

Specialization as an Optimal Strategy Under Varying External Conditions

M. Ani Hsieh, Ádám Halász, Ekin Dogus Cubuk, Samuel Schoenholz, and Alcherio Martinoli

Abstract— We present an investigation of specialization when considering the execution of collaborative tasks by a robot swarm. Specifically, we consider the stick-pulling problem first proposed by Martinoli *et al.* [1], [2] and develop a macroscopic analytical model for the swarm executing a set of tasks that require the collaboration of two robots. We show, for constant external conditions, maximum productivity can be achieved by a single species swarm with carefully chosen operational parameters. While the same applies for a two species swarm, we show how specialization is a strategy best employed for changing external conditions.

I. INTRODUCTION

Advances in embedded processor and sensor technology have accelerated interests in robotic swarms. In recent years, decentralized coordination strategies for large robot teams have mostly followed a swarming paradigm where robots are programmed with simple but identical behaviors realizable with limited on-board computational, communication, and sensing resources [3]–[5]. These works often take inspiration from their biological counterparts [6], [7] and have shown that teams of autonomous agents can stably achieve consensus via local interactions alone, *i.e.* without centralized coordination. While these results have generated much buzz in the multi-robot community, they are only applicable to homogeneous swarms of robots.

Existing works that consider the coordination problem for heterogeneous teams of robots include [8]–[10]. In particular, market-based strategies, where robots must execute various bidding schemes to determine the appropriate allocation based on perceived costs and utilities, have gained much success in many multi-robot applications [11]. However, these existing methods are primarily concerned with solving the allocation problem and, therefore, team heterogeneity is often assumed rather than explicitly designed for.

There are numerous examples of specialization, *i.e.* when individuals are adapted for a special function, in both biology and economics. A biological example is the phenotypic multiplicity of bacteria [12], where a bacteria colony sometimes has sub-populations of individuals that exhibit a phenotype

(different objective features, behavior, etc.) in spite of being genetically identical to the rest of the colony. These sub-populations are adapted to adverse environmental conditions and represent a hedging strategy, whereby they forfeit rapid growth for enhanced survival capability. In economics, specialization is often a strategy used to increase productivity and attain maximum production efficiency.

Some of the earliest works to investigate the role of specialization in robot teams include [13], [14]. A method for composing group level behaviors from basic robot behaviors was proposed in [13], while [14] proposed an information theoretic measure to determine the degree of team heterogeneity. In [1], sensor-based simulations along with microscopic numerical models for a team of collaborative robots were analyzed and specialization was shown to be beneficial for small team sizes. A discrete-time probabilistic model for collaborative manipulation was presented in [2]. Both [1], [2] considered tasks that required the collaboration of two robots. Li *et al.* used the stick-pulling problem proposed in [1], [2] to show how specialization can be learned by the team using both global and local reinforcement signals [15]. Additionally, [15], [16] considered more general tasks which required more than two robots, however specialization was not discussed in [16].

In this paper we investigate the advantages of specialization when executing collaborative tasks. Specifically, we use the stick-pulling problem, first studied in [1], as an abstraction of more general collaborative tasks. We formulate a macroscopic analytical model to describe the swarm dynamics, similar to [16]. Unlike [1], [15], [16], we assume individual robot dynamics are described by a set of distinct states where all state transitions are governed by Poisson processes and robots are grouped based on the time constants of these Poisson distributions. As such, state transitions within the same group can occur at different times as long as the transition rates are drawn from the same distribution. This, in turn, allows us to describe the system dynamics in terms of the number of free robots within each group over time and to model team productivity as a function of the total number of free robots. We show that while specialization is not necessarily optimal, a group with specialized individuals is often better suited to perform under varying circumstances. As such, we present conditions where team performance can be optimized through specialization and present simulation results to support our findings.

This paper is organized as follows: We formulate our collaborative swarm model in Section II with detailed analysis presented in Section III. Section IV outlines our simulation

M. A. Hsieh is with the Mechanical Engineering and Mechanics Department, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104 mhsieh1@drexel.edu

A. Halász is with the Department of Mathematics, P.O. Box 6310, West Virginia University, Morgantown, WV 26506 halasz@math.wvu.edu

E. D. Cubuk and S. Schoenholz are undergraduate students at Swarthmore College, 500 College Ave., Swarthmore, PA 19081, USA ecubuk1,sschoen1@swarthmore.edu

A. Martinoli is with the Distributed Intelligent Systems and Algorithms Laboratory, École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland alcherio.martinoli@epfl.ch

methodology and summarizes our results. The discussion of these results is presented in Section V. We conclude with some directions for future work in Section VI.

