
O. Khatib
K. Yokoi*
O. Brock
K. Chang
A. Casal

Robotics Laboratory
Department of Computer Science
Stanford University, Stanford, CA 94086, USA
ok@Robotics.Stanford.edu

Robots in Human Environments: Basic Autonomous Capabilities

Abstract

This article discusses the basic capabilities needed to enable robots to operate in human-populated environments for accomplishing both autonomous tasks and human-guided tasks. These capabilities are key to many new emerging robotic applications in service, construction, field, underwater, and space. An important characteristic of these robots is the "assistance" ability they can bring to humans in performing various physical tasks. To interact with humans and operate in their environments, these robots must be provided with the functionality of mobility and manipulation. The article presents developments of models, strategies, and algorithms concerned with a number of autonomous capabilities that are essential for robot operations in human environments. These capabilities include: integrated mobility and manipulation, cooperative skills between multiple robots, interaction ability with humans, and efficient techniques for real-time modification of collision-free path. These capabilities are demonstrated on two holonomic mobile platforms designed and built at Stanford University in collaboration with Oak Ridge National Laboratories and Nomadic Technologies.

1. Introduction

A new field of robotics is emerging. Robots are today moving toward applications beyond the structured environment of a manufacturing plant. They are making their way into the everyday world that people inhabit – hospitals, offices, homes, construction sites (Engelberger 1991; Schmidt, Hanebeck, and Fischer 1997; Schraft and Hägele 1997), and other cluttered and uncontrolled environments. While advancing into these new areas, the current generation of service and field robots suffers major shortcomings because of their limited

abilities for manipulation and interaction with humans. Their operations are mostly concerned with transportation, and rarely involve more than the simplest manipulation tasks.

The successful introduction of robots into human environments will rely on the development of competent and practical systems that are dependable, safe, and easy to use. The value of their contribution to the work environment will have to be unquestionable and their task performance must be as reliable as that of a human worker. Typical operations are composed of various tasks, some of which are sufficiently structured to be autonomously performed by a robotic system, while many others require skills that are still beyond current robot capabilities. Today, these tasks can only be executed by a human worker. The introduction of a robot to assist a human in such tasks will reduce fatigue, increase precision, and improve quality; whereas the human can bring experience, global knowledge, and understanding to the execution of task. The synergy of the human/robot team can greatly increase overall performance by fully utilizing their complementary abilities in the completion of the task.

During an assistance task, the robot must be capable of performing basic autonomous operations involving both navigation and manipulation. For more elaborate and delicate operations, the assistant, in its supporting role, must be able to interact and cooperate with the human when performing a guided task. The discussion in this article focuses on the basic capabilities needed for manipulation and posture behaviors, cooperation between multiple robots, interaction with the humans, and efficient techniques for real-time collision-free path modifications.

The development of robots in human environments will depend largely on the full integration of mobility and manipulation. Mobile manipulation is a relatively new research area. There is, however, a large body of work devoted to the study of motion coordination in the context of kinematic redundancy. In recent years, these two areas have begun to

*K. Yokoi is currently with the Dept. of Robotics at MEL, Tsukuba, Japan.
The International Journal of Robotics Research
Vol. 18, No. 7, July 1999, pp. 684-696,
©1999 Sage Publications, Inc.

merge, and algorithms developed for redundant manipulators are being extended to mobile manipulation systems (Yamamoto and Yun 1994; Cameron et al. 1993; Umetani and Yoshida 1989; Ullman and Cannon 1989; Papadopoulos and Dubowsky 1991). Typical approaches to motion coordination of redundant systems rely on the use of pseudo or generalized inverses to solve an under-constrained or degenerate system of linear equations, while optimizing some given criterion. These algorithms are essentially driven by kinematic considerations and the dynamic interaction between the end effector and the robot's self motions are ignored.

Our effort in this area has resulted in a *task-oriented* framework for the *dynamic coordination* (Khatib et al. 1996) of mobile manipulator systems. The dynamic coordination strategy we developed is based on two models concerned with the effector dynamics (Khatib 1987) and the robot self-posture behavior. The *effector dynamic behavior* model is obtained by a projection of the robot dynamics onto the space associated with the task, while the *posture behavior* is characterized by the complement of this projection. To control these two behaviors, a consistent control structure is required. The article discusses these models and presents a unique control structure that guarantees *dynamic consistency* and decoupled posture control (Khatib 1995), while providing optimal responsiveness at the effector.

Another important issue in mobile manipulation concerns the development of effective cooperation strategies for multiple robots (Zheng and Luh 1986; Uchiyama and Dauchez 1988; Hayati 1987; Tarn, Bejczy, and Yun 1987; Adams et al. 1995; Jung, Cheng, and Zelinsky 1997). Our earlier work on multi-arm cooperation established the *augmented object* model, describing the dynamics at the level of manipulated object (Khatib 1988), and the *virtual linkage* model (Williams and Khatib 1993), characterizing internal forces. Effective implementation of cooperative manipulation relies on the availability of a high-rate force sensory feedback from the cooperating robots to the controller. While force feedback is easily accessible for multi-arm systems, the access to this data is difficult for mobile platforms. The article presents a *decentralized* cooperation strategy that is consistent with the *augmented object* and *virtual linkage* models, preserving the overall performance of the system.

A robotic system must be capable of sufficient level of competence to avoid obstacles during motion. Even when a path is provided by a human or other intelligent planner, sensor uncertainties and unexpected obstacles can make the motion impossible to complete. Our research on the artificial potential field method (Khatib 1986) has addressed this problem at the control level to provide efficient real-time collision avoidance. Due to their local nature, however, reactive methods (Khatib 1986; Krogh 1984; Arkin 1987; Latombe 1991) are limited in their ability to deal with complex environments. Using navigation functions (Koditschek 1987) the problems arising from the locality of the potential field approach can

be overcome. These approaches, however, do not extend well to robots with many degrees of freedom, such as mobile manipulators (Carriker, Khosla, and Krogh 1989; Seraji 1993; Yamamoto and Yun 1995). Our investigation of a framework to integrate real-time collision avoidance capabilities with a global collision-free path has resulted in the *elastic band* approach (Quinlan and Khatib 1993), which combines the benefits of global planning and reactive systems in the execution of motion tasks. The concept of elastic bands was also extended to nonholonomic robots (M. Khatib et al. 1997). The article discusses our ongoing work in this area and presents a novel approach, the *elastic strip* (Brock and Khatib 1997), which allows the robot's free space to be computed and represented directly in its workspace rather than in its high-dimensional configuration space. The resulting algorithms are computationally efficient and can easily be applied to robots with many degrees of freedom.

