

## A Comparative Study between Centralized, Market-Based, and Behavioral Multirobot Coordination Approaches

M. Bernardine Dias  
The Robotics Institute  
Carnegie Mellon University  
Pittsburgh, PA 15224  
USA

Anthony Stentz  
The Robotics Institute  
Carnegie Mellon University  
Pittsburgh, PA 15224  
USA

**Abstract** – This paper presents a comparative study between three multirobot coordination schemes that span the spectrum of coordination approaches; a fully centralized approach that can produce optimal solutions, a fully distributed behavioral approach with minimal planned interaction between robots, and a market approach which sits in the middle of the spectrum. Several dimensions for comparison are proposed based on characteristics identified as important to multirobot application domains. Furthermore, simulation results are presented for comparisons along two of the suggested dimension: Number of robots in the team and Heterogeneity of the team. Results spanning different team sizes indicate that the market method compares favorably to the optimal solutions generated by the centralized approach in terms of cost, and compares favorably to the behavioral method in terms of computation time. All three methods are able to improve global cost by accounting for the heterogeneity of the robot team.

**Key words** – Comparison, Multirobot Coordination, Market Approach, and TraderBots.

### I. INTRODUCTION

The growing demand for robotic solutions to increasingly complex and varied problems has dictated that a single robot is no longer the best solution for many application domains; instead, teams of robots must coordinate intelligently for successful task execution. Driven by these demands, many research efforts have focused on the challenge of multirobot coordination. Dias and Stentz [1] present a detailed description of multirobot application domains and their demands, and show that robot teams are more effective than a single robot in many application domains. Simply increasing the number of robots assigned to a task does not necessarily solve a problem more efficiently; multiple robots must cooperate to achieve high efficiency. The difficulty arises in coordinating many robots to perform a complex, global task. Dynamic environments, malfunctioning robots, and multiple user requirements add to the complexity of the multirobot coordination problem. Dias and Stentz [1] explore some of these issues, and evaluate some of the principal efforts in this field of research.

Multirobot coordination mechanisms span a spectrum of approaches ranging from fully centralized approaches to fully distributed approaches. At one end of the spectrum, centralized approaches design the team so that a single robot or central computer acts as a “leader” and is responsible for planning the actions of the entire group. This methodology

usually requires that the group members report most recent status information to the leader who uses this information to coordinate the group. The principal advantage of such centralized approaches is that they allow optimal planning since the decision-making agent has access to all relevant information when planning for the group. However, it is commonly considered that they suffer from several disadvantages including sluggish response to dynamic conditions, intractable solutions for large teams, communication difficulties, and the leader becoming a central point of failure. If the environment is highly dynamic, the leader may not be able to generate plans at a sufficiently fast rate to keep up with changes in status information from the group members. Furthermore, the team members, dependent on the leader for action-decisions, may not be able to avoid dangerous situations in time. Also, the ability to generate optimal solutions for a complex group-coordination problem diminishes drastically with the size of the group. Another important consideration in multirobot systems is communication. A fully centralized approach demands that all group members remain in communication with the leader at all times so that the leader can re-plan whenever a new situation is encountered. Fully centralized approaches also require high-bandwidth communication because the group members have to communicate all their state information to the leader on a regular basis so that the leader can generate informed plans. Finally, in a fully centralized system, the success of the group is tightly coupled to the performance of the leader. Hence, if the leader malfunctions or is disabled, the entire group becomes ineffective. For all of these reasons, fully centralized approaches seem most suited for small groups of robots operating in controlled, static environments with global communication.

On the other end of the spectrum, fully distributed, reactive behavioral approaches address many of the problems inherent to a centralized approach by distributing the planning responsibilities amongst all members of the team. Each robot operates independently, relying on its local sensory information and planning its actions accordingly. Any cooperation between team members are often fortuitous and each robot tends to act as though it is the only member of the team. Many research efforts have modeled such distributed systems inspired by biology, physics, and economics. The principal drawback of these distributed approaches is that they often result in highly sub-

optimal solutions because all plans are based solely on local information and hence, efficient execution of the global goal, by coordinating the resources of the team, isn't prioritized in the local planning. However, to their advantage, fully distributed approaches are often very simple to implement and overcome many of the disadvantages of fully centralized approaches.

Recently, negotiation-based and economy/market-based multirobot coordination has gained popularity. The general concept of these market-based approaches is that the robots act largely independently in terms of planning for themselves, but are able to take into account team resources by trading tasks with team members. Communications are limited to offers and awards of tasks, and bids for tasks, and hence they often consist of low-bandwidth communication. Economic approaches can also maintain the benefits of distributed approaches and opportunistically incorporate pockets of optimal centralized optimization [2]. However, economic approaches are not without their disadvantages. Negotiation protocols, mapping of task domains to appropriate cost functions, and introducing relevant de-commitment penalty schemes can quickly complicate the design of a control-architecture. Furthermore, some negotiation schemes can drastically increase communication requirements. Thus, all of these factors must be considered when designing a market-based architecture.