II. PROBLEM FORMULATION

Assume a swarm of N non-communicating robots and a set of M spatially distributed tasks such that the tasks can only be completed by two collaborating robots. Our objective is to determine the conditions under which specialization would result in better team performance.

Specifically, consider the stick-pulling experiment [1], [2], [15], [16], where M sticks are scattered within a given workspace. Each robot is tasked to wander the workspace and pull as many sticks as possible. However, each stick is large enough that it can only be removed from the ground by two robots pulling in collaboration. While the simple two robot stick-pulling problem may seem trivial, it provides a nice abstraction for more general collaborative tasks such as robotic demining, manufacturing and assembly, and even crop harvesting. In fact, if we assume longer sticks where each robot can only lift $1/\kappa$ of the stick, we would be able to model the more general problem of sequential collaboration among κ robots. Similarly, if we assume heavier sticks such that each stick can only be removed by κ robots pulling simultaneously, we would be able to consider the problem of parallel collaborative among κ robots. In this work, we limit ourselves to the investigation of tasks requiring at most two collaborative robots.

To model this, we assume the robot controller consists of three discrete states: WANDER, HOLD, PULL. Initially, robots WANDER the workspace looking for tasks, or more specifically sticks. In general, the rate at which a robot encounters a stick in the workspace depends on M , the workspace geometry, and possibly other factors. Accordingly, we define the *discovery rate*, denoted by k_D , as the rate a robot encounters a stick normalized for such factors. While k_D can be difficult to model, for a given set of parameters, it is possible to obtain k_D empirically. In this work, we will assume that every time a task is completed, or a stick is pulled, it is immediately replaced with an identical uncompleted task. This will enable us to assume that k_D remains constant throughout the simulated experiment.


When a robot encounters a free stick, it switches from the WANDER mode to the HOLD mode and waits for some time interval τ . Should another robot happen upon the same stick, the two would then switch to the PULL mode and cooperatively remove the stick. However, if no other robot encounters the same stick within the time interval τ , the first robot would relinquish the stick and revert to the WANDER mode. While we assume robots are non-communicating, we do assume that robots have the ability to detect the arrival of another robot when it is in the HOLD mode.

In general, heterogeneity can manifest itself in many ways, *e.g.* hardware, software, controller parameters, etc. In this work, we assume the robots are homogeneous at the hardware level and differ only at the controller parameter level. In contrast to [15], we assume our swarm is composed

of p species, each consisting of multiple individuals, all parameterized by a characteristic *waiting time parameter* (WTP) τ_i for $i = 1, \dots, p$.

Rather than assign a constant deterministic WTP to each species, we assume each characteristic WTP follows a Poisson distribution with an expected value of τ_i . This enables us to describe the rate in which robots of species i relinquish their sticks by $\lambda_i = 1/\tau_i$ and refer to this parameter as the *release rate* for species i . By considering the release rate, rather than the WTP, we will be able to describe the swarm level dynamics using a set of continuous-time differential equations unlike [2], [16]. In this work, since individual species are differentiated based on their respective WTPs, we will use the terms WTP, τ_i , and λ_i interchangeably.

For large enough N and M , we model the dynamics of the stick-pulling problem as a chemical reaction process. Let m_f denote the number of free sticks, *i.e.* sticks with no robots waiting, n_i denote the number of wandering robots of species i , and $m_{h,i}$ denote the number of sticks held by robots of species i . The following reaction processes describe the generation and degeneration of each of these elements:



The first process describes the generation of sticks held by species i robots. The second process describes the degeneration/release of sticks held by robots of species i . The last process describes the generation of pulled sticks, *i.e.* successfully completed tasks. Note, since we assume that every pulled stick is immediately replaced by an identical stick, the encounter of a robot of species j with a stick held by a robot of species i results in one free stick and two wandering robots, one of species i and one of species j . Furthermore, we assume the same discovery rate, k_D , for free and held sticks.

