



May 2002

A Potential Field Based Approach to Multi-Robot Manipulation

Peng Song
University of Pennsylvania

R. Vijay Kumar
University of Pennsylvania, kumar@grasp.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/meam_papers

Recommended Citation

Song, Peng and Kumar, R. Vijay, "A Potential Field Based Approach to Multi-Robot Manipulation" (2002). *Departmental Papers (MEAM)*. 14.
https://repository.upenn.edu/meam_papers/14

Copyright 2002 IEEE. Reprinted from *Proceedings of the IEEE International Conference on Robotics and Automation 2002 (ICRA 2002)*, Volume 2, pages 1217-1222.
Publisher URL: <http://ieeexplore.ieee.org/xpl/tocresult.jsp?isNumber=21842&page=1>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/meam_papers/14
For more information, please contact repository@pobox.upenn.edu.

A Potential Field Based Approach to Multi-Robot Manipulation

Abstract

We describe a framework for controlling and coordinating a group of robots for cooperative manipulation tasks. The framework enables a decentralized approach to planning and control. It allows the robots approach the object, organize themselves into a formation that will trap the object, and then transport the object to the desired destination. Our controllers and planners are derived from simple potential fields and the hierarchical composition of potential fields. We show how these potential field based controllers and planners benefit complex group interactions, specifically for manipulating and transporting objects in the plane. Theoretically, we show how we can derive results on formation stability with potential field based controllers in many cases. Simulation results demonstrate successful application to a wide range of examples without showing sensitivity to parameters. Because the framework is decentralized at both trajectory generation level and the estimation and control agent level, our framework can potentially scale to groups of tens and hundreds of robots.

Comments

Copyright 2002 IEEE. Reprinted from *Proceedings of the IEEE International Conference on Robotics and Automation 2002 (ICRA 2002)*, Volume 2, pages 1217-1222.

Publisher URL: <http://ieeexplore.ieee.org/xpl/tocresult.jsp?isNumber=21842&page=1>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

A Potential Field Based Approach to Multi-Robot Manipulation

Peng Song Vijay Kumar

General Robotics, Automation, Sensing and Perception (GRASP) Laboratory
University of Pennsylvania, 3401 Walnut Street, Philadelphia, PA 19104
E-mail: {pings, kumar}@grasp.cis.upenn.edu

Abstract

We describe a framework for controlling and coordinating a group of robots for cooperative manipulation tasks. The framework enables a decentralized approach to planning and control. It allows the robots approach the object, organize themselves into a formation that will trap the object, and then transport the object to the desired destination. Our controllers and planners are derived from simple potential fields and the hierarchical composition of potential fields. We show how these potential field based controllers and planners benefit complex group interactions, specifically for manipulating and transporting objects in the plane. Theoretically, we show how we can derive results on formation stability with potential field based controllers in many cases. Simulation results demonstrate successful application to a wide range of examples without showing sensitivity to parameters. Because the framework is decentralized at both trajectory generation level and the estimation and control agent level, our framework can potentially scale to groups of tens and hundreds of robots.

1 Introduction

The last few years have seen active research in the field of control and coordination for multiple mobile robots, and application to tasks such as exploration [1], surveillance [3], search and rescue [7], mapping of unknown or partially known environments [6], distributed manipulation [9] and transportation of large objects [16]. An excellent review of contemporary work in this area is presented in [10].

In this paper we consider situations in which there may be no access to any global positioning system and the main sensing modality is vision. Our platform of interest is a nonholonomic car-like robot with a single physical sensor - an omnidirectional camera. Each robot is capable of autonomous operation or following one or two robots. The vision-based controllers used for autonomous operation are described in [2], while the controllers for following other robots are described in [4]. We are particularly interested in problems of cooperative manipulation, where a "rigid" formation may be necessary to transport a grasped object to a prescribed location.

Our main contribution in this paper is the completely decentralized approach to trajectory generation and controller design for coordinated distribution and manipulation. Each robot plans its own trajectory based on the sensory information of its surroundings and chooses from a finite set of control laws that describe its behav-

iors for interaction with the environment. Our framework allows robots to maintain or change formation while following specified trajectories, and to perform cooperative manipulation tasks in a scalable and modular fashion.

