

Improving Cost Estimation in Market-Based Coordination of a Distributed Sensing Task

M. Bernardine Dias, Bernard Ghanem, and Anthony Stentz

Robotics Institute

Carnegie Mellon University

5000 Forbes Ave., Pittsburgh, PA 15213

mbdias@ri.cmu.edu, bsg00@aub.edu.lb, tony@cmu.edu

Abstract – While market-based approaches, such as *TraderBots*, have shown much promise for efficient coordination of multirobot teams, the cost estimation mechanism and its impact on solution efficiency has not been investigated. This paper provides a first analysis of the cost estimation process in the *TraderBots* approach applied to a distributed sensing task. In the presented implementation, path costs are estimated using the D^* path-planning algorithm with optimistic costing of unknown map-cells. The reported results show increased team efficiency when cost estimates reflect different environmental and mission characteristics. Thus, this paper demonstrates that market-based approaches can improve team efficiency if cost estimates take into account environmental and mission characteristics. These findings encourage future research on applying learning techniques for on-line modification of cost estimation and in market-based coordination.

Index Terms – Multirobot Coordination, Market-Based, Cost Estimation

I. INTRODUCTION

The past decade has witnessed a growing emphasis in research relevant to multirobot coordination. This emphasis is generated by the increasing demand for automation in application domains where a single robot is no longer capable of performing the necessary tasks with the required solution quality. However, coordinating multiple robots to execute cooperative tasks is a challenging problem that involves many complexities. A survey of some challenges and solutions in multirobot coordination is presented in a publication edited by Balch and Parker [1].

Market-based approaches have recently been gaining popularity as efficient and versatile coordination mechanisms [6], [3], [9], [8], [2]. These approaches, based on Smith's Contract Net Protocol [10], are designed around the central premise of individual robots in a team estimating their costs to complete a specified set of available tasks and submitting bids, based on these cost estimates, to task auctions in order to accomplish efficient task allocations within the team. Thus, auctions serve as flexible and efficient task allocation mechanisms for the robot team. A comprehensive survey and analysis of market-based approaches applied to multirobot coordination is presented by Dias *et al* [7].

While numerous techniques such as frequent re-trading of tasks, clustered bids, and combinatorial auctions have been used to improve team performance [6], the basic premise of success for market-based approaches still relies on the

ability of individual robots to make reasonably good estimates of costs they must incur to complete tasks and translating these cost estimates into meaningful bids. This paper presents a study of how the mechanism for making these cost estimates impacts solution efficiency in the *TraderBots* approach [6], and discusses the impact of this study on market-based multirobot coordination techniques in general. This paper reports the first study of the interaction between cost estimation and solution efficiency in market-based multirobot coordination. Several experiments are conducted to investigate the extent to which improved cost estimates can improve team performance within the *TraderBots* framework.

The *TraderBots* approach is briefly described in the next section, followed by an explanation of the experimental setup and a presentation and analysis of the results, and concluding with a summary of the presented work and future research directions.

II. TRADERBOTS APPROACH

Dias [6] reports a detailed overview of the *TraderBots* approach; a market-based approach for multirobot coordination inspired by the contract net protocol by Smith [10]. A brief overview of the *TraderBots* approach is presented here to provide context for the reported experimental results and analysis.

Consider a team of robots assembled to perform a particular set of tasks. Consider further, that each robot in the team is modelled as a self-interested agent, and the team of robots as an economy. The goal of the team is to complete the tasks successfully while minimizing overall costs. Each robot aims to maximize its individual profit; however, since all revenue is derived from satisfying team objectives, the robots' self-interest equates to doing global good. Moreover, all robots can increase their profit by eliminating excess cost.

If the global cost is determined by the summation of individual robot costs, each deal made by a robot will result in global cost reduction. Note that robots will only make profitable deals. Furthermore, the individual aim to maximize profit (rather than to minimize cost) allows added flexibility in the approach to prioritize tasks that are of high cost and high priority over tasks that incur low cost but provide lower value to the overall mission.

The competitive element of the robots bidding for different tasks enables the system to decipher the competing local information of each robot, while the currency exchange provides grounding for the competing local costs in terms of the global value of the tasks being performed. The implementation of the TraderBots approach on a team of Pioneer robots engaged in a distributed sensing task, described in detail by Dias et al [5], provides the experimental test-bed for the study reported here.

