

A Free Market Architecture for Coordinating Multiple Robots

Anthony Stentz and M. Bernardine Dias

CMU-RI -TR-99-42

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

December 1999

© 1999 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

The coordination of a large group of robots to solve a specified task is a difficult problem. Centralized approaches can be computationally intractable, brittle, and unresponsive to change. Distributed approaches are not as prone to these problems, but they can be highly sub-optimal. This work introduces a novel approach for coordinating robots based on the free market system. Market economies are a proven way to organize a large number of individuals into a productive group. The free market approach defines revenue and cost functions across the possible plans for executing a specified task. The task is accomplished by dividing it into sub-tasks and allowing the robots to bid and negotiate to carry out these sub-tasks. Cooperation and competition emerge as the robots execute the task while trying to maximize their personal profits. The result promises to be a highly robust multi-robot team that can efficiently exploit resources and opportunistically deal with uncertainties in a dynamic environment. A detailed example of how this model could be applied to a foraging task is presented and the different characteristics of the approach are highlighted in the context of the example. The ability to scale this approach to achieve more complex tasks is also briefly explored. We are in the process of implementing this architecture on a team of ten robots engaged in the task of mapping an interior environment.

Table of Contents

1. Introduction	1
2. Related Work	2
3. The Free Market System	2
3.1 Determining Revenues and Costs	2
3.2 The Role of Price and the Bidding Process	3
3.3 Cooperation vs. Competition	4
3.4 Self Organization	4
3.5 Learning and Adaptation	4
4. A Detailed Example	5
4.1 Gold Finding Mission	5
4.2 Relationship of Computational Capacity to Optimality of Solution	5
4.3 Computational Capacity De-coupled from Implementation	7
4.4 Optimization without Complete Information	7
4.5 Effects of Supply and Demand on Emergent Roles in the Economy	8
4.6 Scaling to A More Complex Task	8
5. Conclusion	9
6. Acknowledgements	10
7. References	10

Table of Figures

<i>Figure 1: Two haulers bid on the loads from two graspers.</i>	6
<i>Figure 2: The winning TSP tour from Robot A.</i>	7
<i>Figure 3: Two sub-tours of two loads each for Robots A and B.</i>	8

1. Introduction

For many applications, a team of robots can be effectively used. A robot team can accomplish a given task more quickly than a single agent can by dividing the task into sub-tasks and executing them concurrently. A team can also make effective use of specialists designed for a single purpose (e.g., scouting an area, picking up objects, hauling payload), rather than requiring that a single robot be a generalist, capable of performing all tasks but expert at no tasks.

The difficulty arises in coordinating all of these robots to perform a single, global task. One approach is to consider the robot team to be a single robot “system” with many degrees of freedom. A central computer coordinates the group optimally to perform the specified task. The problem is that optimal coordination is computationally difficult—the best known algorithms are exponential in complexity. Thus, the approach is intractable for teams larger than a few robots, even if a powerful computer is available. Additionally, the approach assumes that all information about the robots and their environment can be transmitted to a single location for processing and that this information does not change during the time that an optimal plan is constructed. These assumptions are unrealistic for problems in which the environment is unknown and/or changing, communication is limited, and robots behave in unpredictable ways. Another weakness with this approach is that it produces a highly vulnerable system. That is, if the leader (the central planning unit) malfunctions, a new leader must be available or the entire team is disabled.

Local and distributed approaches address the problems that arise with centralized, globally coordinated approaches. The idea is that each robot operates largely independently, acting on information that is locally available through its sensors. A robot may coordinate with other robots in its vicinity, perhaps to divide a problem into multiple sub-problems or to work together on a sub-task that cannot be accomplished by a single robot. The advantage of this approach is that it typically requires little computation, since each robot need only plan and execute its own activities. Also, little communication is required, since the robots only communicate with others in their vicinity. The robots are better able to respond to unknown or changing environments, since they sense and respond to the environment locally. Moreover, the system is more robust since the entire team’s performance no longer depends on the guidance of a single leader. The approach works best for problems that can be decomposed into largely unrelated sub-problems, or problems for which a desired group behavior results from the aggregate of individual behaviors and interactions, as with some biological species such as bees and ants [8, 23]. The problem with biological analogues is that it may be difficult to determine what set of individual behaviors, if any, will produce the global behavior desired, and how close to optimal it will be. Thus, although some biologically inspired systems have been successfully implemented, these systems are limited to performing simple tasks sub-optimally. Enabling these systems to perform highly complex tasks has not been possible thus far.

Typically, the robot designs are constrained to a sufficiently small set. The robots can be programmed to perform individual operations reliably, such as driving from A to B, picking up an object, or scanning an area with a sensor. Furthermore, there may be known algorithms for coordinating at least some of the robots in a rational and productive way. Thus, the parts of the problem that remain unsolved are dealing with unknown and changing environments and globally coordinating robot activities to accomplish a given task.