The rest of this paper is organized as follows. First we give a broad overview of our previous work in Section 2. In Section 3, we describe our approach for decentralized cooperative manipulation, and the three key components of this framework - trajectory generation, coordinated distribution and transportation. Section 4 analyzes stability measures for multi-robot formations and derives sufficient conditions for the construction of potential fields that has global minimum at specified configurations. Section 5 illustrates simulation results for different applications of our framework. Finally, some concluding remarks and directions for future work are provided in Section 6.

2 Background

In this section, we summarize the previous work done by collaborators on experimental platform, control laws and the software architecture for multi-robot coordinations. This section motivates this paper and puts the current work in context of our previous work [2, 5, 14]

Our robots are based on the commercially available Tamiya ClodbusterTM (CB) platform, a radio controlled 1/10 scale model truck. Each CB is equipped with an omnidirectional camera as its sole sensor. Using the controllers discussed in [14], the robots can maintain a prescribed rigid formation. This allows the robots to "trap" objects in their midst and to flow the formation thus guaranteeing that the object is transported (dragged) to the desired position and orientation. In Figure 1, the initial team configuration is centered around the box, with the goal to flow the now encumbered formation along a trajectory generated by the leader. By choosing a constraining formation geometry, the box is kept in contact with all three robots during the formation flow. Several snapshots from a sample run are shown in Figure 1.

The overall framework for control and planning is described in Figure 2. The controllers are nonholonomic formation controllers that allow robots to regulate the shape of the formation. The desired trajectory and shape are provided by a superior level of the control hierarchy, the trajectory generator [5].

In this paper, we address the trajectory generation problem for cooperative manipulation tasks. The trajectory generator can be completely decentralized so that each robot generates its own reference trajectory based

on the information available to it, through its sensors and through the communication network. Alternatively, the robots can designate or elect a leader that plans the trajectory for the group, while the other group members are responsible for organizing themselves and for following the leader. In the next section, we will advocate the decentralized strategy for the trajectory generator of cooperative manipulation tasks.




Figure 1: Cooperative manipulation tasks [14]

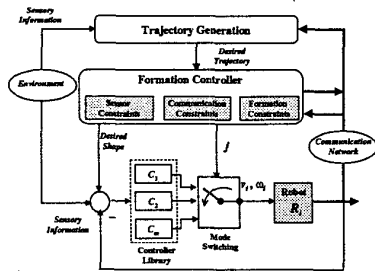


Figure 2: The two-level hierarchy for planning and control. The planned trajectory is described in terms of desired positions relative to neighbors and obstacles in the environment, and translated to a target shape. Geometric constraints, sensor constraints, and information from other robots determine the choice of controller.

3 Decentralized manipulation

We consider the problem of guiding a group of autonomous robots through an obstacle field to surround a target and eventually transport the target to a desired destination. Figure 3 shows an illustrative example in which 16 robots try to reach a circular object and transport it to a new destination.




Figure 3: Decentralized controllers are used to get non-holonomic robots to surround a target. Each robot uses one of the three different controllers (modes) depending on information about the neighbor's state.

Robot models We adopt the simple abstraction for the robot and its omnidirectional sensor shown in Figure 4. The outer most circle delineates the *sensing zone*, the region within which a robot can detect obstacles and

other robots. The *contact zone* is the region that is pre-computed based on the robot's maximum velocity and its ability to brake. It is the set of points in \mathcal{R}^2 that the body of the robot can reach if the robot is traveling at its maximum allowable velocity and is commanded to brake to a stand-still. Obstacles or other robots entering the inner *contact zone* generate forces that influence the robot's planned trajectory and thus its dynamic behavior. The size of the *protected zone* is slightly larger than the dimension of the robot. We model the robot as a rigid core with the shape of the *protected zone*, and a visco-elastic outer shell that replaces the *contact zone*.




Figure 4: The three regions surrounding a robot. The robot is blind outside the sensing zone. Objects in the contact zone generate forces that drive the robot away. Objects entering the protected zone cause collisions.

