

Towards Dynamic Team Formation for Robot Ensembles

T. William Mather, M. Ani Hsieh, and Emilio Frazzoli

Abstract— We present an investigation of dynamic team formation strategies for robot ensembles performing a collection of single and two-robot tasks. Specifically, we consider the abstract “stick and pebble” problem, as a variation of the “stick pulling” problem discussed in the literature. We present a formulation of the dynamic team formation problem that is independent of ensemble size and develop a macroscopic analytical description of the ensemble dynamics. The macroscopic model is then used to determine the optimal teaming strategy for two different performance metrics. We present agent-based simulation results to support the validity of our macroscopic analysis.

I. INTRODUCTION

We are interested in developing dynamic team formation strategies for simultaneous execution of various tasks by a robot ensemble. The objective is to enable robots to autonomously form sub-teams, or *coalitions*, and cooperate on tasks that cannot be accomplished by a single robot. The dynamic coalition formation and multi-robot assignment problem is akin to the resource allocation problem [1] and in general is NP-hard. We seek to develop decentralized strategies where team formation can be achieved dynamically, without requiring knowledge of team size and robot identities. We are also interested in strategies that rely solely on local information that can be obtained with minimal communication and sensing requirements. This is critical for large ensemble sizes since it is often impractical to provide individual robots with access to global information.

Existing works that consider the coalition formation problem include [2], where sensor and actuator sharing between robots is achieved by abstracting individual robot controllers to a set of schemas over the task, sensor, and actuator space. This method then computes the task appropriate ensemble of robotic agents. In [3], heuristic ideas from software coalition forming are adapted to develop a market-based approach to task execution. Both [2], [3] rely on combinatorial decision making and high communication rates to optimize the team formation. The dynamic team formation problem for team of heterogeneous robots is considered in [4] while the distributed dynamic task allocation problem is discussed in [5]. Both [4], [5] formulate the respective problems as variations of the dynamic traveling repairman problem.

T. William Mather is with the Scalable Autonomous Systems Laboratory, Mechanical Engineering & Mechanics, Drexel University, Philadelphia, PA, 19104 twmather@gmail.com

M. Ani Hsieh is with the Scalable Autonomous Systems Laboratory, Mechanical Engineering & Mechanics, Drexel University, Philadelphia, PA, 19104 mhsieh1@drexel.com

Emilio Frazzoli is with the Laboratory for Information and Decision Systems, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA frazzoli@mit.edu

Recently, chemical reaction networks have been employed to model, analyze, and synthesize robot ensemble task allocation strategies [6], [7]. These works employ a multi-level representation of ensemble activity. At the lowest, or *microscopic* level, individual robot behaviors are represented by probabilistic finite state machines. At the highest, or *macroscopic* level, chemical reaction network theory (CRNT) is employed to obtain continuous models with a degree of predictive power to describe the ensemble dynamics. In [6], an adaptive multi-robot task with no explicit inter-agent communication or global knowledge is modeled as a stochastic process. A similar approach towards collaborative manipulation is presented in [8], where a discrete-time macroscopic model is used to analyze the collective behavior of the robot ensemble.

Dynamic assignment and reassignment of a robot swarm to multiple parallel independent tasks using a top-down synthesis approach is presented in [7], [9]. The desired swarm allocation is achieved by using the macroscopic models to optimize the individual robot task preference probabilities. Macroscopic models for a robot swarm executing collaborative tasks are used in [10] to show the effects of controller parameter heterogeneity on team performance when external environmental conditions are unknown. These works focus on the allocation problem and not on the dynamic team formation problem.

Similar to [6]–[9], we propose a multi-level representation of robot ensemble behavior. We follow a top-down design approach towards the synthesis of individual robot controllers that result in dynamic team formation at the agent level. We build on [8] and consider the deployment of a robot ensemble to execute a collection of single and two-robot manipulation tasks. Our proposed approach assumes individual robot dynamics are described by discrete states where state transitions are governed by stochastic processes and robots have the ability to form two-robot teams or operate independently. We develop models to describe the ensemble dynamics using CRNT and optimize the individual robot transition probabilities to maximize ensemble performance. Different from existing methodologies, this approach results in a strategy that is invariant to team size and enables a formulation of the resource allocation problem that scales only in terms of mission complexity. We present the resulting strategies based on two different metrics and present microscopic simulation results to support our findings.