The above reactions result in the following set of rate equations:

$$\dot{n}_i = m_{h,i} \left(\lambda_i + k_D \sum_j n_j \right) - k_D n_i m_f$$

for $i = 1, \dots, p$. However, since N and M are constant, we have the following conservation constraints:

$$\begin{aligned} N_i &= m_{h,i} + n_i & \text{for } i = 1, \dots, p, \\ N &= \sum_i^p N_i, \\ M &= m_f + \sum_i^p m_{h,i}, \end{aligned}$$

where N_i denote the total number of robots of species i . Using these constraints and defining the non-dimensional

variable $\zeta_i = \lambda_i/k_D$, we can further simplify the original set of p rate equations to obtain:

$$\frac{1}{k_D} \dot{n}_i = (N_i - n_i) \left(\zeta_i + \sum_j^p n_j \right) - n_i \left(M + \sum_j^p n_j - N \right) \quad (1)$$

for $i = 1, \dots, p$. The above equation describes the time evolution of the wandering population of species i robots. The state of the system is given by the population variables $n_i(t)$ rather than the individual controller states of the robots. Additionally, while k_D and λ_i are parameters of our macroscopic model, they are also the transition rates that define the transition rules between the controller states for each individual robot. Lastly, We note that while the number of robots and sticks are obviously integers, we treat them as continuous numbers in our formulation. This is justifiable for the large values of N and M we are concerned with.¹

To determine a productivity maximizing strategy, *i.e.* the optimal strategy, we use the average rate at which the team is pulling sticks to measure team productivity. Thus, the productivity of the swarm for this particular set of tasks is given by:

$$E = k_D \sum_i n_i \left(N - \sum_i n_i \right). \quad (2)$$

We note that this function solely depends on the total number of wandering robots.

In this work, we assume the parameter N is fixed, and M may or may not be known *a priori*. The objective is to determine the optimal values of λ_i 's given a particular distribution of the swarm across the p species or n_i 's for a set of λ_i 's (equivalently WTPs) selected *a priori*.

III. ANALYSIS

In this section, we analyze the equilibrium conditions for the system described by (1) and show how an optimal strategy for specialization can be determined by simply considering the equilibrium conditions.

Remark 1: We first note that the swarm performance metric given by equation (2) is simply a quadratic function of the total number of wandering robots, regardless of species. Letting $\bar{n} = \sum_i^p n_i$ and rewriting (2) as $E = k_D \bar{n} (N - \bar{n})$, we see that average team production rate is maximized when $\bar{n} = N/2$.

Next, from (1), the equilibrium condition for each species $i = 1, \dots, p$ is given by:

$$\sum_j n_j = \bar{n} = \frac{n_i(M - N + \zeta_i) - N_i \zeta_i}{N_i - 2n_i} \quad (3)$$

Remark 2: While the above equation is nonlinear in terms of n_i , given N , M and a set of ζ_i and N_i for $i = 1, \dots, p$, we can easily compute the equilibrium values of n_i . In addition, given M , p , and a set of ζ_i 's, we can numerically determine the optimal allocation of the swarm to the p species, *i.e.*

¹To further justify this, we can assume that the original integers N and M are normalized by comparing to some large constant P such that $N_{continuous} = N/P$ and $M_{continuous} = M/P$.




Fig. 1. Theoretical pulling rate vs. waiting time for a single species swarm of 1500 robots with 2500 and 1250 sticks with $k_D = 1$.

the values of the N_i 's, that will enable maximum production from the team since maximum productivity is achieved when $\bar{n} = N/2$.² The same holds true if we assume N_i 's are given and solve for the optimal set of ζ_i 's.

Thus, given a swarm of p species, one can determine the optimal allocation of N robots to the p species or the optimal set of λ_i 's (or equivalently WTPs) to ensure peak performance. We consider the special cases of $p = 1$, *i.e.* a single species swarm, and $p = 2$, *i.e.* two species swarm, in the following sections and determine the optimal strategy, *i.e.* the set of λ_i 's or N_i 's for the group.