Control modes The control agent for decentralized manipulation tasks has three control modes – *approach*, *organization* and *transportation*, as described in Fig 5. In the first *approach* mode, the robots swarm to the object by following an attractive potential field centered at object location. After reaching the object, each robot (independently) enters into an *organization* mode where it moves away from its neighbors while staying near the object. This is done by stacking a repulsive potentials onto the existing approach potential to redistribute the robots. The repulsive potential is designed to organize the robots into the desired formation trapping the object in the process. A more formal definition of trapping based on the concept of *caging* [11, 15] can be found in [18]. Each robot autonomously transitions into the final *transportation* mode after it senses a quorum. In this phase, an added transportation potential similar to the one used in the approach phase but with a much lower intensity and centered at the destination location attracts the robots and the object to the goal position.

A variety of effective attraction and repulsion potential fields are summarized in [8, 17]. Let $r_{ij} = \|\mathbf{r}_i - \mathbf{r}_j\|$ be the Euclidean distance between robot and the object or the goal, a simple quadratic attractive potential function can be expressed as

$$V_{io}^a = \frac{1}{2} k_{io} r_{io}^2. \quad (1)$$

\mathbf{r}_i and \mathbf{r}_o are the position vectors of the robot and the object, respectively. An example of the repulsive potential is

$$V_{ij}^r = k_{ij} / r_{ij}, \quad (2)$$

where $r_{ij} = \|\mathbf{r}_i - \mathbf{r}_j\|$ is the Euclidean distance between robots i and j . In each control mode, the gradient of

the corresponding potential fields and an appropriately designed dissipative function provide the driving force to the robots and the trajectory is calculated by simulating the dynamics of the system.




Figure 5: Control modes for decentralized cooperative manipulation tasks.

The guards for the transitions between different modes are give by:

- Guard 1: $(r_{di} \leq d_0) \wedge (n_i \geq \alpha \frac{n \cdot r_s}{p})$
- Guard 2: $(r_{di} > r_0)$
- Guard 3: $(r_{di} \leq d_0) \wedge (n_i < \alpha \frac{n \cdot r_s}{p})$
- Guard 4: $(r_{di} \leq d_0) \wedge (n_i \geq \alpha \frac{n \cdot r_s}{p})$
- Guard 5: $(r_{di} > d_0)$

where d_0 and α are threshold constants. r_s is the radius of the sensing zone as depicted in Figure 4. n_i is the number of robots inside the sensing zone of robot i . r_{di} is the distance between the object and robot i . The quorum is limited by the perimeter, p , of the object.

Virtual collisions We use rigid body contact dynamics models to allow virtual collisions between the robot and its surroundings instead of avoiding them. We adopt a state-variable based compliant contact model to compute the interaction forces between two contacting objects. The details and variations on the compliant contact model are discussed in [12]. A key feature of this model is that it allows to resolve the ambiguous situations when more than three objects came into contact with one robot.

When objects (including other robots) from the environment enter the sensing zone, their relative position and velocity are estimated by the robot. When they enter the contact zone, the robot simulates contacts between the objects and its visco-elastic shell using a compliant contact model [12] to compute normal and tangential forces exerted on the robot:

$$\lambda = F(\xi) + G(\xi, \dot{\xi}) \quad (3)$$

where ξ is the robot's estimate of the state of the world. The robot simulates its response to the external forces to generate its reference trajectory.

Example We illustrate our approach by an example in which 18 nonholonomic mobile robots are commanded to a goal position within an unknown environment. The snapshots as shown in Figure 6 indicate that the robots start from two groups and decide to split into three smaller groups when they encounter obstacles, and finally merge into one group when they are close to the object. The robots are autonomous – each robot runs its own trajectory generator to get to the goal while avoiding obstacles (including other robots). The control commands, (v_i, ω_i) , are exactly the commands generated by

the trajectory generator. As seen in the figure, the robots are able to navigate and reach the goal position. In the process however, the outer visco-elastic shell of each robot encounters contacts with other robot shells, and with the hard boundaries of the obstacle. The contact model allows the robots to bounce back and head toward its destination.

