

Enhanced Negotiation and Opportunistic Optimization for Market-Based Multirobot Coordination

M. Bernardine Dias
Anthony (Tony) Stentz

CMU-RI -TR-02-18

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

August 2002

© 2002 Carnegie Mellon University

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 Carnegie Mellon University.

Abstract

Multirobot coordination, if made efficient and robust, promises high impact on automation. The challenge is to enable robots to work together in an intelligent manner to execute a global task. The market approach has had considerable success in the multirobot coordination domain. However the implementation of this approach to date restricts the negotiations to two-party, single-task deals which often forces the task allocation solution into a local minimum. This report investigates the effects of introducing multi-party and multi-task negotiations to enhance the market-based approach to multirobot coordination. Multi-party negotiations are enabled by implementing a combinatorial exchange mechanism, while multi-task negotiations are accomplished via clustering of tasks in cost space. Presented results show that global costs can be considerably reduced (on average to within 10% of the optimal solution for the tested scenarios), and hence task allocation can be considerably improved, by enhancing the negotiation capabilities of the robots.

This report also investigates the effects of introducing opportunistic optimization with leaders to enhance market-based multirobot coordination. Leaders are able to optimize within subgroups of robots by collecting information about their tasks and status, and re-allocating the tasks within the subgroup in a more profitable manner. The presented work also considers the effects of introducing pockets of centralized optimization into an otherwise distributed system. The implementations were tested on a variation of the traveling salesman problem. Presented results show that global costs can be reduced, and hence, task allocation can be improved, utilizing leaders. Note the presented work only addresses scenarios where leaders run exchanges to optimize task allocation within a group of robots. Some leaders are also capable of clustering tasks and hence can conduct combinatorial exchanges. But these are not the only opportunities for leaders to optimize within the market. It is also possible to have combinatorial exchanges and leaders as distinct entities within the economy. Leaders could also use other approaches to generate plans for a subgroup of robots. Finally, a leader could simply act as a means of enabling trade between subgroups of robots who are otherwise unable to communicate, thus enriching the set of possible trades. Thus, leaders can enhance the market-based approach by several means including optimizing task-allocation, generating plans, optimizing plans, and enabling better trade opportunities between groups of traders.

Table of Contents

1. <i>Introduction</i>	9
2. <i>The Market Approach</i>	10
2.1. Revenues, Costs, the Role of Price and the Bidding Process	10
2.2. Cooperation, Competition, Learning and Adaptation	10
2.3. Self Organization	10
2.4. Limitations Of Prior Work	11
3. <i>Contribution</i>	11
4. <i>Exploring the Space of Negotiations</i>	11
4.1 Clustering for Multi-Task Processing	11
4.2 Combinatorial Exchange for Multi-Party Optimizations	12
5. <i>Optimizing with Leaders</i>	12
6. <i>Experimentation</i>	13
6.1 Two-Party, Single-Task (TPST) Negotiations	13
6.2 Two-Party, Multi-Task (TPMT) Negotiations	13
6.3 Leaders Performing Multi-Party Single-Task (MPST) Optimizations	14
6.4 Leaders Performing Multi-Party, Multi-Task (MPMT) Optimizations	14
6.5 Multiple Competing Local Groups	14
7. <i>Results and Discussion</i>	14
8. <i>Conclusions and Future Work</i>	17
9. <i>References</i>	18

Table of Figures

<i>Figure 1: 4 robots and 10 tasks for a single run</i>	14
<i>Figure 2: Solution for TSP with 3 overlapping subgroups of 4 robots each and 10 tasks</i>	16
<i>Figure 3: Cost reduction for 8 robots and 30 tasks with 3 sub groups of 4 robots each</i>	17

Table of Tables

<i>Table 1: Performance averaged over 100 randomly generated 2-robot, 10-task TSPs</i>	<i>15</i>
<i>Table 2: Results averaged over 100 randomly generated 4-robot (heterogeneous), 10-task TSPs</i>	<i>15</i>
<i>Table 3: Performance averaged over 100 randomly generated 4-robot (heterogeneous), 20-task TSPs</i>	<i>15</i>
<i>Table 4: Performance averaged over 100 randomly generated 8-robot (heterogeneous), 10-task TSPs with 3 overlapping groups of 4 robots each</i>	<i>16</i>

1. 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 on technology, many research efforts have focused on the challenge of multirobot control.

Multirobot solutions are paramount for several reasons. A single robot cannot perform some tasks alone. In task-decomposable application domains, robot teams can accomplish a given task more quickly than a single robot can by dividing the task into sub-tasks and executing them concurrently. Moreover, a team can make effective use of specialists, can localize itself more efficiently by sharing information, and can produce a wider variety of solutions, and thereby respond opportunistically to dynamic conditions in more creative and efficient ways. Generally, a team of robots provides a more robust solution by introducing redundancy, and by eliminating any single point of failure as long as there is overlap between the robots' capabilities. Thus, for many applications, robot teams are more effective than a single robot. Dias and Stentz [5] present a detailed description of multirobot application domains and their demands on multirobot coordination schemes, and show that robot teams are more effective than a single robot in many of these domains.