II. PROBLEM FORMULATION

Assume an ensemble of N non-communicating robots and two sets of spatially distributed tasks \mathcal{S} and \mathcal{P} . Let \mathcal{S}

denote the set of tasks that can only be completed by two collaborating robots, *i.e.*, two-robot tasks, and \mathcal{P} denote the set of single-robot tasks. The objective is to determine the appropriate team formation strategy to meet some desired ensemble performance metric.

We consider a variation on the stick-pulling experiment [6], [11], where $|\mathcal{S}| = S$ sticks and $|\mathcal{P}| = P$ pebbles are scattered within a given workspace. Each robot is tasked to wander the workspace and remove as many sticks and pebbles as possible. Similar to existing work, we assume that sticks are large enough that they can only be removed by two collaborating robots. Pebbles, on the other hand, can be removed by a single robot. This toy example provides a nice abstraction for more complex missions that are composed of various single and multi-robot tasks, *e.g.* search and rescue. If we assume sticks of varying weights such that some sticks can be removed by κ_i robots while others require κ_j robots, we would be able to consider the general problem of forming different teams sizes within the ensemble. We limit ourselves to the investigation of combinations of single and two-robot tasks in this work.

Assume the individual robot controller consists of five states: 1-WANDER, 2-WANDER, REMOVE-P, REMOVE-S, and HOLD. Robots initially wander the workspace looking for tasks, *i.e.*, sticks or pebbles, either in a 1-robot team, 1-WANDER, or a 2-robot team 2-WANDER. For a uniform distribution of tasks within the workspace, the rate at which a single or a pair of robots encounters a task in the workspace depends on $M = S + P$, $\lambda_S = S/M$ or $\lambda_P = P/M$, the workspace geometry, and possibly other factors. Accordingly, we define the *discovery rate* for a stick or pebble, denoted by k_{d_S} and k_{d_P} , as the rate a robot encounters a stick or pebble normalized by 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 $k_{d_S} = k_{d_P} = k_d$ and every time a task is completed, it is immediately replaced with a task of the same kind randomly placed within the workspace. This will enable us to assume that k_{d_*} remains constant throughout the simulated experiment.

When a single robot encounters a pebble, it switches from the 1-WANDER mode to the REMOVE-P mode. Similarly, if a 2-robot team encounters a stick, the pair switches from the 2-WANDER mode to the REMOVE-S mode. Once either the pebble/stick is removed the 1-robot/2-robot team reverts to either 1-WANDER/2-WANDER mode. However, if a single robot encounters a free stick, it switches from the 1-WANDER mode to the HOLD mode and waits for some time interval τ . Should another single robot happen upon the same stick, the two would then switch to the REMOVE-S mode and cooperatively remove the stick. After removing the stick, the two single robots can decide to remain as a team or split up. Similarly, if a 2-robot team encounters a free pebble, it switches from the 2-WANDER to REMOVE-P mode. Once the pebble has been removed, the 2-robot team can decide to stay as a pair or dissolve into two 1-robot teams. Finally, robots can switch accordingly from 1-WANDER to 2-WANDER and




Fig. 1. The individual robot controller. The k_* represent rates of transition. The θ_* represent a weighting on the rates. $\theta_D + \theta_F = 1$.








vice versa as they wander the workspace and encounter one another. The individual robot controller is shown in Fig. 1.

Rather than choose a constant waiting time interval when a robot encounters a stick, we assume robots draw their τ_i from an exponential distribution with an expected value of τ . This enables us to parameterize the exponential distribution by $k_R = 1/\tau_i$ and refer to this parameter as the *release rate* of the sticks by single robots. Similarly, 1-robot teams will probabilistically decide to form 2-robot teams with propensity θ_F after removing a stick and 2-robot teams will decide to dissolve into two 1-robot teams with propensity θ_D after removing a pebble.