1) *Special Case 1:* Since $p = 1$, we drop the subscript i and equation (1) simplifies to:

$$\frac{1}{k_D} \dot{n} = (N - n)(\zeta + n) - n(M + n - N) \quad (4)$$

whose equilibrium value is given by the roots of the quadratic function: $2n^2 + (M - 2N + \zeta)n - N\zeta = 0$. Since $n = N/2$ maximizes the production rate, given a particular combination of N and M , the optimal ζ is given by $\zeta = M - N$. Therefore, for a swarm of robots with identical WTPs, the best choice for τ is $1/(k_D(M - N))$.

Note, this is only valid when $M > N$ since it is possible for a team to never achieve the theoretical maximum when $N > M$. Furthermore, as $\tau \rightarrow \infty$, one would expect the productivity to saturate since at most $N - M$ robots would be wandering at any given time at equilibrium. This behavior is seen in Fig. 1 where the theoretical productivity, computed from (2), for the single species swarm, is graphed as a function of τ for two values of M . When $M > N$, we see that peak production rate is indeed achieved when $\tau = 1/(k_D(M - N))$ and k_D was chosen to be unitary. These results are consistent with the experimental results in [1], [2], [15].

Lastly, the solution to (4) that maximizes (2) is locally stable and can be easily shown for the single species system.

²The problem of determining the optimal values for N_i can be posed as a convex optimization problem whose design variables are given by the steady state values of n_i and N_i .

2) *Special Case 2*: When $p = 2$, the system dynamics by (1) result in:

$$\begin{aligned} \frac{1}{k_D} \dot{n}_1 &= (N_1 - n_1)(\zeta_1 + n_1 + n_2) \\ &\quad - n_1(M + n_1 + n_2 - N) \end{aligned} \quad (5a)$$

$$\begin{aligned} \frac{1}{k_D} \dot{n}_2 &= (N_2 - n_2)(\zeta_2 + n_1 + n_2) \\ &\quad - n_2(M + n_1 + n_2 - N) \end{aligned} \quad (5b)$$

with the equilibrium conditions given by

$$n_1 + n_2 = \frac{n_i(M - N + \zeta_i) - N_i \zeta_i}{N_i - 2n_i}$$

for $i = 1, 2$. As mentioned previously, when M is known and ζ_1 and ζ_2 are given, we can solve for the desired allocation of the swarm into the two species or vice versa. Similar to the single species case, the solution that maximizes (2) is also locally stable and can be shown by linearizing the system about the maximizing equilibrium values of n_1 and n_2 .

Consider the case when $\tau_1 = \infty$ and $\tau_2 = 0$, *i.e.* $\zeta_1 = 0$ and $\zeta_2 = \infty$. Here, robots of the first species HOLD a stick forever whenever it encounters one, while robots of the second species never hold onto a stick and only pull sticks when another robot is already there. From Remark 1, we know that maximum production rate is achieved when $\bar{n} = N/2$ under the condition $N < M$. Furthermore, the steady state value of n_2 is given by N_2 . Substituting this into (2) and rewriting it as a function of n_1 results in the following maximizing value for n_1 :

$$n_1^* = (N_1 - N_2)/2.$$

This suggests that while productivity is maximized when $\bar{n} = N/2$, in the extreme case when we choose $\tau_1 = \infty$ and $\tau_2 = 0$, productivity is governed by the number of free robots of the first species who have successfully completed their tasks and are in the process of locating new ones.

While one can solve for the optimal allocation given a particular combination of N , M , ζ_1 and ζ_2 , we believe the simple strategy of splitting the swarm such that $N/2$ robots are parameterized by $\tau_1 = \infty$ and the other half by $\tau_2 = 0$, *i.e.* the extreme case, would generally outperform a single species swarm when the exact number of tasks, M , is unknown. This is simply because as M becomes larger, the two species solution can always ensure at least $N/2$ wandering agents. Therefore, specialization can be seen as a hedging strategy in the face of uncertain conditions, in this case, the total amount of work to be carried out. We support this with simulation results presented in the following section.

IV. SIMULATIONS

A. Methodology

1) *States and transitions*: We consider a system of N robots and M sticks and simulate this as a continuous time

Markov chain. If we let h denote the total number of agents holding onto a stick, then the system with N agents can have

$$\sum_{h=0}^{\min(N,M)} \frac{N!}{h!(N-h)!} \quad (6)$$

distinct states, where we distinguish the agents but not the sticks.³ Similar to [17], the transitions of each agent are stochastic and controlled by individual Poisson processes. For example, the transition from WANDER to HOLD is characterized by a transition *rate* given by $k_D(M-h)$. The discovery rate k_D is defined as the probability per unit time for a wandering robot to encounter a *specific* stick which is modeled by a Poisson process with time constant k_D . Since there are $M-h$ free sticks, the probability per unit time for *this robot* to discover any one of them is given by $k_D(M-h)$.

