Mobile Manipulation: Collision-Free Path Modification and Motion Coordination

Oliver Brock Oussama Khatib
Robotics Laboratory, Department of Computer Science
Stanford University, Stanford, California 94305
e-mail: {oli, khatib}@cs.stanford.edu

Abstract

In many robotic applications motions must be executed robustly in dynamic and partially unknown environments. Despite this requirement most motion planning algorithms assume the environment to be known and changes to be predictable. The planning problem in dynamic environments can be decomposed into a planning and an execution phase. In this paper we describe a new framework for the execution of motion plans for robots with many degrees of freedom in dynamic environments. An initial valid trajectory is incrementally modified according to changes in the environment to maintain a collision free path. This framework is particularly well suited for redundant systems and mobile manipulation, as it allows motion specification of a subset of the degrees of freedom of the robot, greatly simplifying the task of robot programming.

1 Introduction

The robust execution of planned robot motion in dynamic environments and in the presence of uncertainty is essential for achieving intelligent robot behavior. In mobile manipulation this problem becomes particularly challenging because the execution of a task has to be integrated with reactive collision avoidance for robots with many degrees of freedom.

Reactive collision avoidance for manipulators or mobile platforms is a well-studied problem [9, 6, 1, 10]. Due to the recent interest in mobile manipulation more advanced methods are being investigated to incorporate obstacle avoidance with limited coordination between manipulator and mobile base [3, 15, 17, 12].

Most of the motion planning algorithms for robots are based on the assumption that changes in the environment during planning and execution are known a priori [10]. However, most applications in robotics require robots to operate in environments that change unpredictably. To allow the execution of motion plans in dynamic environments the global planning process can be augmented with a fast, reactive obstacle avoidance component.

With this augmentation, an initial motion plan can be generated under the assumption that all obstacles are known [16]. During execution the robot deviates from the planned path, guided by potential fields caused by previously unknown or moving obstacles. The planner attempts to rejoin the original path at a later point that remained unaffected by changes in the environment. To escape local minima of the potential function fast local replanning can be applied [4]. These approaches have only been applied to low-dimensional configuration spaces.

Deviating from a pre-planned path, however, can result in local minima that require replanning [4, 16]. To overcome this problem the concept of elastic bands was introduced [14]. A path is represented as a curve in configuration space, called elastic band. It is incrementally modified according to the position of obstacles in the environment to maintain a smooth and collision-free path. The concept of elastic bands was extended to nonholonomic robots [5].

The first planner for dynamic environments that has been shown in practice to achieve real-time performance addresses the problem of planning for two cooperating robot arms [11]. The planning task is decomposed into lower-dimensional problems that can be solved efficiently. The extension to higher-dimensional problems cannot preserve the real-time performance of these methods.

In this paper we introduce a new framework, called the *elastic strip*, that allows the execution of motion plans for robots with many degrees of freedom in dynamic environments [2]. This approach is based on the idea of decomposing the problem of motion planning in dynamic environments [16, 14] into a global planning phase, in which certain assumptions about the environment are made and a reactive execution phase, in which the planned path is updated according to changes in the environment.

The reactive modification of trajectories, however, does not completely eliminate the need for replanning. It allows continuous execution of motion until

a new motion plan has been computed. This new plan can then be used to update the elastic strip.

2 Elastic Strip Framework

2.1 Overview

The basic idea of the elastic strip framework is similar to elastic bands [14]. We assume that a collision-free path for a planning problem has been obtained using an arbitrary planning method. The resulting trajectory is incrementally modified by external forces originating from obstacles in the environment and internal forces applied to shorten and smoothen it.