Simply increasing the number of robots assigned to a task does not necessarily solve a problem more efficiently; multiple robots must cooperate intelligently to achieve efficiency. The difficulty arises in coordinating many robots to collectively 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 [5] explore some of these complexities, and present some of the principal efforts in this field of research.

One approach is to design the team such that a single robot or central computer acts as a "leader" and is responsible for planning the actions of the entire group. An example of the centralized approach is work done by Brummit and Stentz [1]. The principal advantage of such centralized approaches is that they allow optimal planning. However, they often suffer from several disadvantages including sluggish response to dynamic conditions, intractable solutions for large teams, communication difficulties, and single points of failure.

Local and distributed approaches address these problems by distributing the planning responsibilities amongst all members of the team. Each robot operates independently, relying on its local sensor information. Many research efforts have modeled distributed systems inspired by biology [1], physics [3], and economics [17]. The principal drawback of distributed approaches is that they often result in highly sub-optimal solutions because all plans are based solely on local information with minimal inter-robot coordination.

Recently, negotiation-based and economy/market-based multirobot coordination has gained popularity. This work in multirobot coordination draws from the software agents literature that began with Smith's Contract Net Protocol [14] and has since been extended to control a variety of multiagent (and more recently multirobot) systems. Golfarelli and Rizzi [9] proposed a swap-based negotiation protocol for multirobot coordination that restricted negotiations to task-swaps. Rabideau et al. [11] utilize a Contract Net approach to include local rover estimates for path information into their centralized planning approach to solving a Multi-Traveling Salesman Problem for a scenario where multiple rovers are used to sample spectra of rocks on Mars. Stentz and Dias [15] introduced the use of a sophisticated market mechanism to coordinate a team of robots performing a group task, building on work by Smith [14] and Wellman and Wurman [17]. The approach proposes opportunistically injecting pockets of centralized optimal planning into a distributed system, thereby exploiting the desirable properties of both distributed and centralized approaches. Thayer et al. [16], Gerkey and Mataric [7], and Zlot et al. [18] have since produced market-based multirobot coordination results. Economic approaches are not without their disadvantages. Negotiation protocols, mapping of task domains to appropriate cost and revenue functions, and introducing relevant de-commitment penalty schemes can quickly complicate the design of a coordination-architecture. Furthermore, some negotiation schemes can drastically increase communication requirements.

2. The Market Approach

Stentz and Dias [15] first introduced the concept of using a sophisticated market approach to coordinate multiple robots to cooperatively complete a task, building on the Contract Net protocol by Smith [14] and the work on market-aware agents by Wellman and Wurman [17]. This work introduced the methodology of market mechanisms for intra-team robot coordination (i.e., non-competitive environments) as opposed to competitive multirobot domains and competitive inter-agent interactions in domains such as E-commerce. Simulation results using this approach were produced by Dias and Stentz [6], and proven robot results were presented by Thayer et al. [16], and Zlot et al. [18]. A brief introduction to this approach is presented here.

Consider a team of robots assembled to perform a particular set of tasks. Consider further, that each robot in the team is modeled as a self-interested agent, and the team of robots as an economy. The goal of the team is to complete the tasks successfully while minimizing overall costs. Each robot aims to maximize its individual profit (which often translates to minimizing individual cost where possible); however, since all revenue is derived from satisfying team objectives, the robots' self-interest equates to doing global good. Moreover, each robot can only increase its profit by eliminating unnecessary waste (i.e. cost). Hence, if the global cost is determined by the summation of individual robot costs, each deal made by a robot (note that robots will only make profitable deals) will result in global cost reduction.

2.1. Revenues, Costs, the Role of Price and the Bidding Process

Appropriate functions are needed to map possible task outcomes onto revenue values and to map possible schemes for performing the task onto cost values. As a team, the goal is to execute some plan such that the overall profit, the excess of revenue over cost, is maximized. Furthermore, these functions must provide a means for distributing the revenue and assessed costs among the individual robots.

Thus, robots receive revenue and incur costs for accomplishing a specific team task, but the team's revenue function is not the only source of income. A robot can also receive revenue from another robot in exchange for goods or services. The *price* dictates the payment amount for the good or service. A common approach is to *bid* for a good or service in order to arrive at a mutually acceptable price.

2.2. Cooperation, Competition, Learning and Adaptation

Two robots are *cooperative* if they have complementary roles, that is, if both robots can make more profit by working together than by working individually. Conversely, two robots are *competitive* if they have the same role; that is, if the amount of profit that one robot can make is negatively affected by the presence of the other robot. The flexibility of the market-model allows the robots to cooperate and compete as necessary to accomplish a task.

Moreover, the robot economy is amenable to learning new behaviors and strategies as it executes its complex global task. One of the greatest strengths of the market economy is its ability to deal opportunistically with dynamic environments.

