

---

# A Comparative Study of Market-Based and Threshold-Based Task Allocation

Nidhi Kalra<sup>1</sup> and Alcherio Martinoli<sup>2</sup>

<sup>1</sup> Robotics Institute, Carnegie Mellon University [nkalra@cmu.edu](mailto:nkalra@cmu.edu)

<sup>2</sup> Swarm-Intelligent Systems Group, École Polytechnique Fédérale de Lausanne  
[alcherio.martinoli@epfl.ch](mailto:alcherio.martinoli@epfl.ch)

In this paper we compare the costs and benefits of market-based and threshold-based approaches to task allocation in real world conditions, where information and communication may be limited or inaccurate. We have performed extensive comparative experiments in an event-handling domain. Our results indicate that when information is accurate, market-based approaches are more efficient; when it is not, threshold-based approaches offer the same quality of allocation at a fraction of the expense. Additionally, both approaches are robust to low communication and task perception ranges in our experimental domain.

## 1 Introduction

Multirobot coordination has become a popular area of research and advanced significantly in recent years. Researchers have developed a wide range of coordination approaches to harness the many benefits of robot teams (including speed, robustness, flexibility, and the ability to complete a wider range of tasks) in a variety of challenging real-world application domains. Nevertheless, before multirobot systems can be used extensively, it is imperative to understand the tradeoffs between the numerous coordination schemes.

We are interested in particular in understanding the tradeoffs between self-organized approaches and intentional approaches to multirobot task allocation in real-world conditions. Self-organized approaches are fully decentralized and achieve complex collective behavior from the local interactions of many simple individuals. Robots choose their actions independently and asynchronously using positive and negative feedback mechanisms and randomness. Additionally, interactions between agents are often modulated by signs left in the environment or by local, broadcast communication. Threshold-based approaches, in particular, are popular self-organized solutions to multirobot task allocation; in such approaches, a robot's choice of activity is modulated by a perception of stimulus or demand for a task and its response threshold

for that task. Alternatively, in intentional approaches, robots are typically complex and coordinate with the explicit intent of achieving a team goal. Market-based approaches, are popular intentional approaches to multirobot task allocation. In such systems, robots act as self-interested agents participating in a virtual market economy and allocate tasks by buying and selling them over the market. These approaches exploit points of centralization in the form of auctions to produce allocations. In general, self-organized approaches such as threshold-based algorithms consume fewer communication and computation resources to create a division of labor, while intentional approaches such as market-based systems consume more resources but also tend to produce more efficient allocations.

Little work has been done to truly understand the costs and benefits of exploiting communication and points of centralization for real-world task allocation. In this paper, we highlight concepts and results from our comparative study of the performances of market-based and threshold-based task allocation schemes under realistic conditions with limited communication, noisy state and task estimation, and limited task perception. We show how such conditions affect each approach and make to recommendations for using these allocation methods. Our full study is published in a technical report [1].

## 2 Threshold- and Market-Based Task Allocation

Given a team of robots with a common goal and a finite set of resources, the challenge of task allocation is to determine how the team’s global goal be divided up and assigned to individual team members so that some global objective is met.

In threshold-based approaches, each robot has an activation threshold for each task that needs to be performed. It continuously perceives the stimulus for each task; this stimulus reflects the urgency or importance of performing that task. When a robot perceives that the stimulus for a particular task exceeds its threshold, it begins completing the task. When the stimulus falls below this threshold (e.g., when the task is completed), the agent stops executing those behaviors. This response can be deterministic or probabilistic. Threshold-based task allocation has been demonstrated in a number of domains such as foraging [2] and aggregation [3].

Alternatively, in market-based systems, robots act as self-interested agents in pursuit of individual profit. They are paid in virtual money for tasks they complete and must pay in virtual money the value of the resources they consume. Tasks typically are distributed through auctions held by an auctioneer; this auctioneer is either a supervisor agent or one of the robots. Robots compete through bidding to win those tasks that they can complete inexpensively and thus maximize their profit. This price-driven redistribution simultaneously results in better team solutions. Market-based task allocation has also

been proven on a number of domains including as exploration [4] and object manipulation [5].

Prior comparisons of these approaches include work by Gerkey and Mataric [6] in which they discuss the theoretical aspects of these approaches but not their performance in realistic conditions. Campos et. al. [7] and Cicirello and Smith [8] compare the two on the problem of allocating trucks to paint booths; however, they investigate the benefits of encoding the problem as a market-based approach or a threshold-based approach, but not the costs and benefits of centralization and communication. We believe that no prior work compares these approaches along the dimensions in which we are interested.