Instead of a biologically inspired system, consider an economic system for coordinating robots. An economy is nothing more than a population of agents (i.e., citizens) producing a global output. The agents coordinate with each other to produce an aggregate set of goods. The human race has field-tested a variety of economic systems and a winner has emerged. Socialist/communist systems rely heavily on a centralized approach, with five-year plans coordinating the instruments of production. This economic system is prone to the same problems as other centrally organized systems: inability to gather all salient information, uncertainty in how to optimize with it, and unresponsiveness to changing conditions. Such systems are also subject to a more subtle yet equally fatal flaw. Since economic output is divided equally amongst the entire population, individuals have little incentive to work harder or more efficiently than what is required to minimally comply with the economic plan. Individual input is de-coupled from individual output. The net effect is a sluggish, brittle, inefficient economy.

Laissez-faire capitalism, on the other hand, is thriving. The free market is unencumbered by centralized planning; instead, individuals are free to exchange goods and services and enter into contracts as they see fit. *Despite the fact that individuals in the economy act only to advance their own self-interests, the aggregate effect is a highly productive society.* Individuals are in the best position to understand their needs and the means to satisfy them. Thus,

individuals reap the direct benefits of their own good decisions and suffer the direct consequences of their bad ones. At times they cooperate with other members of the society to achieve an outcome greater than that possible by each member alone. At times they compete with other members to provide goods or services at the lowest possible cost, thus eliminating waste and inefficiency. But at every turn, the individual members act solely to reap the greatest profit for themselves. In this report, we describe a method for applying these powerful mechanisms to the task of coordinating a team of robots.

The next section summarizes some of the related work published in recent years. A detailed description of our approach for applying the mechanisms of the free market system to a multi-robot team follows. We then provide a detailed example of an application of a free market model to a multi-robot system, and conclude with a summary of the market-based approach for controlling a multi-robot team and an outline of future research directions that we plan to take.

2. Related Work

The past decade has witnessed a growing focus on multi-agent systems. Mataric [14] presents a comprehensive summary of some of the principal efforts in this area of research. Jensen and Veloso [9], Švestka and Overmars [20], and Brumitt and Stentz [6] are examples of the centralized approach to control a multi-robot system organized hierarchically. A number of researchers have developed biologically inspired, locally reactive, behavior-based systems to carry out simple tasks [1, 2, 3, 5, 13, 14]. These distributed systems have found applications in many different domains. Some behavior-based systems have been extended to more complex task domains; for example see work published by Mataric [13,14] and Arkin et al. [1]. Mataric [13, 14] shows how more complex behaviors can be built from a basic set of behaviors for a multi-robot team. Arkin et al. [1] present a flexible, behavior-based, software architecture for developing mission-specific robot behaviors for urban warfare application.

Other novel approaches have been adopted to control multi-robot teams. Parker [17] explores the advantages and disadvantages of centralized versus distributed control for cooperative agent teams. Tambe [21] introduces a method of enabling flexible teamwork by providing the agents with general models of teamwork. Pagello et al. [15] examine multi-agent cooperation in the soccer domain through implicit communication. Veloso et al. [22] investigate methods of anticipation in order to improve cooperation of multi-robot teams in the soccer domain. Schneider-Fontán and Mataric [18, 19] present an approach of territorial division of tasks for a multi-robot team. Taking a similar approach, Parker [16] introduces a temporal division of tasks to allow fault-tolerant multi-robot cooperation.

Work most similar to our approach has been carried out mainly in the software-agent domain. Faratin et al. [7] present a model for negotiation between autonomous agents in the business process management domain. Bonatti et al. [4] introduce a prototype of a simulator “INFOGEN” for modeling complex distributed multi-agent systems in the information industry domain. Krovi et al. [11] adopt a genetic algorithm approach for exploring the impact of the interaction of different agent behaviors on the process and outcome of negotiations. Other researchers [4, 7, 11] have investigated different negotiation strategies. Johnson et al. [10] examine the volatility and agent adaptability in a “bar-attendance” market model, while Lux and Marchesi [12] use a multi-agent model of financial markets to investigate scaling and criticality in the market. To the best of our knowledge, we are the first to present a free market architecture for controlling a multi-robot team, specifically for the sensing, motion planning, communication, coordination, and task allocation issues pertinent to these mobile machines.

3. The Free Market System

3.1 Determining Revenues and Costs

Consider a team of robots assembled to perform a particular mission. The goal of the team is to perform the mission well while minimizing costs. A function, \mathbf{F}_r , is needed that maps possible mission outcomes onto revenue

values. Another function, F_i , is needed that maps possible schemes for performing the mission onto cost values. As a team, the goal is to execute some plan P such that profit, $F_o(P) - F_i(P)$, is maximized.