2.3. Self Organization

Conspicuously absent from the market is a rigid, top-down hierarchy. Instead, the robots organize themselves in a way that is mutually beneficial. Since the aggregate profit amassed by the individuals is directly tied to the success of the task, this self-organization yields the best results.

Consider a group of ten robots. An eleventh robot, **A**, offers its services as their leader. It does not become their leader by coercion or decree, but by convincing the group that they will make more money by following its advice than by acting individually or in subgroups. **A** does this by investigating "plans" for utilizing all ten robots. If **A** comes up with a truly good plan, it will maximize profit across the whole group. The prospective leader (**A**) can use this large profit to bid for the services of the group members, and of course, retain a portion of the profit for itself. Note that all relevant robots will have to commit to the plan before it can be sold. The leader may be bidding not only against the individuals' plans, but also against group plans produced by other prospective leaders. Note that the leader acts both as a benevolent and a self-interested agent—it receives personal compensation for efforts benefiting the entire group.

But there is a limit to this organization. As the group becomes larger, the combinatorics become intractable and the process of gathering all of the relevant information to produce a good plan becomes increasingly difficult. A leader will realize this when it can no longer convince its subjects (via bidding for their services) to follow its plans.

2.4. Limitations Of Prior Work

One of the key limitations of the implementation of this approach to date is the restriction of negotiations to two-party, single-task deals. In many cases, this restriction limits the global cost reduction, since the robots do not have the negotiation tools to reason their way out of shallow, local cost minima. The work presented here extends these tools to permit multi-party and multi-task deals with better global cost reduction potential. Work to date also does not explicitly deal with robots with heterogeneous capabilities in bid formulation. The presented work explores this heterogeneity and the possibility of enhancing such systems by using leaders.

3. Contribution

The work presented in this paper explores the effects of enriching the negotiation capabilities of the robots such that multi-party and multi-task deals are possible, and also the effects of enabling opportunistic optimization with leaders in market-based multirobot coordination.

4. Exploring the Space of Negotiations

The work presented here enhances the negotiation capabilities of the robots in the economy by introducing clustering to enable multi-task deals, and a combinatorial exchange to handle multi-party deals.

4.1 Clustering for Multi-Task Processing

The capability to negotiate multi-task deals greatly enhances the market approach because it allows a robot to escape some local minima in task allocation solutions. However, if the robots bid on every possible combination of tasks, the number of bids submitted will grow exponentially with the number of tasks. Consequently, processing these bids will be impossible for more than a few tasks. Hence, some form of clustering algorithm is necessary to determine the clusters of tasks to bid on. The possibilities for such clustering algorithms are numerous [10].

Golfarelli and Rizzi [8] presented a clustering algorithm for their task-swap negotiation protocol in which tasks are clustered over spatial and temporal dimensions. The clusters are made by initially forming single-task clusters, and then for each of these clusters, adding all neighbors within a threshold to the cluster. Finally, all clusters fully contained within other clusters are removed. This clustering algorithm has an appealing multi-dimensional capability. However, it subsumes smaller clusters within larger ones and does not guarantee that clusters will be built across the entire task space. A spanning range of cluster sizes (i.e. clusters ranging in size from single-task clusters to a wholly-inclusive cluster) and task membership is important, because a robot cannot necessarily predict the interaction of the clusters it offers with the tasks of other bidders. The chosen clustering algorithm preserves these properties — it operates as follows:

1. Create a list of edges spanning all tasks on offer (N), where each edge joins two tasks and the cost of the edge represents the distance in cost space between the two tasks. A low edge value implies, but does not guarantee, that two tasks can be performed more cost effectively together than apart.
2. Sort the edge list from lowest to highest cost.
3. Form the first group of clusters by creating a single-task cluster for each task on offer.
4. For cluster sizes ranging from 2 to N , recursively form new clusters by adding the next best available edge (an edge is unavailable if it is either already included in a previous cluster or if the edge connects two tasks which are not included in any of the previous clusters) to a cluster in the previous cluster list. (Note, when new clusters are formed, all previous clusters are preserved). Thus, recursively form a forest of minimum spanning trees (MSTs) [1] ranging in size from 1 to N .

This algorithm can be applied to determine cost-effective clusters of tasks, without computing every possible cluster. Suitable variations of this algorithm (or others) can be chosen to enable multi-task

negotiations in different task domains. The presented work is verified on a multi-depot traveling salesman problem (TSP), and hence, the MSTs are decomposed into sub-tours as follows. If a newly added edge breaks the continuity of the sub-tour, the MST is adjusted by removing one of the edges connecting the tour to the newly added edge and adding the necessary edge to preserve the continuity of the tour with the least addition to the cost of the tour. Note that this change still preserves the bounds of the MST, which guarantees that the cost of the tour does not exceed twice the optimal cost, in metric cost spaces where the triangle inequality is preserved.

