A MARKET APPROACH TO TIGHTLY-COUPLED MULTI-ROBOT COORDINATION: FIRST RESULTS

Nidhi Kalra
Anthony Stentz

Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

ABSTRACT

In multi-robot systems, a group of robots cooperate to accomplish tasks more efficiently than could a single robot. In this paper, we apply the market approach to the problem of tightly-coupled, distributed control of a multi-robot team. In this approach, the robots act in a self-interested manner to maximize their profits. Since profit is tied to accomplishment of the group task, the robots act to advance the team objectives. The robots achieve tight coordination by repeatedly evaluating fine-grained actions, reacting to their team members' actions in the process. We present preliminary results on the application of the market approach to the sweeping perimeter problem, a task with implications in security and the mapping of dynamic environments.

INTRODUCTION

A group of collaborating robots performs certain tasks better than a single robot. Multi-robots allow the concurrent execution of decomposable tasks such as mapping and construction. By distributing the task temporally and spatially, the group often completes the task in a more efficient manner than would an individual. They also enable execution of tasks that cannot be completed by a single robot, e.g. lifting a very heavy object or climbing a mountain using a belay system. Dias and Stentz [4] present a survey of multi-robot systems and outline the advantages and complexities associated with them.

Many multi-robot problems can be solved by decomposing the problem into subtasks that are then executed by individual robots. Robots cooperate during the decomposition process but execute subtasks independently. For some problems, the decomposition is not so clean. Multiple robots are needed to execute each given subtask and must closely coordinate during the process.

In multi-robot systems, the effectiveness of the subtask distribution determines the system's ability to find an optimal solution to the original task. In centralized approaches, a single agent or robot decomposes the problem and assigns subtasks to the other team members. Examples of centralized systems can be found in work done by Jensen and Veloso [5] and Chaimowicz et. al. [2]. Centralized approaches have the potential to be optimal since the single planning agent can investigate all possibilities before making the assignments. For large problems, this approach is intractable because the number of possibilities can be exponentially large. Centralized approaches also suffer from poor robustness to uncertainty. If a robot discovers that the world is not as expected, then the 'correct' response in a centralized system is to transmit this information back to the planning agent who can potentially re-assign all of the tasks to achieve optimality. However, this process can be very slow and impractical if communication is poor. On the other hand, distributed systems such as those produced by Arkin and Balch [1], Ostergaard et. al. [6], Chevallier and Payandeh [3], and Dias and Stentz [4] offer robots local control and tasking. Distributed approaches are very fast and responsive to uncertainty because individual agents can change their own plans based on what they see and experience in the world. However, the solutions tend to be sub-optimal since the subtask assignment is done via a computationally simple strategy such as first-come-first-serve or random assignment.

Developed by Stentz and Dias [7], the market approach (now referred to as 'TraderBots') addresses the division of labor in multi-robot systems using an economic model. Robots bid on subtasks to execute, receiving a revenue for successful execution from which a cost is deducted for resources consumed. Robots negotiate tasks to maximize their individual profit, which results in the task accomplishment with minimal resource expenditure. The market approach tries to achieve the best of both centralized and distributed systems by being fast and nimble in response to uncertainty and change. If time permits, it attempts to be optimal as well by investigating plans that better optimize all or part of the task and selling them on the market. If these plans are truly more optimal, they will be adopted since they consume fewer resources and generate more profit for the individual robots.

To date, the market approach has been applied to problems that can be decomposed into independent subtasks. In such cases, revenues and costs are straightforward to define. Thayer et. al [8] and Zlot et. al. [9] employ the market approach for robotic exploration. Stentz and Dias [7] use it to control a group of foraging robots. In this paper, we investigate the application of the market approach to the class of problems that cannot be neatly divided among the robots. Such tasks require a plan selection scheme that permits tight coordination. Secondly, they demand a solution to the credit assignment problem of how individual robots should be paid and penalized for their actions when their successes and failures critically depend on other robots' actions.

The sweeping perimeter task requires the close coordination of multiple robots. For this task, the robots form an expanding perimeter and sweep out an area with their sensors as they move. The robots are instructed not to create any blind spots along the perimeter that would allow a foe to enter the secure area. The task is easy in an open area, but more difficult when there are obstacles to breach that impair line of sight. Robots must tightly coordinate and collaborate with their neighbors to surround obstacles without creating gaps through which something could pass undetected. Solutions to the sweeping perimeter task have obvious importance for security applications. They are also significant for mapping dynamic areas where anything entering or exiting the mapped area must cross the perimeter and therefore will be detected.