The discussion in this article focuses on the various methodologies developed for the integration of mobility and manipulation, the cooperation between multiple robotic platforms, the interaction between humans and robots, and the real-time modification of collision-free paths. The article also presents the implementation of these developments on the Stanford Robotic Platforms, shown in Fig. 1.

2. Integration of Mobility and Manipulation

The ability to interact with the environment is an important capability for robotic systems; grabbing, lifting, pushing, and manipulating objects, while maneuvering to reach, avoid collision, and navigate in the workspace. The control of the two functionalities, mobility and manipulation, must address both their complex kinematic coordination, and their strong dynamic interaction and coupling. Another critical aspect of




Fig. 1. The Stanford robotics platforms: two holonomic platforms, each is equipped with a PUMA 560 arm, various sensors, two computer systems, a multi-axis controller, and sufficient battery power to allow for autonomous operation.

mobile manipulation dynamics is the higher requirements manipulation tasks have on the robot responsiveness compared with those of mobility.

Mobile manipulator systems share many of the characteristics of macro/mini structures (Khatib 1995): coarse and slow dynamic responses of the mobile base (the macro mechanism), and the relatively fast responses and higher accuracy of the manipulator (the mini device). Inspired by these properties of macro/mini structures, we have developed a framework for the coordination and control of mobile manipulator systems. This framework provides a unique control structure for decoupled manipulation and posture control, while achieving optimal responsiveness at the effector. This control structure is based on two models concerned with the effector dynamic behavior and the robot self-posture behavior. The *effector dynamic behavior model* is obtained by a projection of the robot dynamics into the space associated with the effector task, and the *posture behavior* model is characterized by the complement of this projection.

We first present the basic models associated with the end effector. In a subsequent section we present the vehicle arm coordination strategy and posture control behavior.

2.1. Effector Dynamic Behavior

The joint space dynamics of a manipulator are described by

$$A(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{b}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{g}(\mathbf{q}) = \mathbf{\Gamma}, \quad (1)$$

where \mathbf{q} is the vector of n joint coordinates, $A(\mathbf{q})$ is the $n \times n$ kinetic energy matrix, $\mathbf{b}(\mathbf{q}, \dot{\mathbf{q}})$ is the vector of centrifugal and Coriolis joint forces, $\mathbf{g}(\mathbf{q})$ is the vector of gravity, and $\mathbf{\Gamma}$ is the vector of generalized joint forces.

For a nonredundant robot, the effector dynamic behavior is described by the operational space equations of motion (Khatib 1987)

$$\Lambda(\mathbf{x})\ddot{\mathbf{x}} + \mu(\mathbf{x}, \dot{\mathbf{x}}) + \mathbf{p}(\mathbf{x}) = \mathbf{F}, \quad (2)$$

where \mathbf{x} , is the vector of the m operational coordinates describing the position and orientation of the effector, $\Lambda(\mathbf{x})$ 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.

2.2. Effector Interaction with the Environment

The operational space model provides the basis for a unified approach to task-level motion and force control. The operational forces, \mathbf{F} , are produced by submitting the manipulator to the corresponding joint forces, $\mathbf{\Gamma}$, using a simple force transformation

$$\mathbf{\Gamma} = J^T(\mathbf{q})\mathbf{F} \quad (3)$$

where $J(\mathbf{q})$ is the Jacobian matrix.

The use of the forces generated at the end effector to control motions leads to a natural integration of motion and force control. By the nature of coordinates associated with spatial rotations, operational forces acting along rotation coordinates are not homogeneous to moments, and vary with the representation used (e.g., Euler angles, direction cosines, Euler parameters, quaternions). The homogeneity issue is addressed by establishing the end-effector dynamic model in terms of linear and angular velocities and accelerations (Khatib 1987).

With respect to linear and angular motions, the end-effector equations of motion can be written as

$$\Lambda(\mathbf{x})\dot{\boldsymbol{\vartheta}} + \mu(\mathbf{x}, \boldsymbol{\vartheta}) + \mathbf{p}(\mathbf{x}) + \mathbf{F}_{\text{contact}} = \mathbf{F}. \quad (4)$$

The vector $\mathbf{F}_{\text{contact}}$ represents the contact forces acting at the end effector. $\boldsymbol{\vartheta}$ is the vector of end-effector linear and angular velocities and \mathbf{F} is the vector of end-effector forces and moments (Khatib 1987).

Compliant motion and contact operations involve motion control in some directions and force control in other directions. Such tasks are described by the *generalized selection matrix* Ω and its complement $\bar{\Omega}$ associated with motion control and force control, respectively (Khatib 1987).

The end-effector dynamically decoupled motion and force control can be achieved by the control structure

$$\mathbf{F} = \mathbf{F}_{\text{motion}} + \mathbf{F}_{\text{active-force}}, \quad (5)$$

where

$$\mathbf{F}_{\text{motion}} = \hat{\Lambda}(\mathbf{x})\Omega\mathbf{F}_{\text{motion}}^* + \hat{\mu}(\mathbf{x}, \boldsymbol{\vartheta}) + \hat{\mathbf{p}}(\mathbf{x}), \quad (6)$$

$$\mathbf{F}_{\text{active-force}} = \hat{\Lambda}(\mathbf{x})\bar{\Omega}\mathbf{F}_{\text{active-force}}^* + \mathbf{F}_{\text{sensor}}, \quad (7)$$

and $\hat{\cdot}$ represents estimates of the model parameters.