For large enough N and M , we model the dynamics of the pebble and stick removal problem as a chemical reaction process and define the following population variables:

n_R	Single robots
n_{2R}	Two robot teams
m_S	Free sticks
m_P	Free pebbles
m_{SR}	Occupied sticks
$\beta = (S + P)/N$	Ratio of total tasks to agents
$\lambda_S = S/M$	Fraction of sticks in M
$\lambda_P = P/M$	Fraction of pebbles in M

The following *robot reaction processes* describe the production and consumption of each of these elements:



where k_{TF} and k_{TD} denote the team formation and dissolution rates respectively. Process (1a) describes the formation and release of sticks held by 1-robot teams. Process (1b) describes the removal of pebbles by 1-robot teams and process (1c) describes the formation/dissolution of 2-robot teams. Process (1d) and (1e) describe the removal of sticks and pebbles, respectively, by 2-robot teams and their subsequent propensity to stay as 2-robot teams or dissolve into two 1-robot teams. Process (1f) describes the removal of sticks by two 1-robot teams and their subsequent propensity to remain as two single robots or form 2-robot teams. Process (1g) describes the encounter of a 2-robot team with a held stick.

The above reactions result in the set of rate equations shown in Fig. 2. These equations describe the time evolution of the population of 1-robot and 2-robot teams, n_R and n_{2R} , and the populations of free and held sticks, m_S and m_{SR} . The state of the system is given by $c = [n_R, n_{2R}, m_S, m_{SR}]^T$. The system is subject to the conservation constraints $N = n_R + 2n_{2R} + m_{SR}$ and $S = m_s + m_{SR}$ since N and S are constant. While the number of robots, sticks, and pebbles are obviously integers, we treat them as continuous numbers in our formulation. This is justifiable for the large values of N and M .¹ Although $k_d\theta_D$ and $k_d\theta_F$ are parameters of the macroscopic model, they are also the transition rates that define the transition rules between controller states for the individual robots. Lastly, in our formulation, the complexity of the coordination problem depends solely on the complexity of the mission at hand and is invariant with respect to N .

To determine a productivity-maximizing strategy, *i.e.*, an optimal strategy, we define two metrics. The first metric describes the average rate in which the ensemble removes pebbles and sticks from the workspace and is given by

$$E_1 = \alpha_S(k_d m_{SR} n_R + k_d n_{2R} m_S + k_d n_{2R} m_{SR}) + \alpha_P(k_d n_R m_P + k_d n_{2R} m_P) \quad (2)$$

where $\alpha_S, \alpha_P > 0$ are constant weights. The first term in (2) is the average removal rate for a stick while the second term is the average removal rate for a pebble. In our formulation, we assume $\alpha_S = \alpha_P = 1$. The second metric describes the average time a pebble or stick has to wait before being removed and is given by:

$$E_2 = T_S + T_P$$

where $T_S = \frac{m_S}{k_d m_{SR} n_R + k_d n_{2R} m_S + k_d n_{2R} m_{SR}}$

$$T_P = \frac{m_P}{k_d n_R m_P + k_d n_{2R} m_P}. \quad (3)$$

In this work, we assume the N is fixed and S and P are given *a priori*. The objective is to determine the optimal values of n_R and n_{2R} and the related values of $k_d\theta_S$ and $k_d\theta_F$ given S and P .

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 $M_{continuous} = M/P$.

III. ANALYSIS

In this section, we analyze the equilibrium conditions for the system shown in Fig. 2. In the case when an optimal strategy exists in the space of feasible solutions, we show how the optimal mix of 1-robot/2-robot teams can be achieved for any combination of 1-robot/2-robot tasks using two metrics.

