

# A Control Paradigm for General Purpose Manipulation Systems

Roderic A. Grupen  
Thomas C. Henderson

UUCS-87-029  
December 4, 1987

## Abstract

Mechanical end effectors capable of dextrous manipulation are now a reality. Solutions to the high level control issues, however, have so far proved difficult to formulate. We propose a methodology for control which produces the functionality required for a general purpose manipulation system. It is clear that the state of a hand/object system is a complex interaction between the geometry of the object, the character of the contact interaction, and the conditioning of the manipulator. The objective of this work is the creation of a framework within which constraints involving the manipulator, the object, and the hand/object interaction can be exploited to direct a goal oriented manipulation strategy. The set of contacts that are applied to a task can be partitioned into subsets with independent objectives. The individual contacts may then be driven over the interaction surface to improve the state of the grasp while the configuration of the hand addresses the application of required forces. A system of this sort is flexible enough to manage large numbers of contacts and to address manipulation tasks which require the removal and replacement of fingers in the grasp. A simulator has been constructed and results of its application to position synthesis for initial grasps is presented. A discussion of the manipulation testbed under construction at the University of Utah employing the Utah/MIT Dextrous hand is presented.

## 1 Introduction

Many different manipulators have been developed with a variety of configurations and with correspondingly diverse capabilities[11,13,14,19,22,24,25,26]. The premise of our work is that the utility of the human hand is largely its ability to perform as a general purpose manipulation device. Special purpose hardware is always capable of out performing the human hand if the application domain is restricted sufficiently. The work presented here is focused on the development of a control structure which is flexible enough to support general manipulation.

### 1.1 Controlled Interaction with the Environment

Perhaps the first gripper operated by a computer under a feedback control strategy was demonstrated by Ernst [8] in 1961. It was capable of performing simple manipulation tasks using tactile feedback to verify the presence or absence of an object. However, position control alone proved unsatisfactory in the vast majority of assembly operations.

Whitney in 1977 implemented an admittance matrix model to predict the manipulator velocity in response to external forces[27]. Force, position and velocity sensors were used to make small corrections in the trajectories of contacting parts. The system permitted sensed contact forces and geometrical information about the task to drive small relative reorientations of mating parts.

Similar efforts to imbue the manipulator with a compliant character have defined orthogonal subspaces within which either position or force may be controlled. Paul *et al.* in 1976 implemented this strategy by selecting the manipulator joints that are most closely aligned with a command force vector and imposing the appropriate input force with those joints[23]. The joints not involved with the force command were position controlled. Raibert *et al.* in 1981 combined force and position trajectory constraints by specifying a position controlled space and a complementary force controlled space. The components of the error signals in both force and position which map into a joint's workspace contribute to the feedback control of that joint. Mason in 1981 used the notion of natural and artificial constraints[20]. Tasks are modeled by a C-surface (constraint surface), to which forces may be applied along the normal and positions controlled along the tangent.

Salisbury in 1982 suggested that task defined stiffnesses be established[24]. Task constraints are established using the position and force measurements so that system stiffnesses can be adjusted accordingly. The environment is actively sampled along 6 orthogonal axes to measure contact forces resulting from motion. Knowledge of the nominal control stiffness and the contact force are used to evaluate the environmental stiffness.

Geschke in 1983 demonstrated a unique implementation of the position/force controller in the form of his *Robot Servo System* or RSS[9]. A single instruction in RSS initiates a servo process which actively seeks its goal until either cancelled or redefined. Compliant subspaces may be defined which change continually as the task geometry changes.