4.2 Combinatorial Exchange for Multi-Party Optimizations

A combinatorial exchange (a market where bidders can jointly buy and sell a combination of goods and services within a single bid) is chosen to enable multi-party optimizations for a team. Allowing robots to offload an owned cluster when bidding to accept a new cluster of tasks further enhances the bidding capability of the robots. A combinatorial exchange also enables a leader to better optimize the task assignments of a subgroup of robots and to potentially achieve a greater global cost reduction. Many researchers including Sandholm and Suri [13] have presented valuable insight on how to efficiently implement and clear combinatorial exchanges for E-commerce applications. However, many of these tools are relatively complex and are not used in this work for simplicity. Instead, Sandholm's [12] basic recommendation of searching a binary bid tree is applied. The chosen implementation for clearing the combinatorial exchange in this work is a depth first search on a binary tree where each node of the tree represents a bid and the binary aspect of the tree represents accepting or rejecting that bid. The tree is pruned to disallow accepting multiple bids from any single bidder, and to disallow exchanging of any single task more than once. Note that the pruning does not affect the solution except by improving the runtime.

5. Optimizing with Leaders

An important contribution of this work is the preliminary investigation of a "leader" role that allows a robot with the necessary resources to assess the current plans of a group of robots and to provide more optimal plans for the group. Note that the "better plan" sometimes is simply a more cost-effective assignment of execution responsibility for different parts of the existing plan. The leader can gain knowledge of the groups' current situation through communication or some form of observation. A prospective leader can use the profits generated by an optimized plan to bid for the services of the group members, and retain a portion of the profit for itself. The leader may bid not only against the individuals' plans, but also against group plans produced by other prospective leaders. Centralized and distributed approaches are two extremes along a continuum. The introduction of leaders allows the market-based approach to slide along this continuum in the direction of improved profitability in an opportunistic manner.

The leader role in the market approach is implemented as follows. A leader queries surrounding robots to discover what tasks they have to offer and their current states, and re-allocates tasks within the group using the combinatorial exchange mechanism described above. The presented work only addresses scenarios where leaders run exchanges to improve task allocation within a group of robots. Some leaders are also capable of clustering tasks and hence can conduct combinatorial exchanges. It is also possible to have combinatorial exchanges and leaders as distinct entities within the economy. For example, there could be a leader that simply clusters tasks and sells these *cluster plans* to a combinatorial exchange. Note that the leader is not selling the *actual* cluster of tasks—just a plan for which tasks to cluster. The exchange could then buy all of the component tasks, sell off the resultant cluster, and pay a fee to the leader. The presented results indicate that the benefit from the ability to cluster tasks and participate in multi-task negotiations exceeds the benefit from the ability to perform multi-party negotiations. Leaders could also use other approaches to generate plans for a subgroup of robots. Finally, a leader could simply act as a means of enabling trade between subgroups of robots who are otherwise unable to communicate, thus enriching the possible trades. Thus, leaders can enhance the market-based approach by several means including optimizing task-allocation, generating plans, optimizing plans, and enabling better trade opportunities between groups of traders.

6. Experimentation

The proposed multi-task and multi-party enhancements are developed and tested in a simulated distributed sensing task. A group of robots, located at different starting positions in a known simulated world, are assigned the task of visiting a set of pre-selected observation points. This problem is a variation of the multi-depot traveling salesman problem, where the observation points are the cities to visit. The costs are the lengths of the straight-line paths between locations, interpreted as money. Let c_{ij} be the cost for the j^{th} robot to visit the i^{th} city from the $(i-1)^{\text{th}}$ city in its tour (where the 0^{th} city is the starting location).

The robot cost function for the j^{th} robot is computed as follows:

$$\mathbf{rcost}(\mathbf{j}) = \sum_{i=1}^{n_j} c_{ij}$$

Where n_j is the number of cities in the tour for robot j .

The team cost function is:

$$\mathbf{tcost} = \sum_{j=1}^m \mathbf{rcost}(\mathbf{j})$$

Where m is the number of robots in the team.

The team revenue and robot revenue functions are determined by the negotiated prices. Thus, the robots are paid a fixed amount for visiting each city and can negotiate prices for subcontractors. All robots (bidders) adopt the same simplistic strategy of bidding a fixed 10% markup above the cost of completing the task. According to this strategy, if an announced task costs \mathbf{c} to execute, a robot computes its bid \mathbf{b} as $\mathbf{1.1c}$, indicating it will take on the task for no less than $\mathbf{1.1c}$. Thus, the robots bid for each city based on their estimated costs to visit that city. Similarly, if a robot offered up a task and bid to buy the services of another robot to complete that task, the bid price \mathbf{b} is set as $\mathbf{0.9c}$. Thus, the robot announces that it is not willing to pay more than $\mathbf{0.9c}$ to offload the task(s).

