A COMPARISON OF HUMAN AND MARKET-BASED ROBOT TASK PLANNERS

Guido Zarrella, Robert Gaimari and Bradley Goodman
The MITRE Corporation, 202 Burlington Road, Bedford MA 01803, USA

Keywords: Market-based multi-robot planning, intelligent control systems, distributed control, planning and scheduling, tight coordination, task deadlines, decision support.

Abstract: Urban search and rescue, reconnaissance, manufacturing, and team sports are all problem domains requiring multiple agents that are able to collaborate intelligently to achieve a team goal. In these domains task planning and assignment can be challenging to robots and humans alike. In this paper we introduce a market-based distributed task planning algorithm that has been adapted for heterogeneous, tightly coordinated robots in domains with time deadlines. We also report the results of our experiments comparing the robots' decisions with the decisions produced by ten teams of humans performing an identical search and rescue task. The outcome provides insight into the types of problems for which information technology can add value by providing decision support for human problem solvers.

1 INTRODUCTION

There are many modern problems that are not efficiently solved by a single human or robot. In domains like search and rescue, reconnaissance, and RoboCup, any attempt to solve the problem with a single robot may be inefficient, failure prone, or completely impossible. In these circumstances a team of agents must collaborate intelligently and task planning becomes central to the team success.

The extremes of multi-robot task planning and allocation algorithms are centralized and distributed approaches. In a centralized approach one agent plans the actions of the entire team and distributes the orders. In a distributed approach, each robot is responsible for creating its own plan using only local communication among robots. Centralized methods possess the key advantage of having all information needed to generate a globally optimal plan, while distributed approaches tend to be more scalable, robust to failure, and faster to respond to changes in the local environment. The ideal algorithm would combine features of both approaches to create a robust planning mechanism that is able to find a reasonable approximation of the optimal solution.

Past research into decentralized market-based task allocation protocols (Walsh et al., 1998; Dias, 2004; Lagoudakis et al., 2005) has been motivated by an attempt to design one such algorithm. In a market-based algorithm the robots bid against each other for tasks while acting rationally to maximize personal profit based on local calculations of cost and reward. This will move the entire team on average toward a globally efficient solution if the costs and revenues functions are properly constructed (Gerkey and Mataric, 2004). A market-based approach allows robot teams to reason efficiently about task allocation and resource management while preserving the ability of members of the team to adapt rapidly and robustly in the face of a dynamic environment. This technique mimics the flexibility of a free market economy by allowing ad-hoc teams to cooperate or compete opportunistically.

Prior research has demonstrated the effectiveness of variants of Dias' market-based TraderBots in several domains. The approach has been applied to tightly coordinated tasks that require heterogeneous, dynamically formed teams (Jones et al., 2006a). In this work two types of treasure hunting robots collaborate to simultaneously map an environment and detect the treasure within it. The TraderBots approach has also been used for task assignment in domains with time deadlines (Jones et al., 2006b), for example in homogeneous teams of fire-fighting robots completing tasks in which the reward for extinguishing a fire decays as a function of the elapsed time.

Yet another class of problems combines elements from the above domains. Collaborative

Zarrella G., Gaimari R. and Goodman B. (2007).
A COMPARISON OF HUMAN AND MARKET-BASED ROBOT TASK PLANNERS.
In *Proceedings of the Fourth International Conference on Informatics in Control, Automation and Robotics*, pages 149-154
DOI: 10.5220/0001631701490154

Time Sensitive Targeting (TST) is a domain requiring a diverse team of agents able to coordinate in discovering, assessing, prioritizing and solving new tasks within a very limited amount of time. This requires heterogeneous, dynamically formed teams that are both tightly coordinated and capable of reasoning about task deadlines. Search and rescue is one real-world example of a TST problem. For instance an avalanche rescue team's goal might be to "find each buried survivor and dig him or her out of the snow within sixty minutes." In this case searchers and diggers need to form dynamically changing and complementary teams to rescue as many survivors as possible within a limited time.

Time Sensitive Targeting can be a difficult problem solving task for humans as well as robots. Frequently decisions must be made about how to re-evaluate team strategy to make the best use of scarce resources. This makes TST an ideal testbed for a market-based task planning and allocation algorithm.

This paper describes our attempt to design and evaluate the first market-based planning system capable of reasoning in situations requiring tightly coordinated, deadline aware agents. In Section 2 we describe the specifics of our simulated Time Sensitive Targeting domain. We introduce our planning algorithm in Section 3. In Section 4 we discuss our experiments involving teams of humans attempting to solve a TST problem. In Section 5 we contrast the human and robot results, and in Section 6 we present our conclusions about the potential for the application of information technology to benefit teams of human decision makers.