A Poisson process with time constant k fires at random times with the firing probability per unit time given by k . The process is Markov since the firing probability is independent of past history. The distribution of intervals between two firings can be derived analytically and is given by $p(t) = ke^{-kt}$. Thus, one can simulate Poisson transitions in two mathematically equivalent ways. (1) Run iterations with a small time step $\Delta t \ll 1/k$; at each iteration, the probability of transition is $\Delta p = k\Delta t$. The transition is triggered in the current iteration if $r < \Delta p$, where $0 < r < 1$ is a uniformly distributed random number. This implementation is exact in the limit $k\Delta t \rightarrow 0$ and the methodology was employed in [2], [15]. (2) Generate a random number t_r distributed according to $p(t) = e^{-kt}$ and take the transition at time t_r . This second implementation is exact [18].

Systems with a number of discrete states connected by Poisson transitions are continuous time discrete Markov chains. They have been applied in swarm robotics where the different states can represent either physical locations or behaviors (controllers) chosen by each of the agents [2], [17], [19]. In these works, the state of the swarm is labeled by the number of agents in each state and can be simulated using a *macroscopic discrete* algorithm [18] that is mathematically equivalent to either one of the microscopic (agent-level) implementations described above. A related advantage of stochastic Poisson transitions is that in the limit of large numbers of agents, the time constant of the Poisson process is identical to that of the linear term in the differential equation for the average number of agents (or rather, population fraction $f_i = N_i/N_{total}$ per species). The Master Equation describing the time dependence of occupancy probabilities of different states will have the same analytical form.

The bulk of the simulations performed here are *macro-discrete*. We follow the states of N individual agents. These can be either WANDER or HOLD. There are M tasks, or sticks, in the system which can be completed by two collaborating robots. Since the sticks are immediately replaced

³If we did not distinguish between the agents, the number of states of the system would simply be $\min(M, N)$ labeled by the number of agents holding sticks.

when pulled, at any given time there are h robots holding sticks, $M - h$ free sticks, and $N - h$ wandering robots.

There are three types of transitions, all controlled by Poisson processes. The time constants are listed from the point of view of the active robot.

- *Found* – A free robot i may encounter a free stick; then time constant is $k_D(M - h)$. As a result this robot will shift to a HOLD state.
- *Pull* – A free robot i may collide with a stick held by another robot j . The time constant for this process is $k_D h$ since we assume the probability of discovering a stick is the same regardless of whether it is held or not. As a result, both robots i and j revert to WANDER.
- *Release* – A robot i holding a stick releases it. The time constant for release is λ_i and robot i reverts to WANDER.

Unlike [1] where the WTPs are given by constant deterministic time intervals, we implement the release as a Poisson process.

2) *Microscopic simulation algorithm*: Given the Poisson nature of all possible transitions, the microscopic simulation only requires additional memory and thus is faithful to the corresponding Gillespie simulation [18] from a timing perspective. In addition, a central property of Poisson processes is that the superposition of several Poisson processes is also a Poisson process whose propensity is just the sum of those individual processes. This allows us to bundle similar transitions together in the simulation.

At each iteration, we compute the total propensity for a *Found/Pull* transition and for a *Release*:

$$q_{Found/Pull} = k_D M(N - h), \quad (7a)$$

$$q_{Release} = \sum_{holders} \frac{1}{\tau_i}. \quad (7b)$$

We generate a next transition time for both, and implement the one with the smaller magnitude. We select the actual agent to transition either by a uniform random process or by a random selection weighted by λ_i . The state(s) of the respective agents (and sticks) are updated and the time counter is advanced by the selected transition time.

B. Results

Our first simulations consist of a swarm of 1500 robots with two different numbers of sticks: 2500 and 1250. We assume the swarm is composed of a single species of robot, and $k_D = 1$. Fig. 2 shows the simulated and theoretical steady state average team productivity for a range of WTPs (τ 's) for two combinations of N and M with standard deviation bars shown for each experimental data point.