A. Stability of Equilibrium Solutions

Recall, the state of our system is given by $c = [n_R, n_{2R}, m_S, m_{SR}]^T$. We refer to elements of c as *species* and reactions (1a-g) are made up of *complexes*, *i.e.*, the reactants and byproducts of the reactions. The rate equations (Fig. 2) can be represented as a graph, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, whose nodes represent the complexes and the edges denote the reactions. We define $\Psi(c) = [n_R + m_S, m_{SR}, n_R + m_{SR}, n_R^2, n_{2R}, n_{2R} + m_P, n_{2R} + m_S, n_R + m_P, n_{2R} + m_{SR}, n_{2R} + n_R]^T$ as the vector of complexes with the rate equations equivalently expressed as

$$\frac{d}{dt}c = Y A_k \Psi(c). \quad (4)$$

In general, given I species and J complexes, Y is an $I \times J$ matrix such that $Y_{ij} = 1$ if one degree of species i is part of complex j and $Y_{ij} = 0$ otherwise. A_k is the $J \times J$ negative weighted directed Laplacian matrix for \mathcal{G} given by

$$A_{k_{ij}} = \begin{cases} -k_{ij} & (i, j) \in \mathcal{E} \\ \sum_{l \in \mathcal{E}} k_{il} & i = j \\ 0 & \text{otherwise} \end{cases}$$

System (1), represented by the graph \mathcal{G} , is weakly reversible because we assume that every completed task is immediately replaced with an equivalent new task. By the zero deficiency theorem [12], the system given by Fig. 2 with constant N , S , and P has a single stable positive equilibrium. Furthermore, the stability of the equilibrium is independent of the reaction rates. Such a property allows us to select the appropriate team formation and dissolution rates, k_{TF} , k_{TD} , and θ_{*s} , based on the desired ensemble performance metric without affecting the stability of the system.²

B. Team Formation Strategy

In this section, we determine the optimal teaming strategy based on two metrics: average task removal rate (2) and average task wait time (3). We parameterize the workspace using the values: $\beta = (S + P)/N$ which is the proportion of total tasks to robots, $\lambda_P = P/M$, and $\lambda_S = S/M$. For this analysis, all k_{d*} are assumed to be one. From our conservation constraints the domain of the population is

$$\mathcal{D} = \begin{cases} n_R \geq 0 \\ n_{2R} \geq 0 \\ N - n_R - 2n_{2R} \geq 0 \\ \lambda_S \beta N - N + n_R + 2n_{2R} \geq 0 \end{cases} \quad (5)$$

²When completed tasks are *not* replaced by new tasks, the resulting rate equations will no longer be weakly reversible. However, in this scenario, the number of uncompleted tasks will always be decreasing, and as such one can employ a Lyapunov argument to show that the system is stable.

$$\begin{aligned}
\dot{n}_R &= 2k_d\theta_{TD}n_{2R}(m_S + m_{SR} + m_P) + 2k_{TD}n_{2R} + k_R m_{SR} + k_d\theta_{TD}n_R m_{SR} \cdots \\
&\quad - k_d n_R m_S - k_d\theta_{TF}n_R m_{SR} - 2k_d k_{TF} n_R^2 \\
\dot{n}_{2R} &= \frac{1}{2}k_d k_{TF} n_R^2 + k_d\theta_{TF}n_R m_{SR} - k_d\theta_{TD}n_{2R}(m_S + m_{SR} + m_P) - k_{TD}n_{2R} \\
\dot{m}_S &= k_R m_{SR} - k_d n_R m_S + k_d n_R m_{SR} + k_d n_{2R} m_{SR} \\
\dot{m}_{SR} &= k_d n_R m_S - k_R m_{SR} - k_d n_R m_{SR} - k_d n_{2R} m_{SR}
\end{aligned}$$

Fig. 2. Rate equations obtained from the reactions outlined in (1). The analysis contained in this work was done on the system with replacement. This means that the sticks and pebbles that are removed from the space are immediately replaced.

