Carnegie Mellon University
Research Showcase @ CMU

Robotics Institute | School of Computer Science

12-2005

Market-Based Multirobot Coordination: A Comprehensive Survey and Analysis

Nidhi Kalra
Carnegie Mellon University

Robert Zlot
Carnegie Mellon University

M. Bernardine Dias
Carnegie Mellon University

Anthony Stentz
Carnegie Mellon University

Follow this and additional works at: http://repository.cmu.edu/robotics

Part of the Robotics Commons

This Technical Report is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Robotics Institute by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Market-Based Multirobot Coordination: A Comprehensive Survey and Analysis

Nidhi Kalra Robert Zlot M. Bernardine Dias
Anthony Stentz

CMU-RI-TR-05-16

December 2005*

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

© Carnegie Mellon University

*A version of this technical report was published in July 2006 as an article in the IEEE Special Issue on Multirobot Coordination. The original technical report has been updated with select publications that appeared between its original publication date in December 2005 and the article's publication date in July 2006.

Abstract

As robotic technology improves, we charge robots with increasingly varied and difficult tasks. Many of these tasks can potentially be completed better by a team of robots working together than by individual robots working alone. Coordination can lead to faster task completion, increased robustness, higher-quality solutions, and the completion of tasks impossible for single robots. Nevertheless, effective coordination can be difficult to achieve because of a range of adverse real-world conditions including dynamic events, changing task demands, resource failures, and limited deliberation time. The desire to overcome these challenges and harness the benefits of robot teams has made multirobot coordination a vital field in robotics research.

Of the resulting wealth of research, market-based multirobot coordination approaches in particular have received significant attention and are growing in popularity within the community. These approaches harness the principles of market economies—which have successfully governed human coordination for thousands of years—and use them to enable robot coordination. In market-based approaches, robots on the team act as self-interested agents operating in a virtual economy in which tasks and team resources are exchanged over the market in pursuit of individual profit. The essence of market-based approaches is that the process of robots trading tasks and resources with one another to maximize their wealth simultaneously improves the efficiency of the team.

Market-based approaches to multirobot coordination inherit many of the benefits associated with market economies, including flexibility, efficiency, responsiveness, robustness, scalability, and generality. In practice, they have been successfully implemented in a variety of domains ranging from mapping and exploration to robot soccer. The research literature on market-based approaches to coordination has now reached a critical mass that warrants a survey and analysis. This paper meets this need in three ways. First, it provides a tutorial on market-based approaches by discussing the motivating philosophy, defining the requirements and tradeoffs inherent in such approaches, analyzing their strengths and weaknesses, and placing them appropriately in the context of the larger set of approaches to multirobot coordination. Second, this paper surveys and analyzes the relevant literature. Third, it inspires and directs future research on this topic through a discussion of remaining challenges.

Contents

1 Introduction **1**

2 Overview **3**
2.1 Definition of a Market-based Approach 3
2.2 Auctions 4
2.3 Costs, Utilities, and Valuation 4
2.4 The Range of Coordination Approaches 5

3 Planning **6**
3.1 Related Work 6
3.1.1 Planning and Task Allocation 6
3.1.2 Planning and Task Decomposition 10
3.1.3 Planning and Task Execution 11
3.2 Future Challenges 12

4 Quality of Solution **12**
4.1 Related Work 13
4.1.1 Relationship to Operations Research 15
4.2 Future Challenges 17

5 Scalability **17**
5.1 Related Work 17
5.1.1 Computation and Communication Considerations 18
5.1.2 Selective or Opportunistic Centralization 21
5.2 Future Challenges 21

6 Dynamic Events and Environments **21**
6.1 Related Work 22
6.1.1 Robustness and Fluidity 22
6.1.2 Online Tasks 23
6.1.3 Uncertainty 24
6.2 Future Challenges 24

7 Heterogeneous Teams **25**
7.1 Related Work 25
7.2 Future Challenges 26

8 Learning and Adaptation **26**
8.1 Related Work 27
8.2 Future Challenges 27

9 Practical Considerations **27**
9.1 Related Work 27
9.1.1 Flexibility 27
9.1.2 Extensibility 28
9.1.3 Implementation 28
9.1.4 Comparisons 28
9.2 Future Challenges 29

10 Conclusions and Future Directions **29**

A Example Problems and Case Studies **37**
A.1 Basic Aggregation 37
A.1.1 Aggregation of Multiple-Robot Tasks 38
A.2 Market-based Exploration 39

1 Introduction

As robots become an integral part of human life, we charge them with increasingly varied and difficult tasks including planetary exploration, manufacturing and construction, medical assistance, search and rescue, and port and warehouse automation. Like humans, robots working in challenging domains can potentially perform better by working together in teams than by working alone. Ideally, robots will coordinate to redistribute resources amongst themselves in a way that enables them to accomplish their mission efficiently and reliably. Coordination can lead to faster task completion, increased robustness, higher-quality solutions, and the completion of tasks impossible for single robots. However, these domains simultaneously present many obstacles to effective coordination, such as dynamic events, changing task demands, resource failures, the presence of adversaries, and limited time, energy, computation, communication, sensing, and mobility. Therefore, coordinating a multirobot team requires overcoming many formidable research challenges.

Humans have met these coordination challenges for thousands of years with increasingly sophisticated market economies. In these economies, self-interested individuals and groups trade goods and services to maximize their own profit; simultaneously, this redistribution results in an efficient production of output for the system as a whole. Researchers have recently applied the principles of market economies to multirobot coordination. In market-based multirobot systems, robots are designed as self-interested agents that operate in a virtual economy. Both the tasks that must be completed and the available resources are commodities of measurable worth that can be traded. For example, tasks can be assigned to robots via market mechanisms such as auctions. When a robot completes a task, it receives some payment in the form of virtual money for providing a service to the team. However, the robot must also pay for the resources it consumed to complete the task. The essence of market-based approaches is that, in a well-designed system, the process of robots trading tasks and resources with one another to maximize individual profit simultaneously improves the efficiency of the team.

To illustrate this more concretely, consider a team of robots performing a distributed sensing mission on Mars. As illustrated in Figure 1, the robots must gather data from specific sites of interest to scientists while consuming the least amount of energy. One important aspect of completing the mission is to determine which robot should visit each site. We can solve this problem using a market-based approach in which robots compete in auctions for each task of visiting a site. After estimating their resource usage for an offered task and submitting bids based on those expected costs, the robot with the best bid is awarded a contract for that site.

Specifically, suppose that we offer a maximum reward of \$50 for each task and that robots incur a cost of \$2 for each meter of travel (since the resource of concern is energy consumed). This \$50 is a *reserve price* that essentially says that the task should only be attempted if the site can be reached by increasing one's path length by less than 25 meters. Further suppose that a robot A is only 5 meters from a site S. Since A would have to spend \$10 to complete the task, it bids \$10. Meanwhile, a robot B that is 10 meters from the site bids \$20. A is awarded the contract because it can perform the task more efficiently and for less than the reserve price.

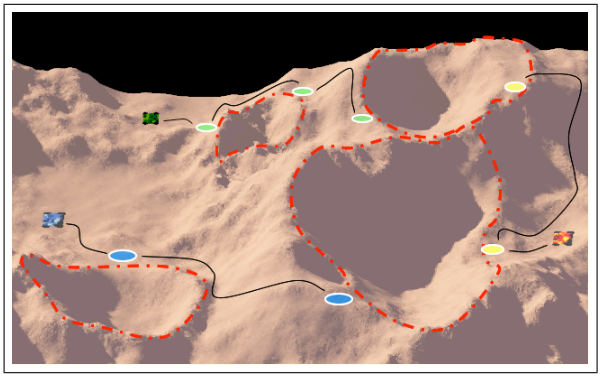

Figure 1: An illustration of three robots exploring Mars. The robots' task is to gather data around the four craters, which can be achieved by visiting the highlighted target sites.

This simple example illustrates the basic mechanism of a market-based approach to coordination. As the problem increases in complexity with the addition of more robots, more resources (*e.g.* time, network bandwidth, computing power, sensors, etc.), added constraints between the tasks, dynamically changing tasks, and so forth, the coordination approach requires added functionality to produce efficient solutions. We use this distributed sensing scenario throughout the remainder of this paper to illustrate the complexities of coordination and the diversity of market-based approaches.

The earliest examples of market-based multiagent coordination appeared in the literature over thirty years ago [20, 62] and have been modified and adopted for multirobot coordination in more recent years. This work is motivated by the growing popularity of market-based approaches and the lack of a comprehensive review of these approaches. This paper makes three contributions to the robotics literature. First, it provides a tutorial on market-based approaches by discussing the motivating philosophy, defining the requirements and tradeoffs inherent in such approaches, analyzing their strengths and weaknesses, and placing them appropriately in the context of the larger set of approaches to multirobot coordination. Second, this paper surveys and analyzes the relevant literature. Third, it inspires and directs future research on this topic through a discussion of remaining challenges. This paper in particular extends our shorter technical report [19]: it presents more concrete results, more thoroughly explores the topics, and includes new topics of discussion. Additionally, the appendix includes detailed example problems that demonstrate the step-by-step application of market-based approaches and a case-study of market-based multirobot exploration that highlights the variety of ways markets can be applied to one domain.

The scope of this paper is limited to market-based approaches for coordinating teams that include robots. Moreover, this review principally considers approaches that actively reason about the existence of other agents when coordinating the team, in contrast to approaches in which agents coexist. Nevertheless, related publications outside the stated scope of this paper are included as necessary to augment the discussion.

The following section provides an introduction to market-based mechanisms for

readers less familiar with the field. This overview is followed by an extensive review of market-based multirobot coordination approaches to date, categorized and analyzed across several relevant dimensions: planning, solution quality, scalability, dynamic events and environments, and heterogeneity. The paper concludes with a summary of the survey and future challenges in this research area. Further material on example problems and case studies is included in Appendix A.

2 Overview

In this section, we discuss key concepts that will provide a foundation for the remainder of the article, including a definition of market-based approaches and an introduction to auctions. We then place market-based approaches in the larger spectrum of coordination approaches.

2.1 Definition of a Market-based Approach

Most market-based multirobot and multiagent coordination approaches share a set of underlying elements. Market theory provides precise definitions for several of these elements. Borrowing from both bodies of literature, we define a market-based multirobot coordination approach based on the following requirements:

- The team is given an objective that can be decomposed into subcomponents achievable by individuals or subteams. The team has access to a limited set of resources with which to meet this objective.
- A global objective function quantifies the system designer's preferences over all possible solutions.
- An individual utility function (or cost function) specified for each robot quantifies that robot's preferences for its individual resource usage and contributions towards the team objective given its current state. Evaluating this function cannot require global or perfect information about the state of the team or team objective. Subteam preferences can also be quantified through a combination of individual utilities (or costs).
- A mapping is defined between the team objective function and individual and subteam utilities (or costs). This mapping addresses how the individual production and consumption of resources and individuals' advancement of the team objective affect the overall solution.
- Resources and individual or subteam objectives can be redistributed using a mechanism such as an auction. This mechanism accepts as input teammates' bids, which are computed as a function of their utilities (or costs), and determines an outcome that maximizes the mechanism-controlling agent's utility (or minimizes the cost). In a well-designed mechanism, maximizing the mechanism-controlling agent's utility (or minimizing cost) results in improving the team objective function value.

2.2 Auctions

Auctions are the most common mechanisms used in market-based approaches. In an auction, a set of items is put on the market by an auctioneer in an *announcement* phase, and the participants can make an offer for these items by submitting *bids* to the auctioneer. Once all bids are received or a pre-specified deadline has passed, the auction is then *cleared* in the *winner determination* phase by the auctioneer who decides which items to award and to whom. In robotic applications, the items for sale are typically tasks, roles, or resources. The bid prices reflect the robots' costs or utilities associated with completing a task, satisfying a role, or utilizing a resource.

The simplest kind of auction is a *single-item* auction in which only one item is offered. In such auctions, each participant submits a bid, and the auctioneer awards the item to the highest bidder[1]. Alternatively, the auctioneer retains the item if no bid beats the auctioneer's price (called a *reserve price*). Bids are usually submitted only to the auctioneer; such *sealed-bid* auctions are in contrast to *open-cry* auctions where bidders have the benefit of overhearing the other bids as they are made. There are two common approaches to determining the sale price of the auctioned item. In a *first-price auction*, the sale price is the same as the winning bid. In a *Vickrey auction*, the sale price is the value of the second-highest bid and is intended to motivate truthful bids from the participants. Some multirobot systems have used Vickrey auctions (*e.g.* [49]), though the resulting allocations are equivalent to first-price auctions if the robots are designed to behave truthfully. Wolfstetter provides an excellent introductory survey into single-item auction theory [71].

Combinatorial auctions are more complex: multiple items are offered and each participant can bid on any combination of *bundles* (*i.e.* subsets) of these items. This allows the bidder to explicitly express the synergies between items. In the context of the Mars distributed sensing scenario, a bidder can express the positive synergy between two sites that are close together by bidding only slightly higher for the bundle containing these tasks than for either task individually. To express the negative synergy between two tasks located far from one another, the bid for the bundle would be much higher than the sum of the individual costs of the tasks. In general, there are an exponential number of bundles to consider in a combinatorial auction, which makes bid valuation, communication, and auction clearing intractable if all bundles are considered [55].

