Adaptive Robotic Communication Using Coordination Costs*

Avi Rosenfeld, Gal A Kaminka, and Sarit Kraus

Bar Ilan University, Ramat Gan, Israel
{rosenfa, galk, sarit}@cs.biu.ac.il

Summary. Designers of robotic groups are faced with the formidable task of creating effective coordination architectures that can deal with changing environment conditions and hardware failures. Communication between robots is one mechanism that can at times be helpful in such systems, but can also create a time and energy overhead that reduces performance. In dealing with this issue, various communication schemes have been proposed ranging from centralized and localized algorithms, to non-communicative methods. In this paper we argue that using a coordination cost measure can be useful for selecting the appropriate level of communication within such groups. We show that this measure can be used to create adaptive communication methods that switch between various communication schemes. In extensive experiments in the foraging domain, multi-robot teams that used these adaptive methods were able to significantly increase their productivity, compared to teams that used only one type of communication scheme.

1 Introduction

Groups of robots are likely to accomplish certain tasks more quickly and robustly than single robots [3, 5, 7]. Many robotic domains such as robotic search and rescue, demining, vacuuming, and waste cleanup are characterized by limited operating spaces where robots are likely to collide. In order to maintain group cohesion under such conditions, some type of information transfer is likely to be helpful in facilitating coherent behavior in robotic group tasks and thus better achieve their task. This is especially true as robotic domains are typically fraught with dynamics and uncertainty such as hardware failures, changing environmental conditions, and noisy sensors.

Questions such as what to communicate and to whom have been the subject of recent study [7, 11, 12]. In theory, communication should always be advantageous–the more information a robot has, the better. However, assuming communication has a cost, one must also consider the resources consumed in communication, and whether the cost of communication appropriately matches the needs of the domain. We believe that different communication schemes are best suited for different environmental conditions. Because no one communication method is always most effective, one

* This work was partially supported by the Israeli Ministry of Science and Technology.

way to improve the use of communications in coordination is to find a mechanism for switching between different communication protocols so as to match the given environment.

This paper provides such an adaptive communications framework using a coordination cost measure that quantifies all resources spent on coordination activities. Our model explicitly includes resources such as the time and energy spent communicating. In situations where conflicts between group members are common, more robust means of communication, such as centralized models, are most effective. When collisions are rare, coordination methods that do not communicate and thus have the lowest overhead, work best.

We present two novel domain-independent adaptive communication methods that use communication cost estimates to alter their communication approach based on domain conditions. In our first approach, robots uniformly switch their communication scheme between differing communication approaches. In this method, robots contain full implementations of several communication methods, and switch between them as needed. In contrast, our second approach represents a generalized communication scheme, that allows each robot to adapt independently to its domain conditions. Each robot creates its own communication range radius (which we refer to as its *neighborhood of communication*), to create a sliding scale of communication between localized to centralized methods. Each robot uses its coordination cost estimate to determine how large its neighborhood should be.

To evaluate these adaptive methods, we performed thousands of trials using an established robotic simulator, in a multi-robot foraging task. We tested groups of varying sizes and communication methods. We found that groups that used the adaptive methods often significantly exceeded the best productivity levels of the non-adaptive algorithms they were based on.

2 Related Work

A major challenge to designers of robotic groups exists in choosing an optimal communication method. Many practical frameworks have been presented for use within robotic teams [3–9, 12] and can generally be assigned to categories of no communication, localized, and centralized approaches.

It is possible to create effective group behavior without any communication [2]. For example, the Stigmergy concept [6] involves group members basing their actions by observing how other group members previously modified their environment. This approach has been shown to be effective in several animal and robotic domains [6] without using any explicit communication. Coordination without communication can potentially facilitate better adaptability, robustness and scalability qualities over methods using communication [11]. Additionally, the lack of communication also allows such methods to be implemented on simpler robots. However, such algorithms often require powerful and accurate sensing capabilities [9]. Also, our results demonstrate that groups implementing these methods did not always provide the highest levels of productivity, especially within dynamic domains where frequent coordination conflicts exist.

