MT) robots, single-robots (SR), time-extended assignment (TA) problem [8]. In the multi-robot domain, market-based approaches have gain much success [7, 25, 10, 5, 19, 15, 16] and can be further improved when learning is incorporated [4]. However, these methods often scale poorly in terms of team size and number of tasks. Furthermore, these approaches often depend on timely communication of the various local cost and utility functions and thus may not be suitable for situations when inter-agent wireless communication is unreliable or extremely limited [6, 9].

In this chapter, we present an ensemble approach towards the modeling, analysis, and design of distributed coordination strategies for the dynamic allocation of a team of homogeneous robots to a collection of spatially distributed tasks. Our approach is similar to existing work where macroscopic continuous models are used to describe the dynamics of a robot swarm [20, 18, 14]. The technique builds upon the representation of the individual robot controllers as probabilistic finite state machines which enables the approximation of team dynamics as a continuous-time Markov process [20, 18, 13]. In [11, 13], macroscopic models were used to synthesize stochastic agent-level control policies to enable the dynamic allocation of a team of robots to multiple locales in predefined proportions without explicit inter-agent wireless communication. Different from [18], the desired allocation was achieved through the selection of the individual robot transition rates and were extended to account for navigation delays in [1].

In our work, we assume a team of identical robots where the individual agent-level control policy is given by a sequential composition of individual task controllers. Different from existing work, we use the lower dimensional macroscopic representation of the ensemble dynamics for analysis and controller synthesis. We present two different approaches towards the design of stochastic transition rules that enable the team to autonomously achieve a desired distribution across a collection of spatially distributed task. In the first approach, mean-field models are used to

analyze and mitigate the effects of unexpected inter-robot interactions on the team's ability to maintain a desired allocation [21]. In the second approach, agent-level control policies that can affect both the mean and the variance of the distribution is obtained by modeling the team dynamics as a polynomial Stochastic Hybrid System (pSHS) [17, 23].

The novelty of the contribution is a team-size invariant approach towards the design of distributed agent-level control policies that can respond to robot failures in a natural way and ensure a graceful degradation of the system. By providing a systems-level view of the team dynamics, the inter-agent communication needs of a desired coordination strategy can be more explicitly accounted for at the controller synthesis stage. The rest of the chapter is organized as follows: Section 2 presents the development of the macroscopic models for an ensemble of robots executing a collection of tasks with deterministic task execution times. Section 3 describes the analysis and design of our ensemble model derived distributed control strategies. We conclude with a brief discussion of future work in Section 4.

2 Problem Formulation

Consider the deployment of N robots to M tasks each located within a different region in the workspace. In this section, we show how continuous macroscopic models can be obtained from a collection of individual robot controllers. The goal is to use these models to design a decentralized control policy to enable the team to autonomously distribute across the M tasks and maintain the desired allocation at the various regions.

2.1 Single Robot Controller

Given a collection of $\{1,\ldots,M\}$ tasks/sites, we use a directed graph, $\mathscr{G} = (\mathscr{V},\mathscr{E})$, to model the pairwise precedence constraints between the tasks. Each task is represented by a vertex in $\mathscr{V} = \{1,\ldots,M\}$. A directed edge exists between two vertices $(i,j) \in \mathscr{V} \times \mathscr{V}$ if task i must precede task j and we denote this relation as $i \prec j$. Then, the set of edges, $\mathscr{E}$, is given by $\mathscr{E} = \{\forall (i,j) \in \mathscr{V} \times \mathscr{V} | i \prec j\}$. We assume $\mathscr{G}$ is a strongly connected graph, *i.e.*, a directed path exists for any $i, j \in \mathscr{V}$.

Given the M tasks, we denote the set of task controllers for each task as $\{U_1,\ldots,U_M\}$ and assume that the single robot controller is obtained through the sequential composition of $\{U_1,\ldots,U_M\}$ such that the precedence constraints specified by $\mathscr{G}$ are satisfied. We represent the robot controller as a finite state automaton where each automaton state i is associated with a task controller U_i. Fig. 1(a) shows an example robot controller where the arrows denote state transitions that satisfy the constraints specified in $\mathscr{G}$.

In this work, we consider the assignment/allocation of the team to M tasks/sites. The team's objective is to maintain the desired allocation of the robots across the various regions. At each site, robots execute U_i for a pre-specified amount of time τ_i. This represents the time required by a robot to complete the task at the given site.