III. EXPERIMENTAL SETUP

The chosen application is a distributed sensing problem where robots are tasked with gathering sensory information from various designated locations of interest (see Figure 1).

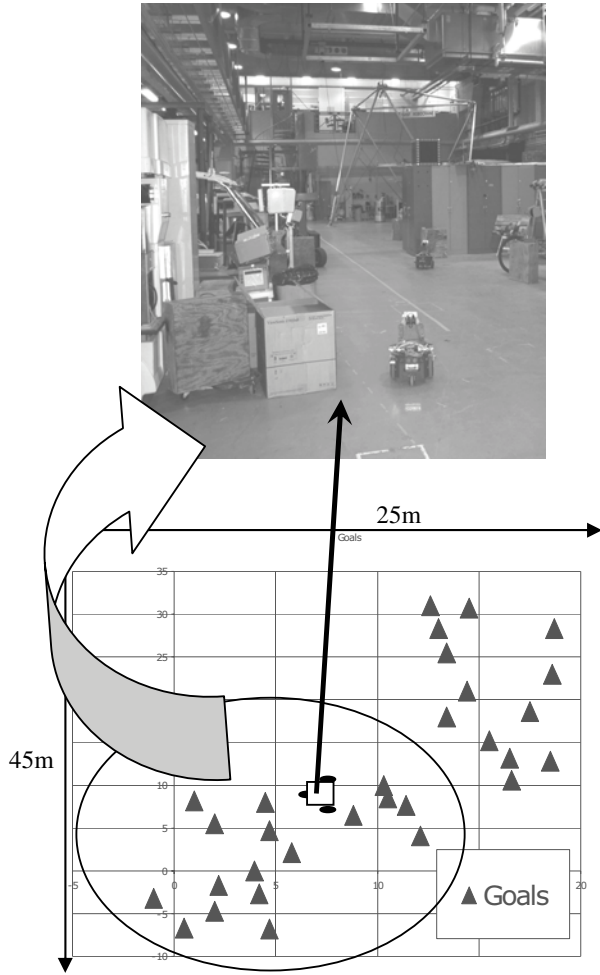


Figure 1: Pioneer robots engaged in a distributed sensing problem in a cluttered environment

This translates into a version of the travelling salesman problem (TSP) with the robots being represented by multiple salesmen following paths instead of tours (i.e. without the requirement that robots need to return to their starting locations) and where all the robots can start from different base locations. This is known as the multi-depot travelling salesman path problem (MD-TSPP). The tasks

translate to cities and the costs are computed as the cumulative distance to travel from city to city.




Figure 2: Robot-Generated Map of Test Area

A task is completed when a robot arrives at a city and the team task is complete when all cities are visited by at least one robot. The team cost is computed as the summation of the individual robot costs, and the team goal is to complete the team task while minimizing the cumulative distance travelled by the team, or equivalently, the number of robot-hours consumed by the team.

A. Robots

Two ActivMedia Pioneer II DX robots are used for these experiments. An 802.11b wireless card is used for ad-hoc communication between robots. Encoder data from the drive wheels is collected onboard from which dead reckoning position is calculated (x, y, θ) and enhanced by using a fiberoptic rate gyroscope to localize each robot. Obstacle detection is accomplished using a 180° SICK scanning laser range finder. Horizontal scan-range-data is incorporated with position data to create a 2-dimensional map. In addition to providing information to the operator, this map is used for local navigation, and a lower-resolution version of this map is used for cost estimation for trading.

B. Trading

Trading is a key component of the TraderBots approach. A trader assigned to each robot (*RoboTrader*) is responsible for participating in auctions to determine the tasks each robot commits to executing. In addition, these traders frequently attempt to subcontract or transfer tasks (that are not currently being executed) to other robots via auctions. A separate trader (*OpTrader*) serves as an interface agent between the operator and the robot team. These traders represent the human operators, and hence do not bid on any tasks for execution.

Each trader maintains a portfolio (list) in which it keeps track of its commitments, schedule, currently executing tasks, and tasks it trades to others. An internal alarm prompts each trader to auction all tasks in its schedule periodically. Note that tasks being executed are removed from the schedule and hence cannot be traded.