A second set of approaches attempt to improve group performance by having robots locally communicate information [7, 9]. For example, work of Jäger and Nebel [7] present a method whereby robots nearing a collision stopped to exchange trajectory information. They then successfully detect and resolve deadlock conditions of two or more robots mutually blocking. However, their trajectory planning method was not able to perform well in groups of over five robots. In contrast, Mataric [9] reported that a local communication scheme scaled well with group size. One key difference seems to lie within the local communication implementations. In Jäger's algorithm, one or more robots must stop moving during trajectory replanning. We believe this led to a breakdown in the system once the group size grew. Mataric's locally communicating robots broadcast information while continuing their foraging task. This allowed for better scalability qualities.

A third type of approach involves the use of some type of central repository of information [12]. This centralized body, which could also be implemented as one "expert" teammate, would then be able to easily share its store of pooled information with other teammates. While this approach allows for free information sharing and can thus improve performance, several drawbacks are evident. First, the centralized mechanism creates a single point of failure. The cost of communication is also likely to be large, and requires hardware and bandwidth suitable for simultaneous communication with the centralized body. While these drawbacks are at times significant, they may be justified given the needs of the domain.

In this work, we assume that representative communication methods from these categories are predefined, and have been implemented with optimal values for their exact parameters given domain conditions. Several approaches exist for finding these parameters within a given coordination method. For example, work by Yoshida et al. [4] presented a framework to derive an optimal localized communication area betweens within groups of robots to share information in a minimum of time. This approach assumes domain conditions such as spatial distributions and the probability of information transmission can be readily calculated. Previously, Goldberg and Mataric [5] focused on *interference* (which they defined as the time robots spent colliding) as a basis for measuring a coordination method's effectiveness. However, they did not address how to create adaptive methods based on interference. Our previous work [10] built upon this interference definition to include all resources spent resolving coordination conflicts including the time spent before and after collisions. We then demonstrated that parameter tweaking is possible through this measure. The advantage to this approach over the work of Yoshida et al. [4] is its ability to allow robots to autonomously adapt, even in dynamic environments. However, in contrast to their work, our previous work [10] did not study communication issues.

In this work, we use coordination cost measures to compare a given set of communication methods and to create adaptive methods based on these existing methods. We explicitly model all resources spent on coordination activities including the resources spent on communication even if they do not detract from the time to complete the task. Our goal was to properly match communication methods to domain conditions, while considering their relative costs. Furthermore, adaptation between communication schemes presents new challenges, since many protocols require standardized communication between all team members. These challenges are addressed in this paper.

3 Using Coordination Costs to Adapt Communications

Our coordination cost measure facilitates identifying which communication method is most suitable given the environment. We model every robot's coordination cost $\mathcal{C}_i$, as a factor that impacts the entire group's productivity. We analyze two cost categories: (i) costs relating to communication and (ii) proactive and/or reactive collision resolution behaviors. We focus on the time and energy spent communicating and in the consequent resolutions behaviors (see Implementation Section for full details). We then combine these factors to create a multi-attribute cost function based on the Simple Additive Weighting (SAW) method [14] often used for multi-attribute utility functions. While methods with no communication have no $\mathcal{C}_i$ for communication, this method could not always successfully resolve collisions and then spent more resources on collision resolution behaviors, or another $\mathcal{C}_i$. Conversely, centralized methods incurred a communication cost $\mathcal{C}_i$ that often eclipsed the needs of the domain and weighed heavily on productivity. Other communication issues, such as bandwidth limitations, can similarly be categorized as additional cost factors as they impact any specific robot. For example, if a robot needed to retransmit a message due to limited shared bandwidth, costs in terms of additional time latency and energy used in retransmission are likely to result.