2 TST SCENARIO

The central element of solving a Time Sensitive Targeting problem is the ability to assess and respond to emerging tasks within a limited window of time. The typical TST task requires a coordinated effort between a large number of specialized information gathering and action taking agents. Furthermore it is essential that the team is able to continually reprioritize its goals as new information arrives from the noisy and rapidly changing environment.

We designed a simulated TST scenario to use in our task planning and problem solving experiments. Our scenario is a type of Search and Rescue problem in which agents attempt to locate, investigate, and rescue six simultaneously moving targets before each target's time deadline expires.

The premise of the scenario is that the Coast Guard is responsible for monitoring three areas of ocean for sick or injured animals. The Coast Guard is provided with a fleet of specialized vehicles such as helicopters, boats, and submarines. The goal is to use these vehicles to find, diagnose, and rescue a series of endangered animals. In our experiments the fleet of vehicles was controlled either by a small team of humans or by our market-based robot task planning algorithm.

Over the course of the 90 minutes of an exercise, the Coast Guard receives messages containing reports of the general locations where distressed animals have been sighted. A message provides the type of animal, an approximate latitude-longitude, a time deadline for task completion (e.g. cure the sick manatee within 30 minutes or it will die), and the relative value of the task (represented by the maximum reward offered for task completion).

The Coast Guard's vehicle fleet includes a heterogeneous collection of robots. There are three main categories of vehicles.

Radar Sensors are planes and boats equipped with radar or sonar sensing capabilities. They are generally very fast and have large sensing range, so they can get to a location quickly, pinpoint where an animal is located, and track an animal as it moves. They can share the information they gather with other teammates. Due to the limitations of radar, this type of sensor is not able to determine an animal's species or diagnose an illness.

Video Sensors include boats and helicopters with visual sensing capabilities. They are able to identify animal types and diagnose diseases. They can also report the information they have gathered to the rest of the team. However they tend to move slowly and have limited sensing range, so they are best used in tandem with other sensors.

Rescue Workers are boats or submarines outfitted with equipment for capturing or curing an animal in distress. This is the only type of vehicle capable of saving an animal once it has been located. They are generally about as fast as radar sensors, but they have no sensors of their own. They must rely on reports from the sensor robots for navigation data. Also, they are only allowed to assist an animal after the proper diagnosis has been made by a video sensor.

The Coast Guard has multiple robots in each group. Even within groups there are variations of individual characteristics such as speed or sensor range. There are 33 vehicles in total, divided between three separate areas of ocean.

Because of the specialization of the robots, they are required to form ad hoc teams to fully complete any task. Each team must, at a minimum, consist of two robots: a video sensor to find the animal and make the diagnosis, and a rescue worker to assist the animal. A radar sensor is not required but its speed and sensor range can greatly reduce the overall time needed for a team to assist an animal.

3 MARKET-BASED ALGORITHM

We chose to develop our market-based multi-robot task planning algorithm within a controlled simulation environment. The entire package was written in Java using the JADE agent framework (Bellifemine et al., 2001). Our agents used only local robot-to-robot communication to implement the task planning protocol. The planning algorithm shares many similarities with TraderBots and other existing market-based approaches. Our goal was to extend the existing approaches to be capable of performing planning in domains with both task deadlines and tightly coordinated ad hoc teams.

The agents in our simulation trade labor for a fictional currency. An agent earns revenue for the successful completion of one of the tasks the team has been asked to perform, but only if the task is completed before the time deadline arrives. The agent incurs costs in the process of doing work to achieve a goal; these costs are proportional to the amount of time spent working towards the task. An agent also condsiders the opportunity cost (Schneider et al., 2005) of agreeing to perform a task. The self interested agents will only bid on a task if the potential revenue outweighs the sum of the impending costs. Agents buy and sell tasks from each other, forming efficient, specialized teams in the process. The cost and revenue functions we have chosen are conducive to fostering teams that solve problems as quickly as possible without over-committing the existing resources.

This section of the paper contains a high-level description of our implementation. (Gaimari et al., 2007) provides more detail about the algorithm.

3.1 Agents

The *TraderAgent* is the building block of the robot economy. Each TraderAgent controls one robot in the simulation environment. A TraderAgent's primary job, as the name implies, is to trade tasks.