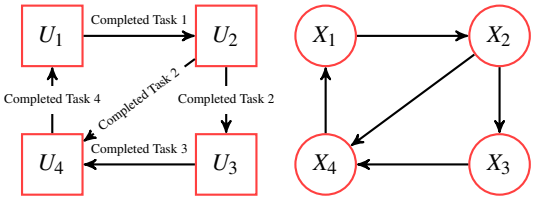

(a) Robot Controller (b) Ensemble Model

Fig. 1 (a) The robot controller. The robot changes controller states based on the guard conditions. (b) Graphical representation of the equivalent chemical reaction network for a robot ensemble executing the tasks.

Once the task has been completed, the robot must navigate to the next adjacent site based on the constraints encoded in $\mathscr{G}$. As such, we assume each robot has complete knowledge of $\mathscr{G}$, the ability to localize within the workspace, and is capable of navigating from one task/site to another while avoiding collisions with other robots in the workspace.

2.2 The Ensemble Model

For a team of N robots, each executing the same sequentially composed controller, *e.g.*, the one in Fig. 1(a), the ensemble dynamics can be modeled as a *polynomial stochastic hybrid system* (pSHS). This enables us to derive lower dimensional macroscopic models that describe the time evolution of the distribution of the team across the various tasks/sites.

Let $X_i(t)$ and $\bar{X}_i$ denote the number of robots executing task i (or being site at i) and the desired number of robots for task i respectively. Then the fraction of the robots at site i is given by $x_i(t) = X_i(t)/N$ with $\bar{x}_i$ denoting the desired fraction of robots. The specification in terms of fractions rather than absolute numbers provides a team size invariant formulation and is practical for scaling purposes. Since the tasks are spatially distributed, the robots will move from one task to another and must avoid collisions with other robots. The variability in robot arrival times at each site is modeled using *transition rates*. For every edge $e_{ij} \in \mathscr{E}$, we assign constant $k_{ij} > 0$ such that k_{ij} defines the transition probability per unit time for one agent from site i to go to site j. Furthermore, we assume the ensemble dynamics is Markov which will allow us to model the dynamics of the robot distribution as a set of linear differential equations. It is important to note that in general $k_{ij} \neq k_{ji}$.

2.2.1 Mean-Field Dynamics

It was shown in [13, 1] that the time evolution of the population fraction executing task i can be modeled as a continuous-time Markov process in the absence of task execution times, *i.e.*,

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

The task execution times can be incorporated by reformulating the above equation as a delayed differential equation of the form

$$\frac{d}{dt}x_i(t) = \sum_{(j,i)\in\mathscr{E}} k_{ji}x_j(t-\tau_j) - \sum_{(i,j)\in\mathscr{E}} k_{ij}x_i(t). \tag{2}$$

We note that these models are mean-field descriptions of the team dynamics where the system state is given by $\mathbf{x}(t) = [x_1(t),\ldots,x_M(t)]^T$. In [13, 2], the desired distribution across the M sites was achieved by using (1) to optimize the k_{ij} terms to meet specific ensemble performance metrics.

2.2.2 Moment Dynamics

While the mean-field formulation provides a model of the time evolution of the fractions of robots at each task location, it is possible to provide a different macroscopic description of the ensemble dynamics by considering the rates of change of the various moments of the robot population distribution. Similar to the use of fractions in the previous section, the specification in terms of the moments of the robot population will also provide a team size invariant formulation. This is achieved by describing the ensemble dynamics using a set of transition rules of the form:

$$X_i \xrightarrow{k_{ij}} X_j \quad \forall\, e_{ij} \in \mathscr{E}. \tag{3}$$

The above expression represents a stochastic transition rule with k_{ij} as the per unit reaction rate and $X_i(t)$ and $X_j(t)$ as discrete random variables. In the robotics setting, equation (3) implies that robots at site i will transition to site j with a rate of $k_{ij}X_i$.

In this formulation, the system states are the random variables $X_i(t)$ with the state vector given by $\mathbf{X}(t) = [X_1(t),\ldots,X_M(t)]^T$. Given the set of stochastic transition rules in (3), the moment equations for the discrete random variable X_i is given by the extended generator of the system [12]. For a real-valued function $\psi(X_i)$, the extended generator is an expression for the time derivative of the expected value of ψ, *i.e.*, $\frac{d}{dt}\mathrm{E}[\psi(X_i)] = \mathrm{E}[L\psi(X_i)]$, and takes the form

$$L\psi(X_i) = \sum_j \left[(\psi(X_i-1)-\psi(X_i))w_{ji} + (\psi(X_i+1)-\psi(X_i))w_{ij}\right]. \tag{4}$$

The right hand side of (4) gives the continuous time derivatives of the system for a discrete change in the state X_i. The expression $[\psi(X_i-1)-\psi(X_i)]$ represents the change in ψ given a unit change in the discrete variable X_i, while w_{ij} represents the

frequency at which the change occurs. For the system given by (3), $w_{ij} = k_{ij}X_i$. To obtain the rate of the change of the expected value of X_i, $\frac{d}{dt}\mathrm{E}[X_i]$, we let $\psi(X_i) = X_i$ in (4). Similarly, to obtain $\frac{d}{dt}\mathrm{E}[X_i^2]$, we let $\psi(X_i) = X_i^2$.