Using this measure we can compare communication methods, but in this paper we focus on using it for online adaptation between communication schemes. In this section we present two types of adaptive methods: (i) uniform communication adaptation (ii) adaptive neighborhoods of communication. Both methods led to significant increases in productivity over static approaches (see Experiments section).

3.1 Uniform Switching Between Methods

In our first method, all robots simultaneously switch between mutually exclusive communication methods as needed. In order to facilitate this form of adaptation, each robot autonomously maintains a cost estimate, V used to decide which communication method to use. As a robot detects no resource conflicts, it decreases an estimate of this cost, V, by an amount W_{down}. When a robot senses a conflict is occurring, the value of V is increased by an amount W_{up}. The values for V are then mapped to a set of communication schemes methods ranging from those with little cost overhead such as those with no communication, to more robust methods with higher overheads such as the localized and centralized methods. As the level of projected conflicts rises (as becomes more likely in larger group sizes) the value of V rises in turn, and the robots use progressively more aggressive communication methods to more effectively resolve projected collisions. While these activities themselves constitute a cost that detracts from the group's productivity, they are necessary as more simple behaviors did not suffice. As different coordination methods often have different costs, $\mathcal{C}_i$ for a given domain, we believed this approach could be used to significantly improve the productivity of the group.

Several key issues needed to be addressed in implementing this method with groups of robots. First, we assumed that all group members are aware of the overheads associated with various coordination methods, and can order them based on their relative complexities. This ordering can be derived from theoretical analysis or through observation (as we do in later in this paper). Second, an approach to quickly set the weights, W_{up}, and W_{down} used within our algorithms is needed. While traditional learning methods, such as Q-learning [13] may converge on an optimal policy, this approach is difficult to implement because of two major reasons. First, Q-learning is based on a a concept of "state" that is not readily definable during task execution. As opposed to clearly defined discrete domains, there is no reward for any given cycle of activity in the robotic domains we studied. Even assuming an optimal policy could be learned, a second, more fundamental problem exists. Robotic domains often contain dynamics that render a learned policy obsolete very quickly. Thus, our approach is to sacrifice finding a globally optimal policy in exchange for finding a locally optimal policy after a much shorter training period for our weights. We used a gradient learning procedure to achieve this result.

Next, it must be noted that uniform adaptation requires all robots to change communication in sync because of the mutual exclusivity of the methods used. For example, it is impossible for one robot to use a centralized method, with others using one without communication, as the centralized approach is based on information from all team members. As a result, once any one robot in the group autonomously decided it needed to switch communication schemes, a communication change must also occur within all other team members. This could force certain members to use a more expensive communication method than it locally found necessary. We relaxed these requirements in the second adaptive method, presented in the next section.

Finally, care must be taken to prevent the robots from quickly oscillating between methods based on their localized conditions. In our implementation, communication adaptation was triggered once one robot's value for V exceeded a certain threshold. After this point, that robot broadcasted which method it was switching to and all group members would change in kind and reinitialize their cost estimates V to this new value. Furthermore, we also used domain specific information, such as prioritizing collisions closer to the home base within our foraging domain. In this fashion, we partially limited the types of triggers to those of importance to the entire group. Once again, our second type of communication adaptation relaxes this requirement and is effective without any such heuristics.

3.2 Adaptive Neighborhoods of Communication

The advantage in our first adaptive approach lies in its simplicity. Our uniform adaptive approach switches between existing coordination methods based on estimated coordination cost. Assuming one analyzes a new domain with completely different communication methods, and can order the communication methods based on their communication costs, this approach will be equally valid as it implements existing methods and reaches the highest levels of productivity from among those methods–whatever they may be.

In contrast, our second adaptation method is a parameterized generalization of the three specific categories of communication methods (No-Communication, Localized, and Centralized). As many robotic domains use elements of these same methods [3, 4, 6–8, 12], we reason that a similar approach is likely to work in these and other domains as well.