A trader initiates an auction by sending out a call for bids. Traders who are within communication range compute and submit bids to this auction. Once the specified deadline

expires, the auctioneer resolves the call by making an allocation based on the bids it received. If a trader receives an award for a bid it submitted, it accepts or rejects that award based on its current state. An award is binding after it has been accepted. Only single-task bids are allowed in this implementation.

C. Bidding and Costing

In order to participate in an auction, a trader needs to estimate the costs of a robot executing the tasks on auction. A trader announcing an auction must determine its reservation price, i.e. the highest price it is willing to pay to subcontract a task. These valuations are based on marginal costs – the difference between the cost of the current task schedule and the cost of the schedule without the tasks on auction. For a single task, an auctioneer’s valuation is the savings resulting from removing that task from its tour. A bidder’s marginal cost for a single task is the estimated cost of inserting the task into its tour.

Each robot is responsible for optimizing its own local schedule (i.e. given a set of tasks, the robots attempt to find the optimal TSPP solution to their local problem instance). In general, the TSPP is NP-hard, so an approximation algorithm is used when the problem instance encountered is large. Additionally, the problem encountered is an online variant of the TSPP – cities are arriving whenever a robot is awarded a task in an auction and are being removed whenever a task is traded to another robot. When adding a task to the tour, it is inserted into the tour at the location that results in the smallest increase in marginal cost. Tours are also optimized as a whole whenever tasks are added or removed. If the number of tasks is at most 12, the optimal solution is computed using a depth first search-based algorithm. If the number of tasks exceeds 12, then a minimum spanning tree-based 2-approximation algorithm is used [4] and the approximated solution is adopted if the resulting tour has a lower cost than the current tour. Tour optimization is also performed whenever a task is completed or failed as costs between cities may have changed due to new map information from recent sensor readings.

In the implemented TSPP scenario, all valuations are derived from inter-point distance costs. These costs are estimated using a D* path planner [12] with the robot’s local map as input. The D* map cells are valued as either FREE (cost of 5 for traversing half the width of the cell and 7 for traversing half the diagonal width of the cell), OBSTACLE (cost of 255), or UNKNOWN. In previous implementations of the TraderBots approach where efficient team solutions are produced, the UNKNOWN cell cost was optimistically valued as FREE [6], [5], [11].

D. Experiments

The goal of the study reported in this paper is to understand the effect on the team performance when cost estimation is altered, and potentially to improve these cost estimates and corresponding bids to enable higher efficiency in team performance. In order to carry out this investigation, different sets of tasks are assigned to two robots operating in our most frequently used test-bed, the Field Robotics Center High-Bay (Figure 1 shows a snapshot of the Field Robotics

Center High-Bay and Figure 2 shows a laser-map of the High-Bay generated by the robots).

During each experimental run, a chosen set of tasks is announced for auction, and the robots are allowed to trade until all the tasks have either been accomplished or rejected as failed. A task fails when a robot attempts to reach the specified location but cannot accomplish this goal within an allocated time period (typically 3-times as long as the time estimated for completion of the task based on available map information).

In order to provide some ground truth for the accuracy of cost estimates, the incurred cost (i.e. the actual distance traveled as reported by the encoders) is recorded. However, this is not a simple procedure. Since trading occurs throughout the execution of the team task, the robot tours can fluctuate during the entire time of operation. Hence, the robots do not frequently execute the tours that their cost estimates are based on. Instead, after a task is added to the robot’s schedule based on a cost estimate for completing that task, the robot’s schedule may change due to the addition or subtraction of other tasks in the schedule due to trading. Therefore, in order to correctly compare cost estimates with corresponding costs incurred, the following procedure is followed. Samples of scheduled tours, the corresponding cost estimates, and snapshots of the map used to estimate those costs are logged during the execution of each set of tasks. After the run is completed, the cost estimates of these sample tours are tested by assigning each sample tour to a single robot together with the corresponding map at the time of that cost estimate. The actual distance covered by the robot for each sample is recorded for comparison with the estimated cost. This procedure is repeated numerous times under various circumstances to yield the results analyzed in this paper. Specifically, five categories of experiments are performed on the robots:

1. Comparison of estimated and incurred costs with UNKNOWN map cells optimistically valued as FREE
2. Impact on team efficiency if UNKNOWN map cells are valued based on empirical measurements
3. Impact on team efficiency if UNKNOWN map cell cost is altered to reflect the possibility of unreachable goals
4. Sensitivity of UNKNOWN map cell valuation to the environment
5. Sensitivity of UNKNOWN map cell valuation to different task sets in the same environment

IV. EXPERIMENTS, RESULTS AND ANALYSIS

For each of the five categories listed in the previous section, the conducted experiments, obtained results, and corresponding analysis are presented next. A discussion of the overall results is also presented.

