Distributed Aggregation in the Presence of Uncertainty: A Statistical Physics Approach

M. Ani Hsieh and T. William Mather
Scalable Autonomous Systems Lab
Drexel University
Philadelphia, PA 19104

Abstract

We present a statistical physics inspired approach to modeling, analysis, and design of distributed aggregation control policies for teams of homogeneous and heterogeneous robots. We assume high-level agent behavior can be described as a sequential composition of lower-level behavioral primitives. Aggregation or division of the collective into distinct clusters is achieved by developing a macroscopic description of the ensemble dynamics. The advantages of this approach are two-fold: 1) the derivation of a low dimensional but highly predictive description of the collective dynamics and 2) a framework where interaction uncertainties between the low-level components can be explicitly modeled and control. Additionally, classical dynamical systems theory and control theoretic techniques can be used to analyze and shape the collective dynamics of the system. We consider the aggregation problem for homogeneous agents into clusters located at distinct regions in the workspace and discuss the extension to heterogeneous teams of autonomous agents. We show how a macroscopic model of the aggregation dynamics can be derived from agent-level behaviors and discuss the synthesis of distributed coordination strategies in the presence of uncertainty.

Introduction

The development of robotic teams that can operate in complex and dynamic environments in support of or cooperating with human agents poses significant challenges. In applications such as automation of distribution warehouses, distributed construction and assembly of large-scale infrastructure, and environmental monitoring, teams must have the ability to autonomously distribute and redistribute to ensure the timely completion of the various aspects of the project. This is similar to the adaptive aggregation problem where agents must assess who, when, and how members should assemble into various collectives in order to maximize the overall performance of the system which may be affected by availability of resources, component failures, and/or changes in the environment.

In the multi-robots domain, this aggregation problem can often be posed as a resource allocation problem and is categorized as the multi-task (MT), single-robots (SR) or multi-robots (MR), time-extended assignment (TA) problem depending on whether tasks required single or multiple robots to perform (Gerkey and Mataric 2004). While market-based approaches (Gerkey and Mataric 2002; Dias et al. 2006) have shown much success, especially when learning is incorporated (Dahl, Mataric, and Sukhatme 2006), these methods often scale poorly in terms of team size and number of tasks (Dias et al. 2006; Golfarelli, Maio, and Rizzi 1997). Furthermore, the performance of these methods often degrade significantly when inter-agent wireless communication is extremely limited, noisy, or completely unreliable.

In this work, we present a statistical physics inspired approach where a macroscopic description of the ensemble dynamics is employed to synthesize distributed agent-level control policies to enable autonomous agents to dynamically aggregate and assemble into different groups/collectives. In recent years, macroscopic continuous models have been employed to model the dynamics of robotic self-assembly (Hosokawa, Shimoyama, and Miura 1994; Napp, Burden, and Klavins 2009) and robotic swarm systems (Martinoli, Easton, and Agassounon 2004; Lerman, Martinoli, and Galstyan 2005; Hsieh et al. 2008). These continuous population models are usually obtained by representing the individual robot controllers as probabilistic finite state machines and approximating the collection of discrete Markov processes as a continuous-time Markov process. The macroscopic models are then used to determine the ensemble effects of a series of microscopic, or agent-level, behaviors.

We formulate the aggregation of the team into distinct groups as a resource allocation problem similar to (Halasz et al. 2007). We first consider the aggregation problem for a team of homogeneous robots where the agent-level control policies are obtained via the sequential composition of individual task controllers. We show how uncertainties in the aggregation process that arise from individual interactions can be incorporated into the macroscopic models and show how these models can be further used to inform the design of agent-level control policies (Hsieh et al. 2008; Berman et al. 2008; Mather and Hsieh 2010). Our strategy is inspired by existing work in modeling and control of molecular dynamics where a polynomial Stochastic Hybrid System (pSHS) is often employed to describe the ensemble dynamics (Feinberg 1979; Higham 2008; Klavins 2010;

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

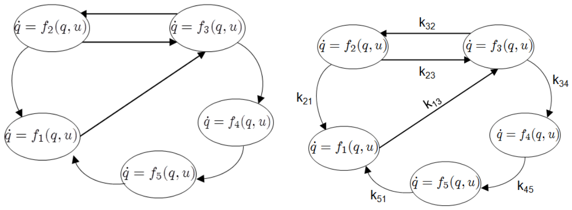

(a) (b)

Figure 1: (a) Robot controller for a set of 5 tasks where $f_i(q, u)$ denotes the open-loop dynamics and u denotes the feedback control strategy of the robot executing task i. This is similar a first order reaction process. (b) Schematic representation of a stochastic hybrid system.