In between these two extremes are multi-item auctions in which multiple items are offered but the participants can win at most one item apiece. The maximum number of awards per auction may also be limited. Multi-item auctions can be thought of as a special case of combinatorial auctions where only bundles of cardinality one are considered; bidding and clearing become tractable, but the resulting solutions are generally much less efficient.

2.3 Costs, Utilities, and Valuation

The example scenario in Section 1 compares robots' suitability for tasks in terms of cost. That is, the auction allocates tasks to the robots with the lowest costs for perform-

[1] We will assume utility maximization here; the case of cost minimization is analogous, with awards going to the lowest-cost bidders.

ing them and the overall goal is to minimize some global cost function. As suggested in Section 2.2, in some systems bids are compared based on *utilities*, in which case the highest bids win auctions and the system attempts to maximize the global utility function. Utilities often encapsulate multiple factors, some representing the benefit or expected quality of task execution and others representing cost estimates. Cost estimates can also include diverse factors such as the time taken to compute solutions and the loss of efficiency caused by transitioning between tasks. As an example of a utility calculation, Gerkey and Matarić [25] propose subtracting a robot's *cost* of performing a task from its expected *quality* of completing the task, assuming the units of cost and quality are directly comparable. Thus, utility and cost functions that combine multiple factors often require finding a reasonable set of weights between the different components considered. In general, utility and cost factors might be combined through any arbitrary, possibly nonlinear function.

The process of estimating costs for bid valuation can also be difficult. Though participants in the market may have well-defined cost or utility functions, these functions still rely on having accurate models of the world state and may require computationally expensive operations. For example, the cost to complete the task of driving to a goal site depends on having an accurate map of the environment; however, the robots may be working in an unknown, partially known, or changing environment. When there are multiple goal locations, determining the cost to perform even one task can require solving multiple path planning problems and an instance of the traveling salesman problem (TSP), the latter being $\mathcal{NP}$-hard. Thus heuristics and approximation algorithms are commonly used, implying that bid prices may not always be entirely accurate. Inaccurate bids can result in tasks not being awarded to the robots best able to complete them. In this case re-auctioning tasks can often improve solution quality.

2.4 The Range of Coordination Approaches

The goal in most robotic application domains is to generate optimal solutions in a timely manner. Unfortunately, many multirobot coordination problems are $\mathcal{NP}$-hard. The challenges are compounded by team considerations that include operation in dynamic and uncertain environments, inconsistent information, unreliable and limited communication, interaction with humans, and various system and component failures. A spectrum of coordination approaches has emerged to negotiate these demands.

At one end of the spectrum, fully centralized approaches employ a single agent to coordinate the entire team. In theory, this agent can produce optimal solutions by gathering all relevant information and planning for the entire team. In reality, fully centralized approaches are rarely tractable for large teams (the associated planning problems typically grow exponentially with team size and mission complexity), can suffer from a single point of failure, have high communication demands, and are usually sluggish to respond to local changes. Thus, centralized approaches are most suited for applications involving small teams and static environments or easily available global information.

At the other end of the spectrum, in fully distributed systems, robots rely solely on local knowledge. Such approaches are typically very fast, flexible to change, and robust to failures but can produce highly suboptimal solutions since good local solutions may not necessarily aggregate to a good global solution. Applications where large

teams carry out relatively simple tasks with no strict requirements for efficiency are best served by fully distributed coordination schemes.

A vast majority of coordination approaches have elements that are centralized and distributed and thus reside in the middle of the spectrum. Market-based approaches fall into this hybrid category, and, in some instances, they can opportunistically adapt to dynamic conditions to produce more centralized or more distributed solutions. Market mechanisms can distribute much of the planning and execution over the team and thereby retain the benefits of distributed approaches, including robustness, flexibility, and speed [24, 17]. Auctions quickly and concisely assemble team information in a single location to make decisions about distributing resources; in some cases they provide guarantees of solution quality [55, 40]. Market-based approaches may also incorporate methods of opportunistically coordinating subteams in a centralized manner [18, 34]. Nevertheless, market-based approaches are not without their weaknesses. In domains where fully centralized approaches are feasible, market-based approaches can be more complex to implement and produce poorer solutions. In domains where fully distributed approaches suffice, market-approaches can be unnecessarily complicated in design and have greater communication and computation requirements.

The sections that follow discuss market-based multirobot coordination in greater detail along the dimensions mentioned in the introduction. Each section introduces the topic and its challenges, defines the goals and appropriate evaluation metrics, reviews the relevant literature, and identifies remaining research challenges.

3 Planning

In multirobot teams, planning can be required to coordinate robots to accomplish the team mission. Unfortunately, optimal planning problems for multirobot systems are typically $\mathcal{NP}$-hard [1]. The challenge then is to have tractable planning that produces efficient solutions. Market-based approaches manage this by distributing planning over the entire team to produce solutions quickly. When required or when resources permit, markets can behave in a more centralized fashion and plan over larger portions of the team to improve solution quality. Here, we consider different layers at which planning arises in a multirobot system and how these planning problems are handled by various market-based approaches.

3.1 Related Work

3.1.1 Planning and Task Allocation

Task allocation is the problem of feasibly assigning a set of tasks to a team in a way that optimizes a global objective function. Many special cases of task allocation appear frequently in the literature; here, we offer a general and formal definition that allows us to discuss and compare them.

Definition 1 *Given a set of robots R, let $\mathcal{R} := 2^R$ be the set of all possible robot subteams. An* ***allocation*** *of a set T of tasks to R is a function, $A : T \rightarrow \mathcal{R}$, mapping each task to a subset of robots responsible for completing it. Equivalently, $\mathcal{R}^T$ is the*

set of all possible allocations of the tasks T to the team of robots R. Let $T_r(A)$ be the set of tasks allocated to subteam r in allocation A.

Definition 2 ***The Multirobot Task Allocation Problem****: Given a set of tasks T a set of robots R and a cost function for each subset of robots $r \in \mathcal{R}$ specifying the cost of completing each subset of tasks, $c_r : 2^T \rightarrow \mathbb{R}^+ \cup \{\infty\}$, find the allocation $A^* \in \mathcal{R}^T$ that minimizes a global objective function $C : \mathcal{R}^T \rightarrow \mathbb{R}^+ \cup \{\infty\}$.*

Gerkey and Matarić [25] provide a taxonomy for some variants of the task allocation problem, distinguishing between: single-task (ST) and multi-task (MT) robots; single-robot (SR) and multiple-robot (MR) tasks; and instantaneous (IA) and time-extended (TA) assignment. In instantaneous assignment robots do not plan for future allocations and are only concerned with the one task they are carrying out at the moment or for which they are bidding. In time-extended assignment robots have more information and can come up with longer-term plans involving task sequences or schedules. Definition 2 encompasses each of the types of task allocation in the taxonomy, but in general describes TA task allocation. IA allocation can be represented as a special case where all cost functions map to infinity for any subsets of tasks with cardinality greater than one. Further, if we allow the sets of tasks T and robots R to be time dependent (*i.e.* $T(t)$, $R(t)$) and require the objective function be minimized at every instant of time or over the entire history, then the definition also covers online and dynamic domains where tasks and robots may be added or removed over time (see Section 6). This definition also implies that task allocation is $\mathcal{NP}$-hard in general, as the multi-depot traveling salesman problem is a special case [1].

Market-based approaches distribute the planning required for task allocation through the auction process: each robot or group of robots locally plans the achievement of the offered tasks, computes its costs, and encapsulates the costs in its bids. This process is illustrated in the introduction of this paper for a distributed sensing task on Mars: each robot determined its own cost of visiting different sites. Most existing market-based approaches fall into the SR-ST category in the task allocation taxonomy. Several assume instantaneous assignment (IA) [24, 39, 60, 65, 68], while others allow for time-extended assignment (TA), introducing an additional layer of planning whereby robots sequence [40, 4, 11, 28, 50, 51, 75] or schedule [26, 43, 58] a list of tasks and can therefore explicitly reason about the dependencies between multiple tasks and upcoming commitments. More recently, market-based systems have addressed the allocation of multiple-robot tasks (MR-ST) [29, 45], including human-robot tasks [33]. Definition 2 does not cover all cases of MR tasks: the cost functions for subteams in the definition are assumed to be independent and do not account for the cost dependencies associated with any subteam members participating in another subteam or performing a task on their own. This can be resolved by conditioning the local cost functions on the other allocations of the teammates. Market-based mechanisms for task allocation can also be differentiated as centralized or distributed. Centralized mechanisms have the ability to find optimal solutions (*e.g.* through combinatorial auctions [55, 4]) or provide bounds on solution quality [40], but sometimes require an exponential amount of computation and communication [55]. Distributed mechanisms [11, 28] can act as anytime algorithms and require less computation and communication resources, but

are not guaranteed to find optimal solutions and have no known approximation bounds. *TraderBots* [11] attempts to find a balance between these two approaches by opportunistically allowing "pockets" of centralized optimization to emerge within subgroups of the team when resources permit. In our distributed sensing example, for instance, the team might begin with a suboptimal allocation of sites, perhaps caused by an inaccurate map of the environment resulting in inaccurate bids. At some point during execution (perhaps when map information is more accurate), a robot might find a better distribution of sites for some subset of its teammates. The robot's motivation for group optimization is that it can pocket the cost difference as profit by winning the tasks from the original holders and subcontracting them to the new holders. Simultaneously, this results in a better team solution. Solution quality and scalability aspects of these different approaches are discussed in more detail in Sections 4 and 5 respectively.

Related problems of allocating constrained subtasks, roles, and multiple-robot tasks can have additional planning requirements:

Allocating Constrained Subtasks In many domains, tasks are temporally constrained with respect to one another. They may be partially ordered or may need to start or finish within a common time frame. For instance, consistency may be important in our Mars distributed sensing task, so we might require that samples from particular sites be collected at the same time. In the case of partially ordered tasks, one can use a central allocator to auction only those tasks whose predecessors have been completed [5, 65]. Alternatively, during assignment, robots can incorporate the cost of meeting constraints into their bids [46]. In terms of Definition 2, a violation of constraints can be modeled as infinite values for local cost functions or the global objective function. Constraints can add another dimension to the bid valuation and auction clearing processes and may thus increase computation requirements. Often, robots must also coordinate during execution to reschedule and accommodate team and task changes that have occurred since the initial allocation [26, 43]. In these cases, robots must be able to determine when and how the rescheduling should occur.

Allocating Roles and Instantaneous Assignment In team games such as robotic soccer one usually assigns *positions* such as "primary offense" or "supporting defense" instead of tasks such as "shoot the ball" or "capture a rebound." These positions can be classified as *roles*. More generally a role defines a collection of related actions or behaviors. Indeed, in many domains it is more natural to think of teammates playing roles than completing distinct tasks. In market-based approaches, role allocation can use the same *auction-bid-award* protocol as task allocation. However, robots can usually take on only one role at any given time (SR-ST-IA or MR-ST-IA) and generate bids by evaluating a fitness function that reflects how well its current state matches the requirements of the role. Once allocated, a robot locally plans the execution of actions and behaviors specified by its role. Market-based role allocation has been demonstrated in robot soccer [39, 68] and treasure hunt [12] domains.

Instantaneous assignment (IA) also arises in cases where the tasks being allocated are short-term partial actions that bring the team goal closer to being realized. Examples of instantaneous assignment include allocating *push* actions in a box-pushing

application [24] and assigning waypoint locations (that do not necessarily have to be reached) in an exploration scenario [60].

Allocating Multiple-Robot Tasks Assigning multiple-robot (MR) tasks is another challenging variant of the task allocation problem. In SR task allocation, a robot's bid for a task generally depends only on its own state (including its existing task commitments), and the auctioneer can consider bids individually when making an allocation. In contrast, multiple-robot tasks depend on more than one robot's state and require some amount of subteam planning unless the tasks can be easily decomposed. MR tasks may require robots to tightly coordinate during execution, an issue we consider separately in Section 3.1.3.