For the case when $M = 2$, the first and second moment dynamics for X_1 are given by

$$\begin{aligned}
\tfrac{d}{dt}\mathrm{E}[X_1] &= \mathrm{E}\left[((X_1+1)-X_1)k_{21}X_2+((X_1-1)-X_1)k_{12}X_1\right] \\
&= k_{21}\mathrm{E}[X_2]-k_{12}\mathrm{E}[X_1] \\
\tfrac{d}{dt}\mathrm{E}[X_1^2] &= \mathrm{E}\left[((X_1+1)^2-X_1^2)k_{21}X_2+((X_1-1)^2-X_1^2)k_{12}X_1\right] \\
&= -2k_{12}\mathrm{E}[X_1^2]+2k_{21}\mathrm{E}[X_1X_2]+k_{21}\mathrm{E}[X_2]+k_{12}\mathrm{E}[X_1].
\end{aligned}$$

When the w_{ij}'s are linear with respect to the system state $\mathbf{X}$, the moment equations are closed. This means that the time derivative for the first moment of X_i, $\frac{d}{dt}\mathrm{E}[X_i]$, is only dependent on the first moments of X_i for $i = 1,\ldots,M$, *i.e.*, $\mathrm{E}[X_1],\ldots,\mathrm{E}[X_M]$, the second moments are dependent on the first and second moments, and so on and so forth. This is important because when the moment equations are closed, the moment dynamics can be expressed as a linear matrix equation.

In general, the ensemble moment dynamics for the system with M tasks/sites is given

$$\begin{aligned}
\tfrac{d}{dt}\mathrm{E}[X] &= \mathbf{K}\mathrm{E}[X] \\
\tfrac{d}{dt}\mathrm{E}[XX^T] &= \mathbf{K}\mathrm{E}[XX^T] + \mathrm{E}[XX^T]\mathbf{K}^T + \Gamma(\alpha, \mathrm{E}[X])
\end{aligned} \tag{5}$$

where $[\mathbf{K}]_{ij} = k_{ji}$ and $[\mathbf{K}]_{ii} = -\sum_{(i,j)\in\mathcal{E}} k_{ij}$. It is important to note that $\mathbf{K}$ is a Markov process matrix and thus is negative semidefinite. This coupled with the conservation constraint $\sum_i X_i = N$ leads to exponential stability of the system given by (5) [11, 17, 23]. Each entry in the matrix of second moments is determined from the moment closure methods shown above where the entries of $\Gamma(\alpha, \mathrm{E}[X])$ are all linear with respect to the k_{ij}'s and the means $\mathrm{E}[X]$. For a system with two states, X_1 and X_2, $\Gamma(\alpha, \mathrm{E}[X])$ is defined as

$$\Gamma(\alpha, \mathrm{E}[X]) = \begin{bmatrix} k_{12}\mathrm{E}[X_1]+k_{21}\mathrm{E}[X_2] & -k_{12}\mathrm{E}[X_1]-k_{21}\mathrm{E}[X_2] \\ -k_{12}\mathrm{E}[X_1]-k_{21}\mathrm{E}[X_2] & k_{12}\mathrm{E}[X_1]+k_{21}\mathrm{E}[X_2] \end{bmatrix}.$$

Similar to the mean-field description, the k_{ij}'s can be chosen to enable a team of robots to autonomously maintain some desired mean steady-state distribution of the team across the various tasks/sites [11, 13, 1]. In both formulations, the k_{ij}'s translate into a set of stochastic guard conditions for the single robot controllers. The result is a set of decentralized agent-level control policies that allow the team to maintain the steady-state mean of the ensemble distribution. Different from the mean-field approach, the formulation of the ensemble dynamics in terms of the moments of the robot population enables us to synthesize distributed control strategies to enable the team to maintain *both the mean and the variance* of the robot team distribution across the various tasks/sites. We describe the approaches in the following sections.