## II. CONTRIBUTION

The goal of this work is to conduct a meaningful comparison between the three approaches discussed above that span the spectrum of multirobot coordination mechanisms. The three chosen approaches are a fully centralized approach, a fully distributed behavioral approach, and our TraderBots market-based approach. Rabideau et al. [3] conduct a similar comparative study between a centralized planner that does not guarantee the optimal solution, a distributed planner, and a single-round-auction Contract Net approach. They evaluate these coordination schemes for a simulated 3-rover, 20-rock, geological science scenario where rovers sample spectra of rocks on Mars. They conclude that the Contract Net approach performed best but took up most CPU cycles. The work presented in this paper differs in that the market approach used is not limited to a single round of bidding, the chosen approaches span the spectrum of those currently used, and the chosen approaches are evaluated along several different dimensions.

Another recently published comparative study is that of Gerkey and Mataric [3]. This work aims to compare different multirobot coordination schemes in terms of complexity and optimality. However, the proposed framework for comparison makes the limiting assumption that only one task can be assigned to any robot. Thus, the framework limits the performance and hence the analysis of some of the approaches under consideration, where the approaches do not limit task assignment to a single task per robot.

The contribution of this paper is two-fold. First, this paper proposes several dimensions along which the

comparison of multirobot coordination methods is useful for multirobot application domains. Second, this work provides simulation results comparing three approaches spanning the spectrum of multirobot coordination schemes along two of the identified dimensions: Number of robots, and Heterogeneity. Comparisons such as these are extremely useful for determining the strengths and weaknesses of these approaches when choosing a suitable approach for a given application.

## III. COMPARISON DETAILS

The scope of this comparison is limited to evaluating the performance of three approaches that span the spectrum of multirobot coordination mechanisms. Multirobot coordination schemes can be evaluated along the following dimensions identified by Dias and Stentz [1] as important to multirobot applications domains:

1. Scalability with large numbers of robots
2. Handling heterogeneous robot teams
3. Dealing with unknown and changing environment and relying on sensor feedback from robots to gather information about the world
4. Handling new tasks from the operator during task execution
5. Dealing with robot death and partial robot malfunction
6. Dealing with the necessity for robots to generate new tasks during execution depending on sensor-feedback and other dynamic occurrences
7. Dealing with a change in execution-priority and quality requirements for tasks during execution
8. Dealing with limited-range and limited-bandwidth communication between robots

Two more comparative dimensions that could be useful for emerging multirobot applications are:

9. Handling cooperative tasks that require at least two robots to complete a given task
10. Handling tightly coordinated task requirements

This paper reports the first results of an evolving comparison between three coordination mechanisms that span the spectrum of available approaches. On one end of the spectrum is the fully centralized approach where a single robot plans for the entire group. The other end of the spectrum is characterized by a fully behavioral approach where each robot completely handles its own planning and has very limited interaction with other agents. Finally, the market-based approach sits in the middle of the spectrum, where the robots are independent and able to plan for themselves, but are also able to interact with other teammates via trading tasks. This paper reports comparison results for the first two dimensions in the above list, where the evaluation metrics are the number of robots and team heterogeneity. Ongoing work continues the comparison across all of the dimensions listed above.

#### IV. SCENARIO AND IMPLEMENTATION DETAILS

The chosen application is that of a distributed sensing problem where robots are tasked with gathering sensory information from various designated locations of interest. This translates into a version of the multi-robot traveling salesman problem (MTSP) with paths instead of tours (i.e. without the requirement that robots need to return to their starting locations). Hence, the tasks can be considered as cities to be visited where the costs are computed as the time taken to traverse the Euclidian distance between cities. A task is completed when a robot arrives at a city. The global task is completed when all cities are visited by at least one robot. The global cost is computed as the summation of the individual robot costs. Thus, the goal is to complete the global task while minimizing the number of robot-hours consumed.

Each robot travels by taking a fixed step (based on its speed) each control-loop-cycle along the straight line connecting its current position to the next city on its task queue. When the robots are not executing tasks, they remain stationary at their current locations. Implementation details for each of the approaches are detailed next. The key challenge in the implementation is to allow each approach to maximize its benefits and minimize its disadvantages while maintaining the fundamental philosophy of the approach and while keeping the playing field level in order to make a useful comparison. There are numerous ways to approach such a comparison, and no matter which perspective is chosen, there will be disadvantages. Hence, this work attempts to choose the least contentious methodology for the comparison. Thus, the team of robots is kept the same across approaches. That is, the comparison is based on the common scenario where a team of robots is available for executing a task and the comparison is based on coordinating the team in different ways to successfully execute the task. In order to keep the playing field level, the algorithms used for the centralized approach are also used for individual tour queue optimization in the market and behavioral approaches. This is explained in more detail in the following section.