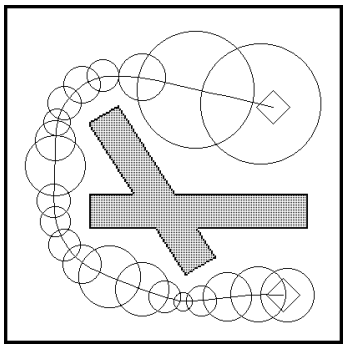

Figure 1: Bubbles on an elastic band

The key to efficiency of the elastic band is its use of a local description of the free space around each configuration. With only one distance computation in the workspace a hypersphere of free configuration space is computed. This hypersphere is called *bubble* [13]. By covering the entire elastic band with overlapping bubbles we can guarantee that the path it represents is collision-free. This is illustrated in Figure 1.
The difficulty with the elastic band lies in its use of the configuration space representation to describe the local free space around a robot configuration. As the number of degrees of freedom of the robot increases, the estimate of the local free space around a configuration becomes extremely conservative. An excessive amount of bubbles are then needed to cover the elastic band, affecting the real-time performance of the algorithm.
To minimize the impact of the dimensionality of the configuration space on the computational complexity of the algorithm and thereby its real-time performance, we choose to represent the trajectory in the robot workspace. This trajectory must be maintained in such a way that the volume swept by the robot during execution of the trajectory remains in free space. This volume can be imagined as a strip of elastic material that is deformed by approaching obstacles and shortens as obstacles recede.

The properties of an elastic strip differ from elastic material. Since we are not restrained by the physical laws of elasticity, we choose the effects of forces on the elastic strip to implement some desired behavior during reactive collision avoidance.

2.2 Rigid Body Representation

During the execution of a trajectory a robot sweeps out a volume in the workspace. The trajectory is collision-free if no obstacle is inside this volume. To warrant collision avoidance the volume has to be computed using a model of the rigid bodies in motion and checked against obstacles in the environment.

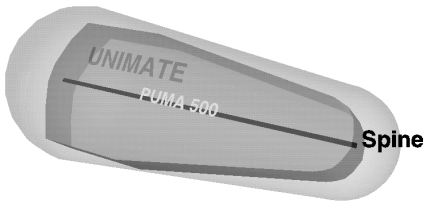

Figure 2: Illustration of a spine

A rigid body can be approximated by a line segment that is parameterized by a varying width. We call that line segment the *spine* of the rigid body. The width specifies the free space required around the line segment for the body to be free of collision. In Figure 2 the spine of a transparent PUMA 560 link is shown as the black line along its central axis; the volume associated with the spine encloses the link.
As a rigid body comes closer to obstacles, the model described above might result in incorrect collision detection. To address this difficulty the representation of rigid bodies can be generalized to describe rigid bodies at an arbitrary level of detail.

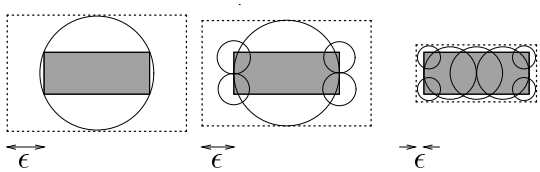

Figure 3: Covering a body with spines

The basic idea of this generalization is to introduce more spines to cover the volume of the rigid body with increasing accuracy as the distance to obstacles decreases. In Figure 3 the rectangular cross section of a body and its circular covering by spines are shown for different resolutions δ_i, indicated by the dotted region. The three pictures could also be interpreted as the evolution of the volume description as the rigid body approaches an obstacle at distance δ_i.
For the sake of simplicity we assume in the remainder of this paper that the volume of a rigid body is described by a single spine.

2.3 Representing Local Free Space

To compute a bubble in the elastic band approach, a distance measurement to a point in the workspace has to be translated into a hypersphere in configuration space. When the free space is represented in the workspace, the distance computation translates directly into a bubble in the workspace. Let $\rho(\mathbf{p})$ be the function that computes the minimum distance from a point $\mathbf{p}$ to any obstacle. The *workspace bubble* of free space around $\mathbf{p}$ is defined as

$$\mathcal{B}(\mathbf{p}) = \{\, \mathbf{q} : \|\mathbf{p} - \mathbf{q}\| < \rho(\mathbf{p})\}.$$

An approximation of the local free space around a rigid body b in configuration q can be computed by generating a set of workspace bubbles centered on the spine. This set of bubbles is called *protective hull* $\mathcal{P}_q^b$. The local free space or protective hull $\mathcal{P}_q^{\mathcal{R}}$ of a robot $\mathcal{R}$ at a configuration q is described by the union of protective hulls of each rigid body of $\mathcal{R}$,

$$\mathcal{P}_q^{\mathcal{R}} = \bigcup_{b \in \mathcal{R}} \mathcal{P}_q^b.$$

Figure 4 shows a protective hull of the *Stanford Mobile Platform*. Note that a single workspace bubble may contain multiple rigid bodies or even the entire robot, implying that for large clearances a simple description of the local free space suffices.