Tasks and robot positions are randomly generated within a 100x100 world, and initial task allocations are made by randomly distributing the tasks among the robots. Heterogeneous robot capabilities are simulated by considering scenarios where some robots can only process single-task (ST) deals (that is, they lack the capability to compute and reason about clusters of tasks), while other robots can process multi-task (MT) deals. Robots capable of playing leader roles possess the additional capability of performing multi-party optimizations via either a single-goods exchange or a combinatorial exchange, depending on their capability. Sections 6.1 through 6.4 further describe the scenarios of robots negotiating in the absence of a leader and the optimization scenarios with leaders. Section 6.5 describes the scenario where robots have limited communication range and hence can only trade within subgroups.

6.1 Two-Party, Single-Task (TPST) Negotiations

Once the initial random task assignments are made, each of the robots, *in turn*, offers all its assigned tasks to all the other robots. Thus, interactions are limited to two parties at any given time. Each bidder then submits a bid for each task. In order to estimate the additional cost of inserting a task into its queue, the bidder uses the cluster generation algorithm described above to generate an MST with its current queue of tasks plus the offered task, and computes the cost difference between the resulting and original queues. The offerer accepts the most profitable bid it receives. The cost of the offerer's resulting queue is computed by removing from its queue the task that was transferred through the winning bid, clustering the remaining tasks using the clustering algorithm, and computing the cost of the resulting queue. Hence, in the TPST scenario, only single-task (ST) deals are considered, and pairs of robots continue to negotiate amongst themselves in round-robin fashion until no new, mutually profitable deals are possible. Therefore, negotiations cease once the system settles into a local minimum for the global cost function.

6.2 Two-Party, Multi-Task (TPMT) Negotiations

In this case, the previous case is repeated with clusters of tasks being the atomic unit of the negotiations. That is, the initial assignments are followed by each of the robots, *in turn*, offering all of its assigned tasks to all the other robots. The robots then bid for clusters of these tasks. Once again, costs are

computed by using the clustering algorithm to cluster all tasks under consideration and compute the cost of the resulting queues. Negotiations are always restricted to occur between two robots in this scenario too.

6.3 Leaders Performing Multi-Party Single-Task (MPST) Optimizations

A leader whose capability is restricted to dealing in single-task deals, is introduced in this case. The leader queries all the robots, and gathers all the tasks of all the robots along with each robot’s state information. The leader then sets up an exchange by formulating single-task bids for the robots in the sub-group based on the gathered information. The exchange used in the MPST scenario is a single-task exchange (i.e. a single bid can contain buying of a single task and selling of another single task). The exchange is then cleared to maximize the leader’s profit. This process is repeated until the exchange cannot produce any further profit, and the corresponding task re-allocation is proposed to the sub-group of robots. If the leader’s plan reduces the global cost, the resulting excess profit can be distributed among the entire subgroup (including the leader) such that the robots in the subgroup accept the leader’s task re-allocation.

6.4 Leaders Performing Multi-Party, Multi-Task (MPMT) Optimizations

Here, the previous case was repeated with the added capability of the leader to process MT bids. That is, the leader sets up and clears a combinatorial exchange to determine the re-allocation of tasks. In a combinatorial exchange, clusters of tasks can be bought and sold within a single bid. Note further that clusters can be bought, re-grouped, and sold as different clusters.

6.5 Multiple Competing Local Groups

This set of experiments involves 8 robots divided into 3 non-disjoint groups of 4 robots each (with the middle group overlapping the other two groups) and 10 tasks. Trading and optimization with leaders are restricted to within the subgroups. This scenario simulates robots with limited range in communication – that is, the robots can only communicate with other robots within their limited communication range. The robots are evenly spread throughout a 2000x2000 world and the cities (tasks) are randomly generated. Scenarios with and without leaders, and with ST-capable and MT-capable robots are considered.

7. Results and Discussion

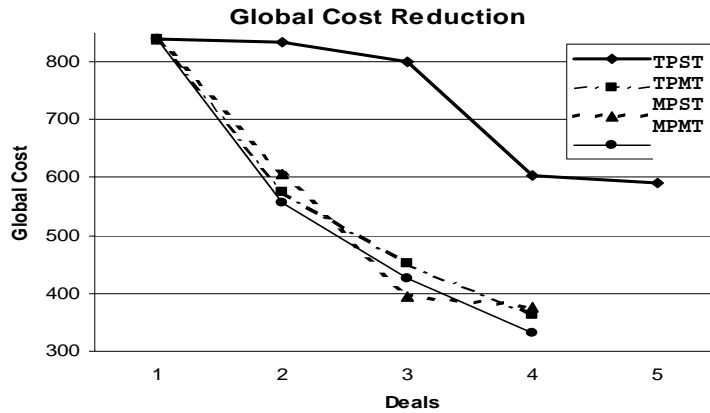


Figure 1: 4 robots and 10 tasks for a single run

The plot in Figure 1 shows the reduction in global cost with each deal for the four cases, TPST, TPMT, MPST, and MPMT, described above, for a single run. While a single case does not provide statistical information, this figure illustrates the point that global cost is reduced with each deal made.