Any agent that owns a task may put it up for auction, announcing the maximum reward it is prepared to pay. Other TraderAgents that wish to bid for the job may do so, and after one round of bidding the seller announces the winner. In standard re-auctioning this passes ownership of the task to the buying agent; this agent must have a robot with the same capabilities as the seller. The selling agent's only responsibility thereafter is to pay the promised bid to the buyer upon completion of the task.

In our system there is an additional type of re-auctioning that occurs. Since the robots must work together in teams to complete the tasks, some re-auctions are for the purpose of teambuilding among agents with different capabilities. In this case both agents retain ownership of the task. For each task a TraderAgent owns, there is a corresponding list of the teammates it is working with. If an agent owns multiple tasks it can belong to multiple teams.

New tasks are given to a special agent that executes the initial auction. This agent does not control a robot in the simulation.

3.2 Bidding on Auctions

When a TraderAgent receives an auction announcement, it performs the following steps:

- It calculates the estimated cost for performing the task. In our scenario the cost is given by the amount of time required to accomplish a task. Since the description of a task provides noisy and imprecise information about the location of an animal, costs cannot be determined exactly in advance. The agent prepares a cost estimate based on its technical abilities and current location.
- It calculates the opportunity cost associated with accepting responsibility for the task. This represents the likelihood that the agent will be able to win hypothetical future tasks. Robots with especially unique abilities will have higher opportunity costs than more common types of robots. Opportunity cost is also affected by the robot's location on the map, as some areas are more desirable for finding work than others.
- It calculates the desired profit margin. This is a function of the opportunity cost and the difference between the offered reward and the estimated cost. Robots with low opportunity costs will lower their desired profit margin in an attempt to increase the chance of winning the current auction.
- It calculates the final bid amount and places a bid if the cost plus the desired profit margin is less than the reward offered by the seller.

3.3 Collecting Payment

Once a task is completed, each TraderAgent reports that fact to the agent it bought the task from, asking to be paid. Domains with tightly coordinated, heterogeneous teams and time deadlines require special handling of payment allocation. In this case the teams are made up of robots that do their jobs at greatly different speeds. Slower robots can lead to much higher costs and lower rewards than a faster robot may have originally estimated. If the cost estimations are too inaccurate, the ability of agents to prioritize different tasks is damaged.

In our system an agent penalizes its teammates when the team underperforms expectations. Each agent requests the amount of payment agreed upon during the bidding process. As the payments are distributed, each agent compares its actual cost to the estimated cost it had initially planned upon. The difference between these is deducted from the amount paid to the next agent. This agent then adds the difference in its own actual cost and original estimate, plus the amount it was penalized by its seller. The penalty moves down the chain in this fashion until it finally ends where it belongs, on the slowest member of the team. These payments reflect the amount of money the original agents would have bid had they had known the true cost of working with slower robots. This penalty system provides feedback that allows the robots to learn improvements to their cost estimation and bidding practices.

4 HUMAN EXPERIMENTS

We tested the performance of teams of people on an isomorphic version of the Coast Guard search and rescue problem. The performance results of these teams of humans are directly comparable to the performance results of our market-based robot teams.

In these experiments, each team consisted of three college educated adults. The teams were mixed sex and made up of computer literate participants between the ages of 28 and 65. The members of the teams were provided with computer tools allowing them to view maps of the environment and control the movements and actions of the simulated robot vehicles. The participants were working in the same room and were permitted to speak with each other but were not allowed to look at the others' computer displays. Each member of the team was randomly assigned a unique and complementary role.

The **Intel Officer** acted as the team leader and was responsible for coordinating the team response to targets assigned to the group. This officer received the messages containing the rumored locations of new targets. The messages also specified a time deadline by which the task had to be completed. The intel officer was then expected to share this new information with the team and monitor the group's progress toward the goal.

The **Sensor Analyst** commanded the fleet of 20 heterogeneous sensor devices, including video equipped helicopters and radar planes. The sensor analyst was responsible for choosing which sensors to use, for ordering changes in sensor paths, and for monitoring the state of each sensor to check for newly detected items.

The **Rescue Worker** commanded a fleet of 13 heterogeneous rescue vehicles. This analyst was responsible for choosing which rescue vehicles to deploy, for ordering changes to each vehicle's path, and for giving the official order to rescue an animal.