Hespanha 2008). The main contribution is a team-size invariant approach towards the design of distributed agent-level control policies that has the ability to respond to robot failures in a natural way, ensuring graceful degradation.

The rest of the paper is structured as follows: We present our methodology for homogeneous agents in Section and describe the extension of the approach for heterogeneous agents in Section . Section discusses the challenges of employing these models for synthesis of distributed control and coordination strategies. We conclude with a discussion of directions for future work in Section in .

Modeling Homogeneous Teams

In general, the explicit modeling of agent-level interactions can be extremely complex because the dimension of the ensemble configuration space is high and the correlation between the dimensions can be even higher. However, a team of autonomous agents operating within a dynamic environment subject to sensor and actuation noise is similar to a chemical reaction process. At the microscopic level, these systems are composed of various stochastically interacting molecules whose individual behaviors are difficult to predict (Feinberg 1979; Gunawardena 2003; Gillespie 2007). At the macroscopic level, the time evolution of the ensemble statistics for these stochastic systems can be accurately modeled using mass-action kinetics.

To illustrate this, consider the problem of enabling a team of autonomous agents to distribute and aggregate around different tasks that are located at M distinct locations in the workspace. For simplicity, we assume that the agents have the ability to autonomously navigate within the workspace and execute the tasks once they aggregate at the correct locations. We represent the single agent controller as a finite state automaton as shown in Fig. 1(a) where q denotes the agent's state, $f_i(q, u)$ denotes the agent dynamics at location/task i, and u denotes the agent's feedback control strategy at location/task i. Agents transition between tasks based on pre-specified guard conditions assigned to the edges in Fig. 1(a). The guard conditions are determined based on criteria that describe the satisfactory completion of the given task.

In general, the time for an agent to travel between locations can be accurately determined, especially in highly structured environments. However, for an N-agent team, the individual travel times can vary significantly since each agent will have to negotiate different traffic patterns as they operate in a common workspace. This is particularly challenging in the distributed control setting since individuals must operate asynchronously. As such, it makes sense to model the time it takes each agent to move from one location to another (and the time it takes to complete a task) as random variables. In other words, we consider the average arrival rate of the agents at each location rather than the actual time taken by each agent. Let $k_{ij} > 0$ be defined as the transition probability per unit time for an agent to complete task i and travel to location j where k_{ij} is simply the inverse of the average task completion time. Fig. 1(b) is the representation of a collective of N agents whose individual behaviors are described by Fig. 1(a). In Fig. 1(b) k_{ij} defines the transition probability per unit time for an agent in state S_i with dynamics $\dot{q} = f_i(q, u)$ to switch to state S_j.

The directed graph, $\mathcal{G}$, in Fig. 1(b), is an example of a *stochastic hybrid system* (SHS) where the discrete states are represented by the nodes and each may be possess distinct continuous dynamics. In this work, we abstract away the continuous dynamics in the models and assume that individual agents have the ability to store and execute the controllers in each discrete state (u for each $f_i(q, u)$) and only consider the switching dynamics, *i.e.*, $f_i(q, u) = i$. Let $x_i(t)$ denote the fraction of agents in state S_i at time t, then the time evolution of the $x_i, \ldots, x_M$ is given by

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

for all $i = 1, \ldots, M$ where $\mathcal{E}$ denotes the edge set of $\mathcal{G}$.

The above system of linear ordinary differential equations (ODEs) describes the average rate of change of the fraction of robots at each task. The specification in terms of fractions rather than absolute robot numbers allows for a team size invariant formulation which is practical for scaling purposes as well as in situations where losses of robots to attrition and breakdown may happen. Equation (1) models the ensemble dynamics of the agents as they aggregate/cluster around different regions of the workspace. If we think of the transition rates as design parameters, instead of describing the variability in transition times between regions, we can shape the steady-state distribution of the team across the various clusters through the selection of the $k'_{ij}s$ (Halasz et al. 2007; Hsieh et al. 2008). The result is a set of agent-level transition rules that can be implemented at the agent-level without the need for communication. In other words, given the set of controllers in each discrete state, complete knowledge of $\mathcal{G}$, k_{ij}'s, and the ability to localize, the agents can automatically aggregate and assemble themselves accordingly across the M locations *without the need for any inter-agent communication*. This is a team-size invariant solution to the dynamic multi-task (MT) robots, single-robot (SR) tasks, time-extended assignment (TA) problem (MT-SR-TA) (Gerkey and Mataric 2004) and can be extended to the dynamic MT-MR-TA (multi-robot tasks) problem through the appropriate selection of the discrete controllers u.