1) *Case 1: Heterogeneous Robots:* We consider the case when we limit 1-robot teams to only execute on 1-robot tasks and 2-robot teams to only execute on 2-robot tasks. This is equivalent to case when (1b-d) and (1g) are the only governing reactions in the system. Then the metrics (2) and (3) simplify to

$$\begin{aligned}
E_1 &= \alpha_S(k_d n_{2R} m_S) + \alpha_P(k_d n_R m_P) \\
E_2 &= \frac{n_{2R} + n_R}{k_d n_{2R} n_R}
\end{aligned}$$

respectively. If $\lambda_P/\lambda_S > 0.5$, a population of all 1-robot teams will maximize E_1 . If $\lambda_P/\lambda_S < 0.5$, all 2-robot teams will maximize E_1 . This is because E_1 rewards completion of 1-robot and 2-robot tasks equally. Since E_2 does not depend on the number of sticks nor the number of pebbles the optimal allocation of results when $n_R = (\sqrt{2}-1)N$, and $n_{2R} = (1 - \sqrt{2}/2)N$.

2) *Case 2: Homogeneous Robots:* In this section we consider the case when robot teams are not limited to specific types of tasks, *i.e.*, when all of the reactions in (1) are valid.

a) *Average Removal Rate:* Under (2), robots have the incentive to remove as many items out of the workspace as possible. The gradient of E_1 with respect to n_R and n_{2R} is given by

$$\nabla E_1 = \begin{bmatrix} N - 2n_R - 2n_{2R} + \lambda_P \beta N \\ \beta N - 2n_R \end{bmatrix}$$

From the above equation, E_1 has an equilibrium point at,

$$\begin{bmatrix} n_R \\ n_{2R} \end{bmatrix} = \frac{N}{2} \begin{bmatrix} \beta \\ (1 - \lambda_S \beta) \end{bmatrix}$$

with the Hessian of E_1 given by

$$\mathcal{H}(E_1) = \begin{bmatrix} -2 & -2 \\ -2 & 0 \end{bmatrix} \Rightarrow \text{saddle point}$$

The saddle equilibrium in the system dictates that the maximal removal rate must lie on the boundary of the domain. Since the Hessian is constant, the location of the saddle point determines the optimal population distribution. Fig. 3 shows the shape of E_1 versus n_{2R} and n_R . The saddle, marked with an \times is visible in the center. The black squares are optimal values of n_R and n_{2R} .

On the boundary there are two possible solutions. When $n_R = 0$, the entire robot ensemble is made up of 2-robot teams with $m_{RS} = 0$ and $n_{2R} = N/2$. The optimal removal rate for the population of $N/2$ 2-robot teams will be $\frac{1}{2}\beta N^2$. For $n_{2R} = 0$ there is a range of possible solutions on the $[N - m_{RS}/2, N]$ line segment. However, n_R cannot dip

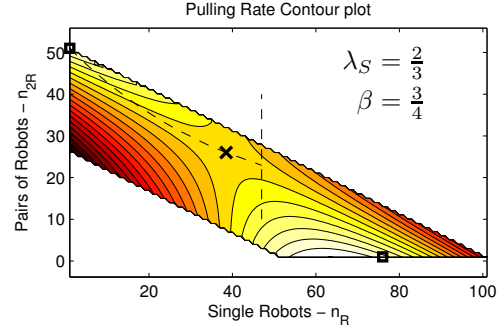


Fig. 3. The removal rate E_1 as a function of n_R and n_{2R} . The saddle point is marked by \times in the center and the optimal distributions of n_R and n_{2R} are marked by the black squares are optimal values of n_R and n_{2R} .

below $N - m_{RS}/2$ because that is the condition that results in single robots staying in the hold position for ever, $k_R = 0$. As the robots search for occupied sticks to clear, they are equally likely to find a new stick as they are to clear an occupied one. The rate of occupying empty sticks is equivalent to the rate of clearing occupied sticks.