Discussion Note that the computation of the trajectory for each vehicle is based on information that is available to it through its sensors or through the communication network. The relevant information is the relative state information of other robots and obstacles in the contact zone. Collisions can occur only if this information is not available, either because of faulty sensors or failed communication channels. Each robot runs a simulator of the world and the same algorithm for computing trajectories. Thus, penetration of two contact zones, for example, will result in both robots being bounced away from each other with equal and opposite contact forces. Obstacles do not have visco-elastic shells. However, obstacles are stationary. Thus the contact zone must be sized and the properties of the visco-elastic shell must be carefully selected. Even if there is a head-on confrontation with an obstacle, it should be possible to completely dissipate the energy of the robot and allow it to come to a complete stand-still without causing it's protected zone to touch the obstacle.

The discussion thus far has not addressed proofs of convergence or stability. In the next section, we attempt to analyze the measures to see what guarantees can be established for stability.




Figure 6: Decentralized trajectory generation and control. ($r_P = 0.15m$, $r_C = 0.25m$ and $r_S = 2.5m$.)

4 Stability analysis

Our goal in this section is to analyze the stability of formations generated by using the potential field based control modes. We will restrict the discussion in this section to an obstacle free environment. Further, we will not consider the transportation mode. If the robots can successfully organize themselves, it is assumed that the transportation phase can be successfully executed.

We consider the problem of organizing a group of n robots around a goal positioned at r_o as depicted in Figure 7. The robots are considered as particles. We develop the organization scheme by superimposing mutual repulsion behaviors upon individual robot or robot groups. The total energy of the system is a composition of the kinetic energy, the approaching potentials and the repulsive potentials.

$$E = K + V = \frac{1}{2} \sum_{i=1}^n \dot{\mathbf{r}}_i \cdot \dot{\mathbf{r}}_i + \sum_{i=1}^n \left(\frac{1}{2} \sum_{j \neq i}^n V_{ij}^r + V_{io}^a \right) \quad (4)$$




Figure 7: Organizing robots around the goal location

4.1 Conditions for force equilibrium

Let \mathbf{q} be the position vector of the system defined as

$$\mathbf{q} = (r_{1o}, r_{2o}, \dots, r_{no}, \theta_{1o}, \theta_{2o}, \dots, \theta_{no})^T$$

The condition for the force equilibrium is given by

$$\nabla \mathbf{q} V = 0. \quad (5)$$

For potential fields give by Equations (1) and (2), the condition for equilibrium can be expressed as

$$k_{io} r_{io} - \sum_{j \neq i}^n \frac{r_{io} - r_{jo} (\cos(\theta_{io} - \theta_{jo}))}{r_{ij}^{3/2}} k_{ij} = 0 \quad (6)$$

$$\sum_{j \neq i}^n \frac{r_{io} r_{jo} \sin(\theta_{io} - \theta_{jo})}{r_{ij}^{3/2}} k_{ij} = 0 \quad (7)$$

We can solve the above equations for k for a given set of r_{io} and θ_{io} . In other words, we can construct the potential fields to make certain configuration of interest the equilibrium. We can also study the stability of the equilibrium by looking at the positive definiteness of Hessian matrix at such a configuration.




Figure 8: Desired configurations of 2- and 3-robot formations




Figure 9: Desired configurations of 4-robot formations

4.2 $n = 2, 3$

The two- and the three-robot problems result in the stable equilibria shown in Figure 8. The three-robot case can have two different solutions depending on the choice of constants k_o and k_{ij} . Given the choice of constants, the equilibrium is unique and globally stable. This is not shown here because of space limitations. However, the four-robot problem is analyzed next in greater detail.

4.3 $n = 4$

To illustrate we consider a 4-robot ($n = 4$) formation with three configurations as shown in Figure 9.