## 1.2 Haptics for Mechanical Manipulators

In addition to the control issues, much work has been done to support the role of the manipulator as a sensory device. It is clear that manipulation cannot be approached in an open loop manner. Uncertainty in the dynamic state models required for manipulation must be managed and the interaction forces must be measured. Tactile technology will, therefore, play a major role in the development of manipulation strategies. Allen describes a system which integrates vision and taction in order to more efficiently describe an object's surface than either sensory modality can do independently [1,2]. Bolles and Horaud [5,6] and recent work by Hansen and Henderson [12] demonstrate systems which hypothesize and subsequently verify an object's orientation by comparing the topology of local sensed features to an extended CAD model. Such methods permit the system to acquire additional features of the object in order to discriminate among competing hypotheses. Grimson uses sparse range or tactile data to disambiguate sensory interpretations and thus determine the orientation of an object [10]. The promising results of such systems, along with the development of systems which optimally combine information from different sensory modalities [4], suggests that the dynamic world models required for manipulation are not unrealistic. The work described in this paper assumes that such models are available.

## 1.3 Quantifying a Grasp

A useful approach to the problem of stability in the hand/object system is the enumeration of all the mechanical impedances imposed on the object by virtue of its contact with each finger. To completely constrain an object, the union over the contact set must span the object's six degrees of freedom. Kobayashi posed this problem in terms of hybrid position and force controllers[16,17]. In order to maintain a stable grasp while actuators move the object, Kobayashi develops the idea of manipulation and free subspaces:

$S_m$  - space in which all fingers remain in contact with the object and may move without violating the mode of contact of another finger, and

$S_f$  - space or degrees of freedom remaining for the object when all actuators are fixed.

The grasp is entirely constrained and stably grasped if the rank of the  $S_f$  space is zero and may be manipulated in the space defined by  $S_m$ . In three dimensions, the objective is a free space with rank zero and a manipulation space of rank 6.

Salisbury expresses the constraint criteria in terms of screw systems [24]:

**wrenches** : an ensemble of coordinate directions about which generalized forces may be exerted (i.e., forces and moments) so that  $\vec{w}_{net} = \cup \vec{w}_i$ , and

**twists** : an ensemble of coordinate directions about which generalized displacements may occur (i.e., translations and rotations) so that  $\vec{t}_{net} = \cap \vec{t}_i$ .

The net wrench applied to a body may be expressed as:

$$\vec{w} = [\vec{w}_1 \vec{w}_2 \dots \vec{w}_n] \vec{c}, \text{ or}$$

$$\vec{w} = \overline{W} \vec{c}$$

where the  $n$   $\vec{w}_i$ 's represent all the principal wrenches available to the current contact system and  $\vec{c}$  represents a command vector of wrench intensities.

The solution for  $\vec{c}$  can be expressed as the sum of a homogeneous solution,  $\vec{c}_h$ , and a particular solution,  $\vec{c}_p$ . The subspace spanned by  $\vec{c}_h$  is referred to as the *null space* of  $\overline{W}$ .

The Grip Transform,  $\overline{G}$ , is constructed by augmenting the matrix  $\overline{W}$  with basis vectors which span the null space of  $\overline{W}$ . The resulting algebraic expression relates the command wrench intensities to the net forces applied to the object. The rank of the Grip Jacobian is equal to the number of contact wrenches in the system. When the rank of the Grip Jacobian grows, so does the rank of its Null space. The solution for the internal force magnitudes is not unique and requires additional knowledge to produce a command wrench intensity vector.

The value of a *General Purpose Manipulation System* is its ability to accommodate an endless variety of object/task combinations. This objective has motivated researchers to enumerate grasps employed by humans with the hope of discovering the intrinsic qualities of successful hand/object interactions [7]. Such a manipulation *syntax* could be used to define an operational paradigm for each object and task, reducing the manipulation problem to one of indexing. Arbib *et al.* proposed the concept of *virtual fingers* [3] as an effective means of reducing the complexity of the human hand and matching the hand's capabilities with the object's surface in light of the task. Tomavic *et al.* are developing the Belgrade hand and the concept of reflex control[26] in order to reduce the dimensionality of the control problem. A problem with this approach is that to reduce the complexity of the hand mechanism before the fact, also restricts the domain of objects and tasks to which that mechanism may be applied after the fact. Our approach to general purpose manipulation requires the management of the complexity of the manipulator and the subsequent reduction of complexity where it is warranted by experience with classes of objects and tasks.