Given that the solution is on either side of the domain, it is useful to define a boundary in the space that delineates equivalent strategies. The following equations give the conditions where E_1 is equivalent for team configurations of “all-pairs”, $n_{2R} = N/2$, and all singles, $n_R = N/2$, $k_R \rightarrow 0$, respectively.

$$\beta = \frac{2\lambda_S - 2}{\lambda_S^2 - 2\lambda_S}, \quad \frac{n_R}{2n_{2R}} = \frac{\lambda_P}{\lambda_S} \quad (6)$$

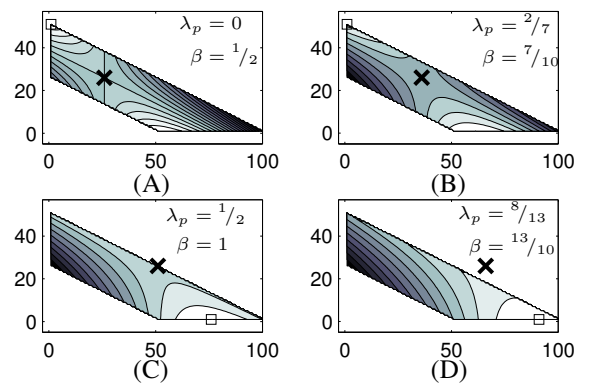


Fig. 4. This shows how the saddle equilibrium, denoted by \times , moves for different values of clutter, β , and proportion of pebbles, λ_P . The squares on each graph represent the max of E_1 .

An interesting result that emerges from this analysis is that the all pairs solution will always be the optimal solution when there are no pebbles. In the case of extremely low β the pulling rate for pairs is still $\beta N^2/2$. The solution for singles has the chance to be much bigger than that. If there are many more agents than sticks, the system could potentially reach a point in which every stick has an agent waiting. This would significantly outperform the pairs solution, $\beta N^2/2 < \beta N^2$. However, due to the nature of the system, all of the sticks being filled is extremely unlikely. When the number of occupied sticks, m_{SR} , equals the number of free sticks, m_S , the chance of a robot clearing the occupied stick and occupying an empty stick are the same. This is because $m_{SR} = m_S$ is a stable equilibrium point for systems with no pairs and high stick waiting times. This means that for single robots the pulling rate is upper bounded by $\beta N^2/2$, the solution for all pairs.

b) *Average Task Waiting Time*:: From a queuing theory perspective, a better behavior is one that minimizes the average amount of time each stick or pebble has been waiting to be pulled out of the space given by (3). Since minimizing E_2 is equivalent to maximizing E_2^{-1} , i.e., the harmonic average of the pebble and stick removal rates, our discussions will focus on the problem of maximizing E_2^{-1} .

There are some plain differences between the two optimizations. First, E_1 is indifferent about what gets removed and thus is maximized by finding the fastest way to clear as many objects as possible. If the system can drag out more pebbles than sticks, then a population of all single robots is the optimal clearing behavior. In fact, that is the optimal solution for systems with high β and $\lambda_P > \lambda_S$. Under these conditions, the optimal teaming strategy is to form no teams and only pick up pebbles. The task waiting time strategy throws out those solutions as it punishes queue instability.




Fig. 5. Harmonic mean of the rates changes with the number of sticks m_S .

It also important to note that the waiting time is independent of the number of pebbles in the space. In the macroscopic rate equations, the pebble pulling rate is linearly dependent on the population variable m_P . This means that the average waiting time for a single pebble is $T_P =$

TABLE I
SUMMARY OF TEAM FORMATION STRATEGIES

$\{n_R, n_{2R}\}$	Case 1		Case 2	
	E_1	E_2	E_1	E_2
$\lambda_P > \lambda_S$	$\{N, 0\}$	(*)	$\{N, 0\}$	$\{0, N/2\}$
$\lambda_S > \lambda_P$	$\{0, N/2\}$	(*)	$\{0, N/2\}$	$\{0, N/2\}$

(*) = $\{(\sqrt{2} - 1)N, (1 - \sqrt{2}/2)N\}$

$[k_d(n_{2R} + n_R)]^{-1}$. By observing 5, it is evident that the high resource tasks, or sticks, dominate the behavior of the waiting time metric. In our model, paired robot teams can execute on any task they find, so for optimizing the average waiting time, focusing on the resource demanding tasks while still maintaining an ability to execute on smaller tasks gives the population of 2-robot teams a big advantage. The average task waiting time is a metric that quantifies how bad the worst case is. Since T_S will always be greater than T_P in this formulation, the results will strongly favor 2-robot teams. We summarize our results in Table I.