We note that in this framework, the complexity of the aggregation problem is only dependent on the complexity of the desired high-level group behavior, *e.g.*, autonomous operation of a warehouse or environmental monitoring, and is encoded in the structure of the SHS. While equation (1) allows us to predict the mean ensemble behavior (Martinoli, Easton, and Agassounon 2004; Lerman, Martinoli, and Galstyan 2005) and can be used to synthesize agent-level control policies (Halasz et al. 2007; Hsieh et al. 2008; Berman et al. 2009), it is based on mass action kinetic principles which provides an approximation of the average value of the system state. In theory, this approximation only becomes exact in the thermodynamic limit, *i.e.*, when population sizes approach infinity. However, recent studies have shown that these models can perform surprisingly well even when the group size is small (Mather and Hsieh 2010). This is because, although the number of agents may be relatively small, the number of potential interactions between the various agents may be quite large.

Modeling Heterogeneous Teams

Consider the automation of a distribution warehouse or factory floor where a fleet of autonomous agents must cooperate in order to accomplish a task. To explicitly model aggregation and interactions among heterogeneous agents, we must move beyond first order linear differential equations. Consider the cooperation between two types of agents: robots and humans. Let x_r, x_h, and x_{rh} be the fraction of available robots, available humans, and human-robot teams. This results in the following system of nonlinear rate equations for the ensemble:

$$\dot{x}_{rh} = k_{rh} x_r x_h - \lambda_{rh} x_{rh} \tag{2}$$

where k_{rh} denotes the likelihood per unit of time for a free robotic agent to encounter a free human agent and λ_{rh} denotes the likelihood for a human-robot team to dissolve into free robot and human agents.

Similar to equation (1), the equilibria of the system given by equation (2) is determined by the transition rates. Furthermore, in this setting, interaction uncertainties are explicitly captured in the model via the transition rates. Different from its linear equivalent, these systems can possess multiple steady-states, limit cycles, and even chaotic behavior depending on the choice of system parameters. Similar to the system given by equation (1), there has been significant work in the last thirty years that relates the structure of the agent-level controllers and the inter-agent interactions to conditions for the existence, uniqueness, multiplicity, and stability of equilibrium points (Feinberg 1979; Gunawardena 2003). The challenge then is to determine the ideal set of system parameters for a given desired outcome.

Achieving Distributed Ensemble Feedback

The ability to characterize the uncertainties that arise from the interactions of a multi-agent team operating in a shared complex and dynamic workspace opens up the possibility for an ensemble approach towards the design of distributed agent-level control and coordination policies for aggregation.

For a system described by (1), it was shown in (Mather, Hsieh, and Frazzoli 2010) that spurious frequency components can be predicted via a frequency domain analysis of the linear model. These frequency components represent the effects of noisy interactions among the agents during the aggregation process. If we only consider the ensemble model and disregard the underlying physical system, then we can leverage on decades of control theoretic approaches to design an appropriate filter to get rid of the spurious behavior. One such example is a notch filter which selectively removes a specific frequency while leaving the rest untouched, effectively reducing the gain of the spurious frequency component. A typical 2^{nd} order notch filter has the transfer function $H(s)$,

$$H(s) = \frac{s^2 + 2\zeta_1 \omega_N - \omega_N^2}{s^2 + 2\zeta_2 \omega_N - \omega_N^2} \tag{3}$$

with ω_N and ζ_1/ζ_2 chosen such that the location and magnitude of the notch are properly located at the spurious frequency. While applying this filter to the ensemble model is straight-forward, it is not clear how such a filtering strategy can be achieved in a distributed fashion without requiring agents to estimate the higher order derivatives of the ensemble states.

A common starting point in the synthesis of distributed controllers is to assume perfect communication among the agents, thus providing individuals with full knowledge of the system states. Such an approach invariably results in a distributed control policy that is, in spirit, equivalent to a centralized one. While not ideal, this exercise helps us understand the information required by each robot to achieve a truly decentralized implementation of the proposed strategy where robots only rely on their local information. For any ensemble derived feedback strategy, the challenge lies in determining the appropriate distributed implementation that minimizes the usage of available network resources. As such, the advantage of this approach is the ability to explicitly take into account network resource needs at the controller synthesis stage by minimizing the number of ensemble states that need to be estimated. For the system described by (1), a communication-less implementation of the notch filter can be achieved by approximating the ideal notch filter using a series of carefully placed artificial delays in the agent-level controller. We refer the interested reader to (Mather and Hsieh 2010) for the specific details.

Future Work

In this work, we presented a statistical physics inspired approach towards modeling the dynamics of a aggregation in robot teams in the presence of uncertainty. The advantage of this approach are two-fold: 1) a team-size invariant description of the ensemble and 2) a communication resource aware approach towards the design of scalable distributed control and coordination strategies for aggregation.