##### A. Centralized

The centralized approach is implemented with one robot in the team acting as the leader who does all the planning for the group. The centralized planner is implemented using an exhaustive depth-first search (DFS) with some pruning, which produces an optimal solution for the task allocation problem. The robots transmit state information to the leader, the leader generates a plan based on this information, and the team executes the instructions of the leader to the best of its ability. In the centralized approach all robots do not start execution until the leader completes planning for the group. Thus, it is assumed that the leader alone does all the planning-related computation.

##### B. Behavioral

The behavioral approach is implemented with each robot planning and executing all its actions independently. Each robot uses the same DFS optimization algorithm used in the centralized approach to plan the optimal scheduling of its own tasks. The robots all plan to execute all tasks since no team coordination is planned. However, to level the playing field somewhat, as each robot completes a task, it announces the completion of that task to all the other robots. When a robot learns that a task has been completed, it removes that task from its queue and re-schedules its task-queue accordingly. In the behavioral approach, each robot stops execution until its own planning/re-planning is complete and then resumes execution of any remaining tasks.

##### C. Market

The market approach is implemented with an OpTrader agent (the OpTrader is assumed to reside on one of the robots) that is responsible for translating and trading on behalf of the operator. The robots can all plan for themselves using the same algorithm as in the behavioral approach. The robots can also trade their tasks with other robots during execution. In this implementation, only single-task trades were allowed and hence the market solution can get stuck in local minima. In the market approach, execution is not started until initial negotiations with, and allocations by the OpTrader are completed. Thereafter the robots keep trading tasks during execution as the opportunities for profitable trades arise. Each robot has a specified frequency at which it trades. In this implementation, all robots traded with the same frequency (once every 5 steps), but at staggered intervals. The OpTrader used a greedy algorithm for auction clearing, where the  $n$  best allocations were made each round, one award per bidder, where  $n$  is the number of robots in the team. Any tasks not auctioned off in a round were re-announced and cleared during subsequent rounds of auctioning. During the inter-robot trading, each robot announced all of its current tasks for auction, and accepted the single most profitable bid it received.

#### V. RESULTS AND DISCUSSION

All experiments used a set of randomly generated tasks in a 2000m by 2000m obstacle-free world. The robots start off in random positions within this world. The operator announces all tasks to the group at the beginning. In all the approaches, the robots stop execution during their planning stage.

##### A. Scalability With Number Of Robots

The goal in this comparison is to study the performance of the three approaches for different team sizes. It is assumed that one of the robots acts as the leader in the centralized approach and as the OpTrader in the market approach. All robots are otherwise assumed to be homogeneous.

Figure 1 shows a plot of the global cost versus the team size for a single run each for different team sizes ranging from 2 to 10. The plot shows data gathered for the same randomly generated set of 10 tasks. As expected the optimal solution cost is achieved by the centralized approach. The behavioral method produces fairly sub-optimal solutions as predicted, and the market method sits in between. (Note that this comparison was only run for one set of randomly generated task distributions because of the high time consumption of the centralized method for larger team sizes.)




Figure 1: Cost Comparison For A Single Run




Figure 2: Total Time Comparison For Single Run

However, if we observe the total time taken, including computing time, then the centralized approach exceeds the cost of the solution generated by the market approach for team sizes exceeding 5, as seen Figure 2. The behavioral method still exceeds the cost of the other two approaches.

The same experiments were then repeated averaged over 100 runs. Here too we see a similar trend when looking at the solution cost (see Figure 3) and the total time taken including the computation time (see Figure 4).

Note once again that we were only able to generate results for up to 5-robot teams for the centralized approach due to time constraints. Hence we were unable to observe how close the total time taken by the centralized and market methods compare as the number of robots grows. However,

based on time performance for a smaller number of runs, it is highly probable that the market approach matches if not exceeds the performance of the centralized approach for larger numbers of robots.




Figure 3: Solution Costs Averaged Over 100 Runs




Figure 4: Total Time Averaged Over 100 Runs

Another interesting aspect to notice is that the global cost increases with the number of robots for the behavioral approach. This is because there are more robots moving towards the same goals until the robot that gets there first eliminates them from the task queues. Interestingly though, if the average execution time per robot is monitored (see Figure 5) it is clear that this reduces with the number of robots for all three approaches. Thus, it is clear that all three approaches are tending towards reducing the number of robot-hours spent.