In Table 1 below, the first row shows the global cost based on an initial random allocation of tasks between the 2 robots, and the “Opt. Error” column indicates the percentage increase from the global cost of the optimal task allocation. Note that on average (averaged over 100 runs) the random task allocation resulted

in a global cost that was 65.6% higher than that of the optimal task allocation. Results of the two cases with no leader (that is, the TPST and TPMT cases) are shown in the next two rows, followed by the two cases with leaders (that is, the MPST and MPMT cases), and the final row shows the results for the optimal task allocation. The “Improved” column in all of these cases shows the percentage improvement in global cost compared to the global cost incurred via the initial random task allocation. The “Itns” column indicates the number of iterations required for the system to reach a solution – i.e. how many rounds of bidding were necessary before a final solution was reached.

	Cost	Itns	Improved	Opt. Error
Random	351	-	0.0%	65.6%
No Leader (TP)				
2 TPST	256	2	25.9%	21.4%
2 TPMT	231	1	33.0%	9.0%
MPST (ST Leader)	245	2	29.0%	16.2%
MPMT (MT Leader)	227	1	34.4%	7.0%
Optimal	212	-	38.6%	0.0%

Table 1: Performance averaged over 100 randomly generated 2-robot, 10-task TSPs

	Cost	Itns	Improved	Opt. Error
Random	411	-	0.0%	124.6%
No Leader (TP)				
4 TPST	230	5	42.7%	27.7%
2 TPST + 2 TPMT	222	5	44.6%	23.3%
1 TPST + 3 TPMT	209	4	47.8%	16.2%
4 TPMT	197	4	50.90%	9.7%
MPST (ST Leader)	218	3	45.8%	21.1%
MPMT (MT Leader)	193	2	51.8%	7.5%
Optimal	183	-	-	0.0%

Table 2: Results averaged over 100 randomly generated 4-robot (heterogeneous), 10-task TSPs

In Table 2 and Table 3, the no-leader cases are broken into four different cases where the level of heterogeneity of the robots in the group is altered. The first of these cases is where all four robots can only make ST deals. The next three cases add an increasing number of robots with the ability to reason about MT deals. Note that at least two robots need to be able to handle MT deals in order for an MT deal to take place. Hence, the 3 TPST + 1 MPST case is equivalent to the 4 TPST case, and hence has not been shown.

As evident in Table 1, Table 2, and Table 3, on average, an MT-capable leader can improve the profit of the group significantly. An ST-capable leader can only improve the profit of the group on average for groups of robots where there are at most 50% MT-capable robots. This observation is consistent for the 2-robot and 4-robot cases, and also for the 10-task and 20-task cases.

	Cost	Itns	Improved
Random	725	-	0.0%
No Leader (TP)			
4 TPST	400	10	44.1%
2 TPST + 2 TPMT	388	9	45.7%
1 TPST + 3 TPMT	359	7	49.8%
4 TPMT	336	5	53.0%
MPST (ST Leader)	373	6	47.7%
MPMT (MT Leader)	322	3	54.9%

Table 3: Performance averaged over 100 randomly generated 4-robot (heterogeneous), 20-task TSPs

Figure 3, Figure 2 and Table 4 illustrate preliminary results for the competing subgroup scenario. The scenario is illustrated in Figure 2, which depicts the results of a single run, and shows each subgroup circled. Figure 3 shows the variation of the global cost, as well as the individual group costs, with each deal made for the four specified cases.

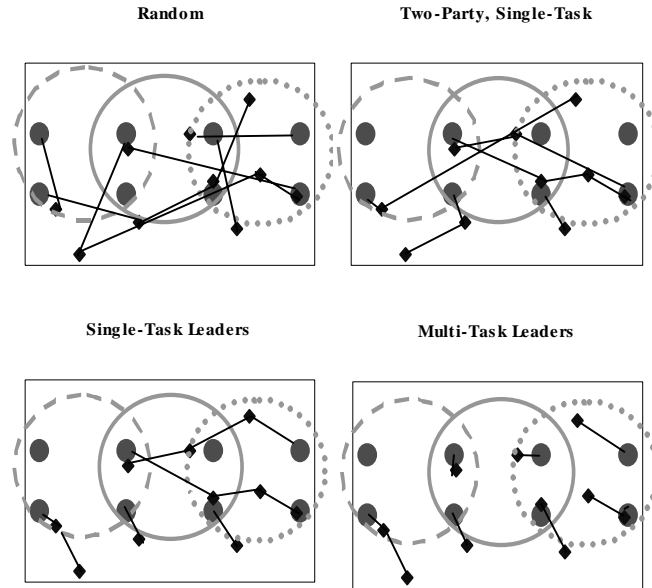


Figure 2: Solution for TSP with 3 overlapping subgroups of 4 robots each and 10 tasks