Approaches for MR task allocation generally require joint planning for each task followed by an assignment of robots to the subcomponents of this plan. Ideally, in a market-based approach, the agent doing the joint planning can do both the planning and assignment based on minimal information encapsulated in teammates' bids. This ensures that the robots only communicate simple bid prices in order to parallelize local planning and individual cost calculations, as well as to reduce communication bandwidth. For instance, one way to handle MR tasks is to assign each task to a single robot that must recruit other teammates to assist it. Chaimowicz *et al.* [9] and Guerrero and Oliver [29] describe two similar approaches using instantaneous assignment in a foraging and object transportation task. In both, robots explore the environment looking for objects. Upon finding one, a robot uses an auction to recruit help in moving the object to a goal location. These helpers are then committed to assist the leader, but can be preempted for another task if the utility is higher. One disadvantage to this type of approach is that upon discovery of a task, it is automatically assigned to a robot that may not be able to find enough assistance. That robot is then committed to the task rather than being able to perform other useful work. Additionally, while this approach can be applied to most situations where robots are discovering or generating the tasks during execution, it is not clear how it would apply to tasks being introduced by external sources (*e.g.* a human operator). Jones *et al.* [33] introduce an auction mechanism that allows robots to solicit cooperation without committing to the tasks. This mechanism is capable of allocating tasks originating from both online robot discovery and external input. In order to avoid premature commitment, each auction round incorporates a second nested round in order to garner information to compute subteam costs. After the initial auction call, each bidder selects one among several alternatives from a list of "plays" that solve the task. Each play consists of several roles and the bidder chooses the lowest cost role for itself. The rest of the roles are then proposed to the other teammates in a non-binding role auction, allowing the initial bidder to subsequently estimate the full cost of the play and submit a bid price to the original auctioneer. Upon being awarded a task, the winner requests commitments from the role-bidders who accept the award unless they have agreed to perform some other task in the time since they submitted their bids. If any role-bidders reject the award, the task winner rejects the award from the initial auctioneer and the task gets re-auctioned. This ensures that no robot accepts an awarded task without a committment for all relevant roles being filled as estimated during bid valuation. Vig and Adams [70] use a different

mechanism where a task is broken down into a list of required resources or services, then agents representing each kind of resource determine the minimum costs from the robots able to provide some amount of that resource. The resource agents then report these costs to the auctioneer, who then decides if the price is low enough to make the task profitable.

An approach by Lin *et al.* [45] requires that robots submit more complete state information to the auctioneer who then arrives at a plan. Here, the robots send in "capability vectors" describing their current resources and abilities. The auctioneer then offers a "pre-award" to the subgroup that can compete the task for the least cost. The members of the subgroup then communicate amongst themselves and can agree to form a subteam for that task given the auctioneer's price. After notifying the auctioneer, the actual award is finally sent out. Because the approach involves communicating detailed state information rather than bid prices, this approach loses some of the advantages of market-based coordination, such as distributing planning activities and low communication requirements. Nonetheless, centralized planning may be able to better deal with problems containing complex inter-robot constraints. A different approach by MacKenzie [46] does most planning locally, but requires more information to be communicated in order to solve temporal inter-robot constraints centrally. This is done by supplementing each bid price with information on projected task start times, and including multiple such bids for each task. From the submitted bids, the auctioneer then determines which subset of teammates can coordinate in time and space to collectively complete the task most efficiently. The most notable algorithmic difference between this work and that of Lin *et al.* [45] is that although in both some joint planning is done centrally, most of the planning in Mackenzie's approach is still done locally when robots compute their set of bids: the auctioneer is merely finding a feasible plan based on a limited set of options provided by the bidders.

3.1.2 Planning and Task Decomposition

Although many approaches to task allocation assume that a list of primitive or simple tasks is input to the system, a complex mission is often more naturally described at a higher level of abstraction. For example, scientists desiring data about Mars may only be concerned with the general regions from which data is collected and not the precise sites. As illustrated in Figure 1, a mission might be phrased as "capture images that collectively show 50% of crater regions A, B, C, and D." In these cases, multirobot systems must also decompose a mission into subtasks, often making use of well-known planners [5, 7] or domain-specific decomposition algorithms [74].

There are two common approaches to this planning problem. In the *decompose-then-allocate* method, a single agent recursively decomposes the task into simple subtasks which are allocated to the team [7, 65]. In the distributed sensing scenario, this amounts to finding a fixed set of observation sites for all crater regions and then allocating these sites to the team. In the *allocate-then-decompose* method, complex tasks are first allocated to robots then each robot locally decomposes its awarded tasks [5]. This corresponds to assigning entire crater regions to robots and letting each robot choose the sites. It is also possible to include instances of both techniques [26, 5].

By decoupling the decomposition and allocation problems, these approaches do

not consider the complete solution space and may find highly inefficient solutions. In general, one cannot decompose a task optimally without knowing which robots will execute the subtasks, nor can one allocate tasks efficiently without knowing how they will be decomposed. One solution is to simultaneously work on both problems by generalizing tasks to *task trees* and trading these explicitly on the market [74]. In this setting, bidding occurs in two stages. In the first pass, bidders simply valuate nodes in the offered task tree according to the auctioneer's plan. Second, the bidders may come up with new decompositions for any abstract nodes in the tree, and use the price of this new plan within their bid if the cost is lower. By using this bidding procedure together with a specialized task tree auction clearing algorithm, both the costs of allocations and plans can be compared in a single auction mechanism. Experiments in complex task domains demonstrate that using task tree auctions can improve solution quality over the aforementioned two-stage approaches [74].

3.1.3 Planning and Task Execution

Many missions including our example distributed-sensing mission on Mars consist of tasks that can be completed independently by individual robots. Such missions often consist of SR tasks and can usually be achieved by a *loosely-coordinated team* in which robots coordinate during task decomposition and allocation but not during execution. Thus, in these domains, planning how tasks should be executed can be done at an individual level without consideration of teammates' actions and is outside the scope of market-based coordination. (An exception is when robots unexpectedly interfere with each other during execution. Azarm and Schmidt [2] address collision avoidance of independent robots during execution using market-based techniques.)

Another class of problems, however, require *tightly-coordinated teams* in which members continuously coordinate throughout execution. This includes tasks such as the collective manipulation of objects (*e.g.* beams to construct a scaffold) and moving in formation (*e.g.* to safely travel between sites of interest). Tasks requiring tight coordination pose a number of significant challenges. Firstly, these tasks tend not to be easily decomposable (*e.g.* the object the team is transporting cannot be split into pieces and carried independently). Secondly, teammates heavily constrain each other's choice of actions (*e.g.* they must all move in concert together to avoid dropping the object). Thirdly, systems including these types of tasks are rarely fault-tolerant since task success depends on the simultaneous success of multiple teammates. In total, teams must essentially solve a tightly-coupled multirobot planning problem, but cannot easily take advantage of the distributed planning and execution that make loose coordination tractable and robust.

In many cases, the role of market-based approaches in these missions is limited to allocating the MR tasks to a subset of the team [61]. Market-based approaches are rarely used to actually coordinate the activities of the subset of robots tightly coordinating to achieve each MR task because, in practice, this coordination can be achieved with reactive or behavior-based approaches that forgo planning and are less expensive in terms of design, computation, and communication.

Nevertheless, some domains greatly benefit from and even require advanced planning of the coordination between robots and cannot be solved with simpler emergent

approaches. For example, we may want our Mars rovers to always stay in communication contact with a base station; doing so requires that they tightly coordinate over large distances and plan paths to sites with each others' actions in mind. Reactive and behavior-based approaches cannot achieve the required planning; instead, market-based approaches have recently been developed to address these domains. The idea is to exploit small pockets of centralized planning by having robots buy and sell tightly-coupled joint execution plans over the market (as opposed to buying and selling individual tasks) [34]. This technique is related to the idea of opportunistic centralization discussed earlier. Although such approaches have higher communication and computational demands than the market-based approaches designed for loosely-coordinated teams, they have been shown to outperform competing approaches to the same problems.

3.2 Future Challenges

Multirobot systems typically must incorporate multiple types of planning for different aspects of the problems they address. Market-based approaches are currently capable of many types of planning, but several challenges remain. First, there has been limited work in domains with many complex constraints between tasks and domains requiring tight coordination. Second, efficient replanning is crucial to working in uncertain environments and relates closely to issues to be raised in Section 6. Task reallocation can be achieved by peer-to-peer trading and some progress has been made in redecomposing complex tasks in market-based systems [74], but significant work remains in replanning for tightly-coordinated teams. A third important and relevant area of research (for which some initial work has been done [12]) is understanding the formation of subteams and enabling their positive interaction using market-based methods. Finally, market-based approaches need better strategies for making use of multirobot planners and providing alternatives to combinatorial auctions for vetting complex plans in the market.

4 Quality of Solution

In this section, we look at the solution quality of existing market-based allocation algorithms. Essentially, all auction-based solutions aim to optimize some global cost or utility. Here we focus on theoretical and experimental results that demonstrate how well some of these algorithms achieve that goal. As discussed in Section 3, a fundamental optimization problem encountered in market-based multirobot systems is the task allocation problem. Since task allocation is $\mathcal{NP}$-hard, system designers face the challenge of choosing market mechanisms that result in the most efficient solutions within a reasonable amount of time. Various global cost objectives (C in Definition 2) appear in market-based systems depending on the application. Most common are minimizing the sum of individual robots' costs ($C(A) = \sum_{r \in \mathcal{R}} c_r(T_r(A))$) [4, 11, 28, 47, 53] or minimizing the maximum individual cost (makespan, $C(A) = \max_{r \in \mathcal{R}} c_r(T_r(A))$) [50, 43]; although others are possible (*e.g.* minimizing the average overall time to complete each

task [67]). In the Mars distributed sensing example, these global objectives correspond to finding the allocation of sites to robots that results in the least amount of fuel expended (sum of costs) or the task being done in the least amount of time (makespan). While it has been demonstrated that inaccuracies in cost models can affect the quality of solutions obtained [13], the results presented in this section are all developed under the assumption that the robots' cost or utility estimates are accurate.

In this section, we make use of the terms *approximation* and *competitive* ratios. Approximation algorithms find suboptimal solutions to (usually $\mathcal{NP}$-hard) offline problems where the factor relating the two solutions can be bounded by a known factor.

Definition 3 *An algorithm A is a ρ-approximation for a minimization problem iff for all instances of the problem, the cost of the solution, $C(A)$, is guaranteed to be at most ρ times the cost of the optimal solution, $C(\text{OPT})$,* i.e. $C(A) \leq \rho * C(\text{OPT})$. *Similarly, for a maximization problem, an algorithm A is a ρ-approximation iff the cost of the solution is guaranteed to be at least $\frac{1}{\rho}$ times the cost of the optimal solution.*

Online problems are those in which the input is not fully specified ahead of time; for example, the Mars scenario becomes an online problem if the site tasks are being introduced when a partial allocation has already been made or even while the robots are executing the task. For online algorithms, a common measure of the solution quality is the *competitive ratio*.

Definition 4 *An algorithm A is ρ-competitive for an online minimization problem iff for all inputs of the problem, the cost of the online solution, $C(A)$ is guaranteed to be at most ρ times the cost of the offline optimal solution, $C(\text{OPT})$,* i.e. $C(A) \leq \rho * C(\text{OPT})$. *Similarly, for an online maximization problem, an algorithm A is ρ-competitive iff the cost of the solution is guaranteed to be at least $\frac{1}{\rho}$ times the cost of the offline optimal solution.*

4.1 Related Work

As described in Section 3, we distinguish between instantaneous (IA) and time-extended (TA) task allocation. The IA model often arises in cases where tasks require time-indefinite exclusive commitments by robots. Positional roles in robot soccer [39, 68] fall into this category as do "persistent" tasks (*e.g.* in a target tracking scenario, robots must follow a target for as long as possible and therefore do not need to consider taking on future tasks). In both cases, robots will carry out a task or role indefinitely unless it gets reassigned to another robot. At the other extreme are instances where the tasks being assigned are short-lived partial actions that bring the team goal closer to being realized. After some of these actions are carried out, new ones are generated by planning from the observed resulting state. For example, Gerkey and Matarić [24] use an IA approach in a box-pushing domain where with high frequency an auctioneer assigns incremental push actions given the box orientation and the goal location. Similarly, in one approach for an exploration scenario, auctions are used to assign a goal point to each robot on a team, then determine a new set of target points once the first robot reaches its goal (the remaining robots are not obligated to reach their initial

goals) [60]. However, sometimes IA approaches are used for simplicity in lieu of a TA approach; they are easier to implement and do not need computationally expensive task sequencing or scheduling algorithms [25, 29, 7, 10, 36, 59, 44]. In these systems, if there are more tasks than robots, the remaining tasks can be allocated once robots complete their previous assignments. Ignoring dependencies between tasks by using an IA approach should theoretically result in inferior solution quality, but we are not aware of any explicit comparative study.

When IA allocation is appropriate, it has been demonstrated that optimal allocation is possible when there are at least as many robots as tasks, although several existing systems use a 2-approximate greedy solution [25]. Additionally, performance guarantees are not always equivalent for cost- and utility-based systems: the greedy algorithm for the metric online variant of the IA task allocation problem is 3-competitive for utility maximization [25, 37] but scales exponentially with the number of robots for cost minimization [37].

TA sequencing approaches have additional planning and scheduling requirements, but, when appropriate, model the problem more accurately and should produce better results. For example, a combinatorial auction can theoretically result in an optimal allocation if the robots compute and submit bids on all possible combinations of tasks (of which there are an exponential number) [55]. In practice, performance guarantees are sacrificed in order to reduce the computation and communication requirements by considering only a relatively small number of task bundles [4, 47, 31, 16] (details in Section 5).

A simpler centralized mechanism is one in which single tasks are iteratively allocated in multiple auctions until all tasks are assigned [28, 50, 75, 5, 7, 53]. In general, these types of auctions are not guaranteed to find the optimal solution; however, because they require less computation and communication than combinatorial auctions and are easier to implement, they are more prevalent in the literature. Additionally, if costs are considerably uncertain investing the time to produce an optimal solution may not be worthwhile since future assessments are likely to invalidate the optimality of the initial solution.