A. Comparison of estimated and incurred costs with UNKNOWN map cells optimistically valued as FREE

A set of 50 samples were chosen from resulting tours when 27 goal points (shown in Figure 8, set 1) are assigned to the two robots. For each sample, the estimated D* tour cost and the incurred tour cost are compared. Figure 3

Reference source not found. shows the variation in the error between the estimated and incurred tour costs with the distance traversed while executing the tour; the error grows significantly as the actual distance covered increases.

At short distances, within the sensor range for the SICK laser, the error incurred remains insignificant, since the immediate environment is accessible to the initial laser scan and hardly any unknown cells are encountered. However, as the tour extends beyond the sensed range, the error increases since the estimate is based on a tour that encounters a larger number of unknown cells. Since the traders are optimistic in their valuation of unknown cells, the cost of the tours are significantly underestimated.

B. Impact on team efficiency if UNKNOWN map cells are valued based on empirical measurements

The next step is to evaluate the impact of a less optimistic valuation of the unknown cells. To compute an appropriate valuation, an average incurred cost is computed based on empirical measurements. First, the incurred tour cost for each sample is converted to the corresponding D* cost by dividing the distance covered by the map resolution (0.2m) and multiplying by the average D* cell cost (computed as 12). Note that the traversed cells were assumed to be cells with a value of FREE and hence the average cost of traversing a cell (a cost of 14 along the diagonal and a cost of 10 along the width) is used to convert between distance and D* cost. This computed cost is then divided by the number of traversed cells to discover the cost per cell. Figure 3 shows the incurred and estimated cost per cell for each of the sample tours. The difference of these two costs forms the error graph shown previously in Figure 2.




Figure 3: Variation of estimated and incurred D* cost with distance traversed

As previously noted, the under-estimation of the cost grows significantly with increasing distance traversed in the tours. However, the most significant underestimates are due to tasks being unreachable; 3 of the assigned 27 goals were unreachable due to obstacles. Unreachable goals are often encountered during experiments, and hence are included for realism. Figure 4 shows the 43 samples from Figure 3 that did not include any unreachable goals.




Figure 4: Variation of incurred and estimated D* cost with distance traversed, excluding tours containing unreachable goals

As a first attempt to improve the estimated value of unknown cells, an average error per cell is computed excluding tours with unreachable goals. This average, computed over 43 sample runs (excluding samples containing any unreachable goals), is approximately 7 D* cost units (D* uses integer costs for higher computation speed). Thus, the unknown cells are initialized to a value of FREE + 7 (width cost of 24 and diagonal cost of 34) and the original 27 tasks are assigned to the robots again. Ten sample tours were analyzed as before, with the new value for UNKNOWN cells. The real and estimated costs for these 10 samples are compared below in Figure 5.




Figure 5: Variation of estimated and incurred cost with distance where unknown cells are valued as FREE + 7

The accuracy of estimated costs is notably improved using the new value of FREE + 7 for UNKNOWN cells. In fact, the team performance (measured as the cumulative cost accrued by the team) improved by 17% on average compared to the case where UNKNOWN cells are valued as FREE. Thus, a significant performance improvement is possible by taking into account environmental clutter when estimating costs.

C. Impact on team efficiency if UNKNOWN map cell cost is altered to reflect the possibility of unreachable goals

In order to incorporate information for typical runs where unreachable goals are assigned with some frequency, the above experiment is repeated while taking into account the effect of unreachable goals.




Figure 6: Variation of estimated and incurred cost with distance where unknown cells are valued as FREE + 11

Thus, a weighted average of the error per cell for all 50 original samples is computed and the corresponding error was determined to be 11 D* cost units. This value increased the value of UNKNOWN to 16 units of cost per half cell width. Once again, the same 27 tasks are assigned to the robots and ten sample tours are analyzed as before. Figure 6 shows the resulting estimated and incurred costs. This refinement of the value of an UNKNOWN cell resulted in an improvement of team efficiency by 29% compared to the optimistic valuation of UNKNOWN cells as FREE.