3 Methodology

Regardless of the approach, both (1) (or (2)) and (5) are macroscopic models of the ensemble activity. In this section, we show how these models can be used to analyze the effects of uncertainty on the performance of the team of robots servicing a collection of spatially distributed tasks. Furthermore, we will show how these models can be used to design distributed coordination strategies to improve the team's performance in the presence of the these uncertainties.

3.1 Characterizing and Filtering the Ensemble Noise

Berman *et al.* showed that when the task execution times are stochastic, the ensemble dynamics given by (2) can be approximated using an equivalent expanded linear system [1], *i.e.*, a Multi-Pole approximation. This is achieved by introducing additional dummy transitions between states to approximate the effects of the stochastic delay times. When delay times are deterministic or near deterministic, Mather *et al.* showed that Padé approximants employed in the frequency domain do a better job of capturing the effects of the delay in (2) [22].

To determine the effects of the deterministic (or near deterministic) delays on the overall system performance, we can analyze (2) in the frequency domain [22]. In the frequency domain, the time delay is modeled as an exponential variable. As the frequency increases, the output signal is delayed by more and more periods, in effect, worsening the phase error. The advantage of the Padé approximation is the ability to more accurately capture these effects while retaining the algebraic structure of the differential equations, *i.e.*, (2), in the frequency domain.

Consider the deployment of an ensemble of 10 robots moving in the plane to 2 distinct locations/sites. Initially, robots are randomly assigned to each of the two sites. To simulate the execution of a task at a given site, each robot is tasked to circle the site in a clockwise direction for a fixed time $\tau_i = \tau$. Once the task has been executed, the robot moves to the next site and performs the same task at the new site. The variability in each robot's site-to-site navigation times depends on the amount of traffic on the road, which is affected by the number of collision avoidance maneuvers each robot must execute. Fig. 2(a) shows the Fast Fourier Transform (FFT) of the average output, *e.g.*, x_i in (2), of 54 agent-based simulations performed in USARSim [24]. The frequency response of the agent-based simulations was obtained by logging the population fractions at each site over time and applying the FFT to these variables for each run. The FFT results were then averaged over all 54 runs. The agent-based system exhibits a maximum gain at approximately 7.5 mHz while both the Padé and Multi-Pole macroscopic models exhibit peaks at approximately the same frequency. However, the Padé model shows larger gain [22].

The spurious frequency components shown in Fig. 2(a) manifest as oscillations in the ensemble states in the time domain and is an effect of robots clustering together as they travel from one site to another. These clusters form because too many robots are traveling between tasks resulting in more collision avoidance maneuvers and thus further delaying the arrival of robots at their next task. This leads to degraded

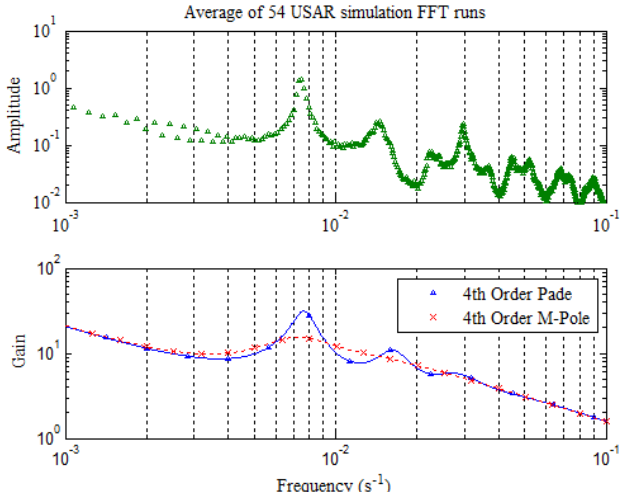

(a) Ensemble Response

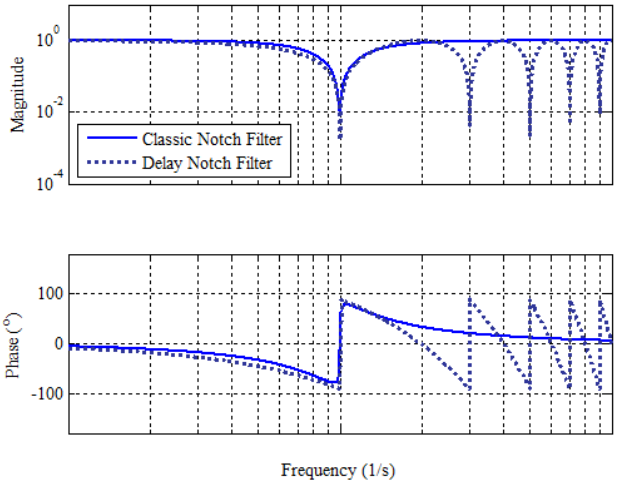

(b) Notch Filter Response

Fig. 2 (a) Top: Average of the FFT of the population fraction at site 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, and a 4^{th} order Multi-Pole macro-continuous systems. (b) Frequency response of the classical 2^{nd} order notch filter and the delayed notch filter given by $H_\tau(s) = \frac{1}{2}(1 + e^{-s\omega_\tau})$.

performance as the average transit time between sites will increase due to these traffic concerns. The Padé approximated macro-continuous model has the ability to better predict the spurious frequency component that is present in the agent-based simulations and can provide insight into the synthesis of agent-level controllers to filter out these spurious frequencies [22].