Case 1: Line shape The position vector for the robots in the straight line shape is given by

$$\mathbf{q}_{line} = (2r_o, r_o, r_o, 2r_o, \theta_o, \theta_o, \theta_o + \pi, \theta_o + \pi)^T.$$

The force equilibrium for the line shape can be achieved by setting the intensities of the potential fields to

$$\begin{cases} k_{io} = k_o & \text{for } i = 1, \dots, 4 \\ k_{42} = k_{31} \\ k_{43} = k_{21} \\ k_{21} = \frac{16k_o r_o^3 + 4k_{32} - k_{41}}{32} \\ k_{31} = \frac{16k_o r_o^3 + 4k_{32} - k_{41}}{32} \end{cases} \quad (8)$$

The eigenvalues of Hessian matrix the for the line shape configuration at the force equilibrium is give by

$$\begin{pmatrix} k_o \\ 3k_o \\ \frac{144k_o r_o^3 + 40k_{32} - 5k_{41}}{48r_o^3} \\ \frac{30k_o r_o^3 + 2k_{32} - k_{41}}{6r_o^3} \\ \frac{-8k_{32} + k_{41}}{12r_o} \\ \frac{-10k_{32} + 5k_{41} - \sqrt{(10k_{32} - 5k_{41})^2 - 96r_o(96k_o^2 r_o^6 + 16k_o r_o^2 k_{32} - 8k_o r_o^2 k_{41})}}{48r_o} \\ \frac{-10k_{32} + 5k_{41} + \sqrt{(10k_{32} - 5k_{41})^2 - 96r_o(96k_o^2 r_o^6 + 16k_o r_o^2 k_{32} - 8k_o r_o^2 k_{41})}}{48r_o} \\ 0 \end{pmatrix}$$

For the Hessian to be positive semi-definite, all of the eigenvalues have to be nonnegative or equivalently, all of the principal minors of the Hessian need to be greater than or equal to zero. We can show that Equation (8) plus the following constraint provide a sufficient condition for the straight line shape to be a stable configuration

$$8k_{32} \leq k_{41} \leq 152k_{32}$$

Case 2: Star shape For the star shape configuration

$$\mathbf{q}_{star} = (0, r_o, r_o, r_o, 0, \theta_o, \theta_o + 2\pi/3, \theta_o + 4\pi/3)^T.$$

Based on the same token in Case 1, we can show that if we construct the potential fields based on the following conditions, the star shape configuration is stable.

$$\begin{cases} k_{io} = k_o & \text{for } i = 1, \dots, 4 \\ k_{21} = k_o r_o^3 - k_{42}/\sqrt{3} \\ k_{31} = k_o r_o^3 - k_{42}/\sqrt{3} \\ k_{41} = k_o r_o^3 - k_{42}/\sqrt{3} \\ k_{32} = k_{42} \\ k_{43} = k_{42} \\ k_{42} = \sqrt{3}(1-c)k_o r_o^3, \quad 0 < c = \text{const.} < 1 \end{cases} \quad (9)$$

Case 3: Square shape For the square shape, the configuration is given by

$$\mathbf{q}_{square} = (r_o, r_o, r_o, r_o, \theta_o, \theta_o + \pi/2, \theta_o + \pi, \theta_o + 3\pi/2)^T$$

Note that this shape is particularly useful for the cooperative manipulation tasks. Because at end of the organization phase, the object is expected to be surrounded by

robots. This will, in general, rule out the line shape and the star shape. For this configuration, a sufficient condition of stability can be obtained as the following:

$$\begin{cases} k_{i0} = k_o & \text{for } i = 1, \dots, 4 \\ k_{32} = 4k_o r_0^3 - 2\sqrt{2}k_{42} \\ k_{41} = 4k_o r_0^3 - 2\sqrt{2}k_{42} \\ k_{21} = k_{42} \\ k_{31} = k_{42} \\ k_{43} = k_{42} \\ k_{42} = c k_o r_0^3, & 0 < c = \text{const.} < \sqrt{2} \end{cases} \quad (10)$$

4.4 N-robot extension

For a N-robot formation surrounding an object, the star shape and line shape are no longer of interest. Instead, we are interested in formations in which the robots are symmetrically organized around the object:

$$\mathbf{q} = \left(r_0, \dots, r_0, \theta_0, \theta_0 + \frac{2\pi}{N}, \dots, \theta_0 + 2\frac{N-1}{N}\pi \right)^T \in \mathbb{R}^{2N \times 1}$$

Note that the repulsive potential between two neighboring robots is effective only when the distance between them is less than the radius of sensing zone. This will simplify the expression for the total potential functions, and similar stability results can be obtained by following the steps in the previous subsection.