	Cost	Iterations	Improved
Random	9091	-	0.0%
No Leader (TP)			
4 TPST	4598	8	48.9%
2 TPST + 2 TPMT	4379	9	51.2%
MPST (ST Leader)	4312	6	52.1%
MPMT (MT Leader)	3687	6	58.9%

Table 4: Performance averaged over 100 randomly generated 8-robot (heterogeneous), 10-task TSPs with 3 overlapping groups of 4 robots each

In the case of the competing sub-groups, optimization within a group always reduces intra-group task costs and the global task cost as seen in Figure 3. (A group's cost is simply the sum of the task costs of the robots within the group). Note that the subgroups exchange tasks sequentially. However, in the case of overlapping groups, a group's intra-group cost can rise while an overlapping group is optimizing because new tasks can enter into the first group from the overlapping group's deals. (The profit of each group however will always stay the same or rise, and never fall, because each robot will only make profitable deals).

Table 4 reports the performance averaged over 100 randomly generated task distributions. Again, the results show that on average the local optimization with leaders improves the global profit.

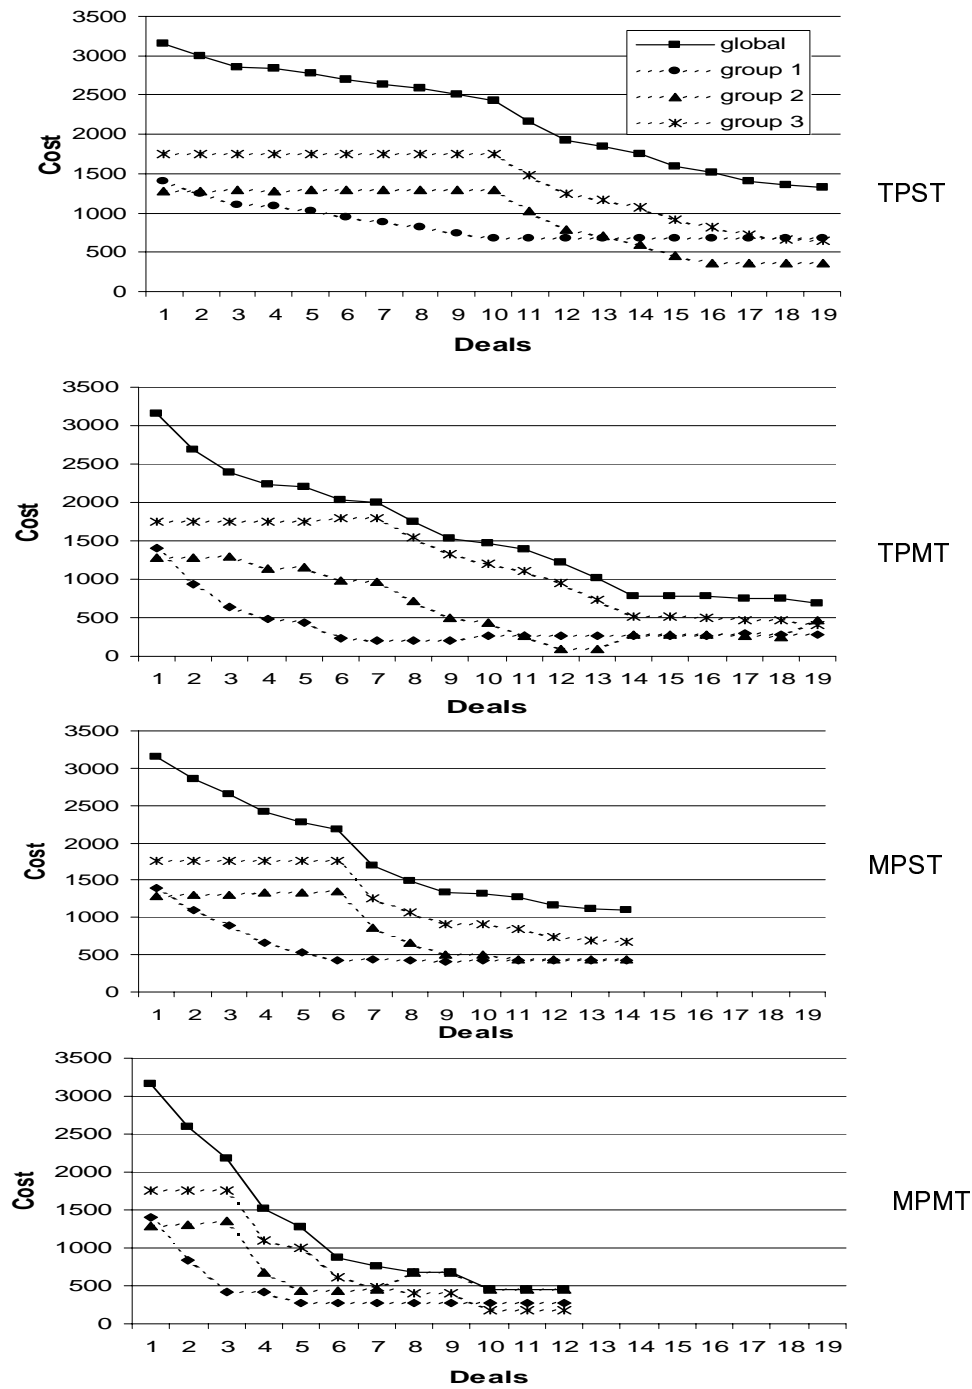


Figure 3: Cost reduction for 8 robots and 30 tasks with 3 sub groups of 4 robots each