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 is related by the shape of the frequency spectrum of the noise input to the system.

3.1.1 Distributed Filtering

To smooth the response of the system, a common approach is to implement a notch filter to get rid of the spurious behavior. A notch selectively filters out a specific frequency while leaving other frequency components unchanged, effectively reducing the gain of the single spurious frequency component. A typical 2^{nd} order notch controller has the transfer function $H_1(s)$ given by

$$H_1(s) = \frac{s^2 + 2\zeta_1\omega_N + \omega_N^2}{s^2 + 2\zeta_2\omega_N + \omega_N^2}$$

where ω_N, ζ_1, and ζ_2 set the location and magnitude of the notch. However, careful inspection of the closed-loop time domain equations suggest implementation of the filter will require individual robots to estimate the higher order derivatives of the populations at the various sites. Instead, we propose an approximate solution, where the spurious frequency response can be removed without extra knowledge of the system states by the individual robots.

This can be achieved by splitting the team into two sub-teams where one team purposely adds an additional delay at a given site for each cycle path in $\mathscr{G}$. This approach can eliminate a frequency by adding a signal to a copy of itself, 180^o degrees out of phase. The transfer function for the proposed notch filter is $H_\tau(s) = \frac{1}{2}(1 + e^{-s\omega_\tau})$. The advantage of this approach is that it can be implemented in a completely distributed way without requiring any inter-agent communication. The frequency response plot for this distributed delay notch filter and the classical 2^{nd} order filter are shown in Fig. 2(b). The difference is that the delay filter cancels every odd harmonic of the primary notched frequency [21].

The addition of a single notch filter will suppress a single spurious population behavior. If the task precedence graph has multiple cycles with spurious loops, multiple notches are required to eliminate spurious behavior. While introduction of a delay into a system with feedback can be dangerous since it can lead to enough phase lag to turn negative feedback into positive feedback resulting in unstable

oscillations, the systems discussed here fall into a family of systems that are stable independent of delay [3]. Thus, no amount of extra phase delay can drive the system unstable.

Fig. 3(a) shows the frequency response of the average of 50 agent-based simulations for a team of 10 with two tasks with and without the delay notch filter. Our results show that the distributed notch filter suppressed the undesired frequency component by 70%. Fig. 3(b) shows the frequency response of the average of three experimental trials for a team of 10 robots with and without the delay notch filter. Each experiment ran for roughly 45 to 50 minutes with the robots executing 750 state transitions.

3.2 Controlling the Ensemble Moment Dynamics

While the mean-field approach enables us to design decentralized coordination strategies that can be implemented without any inter-agent wireless communication, it is limited since it can only affect the mean of the distribution. On the other hand, the explicit modeling of the moment dynamics of the robot population distribution using (5) gives the ability to devise ensemble feedback strategies that enables the team to affect the mean, the variances, and any higher order moments of the ensemble distribution.

As shown with equation (3), the rate in which agents in state X_i transition to X_j depends on the population in state X_i. As such, the more agents in state X_i, the faster they transition to X_j. However, Klavins recently showed that if we allow for both positive and negative transition rates, it is possible to shape both the mean and the variance of the ensemble distribution [17]. In other words, by introducing a negative feedback rate, it is possible to slow the population growth at a given state and thus affect the population variance in that state.