But it is not enough to define just the revenue and cost functions for the team. These functions must provide a means for distributing the revenue and assessing costs to individual robots. Preferably, these individual revenues and costs are assigned based on factors over which the individuals have direct control. For example, if the mission is to find and retrieve a scattering of objects, the team's revenue, F_o , could be the number of objects retrieved (converted to a "cash" value), and the team's cost, F_i , could be the amount of energy consumed by the entire team to find the objects (again, converted to a cash value). The individual revenues and costs, f_o and f_i , could be the cash value of the number of objects turned in and the energy expended, respectively, by *that* individual.

Therefore, the sum of the individual revenues and costs equals the team's revenues and costs. However, the distribution is not even: individuals are compensated in accordance with their contribution to the overall mission, based on factors that are within the control of the individual. An individual that maximizes its own personal production and minimizes its own personal cost receives a larger share of the overall profit. Therefore, by acting strictly in their own self-interests, individuals maximize not only their own profit but also the overall profit of the team.

3.2 The Role of Price and the Bidding Process

Robots receive revenue and incur costs for accomplishing a specific team mission, 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. For example, a robot may not be equipped to find objects for which the team function provides revenue, but it can transport the objects to the goal once they have been found. Therefore, this haulage robot provides a service to the robots that find the objects, and it receives payment for performing such a service.

The *price* dictates the payment amount for the good or service. How is the price determined? Assume that robot **A** would like to purchase a service from robot **B**. Robot **B** incurs a cost Y for performing the service. Robot **A** can make an additional revenue of X if **B** performs the service for it. Therefore, if $X > Y$, then both parties have an incentive to execute the deal. But how should the composite profit, $X - Y$, be divided amongst the two parties? It may sound fair to split the winnings $(X - Y) / 2$ by setting the price at $(X + Y) / 2$. But robots **A** and **B** may have other opportunities—they may be considering other deals that contend for the same money and resources. Robot **A** may have another way of spending its money that nets it more than $(X - Y) / 2$ dollars. Likewise, robot **B** may have another opportunity to provide its service for a net exceeding $(X - Y) / 2$ dollars.

Since all of the factors that contribute to a robot's decision about how much it can charge or pay for a given good or service may be hidden or difficult to reason about, a common approach is to *bid* for a good or service until a mutually acceptable price is found. For example, robot **A** could start by bidding a price of Y (i.e., robot **A** receives the entire profit). Robot **B** could decline and counter with a bid of X (i.e., robot **B** receives the entire profit). The idea is to start by bidding a price that is personally most favorable, and then successively retreat from this position until a price is mutually agreed upon.

Note that a given robot can negotiate several potential deals at the same time. It begins by bidding the most favorable price for itself for all of the deals, successively retreats from this position with counter bids, and closes the first deal that is mutually acceptable. Note also that a deal can be multi-party, requiring that all parties agree before any part of the deal is binding.

The negotiated price will tend toward the intersection of the supply and demand curves for a given service. If a service is in high demand or short supply, the price will be high. This information will prompt other suppliers to enter the fray, driving the price down. Likewise, if demand is low or supply high, the low price will drive suppliers into another line of business. Thus, price serves to optimize the matching of supply to demand.

Finally, it is important to note that price and bidding are low bandwidth mechanisms for communicating aggregate information about costs. When consumers decide between purchasing apple juice or orange juice for breakfast, they do not analyze land acreage dedicated to both crops, the costs of producing each, the demand for each, and the impact of weather and pest infestations. Instead, they merely look at the price of each and weigh them against their own personal preferences. Yet the price *encodes* all of these factors in a concise fashion that enables them to make a locally optimal decision based on low-bandwidth information available at the point of sale.

3.3 Cooperation vs. Competition

As described in the previous section, robots interact with each other to exchange goods and services. 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. Generally, robot teams foster cooperation between members of different types (heterogeneous). For instance, a robot able to grasp and lift objects and a robot able to transport objects could team together to provide a pick-and-place service that neither one could offer independently.

Conversely, two robots are competitive if they have the same role, that is, if the amount of profit that one can make is negatively affected by the presence of the other robot. Generally, robot teams foster competition amongst members of the same type (homogeneous). For instance, two robots that are able to transport objects compete for the services of a given grasping robot, thus driving the price down. Either one could charge more money if the other were not present.

These delineations are not strict however. Subgroups of heterogeneous robots could form that provide a given service. These subgroups would compete with each other, thus providing an example where robots of different types compete rather than cooperate with each other. Heterogeneous robots could also compete if the same task can be accomplished in different ways. Conversely, two robots of the same type may cooperate by agreeing to segment the market. This can make sense if the relative positions of the two robots make for a cost-effective division of labor (i.e. if the robots cooperate they could both make more profit than if they compete). In this sense, even homogeneous robots can be considered heterogeneous since they *differ* in location or situation. Homogeneous robots can also cooperate if accomplishing a specific task requires more than one robot. For example, several robots with grasping capability may need to cooperate in order to move a heavy object, or two robots with cameras may need to cooperate to provide stereo-mapping data. The flexibility of the market-model allows the robots to cooperate and compete as necessary to accomplish a task, regardless of the homogeneity or heterogeneity of the team.