## 2 General Purpose Manipulation

We propose an effective functional integration of the hand and the object which recognizes the independent nature of the forces that the object's surface may transmit, and the manipulator's ability to generate those forces. The approach described here enforces functional priorities among competing prerogatives: the object prefers to be grasped on surface elements which will transmit wrenches that span the space occupied by the task, the hand prefers to configure the set of fingers for the task defined forces and velocities, while an individual finger requires that the proposed state does not violate its workspace.

### 2.1 Control Hierarchy

A perspective on the ideas that support the control structure presented in this section is presented by Minsky in his discussions of the Society of Mind[21]. The mind is viewed as a society of separate, independent agents, each with its private agenda, contending for the resources available. The structure described allows the Mind to maintain the Body

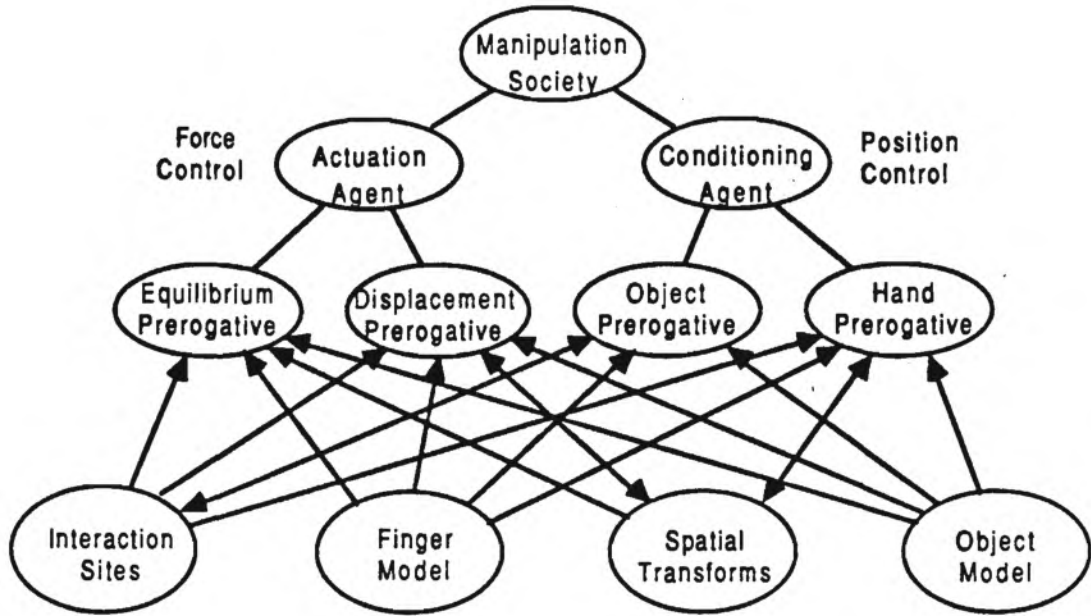


Figure 1: The Manipulation Society

on an interrupt type basis, to be resourceful when confronted with strange circumstances, and to learn about its environment.

The flexibilities of the society structure can be used to implement general purpose manipulation. The *Manipulation Society* is illustrated in Figure 1. The top level of this society serves to focus attention on contending agents. The *Actuation* agent is indulged until the system violates the task defined conditioning envelope, at which time, the *Conditioning* agent is recognized. The next level, the prerogative level, expresses the various independent interests of the manipulation society. The compound objectives of manipulation are distilled into independent tasks which are prioritized and managed by the agents above them. The bottom level of this tree is a dynamic state model for the manipulation system. It represents the geometry as well as the relative position and orientation of the hand and object and the positions through which interactions will take place. The arrows in Figure 1 designate who uses the data at the prerogative level, and define who is responsible for changing the dynamic state of the system at the bottom level. The balance of this section describes the details of the various agents presented in the Conditioning agent half of the Manipulation Society.