### 3 Formulation of Study

In this section we describe several key components of our study, including the domain, axes and metrics for evaluation, and our implementation of the algorithms.

#### 3.1 Event-Handling Domain

We use the event-handling domain as the framework for our study. In this domain, events occur at unpredictable times and locations throughout the environment and must be handled by the robots. In real-world scenarios, these events might correspond to machine malfunctions in a factory or requests for package pick-up in a delivery service. Respectively, handling could involve robots fixing broken machines and making deliveries. Event-handling can also be mapped onto domains that are currently solved by market-based approaches (*e.g.*, exploration [4]) and threshold-based approaches (*e.g.*, foraging [2]).

Because market-based approaches are primarily useful when task and robot state knowledge is available [1], we define event-handling precisely: robots can sense the locations of individual tasks and can estimate their own positions. The resulting task allocation problem is to assign particular events to individual robots. In some market-based approaches, a single robot may be allocated multiple tasks and keeps a schedule of these tasks. However, in threshold-based approaches, robots are almost always limited to one-task-per-robot (OTPR) allocation because there is no clear way to enable scheduling with response thresholds (this would be an interesting area of future work). Thus, in fairness, we must use OTPR allocation in the market-based approach as well. Indeed, this formulation is frequently used for simplicity in many market-based approaches (see Dias et. al. [9] for examples).

#### 3.2 Axes, Metrics, and Experiment Parameters

We are interested in comparing the two candidate approaches under a variety of real-world conditions. Specifically, we consider the effect of poor accuracy


in state estimation, poor accuracy in task localization, reduced perception range, and reduced communication range. Our goal is to compare both the quality of allocations produced and the cost of producing those allocations. We measure quality in terms of the number of events handled and the total distance traveled by the team over a fixed time window. We measure cost in terms of the computation requirements (*i.e.* algorithm complexity and empirical running time) and communication requirements (*i.e.* number of messages and the size of messages).

In all of our experiments, we use a  $100 \times 100$  unit environment and run each trial to 100 iterations. There are a fixed number of robots that must handle a constant number of events; whenever an event is handled, a new one randomly appears to replace it. All our experiments were conducted using a point simulator in which robots are represented by their current locations and were run on a Linux machine with a 2.4 GHz processor and 512 Mb of memory.

### 3.3 Algorithms

Both market-based and threshold-based approaches have a number of parameters that can affect performance. We conducted numerous experiments to find the best formulation for each approach. Due to space considerations, we limit our discussion to resulting algorithms; full experiments can be found in our tech report [1].

In the market-based approach, our goal is to allocate an event to the robot that is closest to it. We compared two auction/bidding strategies. In the first, called  $M_1$ , we auctioned events in random order and robots bid their distances to the event being auctioned. In the second, called  $M_2$ , events were auctioned in order of increasing distance to their nearest robot. The corresponding bidding function is straightforward: a robot  $r$  without an assigned event bids a pair  $\{e, d\}$ , where  $e$  is  $r$ 's closest unassigned event and  $d$  is its distance to  $e$ . The robot that submit the bid with minimum distance  $d$  is assigned the corresponding event. In both approaches, only robots without assignments and with information about the event being auctioned place bids, rounds of auctions are held until all events are allocated or until no bids are received, and a new round of auctions is held whenever one event is completed (this provides a reallocation method). Furthermore, robots move with perfect actuation directly to the allocated event; if no event is allocated to a robot, it does nothing until the next auction round. Our results showed that ordered auctions improve performance significantly but also consume more computation resources. In the remainder of our experiments, we use  $M_2$  as our representative market-based algorithm. We also explored the use of a reserve price, *i.e.* the maximum value for  $d$  at which the auctioneer will award an event. This reserve keeps events from being allocated to very far away robots when all nearby robots are assigned to other events; the idea is that this event will be handled sooner if we simply wait for a nearby robot to finish with its



**Fig. 1.** Baseline comparison of the number of events handled vs. number of robots (left) and distance traveled vs. number of robots (right). We use a  $100 \times 100$  environment with 20 events and compare over 100 trials with 100 iterations each. Robots have perfect task information, perfect actuation, and infinite communication range. Error bars here and throughout this paper represent one standard deviation.

current event. We found that a reserve price equal to the expected distance between robots results in the best performance.