8. Conclusions and Future Work

Presented results show that enhanced negotiation capabilities improve the performance of market-based task allocation, and that leaders can considerably reduce global costs in market-based multirobot

coordination. Initial experiments for optimizing within robot sub-groups with leaders also proved promising. Future work includes implementing these capabilities on a robot team and further extensions of the market approach. Proposed enhancements include more detailed analysis of optimizing with leaders, dealing with time constraints, and experimentation with different task domains. The goal of this work is to produce a robust and efficient market-based multirobot coordination architecture.

9. References

- [1] Arkin, R. C., Balch, T., 1997. *AuRA: Principles and Practice in Review*, Journal of Experimental & Theoretical Artificial Intelligence, Vol. 9, No. 2/3, pp.175-188.
 - [2] Brumitt, B. L., Stentz, A. 1996. *Dynamic Mission Planning for Multiple Mobile Robots*. Proceedings of the IEEE International Conference on Robotics and Automation.
 - [3] Chevallier, D., and Payandeh, S., 2000. *On Kinematic Geometry of Multi-Agent Manipulating System Based on the Contact Force Information*, The 6th International Conference on Intelligent Autonomous Systems (IAS-6), pp.188-195.
 - [4] Cormen, T. H., Leiserson, C. E., and Rivest, R. L. eds. 1990. *Introduction To Algorithms*. Cambridge, MA: The MIT Press/McGraw-Hill Book Company.
 - [5] Dias, M. B. and Stentz, A. 2001. *A Market Approach to Multirobot Coordination*. Technical Report, CMU-RI - TR-01-26, Robotics Institute, Carnegie Mellon University.
 - [6] Dias, M. B. and Stentz, A., 2000. *A Free Market Architecture for Distributed Control of a Multirobot System*. Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6).
 - [7] Gerkey, B. P. and Mataric, M. J. Submitted 2001. *Sold! Market methods for multi-robot control*. To appear in IEEE Transactions on Robotics and Automation, Special Issue on Multi-robot Systems, 2002.
 - [8] Golfarelli, M. and Rizzi, S. 2000. *Spatio-temporal clustering of tasks for swap-based negotiation protocols in multi-agent systems*. Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6), 172-179.
 - [9] Golfarelli, M., Maio, and D., Rizzi, S. 1997. *A Task-Swap Negotiation Protocol Based on the Contract Net Paradigm*. Technical Report, 005-97, CSITE (Research Center For Informatics And Telecommunication Systems, associated with the University of Bologna, Italy).
 - [10] Papadimitriou, C. H., and Steiglitz, K. eds. 1998. *Combinatorial Optimization: Algorithms and Complexity*. Mineola, NY: Dover Publications, Inc.
 - [11] Rabideau, G., Estlin, T., Chien, T., and Barrett, A. 1999. *A Comparison of Coordinated Planning Methods for Cooperating Rovers*. Proceedings of the American Institute of Aeronautics and Astronautics (AIAA) Space Technology Conference.
 - [12] Sandholm, T. 2002. *Algorithm for Optimal Winner Determination in Combinatorial Auctions*. Artificial Intelligence, 135, 1-54.
 - [13] Sandholm, T., and Suri, S. 2000. *Improved Algorithms for Optimal Winner Determination in Combinatorial Auctions and Generalizations*. Proceedings of the National Conference on Artificial Intelligence (AAAI).
 - [14] Smith, R. 1980. *The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver*. IEEE Transactions on Computers C-29 (12).
 - [15] Stentz, A. and Dias, M. B. 1999. *A Free Market Architecture for Coordinating Multiple Robots*. Technical Report, CMU-RI-TR-99-42, Robotics Institute, Carnegie Mellon Univ.
 - [16] Thayer, S. M., Dias, M. B., Digney, B. L., Stentz, A., Nabbe, B., and Hebert, M. 2000. *Distributed robotic mapping of extreme environments*. In the proceedings of SPIE: Vol. 4195: Mobile Robots XV and Telemanipulator and Telepresence Technologies VII.
 - [17] Wellman, M. P., and Wurman, P. R., 1998. *Market-Aware Agents for a Multiagent World*, Robotics and Autonomous Systems, Volume 24, pp.115-125.
 - [18] Zlot, R., Stentz, A., Dias, M. B., and Thayer, S. 2002. *Multi-Robot Exploration Controlled By A Market Economy*. Proceedings of the IEEE International Conference on Robotics and Automation.
-