The vectors $\mathbf{F}_{\text{motion}}^*$ and $\mathbf{F}_{\text{active-force}}^*$ represent the inputs to the decoupled system. With perfect estimates of the dynamic parameters and perfect sensing of contact forces (i.e., $\mathbf{F}_{\text{sensor}} = \mathbf{F}_{\text{contact}}$), the closed loop system is described by the following two decoupled sub-systems:

$$\Omega\dot{\boldsymbol{\vartheta}} = \Omega\mathbf{F}_{\text{motion}}^*, \quad (8)$$

$$\bar{\Omega}\dot{\boldsymbol{\vartheta}} = \bar{\Omega}\mathbf{F}_{\text{active-force}}^*. \quad (9)$$

The above control structure provides a *basic primitive* for effector-level motion and force control. This primitive, the control vector \mathbf{F} of eq. (5), is in fact dependent on the location of the operational point, \mathbf{x} , the selections made for the compliant direction, Ω , and the desired motion and forces. The control vector \mathbf{F} can be viewed as a function of these task specification parameters. By selecting these parameters appropriately, one can instantiate this basic control model in many different ways to adapt to specific tasks.

2.3. Effector Dynamics for Vehicle/Arm System

An important characteristic of mobile manipulator systems is the macro/mini structure they possess. Our study has shown

(Khatib 1995) that, in any direction, the inertial properties of a macro/mini-manipulator system are smaller than or equal to the inertial properties associated with the mini structure in that direction. A more general statement of this *reduced effective inertial* property is that the inertial properties of a redundant robot are bounded above by the inertial properties of the structure formed by the smallest distal set of degrees of freedom that span the operational space.

The *reduced effective inertial* property states that the dynamic performance of a vehicle/arm system can be made comparable to and, in some cases, better than that of the manipulator arm alone. A dynamic coordination strategy that allows full utilization of the mini structure's high bandwidth is essential for achieving effective task performance, particularly in compliant motion operations.

The dynamic behavior at the end-effector of a mobile manipulator is obtained by the projection of its joint-space dynamics (1) into operational space

$$\bar{J}^T(\mathbf{q}) [A(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{b}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{g}(\mathbf{q}) = \mathbf{\Gamma}] \implies \Lambda(\mathbf{q})\ddot{\mathbf{x}} + \mu(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{p}(\mathbf{q}) = \mathbf{F}; \quad (10)$$

where

$$\bar{J}(\mathbf{q}) = A^{-1}(\mathbf{q})J^T(\mathbf{q})\Lambda(\mathbf{q}); \quad (11)$$

$\bar{J}(\mathbf{q})$ is the *dynamically consistent generalized inverse*, (Khatib 1995), which minimizes the robot kinetic energy, and

$$\Lambda(\mathbf{q}) = [J(\mathbf{q})A^{-1}(\mathbf{q})J^T(\mathbf{q})]^{-1}. \quad (12)$$

In the case of non-redundant manipulators, the matrix $\bar{J}^T(\mathbf{q})$ reduces to $J^{-T}(\mathbf{q})$.

The increase in the responsiveness of the robotic system is achieved by a control structure identical to the one used in the nonredundant case (5). For redundant robots, this control structure produces joint motions that minimize the robot's instantaneous kinetic energy. As a result, a task at the effector will be carried out by the combined action of the set of joints that possess the smallest effective inertial properties. This gives a prominent role to the arm of a mobile manipulator for performing the effector task. However, typical operations of a mobile manipulator extend much beyond the limited workspace of the arm, giving the mobile base an important role in providing coverage of wide areas of the workspace.

2.4. Posture Control Behavior

The posture, the robot's self configuration, is key to extending the workspace of a mobile manipulator. An important consideration in the development of posture control behaviors is the interaction between the posture and the effector. It is critical for the effector to maintain its responsiveness and to be dynamically decoupled from the posture behavior. The posture

can then be treated separately from the effector task, allowing intuitive task specifications and effective robot control. In our approach, the overall control structure for the integration of mobility and manipulation is based on the following decomposition of joint torques

$$\mathbf{\Gamma} = J^T(\mathbf{q})\mathbf{F} + N^T(\mathbf{q})\mathbf{\Gamma}_{\text{posture}}, \quad (13)$$

with

$$N(\mathbf{q}) = [I - \bar{J}(\mathbf{q})J(\mathbf{q})]. \quad (14)$$

This relationship provides a decomposition of joint forces into two control vectors: joint forces corresponding to forces acting at the effector, $J^T\mathbf{F}$, and joint forces that only affect the robot posture, $N^T\mathbf{\Gamma}_{\text{posture}}$. To control the robot for a desired posture, the vector $\mathbf{\Gamma}_{\text{posture}}$ will be selected as the gradient of a potential function constructed to meet the desired posture specifications. The interference of this gradient with the end-effector dynamics is avoided by projecting it into the dynamically consistent null space of $J^T(\mathbf{q})$, i.e. $N^T(\mathbf{q})\mathbf{\Gamma}_{\text{posture}}$.

For instance, the robot posture can be controlled to maintain the joint position of the arm at their mid-ranges. This corresponds to the potential function

$$V_{\text{mid-range}}(\mathbf{q}) = k \sum_{i=n_M+1}^n (q_i - q_{\text{mid}(i)})^2; \quad (15)$$

where k is a constant gain, $q_{\text{mid}(i)}$ is the mid-range position of joint i , and n_M is the number of degrees of freedom associated with the mobile base. The gradient of this function

$$\mathbf{\Gamma}_{\text{posture}} = -\nabla V_{\text{mid-range}}, \quad (16)$$

provides the required attraction to the mid-range joint positions of the manipulator. Other posture behaviors can be similarly designed (Russakow and Khatib 1992). In addition, collision avoidance can be also integrated in the posture control as discussed in section (4). With this posture behavior, the explicit specification of the associated motions is avoided, since desired behaviors are simply encoded into specialized potential functions for various types of operations. This is illustrated in the simulation results for a 24-degree-of-freedom humanoid system shown in Fig. 2, whose task was generated from simple manipulation and posture behaviors.