IV. SIMULATION RESULTS

The macroscopic description of the ensemble behavior is an approximation of the average behavior of the microscopic system and only become exact when populations tend toward infinity. To show that the previous analysis apply for robotic systems, we present agent-based simulation results to support our analysis and to claim that these macroscopic models can be useful for analyzing and synthesizing collective behaviors.

In our agent-based simulations, robots are treated as point masses with a fix sensing radius to model simple, non-communicating robots. Each robot moves at unit speed with basic collision avoidance protocol. They move in straight lines until they encounter other robots, tasks, or the workspace boundary. As the robots move into each other's collision range, they take a noisy right turn to move around each other. When the robots hit the workspace boundary, they turn around with a preference towards the farther wall to avoid getting stuck in the corners. When two single robots meet, they can join a team or stay as 1-robot teams based on the team formation and dissolution rates obtained from the macroscopic analysis. When a 2-robot team clears a stick, it can go about its way or it can split up depending on the chosen rates. For example, if $\theta_{TF|S} = p$, when a pair clears a stick, the 2-robot team will randomly choose with probability p to stay as a pair or separate. The following results correspond to agent-based simulations of an ensemble of 30 robots, operating within a 10×10 workspace, possessing a sensing radii of 0.3 units for different values of S and P .

Fig. 6 shows the correspondence of the agent-based simulations (top) with the macroscopic results (bottom). Each shaded block in the top graph represents an (n_R, n_{2R}) pair and the steady-state removal rate, E_1 , attained in the agent-based simulation. In these figures the lighter the shade the higher the value of E_1 .

V. DISCUSSION

From our analysis, regardless of the metric employed, maximal performance for the ensemble is achieved when

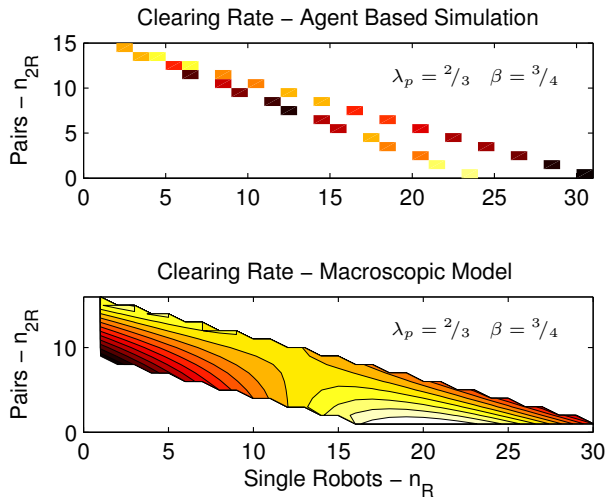


Fig. 6. Top: Steady-state removal rates, E_1 , for different pairs of (n_R, n_{2R}) in the agent-based simulations. Bottom: Corresponding steady-state removal E_1 predicted by the macroscopic models.

either the entire ensemble is tasked to operate as either 1-robot or 2-robot teams. From a utility-based point-of-view, *i.e.*, when one considers E_1 , it makes sense that either sticks or pebbles are seldom picked up since there is no penalty imposed for ignoring hanging tasks. From a queuing perspective, such a solution results in queue instability. However, if one considers the average task waiting time, E_2 , there is a tendency to over-penalize the ensemble for not finding the less likely tasks. In other words, if $S \ll P$, using E_2 results in a strategy that too strongly favors the formation of two-robot teams. Such a behavior seems to suggest that the metrics employed in our study are sensitive to “outliers”. This is of significance since the 2-robot behaviors are often more complex and difficult to implement in hardware.