The basis of this approach is introducing a parameter to control how large a radius of communication is used by each robot. This method uses a distance d inside which robots exchange information, which we term its communication neighborhood. Formally, this radius of communication could be considered a neighborhood Γ of size d, created from robot v and includes all teammates, u, inside this radius. As such, we represent the neighborhood as $\Gamma_d(v) = \{u| \ u$ robot, $dist(u, v) \leq d\}$.

Adjusting the value of d in Γ_d can be used to approximate the previously studied communication categories. Assuming d is set to zero, no communication will ever be exchanged and this method is trivially equivalent to the No-Communication method. Assuming d is set to some small amount, ε, this method will become similar to the Localized method and information will be exchanged only with the robot it is about to collide with. If d is set to the radius of the domain, the neighborhood of communication encompasses all teammates this method becomes similar to the Centralized method. Thus, the degree of centralization exclusively depends on the value of d.

4 Implementation Details

We used the Teambots [1] simulator to implement communication schemes involving no communication, localized and centralized approaches within groups of Nomad N150 foraging robots. The foraging domain is defined as locating target items from a search region *S*, and delivering them to a goal region *G* [5]. Foraging robots often collide as they approach the home base(s) within their area of operation. In our domain there was only one goal region, *G*, which was located in the center of the operating area. In our implementation, there were a total of 60 target pucks spread throughout an operating area of approximately 10 by 10 meters. We measured how many pucks were delivered to the goal region within 9 minutes by groups of 2–30 robots using each communication type. We averaged the results of 100 trials for each group size with the robots being placed at random initial positions for each run. The number of trials performed and the relatively large group sizes would have been difficult to implement on physical robots.

We created experiment sets measuring the time and energy spent in two coordination categories–communication and collision resolution. The coordination costs in our first set of experiments involved the time spent in communication and collision resolution behaviors out of each trial's total time of 9 minutes. In our second set of experiments, we allocated each robot 500 units of fuel. We assumed most of the fuel was used by the robots to move, with a smaller amount (1 unit per 100 seconds) used to maintain basic sensors and processing. For the time based experiments, we assumed robots pairs stopped for 1/5 of a second to communicate, representing some methods [7] where robots stop to exchange information. In the energy based localized experiments, we assumed robots did not stop to communicate, as is the case with other methods [9], but each robot still spent 0.3 units of fuel per communication exchange. Our coordination cost involved the amount of fuel that was used in communication and repulsion behaviors.

All three communication schemes were similar in that they resolved collisions by mutually repelling once they sensed a teammate within a certain safe distance ε, which we set to 1.5 robot radii. Once within this distance, robots acted as they were in danger of colliding and used repulsions schemes to resolve their collision(s). The No-Communication method was unique in that robots never used time or fuel to communicate, and thus only had costs relating to the repulsion behaviors robots engaged in. However, this method assumed domain specific information, namely it used the robot's autonomously computed scalar distance, $\mathcal{S}$, from its location to the home base in the domain. Robots used a function of this distance, which we implemented to be 5 times $\mathcal{S}$ and rounded to the closest second, as the time to repel from its teammate(s) after a projected collision.

Our localized method used less domain specific information and is similar to the localized methods previously proposed [7, 9]. Communication between robots was initiated once it was in danger of colliding–a teammate came within the ε distance. After this event, these group members would exchange information above their trajectories (here their relative distances from their typical target, their home base). The closer robot then moved forward, while the other robot repelled for a fixed period of 20 seconds.

Our final method, *Centralized*, used a centralized server with a database of the location of all the robots similar to other centralized methods [12]. Within this method, one of two events triggered communication. First, as with the localized method, robots dropping within the ε distance initiated communication by reporting its position, done here with the centralized server. The server then reported back a repel value based on its relative position to all other teammates. However, in order for the server to store a good estimate of the positions of all robots, a second, often more frequent type of communication was needed where each robot reported its position to the server with frequency L. If this communication occurred too frequently, this central database would have the best estimate of positions, but the time or energy spent on communication would spike, and productivity would plummet. If communication was infrequent, the latency of the information stored on the server would create outdated data. This in turn would reduce the effectiveness of this method, and result in more collisions. We found that a latency time of 1 second yielded the highest productivity in the time based experiments, and a latency value of 5 seconds yielding the highest productivity within the energy based experiments.