Tovey *et al.* [67] suggest a hill-climbing heuristic for generating bidding rules for single-task auctions with various global objective functions: essentially each robot bids the difference in *team* cost between the existing allocation and the allocation that would result if the robot wins the offered task. It turns out that for common objective functions robots can calculate these bids using only local information—without requiring knowledge of the states or current assignments of its teammates. This method gives some justification for the typically utilized strategies, as one can derive the commonly used bidding rules encountered in market-based systems: for minimizing total cost (called MINISUM by the authors), this rule advises that bidders should base their bids on the marginal costs of the offered tasks [4, 11, 28, 47, 53]; for makespan minimization (MINIMAX), load balancing can be better achieved if participants instead bid based on their total costs [50]. A third objective function, minimum average latency (MINIAVE) is also suggested, and the generation heuristic suggests that robots should bid the difference between the minimum sum of per target costs with and without the task under consideration. Experimentally, Tovey *et al.* found that the bidding rule derived for each of the three objective functions results in better solutions than the rules

derived for the other objectives. By modeling multirobot problems as vehicle routing problems [8, 27, 42], Lagoudakis *et al.* [40] provide a set of approximation bounds for the same rules and objective functions. They present six results for the three objective functions with two bidding rules derived for each for a total of eighteen approximation bounds. The results prove that bids based on individual *marginal costs* when applied to a sum-of-costs objective results in a 2-approximation, while bids based on individual *total costs* applied to a makespan objective yields an approximation algorithm that scales linearly with the number of robots (which is a worst-case result for *any* makespan algorithm). Examples are also given to show lower bounds for each algorithm, which in some cases show that the approximation bounds are tight.

Task reallocation is also possible by introducing peer-to-peer auctions [11, 28, 43, 53, 14]. In this case, there is some initial allocation, and any robot on the team is capable of holding auctions in order to reallocate tasks to robots that are better suited to perform them. In static environments, distributed trading can improve inefficient initial allocations resulting from the use of faster but suboptimal mechanisms. In unknown or partially known environments where costs constantly change as new observations are made, initial solutions may no longer maintain optimality guarantees or even be reasonably efficient; in such a case the use of peer-to-peer auctions can be used to repair undesirable allocations. Peer-to-peer trading can be viewed as a local search and thus is subject to local optima. Sandholm proves that by using a sufficiently expressive set of contract types (single-task, multi-task, swap, and multi-party), the global minimum can be reached in a finite (but possibly large) number of steps [54], while experiments by Andersson and Sandholm demonstrate that more practical systems that include just single- and multi-task contracts (*e.g.* [16]) find the most efficient solutions given a limited number of rounds [1]. Another interesting result by Vidal [69] shows that by not requiring agents to be purely selfish (*i.e.* some agents may be worse off after some trades) the local search algorithm can circumvent some local optima and in the long run find better solutions. Dias *et al.* [14] look at initializing the team allocation by holding central greedy multi-task auctions (multiple tasks can be awarded in each but at most one task is awarded per robot per auction) before distributed trading begins. They find that increasing the number of tasks awarded per auction can have a negative effect on the resulting solution quality but requires less time (fewer auctions are held) to find a solution. Table 1 gives a summary of the results presented in this section.

4.1.1 Relationship to Operations Research

As has been noted in many publications (*e.g.* [4, 15, 25, 40, 50, 74]), multirobot task allocation problems can be modeled as well-studied problems from the field of operations research (OR), and therefore many of the existing techniques and solutions can be applied to multirobot domains. A simple example is the use of the Optimal Assignment Problem (OAP), which looks at assigning a set of jobs to a set of workers, to model SR-ST-IA task allocation [25].

As mentioned previously, another common manifestation of OR-type problems in multirobot domains is in treating multirobot routing problems as Vehicle Routing Prob-

Table 1: Summary of solution quality results.

Approach	Theoretical guarantees	Experimental results
Combinatorial auctions [4, 47]	Optimal (if all bundles are considered) [55]	Good solutions with limited number of task bundles [4, 47, 31, 16]
Central single task iterated auctions [40, 67]	Approximation bounds for 18 cases (3 objective functions, 6 bidding rules) [40]	Close to optimal results when using the appropriate bidding rules [67]
Central instantaneous assignment (IA) [24, 39]	Optimal possible; commonly used greedy algorithm is a 2-approximation; greedy algorithm for online version is 3-competitive [25]	
Peer-to-peer trading [11, 28, 50, 75, 53]	Optimal solution possible in a finite number of trades with a sufficiently expressive set of contract types [54]	In a limited number of rounds, a combination of single- and multi-task trades outperforms all other combinations of single-task, multi-task, swap, and multi-party contracts [1]; allowing non-individual rational trades can lead to better solutions [69]
Central multi-task auctions followed by peer-to-peer trading [14]		Increasing the maximum number of tasks awarded per multi-task auction results in poorer solution quality [14]

lems (VRP)[2] [27]. Variants of the Traveling Salesman Problem (TSP), a type of VRP, have been mentioned as problem domains in several market-based coordination publications [4, 15, 40, 50, 74]. The multirobot variants of the TSP require a team of robots to visit a set of target points (each point must be visited by at least one robot) while incurring a minimum team travel cost. In the simplest case the relevant combinatorial optimization problems are the k-TSP [21, 52] and the multi-depot TSP [8, 42]. The distinction between the two problems is that in the k-TSP the robots must all start at the same point, whereas in the multi-depot version the robots can start at different locations. Another consideration is whether the robots are required to return to their start

[2]The problem does not have to literally be a transportation-type problem. For example, any problem in which the cost to perform a task depends only on the state of the robot upon completing the previous task is analogous to a Traveling Salesman Problem.

points upon completion of their tasks or not. In the latter case, the problem is called the Traveling Salesman Path Problem (TSPP) [30]. Other variants of the VRP such as consideration of finite vehicle capacities, global vehicle fuel constraints, or precedence constraints between tasks (as in the existence of pickup and delivery points) may well be useful in future multirobot research.

In some cases, existing algorithms from the OR literature may give us insight on auction-based algorithms used in the multirobot community. For example, Gerkey and Matarić [25] demonstrated that a 2-approximate greedy assignment algorithm is being used in several existing multirobot systems (*e.g.* [18, 68]) despite the fact that a polynomial-time optimal algorithm already exists. In other instances, one may find an existing algorithm with desirable theoretical properties that can be readily converted to a distributed auction algorithm. For example, Cerdeira's multi-depot TSP 2-approximation [8] can be easily converted to an auction-based algorithm [41].

4.2 Future Challenges

While some theoretical guarantees for simple auctions are known, future work should address the more complex mechanisms that are present in implemented systems which can include online, multi-task, peer-to-peer, simultaneous, and overlapping auctions as well as task and scheduling constraints. Additionally, solution quality depends on accurate cost and utility measures which may be very challenging to aquire. Although some progress has been made in methods for learning [58] and improving [13] these estimates, further work is required.

5 Scalability

Scalability is an important consideration for any multirobot coordination approach. In general, a system is scalable if it can operate effectively even as the number of inputs or the size of inputs increases arbitrarily. The scalability of a multirobot coordination approach is typically evaluated by its ability to produce efficient solutions as the team size or the task complexity increase. For example, in the Mars distributed sensing scenario, a scalable coordination approach will continue to produce efficient task allocations as the number of robots in the team and the number of sensing tasks assigned to the team increase. Scalability in some market-based approaches may be limited by the computation and communication needs that arise from increasing auction frequency, bid complexity, and planning demands. However, market-based approaches can scale well in applications where the team mission can be decomposed into tasks that can be independently carried out by small subteams.

5.1 Related Work

In Section 4 we highlighted the tradeoffs between scalability and solution quality in market-based systems. In this section, we first elaborate these tradeoffs with an analysis of the computation and communication requirements of the different types of auctions we initially introduced in Section 2. We then describe how market-based ap-

proaches can more effectively negotiate this tradeoff by dynamically responding to the changing demands of the task and adaptively utilizing communication and computation resources.

5.1.1 Computation and Communication Considerations

Single-item auctions are usually computationally feasible and light on communication, but they produce suboptimal solutions. Only one task, resource, or role description is required to be included in the auction call, to be planned for and bid on, and to be considered during auction clearing and awarding. Combinatorial auctions can produce optimal solutions, but can require an exponential amount of computation and communication. This is due to the fact that there are an exponential number of bundles for the bidders to consider making bid valuation, bid submission, and winner determination potentially intractable. Multi-item auctions are also computationally manageable, but produce inferior solutions as compared to single-item and combinatorial auctions. Although there are multiple items included in an auction in this case, bids are only considered for each item independently. Therefore, there is only a linear increase in communications and bid valuation. One advantage of multi-item auctions is that more items are awarded per auction so items can be allocated quickly since less auctions are required. When considering these tradeoffs it is also important to consider the problem domain: for highly uncertain or dynamic environments, it may not be worth spending the time to compute an optimal solution if that solution will constantly be changing as more information is gathered; or, if there are hard real-time constraints, there may not be enough time to compute an exact solution.

Table 2 summarizes the time complexities of the important phases of several auction types. For each protocol, the table lists the maximum number of bid valuations, the computation times of the best-known auction clearing algorithms, and the number of auctions required to allocate all items to the team (if the objective is to offload all items from the initial auctioneer—this may not be the aim in peer-to-peer auctions as the auctioneer may retain some items or some reserve prices may not be met). Table 3 similarly gives a summary of communication costs. Of particular interest in these tables are the exponential expressions for combinatorial auctions in the areas of bid valuation, winner determination, and bid submission. As discussed previously, in order to make combinatorial auctions scalable, the common approach is to limit the number of bundles that are considered during bid valuation. This in turn reduces the number of bids that need to be communicated and also allows the auction to be cleared quickly in practice. Indeed, Sandholm's optimal clearing algorithm, CABOB [55], relies on a sparse bid set in order to find solutions quickly (although the resulting allocation is still likely to be suboptimal given that not all item bundles are considered). The number of bundles can be reduced by the auctioneer, the bidders, or both. The auctioneer may offer only a limited set of bundles, for example, by grouping items that may require similar resources [31], or by exploiting hierarchical problem structure to form the bundles [74]. On the bidders' side heuristic clustering algorithms (*e.g.* nearest neighbor) are often used [4, 16]. Another strategy is to consider only those bundles that are smaller than a given size [4, 47]. Berhault *et al.* [4] compare four clustering algorithms for goal point tasks and find that one based on repeated graph cuts outperforms

a nearest-neighbor algorithm and two algorithms based on limiting cluster size.

Table 2: Comparison of time complexities of various auction types. n is the number of items, r is the number of bidders, b is the number of bids, and $m \leq r$ is the maximum number of awards per auction (for multi-item auctions). v and V represent the (domain-dependent) amount of time required to perform a valuation for a single item (v) or set of items (V).

Auction type	Bid valuation	Winner determination	Number of auctions
Single-item	v	$O(r)$	n
Multi-item (greedy)	$O(n \cdot v)$	$O(n \cdot r \cdot m)$	$\lceil n/m \rceil$
Multi-item (optimal)	$O(n \cdot v)$	$O(r \cdot n^2)$ [25]	$\lceil n/m \rceil$
Combinatorial	$O(2^n \cdot V)$	$O((b+n)^n)$ [55]	1

Table 3: Comparison of communication complexities of various auction types. Notation is the same as in Table 2. This assumes constant message space for each task description.

Auction type	Auction call	Bid submission	Award	Award (+ losers)
Single-item	$O(r)$	$O(r)$	$O(1)$	$O(r)$
Multi-item	$O(r \cdot n)$	$O(r \cdot n)$	$O(m)$	$O(r)$
Combinatorial	$O(r \cdot n)$	$O(r \cdot 2^n)$	$O(n)$	$O(r+n)$

Bid valuation itself may be computationally expensive as the process usually involves some amount of local planning. That is, the expressions v and V in Table 2 may in reality represent the running time of algorithms that must solve a difficult or even $\mathcal{NP}$-hard problem in order to estimate costs [4, 11, 28, 50, 53]. Additionally, expensive task decompositions may be required at the bidding stage [74]. In the Mars exploration scenario, bid valuation can be expensive since evaluating the cost of a set of target points may require solving two traveling salesman problems (TSP): one to determine the cost of the rover's path after including the offered tasks, and one to find the cost without the new tasks. Heuristics and approximation algorithms can help deal with the $\mathcal{NP}$-hard problems (*e.g.* TSP approximation algorithms for task sequencing [18]), although when there are many items to consider simultaneously—either from auctions offering many items [4, 11, 28, 43, 74, 47] or from multiple robots holding auctions simultaneously [11, 74]—bidders can still be overburdened with valuation problems. As a result, system designers must ensure that bidders are able to meet auction deadlines and do not tax their processors to the point of compromising real-time requirements.

As mentioned in Section 4, bid valuation computation can be less costly for instantaneous allocation problems (IA) as there is no need for an explicit task sequencing or scheduling algorithm. However, if the problem is inherently a time-extended one (TA), ignoring scheduling in general results in poorer quality solutions.