Of particular interest is the development of additional metrics that can result in mixed team initiatives that varies as a ratio of the different tasks within the workspace. Specifically, given the individual removal rates or task wait times, it may be possible to develop metrics that are less sensitive to outliers through the use of robust statistical methods. Another incentive for considering mixed team initiatives is in situations when the total number of tasks or the exact mix of tasks is unknown. A mixed team ensemble can provide added robustness to unforeseen or unmodeled environment conditions. Lastly, we note that in situations where dynamic team formation is desired, our macroscopic models are capable of accurately describing the ensemble behavior through the appropriate selection of the formation and dissolution rates given by k_{TF} , k_{TD} , and θ_* .

VI. FUTURE WORK

In this work we presented a study of the applicability of the chemical reaction network models to the study of dynamic team formations in robot ensembles. We show that the CRNT framework enables us to model, analyze, and

design for teaming strategies that are independent of team size and scales solely in terms of the mission complexity. Our simulation results confirmed our analysis of the macroscopic models and their ability to predict the behavior of the agent-based simulations.

There are numerous directions for future work. Of particular interest is the further refining of our agent-based simulations by incorporating more refined agent-based controllers for task servicing similar to those presented in [5]. This will also require the development of new techniques to incorporate explicit modeling of the task service times and the effects of more deterministic navigation controllers into the macroscopic models. Another direction for future work is the development of on-line estimation of task distribution and composition within the workspace to enable individual robots to adapt the different transition rates to varying external conditions. Lastly, we are interested in validating the macroscopic models against experimental data obtained from an actual multi-robot testbed. From our agent-based simulations, it is clear that task discovery rates vary significantly for different team sizes and numbers of tasks. This is because as the number of robots increase within the same workspace, robots will spend most of their time executing collision avoidance maneuvers rather than completing tasks.

REFERENCES

- [1] B. P. Gerkey and M. J. Mataric, “A formal analysis and taxonomy of task allocation in multi-robot systems,” *Int. Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, September 2004.
- [2] L. E. Parker, “Building multirobot coalitions through automated task solution synthesis,” *Proceedings of the IEEE*, July 2006.
- [3] L. Vig and J. A. Adams, “Coalition formation: From software agents to robots,” *Journal of Intelligent Robotic Systems*, 2007.
- [4] S. Smith and F. Bullo, “The dynamic team forming problem: Throughput and delay for unbiased policies,” *Systems & Control Letters*, June 2009.
- [5] M. Pavone, E. Frazzoli, and F. Bullo, “Adaptive and distributed algorithms for vehicle routing in a stochastic and dynamic environment,” *IEEE Trans Automatic Control*, 2009.
- [6] 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.
- [7] S. Berman, A. Halasz, M. A. Hsieh, and V. Kumar, “Optimized stochastic policies for task allocation in swarms of robots,” *IEEE Trans on Robotics*, vol. 25, no. 4, pp. 927–937, Aug 2000.
- [8] A. Martinoli, K. Easton, and W. Agassounon, “Modeling of swarm robotic systems: a case study in collaborative distributed manipulation,” *International Journal of Robotics Research: Special Issue on Experimental Robotics*, vol. 23, no. 4-5, pp. 415–436, 2004.
- [9] M. A. Hsieh, A. Halasz, S. Berman, and V. Kumar, “Biologically inspired redistribution of a swarm of robots among multiple sites,” *Swarm Intelligence*, December 2008.
- [10] M. A. Hsieh, A. Halasz, E. D. Cubuk, S. Schoenholz, and A. Martinoli, “Specialization as an optimal strategy under varying external conditions,” in *Proc Int Conf on Robotics and Automation (ICRA07)*, Kobe-Japan, May 2009.
- [11] 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.
- [12] M. Feinberg, “Lectures on chemical reaction networks,” in *Notes of lectures given at the Mathematics Research Centre, University of Wisconsin*, Madison, WI, 1979.