For the M state system described by (5), consider the following ensemble feedback controller

$$\mathbf{u} = -\mathbf{K}_\beta \mathrm{E}[X] \quad \mathbf{K}_\beta^{ij} = \begin{cases} \beta_{ji} & \forall (i,j) \in \mathscr{E} \\ -\sum_{(i,j)\in\mathscr{E}} \beta_{ji} & \forall i = j \\ 0 & \text{otherwise} \end{cases}. \tag{6}$$

Expression (6) can be seen as a form of linearizing feedback control that inhibits transitions from X_i to X_j as X_j becomes larger than X_i. This results with the following closed-loop moment dynamics

$$\begin{aligned} \tfrac{d}{dt}\mathrm{E}[X] &= (\mathbf{K}_\alpha + \mathbf{K}_\beta)\mathrm{E}[X] \\ \tfrac{d}{dt}\mathrm{E}[XX^T] &= (\mathbf{K}_\alpha + \mathbf{K}_\beta)\mathrm{E}[XX^T] + \mathrm{E}[XX^T](\mathbf{K}_\alpha + \mathbf{K}_\beta)^T \\ &\quad + \Gamma(\alpha, \beta, \mathrm{E}[X]). \end{aligned} \tag{7}$$

The steady-state values of $\mathrm{E}[X_i]$ and $\mathrm{E}[X_i X_j]$ can be independently set by adjusting parameters α and β. The above equations are obtained by simply substituting

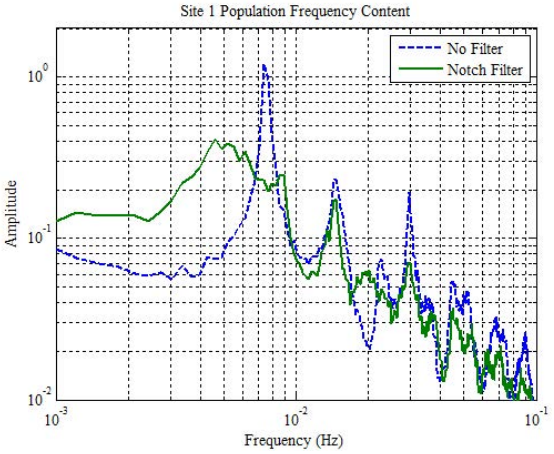

(a)

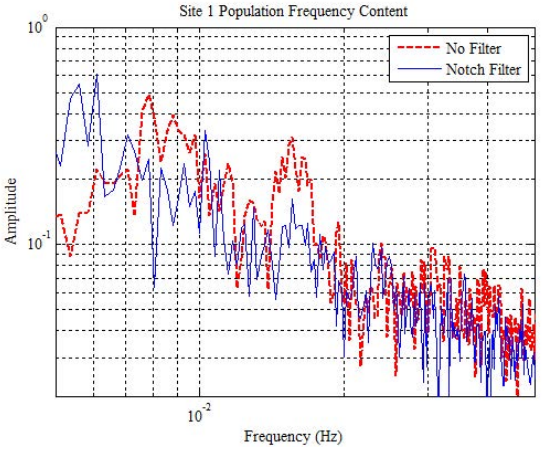

(b)

Fig. 3 (a) Frequency response results for the simulated system. This plot shows the average response of 50 simulations of the unfiltered (no notch) system and 50 simulations of the system with the delay notch filter. The system with the distributed filter shows a depression in the frequency response at the active notching point. (b) Frequency response of the initial experimental trials. The plots shows the frequency content of 3 averaged response for the unfiltered and notched system. The peak, though small, is properly located according to the transition times. The high peak at 6.1mHz in the notched response is due to the round trip time if all agents went along the delay route that includes the τ_{NOTCH} delay.

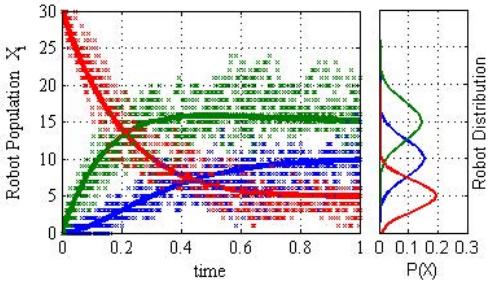

(a)

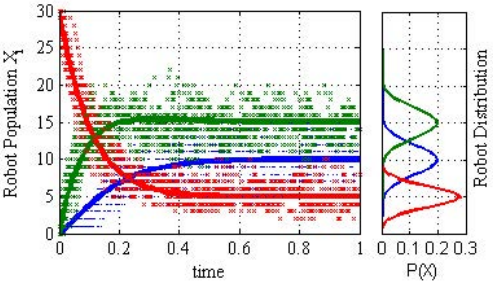

(b)

Fig. 4 These plots compare the steady state distributions and the convergence rate of the system with and without ensemble feedback. Each left side plot shows the transient behavior from an initial condition of $X = [0,0,30]$. The solid lines denote the numerical solutions of the first moment dynamics and the data points are 10 representative stochastic simulation runs. The right side plots are the steady state distributions represented as Gaussians.

$k_{ij} = \alpha_{ij} - \beta_{ij}\frac{X_j}{X_i}$ in the reactions given by (3) and applying the extended generator to $\psi(X_i) = X_i$ [23].

