Measuring the Cost of Robotic Communication

Avi Rosenfeld, Gal A Kaminka, Sarit Kraus*
Department of Computer Science
Bar Ilan University, Israel
{rosenfa, galk, sarit}@cs.biu.ac.il

1 Introduction

Groups of robots are likely to accomplish certain tasks more quickly and robustly than single robots, [Jager and Nebel, 2001]. Many robotic domains such as robotic search and rescue, 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 useful between members of the team. 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 [Sen *et al.*, 1994], [Jager and Nebel, 2001], [Tews, July 2001]. At times, forms of implicit coordination have been shown to allow agents better adaptability, robustness and scalability qualities [Sen *et al.*, 1994]. In theory, the lack of communication allows such methods to be implemented on simpler robots. A second series of approaches attempt to improve group performance by having robots locally communicate information [Jager and Nebel, 2001]. A third type of approach involves the use of some type of central planner [Tews, July 2001]. We believe that each type of communication framework is best suited for different environmental conditions. A mechanism is needed to match the proper system to the given environment. This paper attempts to provide such a framework with its use of a coordination cost measure.

We measure all coordination costs including the time and energy spent on coordination. We use this measure to evaluate the cost of communication. This result also allows us to create robots that alter their communication scheme when faced with dynamic domain conditions.

2 Talking the Right Amount

We model every robot's coordination cost $\mathcal{C}_i$, as a factor that impacts the entire group's productivity. This cost includes the time and energy used during communication and also proactive and / or reactive collision resolution behaviors. We found that the use of communication, or lack thereof, can impact the time or energy used in the second type of behaviors.

*This material is based upon work supported in part by NSF grant #IIS-0222914, BSF grant #2002401, and the Israeli Ministry of Science and Technology. Sarit Kraus is also affiliated with UMIACS.

We also are able to create adaptive coordination methods that switch between different communication methods based on each robot's estimate of its coordination cost. When robots detect no resource conflicts, it decreases an estimate of this cost, V, by a certain amount W_{down}. When a robot senses a conflict is occurring, the value of V is increased by a certain amount W_{up}. Thus, the value V is constantly in flux based on the robot's perception of its environment. The values for V are used to switch between 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. This approach is novel in that robots individually measure these values online during task execution and can thus adapt their communication method as needed in response to unforseen changes such as group or environment changes.

We proceeded to implement variations of three types of communication schemes – implicit, localized, and centralized methods. In our implementation, all communication types are similar in that they resolve collisions by mutually repelling once they sensed a teammate within a certain safe distance ε, which we set to one robot diameter. Once robots came within this distance, they acted as they were in danger of colliding and used one of the resolution behaviors involving various types of communication schemes. Different types of communication were used to determine the length of time a robot would repel after this point.

We used the Teambots [Balch, 2000] simulator to implement these types of communication within groups of Nomad N150 robots. We measured how many pucks (out of 60) were delivered to the goal region within 9 minutes by groups of 2 – 25 robots using each communication type. We repeated each trial 100 times for statistical significance.

Within the implicit method, no time or energy was lost due to communication. This method assumed that robots were able to autonomously compute their distance from the home base. Once robots detected an impending collision, they used a function of this distance (scalar distance * 5) as the time to repel from its teammate(s). For these robots, the coordination cost was equal to the time or the energy spent in these repulsion behaviors. In our first set of experiments we measured the time spent in these behaviors. 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 capabilities. Our coordination cost involved the amount of fuel used for resolving coordination issues while robots engaged in repulsion behaviors.

Our localized method initiated communication only between the robots within the ε distance. After this event, these group members would exchange information above 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. In our first set of experiments, we measured the time these robots engaged in communication and repulsion behaviors. We assumed that 1/5 of a second was needed for this localized information transfer. In our second group of experiments, we measured our coordination cost as a function of energy spent on coordination. We again calculated this amount as a function of the distance traveled during repulsion activities with a smaller amount on sensors. In addition, we assumed each robot spent 0.3 units of fuel per localized communication exchange.

Our final method, *Latency*, used a centralized server with a database of the location of all robots. 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, only in this case with a centralized server. The server then reported back a repel value based on its relative position to 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. We used the same time and energy usage parameters from the localized experiment to measure the cost of coordination.

3 Analyzing the Cost of Communication

Our results support our claim that the best method of communication does change with domain conditions. Figure 1 contains the results from the time based coordination cost trials. The X-axis represents the group size, and the Y-axis the number of pucks successfully retrieved within each group. The implicit approach worked best in small groups where collisions were less likely. In medium sized groups, the localized approach worked better. As collisions became frequent, the large amount of communication inherent in the centralized method became justified, and this group performed significantly better. In these experiments, we found that a latency time of 1 second yield the highest productivity from various latency variations we studied.

When comparing equally sized groups, the team's productivity strongly negatively correlated to its time and energy coordination costs. On average over groups of 2 – 25 robots, we found a high statistical correlation of -0.96 between these groups' productivity and their total coordination cost. In our equivalent energy experiment, we found a correlation value of -0.95. This supports our hypothesis that coordination costs measure based on time or energy can effectively model the most effective communication method in groups.

Our measure was useful for communication adaptation as well. In our implementation, adaptation was triggered once a robot's autonomously measured coordination cost threshold,

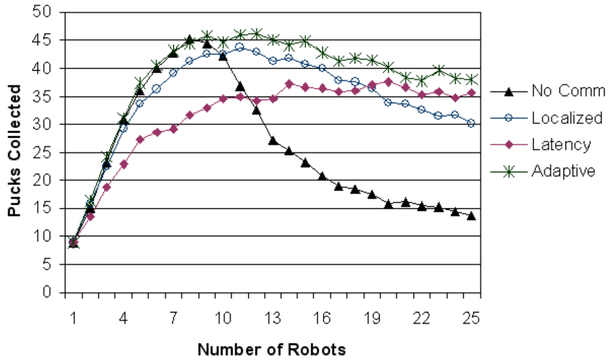

Figure 1: Adaptation Using Coordination Costs

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. Figure 1 shows the results from our time based adaptive method. We conducted the two tailed t-test to confirm the statistical significance of for our adaptive time and energy methods (not shown). In both cases we found p-values well below the 0.05 confidence interval. We believe the success of these approaches was based on the ability of these method to switch their basic communication strategy as the probability of collisions vacillated. This enabled these robots to improve their productivity in a significant fashion regardless of the group size.

4 Conclusion and Future Work

This work demonstrates how coordination costs can model the relative effective of robotic communication methods. Our measure focuses on the time and fuel spent communicating and resolving collisions. It facilitates effective comparison between implicit, localized and centralized communication methods. Using this information we are able to select the most effective communication scheme for a given domain for a group of robots. We are even able to dynamically switch between these methods in a quick, online fashion. For future work, we believe that additional expansions to our measure will facilitate comparison and adaptation even within heterogeneous groups of robots with diverse qualities.

References

[Balch, 2000] Tucker Balch. www.teambots.org, 2000.

[Jager and Nebel, 2001] Markus Jager and Bernhard Nebel. Decentralized collision avoidance, deadlock detection, and deadlock resolution for multiple mobile robots. In *IROS*, pages 1213–1219, 2001.

[Sen *et al.*, 1994] Sandip Sen, Mahendra Sekaran, and John Hale. Learning to coordinate without sharing information. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 426–431, Seattle, WA, 1994.

[Tews, July 2001] A. Tews. Adaptive multi-robot coordination for highly dynamic environments. In *CIMCA 2001*, July 2001.