While existing work have shown the potential of this approach, fundamental questions still remain. Specifically, what are the limitations of these abstract ensemble models? Given an ensemble feedback controller, how do we distribute the sensing and control to achieve the prescribed control strategy? Are there appropriate ensemble metrics that will enable us to further optimize the topology of the high-level description of the agent-level control policies in order to simply the distributed implementation? Another direction for future work is to determine how one selects the appropriate analysis and design techniques as agent-level interactions become more complex. This is of particular importance since complex interactions between heterogeneous agents would result in more and more nonlinear components in the ensemble dynamics. Finally, the proposed models provide a description of the average ensemble behavior in the presence of uncertainty. A final question of paramount interest is whether we can leverage the inherent ability of noise for self-organization to further develop minimal effort agent-level control and coordination strategies.

Acknowledgments

We gratefully acknowledge the support of NSF grant CNS-1143941.

References

Berman, S.; Halasz, A. M.; Hsieh, M. A.; and Kumar, V. 2008. Navigation-based optimization of stochastic strategies for allocating a robot swarm among multiple sites. In *Proc. 2008 IEEE Conf. on Decision & Control (CDC'08)*, 4376–4381.

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

Dahl, T. S.; Mataric, M. J.; and Sukhatme, G. S. 2006. A machine learning method for improving task allocation in distributed multi-robot transportation. In Braha, D.; Minai, A.; and Bar-Yam, Y., eds., *Understanding Complex Systems: Science Meets Technology*. Berlin, Germany: Springer. 307–337.

Dias, M. B.; Zlot, R. M.; K., N.; and Stentz, A. 2006. Market-based multirobot coordination: a survey and analysis. *Proc. of the IEEE* 94(7):1257–1270.

Feinberg, M. 1979. Lectures on chemical reaction networks. In *Notes of lectures given at the Mathematics Research Centre, University of Wisconsin*.

Gerkey, B. P., and Mataric, M. J. 2002. Sold!: Auction methods for multi-robot control. *IEEE Trans. on Robotics & Automation* 18(5):758–768.

Gerkey, B. P., and Mataric, M. J. 2004. A formal framework for the study of task allocation in multi-robot systems. *Int. J. of Robotics Research* 23(9):939–954.

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

Golfarelli, M.; Maio, D.; and Rizzi, S. 1997. Multi-agent path planning based on task-swap negotiation. In *Proc. 16th UK Planning & Scheduling SIG Workshop*. Durham, England: PlanSIG.

Gunawardena, J. 2003. Chemical reaction network theory for in-silico biologists. Technical report, Harvard University, Cambridge, MA.

Halasz, A.; Hsieh, M. A.; Berman, S.; and Kumar, V. 2007. Dynamic redistribution of a swarm of robots among multiple sites. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'07)*, 2320–2325.

Hespanha, J. P. 2008. Moment closure for biochemical networks. In *Proc. of the 3rd Int. Symp. on Control, Communications and Signal Processing*.

Higham, D. J. 2008. Modeling and simulating chemical reactions. *SIAM Review, Education Section* 50:347–368.

Hosokawa, K.; Shimoyama, I.; and Miura, H. 1994. Dynamics of self assembling systems: Analogy with chemical kinetics. *Artificial Life* 1(4):413–427.

Hsieh, M. A.; H., A.; Berman, S.; and Kumar, V. 2008. Biologically inspired redistribution of a swarm of robots among multiple sites. *Swarm Intelligence*.

Klavins, E. 2010. Proportional-integral control of stochastic gene regulatory networks. In *Proc. 2010 IEEE Conf. on Decision & Control (CDC'10)*.

Lerman, K.; Martinoli, A.; and Galstyan, A. 2005. A review of probabilistic macroscopic models for swarm robotic systems. In Sahin, E., and Spears, W. M., eds., *Swarm Robotics*, volume 3342 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 143–152.

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

Mather, T. W., and Hsieh, M. A. 2010. Distributed filtering for time-delayed deployment to multiple sites (best paper award winner). In *10th Int. Symp. on Distributed Autonomous Robotics Systems (DARS 2010)*.

Mather, T. W.; Hsieh, M. A.; and Frazzoli, E. 2010. Towards dynamic team formulation for robot ensembles. In *Proc. of the 2010 IEEE Int. Conf. on Robotics & Autmoation*.

Napp, N.; Burden, S.; and Klavins, E. 2009. Setpoint regulation for stochastically interacting robots. In *Robotics: Science and Systems V*. MIT Press.