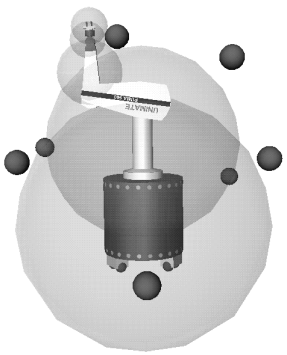

Figure 4: Protective hull of the Stanford Mobile Platform amidst spherical obstacles

2.4 Connectedness of Elastic Strip

An elastic strip $\mathcal{S}_{\mathcal{T}}^{\mathcal{R}} = (q_1, q_2, q_3, \cdots, q_n)$ is a sequence of configurations q_i on the trajectory $\mathcal{T}$ of the robot $\mathcal{R}$. The local free space of a configuration is described by the protective hulls $\mathcal{P}_i^{\mathcal{R}}$.
Since each configuration q_i is guaranteed to be free of collisions by the protective hull $\mathcal{P}_{q_i}^{\mathcal{R}}$, it remains to be shown that the union of all protective hulls contains the volume $V_{\mathcal{T}}^{\mathcal{R}}$ swept by the robot along the trajectory. The condition of feasibility of trajectory $\mathcal{T}$ described by $\mathcal{S}_{\mathcal{T}}^{\mathcal{R}}$ is

$$V_{\mathcal{T}}^{\mathcal{R}} \subseteq V_{\mathcal{S}}^{\mathcal{R}} = \bigcup_{1 \le i \le n} \mathcal{P}_i^{\mathcal{R}}. \qquad (1)$$

This is illustrated in Figure 5, where three consecutive protective hulls cover the trajectory of the robot. The initial and the final configuration are shown. An obstacle is reducing the size of the intermediate protective hull.

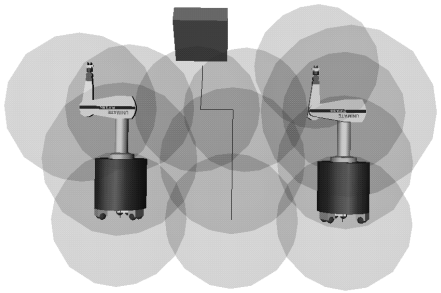

Figure 5: Protective hulls covering a trajectory

It suffices to describe a procedure that verifies the existence of a path between two consecutive protective hulls $\mathcal{P}_i^{\mathcal{R}}$ and $\mathcal{P}_{i+1}^{\mathcal{R}}$. By applying this procedure repeatedly the condition of feasibility (1) can be ensured.
We will make the assumption that every point on a rigid body b moves on a straight line as b transitions from q_i to q_{i+1}. This ignores the effect of rotation. However, this effect can be bounded and taken into account at a computational expense, when computing the protective hull of b. The justification for this assumption is that two adjacent configurations will be similar enough for this effect to be insignificant when the robot is close to an obstacle. This is a simplification but not an inherent limitation of the approach.
Using this assumption the path of each rigid body b can be examined independently. If a trajectory between q_i and q_{i+1} exists for all rigid bodies $b \in \mathcal{R}$, one exists for $\mathcal{R}$.
The existence of a trajectory $\mathcal{T}_{i,i+1}$ for a rigid body b from configuration q_i to q_{i+1} is guaranteed if the volume $V_{\mathcal{T}_{i,i+1}}^b$ swept by b along $\mathcal{T}_{i,i+1}$ is contained within the protective hulls of the configuration q_i and q_{i+1},

$$V_{\mathcal{T}_{i,i+1}}^b \subseteq \left(\mathcal{P}_i^b \cup \mathcal{P}_{i+1}^b\right). \qquad (2)$$