D. Sensitivity of UNKNOWN map cell valuation to the environment

The next experiment carried out investigates the dependence of the UNKNOWN cost on the obstacle density (clutter) of the environment. A second set of 25 tasks (including 4 unreachable goals), spread over a more cluttered and larger environment, are assigned to the robots. Once again, 50 sample tours, shown in Figure 7, are chosen and analyzed.




Figure 7: Variation of estimated and incurred cost with distance in a more complex environment

With this configuration, the new value for UNKNOWN cells is evaluated to be 78, which is a large variation from the previous configuration. Thus, a significant change in environment requires a new valuation of UNKNOWN cells as might be expected. This result demonstrates the need for adapting cost models to reflect the nature of the environment to improve team efficiency. In fact, frequent refinement of cost models to incorporate newly acquired information during execution is desirable. Furthermore,

dynamic environments will require updating the valuation of previously “known” map cells.

E. Sensitivity of UNKNOWN map cell valuation to different task sets in the same environment

The previous experiment demonstrates that the environment can drastically impact the cost models and hence, team performance. Thus, a final set of experiments is carried out to determine the adaptability of the previously computed values for UNKNOWN cells, given different task sets. Four additional sets of tasks are assigned in the same environment used in sections A through C. These four task sets are shown in Figure 8. The first set of tasks spreads evenly across the entire area. The second set of tasks is smaller and constrained to a region with higher obstacle density. The third set of tasks is spread over a relatively obstacle-free region and contains no unreachable goals. Finally, the fourth set is spread over the same region as the third set, but includes some unreachable goals.




Figure 8: Map of environment in task sets 1-4; black regions represent UNKNOWN cells, grey regions represent FREE cells, and white regions represent OBSTACLE cells.

The same experiments are performed again with three possible values for UNKNOWN (FREE, FREE + 7 and FREE + 11). Average performance over three runs is measured, compared, and reported in Table 1.

Unknown cell value	Set 1	Set2	Set3	Set4
FREE + 7	14%	5%	10%	34%
FREE + 11	21%	8%	2%	37%

Table 1: Percentage improvement in team performance (compared to an optimistic valuation of UNKNOWN cells as FREE) when UNKNOWN cells are valued as FREE + 7 and as FREE + 11 for different task sets

In general, an optimistic valuation of UNKNOWN cells as FREE is outperformed by the two refined valuations in all four experiments. This is to be expected since none of the environments were obstacle free. Moreover, the refinement that accounted for unreachable goals performed best in all cases where unreachable goals were assigned.

Perhaps a more surprising result is the lower gain in team-efficiency for sets 2 and 3. This observation is less surprising for set 3 where the goals were situated in a relatively obstacle-free area, and hence the improvement over an optimistic valuation of unknown cells is low. Furthermore, since there were no unreachable goals in this set, reflecting a risk factor for unreachable goals in the cost estimates results in poorer team performance. In set 2 however, the goals are assigned in a smaller but more cluttered area. In this case, perhaps the low improvement of team performance over optimistic valuation of unknown cells results due to the lower number of unknown cells due to the smaller designated area.

F. Discussion of Overall Results

The results presented in the last five sections argue for frequent updates to cost estimation metrics based on several factors including sensing capability, environmental clutter, the risk of unreachable goals, and the spread of assigned tasks. This is not however a complete list of factors that impact cost estimates and the resulting team performance. In fact, the results reported in this paper show the inadequacy of formulating bids based simply on an accurate representation of obstacles in the environment; the incorporation of an “unreachable goal risk factor” improved team performance on average.

Added complexities to the coordination problem arise from team interactions, communication limitations and robot malfunctions, dynamic environments, various bidding strategies and profit maximization schemes, and local optima produced by single-task bids and approximated solutions to tour optimization. For example, if a robot increases its cost estimations to reflect the current environment, but one of its teammates does not make this cost adjustment, the teammate will bid lower and win all the tasks. Similarly communication limitations can prevent participation in auctions and hence team performance can be adversely affected. The experiments reported in this paper only consider interactions between two robots and a single task domain. Even in the two-robot case, the improved cost estimation positively impacts team performance. In larger teams the improvement in efficiency should be greater, thus enhancing the scalability of market-based coordination approaches. Furthermore, a variety of task domains should benefit from improved cost estimation since many task domains rely on efficient interaction with uncertain environments. While team performance can be improved by accounting for some factors by empirical or other means, the application of learning techniques to alter cost models to reflect “current conditions” appears to be a more powerful strategy. The authors are currently researching relevant learning techniques to address this problem.