3.4 Self Organization

Conspicuously absent from the free market system is a rigid, top-down hierarchy. That is not to say that the system is in anarchy; 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 mission, 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 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. The leader may be bidding not only against the individuals’ plans, but also against group plans produced by other prospective leaders. This self-organization need not be limited to two levels. A leader could organize a group of subgroups, each of which has its own leader.

From this example, it is evident that leadership roles are generally reserved for the “deep thinkers”, the robots with the computational capacity to reason about potentially complicated group interactions. If the team has no deep thinkers, the mission may still be achievable but in a less optimal fashion. The deep thinkers offer the opportunity to milk every last bit of efficiency out of the team. 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.

3.5 Learning and Adaptation

The robot economy is able to learn new behaviors and strategies as it executes its mission. This learning applies to both individual behaviors and negotiations as well as to the entire team. Individual robots may learn that certain strategies are not profitable, or that certain robots are apt to break a contract by failing to deliver the goods or proper payment. Individuals may also learn successful bidding strategies or which deals to offer when. The robot team may learn that certain types of robots are in over-supply, indicated by widespread bankruptcy or an inability to

make much money. Conversely, the robot team may learn that certain types of robots are in under-supply, evidenced by excessive profits captured by members of the type. Thus, the population can learn to exit members of one type and enter members of another. Moreover, in this approach, successful agents are able to accumulate wealth and perpetuate their winning strategies because of their ability to offer higher payments to other agents.

One of the greatest strengths of the market economy is its ability to deal successfully with changing conditions. Since the economy does not rely on a hierarchical structure for coordination and task assignment, the system is highly robust to changes in the environment, including malfunctioning robots. Disabling any single robot should not jeopardize the system's performance. By adding escape clauses for "broken deals", any tasks undertaken by a robot that malfunctions can be re-bid to other robots, and the entire mission can thus be accomplished. Furthermore, if a robot is partially disabled in some manner (e.g., a robot is immobilized due to losing a wheel), it can still be useful to the economy by seeking a new role that makes use of its remaining resources (e.g., computation or communication resources). Thus, the market model allows the robots to deal with the uncertainties of a dynamic environment in an opportunistic and adaptive manner.

4. A Detailed Example

4.1 Gold Finding Mission

In this section, we illustrate some of the advantages of the free market approach via a detailed example. Consider a team of robots chartered to find and retrieve pieces of gold. The pieces of gold are randomly scattered about a given area at unknown locations. There are "banks" where the gold can be cashed in for a fixed amount of money. The robots consume fuel in order to find and transport the gold to the banks. A productive economy in this case is one that maximizes profit, defined to be the total money paid out by the banks for the retrieved pieces of gold minus the total cost of the fuel consumed by all robots to retrieve the gold. Note that the maximum profit may be achieved by finding fewer than all of the pieces of gold. For example, one piece of gold may be so far away that it is not possible to retrieve it without accruing more fuel costs than would be recovered by selling the gold to the bank.

There are two types of robots on the team: grasping robots and haulage robots. The grasper is able to use its sensors to search a given area and pick up and carry single pieces of gold. Haulers are able to carry up to four pieces of gold at a time. They are loaded by graspers and can "dump" their load at a bank for cash payment. Both robots incur a cost from driving around. The haulage robots incur 1.5 times the cost per unit distance as the grasping robots, but they can carry four pieces.

Grasping and haulage robots are complementary. Graspers provide a function that haulers cannot: they can find and pick up pieces of gold. Without an ability to get gold into its bin, the hauler provides no marketable service. However, the grasper robots can make money by finding and hauling gold all of the way to the bank for sale. But the haulage robots offer a cheap mode of transportation if the grasper finds more than one piece of gold. Thus, the two robots are complementary: if they work together both can make more money than either could individually.

Grasping and haulage robots compete with other robots of the same type. Two graspers can both find pieces of gold. The one that drives the least distance can "produce" gold at the lower cost. Therefore, for a given price offered to a haulage robot, it can make more profit. In fact, it is better equipped to win a price war in the event that haulage robots are scarce or in high demand. This competition favors the more efficient provider.

4.2 Relationship of Computational Capacity to Optimality of Solution

Figure 1 shows two haulage robots, **A** and **B**, positioned near a bank. They know about two piles of gold (two pieces in each pile) deposited by a pair of grasping robots. Assume that robots **A** and **B** are fairly unsophisticated. They do not have enough computational capacity to calculate efficient tours for picking up the loads; instead, they just compute the cost of driving to a single load, picking it up, and returning it to the bank. Based on their calculation, they both make initial (high) offers to graspers **1** and **2** for picking up their loads and cashing them in. Robots **1** and **2** allow **A** and **B** to get into a bidding war for their gold. For cost reasons, **B** must drop out of the running for **1** and **A** must drop out for **2**. Note that this is a reasonable solution, given **A** and **B**'s limited

computational ability to optimize their own resources. The robots avoided the inefficient solution where **A** retrieves **2** and **B** retrieves **1**.