As in the robot experiments, the teams were expected to locate and positively identify each target using their sensors before rescuing the animal. The experiment was an exercise in communication and team problem solving. Successful prosecution of a target was dependent on the participants' ability to 1) share relevant information without distracting each other from the task at hand, 2) interpret the state of the environment in a timely fashion, and 3) choose appropriate actions to execute. The simulation was developed as a simplification of real world exercises performed by similar teams of TST analysts (Goodman et al., 2005).

5 EXPERIMENTAL RESULTS

We evaluated the performance of the human and robot teams on our search and rescue TST scenario. Ten teams of three people attempted the problem. Each experiment lasted for 90 minutes. During this time, six targets were assigned to each team. The first three targets were assigned at 15 minute intervals, and the last three targets were assigned at 5 minute intervals. Each target had a time deadline between the 80th and 90th minutes of the experiment.

Table 1 shows the number of tasks completed by each human team. The best teams completed four of the six tasks before the time deadline. The worst teams were unable to successfully complete any of the tasks. The average number of tasks completed by the ten teams was 1.9, and the median was 2. In all, the teams of humans completed 32% of the tasks.

Table 1: The number of tasks, out of 6, completed by each group before the time deadline.

Team #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
# Tasks Finished on Time	2	3	4	1	3	0	0	0	4	2

Table 2 shows the number of teams that were able to successfully complete each task before the time deadline. Note that two of the tasks (#4, #6) were not completed by any of the teams. Another task (#2) was completed by all teams except for those groups that did not complete any tasks. These figures indicate that in general the tasks were not trivial to solve by teams of humans attempting the assignment, and that there was a good mix of difficulty levels in the problems presented to the teams.

Table 2: The number of teams, out of 10, that completed each task before the deadline.

Task #	1	2	3	4	5	6
# of Successful Teams	4	7	3	0	5	0

Table 3: The tasks completed by the autonomous robots before the deadline.

Task #	1	2	3	4	5	6
Solved before Deadline?	Y	Y	Y	Y	Y	N

The results of the robot team are displayed in Table 3. The team of robots completed five of the six tasks, a success rate better than best of the human teams. This demonstrates the ability of the robots to apply effective team building and task assignment strategies. We also use length of time before solution to compare robot team performance to human team performance, shown in Figure 1.

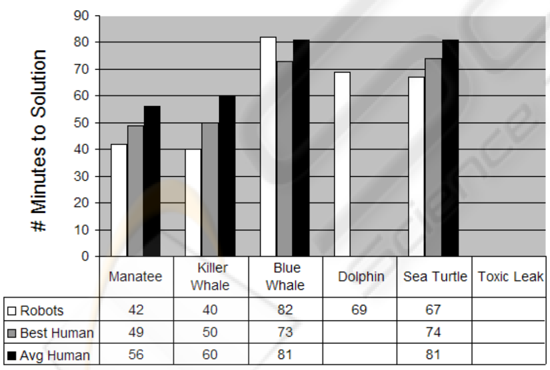

Figure 1: A comparison of the time to solution for the robot team, average human team, and best human team. Lower is better.

The robot teams compare very favorably to the human teams. The simulated agents were much faster than the best human teams in three of the four tasks that were solved by both humans and robots. The agents also were able to complete the Dolphin task, which none of the ten human teams had successfully accomplished within the time deadline. The simulated agents did fail to complete one task, but none of the human teams were able to successfully complete that task either.

6 APPLICATIONS AND CONCLUSION

The ad hoc teams of distributed market-based task planners demonstrated better performance on a simplified Time Sensitive Targeting task than the teams of humans attempting the same task. This result demonstrates that it is feasible to use our planning algorithms on tightly coordinated and time constrained tasks. The result is especially interesting in light of the fact that real-life problem solvers, such as military TST analysts, are humans collaborating in ad-hoc teams to attempt to combine forces into one integrated, efficient system.

What are the reasons for these differences in performance, and how can we use advances in information technology to improve human efficiency? Humans have an advantage over computers in that a lifetime of interactions with other humans allow them to plan and coordinate actions without the need for a formal communication and negotiation structure. Humans are also naturally able to integrate new information into the planning process in an online manner. Therefore the results described above are at least moderately surprising.

However, one key limiting factor on human performance is that humans have limited attention resources. It isn't possible for a single person to attend to the output of all twenty sensors simultaneously. As the number of concurrent tasks increase, human teams can suffer from increased cognitive load, which can dramatically affect a team's ability to respond to new information in a timely manner. One example of this can be seen in our human TST experiments, in which the average time delay between receiving and reading a new e-mail message increased steadily as more concurrent tasks were added.