V. CONCLUSION AND FUTURE WORK

The results reported in this paper make a strong case for adapting cost estimates to reflect environmental clutter, the risk of being assigned unreachable goals, and the spread of goals in relation to the sensing footprint, in order to improve team performance in market-based coordination. Team

efficiency is notably improved with more relevant cost estimates. Future work will include the application of on-line learning techniques to appropriately adapt cost models during execution, while taking into account factors such as environmental clutter, frequency of unreachable tasks, and projected path distance. Also, a variety of bidding strategies will be tested to account for team interactions. Different application domains, cost functions, and task complexities will add further dimension to future research in this area.

ACKNOWLEDGMENT

The authors wish to acknowledge valuable discussions and feedback provided by Drew Bagnell, Robert Zlot, and Brett Browning, which contributed to the research and results reported in this paper. The authors would also like to acknowledge Robert Zlot, Marc Zinck, and Juan P. Gonzalez for their important contributions in developing the robot hardware and software. This work is sponsored by a research grant from The Boeing Company.

REFERENCES

- [1] Balch, Tucker, and Parker, Lynne E., (eds), “Robot Teams: From Diversity to Polymorphism” (book), Published by AK Peters, 2001.
- [2] Bererton, Curt, Gordon, Geoff, and Thrun, Sebastian, “Auction Mechanism Design for Multi-Robot Coordination”, *Proceedings of Conference on Neural Information Processing Systems*, 2003.
- [3] Berhault, Mark, Huang, He, Keskinocak, Pinar, Koenig, Sven, Elmaghraby, Wedad, Griffin, Paul, and Kleywegt, Anton, “Robot Exploration with Combinatorial Auctions”, *Proceedings of the International Conference on Intelligent Robots and Systems*, 2003.
- [4] Cormen, Thomas H., Leiserson, Charles E., Rivest, Ronald L., and Stein, Clifford, “Introduction to Algorithms (Second Edition)”, MIT-Press, 2001.
- [5] Dias, M. Bernardine, Zlot, Robert M., Zinck, Marc, Gonzalez, Juan P., and Stentz, Anthony, “A Versatile Implementation of the TraderBots Approach for Multirobot Coordination”, *Proceedings of the 8th Conference on Intelligent Autonomous Systems IAS*, 2004.
- [6] Dias, M. Bernardine, “TraderBots: A New Paradigm for Robust and Efficient Multirobot Coordination in Dynamic Environments,” *Ph.D. Thesis*, CMU-RI-TR-04-30, January 2004.
- [7] Dias, M. Bernardine, Zlot, Robert M., Kalra, Nidhi, and Stentz, Anthony, “Market-Based Multirobot Coordination: A Comprehensive Survey and Analysis”, *IEEE special edition on multirobot coordination*, submitted, 2005.
- [8] Frias-Martinez, Vanessa, Sklar, Elizabeth, and Parsons, Simon, “Exploring auction mechanisms for role assignment in teams of autonomous robots”, *Proceedings of the Robocup International Symposium*, 2004.
- [9] Gerkey, Brian P., and Mataric, Maja J., “Sold! Auction methods for multi-robot coordination”, *IEEE Transactions on Robotics and Automation, Special Issue on Multi-robot Systems*, 18(5):758-768, 2002.
- [10] Smith, Reid G., “The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver”, *IEEE Transactions on Computers*, Vol. C-29, No. 12, 1980.
- [11] Stentz, Anthony, Dias, M. Bernardine, Zlot, Robert M., and Kalra, Nidhi, “Market-Based Approaches for Coordination of Multi-Robot Teams at Different Granularities of Interaction,” *Proceedings of the ANS 10th International Conference on Robotics and Remote Systems for Hazardous Environments*, March, 2004.
- [12] Stentz, Anthony, “Optimal and Efficient Path Planning for Partially-Known Environments”, *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1994.