To verify condition (2) the union $\mathcal{U} = \mathcal{P}_i^b \cup \mathcal{P}_{i+1}^b$ is examined. If b can pass through $\mathcal{U}$ on a straight line trajectory form q_i to q_{i+1} then the existence of a trajectory $V_{\mathcal{T}_{i,i+1}}^b$ contained within $\mathcal{P}_i^b \cup \mathcal{P}_{i+1}^b$ is guaranteed.
If for all rigid bodies $b \in \mathcal{R}$ the union of their protective hulls $\mathcal{P}_i^b \cup \mathcal{P}_{i+1}^b$ is large enough to allow a straight-line trajectory, we say that two consecutive protective hulls $\mathcal{P}_i^{\mathcal{R}}$ and $\mathcal{P}_{i+1}^{\mathcal{R}}$ are *connected*.

2.5 Forces Acting on the Strip

Whereas elastic material is homogeneous and its principal physical properties do not vary over its volume,

this is not true for an elastic strip. In this framework an elastic strip can be seen as a two-dimensional grid of links and springs. Figure 6 illustrates that for the elastic strip $\mathcal{S} = (q_1, q_2, q_3)$ for an arm mounted on a mobile base.

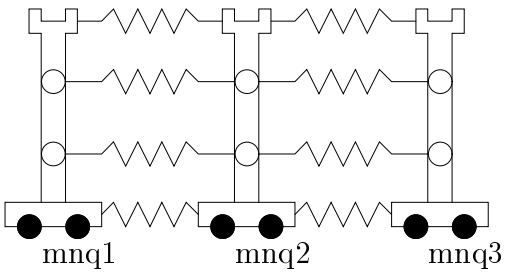

Figure 6: Principal structure of elastic strip

The internal forces acting on the elastic strip are generated by the virtual springs attached to control points in subsequent configurations along the trajectory. Let $\mathbf{p}_j^i$ be the position vector of the origin of the frame attached to the j-th joint of the robot in configuration q_i. We use these points as control points. The internal contraction force $\mathbf{F}_{i,j}^{int}$ caused by the springs attached to joint j is defined as

$$\mathbf{F}_{i,j}^{int} = k_c \left(\frac{d_j^{i-1}}{d_j^{i-1} + d_j^i} (\mathbf{p}_j^{i+1} - \mathbf{p}_j^{i-1}) - (\mathbf{p}_j^i - \mathbf{p}_j^{i-1}) \right),$$

where d_j^i is the distance $\|\mathbf{p}_j^i - \mathbf{p}_j^{i+1}\|$ in the initial, unmodified trajectory and k_c is a constant determining the contraction gain of the elastic strip.
These forces cause the elastic strip to contract, maintaining a constant ratio of distances between every three consecutive configurations. Note that the force acting on the control points depends only on the local curvature of the elastic strip and not on its elongation.
The external forces are caused by a repulsive potential associated with the obstacles. For a point $\mathbf{p}$ this potential function is defined as

$$V_{ext}(\mathbf{p}) = \begin{cases} \frac{1}{2} k_r (d_0 - d(\mathbf{p}))^2 & \text{if } d(\mathbf{p}) < d_0 \\ 0 & \text{otherwise} \end{cases},$$

where $d(\mathbf{p})$ is the distance from $\mathbf{p}$ to the closest obstacle, d_0 defines the region of influence around obstacles, and k_r is the repulsion gain.
The external force $\mathbf{F}_{\mathbf{p}}^{ext}$ acting at point $\mathbf{p}$ is defined by the gradient of the potential function at that point:

$$\mathbf{F}_{\mathbf{p}}^{ext} = -\nabla V_{ext} = k_r (d_0 - d(\mathbf{p})) \frac{\mathbf{d}}{\|\mathbf{d}\|},$$

where $\mathbf{d}$ is the vector between $\mathbf{p}$ and the closest point on the obstacle.