3. Cooperative Manipulation

The development of effective cooperation strategies for multiple robot platforms is an important issue for both the operations in human environments and the interaction with humans. Human guided motions may involve tightly constrained cooperation performed through compliant motion actions or less restricted tasks executed through simpler free-space motion commands. Several cooperative robots, for instance, may




Fig. 2. Manipulation and Posture Behaviors: a sequence of three snapshots from the dynamic simulation of a 24-degree-of-freedom humanoid system, whose task is generated from simple manipulation and posture behaviors.

support a load while being guided by the human to an attachment, or visually following the guide to a destination. In this section, we focus on constrained cooperation between multiple robots and describe our approach for a decentralized strategy for robot cooperation.

Our approach is based on the integration of two basic concepts: The *augmented object* (Khatib 1988) and the *virtual linkage* (Williams and Khatib 1995). The *virtual linkage* characterizes internal forces, while the *augmented object* describes the system's closed-chain dynamics. These models have been successfully used in cooperative manipulation for various compliant motion tasks performed by two and three fixed-base PUMA 560 manipulators (Williams and Khatib 1995). First we will present these two models and the corresponding cooperation control strategy. The extension to mobile manipulators, presented in a subsequent section, is based on a decentralized cooperation strategy that is consistent with the *augmented object* and *virtual linkage* models.

3.1. Augmented Object and Virtual Linkage

The *augmented object* model provides a description of the dynamics at the operational point for a multi-arm robot system. The simplicity of the equations associated with this model is the result of an additive property that allows us to obtain the system equations of motion from the the dynamics of the individual mobile manipulators. The *augmented object* model is

$$\Lambda_{\oplus}(\mathbf{x})\ddot{\mathbf{x}} + \mu_{\oplus}(\mathbf{x}, \dot{\mathbf{x}}) + \mathbf{p}_{\oplus}(\mathbf{x}) = \mathbf{F}_{\oplus} \quad (17)$$

with

$$\Lambda_{\oplus}(\mathbf{x}) = \Lambda_{\mathcal{L}}(\mathbf{x}) + \sum \Lambda_i(\mathbf{x}) \quad (18)$$

$$\mu_{\oplus}(\mathbf{x}) = \mu_{\mathcal{L}}(\mathbf{x}) + \sum \mu_i(\mathbf{x}) \quad (19)$$

$$\mathbf{p}_{\oplus}(\mathbf{x}) = \mathbf{p}_{\mathcal{L}}(\mathbf{x}) + \sum \mathbf{p}_i(\mathbf{x}) \quad (20)$$

where $\Lambda_{\mathcal{L}}(\mathbf{x})$ and $\Lambda_i(\mathbf{x})$ are the kinetic energy matrices associated with the object and the i th effector, $\mu_{\mathcal{L}}(\mathbf{x})$ and $\mu_i(\mathbf{x})$

are the centrifugal and Coriolis vectors for the object and the i th effector, and $\mathbf{p}_{\mathcal{L}}(\mathbf{x})$ and $\mathbf{p}_i(\mathbf{x})$ are the gravity vectors for the object and the i th effector, respectively. \mathbf{F}_{\oplus} also has the same additive property shown above for $\Lambda_{\oplus}(\mathbf{x})$, $\mu_{\oplus}(\mathbf{x}, \dot{\mathbf{x}})$ and $\mathbf{p}_{\oplus}(\mathbf{x})$.

Object manipulation requires accurate control of internal forces. We have proposed the *virtual linkage* (Williams and Khatib 1995), as a model of object internal forces associated with multi-grasp manipulation. In this model, grasp points are connected by a closed, nonintersecting set of virtual links (Fig. 3.)

For an N -grasp manipulation task, the *virtual linkage* model is a $6(N - 1)$ degree of freedom mechanism that has $3(N - 2)$ linearly actuated members and N spherically actuated joints. By applying forces and moments at the grasp points we can independently specify internal forces in the $3(N - 2)$ linear members, along with $3N$ internal moments at the spherical joints. Internal forces in the object are then characterized by these forces and torques in a physically meaningful way. The relationship between applied forces, their resultant, and internal forces is

$$\begin{bmatrix} \mathbf{F}_{res} \\ \mathbf{F}_{int} \end{bmatrix} = \mathbf{G} \begin{bmatrix} \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_N \end{bmatrix}, \quad (21)$$

where \mathbf{F}_{res} is the resultant forces at the operational point, \mathbf{F}_{int} is the internal forces, and \mathbf{f}_i is the forces applied at the grasp point i . \mathbf{G} is the grasp description matrix. It relates forces applied at each grasp to the resultant and internal forces in the object. The matrix \mathbf{G} can be written as




Fig. 3. The Virtual Linkage: for a three-grasp manipulation task, a twelve-degree-of-freedom mechanism (three spherical joints and three prismatic joints) is used to describe the internal forces.

$$\mathbf{G} = [\mathbf{G}_1 \mathbf{G}_2 \dots \mathbf{G}_N], \quad (22)$$

with

$$\mathbf{G}_i = \begin{bmatrix} \mathbf{G}_{res,i} \\ \mathbf{G}_{int,i} \end{bmatrix}, \quad (23)$$

where \mathbf{G}_i represents the contribution of the i th grasp to the resultant and internal forces felt by the object; and where $\mathbf{G}_{res,i}$ and $\mathbf{G}_{int,i}$ are respectively the contribution of \mathbf{G}_i to the resultant forces on the object and to the internal ones.

The inverse of the grasp description matrix, \mathbf{G}^{-1} , provides the forces required at the grasp points to produce the resultant and internal forces acting at the object:

$$\begin{bmatrix} \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_N \end{bmatrix} = \mathbf{G}^{-1} \begin{bmatrix} \mathbf{F}_{res} \\ \mathbf{F}_{int} \end{bmatrix}, \quad (24)$$

with

$$\mathbf{G}^{-1} = \begin{bmatrix} \bar{\mathbf{G}}_1 \\ \vdots \\ \bar{\mathbf{G}}_N \end{bmatrix} \text{ and } \bar{\mathbf{G}}_i = [\bar{\mathbf{G}}_{res,i} \quad \bar{\mathbf{G}}_{int,i}], \quad (25)$$

where $\bar{\mathbf{G}}_{res,i}$ represents the part of $\bar{\mathbf{G}}_i$ corresponding to resultant forces at the object. $\bar{\mathbf{G}}_{int,i}$ represents the part of $\bar{\mathbf{G}}_i$ that contributes toward the object's internal forces.