It is important to stress that the focus of our work is switching between categories of communication methods, and not to find optimal parameters within any one given communication method. We refer the reader to previous work [4, 10] on how to theoretically or empirically derive parameters within one communication method. Our work is based on the understanding that a high negative correlation exists between each groups' productivity and our coordination cost, regardless of the exact implementation for the parameters used in the No-Communication, Localized and Centralized methods. For example, we studied 7 latency variations within the Centralized method in both experiment sets. Our groups enforced maximal latency periods of L set to 0.1, 0.3, 1, 5, 10, 30 and 60 seconds. While the optimal latency value was different across experiment sets, in both cases the productivity of these variations was highly negatively correlated with their relative coordination costs. In the first case, we found a correlation of -0.95 between these latency variations and the corresponding coordination cost based on time. In the trials based on fuel, this value was -0.97. Similarly, factors such as the exact energy spent on communication exchanges (0.3 units of fuel per exchange) or the time spent on communication (1/5 seconds per exchange) could vary across domains, or the distance between communication partners. However, we consistently found that the resources spent on these communication exchanges was strongly negatively correlated with those groups' productivity.

While we consider the neighborhood communication approach to be a parameterized generalization of the three previously described categories, some implementation details differ in this method over the static ones it emulates. Within this method, once any robot A, detects another robot within the ε distance, it initiates communication with all robots found within the $\Gamma_d(A)$ area. All robots in $\Gamma_d(A)$ must then report back to Robot A with their projected trajectories. Robot A then sorts all robots' trajectories by their relative distances from the home base in the domain. This robot then reports back to every robot within $\Gamma_d(A)$ a repel value based on that robot's relative position in the neighborhood. All robots, including the initiating robot (robot A), then accept this value and immediately engage in repel behaviors for the dictated length of time. It is possible that a robot may be a member of more than one neighborhood. In such cases, robots accept the larger repel value regardless of the sender.

While the repel amounts of the robot initiating communication (Robot A) are calculated in a similar fashion to the previously described centralized method, here these values are calculated by members of the team, instead of one centralized server. The radius of communication in the centralized approach is the full width of the domain, while the Γ_d radius is typically much smaller. However, the biggest difference in implementing this approach is how repel values are obtained. Robots in previous methods only repelled based on communication received after dropping within the ε distance. In this method, robots may repel if they enter the Γ_d radius even if they are not in immediate danger of colliding. The reason for this is as follows. As robots within the Γ_d radius are typically close to each other, we found that these robots often would soon initiate their own radii of communication. In other methods this was not a concern, as other teammates were not effected by this phenomenon. However, here this would create multiple neighborhoods involving the same teammates. Thus, proactively assigning repel values was crucial for containing communication costs as Γ_d grew.

5 Experimental Results

The first set of experiments attempts to first lend support to the underlying hypothesis, that the combined coordination cost measure is in fact correlated to the productivity of the different groups. Our results from experiments involving time and energy costs support the claim that the best method of communication does change with domain conditions (see figure 1). In the time experiments, we found an average correlation of -0.96 between the average productivity found in groups of 2–30 robots and the group's corresponding average cost. In the equivalent energy based experiments, we found a value of -0.95.