Busquets and Simmons [6] offer an offline learning approach to improve the scalability of multi-item auctions: by keeping track of histories, the auctioneer can reduce the number of tasks offered if it believes it is unlikely to get any bids for some of them, while a bidder can reduce the number of bids submitted by removing those for tasks it believes it is unlikely to win. Reducing the number of offered tasks in a multi-item auction decreases the amount of communication in the auction call and bid submission, plus the amount of computation in bid valuation and winner determination. Reducing the number of bids decreases bid submission communications and winner determination computation. In this approach, each bidder computes its bid price on a task and then stochastically decides whether to submit its bid to the auctioneer. The probability of submission is equal to the proportion of previous bids at or below the current bid value that successfully won their auctions. Similarly, an auctioneer stochastically decides to offer a task based on the proportion of past bids that have exceeded the reserve price. In a variation of the Mars distributed sensing scenario, this learning algorithm results in significant scalability improvements by reducing the number of messages communicated and the number of task valuations performed.

A scalability comparison of market-based, behavior-based and centralized approaches is presented by Dias and Stentz [11] on a distributed sensing task in which the goal is to visit the last observation site in as little time as possible (*i.e.* a makespan objective). Simulation experiments demonstrate that the market approach can provide significantly higher-quality solutions than behavior-based approaches while using significantly less computation time than centralized approaches. Xu *et al.* [72] compare market- and token-based coordination approaches, and further introduce a hybrid approach in which auction calls (tokens) are sent only to the agents most likely to submit the best bids. This requires agents to maintain a model of the team state in order to make intelligent token-routing decisions. Their experiments show that the hybrid approach offers an interesting tradeoff, producing slightly less efficient solutions than the full auction approach, but with much reduced communications (slightly worse than the token approach).

Tilley and Williams [66] study several scalability effects in a simple parts manufacturing system. In this application, agents bid for machining tasks for which execution times do not depend on task order and a simple scheduling heuristic is used: each part is processed in the order it is allocated. One method they employ to reduce task evaluation time is to limit the number of tasks considered at any time by each machine; that is, if the number of current tasks plus outstanding bids exceeds some small constant that machine ignores all incoming auctions. Decreasing this limit decreases the average number of auction participants and increases the number of auctions that go unbid for (thus decreasing the solution quality), but also lowers the required auction deadlines (time required for all bids to be submitted). They also conclude that the time required to valuate a task has a greater effect on required auction deadlines than does system task load, and that task valuation times should be much smaller than task inter-arrival time for efficient system performance.

5.1.2 Selective or Opportunistic Centralization

Centralized planning has the potential to produce optimal team solutions, but most centralized algorithms have complexity exponential in the number of robots and quickly become intractable for more than a few team members. Nevertheless, market-based approaches can selectively use centralized planning to scalably improve solutions over subparts of the team. For example, in the *TraderBots* architecture, a leader robot may replan the allocation of a subset of tasks contracted to some subset of the team when time and computation resources permit. If this leader discovers a reallocation that is better than the current allocation, it can purchase the tasks from their current holders and subcontract them according to its new allocation. The difference in value between the two allocations minus the required payoffs can be pocketed by the leader as profit. This is similar to the coalition formation problem addressed by Sandholm and Lesser [56] in which agents representing trucking depots could form coalitions to pool their delivery tasks together in order to find more efficient solutions for both parties.

The Hoplites framework uses a similar approach to solve tightly-coupled problems in which robots continuously constrain each other's actions (*e.g.* to remain in line-of-sight contact). Robots begin with a simple coordination strategy which is light on communication and computation and allows teammates to iteratively respond to each other's actions without directly influencing them. When this traps robots in local minima, a more powerful coordination mechanism improves solutions by enabling robots to influence each other directly by purchasing each other's participation in complex plans over the market. This consumes more resources but is also capable of producing much better solutions. Hoplites selectively injects pockets of complex centralized coordination into the system only when necessary; as a result, it provides improvements in solution quality while remaining computationally competitive with much less sophisticated approaches.

5.2 Future Challenges

While much is known theoretically about the scalability of various auction mechanisms, market-based approaches have yet to be implemented on teams of more than a few robots. Larger teams have been used in simulation however; for example, Xu *et al.* [72] consider teams of 100 agents. Further challenges exist in improving opportunistically centralized approaches' means of selecting task clusters and team members to reduce unnecessary computation. The challenge of dealing with limited computation when faced with an excess of solicited bid valuations is also largely unaddressed. The recent work of Busquets and Simmons [6] has made some progress in this area.

6 Dynamic Events and Environments

Operations in dynamic and uncertain environments pose a variety of challenges to team coordination such as ensuring graceful degradation of solution quality with failures, enabling team functionality despite imperfect and uncertain information, maintaining effective response speed to dynamic events, and accommodating evolving conditions

and constraints. Successful team operations in dynamic and uncertain environments is therefore highly dependent on achieving robustness. Benchmarks for the robustness of a coordination approach must take into account the diversity of failures the team can accommodate, the requirements for quantity, quality, and certainty in information, the team's response speed to dynamic events, the fluidity of the team (that is, the ability of the team to accommodate the addition of new members and the loss of current members), and the overall solution quality produced by the team in the face of dynamic events. In this section we examine the different ways in which market-based multirobot coordination approaches to date deal with these conditions.

6.1 Related Work

The application of market-based coordination to dynamic domains is in its early stages. However, several research groups have already made early contributions in this area. We discuss related work here in three categories. The first section on robustness and fluidity deals with a team's ability to handle agents dynamically joining the team. We then discuss online tasks which require a team to adapt to tasks being added and removed from the mission during task execution. Finally, we explore related work that deals with a team's ability to operate with limited information.

6.1.1 Robustness and Fluidity

Robotic systems are often complex and their physical interactions with the environment make them highly prone to failures. Thus, a successful coordination approach must gracefully degrade solution quality when failures occur. For example, in the Mars distributed sensing scenario, it is likely that one or more of the rovers will suffer some form of hardware or software failure. A rover can wander outside the range of communication from other rovers, damage a scientific instrument when performing a sensing task, or become completely disabled due to a rockslide. A successful coordination approach will allow the team to accomplish the mission despite these failures if the remaining resources of the team allows for mission completion. Three principal categories of faults that coordination approaches must consider are communication failures, partial malfunctions, and robot death [17]. A variety of strategies are employed by market-based approaches to handle these failures. These strategies are explored in the following sections.

Teams can often perform more effectively if teammates can communicate [75]. However, communication failures occur in a variety of domains and range from occasional loss of messages to loss of all communication. In the *TraderBots* approach, if the team is informed of a common task, consequent disruptions in communication are gracefully handled by using opportunistic auctioning solely for improving solution quality [17, 75]. Sheng *et al.* [59] also consider disruptions due to limited-range communication in the context of multirobot exploration. Here, utility functions incorporate a term that depends on inter-robot distance in an attempt to actively keep the robots closer together and limit the amount of resulting communication loss. As with communication disruptions, partial malfunctions limit a robot's capability but retain the robot's planning ability. *TraderBots* employs active reasoning about failed resources to allow

robots to reallocate tasks that they can no longer complete due to malfunctions [17]. Similarly, MURDOCH relies on monitoring progress in short-duration tasks to detect and respond to faults [24]. In the case of robot death, the affected robot cannot aid in the recovery process. However, robots can monitor progress or heartbeats of teammates and re-auction tasks previously assigned to the dead robot [24, 17]. A further improvement is to allow repair of malfunctioning robots and enable their return to the team. In order to accomplish this, a coordination approach must accommodate both the exit of malfunctioning robots and the entrance of the repaired robots. Bererton *et al.* demonstrate reasoning about assisting malfunctioning robots and towing of disabled robots to a base station for repair [3], while Dias *et al.* [17, 18] and Gerkey and Matarić [24] demonstrate both the ability to accommodate the loss of a robot and the re-entry of a repaired robot or the entry of a new robot to the team.

When introducing robustness under dynamic and uncertain conditions, coordination approaches must wrestle with several tradeoffs. For example, if detecting robot death relies on monitoring a heartbeat, a robot is presumed dead if its heartbeat is not received by its teammates within a pre-determined interval. However, if this interval is too large, the time to detect and respond to a failure is increased and the solution quality can degrade as a result. Instead, if the interval is too short solution quality can still degrade due to false positives that can arise if the robot temporarily drops out of communication range.

Another important design consideration when operating under dynamic conditions is choosing auction deadlines. A common practice is to specify a deadline at which time the auction is cleared regardless of whether or not all bids have been submitted. This is required in realistic situations where communications are imperfect, or team members may fail, enter, or leave the team. This may also be necessary in situations where bidders do not bid on all offered tasks, which can be the case for a number of reasons. First, a robot may not be able to perform an offered task due to limited capabilities. Second, a robot may not be able to perform a task for less than a reserve price or for less than a known competitor's price; for example, another bidder's bid in an open-cry auction. Finally, bidders may choose not to bid on tasks they are unlikely to win. The tradeoff to consider with respect to auction deadlines is that the deadline must be long enough for all interested participants to submit their bids, while not being overly long such that the slow turnaround time for an auction results in inefficiency due to an insufficient amount of task reallocation.

Determining an appropriate frequency for running auctions is another design consideration that involves a tradeoff. If auctioning happens too frequently, communication networks can get overloaded and robots may spend more time running their own auctions than effectively participating in other on-going auctions. However, if auctions are not held with sufficient frequency, the team may be slow to respond to dynamic conditions. Effectively navigating these tradeoffs and their implications remains an open research challenge.

6.1.2 Online Tasks

In many dynamic application domains, the demands on the robotic system can change during operation. Operators of a multirobot system may submit new tasks or alter

or cancel existing tasks during operation [18]. Alternatively, robots may generate new tasks during execution as they observe new information about their surroundings. In the Mars scenario, a scientist may choose to add new sites to explore or eliminate existing sites while reviewing incoming data, or the robots themselves may have the capacity to make such decisions. Market-based approaches can often seamlessly incorporate online tasks by auctioning new tasks as they are introduced by an operator [11] or as they become available due to the completion of preceding tasks [24, 60, 5, 7]. In some cases new tasks can be generated by the robots themselves and inserted into their plans to be executed or subsequently traded [18, 51, 75, 29]. The solution quality of some simple online auction-based task allocation algorithms are given in Section 4. In general, online versions of optimization problems are more difficult to solve than their offline counterparts.

6.1.3 Uncertainty

Most real-world multirobot applications require operation with only partial or changing information about the environment, the team, and the task. For instance, in the Mars example, it is likely that the rovers will not have access to a complete detailed map of their environment before they begin their mission. Fortunately, market-based approaches have few requirements for prior information and can accommodate new information through frequent auctioning of tasks and resources. The *TraderBots* approach demonstrates that robots can execute tasks with no initial map information and dynamically reallocate tasks when new map information is gathered [18, 75].

6.2 Future Challenges

While much can be done to improve the operation of market-based approaches in dynamic environments, a few key challenges are paramount. Effective information sharing among team members in market-based approaches is one necessary area of research. If a robot discovers a task is expensive because of new environmental information it has gathered, it can potentially allocate the task to another robot that does not have that information. This robot will then try to execute the task until it, too, perceives the new information; then it tries to allocate the task to another robot. This can continue until all robots have attempted to perform the task, causing tremendous inefficiency.

Characterizing the ability of both the individual robots and the team to respond quickly to dynamic conditions using market-based coordination approaches is another important challenge. The authors are not aware of any study of individual or team response speed for any market-based multirobot coordination approach. Other challenges for improving robustness are developing more sophisticated methods for cooperative handling of partial malfunctions and repairs, evaluating robustness to a variety of failures, incorporating contract breaches with appropriate penalties, and incorporating sliding autonomy into market-based approaches to allow robots to request assistance when appropriate.

Formalizing design principals to achieve guaranteed performance within specified bounds despite dynamic conditions is also an important research challenge for market-based team coordination. Specifically, we need more principled methods for navi-

gating the various tradeoffs encountered when designing a market-based coordination approach that must perform effectively under dynamic conditions.

7 Heterogeneous Teams

A team is heterogeneous if not all of its members are equally capable of performing all the tasks (*e.g.* because of hardware or software differences) or if its members play different roles (*e.g.* in team games where robots play different positions). In contrast, the members of a homogeneous team have identical skills or are generalists that can perform all necessary tasks. Heterogeneity is highly advantageous for several reasons. First, complex missions often have many different functional requirements and can be achieved more effectively by a team of specialists rather than by a team of generalists that perhaps cannot perform any single function very well. For example, in our Mars scenario we may want high-resolution images taken, rock samples collected, and core samples taken from the ground. A complex mission such as this is often better achieved with robots that specialize in particular tasks: some robots can take samples from both rocks and the ground while others only capture images. Second, it is often more practical to design robots that specialize in only a small set of skills than to design robots that are capable of all skills. Indeed, in many domains, it may be infeasible to construct robots that can do everything, for example because of limitations in budget, form factor, or on-board power. Third, by being able to coordinate heterogeneous teams, we can reuse robots across multiple applications. Ultimately, a truly heterogenous team is not limited to robots in its membership. Human-robot-agent teams must operate seemlessly and efficiently in several application domains. Although heterogeneous team coordination is challenging, a successful approach should accommodate any team composition.