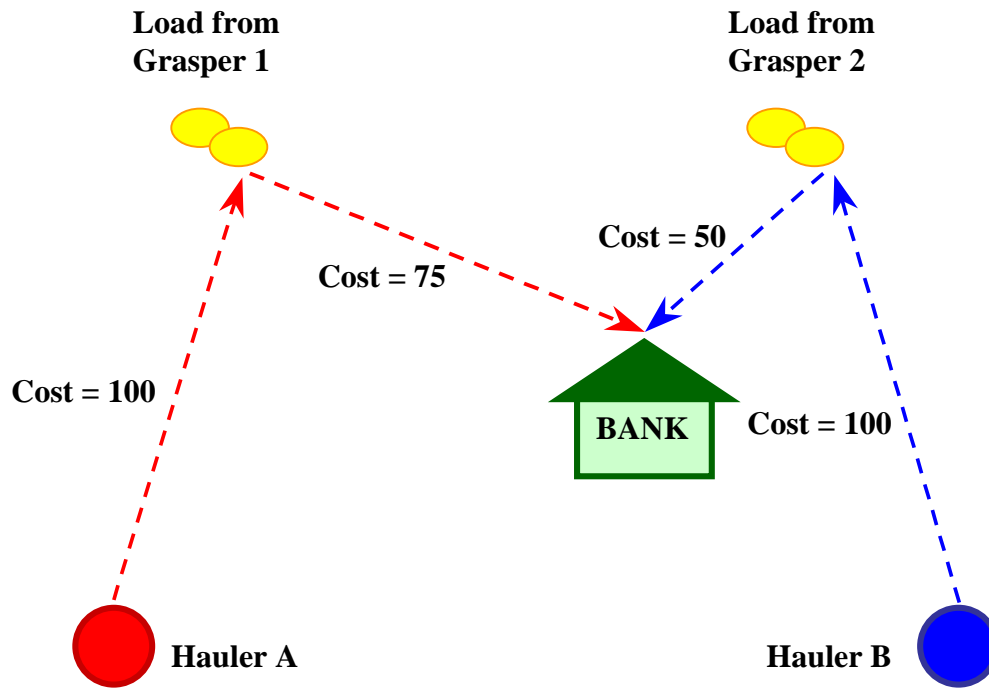


Figure 1: Two haulers bid on the loads from two graspers.

Now assume the same situation, except that **A** and **B** are capable of reasoning about optimal traveling salesman problem (TSP) tours, for a reasonably small number of loads. Both calculate a tour that visits both loads and returns to the bank. Robot **A** is able to perform the task for a lower cost than robot **B** (see Figure 2). Therefore, in the bidding process, as **A** and **B** lower their prices offered to robots **1** and **2** for their loads, robot **A** is able to win the bidding war since it can still make money at a lower price than **B** can. Note that this is the globally optimal solution, and was made possible by *smarter* participants.

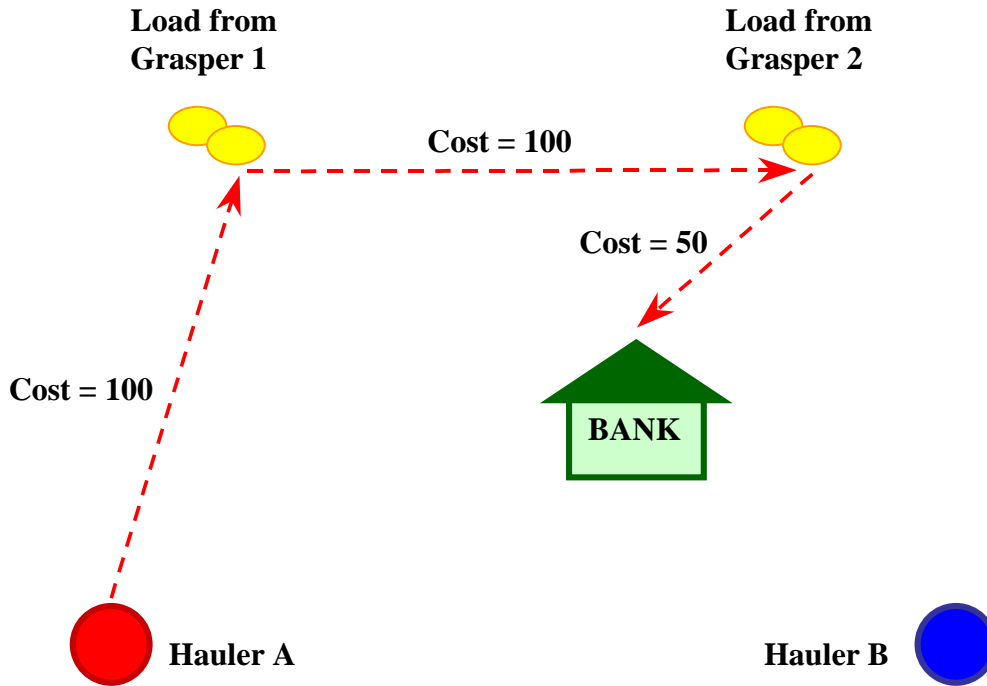


Figure 2: The winning TSP tour from Robot A.