2.6 Elastic Strip Modification

Let $\mathcal{S} = (q_1, q_2, q_3, \cdots, q_n)$ be an elastic strip. When $\mathcal{S}$ is subjected to the forces described in section 2.5, it is deformed by altering each of the configurations q_i in turn. To change a configuration according to the internal and external forces, these forces have to be mapped to joint torques.
For collision avoidance in the absence of a task requirement, we use the Jacobian $J_{\mathbf{p}}$ associated with the point $\mathbf{p}$ at which the force $\mathbf{F}_{\mathbf{p}}$ is acting for this mapping. The joint torques $\mathbf{\Gamma}$ caused by $\mathbf{F}_{\mathbf{p}}$ are given by

$$\mathbf{\Gamma} = J_{\mathbf{p}}^T \mathbf{F}_{\mathbf{p}}. \tag{3}$$

Joint limits can be avoided using a potential field function [6].
The dynamic model of the system can be used to compute the joint displacements caused by the joint torques. The displacements for a configuration q_i define the new configuration q_i', resulting in the modified elastic strip $\mathcal{S}' = (q_1, \cdots, q_i', \cdots, q_n)$. $\mathcal{S}'$ represents a valid trajectory, only if the protective hulls $\mathcal{P}_{i-1}$, $\mathcal{P}_i'$, and $\mathcal{P}_{i+1}$ are connected. This is verified using the procedure described in section 2.4.
If $\mathcal{P}_i'$ and $\mathcal{P}_{i+1}$ are not connected, the elastic strip $\mathcal{S}'$ becomes invalid. This means that the trajectory represented by $\mathcal{S}'$ cannot be proven to be collision-free, using the representation of local free space associated with $\mathcal{S}'$. In order to reconnect $\mathcal{P}_i'$ and $\mathcal{P}_{i+1}$ intermediate protective hulls are inserted into the elastic strip.
As obstacles recede from the vicinity of the elastic strip, the protective hulls of configurations increase in volume and potentially move closer together. This can result in protective hulls $\mathcal{P}_{i-1}$ and $\mathcal{P}_{i+1}$ to be connected. In that case q_i is redundant and can be removed from $\mathcal{S}$.

2.7 Motion Behavior

Given a planned motion, the elastic strip allows a robot to dynamically modify its motion to accommodate changes in the environment. For redundant mechanisms this modification is not uniquely determined and may be chosen depending on the task. A transportation task for a mobile manipulator, for instance, can be described by the motion of the mobile base, while only a nominal posture of the arm and load are specified. For a manipulation task, the description consists of the motion of the end effector and its contact forces, while only a nominal posture of the mobile base and arm is given. In both cases some degrees of freedom are used for task execution, while others can be used to achieve task-independent motion behavior.
The elastic strip also provides an effective framework for executing partially described task. If only those

degrees of freedom necessary for execution have been specified, reactive obstacle avoidance combined with an attractive potential to the desired posture can complete the robot control in real-time. With a partial plan, however, the elastic strip can be subjected to local minima.

3 Motion Execution

The framework for combining motion behavior and task execution relies on the operational space dynamic control structure of redundant robots, summarized below. The operational dynamics for a robot are [7]

$$\Lambda(\mathbf{x})\ddot{\mathbf{x}} + \mu(\mathbf{x}, \dot{\mathbf{x}}) + \mathbf{p}(\mathbf{x}) = \mathbf{F}, \tag{4}$$

where $\mathbf{x}$ is the vector of the m operational coordinates describing the position and orientation of the effector, and $\Lambda(\mathbf{x})$ is the $m \times m$ kinetic energy matrix associated with the operational space. $\mu(\mathbf{x}, \dot{\mathbf{x}})$, $\mathbf{p}(\mathbf{x})$, and $\mathbf{F}$ are respectively the centrifugal and Coriolis force vector, gravity force vector, and generalized force vector acting in operational space. These equations of motion describe the dynamic response of a manipulator to the application of an operational force $\mathbf{F}$ at the end effector. For non-redundant manipulators, the relationship between operational forces $\mathbf{F}$ and joint forces $\mathbf{\Gamma}$ is