In essence, the teams of humans are exhibiting the same drawbacks of a centralized multi-robot planning algorithm. Information from sensors must propagate to the top of the chain of command before a plan can be implemented that reflects changes in the state of the task. For some domains this is an adequate solution; unfortunately humans do not "scale" well to larger scenarios in which attention resources must be divided between larger numbers of targets. The results of our experiments demonstrate that TST teams can struggle when forced to make decisions about which targets are most worth pursuing given limited attention and resources. Real world teams are routinely forced into this situation. At SIMEX, a realistic TST simulation that uses real analysts from various government forces, 145 vehicles are manned by 30 operators pursuing any number of targets (Loren, 2004).

The market-based robot planning system, in these situations, is able to benefit from its distributed nature. As each autonomous agent receives updates on the state of the environment, this information is immediately propagated to the affected agents. This means that new tasks or newly sensed targets are promptly incorporated into the team plan. In the robot teams, the performance bottleneck is the quality of the decision making process rather than the availability of relevant data.

It is unreasonable to suggest that intelligent agents can replace the human decision makers in high risk Time Sensitive Targeting environments. The results from our simplified and noise-free environment can't necessarily be extrapolated to apply in far more complex real-world situations. The research does however indicate that there is value in applying intelligent control systems and other information technology to complement human decision makers by mitigating human weaknesses.

Our future work in this domain is focused on incorporating the task planning agents into an intelligent cognitive aide. The aide will draw attention to relevant events and changes in the environmental state. We could also use this cognitive aide to improve training methods by teaching decision makers to focus their attention on the most critical plan-changing events.

We have shown it is possible to use intelligent control systems to improve upon the results exhibited by teams of human decision makers. Our hope for the future is that it is possible to combine human and robotic planning methods to yield even better results.

ACKNOWLEDGEMENTS

The MITRE Technology Program supported the research described here. We are also grateful for the assistance of Brian C. Williams and Lars Blackmore at the Massachusetts Institute of Technology.

REFERENCES

Bellifemine F., Poggi A., Rimassa G. (2001). Developing multi-agent systems with a FIPA-compliant agent framework. *Software-Practice and Experience*, *31*, 103–128.

Dias M.B. (2004). TraderBots: A New Paradigm for Robust and Efficient Multirobot Coordination in Dynamic Environments. Doctoral dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA.

Gaimari, R., Zarrella, G., Goodman, B. (2007). Multi-Robot Task Allocation with Tightly Coordinated Tasks and Deadlines using Market-Based Methods. Proceedings of *Workshop on Multi-Agent Robotic Systems (MARS)*.

Gerkey B. and Mataric M. (2003). Multi-robot Task Allocation: Analyzing the Complexity and Optimality of Key Architectures. Proceedings of *IEEE Conference on Robotics and Automation*.

Goodman B., Linton F., Gaimari R., Hitzeman J., Ross H., Zarrella G. (2005). Using Dialogue Features to Predict Trouble During Collaborative Learning. *User Modeling and User-Adapted Interaction*, *15* (1-2), 85-134.

Jones E., Browning B., Dias M. B., Argall B., Veloso M., Stentz A. (2006a). Dynamically Formed Heterogeneous Robot Teams Performing Tightly-Coordinated Tasks. Proceedings of *International Conference on Robotics and Automation*.

Jones E., Dias M. B., Stentz A. (2006b). Learning-enhanced Market-based Task Allocation for Disaster Response. Tech report CMU-RI-TR-06-48, Robotics Institute, Carnegie Mellon University.

Lagoudakis M., Markakis E., Kempe D., Keskinocak P., Kleywegt A., Koenig S., Tovey C., Meyerson A., Jain S. (2005). Auction-Based Multi-Robot Routing. *Robotics: Science and Systems*. Retrieved at http://www.roboticsproceedings.org/rss01/p45.pdf

Loren, Lew. (2004, Fall). Experimentation and Prototyping Laboratories Forge Military Process and Product Improvements. *Edge Magazine*. Retrieved at www.mitre.org/news/the_edge/fall_04/loren.html

Schneider J., Apfelbaum D., Bagnell D., Simmons R. (2005). Learning Opportunity Costs in Multi-Robot Market Based Planners. Proceedings of *IEEE Conference on Robotics and Automation*.

Walsh W., Wellman M., Wurman P., MacKie-Mason J. (1998). Some Economics of Market-Based Distributed Scheduling. Proceedings of *International Conference on Distributed Computing Systems*.