MARKET APPROACH

The market approach models a market economy: a group of robots simulate an economic system in which individual robots act out of self-interest. The completion of each subtask requires some cost expenditure and generates some revenue for a robot. Robots offer up subtasks, evaluate the available subtasks, and bid against each other for those subtasks that generate maximum profit. In general, robots receive revenue when they take steps that bring the system closer to its goal and incur costs when their actions move the system away from its goal. Payment-penalty functions are designed so that the overall system approaches a solution as the individual robots try to maximize their personal profits.

The market approach provides a distributed solution to multi-robot control which is generally computationally faster and more responsive to environmental changes than a centralized solution. In the sweeping perimeter task, an individual robot must maintain and expand the perimeter it forms with its neighbors, so a distributed approach fits the local nature of the problem. The market provides the mechanisms necessary to facilitate close coordination by enabling individual robots to buy their neighbors' participation in a subtask at the level of fine-grained actions.

COORDINATION ALGORITHM

In the sweeping perimeter task, each robot in the group has two neighboring robots. If two neighboring robots can see the same location, then the perimeter between them is intact. The group's total perimeter consists of intersecting fields of view of all neighboring robots. Additionally, an area is secure if there are no two neighboring robots between which the perimeter is broken. Consequently, nothing can leave or enter the secure area without being seen by at least one member of the group.

Our preliminary simulation allows the robots to tightly coordinate by making the subtasks in the system spatially small and short in duration. The robots begin with a closed perimeter and consider everything in that area secure. When choosing its next action, a robot only chooses a move for the next small increment of time; it makes no extended plans. Moreover, each robot only evaluates unit steps in the set of eight cardinal and intermediate directions. These limitations on step size and plan length ensure that robots tightly coordinate by frequently and regularly considering each other's positions when choosing their future positions.

In our first implementation, we address the credit assignment problem by attempting to decouple the actions of the robots. We evaluate each robot's move under the assumption that the other robots remain stationary over the move. Thus, each robot's move has a distinct effect on the perimeter. This setup makes credit assignment manageable: each robot receives payment proportional to its individual contribution to the area. Conversely, if the robot's move decreases the area, it is penalized. Additionally, robots are heavily penalized for breaking the perimeter; in simulation, this penalty is high enough that the robots never break the perimeter.

PRELIMINARY RESULTS

Even though the current system both limits the robots' possible choices for moves and simplifies the representation of their surroundings, the robots are able to expand their perimeter in a number of simple environments. Figure 1 illustrates that, without obstacles, the robots easily expand to reach a local maximum of the area swept. The robots end up stretched to the limit in a near-circular configuration. Any move would either break the perimeter or decrease the area secured. When a robot considers its next move, several steps may generate the same revenue because the moves are discretized. Robots randomly choose between these steps, showing no preference for those that place it farthest from its neighbors. While these steps are locally equivalent, they uniquely affect the shape of the perimeter. Consequently, these random steps may prevent the robots from arriving at the optimal circular solution.

The system also overcomes simple obstacles in the environment. In Figure 2, the environment contains an obstacle that the robots have enveloped without breaking the perimeter. Though the robots do not have a long term plan, the obstacles are simple enough that the robots can establish a new perimeter on the far side of the obstacles without breaking the perimeter on the near side. Forming this new perimeter allows robots to increase the area secured and generate revenue, making it advantageous for the robots to move around the obstacle.

Figures 3 and 4 illustrate the system's effectiveness as the environment's complexity increases to three obstacles. In Figure 3, the robots' initial moves were fortuitous: when they arrived at the obstacles, they had enough overlapping sensor coverage to envelop them. Conversely, the robots in Figure 4 cannot successfully surround the leftmost obstacle without breaking the perimeter, due to maximum sensor range limitations. These figures illustrate two ways robots might cooperate. First, robots might offer their neighbors incentives to follow them and provide more leverage to move without breaking the perimeter. Alternately, robots could rearrange neighbors, in a sense short-circuiting the perimeter, to free up robots that could move around