Thus far, it appears that the centralized and market approaches perform best. However, one benefit of the behavioral approach is that it requires very low computation time. Therefore, it is informative to observe the average computation time per robot for the presented results.

Figure 6 shows the variation of the computation time for the different approaches plotted against the number of robots. This figure clearly shows the centralized method at a disadvantage compared to the other two methods. Although the difference is small in comparison to that with the centralized approach, the behavioral approach on average takes slightly less time than the market approach to

generate solutions. This makes sense because in addition to computing the optimal arrangement of its tasks, the robots in the market approach spend some computation time on announcing, clearing, and submitting bids to task auctions.




Figure 5: Average Execution Time Per Robot




Figure 6: Average Computation Time Per Robot

Thus, overall, the market method can provide solutions with significantly better global costs than the behavioral approach using computation time significantly less than the centralized approach. Note however that for applications where the optimal solution must be generated, the centralized approach will suit best. Also, for applications where optimality of the solution is of no consequence, the behavioral method is fast and simple and hence becomes the approach of choice. Note that the market approach can generate better solutions if allowed to trade clusters of tasks [2]. However this improvement will increase computation time somewhat.

### B. Handling Heterogeneous Robot Teams

The goal in this comparison is to study the performance of the three approaches when the group of robots is heterogeneous. The chosen aspect of heterogeneity is robot

speed. Since the optimized quantity is the number of robot-hours, the speed of each robot will affect the quality of the solution.

Figure 7 shows the variation in total time for solutions produced by the three approaches for teams of 4 robots with varying degrees of heterogeneity where robots can have speeds of either 1m/s or 5m/s. On the far left all 4 robots have the minimum speed of 1m/s and on the far right all 4 robots have the maximum speed of 5 m/s. The 3 points in the middle represent the cases of mixed speeds where the points from left to right indicate scenarios where 1, 2, and 3 robots have a speed of 5m/s respectively, while the others have speeds of 1m/s.




Figure 7: Heterogeneous Teams

	5 Runs	100 Runs
BEH	3451	3049
MKT	1834	1545
CNT	1206	1129

Table 1: Results For Heterogeneous Teams With Speeds Randomly Allocated Between The Minimum Speed Of 1m/S And The Maximum Speed Of 5m/S

Table 1 shows a summary of results where each robot is assigned a random speed between the max and min. All three approaches are able to improve cost by taking into account heterogeneity. The performance ranking of the three approaches remains the same as in the previous section.

## VI. CONCLUSION AND FUTURE WORK

This paper presents a comparative study between three multirobot coordination schemes that span the spectrum of coordination approaches; a fully centralized approach that can produce optimal solutions, a fully distributed behavioral approach with minimal planned interaction between robots, and a market approach which sits in the middle of the spectrum. Several dimensions for comparison are proposed based on identified characteristics important to multirobot application domains. Furthermore, simulation results are

presented for comparisons along two of the suggested dimension: Number of robots in the team and Heterogeneity of the team. Results spanning different team sizes indicate that the market method compares favorably to the optimal solutions generated by the centralized approach in terms of cost, and compares favorably to the behavioral method in terms of computation time. All three methods are able to improve global cost by accounting for the heterogeneity of the robot team. Ongoing work continues the comparative study across all of the proposed dimensions. The study is aimed at providing useful information for determining the strengths and weaknesses of these approaches when choosing a suitable approach for a given multirobot application. Future work includes completing comparison along remaining dimensions some of which is completed but was not possible to include in this paper.

#### VII. ACKNOWLEDGEMENTS

This research was sponsored in part by NASA, under contract "Heterogeneous Multi-Rover Coordination for Planetary Exploration" (contract number NCC2-1243). The views and conclusions contained in this document are those of the authors and should not be interpreted as

representing the official policies or endorsements, either expressed or implied, of the U.S. Government.

#### VIII. REFERENCES

- [1] M. B. Dias and A. Stentz, "A Market Approach to Multirobot Coordination", *Technical Report, CMU-RI - TR-01-26, Robotics Institute, Carnegie Mellon University*, August 2001.
- [2] M. B. Dias and A. Stentz, "Opportunistic Optimization for Market-Based Multirobot Control", in *Proceedings of 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- [3] B. P. Gerkey and M. J. Mataric, "Multi-Robot Task Allocation: Analyzing the Complexity and Optimality of Key Architectures", in *Proceedings of 2003 IEEE International Conference on Robotics and Automation*.
- [4] G. Rabideau, T. Estlin, S. Chien, and A. Barrett, "A Comparison of Coordinated Planning Methods for Cooperating Rovers", in *Proceedings of AIAA 1999 Space Technology Conference, Albuquerque, NM*.