The grasp description matrix contains a model of the internal force representation as well as the relationship between applied grasp forces and object resultant forces, thus is central to the control scheme employed by the *virtual linkage* model.

Compared to other methods used to characterize internal forces, the virtual linkage has the advantage of providing a physical representation of internal forces and moments. This allows control of non-zero internal forces to have a physically meaningful effect on the manipulated object.

3.2. Centralized Control Structure

The *virtual linkage* and *augmented object* models have been successfully used in the cooperative control of two and three fixed PUMA arms. For these fixed-base (nonmobile) robots, the control structure was implemented using a *centralized* control scheme. In a centralized control setup, each arm sends its sensory data to a central controller which then commands the motion of each arm based on information from all the arms in the system. However, this type of control is not suited to the more autonomous nature inherent in mobile manipulation systems, where a *decentralized* control scheme is more appropriate. Before presenting the decentralized implementation, we begin with a brief summary of the centralized control structure.

The overall structure of the centralized implementation is shown in Fig. 4. The force sensed at the grasp point of each robot, $\mathbf{f}_{s,i}$, is transformed, via \mathbf{G} , to sensed resultant forces, $\mathbf{F}_{res,s}$, and sensed internal forces, $\mathbf{F}_{int,s}$, at the operational point, using eq. (21)

$$\begin{bmatrix} \mathbf{F}_{res,s} \\ \mathbf{F}_{int,s} \end{bmatrix} = \mathbf{G} \begin{bmatrix} \mathbf{f}_{s,1} \\ \vdots \\ \mathbf{f}_{s,N} \end{bmatrix}.$$

The centralized control strategy consists of (1) a unified motion and contact force control structure for the *augmented object*, \mathbf{F}_{res} ; and (2) \mathbf{F}_{int} , corresponding to the control of internal forces acting on the *virtual linkage*. The first part of the controller, associated with motion and contact force control, is

$$\mathbf{F}_{res} = \mathbf{F}_{motion} + \mathbf{F}_{contact}, \quad (26)$$

where

$$\mathbf{F}_{motion} = \hat{\Lambda}_{\oplus} \Omega \mathbf{F}_{motion}^* + \hat{\mu}_{\oplus} + \hat{\mathbf{p}}_{\oplus} \quad (27)$$

and

$$\mathbf{F}_{contact} = \hat{\Lambda}_{\oplus} \bar{\Omega} \mathbf{F}_{contact}^* + \mathbf{F}_{contact,s}. \quad (28)$$

$\hat{\Lambda}_{\oplus}$, $\hat{\mu}_{\oplus}$ and $\hat{\mathbf{p}}_{\oplus}$ represent the estimates of Λ_{\oplus} , μ_{\oplus} , and \mathbf{p}_{\oplus} . The vector \mathbf{F}_{motion}^* and $\mathbf{F}_{contact}^*$ represent the inputs to the unit mass, decoupled system. Ω is the *generalized selection matrix* associated with motion control and its complement, $\bar{\Omega}$, is associated with contact force control. The control structure for internal forces is

$$\mathbf{F}_{int} = \hat{\Lambda}_{\oplus} \mathbf{F}_{int}^* + \mathbf{F}_{int,s}, \quad (29)$$

where the vector \mathbf{F}_{int}^* represents the input to the decoupled, unit mass system. A suitable control law can be selected to obtain \mathbf{F}_{motion}^* , $\mathbf{F}_{contact}^*$, and \mathbf{F}_{int}^* . The control forces of the individual mobile manipulator at its grasp point, \mathbf{f}_i , are given by eq. (24),

3.3. Decentralized Control Structure

For systems of a mobile nature, a *decentralized* control structure is needed to address the difficulty of achieving high-rate communication between platforms. In the decentralized control structure, the object level specifications of the task are




Fig. 4. Centralized Control Structure

transformed into individual tasks for each of the cooperative robots. Local feedback control loops are then developed at each grasp point. The task transformation and the design of the local controllers are accomplished in consistency with the *augmented object* and *virtual linkage* models (Khatib 1988; Williams and Khatib 1995). The overall structure is shown in Fig. 5. The local control structure at the *i*th grasp point is

$$\mathbf{f}_i = \mathbf{f}_{motion,i} + \mathbf{f}_{force,i}. \quad (30)$$

The control vectors, $\mathbf{f}_{motion,i}$, are designed so that the combined motion of the various grasp points results in the desired motion at the object operational point. In addition, the vectors $\mathbf{f}_{force,i}$ create forces at the grasp points, whose combined action produces the desired internal forces on the object. The motion control at the *i*th grasp point is

$$\mathbf{f}_{motion,i} = \widehat{\Lambda}_{\mathcal{O},i} \Omega \mathbf{f}_{motion,i}^* + \widehat{\mu}_{\mathcal{O},i} + \widehat{\mathbf{p}}_{\mathcal{O},i}, \quad (31)$$

with
$$\widehat{\Lambda}_{\mathcal{O},i} = \widehat{\Lambda}_{g,i} + \overline{\mathbf{G}}_{res,i} \widehat{\Lambda}_{\mathcal{L}} \overline{\mathbf{G}}_{res,i}^T, \quad (32)$$

where $\widehat{\Lambda}_{g,i}$ is the kinetic energy matrix associated with the *i*th effector at the grasp point. The second term of eq. (32) represents the part of $\widehat{\Lambda}_{\mathcal{L}}$ "assigned" to the *i*th robot, described at its grasp point. The vector, $\widehat{\mu}_{\mathcal{O},i}$, of centrifugal and Coriolis forces associated with the *i*th effector is

$$\widehat{\mu}_{\mathcal{O},i} = \widehat{\mu}_{g,i} + \overline{\mathbf{G}}_{res,i} \widehat{\mu}_{\mathcal{L}}; \quad (33)$$