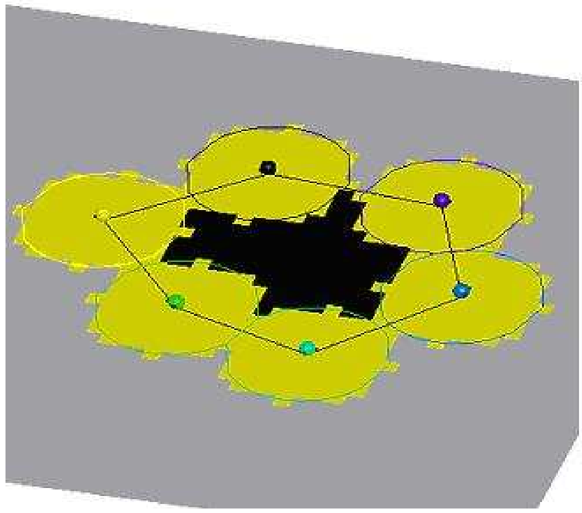

Figure 1: Seven robots completely expand to cover an area without obstacles. The seven robots are shown as spheres. The gray area is unswept, the black area is swept, and the yellow area is currently visible to one or more robots.

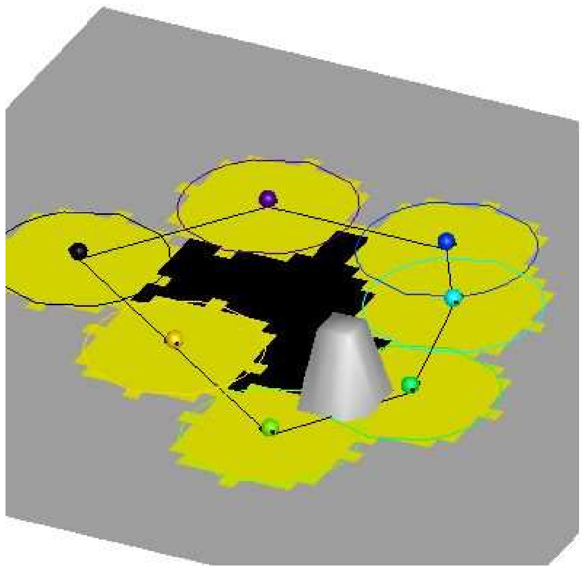

Figure 2: Seven robots can envelop a single obstacle in the environment since they can establish a perimeter on the far side without breaking the perimeter on the near side.

obstacles. Under the market approach, one robot could convince another to detach from the perimeter and relocate.

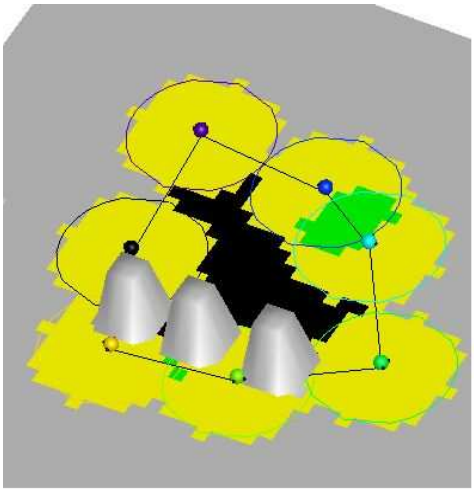

Figure 3: These seven robots envelop three obstacles in the environment. Their initial moves give them enough flexibility later to surround the obstacles without breaking the perimeter.

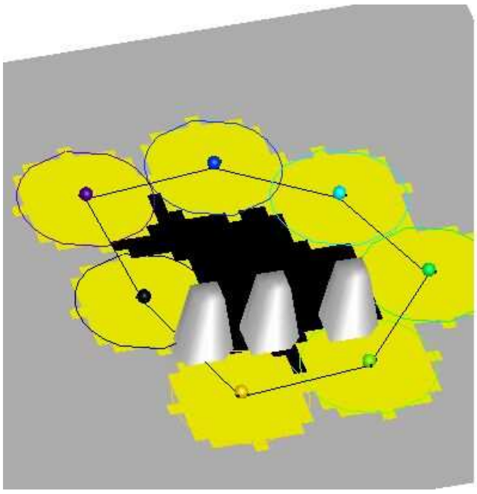

Figure 4: These seven robots cannot envelop three obstacles in the environment without breaking the perimeter. Their initial moves later minimized their overlapping sensor coverage.