In the threshold-based approach, our goal is for each robot to handle the event that it is closest to, without duplicating another robot’s work. Thus, we formulated the stimulus  $\sigma(r, e)$  produced by an event  $e$  for robot  $r$  as in Eq. 1.a, where  $d(r, e)$  is the distance between the robot and the event.

$$(1.a) \quad \sigma(r, e) = \frac{1}{d(r, e)} \quad (1.b) \quad \theta_e = \frac{1}{\langle D_r \rangle} \quad (1.c) \quad p_e = \frac{\sigma(r, e)^n}{\sigma(r, e)^n + \theta_e^n}$$

We also explored the use of different thresholds and found that the best performance was achieved when the threshold  $\theta_e$  for every event  $e$  was equal to the inverse of expected distance  $\langle D_r \rangle$  between robots (Eq. 1.b). This threshold value is consistent with the best reserve price for the market. We then explored the use of deterministic versus probabilistic response. In the deterministic response, robots respond deterministically to the event  $e$  that has the largest stimulus above the threshold. In probabilistic response, they respond with probability  $p_e$  as in Eq. 1.c, where  $\theta_e$  is the threshold and  $n$  is the nonlinearity of the response. We found that deterministic response outperformed all probabilistic response methods (*i.e.* for all finite values of  $n$ ). We also experimented with deterministic versus random movement. We found that deterministic movement is always preferable, except when a robot is very close to its teammate and chances are high that they are handling the same event. Then, random movement for a short period of time reintroduces randomness into the system and produces better results.

To summarize, the best market-based approach uses ordered auctions and a non-zero reserve price; the best threshold-based approach uses a non-zero

**Table 1.** Comparing the computational and communication complexity. Communication complexity refers to the number of packets of information that must be exchanged between teammates. The size of each packet is constant. Here,  $r$  is the number of robots and  $e$  is the number of events.

Algorithm	Computational Complexity	Communication Complexity
Threshold	$O(re)$	–
Market	$O(re \log(e))$	$O(re)$

threshold, deterministic response, and a mixed actuation strategy. Figure 1 presents a baseline comparison of our final threshold-based and market-based algorithm in terms of the number of events handled, the total distance traveled by the team, and the computation time required. The market-based approach always handles more events while consuming less energy traveling than the threshold-based approach; this benefit comes at the cost of significantly more computation time.


Table 1 summarizes the algorithm and communication complexity for these approaches. In the threshold-based approach, each robot computes its distance to each event once per allocation, so the complexity is  $O(re)$  per allocation round, where  $r$  is the number of robots and  $e$  is the number of events. The same is true for the market-based approach; however, robots must also sort the events in order of increasing distance so they always bid on the closest event. Since most sorting algorithms run in logarithmic time of the number of items to sort, the total complexity is  $O(re \log(e))$ , and we expect that the linear components overshadow the logarithmic component. This is supported by our empirical results as well. In terms of communication requirements, the complexity (e.g. the number of packets that must be sent) for the market-based approach is  $O(re)$  because each robot sends one bid per event and the auction sends one award announcement to each robot per event. Additionally, the size of the messages sent between robots is constant. The threshold-based approach as we are currently using it has no communication requirements.

## 4 Comparative Results and Discussion

In this section, we explore the effects of imperfect state and task estimation, reduced communication range, and reduced perception range on market-based and threshold-based algorithms.

### 4.1 Imperfect State and Task Estimation

Local state estimation is imperfect in all real-world domains and usually affects the quality of allocations. To quantify the effects of poor state estimation in the event handling domain, we add Gaussian noise to each robot’s position estimate during allocation. We still allow perfect actuation and perfect




**Fig. 2.** Comparison of the number of events handled (left) and total distance traveled (right) vs. the standard deviation of Gaussian noise added to the robots’ local position estimation. These results are almost identical to those obtained when adding the same noise to task position estimation (not shown). We use a  $100 \times 100$  environment with 20 events and 10, 20, and 40 robots and compare over 100 trials with 100 iterations each.

task information to ensure that results are purely a product of localization error. This scenario is feasible if there is an overseeing agent that has access to task information and communicates that information to the robots; the robots themselves may have very poor perception.