The advantage of the proposed ensemble feedback strategy, over any other negative feedback strategy, is that it maintains the moment closure property for the closed-loop system. This enables us to show that the close-loop moment dynamics remain stable when $\beta_{ij}X_j \leq \alpha_{ij}X_i$ [23]. When $\beta_{ij}X_j > \alpha_{ij}X_i$, the system experiences a backwards flow. As such, in practice, we restrict this rate to be greater than or equal to zero. The addition of these saturation effects will slightly complicate the stability analysis [23].

While the feedback strategy (6) gives robots in state X_i the ability to set their own state transition rates to be independent from the number of robots X_i, it requires robots at task i to know how many robots are at adjacent sites, *i.e.*, robots in X_j where $e_{ij} \in \mathscr{E}$. This differs from the delayed notch filter presented in Section 3.1 which can be implemented with no communication. As such, the implementation of (6) will depend on the available communication infrastructure.

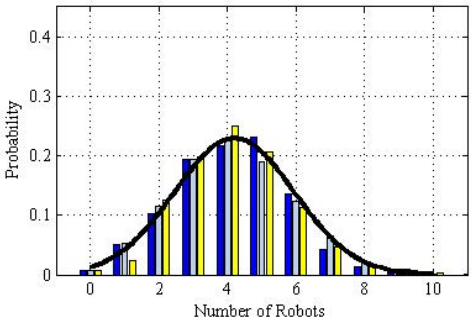

(a)

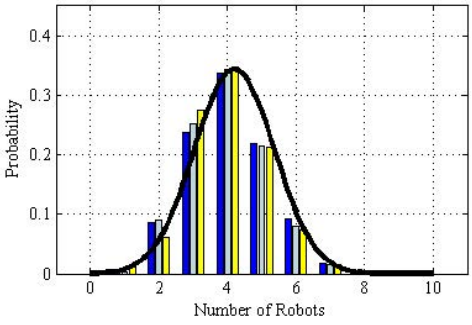

(b)

Fig. 5 Probability distribution of the robot ensemble at each site, $\{X_1, X_2, X_3\}$. (a) Without ensemble feedback and (b) with ensemble feedback using local communication.

If we endow each task location with the computational capability to track the number of robots at the site and the ability to communicate with adjacent task sites, then the estimation of the ensemble states would be similar to having a single global estimator. When individual robots arrive at a location, the information is updated and broadcasted to all adjacent task sites. Fig. 4 shows the first moments of a three state system ($M = 3$ tasks) over time with and without the ensemble feedback strategy. Note how the system with ensemble feedback has both faster convergence and smaller variance on its populations [23].

In practice, not only is it unreasonable to assume full and perfect communication among the robots, it is often unreasonable to assume full and perfect communication between the task sites. This is especially true when sites are distributed across vast geographic regions or in situations where long-range communication is difficult/impossible, *e.g.*, underground/underwater environments. A distributed implemention of (6) can be achieved through local inter-robot communication alone. We assume robots have finite communication ranges and can only communicate with other robots that are co-located at the same site and/or within each other's

communication range. As robots move from one site to another, they can exchange information with other robots they encounter and construct their own estimates of the population levels at the various sites. Fig. 5 shows the allocation of the 15 robots to three task sites with and without the distributed implementation of the ensemble feedback strategy.

4 Discussion and Outlook

In this chapter, we presented a method for synthesizing distributed ensemble feedback control strategies through the development and analysis of an appropriate macroscopic description of the ensemble dynamics. In one case, mean-field models allowed for the identification of the spurious interactions between robots as they moved within a workspace executing a collection of spatially distributed tasks. The macroscopic analysis lead to the development of a distributed filtering strategy that could be implemented without requiring any inter-robot wireless communication nor estimation of population variables. In a second case, moment closure techniques where used to model the dynamics of a team of robots servicing a collection of spatially distributed tasks. The analysis provided a linearizing ensemble feedback strategy to enable the team to maintain the mean and the variance of the robot population distribution across the various tasks.

The key advantage of this approach is a lower dimensional parameterization of the ensemble dynamics that retains the salient features of the underlying agent-based system. These techniques are particularly well-suited for analyzing the effects of uncertain interactions on overall system performance in multi-agent robotic systems. Specifically, these techniques enable the analysis of highly redundant systems in a lower dimensional space while simultaneously retaining the systems-level view of the dynamics. Furthermore, since interaction uncertainties can be explicitly encoded in these models, the feedback strategies developed using these techniques would be robust to any changes in population sizes.