When $N < M$ (see Fig. 2(a)), maximum productivity occurs when $\tau = 1e - 3$. As τ increases, we see a drop in the average production rate of the group since robots spend less time in the WANDER state which decreases the overall population of wandering robots at equilibrium. Similarly, when $N > M$ (see Fig. 2(b)), the production rate saturates as τ increases. These results are consistent with those described in [15] for teams of 2 to 6 robots.




Fig. 2. Simulated and theoretical pulling rate vs. WTP for a single species swarm of 1500 robots with (a) 2500 and (b) 1250 sticks with $k_D = 1$. The +’s are used to denote the standard deviation for the average team productivity associated with each WTP.

Our next simulations consist of a single and two species swarm of 2000 robots. The simulated steady state average team productivity as a function of different numbers of sticks is shown in Fig. 3. We set $\tau = 0.1$ for the single species swarm and $\tau_1 = \infty$ and $\tau_2 = 0$ with $N_1 = N_2 = N/2$ for the two species swarm with $k_D = 1$ in both scenarios. When $N > M$, the single species swarm consistently outperforms the two species swarm even when τ was specifically chosen to be sub-optimal. However, the two species swarm outperforms the single species for a larger range of M when $N < M$.

V. DISCUSSION

In summary, we used the collaborative stick-pulling problem, first proposed in [1], [2] to investigate the advantages of specialization. We use a macroscopic analytical model to describe the dynamics of the system and derive the equilibrium conditions for a swarm consisting of p species.

For the single species case, we showed that there is an optimal WTP for the swarm when N and M are known and $N < M$. Furthermore, as WTP increases, the productivity decreases since higher WTP tends to decrease the number of wandering robots at equilibrium leaving robots in the HOLD state much longer than necessary. This is seen in both our theoretical and simulated experimental results. From our




Fig. 3. Experimental pulling rate vs. number of sticks for a single and two species swarm of 2000.

results we see that while specialization is not necessarily optimal, a two species swarm with non-optimal WTPs perform better than a single species swarm for a larger range of M . Consequently, in situations when the operating conditions are not well known, it makes sense for the team to specialize. Specialization can be seen as way of mitigating poor performance due to lack of information.

While one may consider specialization beyond $p = 2$, in our current framework, it may be unnecessary to consider more than two species since productivity is solely dependent on the total number of wandering robots. Our task requires the collaboration of two robots, which suggests that only two species is required: wanderers and holders. This may be due, in part, to our stochastic modeling, since [15] showed better performance with multiple species with constant deterministic WTPs.⁴ In general, it seems reasonable to expect that the degree of specialization would have a dependency on the task specifications, which then begs the question of whether it is possible to automatically determine the appropriate level of specialization for a given set of tasks. This is an immediate direction for future work.

VI. CONCLUSION AND OUTLOOK

We have presented an initial investigation of specialization in the execution of collaborative tasks. Specifically, we developed a macroscopic analytical model to describe the dynamics of a swarm executing a set of collaborative stick-pulling tasks. We show when external parameters are well understood, it is unnecessary to specialize since we can design our controller parameters to meet the specific demands. However, specialization is a strategy best employed when these external parameters are affected by uncertainty.

Our initial investigation has opened up many avenues for future work. In this work, we have only considered the simple case where the task only required two collaborating robots. An immediate direction for future work is to determine the global stability properties of our single and two species model with the hopes of generalizing these results

⁴It is also important to note that in [15], each species consists of a single robot.

for the p species model. Another direction is to extend our model to tasks which require the collaboration of k robots. Such analyses will enable us to better understand the trade-offs between single and many species swarms.

Finally, our preliminary investigation suggests it may be possible for a homogeneous swarm to, over time, learn the proper level of specialization for a particular task. Our current efforts are focused on refining the learning strategy to gain further insight into how best to design our systems to be resilient to uncertainty.