As shown in Figure 5, the robots are unable to cover more than three obstacles, particularly when they are in a complex configuration rather than in a straight line. Covering a complex environment requires extended planning: robots need to recognize that they are coming upon a group of obstacles and provide incentives for their neighbors to follow them.

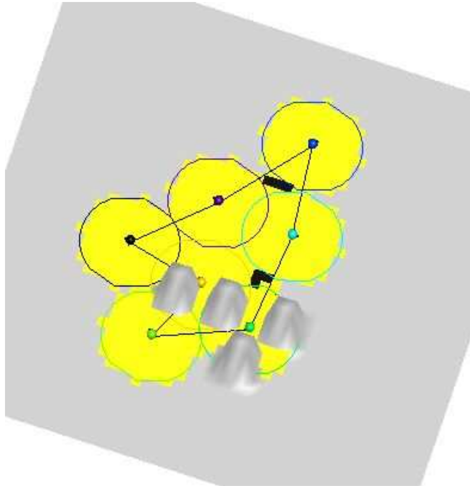

Figure 5: Seven robots cannot envelop obstacles that have a more complicated configuration . Covering these complex environments requires extended planning.

CONCLUSIONS AND FUTURE WORK

Thus far, our work is preliminary and we expect future work to incorporate more market mechanisms. Our first implementation achieves coordination implicitly rather than explicitly. The subtasks are treated as fully decomposed and independent in the market. Coordination occurs because the subtasks are very small moves. The robots effectively coordinate with each other by "reacting" to their neighbors' new positions at each time step. This is effective for covering clear areas and areas with simple, isolated obstacles.

Often, the group's failures result from robots limiting each others' movements in an effort to keep the perimeter intact. To alleviate this problem, robots could offer their neighbors incentives for following them and providing them more freedom to move. The robot following would then receive a portion of the revenue generated by the new area covered, even though it did not itself increase the area. Alternately, robots might recognize when they are coming upon an obstacle and offer their neighbors incentives for moving immediately along the other side of the obstacle.

In a more complex scenario, a robot encountering a group of obstacles could offer incentives to another robot to disconnect from the perimeter, travel around to the other side of the obstacles, and cover the area from the "blind" side. Again, this robot would not itself increase the area contained. However, by aiding other robots, it would be paid by those robots for its efforts.

We expect to test the results of these simulations on a group of robots in a large area with many obstacles. These robots will be tasked with securing as large an area as possible. We expect our research will extend the applications of the market approach to a new class of problems and provide novel distributed solutions to problems that are traditionally centralized.

ACKNOWLEDGMENTS

This work was sponsored by the U.S. Army Research Laboratory, under contract **Robotics Collaborative Technology Alliance** (contract number DAAD19-01-2-0012). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U. S. Government. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

References

[1] R. Arkin and T. Balch. Aura: Principles and practice in review. *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2/3):175–188, 1997.

[2] L. Chaimowicz, T. Sugar, V. Kumar, and M. Campos. An architecture for tightly coupled multi-robot cooperation. In *Proceedings of the International Conference on Robotics and Automation, Seoul, Korea, May 21-26.*, 2001.

[3] D. Chevallier and S. Payandeh. On kinematic geometry of multi-agent manipulating system based on the contact force information. In *Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6)*, pages 188–195, 2000.

[4] M. Dias and A. Stentz. A market approach to multirobot coordination. Technical report, Robotics Institute, Carnegie Mellon University, 2001.

[5] R. M. Jensen and M. M. Veloso. Obdd-based universal planning: Specifying and solving planning problems for synchronized agents in non-deterministic domains. *Lecture Notes in Computer Science*, 1600, 1999.

[6] E. H. Ostergaard, M. J. Mataric, and G. S. Sukhatme. Distributed multi-robot task allocation for emergency handling. In *Proceedings of the IEEE/RSJ International Conference on Robots and Systems (IROS)*, 2001.

[7] A. Stentz and M. Dias. A free market architecture for coordinating multiple robots. Technical report, Robotics Institute, Carnegie Mellon University, 1999.

[8] S. Thayer, M. Dias, B. Digney, A. Stentz, B. Nabbe, and M. Hebert. Distributed robotic mapping of extreme environments. In *Proceedings of SPIE*, volume 4195: Mobile Robots XV and Telemanipulator and Telepresence Technologies VII, 2000.

[9] R. Zlot, A. Stentz, M. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *Proceedings of the IEEE International Converence on Robotics and Automation*, 2002.