4.3 Computational Capacity De-coupled from Implementation

Now let's assume the same situation, except that robot **A** has very limited computational capacity and is unable to reason about TSP tours, and **B** is able to do such reasoning. At first, it may seem like the end result will be that robot **B** will pick up both loads, since robot **A** will base its offers on the cost of driving out twice to pick up each load individually, and **B** will win the price war. But since **B** is smart, it will realize that **A** can pick up both loads for less cost than it can. *Therefore, robot B sells an optimal plan to robot A in exchange for a cut of the take.* Robots **1** and **2** agree to it, since robot **B**'s plan is more cost effective and they will net a higher price for their gold. Robot **A** agrees to it, since **B**'s plan is better and consumes less fuel for the same yield. Thus, all parties favor the deal because it is globally optimal with respect to costs: the additional profit can be distributed amongst the parties, including robot **B**, who acted solely as a *broker* for the deal. Therefore, individuals have an incentive to find and eliminate pockets of inefficiency, even if they merely devise rather than implement the solution.

4.4 Optimization without Complete Information

Consider once again the same situation, except that there are additional loads to be picked up (see Figure 3). These loads are not known to all parties. Robot **A** knows about Load **3**, and Robot **B** knows about Load **4**. This situation is typical, since any one robot cannot know the state of the whole world. The amount of information is too immense and it is constantly changing. Each robot is able to reason about an optimal tour of up to two loads. One strategy would be for each robot to bid on tours of one load each. As soon as a robot wins the job for picking up a load, it can bid on tours of two loads that include the first load. The figure shows the optimal recovery strategy for the four loads using optimal sub-tours of length two. Note that the presence of Loads **3** and **4** results in a different pick-up strategy for Loads **1** and **2**—in this case the loads are split between the two robots. In this way, the bidding process communicates information about the presence of a load that a given robot does not even know about.

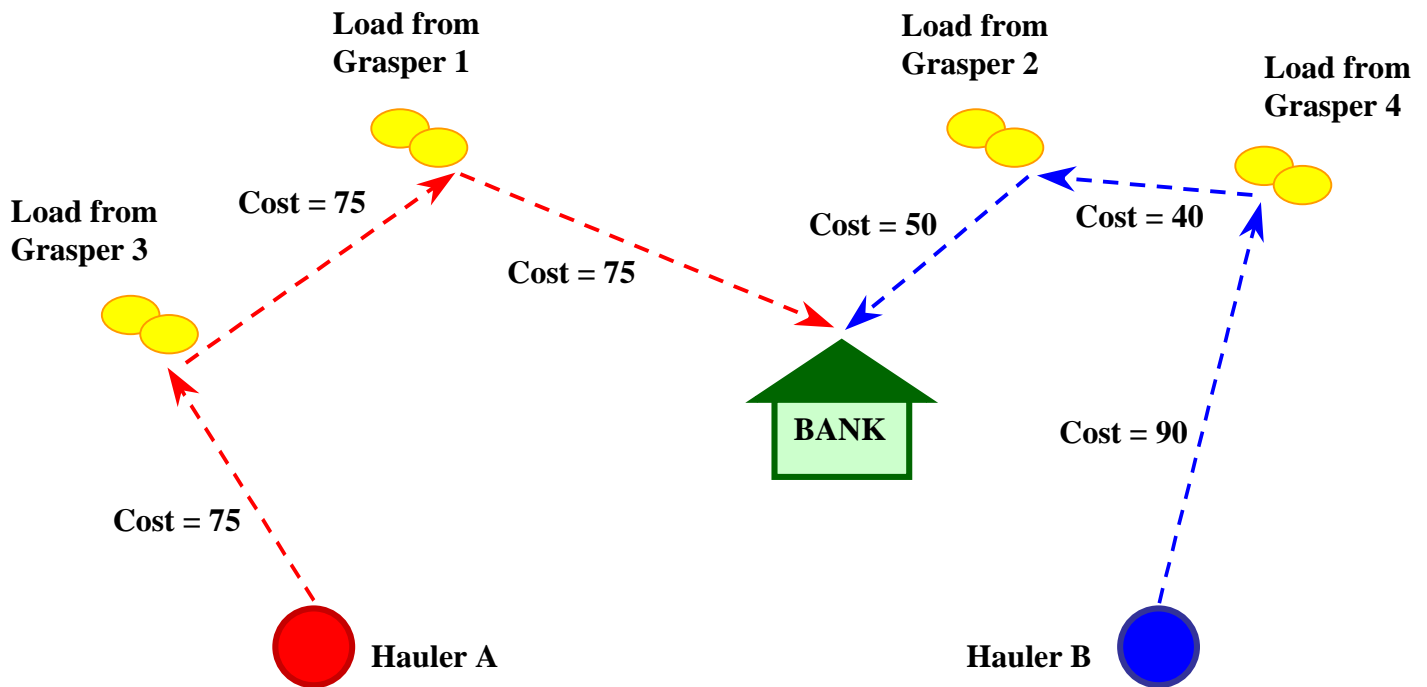


Figure 3: Two sub-tours of two loads each for Robots A and B.