Figure 10: N robots surrounding a object

5 Simulation results

Organization of the robots We first study the approach and organization modes. In each case, we look at the robots' ability to approach the object and organize themselves into an appropriate formation that can trap (or cage) the object. mode. The main purpose for this test is to examine the stability of the formations. In this experiment, 10 robots from randomly assigned initial positions approach rectangular objects with different aspect ratios and organize themselves into a formation around the object. The simulation results are shown in Figure 11 with 6 different aspect ratios of the object ranging from 1/7 to 6/2. A unified switching strategy described in Section 3 is used for all the simulations.

Since we assume our robots are identical. The number of robots within the sensing zone of an individual robot will increase as the total number of the robot increases. Also, all mobile robots have velocity bounds, motion for the robot could be specified beyond its performance capabilities if we do not consider such limits when construction the potential fields. To leverage the total potential forces acting on an individual robot, we

use the following heuristic expression of the repulsive potential intensities:

$$k_{ij} = \begin{cases} c_1 (e^{-c_2(n-1)} - e^{-c_2(n_0-1)}), & n \leq n_0 \\ 0, & n > n_0 \end{cases} \quad (11)$$

where $c_{1,2,3}$ are constants depending on the number of the robots and the size of the object. Figure 12 shows an design example of k_{ij} used in our simulations.

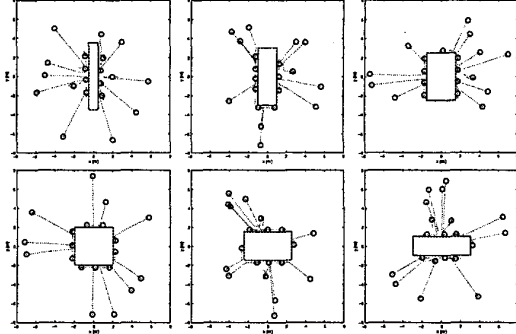


Figure 11: Approach and organization control modes for rectangular objects with different aspect ratios. The constants used in the distribution phase are $r_s = 2.5$, $\alpha = 4$, $r_s = 2.5$, $d_0 = 2$, $p = 16$




Figure 12: Intensity of the repulsive potentials vs. the number of robots

Manipulation We simulate all of the three control modes shown in Figure 5 for $n = 1, \dots, 11$. A previously developed package [13] is used to simulate the dynamics of multiple contact interactions between the object and the robots during a manipulation task. Sample trials are shown in Figure 13. In all of the simulations, the ratio between the intensities of the attractive potential fields for the approach and transportation modes are set to 10.

Performance Based on this framework, an individual robot may have a tendency to fail in any of the three modes. For example, it could get stuck by the obstacles in the approach mode or the transportation mode, or get repelled by other robots in the organization mode. However, the reliability of the system (successful manipulation) improves with the number of robots. This described in Figure 14




Figure 13: Sample trials by using decentralized controllers to get autonomous robots to surround an object and push the object to a desired destination

6 Discussion

In this work we described a completely decentralized framework for the development of intelligent multi-robot manipulation tasks. We assume that each robot has approximate information about the object position, its goal position, and the number of team members, and is equipped with an omnidirectional camera. The sensor has a limited range, but the robots can see the neighbors in this range. We show that decentralized, potential-field based controllers can be used to approach the object, organize the robots into a formation that will trap the object, and then transport the object to the destination. This paper complements experimental studies conducted by Spletzer *et al* [14] on 3- and 4-robot teams.

Theoretical guarantees are harder to come by. We show how we can derive results on formation stability with potential field based controllers in many cases. Simulation results demonstrate successful application to a wide range of examples without showing sensitivity to parameters. Because the framework is decentralized at both trajectory generation level and the estimation and control agent level, our framework can potentially scale to groups of tens and hundreds of robots.