where $\widehat{\mu}_{g,i}$ is the centrifugal and Coriolis vector of the *i*th robot alone at the grasp point and $\overline{\mathbf{G}}_{res,i} \widehat{\mu}_{\mathcal{L}}$ represents the added part due to the load. Similarly, the gravity vector is

$$\widehat{\mathbf{p}}_{\mathcal{O},i} = \widehat{\mathbf{p}}_{g,i} + \overline{\mathbf{G}}_{res,i} \widehat{\mathbf{p}}_{\mathcal{L}}, \quad (34)$$

where $\widehat{\mathbf{p}}_{g,i}$ is the gravity vector associated with the *i*th end effector at the grasp point. The total sensed forces at the *i*th grasp point, $\mathbf{f}_{s,i}$, combine the contact and internal forces felt at the *i*th grasp point, together with the acceleration force acting at the object. The sensed forces associated with the contact and internal forces alone, $\mathbf{f}_{\bar{s},i}$, are therefore obtained by subtracting the acceleration effect of the object from the total sensed forces

$$\mathbf{f}_{\bar{s},i} = \mathbf{f}_{s,i} - \overline{\mathbf{G}}_{res,i} (\widehat{\Lambda}_{\mathcal{L}} \ddot{\mathbf{x}}_d + \widehat{\mu}_{\mathcal{L}} + \widehat{\mathbf{p}}_{\mathcal{L}}). \quad (35)$$




Fig. 5. Decentralized Control Structure

Here, the object desired acceleration has been used instead of the actual acceleration, which would be difficult to evaluate. The force control part of eq. (30) is

$$\mathbf{f}_{force,i} = \widehat{\Lambda}_i \mathbf{f}_{force,i}^* + \mathbf{f}_{des,i}. \quad (36)$$

The vector $\mathbf{f}_{des,i}$ represents the desired force assigned to the *i*th mobile manipulator. This vector is

$$\mathbf{f}_{des,i} = \overline{\mathbf{G}}_{int,i} \mathbf{F}_{int,des} \quad (37)$$

where $\mathbf{F}_{int,des}$ is the desired object internal force. $\mathbf{f}_{force,i}^*$ represents the input to the decoupled, unit mass system associated with the internal forces. It can be achieved by selecting

$$\mathbf{f}_{force,i}^* = -\mathbf{K}_f (\mathbf{f}_{\bar{s},i} - \mathbf{f}_{d,i}) - \mathbf{K}_{vf} \dot{\mathbf{f}}_{\bar{s},i}. \quad (38)$$

The above control structure is consistent with the augmented object and virtual linkage model under the assumptions of no slippage at the grasp points. Significant flexibilities and gripper slip in the real system will result in errors in the grasp kinematic computation and inconsistencies with the *virtual linkage* model. To compensate for these effects, some level of communication between the different platforms is required for updating the robot state and modifying the task specifications. However, the rate at which this communication is required is much slower than the local servo control rate and can be achieved over a radio Ethernet link.

4. Collision-Free Path Modification Behaviors

To perform motion tasks, a robot must combine the abilities of planning motions and executing them. Since a planned motion is based on a priori knowledge of the environment, it is difficult to carry out such a motion when uncertainties and unexpected obstacles are to be considered. Reactive behaviors sought to deal with dynamic environments are, by their local nature, incapable of achieving global goals. Our investigation of a framework to connect real-time collision avoidance capabilities with a global planning system has resulted in a new approach based on the *elastic band* concept (Quinlan and Khatib 1993).

The key to the efficiency of the elastic band is the representation of free space around the path as a series of hyperspheres, called *bubbles*. A bubble represents a region of configuration space that is free of collision. Covering the path with those bubbles, a channel of free space is formed through which the robot's trajectory can be executed. This is illustrated in Fig. 6. The effectiveness of this approach has been demonstrated experimentally on different robotic systems (Quinlan and Khatib 1993).

This approach becomes computationally demanding, however, as the dimension of the configuration space associated with the robot increases. The specification of tasks for robots is most naturally done in workspace. Elastic bands, however, represent a path in the configuration space.




Fig. 6. Bubble Implementation of Elastic Band: as an obstacle moves, the bubbles also move to minimize the force on the elastic band. If needed, bubbles are inserted and deleted to maintain a collision free path.

The *elastic strip* (Brock and Khatib 1997) operates entirely in the workspace. The characterization of free space becomes more accurate in the workspace than in configuration space, resulting in a more efficient description of trajectories. In addition, by avoiding configuration space computation, the framework becomes applicable to robots with many degrees of freedom. The trajectory and the task are both described in workspace. An elastic strip represents the workspace volume that is swept by the entire robot along its trajectory. The basic idea of the elastic strip is to incrementally modify this workspace volume as if it were elastic, expanding and contracting to maintain a short and smooth path. Objects in the environment exert repulsive forces, ensuring a safe distance to obstacles.

To represent the free space associated with the elastic strip, we propose a series of three-dimensional spheres in the workspace around some configurations along the elastic strip. A single configuration is covered with a set of such spheres forming a *protective hull* of that configuration. This is illustrated in Fig. 7. The overlapping protective hulls along the trajectory form an *elastic tunnel*, which represents the local free space along the entire path.

4.1. Local Free-Space Representation

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

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

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

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

Figure 7 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.

An elastic strip $\mathcal{S}_{\mathcal{T}}^{\mathcal{R}} = (q_1, q_2, q_3, \dots, q_n)$ 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 $\mathcal{P}_i^{\mathcal{R}}$, 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

$$V_{\mathcal{T}}^{\mathcal{R}} \subseteq V_{\mathcal{S}}^{\mathcal{R}} = \bigcup_{1 \leq i \leq n} \mathcal{P}_i^{\mathcal{R}}. \quad (39)$$

This is illustrated in Fig. 8, 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.

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 (39) 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.




Fig. 7. Protective Hulls: an example illustrating the protective hulls for 3 configurations of the Stanford Mobile Platform amidst spherical obstacles.