Figure 2 plots the number of events handled by each approach (left) and the total distance traveled (right) against the standard deviation of the Gaussian noise. Qualitatively, the allocations of both algorithms degrade in the same way. Quantitatively, we see that the threshold-based approach and market-based approach have the same performance in terms of the number of events handled once the standard deviation of the error exceeds approximately 8. This ten percent error is a reality for many mobile robotic platforms, particularly smaller platforms without accurate proprioceptive sensors. The market-based approach is consistently less-costly in terms of distance when there are equally many or more robots than events. When there are fewer, the two approaches perform similarly because all robots are almost always moving to handling events.

Robots may also not have perfect task information, for instance when the supervisor or user does not have accurate information about the location of tasks but the tasks must still be allocated. In another set of experiments, we allow robots to have perfect actuation and perfect localization. This scenario is feasible if robots have poor or no task perception and the supervisor’s task perception is also not perfect. However, the robots do have accurate proprioceptive sensors that provide accurate localization. Both qualitatively and quantitatively, the results (not shown) are nearly identical to the results when we compared the effect of localization error. That the effect of poor position estimation is the same regardless of whether it is the robots’ or tasks’ positions highlights the symmetry of our problem. That is, allocating robots to events (as the market-based approach does) or events to robots (as the threshold-based approach does) results in almost the same allocation problem.



**Fig. 3.** Comparison of the number of events handled (left) and the total distance traveled (right), vs. the radius of the communication range. We use a  $100 \times 100$  environment with 20 events and 10, 20, and 40 robots and compare over 100 trials with 100 iterations each.

Finally, when we consider these results and the cost of market-based allocations, it's clear that market-based allocations are usually not worth the computation and communication cost unless accurate local state and task information is available.

## 4.2 Communication

Communication is essential for market-based approaches to perform auctions. However, in the real world, market-based approaches must be robust to reduced communication. Indeed, it is likely that not every member of a team will be able to communicate with every other member or with a central supervisor, if there is one. We experiment with the communication range to test the robustness of the market-based approach in this respect. Specifically, we allow auctions among only a connected set of robots. Within a connected set, we simulate a leader holding auctions for all the tasks for which it is aware. This is analogous to robots having perfect, long-range perception of the events in the environment (and no supervisor) but having limited communication range.

Figure 3 compares the performances of the market-based and threshold-based approaches with respect to the number of events handled (left) and the total distance traveled (right) by the team. First notice that the market-based approach quickly reaches its best performance in terms of events with a fairly short communication range. For 40 robots and 20 events, a range of 10-15 units suffices to produce the best results. Markets also reach their minimum total distance traveled with a fairly short communication range (about 20 units for 40 robots). The best communication range for a scenario depends on the robot and event densities. By comparing the minimum communication



range required to achieve the maximum performance (in terms of events) and the size of the auctions at that range (not shown for space considerations), we found that maximum performance can be achieved with less than a third of the total number of robots participating on average in each auction. Specifically, a team of 10, 20, and 40 robots achieved maximum performance when the average auction size was about 2, 6, and, 15. This highlights that this market-based approach is fairly robust to short communication ranges.

Notice that, in this domain, the strongest candidates for a particular event are likely to be in close proximity to each other and to that event. Therefore, these candidate robots are also likely to be able to communicate with each other and thus perform an allocation of the task even when the overall communication range is low. This suggests that the robustness demonstrated here may be somewhat specific to spatial domains such as this in which spatial proximity to a teammate correlates highly with the likelihood of competing with that teammate for tasks. This important property is common to many domains including exploration, mapping, and aggregation.


Thirdly, when market-based approaches have no communication, their performance degrades significantly. When a robot cannot communicate with any teammates, it will be the only member of the communicating set and thus will hold its own auctions. However, the only bid in the auction will be its own. In effect, the robot will be moving towards its closest event, provided that closest event is within the reserve price. Without communication, the market essentially becomes the threshold-based approach, but with the added expense of internal auctions to achieve the same result.

### 4.3 Task Perception

Often, it is not possible for a supervising agent to be aware of the tasks in the environment, for instance in exploration or search and rescue. Thus, it is important for multirobot task allocation schemes to be able to function without *a priori* task knowledge and exploit task perception during execution.

We experimented with low task perception by varying the range at which robots can detect events. Figure 4 plots the quality of allocation versus the radius of perception. Here, we compare the threshold-based approach against three market-based approaches that assume different communication scenarios. Firstly, we assume very poor communication (range of 5), but allow robots to share perceived event information. Secondly, we allow infinite communication and again allow robots to share event information. Thirdly, we allow infinite communication but only allow robots to bid on events that they perceive on their own. These three scenarios highlight how poor perception might be improved by better communication and information sharing.