Figure 14: Reliability improves with the number of robots, where μ is the coefficient of friction

Acknowledgments

This work was supported by the DARPA ITO MARS Program, grant no. 130-1303-4-534328, AFOSR grant no. F49620-01-1-0382, and NSF grant no. CDS-97-03220. We thank Aavek Das, Rafael Fierro and Zhi-dong Wang for discussions on multi-robot cooperation and control.

References

- [1] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 476–481, San Francisco, CA, April 2000.
- [2] A. Das, R. Fierro, V. Kumar, J. Ostrowski, J. Spletzer, and C. Taylor. A framework for vision based formation control. Submitted to *IEEE Trans. Robot. Automat.*, March 2001.
- [3] J. Feddema and D. Schoenwald. Decentralized control of cooperative robotic vehicles. In *Proc. SPIE Vol. 4364, Aerosense*, Orlando, Florida, April 2001.
- [4] R. Fierro, A. Das, V. Kumar, and J. P. Ostrowski. Hybrid control of formation of robots. *IEEE Int. Conf. Robot. Automat.*, May 2001.
- [5] R. Fierro, P. Song, A. Das, and V. Kumar. Cooperative control of robot formations. In *Series on Applied Optimization*, R. Murphey and P. Pardalos (eds.), Kluwer Academic Press, 2001.
- [6] L. Iochhi, K. Konolige, and M. Bayracharya. A framework and architecture for multi-robot coordination. In *Proc. Seventh Int. Symposium on Experimental Robotics (ISER)*, Honolulu, Hawaii, Dec. 2000.
- [7] J. S. Jennings, G. Whelan, and W. F. Evans. Cooperative search and rescue with a team of mobile robots. *Proc. IEEE Int. Conf. on Advanced Robotics*, 1997.
- [8] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *International Journal of Robotics Research*, 5:90–98, 1986.
- [9] M. Mataric, M. Nilsson, and K. Simsarian. Cooperative multi-robot box pushing. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 556–561, Aug 1995.
- [10] L. E. Parker. Current state of the art in distributed autonomous mobile robotics. In L. E. Parker, G. Bekey, and J. Barhen, editors, *Distributed Autonomous Robotic Systems*, volume 4, pages 3–12. Springer, Tokyo, 2000.
- [11] E. Rimon and A. Blake. Caging 2D bodies by one-parameter two-fingered gripping systems. In *IEEE Int. Conf. on Robotics and Automation*, pages 1458–1464, Minneapolis, MN, Apr 1996.
- [12] P. Song, P. Kraus, V. Kumar, and P. Dupont. Analysis of rigid-body dynamic models for simulation of systems with frictional contacts. *ASME Journal of Applied Mechanics*, 68:118–128, January 2001.
- [13] P. Song, M. Yashima, and V. Kumar. Dynamic simulation for grasping and whole arm manipulation. In *Proc. of the 2000 IEEE Int'l Conf. on Robotics and Automation*, Apr. 2000.
- [14] J. Spletzer, A. Das, R. Fierro, C. J. Taylor, V. Kumar, and J. P. Ostrowski. Cooperative localization and control for multi-robot manipulation. *IEEE/RSJ Int. Conf. Intell. Robots and Syst., IROS2001*, March 2001.
- [15] A. Sudsang and J. Ponce. A new approach to motion planning for disc-shaped robots manipulating a polygonal object in the plane. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 1068–1075, San Francisco, CA, April 2000.
- [16] T. Sugar and V. Kumar. Control and coordination of multiple mobile robots in manipulation and material handling tasks. In P. Corke and J. Trevelyan, editors, *Experimental Robotics VI: Lecture Notes in Control and Information Sciences*, volume 250, pages 15–24. Springer-Verlag, 2000.
- [17] R. Volpe and P. Khosla. Manipulator control with superquadric artificial potential functions: Theory and experiments. *IEEE Trans. on Syst., Man, and Cyber.*, 20(6):1423–1436, 1990.
- [18] Z. Wang and V. Kumar. Object closure and manipulation by multiple cooperating mobile robots. *2002 IEEE Int. Conf. Robot. Automat.*, to appear, May 2002.