Fig. 8. Elastic Tunnel: the protective hulls covering a trajectory for the Stanford platforms form an elastic tunnel of free space.

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} ,

$$V_{\mathcal{T}_{i,i+1}}^b \subseteq (\mathcal{P}_i^b \cup \mathcal{P}_{i+1}^b). \quad (40)$$

If this condition is verified for all rigid bodies of the robot, a trajectory $\mathcal{T}_{i,i+1}$ exists for the robot.

4.2. Forces Acting on the Strip

An elastic strip can be seen as a grid of links and springs. 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.

4.3. Elastic Strip Modification

Let $\delta = (q_1, q_2, q_3, \dots, q_n)$ be an elastic strip. When δ is subjected to the forces described in section 4.2, 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_p associated with the point p at which the force F_p is acting. The joint torques Γ caused by F_p are given by

$$\Gamma = J_p^T F_p. \quad (41)$$

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 $\mathcal{S}' = (q_1, \dots, q'_i, \dots, q_n)$. \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.

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}' . 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 \mathcal{P}_i is redundant and can be removed from \mathcal{S} .

4.4. Motion Behaviors

Given a planned motion, the elastic strip allows a robot to dynamically modify its motion to accommodate changes in the environment. For a mobile manipulator 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 approach for executing partially described tasks. 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.

The framework for combining motion behavior and task execution relies on the effector/posture control structure discussed above. 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 while the redundant degrees of freedom are being used for obstacle avoidance.

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. This is achieved by selecting

$$\Gamma_{\text{posture}} = -\nabla(V_{\text{desired-posture}} + V_{\text{obstacle-avoidance}}) \quad (42)$$

and projecting these torques in the dynamically consistent null space to guarantee that the posture control torques will not alter the end-effector's dynamic behavior.

An example of the the elastic strip implementation is shown in Fig. 9. 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.

5. Stanford Mobile Platforms

In collaboration with Oak Ridge National Laboratories and Nomadic Technologies, we designed and built two holonomic mobile manipulator platforms. Each platform is equipped with a PUMA 560 arm, various sensors, two computer systems, a multi-axis controller, and sufficient battery power to allow for autonomous operation. The base consists of three "lateral" orthogonal universal-wheel assemblies (Pin and Kil-lough 1994) which allow the base to translate and rotate holonomically in relatively flat office-like environments.

The Stanford Robotic Platforms have been used in the implementation and verification of the different strategies discussed above. We have demonstrated real-time collision avoidance with coordinated vehicle/arm motion, and cooperative tasks involving operator-directed compliant motion (M. Khatib et al. 1997).

The elastic strip framework was also implemented and tested on the Stanford platforms. For example, one robot was commanded to perform a straight line motion, while keeping the arm's posture. During the execution of this plan an unforeseen obstacle, the second platform, forces the first robot to deviate from its original plan. Two different perspectives of the simulated modification of the trajectory are shown in Fig. 10. A sequence of snapshots from the execution on the real robot can be seen in Fig. 11.

The Stanford platforms have been also used in a variety of mobile manipulation tasks including ironing, opening a door, and vacuuming, as illustrated in Fig. 12. The dynamic strategy for integrated mobility and manipulation discussed above has allowed full use of the bandwidth of the PUMA manipulator. Object motion and force control performance




Fig. 9. Elastic Strip: the initial plan for the Stanford Platform is incrementally modified by a moving obstacle.




Fig. 10. Interaction between the Two Platforms: the elastic strip of the first platform is modified incrementally to maintain a valid path while avoiding the second moving platform.



Fig. 11. Experimental Execution of a Plan: the path of the first platform is modified in real-time to avoid the second moving platform.




Fig. 12. Experiments with the Stanford Mobile Platforms: vacuuming, opening a door, and ironing are examples of tasks demonstrated with the Stanford Mobile Platforms.

with the Stanford mobile platforms are comparable with the results obtained with fixed base PUMA manipulators.

6. Conclusion

Advances toward the challenge of robotics in human environments depend on the development of the basic capabilities needed for both autonomous operations and human/robot interaction. In this article, we have presented methodologies for the integration of mobility and manipulation, the cooperation between multiple robots, the interaction between human and robots, and the real-time modification of collision-free path to accommodate changes in the environment.

For vehicle/arm coordination and control, we presented a framework that provides the user with two basic task-oriented control primitives, end-effector task control and robot posture control. The major characteristic of this control structure is the dynamic consistency it provides in implementing these two primitives: the robot posture behavior has no impact on the end-effector dynamic behavior. While ensuring dynamic decoupling and improved performance, this control structure provides the user with a higher level of abstraction in dealing with task specifications and control.

For cooperative operations between multiple platforms we have presented a decentralized control structure. This structure relies on the integration of the *augmented object* which describes the system's closed-chain dynamics, and the *virtual linkage* which characterizes internal forces. This decentralized cooperation approach provides the basis for an effective strategy for human/robot interaction.

The notion of an elastic strip encapsulates what must be known about the environment for both executing global motions and adjusting them to dynamic changes and unforeseen circumstances quickly and safely. The generality of this notion makes it the appropriate abstraction at all levels in the

control of a team of cooperating robots. An elastic strip represents the workspace volume swept by a robot along a preplanned 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.

Vehicle/arm coordination, cooperative operations, robot/human interaction, and the elastic strip approach have been demonstrated on the mobile manipulator platforms developed at Stanford University.

Acknowledgments

The financial support of Boeing, General Motors, and NSF (grant IRI-9320017) is gratefully acknowledged. Many thanks to Andreas Baader, Alan Bowling, Bob Holmberg, Francois Pin, James Slater, John Slater, Stef Sonck for their valuable contributions to the design and construction of the Stanford robotic platforms.