7.1 Related Work

In a heterogeneous team, robots have different abilities to perform different tasks. Task and role allocation in heterogeneous teams becomes challenging because it requires reasoning about and comparing different robots' capabilities. Market-based approaches are well suited to meet this challenge because auctions can simplify the problem of reasoning about team skills. When a task or role is auctioned, each robot's bid encapsulates its ability to complete the task in terms of resource usage, the estimated solution quality afforded by these resources, or even the opportunity cost of forgoing other tasks [58]. Additionally, robots can abstain from bidding on tasks for which they do not have sufficient resources, thereby reducing the computational burden for both bidder and auctioneer [24, 31]. The bids can also encode the solution quality afforded by each member's resources, and even the opportunity cost of forgoing other tasks [58]. The auctioneer can award roles or tasks to team members according to the best bid, without requiring knowledge of individual capabilities. Thus, market-based approaches only require each team member to recognize its own skills and resources but not necessarily those of teammates. However, auctions introduce a new difficulty: it is not always clear how to compute and compare the cost of performing a task between different types of

robots who may perform the task in very different ways. One idea is to allow robots to swap tasks directly whenever such a trade results in a mutually beneficial outcome [28]. This circumvents the pricing problem but also severely restricts the number of possible solutions. Thus, the pricing problem remains an open research challenge.

Market-based coordination of heterogeneous teams has been demonstrated on physical robots in a few applications: in automated assembly using three robots with very different physical configurations and capabilities [61]; in box-pushing, where a resource-addressed messaging protocol allows robots to determine in which auctions they should participate [24]; and in treasure hunt, where a human, two Pioneer IIDX robots with laser scanners, and a Segway RMP robot with vision sensing cooperate to seek objects of interest in an unknown environment [33]. In simulation, market-coordinated robots with different science instruments characterize multiple rock types in a space application similar to our Mars scenario [11, 58], and market-based role allocation mechanisms have appeared in heterogeneous robot soccer teams [39, 23]. In work by Lin *et al.* [45] each robot submits a list of its capabilities to the auctioneer who then comes up with a joint plan and sends an offer to the best subteam found who can then accept or reject the task at that price.

7.2 Future Challenges

Ultimately, coordination approaches must accommodate three levels of heterogeneity: heterogeneous robot teams, human-robot teams, and highly heterogeneous teams of humans, robots, and other agents. Future research challenges include modeling human preferences using appropriate reward functions, developing techniques for consistently computing different robots' costs for completing tasks, enabling *pickup teams*, *i.e.* dynamically formed heterogeneous teams where little may be known a priori about the task, the robots, or the environments [12], and addressing the challenges of human-robot teams where tasks are understandable to humans and robots and both participate in task allocation and execution [12].

8 Learning and Adaptation

While a generalized system is useful, its application to specific domains usually requires some adaptation. Online opportunistic adaptation is therefore a highly relevant and useful feature. However, in dynamic environments where teams can fluidly change in size, where interaction strategies can be continuously modified, and where external conditions can be unexpectedly changed, *a priori* definitions of best trading and coordination strategies can be very difficult, and sometimes impossible. Consequently, the robots require not only the ability to quickly adapt their behaviour in response to dynamic events and to changes in the other agents' behaviour, but also the ability to determine when and how this adaptation should take place. Hence, the integration of learning techniques can be a very powerful feature.

8.1 Related Work

The application of learning techniques in market-based coordination is currently at a very early stage. One big debate is whether learning should be applied at the team level or at the individual level or some combination of the two. Another important question to be answered is how to deal with team interactions – should other agents be dealt with as environmental factors or should they be dealt with in a special way? Oliveira *et al.* [48] present a detailed dicussion of the issues relevant to the application of learning in dynamic markets. The role that learning can play in market-based multirobot coordination is also discussed briefly by Stentz and Dias [64].

The authors are unaware of any learning techniques implemented on a team of physical robots coordinated using a market-based approach. However, publications are starting to emerge in the application of learning techniques for market-based coordination of simulated robot teams. Notably, learning techniques are applied to learn bidding strategies in dynamic markets [48], opportunity costs in a simulated distributed sensing task [58], and role assignment [39] and bidding strategies [22] in simulated robot soccer.

8.2 Future Challenges

The application of relevant learning techniques to market-based coordination of robot teams is a wide open research area with tremendous potential for improving team performance in dynamic environments, reducing the requirement for accuracy in cost estimation and *a priori* knowledge, and enabling easy portability to different domains and environments.

9 Practical Considerations

Many practical considerations affect the overall impact of coordination approaches. Approaches that are general and applicable to a variety of domains are more useful in the real world. A general approach will be flexible across application domains and extensible to enable portability and easy enhancement of functionality. Other important considerations include implementation guidelines for different domains and useful comparisons of different approaches to guide the selection of the most effective coordination scheme for a given application.

9.1 Related Work

Here, we discuss several groups of work that are related to the idea of generality and practical impact.

9.1.1 Flexibility

Since different applications will have different requirements, a widely applicable coordination approach will need to be easily configurable for the different problems it proposes to solve. Instructions and advice on how to reconfigure the mechanism for

different applications will also be useful. Identifying important parameters that need to be changed based on the application requirements, instructions on how to change them, identifying components of the mechanism that need to be added/changed based on application requirements, and instructions on how to make these alterations are all important elements of a successful coordination mechanism. A further bonus will be well-designed user interfaces and tools that allow plug-and-play alterations to the coordination mechanism and automated methods for parameter tuning. The authors are aware of only three market-based approaches, MURDOCH [24], Hoplites [34, 35] and *TraderBots* [11, 75, 73], that have been demonstrated in more than one application. However, there is much that still needs to be done in terms of providing a flexible market-based multirobot coordination approach.

9.1.2 Extensibility

The ability to easily add and remove functionality is a key characteristic to building a generalized system that can evolve with the needs of the different applications. A common approach to incorporating extensibility is to build the system in a modular fashion so that different modules can be altered or replaced relatively easily according to the requirements of the specific application. In market-based approaches, it is best to modularize and isolate cost and reward functions as much as possible from task and role specifications, communication protocols, and task executives.

9.1.3 Implementation

As with any claim, a proven implementation is most convincing. Moreover, successful implementation of a coordination mechanism on a robotic system requires discovering and solving many details that are not always apparent in theory, simulation and software systems. Finally, implementation of an approach on many different platforms in a variety of application domains provides valuable insights and guidelines on how to design and implement different components of the approach in a extensible and flexible manner. Although several have been implemented on physical robot teams (*e.g.* [11, 24, 75, 73, 34, 61, 60]), market-based approaches have yet to be proven in a wide variety of domains.

9.1.4 Comparisons

Comparisons are important to provide guidelines on how to evaluate different coordination approaches when deciding which approach is best for a given application. However, comparing different coordination approaches is a highly challenging endeavor since many considerations need to be addressed.

Some of the challenges in providing a comparative framework for coordination approaches are explored by Gerkey and Matarić [25] who provide an initial framework for evaluating task allocation schemes in terms of complexity and optimality, showing that market-based methods perform favorably in terms of computation and communication requirements. Dias and Stentz [11] compare a centralized optimal approach,

a distributed behavioral approach, and a market-based approach, evaluated in a distributed sensing scenario. Simulation results compare the three approaches in scalability and heterogenity, and show that all three approaces perform well with heterogeneous teams, and the market method performs best overall in scalability. Rabideau *et al.* [50] also conduct a similar comparative study between a centralized planner that does not guarantee optimality, a distributed planner, and a single-task auction approach. They conclude that the auction approach performs best but takes up the most CPU cycles.

More recently, Xu *et al.* [72] compare a market-based, a token-based, and a market-token-hybrid approach. They find that a market-based solution finds more efficient solutions than a token-based one, but requires more communication. The hybrid approach falls somewhere in between the pure market and token approaches in both areas. Kalra and Martinoli [36] compare the performance of a market-based and a threshold-based approach to IA task allocation along several dimensions. They find that the accuracy of task and robot state information can play an important role in determining the relative effectiveness of these approaches. In sum, market-based approaches have performed well in comparative studies; nevertheless, these studies are fairly limited and broader studies are in high demand.

9.2 Future Challenges

Market-based multirobot coordination approaches have only been implemented and tested in a few application domains to date. Thus, understanding and implementing generality in market-based approaches still requires significant work. However, the growing popularity of market-based methods for coordinating robot teams will be a large contributing factor to inspiring generality in this research area.

10 Conclusions and Future Directions

The vision that drives research in multirobot systems is that teams of robots will inevitably be an integral part of our future. To realize this vision, robots must be capable of executing complex tasks as a team. While many multirobot coordination approaches have been proposed by the research community, market-based approaches in particular have been proven effective; their resulting increase in popularity over the past few years warrants a survey of the field. many multirobot coordination approaches have been proposed by the research community, market-based approaches in particular have been proven effective and as a result have grown in popularity over the past few years, consequently warranting a survey of the field. We address this need by providing the first survey of the state of the art in market-based multirobot coordination approaches with three contributions to the multirobot literature: a tutorial on market-based multirobot coordination approaches, a review and analysis of the relevant literature, and a discussion of remaining challenges in this research area.

The existing work in market-based multirobot coordination ranges from theoretical formulations to conceptual design frameworks to implementations in simulation and on physical robot teams. The chosen application domains span a wide range and include distributed sensing, mapping, exploration, surveillance, perimeter sweeping, assembly,

box-pushing, reconnaissance, soccer, and treasure hunt. However, this is still a relatively new area of research, and hence many research challenges still remain. Here, we discuss some of the overall challenges in the field.

A first important need is a principled formalization of market-based coordination approaches. Much research is needed to further our understanding of how components such as cost and reward functions, bidding strategies, and auction clearing mechanisms can be designed, implemented, and used effectively in different multirobot application domains. Understanding the tradeoff between solution quality and scalability when designing and implementing coordination mechanisms is also important.

Additionally, much work still remains in rigorously comparing different coordination approaches to enable users to select an appropriate approach given a particular multirobot scenario. First, a relevant set of benchmarks must be defined for effective comparison of different coordination approaches. Second, comparisons must occur between a wider range of approaches, with a greater variety of tasks and team compositions, and across a broader set of metrics. Third, approaches must be compared on physical robot teams as well as in simulation to validate their performance under real-world conditions.

A final challenge is to demonstrate long-term, reliable, and robust operation of larger robot teams in the real world. This will require simultaneous use of many of the techniques discussed in this paper, including learning and adaptation, scalability, fault detection and tolerance, handling uncertainty, and enabling heterogeneity. While many of these features have been demonstrated on a small scale, holistic implementations covering the spectrum of these features are now required.

Despite the many challenges ahead, market-based techniques are proving to be versatile and powerful coordination schemes for groups of robots executing complex tasks as part of a team. Different application requirements and tradeoffs in implementation make it difficult to construct a single market-based approach that can be successful in all domains. Nevertheless, a well-designed market-based approach with sufficient plug-and-play options for manually or automatically altering different tradeoffs can be successful in a wide range of applications. And, with further research, market-based approaches promise to significantly further our vision of robots playing an integral role in human life.

Acknowledgments

This work is sponsored in part by the Boeing Company Grant CMU-BA-GTA-1, in part by the U.S. Army Research Laboratory, under contract **Robotics Collaborative Technology Alliance** (contract number DAAD19-01-2-0012), and in part by the Qatar Foundation for Education, Science and Community Development. 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 Boeing Company, the Army Research Laboratory, the U.S. Government, or the Qatar Foundation.

References

[1] M. Andersson and T. Sandholm. Contract type sequencing for reallocative negotiation. In *International Conference on Distributed Computing Systems*, 2000.

[2] K. Azarm and G. Schmidt. A decentralized approach for the conflict free motion of multiple mobile robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1996.

[3] C. Bererton, G. Gordon, S. Thrun, and P. Khosla. Auction mechanism design for multi-robot coordination. In *Advances in Neural Information Processing Systems*, 2003.

[4] M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghraby, P. Griffin, and A. Kleywegt. Robot exploration with combinatorial auctions. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.

[5] S. S. C. Botelho and R. Alami. M+: A scheme for multi-robot cooperation through negotiated task allocation and achievement. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.

[6] D. Busquets and R. Simmons. Learning when to auction and when to bid. In *Proceedings of the International Symposium on Distributed Autonomous Robotic Systems (DARS)*, 2006.

[7] P. Caloud, W. Choi, J. C. Latombe, C. L. Pape, and M. Yim. Indoor automation with many mobile robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1990.

[8] J. O. Cerdeira. The multi-depot traveling salesman problem. *Investigação Operacional*, 12(2), 1992.

[9] L. Chaimowicz, M. F. M. Campos, and V. Kumar. Dynamic role assignment for cooperative robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2002.

[10] P. Chandler and M. Pachter. Hierarchical control for autonomous teams. In *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2001.