$$\mathbf{\Gamma} = J^T(\mathbf{q})\mathbf{F}, \tag{5}$$

where $J(\mathbf{q})$ is the Jacobian matrix. For redundant systems the above relationship becomes

$$\mathbf{\Gamma} = J^T(\mathbf{q})\mathbf{F} + \left[I - J^T(\mathbf{q})\overline{J}^T(\mathbf{q})\right]\mathbf{\Gamma}_0, \tag{6}$$

with

$$\overline{J}(\mathbf{q}) = A^{-1}(\mathbf{q})J^T(\mathbf{q})\Lambda(\mathbf{q}), \tag{7}$$

where $\overline{J}(\mathbf{q})$ is the dynamically consistent generalized inverse, $A(\mathbf{q})$ the joint space mass matrix, and $\Lambda(\mathbf{q})$ the mass matrix associated with the operational space [7, 8].

Equation 6 provides a decomposition of the joint torques into those caused by forces at the end effector ($J^T\mathbf{F}$) and those that only affect internal motion $\left(\left[I - J^T(\mathbf{q})\overline{J}^T(\mathbf{q})\right]\mathbf{\Gamma}_0\right)$. This decomposition can be used to achieve different kinds of task behavior. Simple obstacle avoidance without the incorporation of task behavior can be achieved by using equation 3 to map internal and external forces to joint torques.

To ensure the execution of a task specified in a particular task frame f, the internal and external forces are mapped into the null space of the Jacobian J_f associated with the task frame. This corresponds to the sets of tasks where the end effector is required to move on a certain trajectory and the redundant degrees of freedom can be used for obstacle avoidance.

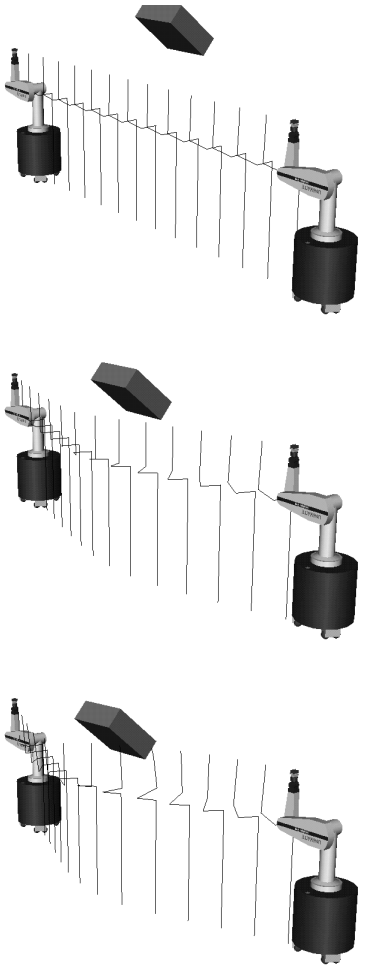

Figure 7: Elastic strip for a 9-dof robot incrementally modified by a moving obstacle

Simple obstacle avoidance behavior can be easily augmented by specifying a desired posture for the robot. This posture can be chosen according to some optimization criterion, for example, to minimize the torques to support the load.

4 Implementation

The results of a preliminary implementation are shown in Figure 7. The framework is applied to the Stanford Mobile Platform, a PUMA 560 robot arm mounted on a holonomic mobile base with a total of nine degrees of freedom. In this example, all links of the robot are subjected to the moving obstacle. The elastic strip is represented by a set of intermediate configurations, displayed as lines connecting joint frames.

The approaching obstacle deforms the elastic strip to ensure obstacle avoidance. As the obstacle moves away, internal forces cause the elastic strip to assume the straight line trajectory. Our current implemen-

tation achieves an update rate for the elastic strip in Figure 7 of 10 to 15 Hz on an SGI Indigo2.