References

- Adams, J. A., R. Bajcsy, J. Kosecka, V. Kumar, R. Mandelbaum, M. Mintz, R. Paul, C. Wang, Y. Yamamoto, and X. Yun (1995). Cooperative material handling by human and robotic agents: Module development and system synthesis. In *Proc. of the Int. Conf. on Intelligent Robots and Systems*, pp. 200–205.
- Arkin, R. C. (1987). Motor schema-based mobile robot navigation. *Int. J. of Robotics Research* 8(4): 92–112.
- Brock, O. and O. Khatib (1997). Elastic strips: Real-time path modification for mobile manipulation. In *Proc. of the Int. Symp. of Robotics Research*, pp. 117–122. Preprints.
- Cameron, J. M., D. C. MacKenzie, K. R. Ward, R. C. Arkin, and W. J. Book (1993). Reactive control for mobile

- manipulation. In *Proc. of the Int. Conf. on Robotics and Automation*, Volume 3, pp. 228–235.
- Carriker, W. F., P. K. Khosla, and B. H. Krogh (1989). An approach for coordinating mobility and manipulation. In *Proc. of the Int. Conf. on Systems Engineering*, pp. 59–63.
- Engelberger, J. F. (1991). *Robotics in Service*. Biddles, Guilford.
- Hayati, S. (1987). Hybrid position/force control of multi-arm cooperating robots. In *Proc. of the Int. Conf. on Robotics and Automation*, pp. 1375–1380.
- Jung, D., G. Cheng, and A. Zelinsky (1997). Experiments in realizing cooperation between autonomous mobile robots. In *Proc. of the Int. Symp. on Experimental Robotics*, pp. 513–524.
- Khatib, M., H. Jaouni, R. Chatila, and J.-P. Laumond (1997). How to implement dynamic paths. In *Proc. of the Int. Symp. on Experimental Robotics*, pp. 225–36. Preprints.
- Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *Int. J. of Robotics Research* 5(1): 90–8.
- Khatib, O. (1987). A unified approach to motion and force control of robot manipulators: The operational space formulation. *Int. J. of Robotics Research* 3(1), 43–53.
- Khatib, O. (1988). Object manipulation in a multi-effector robot system. In R. Bolles and B. Roth (Eds.), *Robotics Research 4*, pp. 137–144. MIT Press.
- Khatib, O. (1995). Inertial properties in robotics manipulation: An object-level framework. *Int. J. of Robotics Research* 14(1), 19–36.
- Khatib, O., K. Yokoi, K.-S. Chang, and A. Casal (1997). The Stanford robotic platforms. In *Video Proc. of the Int. Conf. on Robotics and Automation*.
- Khatib, O., K. Yokoi, K.-S. Chang, D. Ruspini, R. Holmberg, and A. Casal (1996). Coordination and decentralized cooperation of multiple mobile manipulators. *J. of Robotic Systems* 13(11), 755–64.
- Koditschek, D. E. (1987). Exact robot navigation by means of potential functions: Some topological considerations. In *Proc. of the Int. Conf. on Robotics and Automation*, pp. 1–6.
- Krogh, B. H. (1994). A generalized potential field approach to obstacle control. In *Proc. Robotics Research, The Next Five Years and Beyond*, Bethlehem, PA.
- Latombe, J.-C. (1991). *Robot Motion Planning*. Boston: Kluwer Academic Publishers.
- Papadopoulos, E. and S. Dubowsky (1991). Coordinated manipulator/spacecraft motion control for space robotic systems. In *Proc. of the Int. Conf. on Robotics and Automation*, pp. 1696–1701.
- Pin, F. G. and S. M. Killough (1994). A new family of omnidirectional and holonomic wheeled platforms for mobile robots. *IEEE Trans. on Robotics and Automation* 10(4): 480–489.
- Quinlan, S. and O. Khatib (1993). Elastic bands: Connecting path planning and control. In *Proc. of the Int. Conf. on Robotics and Automation*, Volume 2, pp. 802–7.
- Russakow, J. and O. Khatib (1992). A new control structure for free-flying space robots. In *Proc. of the Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*, pp. 395–403.
- Schmidt, G., U. D. Hanebeck, and C. Fischer (1997). A mobile service robot for the hospital and home environment. In *Proc. of the IARP Second Int. Workshop on Service and Personal Robots: Technologies and Applications*.
- Schraft, R. D. and M. Hägele (1997). Servicerobots: Examples of the state-of-art in products, prototypes, and product vision. In *Proc. of the IARP Second Int. Workshop on Service and Personal Robots: Technologies and Applications*.
- Seraji, H. (1993). An on-line approach to coordinated mobility and manipulation. In *Proc. of the Int. Conf. on Robotics and Automation*, Volume 1, pp. 28–35.
- Tarn, T.-J., A. K. Bejczy, and X. Yun (1987). Design of dynamic control of two cooperating robot arms: Closed chain formulation. In *Proc. of the Int. Conf. on Robotics and Automation*, pp. 7–13.
- Uchiyama, M. and P. Dauchez (1988). A symmetric hybrid position/force control scheme for the coordination of two robots. In *Proc. of the Int. Conf. on Robotics and Automation*, pp. 350–356.
- Ullman, M. and R. Cannon (1989). Experiments in global navigation and control of a free-flying space robot. In *Proc. of the ASME Winter Annual Meeting*, Volume 17, pp. 37–43.
- Umetani, Y. and K. Yoshida (1989). Experimental study on two-dimensional free-flying robot satellite model. In *Proc. of the NASA Conf. on Space Telerobotics*.
- Williams, D. and O. Khatib (1993). The virtual linkage: A model for internal forces in multi-grasp manipulation. In *Proc. of the Int. Conf. on Robotics and Automation*, Volume 1, pp. 1025–30.
- Williams, D. and O. Khatib (1995). Multi-grasp manipulation. In *Video Proc. of the Int. Conf. on Robotics and Automation*.
- Yamamoto, Y. and X. Yun (1994). Coordinating locomotion and manipulation of a mobile manipulator. *IEEE Trans. on Automatic Control* 39(6): 1326–32.
- Yamamoto, Y. and X. Yun (1995). Coordinated obstacle avoidance of a mobile manipulator. In *Proc. of the Int. Conf. on Robotics and Automation*, Volume 3, pp. 2255–60.
- Zheng, Y. F. and J. Luh (1986). Joint torques for control of two coordinated moving robots. In *Proc. of the Int. Conf. on Robotics and Automation*, pp. 1375–1380.