4.5 Effects of Supply and Demand on Emergent Roles in the Economy

In this example, the availability of gold in relation to the number of graspers and haulers determines the supply and demand curves for commodities and services. A high supply of gold increases the demand for graspers to locate the gold and for haulers to carry the gold back to the bank. In a market-based economy, when the demand is high for a particular service, the price is high. Consequently, more suppliers of that service emerge because of the promise of larger profits, and because even less efficient providers are able to make money. Conversely, if a particular service is in low demand, the price is low. Less efficient suppliers are forced to seek other roles that make money. Similar scenarios can be envisioned in this foraging example. If the supply of gold is low in comparison to the number of graspers, and the gold is all located in areas close to the bank, the demand for haulers is low. Graspers are able to locate only one or two pieces of gold each, and have to travel only a short distance to the bank. Hence, it is more cost effective for the graspers to deliver the gold to the bank themselves. On the other hand, if the gold is located far away from the bank, the demand for haulers increases since the haulers can transport the gold over long distances at lower cost. Furthermore, if the number of haulers available is very low, and the number of graspers is high in comparison to the supply of gold, some of the graspers can potentially adopt “delivery” roles, even though they are less efficient at doing so. The high price for the hauling service enables them to assume this role and still turn a profit. The flexibility of the market-based architecture allows the robots to assume roles as necessary in order to accomplish the task in an efficient manner.

4.6 Scaling to A More Complex Task

It is easy to see how this model can be extended to a more complex task. For example, assume a team of robots is required to map a designated territory with unknown internal geometry. The environment could be hazardous and hence inaccessible to humans (and maybe even to some robots). Hence, the team must be robust to the loss of some robots, and also able to opportunistically explore the unknown territory. If each robot were equipped with localization, mapping, communicating, and computing ability, then the team would start by negotiating which areas to be mapped (i.e. each robot would try to travel in a different direction of the designated territory). Depending on the capability of the robots, different teams could form where robots played different roles to accomplish the task efficiently. For example, a robot could purchase the service of a second robot to transfer

messages and data to other robots that are beyond the communication range of the first robot. Similarly, a robot that computes a plan, that enables a sub-group of the robots to map a region more efficiently, can sell its plan to the sub-group. Also, if two robots have only one camera each, they could cooperate to acquire stereo data which would perhaps have greater value on the market. If a robot has few available computational cycles due to other responsibilities, it could purchase the services of a second robot to generate a map based on the data acquired by the first robot. In this manner, the market-based approach can facilitate different robot roles such as planners, communication relays, map builders, and visual data gatherers to emerge in order to accomplish the task efficiently.

Other advantages can be gleaned from the economic approach. Robots that successfully navigate through a dangerous area could sell information about the location and nature of the dangers and how to avoid them to other robots that have not yet passed through the area, thus enabling cooperative learning. Furthermore, a group of robots could cooperate to move heavy obstacles out of the way, thereby creating an access path to an area previously inaccessible, and then charge other robots a fee for the use of the passage. Similarly, if a robot gets trapped in some location, it could purchase the services of other robots to rescue it. For a high enough price, a robot could even be convinced to sacrifice itself for the success of the mission. Regardless of the service, if a demand for it arises, robots will spring into action to supply it, and the winners will be the ones that can provide it most cost effectively.

5. Conclusion

This paper introduces a free market system approach for coordinating a group of robots to achieve a given objective. The objective is achieved by individual robots cooperating and competing with each other to further their own self-interests. We expect a multi-robot system controlled via this market system based approach to be highly robust due to the lack of centralized control. The system should be able to efficiently utilize resources and opportunistically deal with uncertainties in a dynamic environment. A detailed example was presented to illustrate some of the mechanisms available to solve a resource allocation problem. Future work will investigate a formal structure for a “deal”, bidding strategies, hierarchical arrangements via self-organizing groups, revenue and cost functions, and bankruptcy and escape clauses. We will implement this architecture on a group of ten robots for mapping an urban environment.

6. Acknowledgements

This research was sponsored in part by DARPA, under contract "Cognitive Colonies" (contract number N66001-99-1-8921, monitored by SPAWAR). 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. The authors thank the members of the Cognitive Colonies group for their valuable contribution: Scott Thayer, Bruce Digney, Martial Hebert, and Bart Nabbe.