REFERENCES

- [1] A. Ijspeert, A. Martinoli, A. Billard, and L. M. Gambardella, "Collaboration through the Exploitation of Local Interactions in Autonomous Collective Robotics: The Stick Pulling Experiment," *Autonomous Robots*, vol. 11, no. 2, pp. 149–171, 2001.
- [2] A. Martinoli, K. Easton, and W. Agassounon, "Modeling Swarm Robotic Systems: A Case Study in Collaborative Distributed Manipulation," *Int. Journal of Robotics Research*, vol. 23, no. 4, pp. 415–436, 2004, special Issue on Experimental Robotics, B. Siciliano, editor.
- [3] L. Chaimowicz, N. Michael, and V. Kumar, "Controlling swarms of robots using interpolated implicit functions," in *Proceedings of the 2005 International Conference on Robotics and Automation (ICRA05)*, Barcelona, Spain, April 2005, pp. 2487–2492.
- [4] M. A. Hsieh, S. Loizou, and V. Kumar, "Stabilization of multiple robots on stable orbits via local sensing," in *Proceedings of the International Conference on Robotics and Automation (ICRA) 2007*, Rome, Italy, April 2007, pp. 2312–2317.
- [5] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Flocking in fixed and switching networks," in *Transactions on Automatic Control*, vol. 52, no. 5, May 2007, pp. 863–868.
- [6] J. K. Parrish, S. V. Viscido, and D. Grunbaum, "Self-organized fish schools: An examination of emergent properties," *Biol. Bull.*, vol. 202, pp. 296–305, June 2002.
- [7] S. C. Pratt, "Quorum sensing by encounter rates in the ant *temnothorax albigipennis*," *Behavioral Ecology*, vol. 16, no. 2, pp. 488–496, 2005.
- [8] U. Tatlıdede, C. Mericli, K. Kaplan, and H. L. Akin, "Q-learning based market-driven multi-agent collaboration in robot soccer," in *Proc. Turkish Symposium On Artificial Intelligence and Neural Networks (TAINN'04)*, Izmir, Turkey, June 10-11 2004, pp. 219–228.
- [9] D. Vail and M. Veloso, "Multi-robot dynamic role assignment and coordination through shared potential fields," in *Multi-Robot Systems*, A. Schultz, L. Parker, and F. Schneider, Eds. Kluwer, 2003.
- [10] B. P. Gerkey and M. J. Mataric, "Sold!: Auction methods for multi-robot control," *IEEE Trans. on Robotics & Automation: Special Issue on Multi-robot Systems*, vol. 18, no. 5, pp. 758–768, Oct 2002.
- [11] M. B. Dias, R. M. Zlot, N. Kalra, and A. T. Stentz, "Market-based multirobot coordination: a survey and analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, July 2006.
- [12] E. Kussell and S. Leibler, "Phenotypic diversity, population growth, and information in fluctuating environments," *Science*, vol. 309, no. 5743, pp. 2075–8, Sep 23 2005.
- [13] M. J. Mataric, "Designing and understanding adaptive group behavior," *Adaptive Behavior*, vol. 4, pp. 50–81, 1995.
- [14] T. Balch, "Hierarchic social entropy: An information theoretic measure of robot team diversity," *Autonomous Robots*, vol. 8, no. 3, pp. 209–238, July 2000.
- [15] L. Li, A. Martinoli, and Y. Abu-Mostafa, "Learning and Measuring Specialization in Collaborative Swarm Systems," *Adaptive Behavior*, vol. 12, no. 3–4, pp. 199–212, 2004, special issue on Mathematics and Algorithms of Social Interactions, C. Anderson & T. Balch (Eds.).
- [16] K. Lerman, A. Galstyan, A. Martinoli, and A. J. Ijspeert, "A Macroscopic Analytical Model of Collaboration in Distributed Robotic Systems," *Artificial Life*, vol. 7, no. 4, pp. 375–393, 2001.
- [17] M. A. Hsieh, A. Halasz, S. Berman, and V. Kumar, "Biologically inspired redistribution of a swarm of robots among multiple sites," *Swarm Intelligence*, vol. 2, no. 2, pp. 121–141, 2008.
- [18] D. Gillespie, "Stochastic simulation of chemical kinetics," *Annual Review of Physical Chemistry*, vol. 58, pp. 35–55, 2007.
- [19] K. Lerman, C. V. Jones, A. Galstyan, and M. J. Mataric, "Analysis of dynamic task allocation in multi-robot systems," *International Journal of Robotics Research*, vol. 25, no. 4, pp. 225–242, 2006.