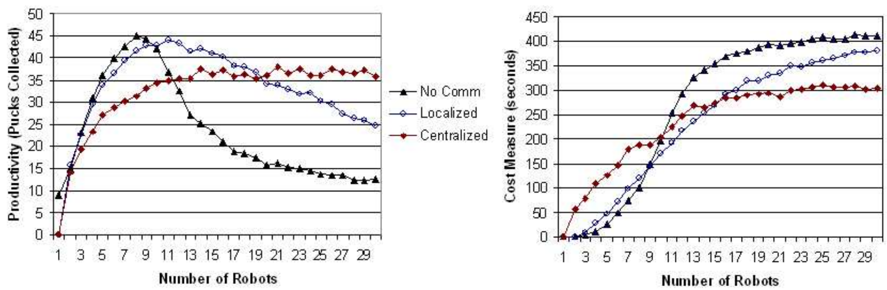

Fig. 1. Comparing the productivity levels of three communication types with the coordination costs based on the time spent on communication relative to different group sizes. Results averaged from 100 trials per datapoint.

Similarly, we found that no one neighborhood size always fared best. We compared the productivity levels of foraging groups where d was set to 1, 2, 3, 5 and 50 robot lengths. Recall that ε is approximately 1 robot length (1.5 radii). Thus Γ_1 represents the nearly localized variation with Γ_{50} corresponding to the nearly centralized version of this method.

Figure 2 represents the relative productivity levels for these static neighborhood groups relative to the energy costs levels measured in these groups. Notice how in small groups, Γ_1 yielded the highest average productivity. As we have seen, when possible, resources spent on coordination, here by creating large communication neighborhoods, should be avoided when possible. As small areas of communication sufficed in small groups, this approach had the highest productivity. As the group size grew, additional communication was necessary to maintain high productivity levels. As a result, larger neighborhoods were necessary and groups with Γ_5 resulted in the highest productivity. However, forcing too much communication when not necessary created communication costs that reduced productivity to levels found in methods that spend too few resources on communication. In this method, the productivity level of the Γ_{50} method, which created too large a neighborhood, approached those of Γ_1, which did not create a large enough one. We again found a strong correlation between the various Γ_d variations and the groups' corresponding coordination costs and productivity with an average negative correlation of -0.96.

Based on the confirmed hypothesis, that the cost measure is indeed correlated (negatively) with performance, the next set of experiments evaluated the performance of the two adaptive methods compared to the static methods on which they were based. Figure 3 shows the results from these experiments. Notice that both adaptive approaches approximated or significantly exceeded the highest productivity levels

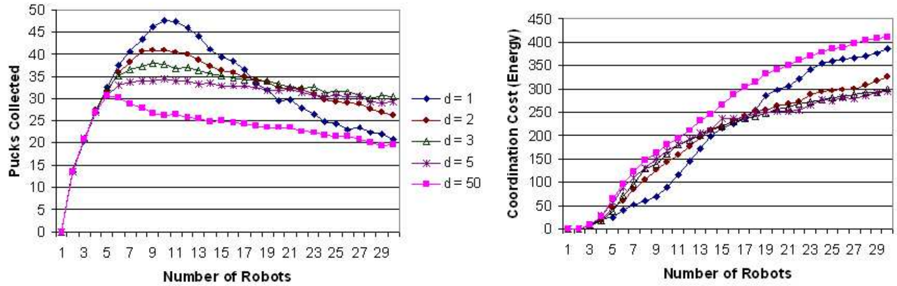

Fig. 2. The impact of varying neighborhood sizes (d) on productivity levels and costs in energy experiments. Results averaged from 100 trials per datapoint.

of the static methods (No Communication, Local, and Centralized methods) they were based on, especially in medium to large groups. We attribute the success of both methods to their ability to change communication methods to the needs of the domain. We believe that the neighborhood method outperformed the uniform one as it was allowed to create locally different neighborhood sizes, something none of the static neighborhood methods were capable of. This in turn facilitated better adaptation and higher productivity.