Notice that again the minimum perception range required to achieve maximum performance is quite short, a range of 20 suffices for all approaches under all conditions. We believe this robustness is also somewhat specific to spatial domains such as this. In the threshold-based approach, as long as a robot



**Fig. 4.** Comparison of the number of events handled (left) and the total distance traveled (right) vs. the radius of the perception range, for the threshold based approach and three variants of the market based approach: one with poor communication and local bidding, one with infinite communication and local bidding, and one with infinite communication and global bidding.

can perceive its single closest event, it can begin handling that event. For the market, there is additional benefit in perceiving other nearby events in case a robot does not win the auction for its single closest event. In these experiments, a robot’s expected distance to its closest event is approximately 12. Thus, significantly greater perception range is not required for sufficient performance. This conclusion is further supported by the fact that performance is not improved in the market-based approach by robots sharing information about perceived tasks and bidding on all events rather than just bidding on locally perceived events. The idea is that if a robot did not perceive an event, it is probably not the best candidate to handle that event. Instead, we attribute the difference between the market-based approaches seen here to the increased communication range.

## 5 Conclusions

In this paper we have compared the performances of a self-organized and an intentional approach to multi-robot task allocation. Specifically, we compare the costs and benefits of exploiting points of centralization as done in market-based approaches to fully distributed threshold-based approaches in real world conditions. We evaluate how accuracy in local state and task estimation, communication range, and task perception range affect the quality of allocations on an event handling domain.

Our results indicate that when information is accurate, market-based approaches are more efficient (though threshold-based approaches with communication must still be considered). When information about tasks and local state is not accurate, market-based approaches may not be worth the added expense; rather, threshold-based approaches offer the same quality of allocation at a fraction of the expense. Additionally, both approaches are robust to low communication and task perception ranges. We hypothesize that this

is in part due to the spatiality of our experimental domain, a property that is common to a number of real-world domains. Furthermore, although our experiments use specific instances of these algorithms, we believe the results can be generalized to other variations since many of those we explored (e.g. probabilistic response thresholds) [1] follow similar trends.

In the near future, we hope to verify the results of our microscopic simulations with a realistic simulator. In the longer term, we hope to perform experiments with real robots.

## Acknowledgments

This work was carried out at EPFL where Nidhi Kalra was sponsored by an EPFL graduate student fellowship. Alcherio Martinoli is funded by a Swiss National Science Foundation Grant, contract No. PP002-68647.

## References

1. N. Kalra and A. Martinoli, “A comparative study of market-based and threshold-based multirobot task allocation,” EPFL, Lausanne, Switzerland, Tech. Rep. SWIS-IP1, February 2006.
2. M. J. B. Krieger and J.-B. Billeter, “The call of duty: Self-organized task allocation in a population of up to twelve mobile robots,” *Robotics and Autonomous Systems*, vol. 30, no. 1-2, pp. 65–84, 2000.
3. W. Agassounon and A. Martinoli, “Efficiency and robustness of threshold-based distributed allocation algorithms in multi-agent systems,” in *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 2002, pp. 1090–1097.
4. R. M. Zlot, A. Stentz, M. B. Dias, and S. Thayer, “Market-driven multi-robot exploration,” Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-02-02, January 2002.
5. B. P. Gerkey and M. J. Mataric, “Sold!: Auction methods for multi-robot coordination,” *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 758–768, Oct. 2002.
6. B. P. Gerkey and M. J. Mataric, “A formal analysis and taxonomy of task allocation in multi-robot systems,” *International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.
7. M. Campos, E. Bonabeau, G. Theraulaz, and J.-L. Deneubourg, “Dynamic scheduling and division of labor in social insects,” *Adaptive Behavior*, vol. 8, no. 2, pp. 83–94, 2001.
8. V. Cicirello and S. Smith, “Wasp-like agents for distributed factory coordination,” *Journal of Autonomous Agents and Multi-Agent Systems*, vol. 8, no. 3, pp. 237–266, 2005.
9. M. B. Dias, N. Kalra, R. Zlot, and A. Stentz, “Market-based multirobot coordination: A survey and analysis,” *IEEE Special Issue on Multirobot Systems*, 2006 (forthcoming).