5 Conclusion

The elastic strip framework is an efficient approach to real-time obstacle avoidance for robots with many degrees of freedom. It is particularly well suited for mobile manipulation, since it allows the integration of task-level behavior with collision avoidance in dynamic and uncertain environments.
An elastic strip represents the workspace volume swept by a robot along a pre-planned trajectory. This representation is incrementally modified by external, repulsive forces originating from obstacles to maintain a collision-free path. Internal forces act on the elastic strip to shorten and smoothen the trajectory. To represent the volume swept by the robot the notion of a protective hull is introduced. A protective hull represents the local free space around a configuration and can be computed efficiently. A sequence of these hulls is used to contain the robot at any point along the trajectory, thus ensuring collision avoidance. The effectiveness of the elastic strip framework is demonstrated in real-time path modification for the Stanford Mobile Platform, a nine degree of freedom manipulator, in a dynamic environment.

Acknowledgments

The authors would like to thank Diego Ruspini and Kyong-Sok Chang for their helpful insights and discussion in preparing this paper. The financial support of Boeing, Nomadic Technologies, General Motors, and NSF (grant IRI-9320017) is gratefully acknowledged.

References

[1] Ronald C. Arkin. Motor schema-based mobile robot navigation. *Int. J. of Robotics Research*, 8(4):92–112, August 1987.

[2] Oliver Brock and Oussama Khatib. Elastic strips: Real-time path modification for mobile manipulation. In *Proc. of the 8th Int. Symp. of Robotics Research*, pages 117–122, 1997. Preprints.

[3] Wayne F. Carriker, Pradeep K. Khosla, and Bruce H. Krogh. An approach for coordinating mobility and manipulation. In *IEEE Int. Conf. on Systems Engineering*, pages 59–63, 1989.

[4] Wonyun Choi and Jean-Claude Latombe. A reactive architecture for planning and executing robot motions with incomplete knowledge. In *Proc. of the IEEE/RSJ Int. Workshop on Intelligent Robots and Systems*, volume 1, pages 24–29, 1991.

[5] Maher Khatib, H. Jaouni, R. Chatila, and J.-P. Laumond. How to implement dynamic paths. In *Proc. of the Int. Symp. on Experimental Robotics*, pages 225–36, 1997. Preprints.

[6] Oussama Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *Int. J. of Robotics Research*, 5(1):90–8, 1986.

[7] Oussama Khatib. A unified approach to motion and force control of robot manipulators: The operational space formulation. *IEEE J. of Robotics and Automation*, 3(1):43–53, 1987.

[8] Oussama Khatib, Kazu Yokoi, Kyong-Sok Chang, Diego Ruspini, Robert Holmberg, and Arancha Casal. Coordination and decentralized cooperation of multiple mobile manipulators. *Journal of Robotic Systems*, 13(11):755–64, 1996.

[9] Bruce H. Krogh. A generalized potential field approach to obstacle avoidance control. In *Robotics Research*, pages MS84–484/1–15, 1984.

[10] Jean-Claude Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.

[11] Tsai-Yen Li and Jean-Cleaude Latombe. On-line manipulation planning for two robot arms in a dynamic environment. *Int. J. of Robotics Research*, 16(2):144–167, 1997.

[12] U.M. Nassal. Motion coordination and reactive control of autonomous multi-manipulator systems. *J. of Robotic Systems*, 13(11):737–54, 1996.

[13] Sean Quinlan. *Real-Time Modification of Collision-Free Paths*. PhD thesis, Stanford University, 1994.

[14] Sean Quinlan and Oussama Khatib. Elastic bands: Connecting path planning and control. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, volume 2, pages 802–7, 1993.

[15] Homayoun Seraji. An on-line approach to coordinated mobility and manipulation. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, volume 1, pages 28–35, 1993.

[16] John P. H. Steele and Gregory P. Starr. Mobile robot path planning in dynamic environments. In *Proc. of the IEEE Int. Conf. on Systems, Man, and Cybernetics*, volume 2, pages 922–925, 1988.

[17] Yoshio Yamamoto and Xiaoping Yun. Coordinated obstacle avoidance of a mobile manipulator. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, volume 3, pages 2255–60, 1995.