7. References

1. Arkin, R.C., Collins, T.R., and Endo, T., "Tactical Mobile Robot Mission Specification and Execution", *Mobile Robots*, September 1999.
2. Arkin, R. C., Balch, T., "AuRA: Principles and Practice in Review", *Journal of Experimental & Theoretical Artificial Intelligence*, Vol. 9, No. 2/3, pp.175-188, 1997.
3. Arkin, R. C., "Cooperation without Communication: Multiagent Schema-Based Robot Navigation", *Journal of Robotic Systems*, Vol. 9, No.3, pp. 351-364, 1992.
4. Bonatti, M., Ermoliev, Y. M., Gaivoronski, A. A., "Modeling of Multi-Agent Market Systems in the Presence of Uncertainty: The Case of Information Economy", *Robotics and Autonomous Systems*, Vol. 24, pp. 93-113, 1998.
5. Brooks, R. A., "Elephants Don't Play Chess", *Robotics and Autonomous Systems*, Vol. 6, pp.3-15, 1990.
6. Brumitt, B. L., Stentz, A., "Dynamic Mission Planning for Multiple Mobile Robots", *Proceedings of the IEEE International Conference on Robotics and Automation*, No. 3, pp. 2396-2401, 1996.
7. Faratin, P., Sierra, C., Jennings, N. R., "Negotiation Decision Functions for Autonomous Agents", *Robotics and Autonomous Systems*, Vol. 24, pp. 159-182, 1998.
8. Hölldobler, B., Wilson, E. O., *The Ants*, Belknap Press of Harvard Press, Cambridge, MA, pp. 1-3, 227-228, 244, 250, 260-262, 265, 268, 271-272, 1990.
9. Jensen, R. M., Veloso, M. M., "OBDD-based Universal Planning: Specifying and Solving Planning Problems for Synchronized Agents in Non-Deterministic Domains", *Lecture Notes in Computer Science*, No. 1600, pp. 213-248, 1999.
10. Johnson, N. F., Jarvis, S., Jonson, R., Cheung, P., Kwong, Y. R., Hui, P. M., "Volatility and Agent Adaptability in a Self-Organizing Market", *Physica A*, Vol. 258, No. 1-2, pp. 230-236, 1998.
11. Krovi, R., Graesser, A. C., Pracht, W. E., "Agent Behaviors in Virtual Negotiation Environments", *IEEE Transactions on Systems, Management, and Cybernetics – Part C: Applications and Reviews*, Vol. 29, No. 1, pp. 15-25, 1999.
12. Lux, T., Marchesi, M., "Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market", *Nature*, Vol. 397, No. 6719, pp. 498-500, 1999.
13. Mataric, M. J., "Coordination and Learning in Multi-Robot Systems", *IEEE Intelligent Systems*, pp. 6-8, 1998.
14. Mataric, M. J., "Issues and Approaches in the Design of Collective Autonomous Agents", *Robotics and Autonomous Systems*, Vol. 16, pp. 321-331, 1995.
15. Pagello, E., D'Angelo, A., Montsello, F., Garelli, F., Ferrari, C., "Cooperative Behaviors in Multi-Robot Systems through Implicit Communication", *Robotics and Autonomous Systems*, Vol. 29, No. 1, pp. 65-77, 1999.
16. Parker, L. E., "ALLIANCE: An Architecture for Fault Tolerant Multi-Robot Cooperation", *IEEE Transactions on Robotics and Automation*, Vol. 14, No.2, pp. 220-240, 1998.
17. Parker, L. E., "Designing Control Laws for Cooperative Agent Teams", *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 582-587, 1993.
18. Schneider-Fontán, M., Mataric, M. J., "Territorial Multi-Robot Task Division", *IEEE Transactions on Robotics and Automation*, Vol. 14, No. 5, 1998.
19. Schneider-Fontán, M., Mataric, M. J., "A Study of Territoriality: The Role of Critical Mass in Adaptive Task Division", *Proceedings, From Animals to Animats 4, Fourth International Conference on Simulation of Adaptive Behavior (SAB-96)*, MIT Press/Bradford Books, pp. 553-561, 1996.

20. Švestka, P., Overmars, M. H., "Coordinated Path Planning for Multiple Robots", *Robotics and Autonomous Systems*, Vol. 23, No. 4, pp. 125-133, 1998.
21. Tambe, M., "Towards Flexible Teamwork", *Journal of Artificial Intelligence Research*, Vol. 7, pp. 83-124, 1997.
22. Veloso, M., Stone, P., Bowling, M., "Anticipation: A Key for Collaboration in a Team of Agents", Submitted to the 3rd International Conference on Autonomous Agents, pp. 1-16, 1998.
23. von Frisch, K., The Dance Language and Orientation of Bees, The Belknap Press of Harvard University Press, Cambridge, MA, pp. 57-63, 236-256, 1967.