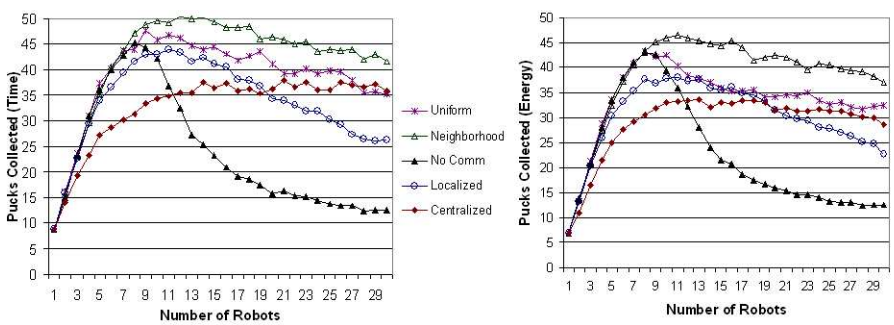

Fig. 3. Comparing adaptive communication methods based on time and energy costs to static methods. Results averaged from 100 trials per datapoint.

To evaluate the statistical significance of these results, we conducted the two tailed t-test and a 1-factor ANOVA test comparing our adaptive groups and the three static groups they were based on. In all cases, in both time and energy categories, the null hypothesis p values were below 0.001. This confirms our hypothesis that we can improve productivity through creating adaptive methods based on communication costs.

6 Conclusion

This work demonstrates how coordination costs can account for the relative effectiveness of robotic communication methods. Our measure focuses on the time and energy spent communicating and resolving collisions. We demonstrate the effectiveness of our methods in comparing between very different communication methods falling within categories of no communication, localized and centralized communication methods. By using this information we are able to match the most effective communication scheme to a given robotic domain. We present two general adaptive communication algorithms, uniform and neighborhood methods. We show, in thousands of foraging experiments, that coordination cost is indeed negatively correlated with productivity, and that the use of our adaptive methods leads to significant performance boosts. While we find the neighborhood adaptive method to be more effective in the robotic foraging domain we studied, both approaches are likely to be applicable to many other domains [3, 7, 8, 12]. It is possible that the uniform method is easier to implement or will yield better adaptive qualities in other domains.

References

1. T. Balch. www.teambots.org, 2000.
2. T. Balch and R. Arkin. Behavior-based formation control for multirobot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939, December 1998.
3. G. Dudek, M. Jenkin, and E. Milios. A taxonomy for multi-agent robotics. *Robot Teams: From Diversity to Polymorphism, Balch, T. and Parker, L.E., eds., Natick, MA: A K Peters*, 3:3–22, 2002.
4. M. Y. J. O. Eiichi Yoshida, Tamio Arai. Local communication of multiple mobile robots: Design of optimal communication area for cooperative tasks. *Journal of Robotic Systems*, 15(7):407–427, 1998.
5. D. Goldberg and M. Matarić. Design and evaluation of robust behavior-based controllers for distributed multi-robot collection tasks. In *Robot Teams: From Diversity to Polymorphism*, pages 315–344, 2001.
6. O. Holland and C. Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artif. Life*, 5(2):173–202, 1999.
7. M. Jager and B. Nebel. Decentralized collision avoidance, deadlock detection, and deadlock resolution for multiple mobile robots. In *IROS*, pages 1213–1219, 2001.
8. G. A. Kaminka and R. Glick. Towards robust multi-robot formations. In *ICRA-06*, 2006.
9. M. J. Matarić. Using communication to reduce locality in multi-robot learning. In *AAAI/IAAI*, pages 643–648, 1997.
10. A. Rosenfeld, G. Kaminka, and S. Kraus. Adaptive robot coordination using interference metrics. In *The Sixteenth European Conference on Artificial Intelligence*, pages 910–916, August 2004.
11. S. Sen, M. Sekaran, and J. Hale. Learning to coordinate without sharing information. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 426–431, Seattle, WA, 1994.
12. A. Tews. Adaptive multi-robot coordination for highly dynamic environments. In *CIMCA*, 2001.
13. C. J. C. H. Watkins. Learning from delayed rewards. *Ph.D. Dissertation, Kings College*, 1989.
14. K. Yoon and C. Hwang. *Multiple attribute decision making: an introduction*. Prentice Hall, Thousand Oaks: Sage, 1995.