[11] M. B. Dias. *TraderBots: A New Paradigm for Robust and Efficient Multirobot Coordination in Dynamic Environments*. PhD thesis, Robotics Institute, Carnegie Mellon University, January 2004.

[12] M. B. Dias, B. Browning, M. M. Veloso, and A. Stentz. Dynamic heterogenous robot teams engaged in adversarial tasks. Technical Report CMU-RI-TR-05-14, Robotics Institute, Carnegie Mellon University, 2005.

[13] M. B. Dias, B. Ghanem, and A. Stentz. Improving cost estimation in market-based coordination of a distributed sensing task. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2005.

[14] M. B. Dias, D. Goldberg, and A. Stentz. Market-based multirobot coordination for complex space applications. In *the 7th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2003.

[15] M. B. Dias and A. Stentz. A free market architecture for distributed control of a multirobot system. In *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS)*, 2000.

[16] M. B. Dias and A. Stentz. Opportunistic optimization for market-based multirobot control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.

[17] M. B. Dias, M. Zinck, R. Zlot, and A. Stentz. Robust multirobot coordination in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

[18] M. B. Dias, R. Zlot, M. Zinck, J. P. Gonzalez, and A. Stentz. A versatile implementation of the *TraderBots* approach to multirobot coordination. In *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS)*, 2004.

[19] M. B. Dias, R. M. Zlot, N. Kalra, and A. T. Stentz. Market-based multirobot coordination: A survey and analysis. Technical Report CMU-RI-TR-05-13, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, April 2005.

[20] D. J. Farber and K. C. Larson. The structure of a distributed computing system – software. In *Proceedings of the Symposium on Computer-Communications Networks and Teletraffic*, 1972.

[21] G. N. Fredrickson, M. S. Hecht, and C. E. Kim. Approximation algorithms for some vehicle routing problems. *SIAM Journal on Computing*, 7(2), 1978.

[22] V. Frias-Martinez and E. Sklar. A team-based co-evolutionary approach to multi agent learning. In *In Proceedings of the 2004 AAMAS Workshop on Learning and Evolution in Agent Based Systems*, 2004.

[23] V. Frias-Martinez, E. Sklar, and S. Parsons. Exploring auction mechanisms for role assignment in teams of autonomous robots. In *Proceedings of the RoboCup Symposium*, 2004.

[24] B. P. Gerkey and M. J. Matarić. Sold!: Auction methods for multi-robot control. *IEEE Transactions on Robotics and Automation Special Issue on Multi-Robot Systems*, 18(5), 2002.

[25] B. P. Gerkey and M. J. Mataric. A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research*, 23(9), 2004.

[26] D. Goldberg, V. Cicirello, M. B. Dias, R. Simmons, S. Smith, and A. Stentz. Market-based multi-robot planning in a distributed layered architecture. In A. Schultz, L. Parker, and F. Schneider, editors, *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings from the 2003 International Workshop on Multi-Robot Systems*, volume 2. Kluwer Academic Publishers, 2003.

[27] B. L. Golden and A. A. Assad, editors. *Vehicle Routing: Methods and Studies*, volume 16 of *Studies in Management Science and Systems*. Elsevier Science Publishers, Amsterdam, 1988.

[28] M. Golfarelli, D. Maio, and S. Rizzi. A task-swap negotiation protocol based on the contract net paradigm. Technical Report 005-97, CSITE (Research Center for Informatics and Telecommunication Systems), University of Bologna, 1997.

[29] J. Guerrero and G. Oliver. Multi-robot task allocation strategies using auction-like mechanisms. In *Sixth Congress of the Catalan Association for Artificial Intelligence (CCIA)*, 2003.

[30] J. A. Hoogeveen. Analysis of christofides' heuristic: Some paths are more difficult than cycles. *Operations Research Letters*, 10, 1991.

[31] L. Hunsberger and B. J. Grosz. A combinatorial auction for collaborative planning. In *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS)*, 2000.

[32] M. Jia, G. Zhou, and Z. Chen. Arena – a centralized framework for multi-robot exploration, 2004. submitted to Robotics and Autonomous Systems.

[33] E. G. Jones, B. Browning, M. B. Dias, B. Argall, M. Veloso, and A. Stentz. Dynamically formed heterogeneous robot teams performing tightly-coupled tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2006.

[34] N. Kalra, D. Ferguson, and A. Stentz. Hoplites: A market-based framework for complex tight coordination in multi-robot teams. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

[35] N. Kalra, D. Ferguson, and A. T. Stentz. Constrained exploration for studies in multirobot coordination. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2006.

[36] N. Kalra and A. Martinoli. A comparative study of market-based and threshold-based task allocation. In *Proceedings of the International Symposium on Distributed Autonomous Robotic Systems (DARS)*, 2006.

[37] B. Kalyanasundaram and K. Pruhs. Online weighted matching. *Journal of Algorithms*, 14(3), 1993.

[38] S. Koenig, C. Tovey, and W. Halliburton. Greedy mapping of terrain. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2001.

[39] H. Köse, U. Tatlidede, Çetin Meriçli, K. Kaplan, and H. L. Akin. Q-learning based market-driven multi-agent collaboration in robot soccer. In *The Turkish Symposium on Artificial Intelligence and Neural Networks*, 2004.

[40] M. Lagoudakis, E. Markakis, D. Kempe, P. Keskinocak, A. Kleywegt, S. Koenig, C. Tovey, A. Meyerson, and S. Jain. Auction-based multi-robot routing. In *Robotics: Science and Systems*, 2005.

[41] M. G. Lagoudakis, M. Berhault, S. Koenig, P. Keskinocak, and A. J. Kleywegt. Simple auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004.

[42] G. Laporte, Y. Nobert, and H. Mercure. The multi-depot travelling salesman problem. *Methods of Operations Research*, 40, 1981.

[43] T. Lemaire, R. Alami, and S. Lacroix. A distributed tasks allocation scheme in multi-UAV context. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

[44] L. Lin, W. Lei, Z. Zheng, and Z. Sun. A learning market based layered multi-robot architecture. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

[45] L. Lin and Z. Zheng. Combinatorial bids based multi-robot task allocation method. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

[46] D. C. MacKenzie. Collaborative tasking of tightly constrained multi-robot missions. In *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings of the 2003 International Workshop on Multi-Robot Systems*, volume 2. Kluwer Academic Publishers, 2003.

[47] R. Nair, T. Ito, M. Tambe, and S. Marsella. Task allocation in the rescue simulation domain: A short note. In *RoboCup-2001: The Fifth Robot World Cup Games and Conferences*. Springer-Verlag, 2002.

[48] E. Oliveira, J. M. Fonseca, and N. R. Jennings. Learning to be competitive in the market. In *Proceedings of the AAAI Workshop on Negotiation: Settling Conflicts and Identifying Opportunities*, 1999.

[49] A. Pongpunwattana, R. Rysdyk, J. Vagners, and D. Rathbun. Market-based coevolution planning for multiple autonomous vehicles. In *Proceedings of the 2nd AIAA Unmanned Unlimited Systems, Technologies and Operations Conference*, 2003.

[50] G. Rabideau, T. Estlin, S. Chien, and A. Barrett. A comparison of coordinated planning methods for cooperating rovers. In *Proceedings of the AIAA 1999 Space Technology Conferece*, 1999.

[51] I. M. Rekleitis, A. P. New, and H. Choset. Distributed coverage of unknown/unstructured environments by mobile sensor networks. In *3rd International NRL Workshop on Multi-Robot Systems*, 2005.

[52] D. J. Rosenkrantz, R. E. Stearns, and P. M. L. II. Approximate algorithms for the traveling salesman problem. In *Proceedings of the 15th Symposium on Switching and Automata Theory*, 1974.

[53] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*, 1993.

[54] T. Sandholm. Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium: Satisficing Models*, 1998.

[55] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1), 2002.

[56] T. Sandholm and V. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence, Special Issue on Economic Principles of Multiagent Systems*, 94(1), 1997.

[57] S. Sariel and T. Balch. Efficient bids on task allocation for multi-robot exploration. In *The 19th International FLAIRS Conference*. Florida Artificial Intelligence Research Society, 2006.

[58] J. Schneider, D. Apfelbaum, D. Bagnell, and R. Simmons. Learning opportunity costs in multi-robot market based planners. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

[59] W. Sheng, Q. Yang, S. Ci, and N. Xi. Multi-robot area exploration with limited-range communications. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004.

[60] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2000.

[61] R. Simmons, S. Singh, D. Hershberger, J. Ramos, and T. Smith. First results in the coordination of heterogeneous robots for large-scale assembly. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*, December 2000.

[62] R. G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, 29(12), 1980.

[63] A. Solanas and M. A. Garcia. Coordinated multi-robot exploration through unsupervised clustering of unknown space. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004.

[64] A. Stentz and M. B. Dias. A free market architecture for coordinating multiple robots. Technical Report CMU-RI-TR-99-42, Robotics Institute, Carnegie Mellon University, December 1999.

[65] G. Thomas, A. M. Howard, A. B. WIlliams, and A. Moore-Alston. Multi-robot task allocation in lunar mission construction scenarios. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, 2005.

[66] K. J. Tilley and D. J. Williams. Modeling of communications and control in an auction-based manufacturing control system. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1992.

[67] C. Tovey, M. G. Lagoudakis, S. Jain, and S. Koenig. The generation of bidding rules for auction-based robot coordination. In *Proceedings of the 3rd International Multi-Robot Systems Workshop, Naval Research Laboratory*, 2005.

[68] D. Vail and M. Veloso. Multi-robot dynamic role assignment and coordination through shared potential fields. In A. Schultz, L. Parker, and F. Schneider, editors, *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings from the 2003 International Workshop on Multi-Robot Systems*, volume 2. Kluwer Academic Publishers, 2003.

[69] J. M. Vidal. The effects of cooperation on multiagent search in task-oriented domains. In *Proceedings of the Autonomous Agents and Multi-Agent Systems Conference*, 2002.

[70] L. Vig and J. A. Adams. Market-based multi-robot coalition formation. In *Distributed Autonomous Robotic Systems (DARS)*, 2006.

[71] E. Wolfstetter. Auctions: An introduction. *Journal of Economic Surveys*, 10(4), 1996.

[72] Y. Xu, P. Scerri, K. Sycara, and M. Lewis. Comparing market and token-based coordination. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 2006.

[73] R. Zlot and A. Stentz. Complex task allocation for multiple robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

[74] R. Zlot and A. Stentz. Market-based multirobot coordination for complex tasks. *International Journal of Robotics Research Special Issue on the 5th International Conference on Field and Service Robotics*, 25(1), January 2006.

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

A Example Problems and Case Studies

In this section, we present a couple of real-world example scenarios and demonstrate how they can be addressed using market-based techniques. First, in Section A.1, we present an aggregation problem in which a team of robots must collect a set of spatially distributed items and transport them to a common location. We concretely illustrate how a market-based approach can be formulated for this domain, and we consider problem variations that highlight alternatives and extensions to the basic solution. Second, in Section A.2, we present a case study in market-based multirobot exploration, a problem that has been addressed in a variety of ways in the literature. Thus it allows us to highlight a range of approaches to a single problem, to analyze and compare some of the popular existing methods, and to suggest areas for future work.

A.1 Basic Aggregation

Consider a scenario in which we have a team R of robotic forklifts that must move all of the wooden pallets T scattered in a warehouse to a loading bay for transportation to another site. Their goal is to bring all the pallets into the loading bay as quickly as possible so that the shipment can be made on time. Assume that each forklift can carry up to P pallets at a time; once a forklift has P pallets, it must go to the loading bay to unload. Here, each pallet maps to a single task in the mission and the team's objective of collecting all the pallets as quickly as possible can be described by the global cost function C:

$$C(A) = \max_{r \in R} c_r(T_r(A)) \tag{1}$$

where A is a particular allocation of tasks to robots, $T_r(A)$ is the set of pallets allocated to r for transport in allocation A, and $c_r(X)$ is the fastest time that r can transport all pallets in X to the bay. Our goal is to find an allocation A^* that minimizes this global cost function.

One approach to task allocation is to hold sequential, single-task auctions (*e.g.* one pallet is allocated in each auction and only one auction occurs at a time). We could use the following bidding function B:

$$B(r, p_j) = c_r(T_r(A_{curr}) \cup \{p_j\}) \tag{2}$$

where A_{curr} is the current allocation and p_j is the pallet being auctioned. Essentially, the robot bids its total time for moving its currently allocated set of pallets and the pallet being auctioned. To find the optimal time, the robot would have to solve an instance of a multi-depot capacitated vehicle routing problem (MD-CVRP). In reality, this is intractable for more than a few tasks, so we might use some heuristics or limit our search space to find a feasible or approximate solution. Naturally, this is not guaranteed to result in the best possible solution. As we discussed in Section 4, Tovey *et al.* [67] demonstrate for $P = |T|$ that such a bidding rule results in better performance than using other bidding rules such as bidding the marginal increase in time.