Despite these advantages, further investigation is needed to determine the classes of multi-agent coordination problems that are amenable to these macroscopic modeling and controller synthesis techniques. Specifically, we are interested in investigating the viability of these techniques in modeling and controlling multi-agent robotic systems executing highly coupled tasks. For any ensemble derived feedback strategy, there is also the added challenge of determining the appropriate distributed implementation. However, this presents an opportunity for network resource aware synthesis of distributed coordination and control strategies for multi-agent systems. We are interested in investigating ensemble controller synthesis techniques that can take into account the trade-off between more precise control and the need for estimating ensemble states.

Acknowledgements. We gratefully acknowledge the support of NSF grant DGE-0947936 and CNS-1143941.

References

[1] Berman, S., Halasz, A., Hsieh, M.A., Kumar, V.: 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 (2008)

[2] Berman, S., Halasz, A., Hsieh, M.A., Kumar, V.: Optimized stochastic policies for task allocation in swarms of robots. IEEE Transactions on Robotics 25(4), 927–937 (2009)

[3] Chen, J., Latchman, H.: Frequency sweeping tests for stability independent of delay. IEEE Transactions on Automatic Control 40(9), 1640–1645 (1995), doi:10.1109/9.412637

[4] Dahl, T.S., Matarič, M.J., Sukhatme, G.S.: 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, pp. 307–337. Springer, Berlin (2006)

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

[6] Dias, M.B., Zlot, R.M., Kalra, N., Stentz, A.T.: Market-based multirobot coordination: a survey and analysis. Proceedings of the IEEE 94(7), 1257–1270 (2006)

[7] Gerkey, B.P., Matarić, M.J.: Sold!: Auction methods for multi-robot control. IEEE Transactions on Robotics & Automation 18(5), 758–768 (2002)

[8] Gerkey, B.P., Matarić, M.J.: A formal framework for the study of task allocation in multi-robot systems. International Journal of Robotics Research 23(9), 939–954 (2004)

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

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

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

[12] Hespanha, J.P.: Moment closure for biochemical networks. In: Proc. of the Third Int. Symp. on Control, Communications and Signal Processing (2008)

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

[14] Hsieh, M.A., Halasz, A., Cubuk, E.D., Schoenholz, S., Martinoli, A.: Specialization as an optimal strategy under varying external conditions. In: Accepted the International Conference on Robotics and Automation, ICRA 2007, Kobe-Japan (2009)

[15] Jones, E.G., Browning, B., Dias, M.B., Argall, B., Veloso, M., Stentz, A.T.: Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks. In: Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA 2006), pp. 570–575. IEEE, Los Alamitos (2006)

[16] Jones, E.G., Dias, M.B., Stentz, A.: Learning-enhanced market-based task allocation for oversubscribed domains. In: Proceedings of the Conference on Intelligent Robot Systems (IROS 2007), pp. 2308–2313. IEEE, Los Alamitos (2007)

[17] Klavins, E.: Proportional-integral control of stochastic gene regulatory networks. In: Proc. of the 2010 IEEE Conf. on Decision and Control (CDC 2010), Atlanta, GA, USA (2010)

[18] Lerman, K., Jones, C., Galstyan, A., Matarić, M.J.: Analysis of dynamic task allocation in multi-robot systems. International Journal of Robotics Research (2006)
[19] Lin, L., Zheng, Z.: Combinatorial bids based multi-robot task allocation method. In: Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA 2005), pp. 1145–1150. IEEE, Los Alamitos (2005)
[20] Martinoli, A., Easton, K., Agassounon, W.: 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 (2004)
[21] Mather, T.W., Hsieh, M.A.: Distributed filtering for time-delayed deployment to multiple sites (best paper award winner). In: 10th International Symposium on Distributed Autonomous Robotics Systems (DARS 2010), Lausanne, Switzerland (November 2010)
[22] Mather, T.W., Hsieh, M.A.: Analysis of stochastic deployment policies with time delays for robot ensembles. International Journal of Robotics Research: Special Issue on Stochasticity in Robotics & Biological Systems, Part 1 30(5) (2011)
[23] Mather, T.W., Hsieh, M.A.: Distributed robot ensemble control for deployment to multiple sites. In: 2011 Robotics: Science and Systems, Los Angeles, CA USA (2011)
[24] USARSim, Unified system for automation and robot simulation (2005), `http://usarsim.sourceforge.net`
[25] Vail, D., Veloso, M.: Multi-robot dynamic role assignment and coordination through shared potential fields. In: Schultz, A., Parker, L., Schneider, F. (eds.) Multi-Robot Systems, pp. 87–98. Kluwer (2003)