A.1.1 Aggregation of Multiple-Robot Tasks

To extend example A.1, suppose that we have pallets of different sizes and that several forklifts might have to work together to jointly carry a pallet depending on its size. For example, small or S-type pallets might require only a single forklift, while medium-sized or M-type pallets require the efforts of two forklifts, and that large or L type pallets require the efforts of three forklifts. Then, the M and L type pallets are multiple-robot tasks while the S type pallets are single-robot tasks. Further suppose that a forklift can carry up to P of the S-type pallets at a time but only one M- or L- type pallet. Now, our objective of collecting the pallets as quickly as possible remains the same, so our global cost function C remains the same as in Equation 1. However, note that now the function c_r will often depend not just on the tasks allocated, but also on the activities of other teammates. Several types of auctions and bidding approaches are possible. We present three such approaches and hypothesize about the benefits and the drawbacks of each.

Instantaneous allocation. A first approach might be to use an instantaneous allocation as we discussed in Section 3.1.1 to simplify the problem of MR task allocation. That is, an auctioneer might hold sequential, single-task auctions and each unassigned robot bids its cost to complete that task:

$$B(r, p_j) = c_r(p_j)$$

where $c_r(p_j)$ is simply the time it would take for robot r to move pallet p_j. If the pallet being auctioned is an S pallet, the auctioneer awards the task to the robot with the smallest bid value. If the pallet being auctioned is an M or L pallet, then the auctioneer would award the task to the set of robots with the smallest collective set of bid values. The advantage to this approach is that it is very simple and will keep most robots working most of the time, thereby reducing the total time taken to move the pallets. However, without sequencing, each robot is limited to moving a single pallet at a time even though it has the capability to move up to P pallets at a time, thus creating potentially significant inefficiencies.

Mission decomposition. A second approach that would allow some sequencing might be to treat the problem as three different missions: the first mission is to move all the S-type tasks, the second is to move all the M-type tasks, and the third is to move all the L-type tasks. Thus, the first portion of the mission would be identical to Example A.1. In the second portion of the mission, the auctioneer could hold sequential single-item auctions and collect bids for each robot that indicate the earliest each robot could complete that item and the corresponding cost. It would award the item to the cheapest pair of bids.

For example, suppose that for a robot r, $T_r(A_{curr}) = (p_{M2}, p_{M1}, p_{M4})$, and this is also the order in which it will complete those tasks. Now, r is *obligated* to work on those tasks in that order because other robots have coordinated their schedule with r through the bidding. So, when a task p_{M6} is auctioned, r can only consider working on it after it completes p_{M4}, say at time t, and so submits a bid for time t with the cost

of that new schedule. The auctioneer awards p_{M6} to the pair of robots that have similar times of completion and that collectively have the least-costly completion time. Once the M-type tasks are all allocated and completed, an analogous approach is taken for the L-type tasks.

The benefit here is that we have significantly simplified the problem into three shifts and can easily harness the benefits of task sequencing. The obvious drawbacks are that some robots may be idle when there are fewer tasks than robots even though there are other types of tasks to be completed, and the solution space is still limited to scheduling of tasks of the same type.

Full sequencing. A third, more complex approach is to have the auctioneer hold sequential, single-item auctions for any type of task at any time. We might use two types of bidding depending on the type of pallet being auctioned. If the item currently up for auction is an S-type pallet, then all robots submit a single bid as in Equation 2. However, when the task is an M-type or a L-type pallet, robots would submit a set of bids which indicate not just the cost, but also the time that a robot could help move that pallet. The auctioneer would award the task to the strongest set of two or three bids, respectively. This approach of submitting multiple bids is similar to that used by MacKenzie [46] which we discussed in Section 3.1.1.

For example, suppose that for a robot r, $T_r(A_{curr}) = (p_{S2}, p_{S1}, p_{S4}, p_{S10})$ where S, M, and L indicate the type of pallets. Further suppose that this is the order in which r expects to complete the tasks. Then, when a pallet p_{L11} is auctioned, r can consider completing it after any of the tasks in its current schedule. It bids a set of pairs $\{t_k, c_r(T_r(A_{curr}) \cup p_j, t_k)\}$ where t_k is some time that r could arrive at p_{L11} to complete it and $c_r(T_r(A_{curr}) \cup p_j, t_k)$ would be the resulting cost of that schedule. The auctioneer would receive many bids for completing this task at different times and the corresponding costs. Its goal is to find the set of three bids that have the smallest total cost and that have the same or very similar arrival times. If such a set exists, p_{L11} is awarded to those robots under the agreement that they will be ready to move P_{L11} at that time. So, r's schedule $T_r(A_{curr})$ might now be $(p_{S2}, p_{S1}, p_{S4}, p_{L11}, p_{S10})$. The next time a pallet is auctioned, r can only submit bids that will still enable it to be ready to complete p_{L11} at the time agreed upon.

The benefit here is that we have greatly enlarged our solution space and will probably find a better solution, but the drawback is the increased complexity of the approach: robots may have to perform complex operations to submit a set of bids, and the auctioneer's clearing algorithm may computationally intensive given the large number of bids per auction. Note that there are other ways in which a market-based solution can be constructed for this problem. For example, hierarchical auctions can be used to allocate both abstract tasks and the decomposed primitive tasks. The detailed solutions explore a few options for applying market-based techniques to the aggregation problem.

A.2 Market-based Exploration

Exploration is a fundamental multirobot problem, and makes for an interesting case study since there are a few different existing market-based solutions [60, 59, 75]. Ex-

ploration requires the team to acquire as much information about the environment as possible (*e.g.* by building a map) within the shortest amount of time or distance traveled[3]. Additionally, limiting repeated coverage demonstrates an efficient team solution. Typically, these requirements are manifested in the utility function, which tries to balance information gain (revenue) versus travel distance, time, or other cost factors. Here we describe three market-based approaches, specifically by looking at the auction mechanism and utility functions in each.

One approach described by Simmons *et al.* [60] uses a central greedy instantaneous assignment (IA) algorithm. The idea is that each robot bids on a list of goal point tasks, and each is assigned exactly one task by the auctioneer. The map representation used is an occupancy grid, in which each cell contains a probability of containing an obstacle and can be classified as free, occupied, or unknown by a simple thresholding. The goal points selected are frontier cells, which are known free cells adjacent to unknown cells. Each bid contains an estimated cost and information gain associated with each frontier point. (The reason for including both values separately will become apparent shortly.) The costs are based on the distance of the optimal path from the current robot position to the goal. The information gain is based on the number of unknown cells expected to be visible from the goal location. The utility for each task t is calculated by the auctioneer during auction clearing as: $U_i(t) = I_i(t) - d_i(t)$, where $I_i(t)$ is robot i's expected information gain from task t, and $d_i(t)$ is the distance between the robot and the task. (Although not used in this approach, a scaling factor may be required for this type of utility function since the units of information gain and cost are different. Such a weight would specify the point at which it is no longer profitable to incur cost to gain more information.)

The auction clearing algorithm proceeds as follows. First, the goal with the maximum utility is allocated to the highest bidder. Then, the auctioneer discounts all remaining information gains according to the expected overlap around the assigned frontier region and the outstanding goals before making the next allocation. This is continued until all robots are assigned a goal or no goals remain. During execution, the robots are retasked whenever there is a map change; thus robots are not always required to reach their goals. One drawback to this approach is that it requires communication with a central agent, as well as considerable information sharing in order to ensure all robots maintain the same map. Additionally, the use of an IA allocation algorithm means that the robots behave myopically, although it is not clear if that is an undesirable trait for multirobot exploration [38].

Sheng *et al.* [59] essentially suggest several modifications and enhancements to the approach of Simmons *et al.* One difference is that they use a distributed auction protocol that does not require a central auctioneer or map. Robots in their system attempt to reach their goals, and upon task completion a robot determines the highest-utility goal from a list of frontier cells and broadcasts a single-task auction for that task. After a predetermined amount of time, if no other robot has submitted a higher bid for the task then it declares itself the winner and starts execution. Because each auction is only for a single task, a robot may have to call multiple auctions if it is continuously

[3] "The multirobot exploration problem" has also been used to describe the multi-depot traveling salesman problem in some publications [4, 41, 57], which is different from the problem we describe here.

outbid. Using multi-task auction may be an effective way to speed up the allocation process by reducing the number of auctions required in such cases. Utility functions are also modified to include a "nearness factor", describing how close the robot is to other teammates: $U_i(t) = \omega_1 I_i(t) - \omega_2 D_i(t) + \omega_3 \lambda_i$. The intention is that this additional cost factor will tend to keep robots further apart to emphasize exploration breadth over depth. Based on simulation results, it appears that including the nearness factor may improve exploration time slightly. Since there is no central auctioneer, the robots try to maintain consistent map information in order to reduce repeated coverage. They do this by broadcasting their map updates upon reaching target points and use an information sharing protocol to reduce repeated data transmission during encounters between previously communication-isolated subgroups. If a robot receives a map update while its auction is still pending, it recalculates its utilities and starts a new auction.

Zlot *et al.* [75] introduce a distributed approach that uses peer-to-peer auctions, time-extended allocation (TA), and require no central agent. In this approach, each robot applies one or more of a set of goal point generation heuristics to produce a list of target points to visit. Although frontier points are possible as a goal generation heuristic, three simpler ones are used: choose a random point in an unknown region of the map; choose the center of the nearest unexplored region; and one based on spatial decomposition using quadtrees. At startup and subsequently upon completion of a task, a robot generates new goal points using one of these three algorithms and inserts them into its schedule until it reaches a predefined maximum length. For each heuristic, if a resulting goal is within a minimum distance of an existing goal point in the robot's schedule or list of done tasks, it is thrown away. Robots periodically call single-award multi-task auctions to rid themselves of any goals that they may have generated that can be better-handled by another teammate. Bids are calculated as expected information gain minus expected weighted cost ($B_i(t) = I_i(t) - \omega c_i(t)$, where $c_i(t)$ is the marginal distance increase from inserting task t into robot i's schedule and ω is a scaling factor). In addition to submitting bids, participants can inform the auctioneer if they are already in possession of a similar task, in which case the auctioneer can eliminate that goal entirely. A limited amount of information sharing is also implemented in order to further reduce repeated coverage. In contrast to the system of Simmons *et al.* and Sheng *et al.* , this approach does not rely on maintaining consistent maps or having a central agent for task allocation, and thus can run in more communication-limited situations. In addition, the inclusion of task sequencing eliminates the myopic behavior of the robots in the other approaches above, although no direct comparisons have been made between the approaches to measure any impact this difference might have.

Unfortunately, since the above approaches have many differences, it is difficult to say which one is the best for a given problem instance without implementing each and running a thorough experiment. However, Table 4 summarizes several isolated features, allowing us to compare some aspects of the algorithms. Listed are the type of allocation agent, the category of task allocation, the auction type, and the utility function. The type of allocator reflects on the robustness of the system: as discussed in Sections 2 and 6 distributed allocation mechanisms can be more resilient to individual robot or communication failures. The allocation type refers to the taxonomy of Gerkey and Matarić [25] as discussed in Section 3. The fact that there are both instantaneous and time-extended solutions for multirobot exploration systems simply tells us that the

problem can be modeled in different ways—determining which choice is better would require empirical or theoretical results that are not currently available. Sections 4 and 5 tell us about the tradeoffs in solution quality and scalability of the auction types used in these approaches which are listed next in the table. As mentioned above, using a multi-item auction instead of a single-item auction in the approach of Sheng *et al.* may speed up the allocation stage with minimal added computational and communications overhead. Finally, Table 4 lists the utility functions used by the three approaches. A fieldable multirobot exploration system would likely require a carefully tuned (by hand or by learning) combination of multiple factors.

Table 4: Summary of market-based exploration approaches. Notation used for utility functions are described in the text of Appendix A.2.

Approach	Allocator	Allocation type	Auction type
Simmons *et al.* [60]	centralized	ST-SR-IA	multi-item
Sheng *et al.* [59]	distributed	ST-SR-IA	single-item
Zlot *et al.* [75]	distributed	ST-SR-TA	multi-item

Approach	Utility Function
Simmons *et al.* [60]	$U_i(t) = I_i(t) - d_i(t)$
Sheng *et al.* [59]	$U_i(t) = \omega_1 I_i(t) - \omega_2 D_i(t) + \omega_3 \lambda_i$
Zlot *et al.* [75]	$U_i(t) = I_i(t) - \omega c_i(t)$

Other systems that are not explicitly market-based but use similar approaches are also worth mentioning here as the utilized utility functions can be used in any of the above solutions. Solanas *et al.* [63] try to explicitly decompose the environment among the team using a k-means clustering algorithm periodically. Robots then add fixed penalty terms to their cost functions if a target point is outside their assigned region, or a smaller variable penalty term based on distance to the region centroid if the point is within the region. This encourages the robots to explore their region and discourages repeated coverage. Jia *et al.* [32] use a more complicated utility which is calculated as the ratio between the information gain and the cost (travel time plus observation time) multiplied by a term representative of local map topology.

The multirobot exploration problem demonstrates that there is not always one particular way to design a market-based approach to solve a problem: specific requirements of the domain and abilities of the robots may play a significant role in the design process. Learning the most effective utility function might provide a more solid grounding over using one of the several existing heuristic utility functions.