## 2.2 Conditioning Agent

The *Conditioning* agent is responsible for developing an interaction strategy for the hand/object system. It is not concerned with the application of forces or velocities, but

restricts its effort to the geometry of the interaction. To do this, it examines two agents: the *Object Prerogative* agent and the *Hand Prerogative* agent. The Object Prerogative agent seeks reachable positions on the surface of the object that improve the state of the system relative to the task. The Hand Prerogative agent drives the hand frame to positions and orientations which best address the task. For general purpose manipulation, it is the policy of the Conditioning agent to comply with the Object's Prerogative whenever possible, and to compromise the hand's conditioning if necessary to accomplish the task. Crosstalk between the two takes place since each is aware of the other's actions through the changing state of the system.

### The Object Model

The object model uses a contact model to determine a family of wrenches that may be transmitted through that surface element. The contact model used currently by the system is a point contact with friction. Figure 2 (a) depicts a set of contact forces that may be transmitted to the object using this friction model. The forces,  $F_1$  through  $F_5$ , are not independent, however, since the tangentially applied forces are dependent on the normally applied force,  $F_1$ . The object prerogative models the force system using a unit normal force and scaled (by the static coefficient of friction) tangential forces. The maximal forces in this system that may be applied independently are then;

$$F_1, F_1 + F_2, F_1 + F_3, F_1 + F_4, \text{ and } F_1 + F_5.$$

The proximity of the contact to the object's center of mass allows the force system to be replaced by a set of wrench space basis vectors. By doing so, we have included local surface properties (surface normal), contact friction properties and global object properties (surface position relative to the center of mass) in a six dimensional subspace representing the degree of constraint due to this contact from the perspective of the object. The object model used to support manipulation consists of an orthonormal basis for this wrench space and the net change in the contact wrench with respect to changes in the surface coordinates,  $u$  and  $v$  as illustrated in Figure 2 (b).

### The Finger Model

The Finger's Prerogative contends that each finger in the contact system should be well conditioned for the task. This requires that the finger be capable of generating forces in some cases and displacements in others. Since the fingers in the Utah/MIT Dextrous Hand are redundant<sup>1</sup>, we express the finger's prerogative by computing (offline) a joint space solution for all positions in the finger's workspace which maximizes a particular conditioning metric. Several such measures have been suggested [15,18,24]. The inverse kinematic models are intended to suggest configurations of the finger which condition it effectively for the generation of grasp forces and allow sufficient workspace in arbitrary directions for the finger to comply to the geometry of the object. These configurations are of course, only idealizations, since unintended contact with the object or the environment will cause the finger to depart from these ideal configurations. This eventuality must

---

<sup>1</sup>the manipulator has more controllable degrees of freedom (i.e., 4 joint angles) than does its workspace (i.e., Cartesian 3 space - no moments, point load).




Figure 2: The Object Model Representing a Point Contact with Friction

be signalled by tactile feedback and accommodated during the compliant conditioning phase.

The quantification of the manipulator in terms of its ability to control the application of forces and velocities to an object have been described in terms of the so-called manipulability ellipsoid[28,29]. This volume of solution space can be generated for any system of linear equations of the form,

$$A \vec{x} = \vec{b},$$

where:

$A = N \times M$  transform,

$\vec{x} = M \times 1$  output space vector, and,

$\vec{b} = N \times 1$  input space vector.

The input and output space nomenclature is used here to describe the (input) joint angle space of the manipulator and the (output) Cartesian forces or velocities. For the case of the 4 DOF finger, the input space is the 4 DOF joint space, ( $N = 4$ ), and the output space is Cartesian 3 DOF space, ( $M = 3$ ). The manipulator kinematics then produce a relationship of the form:

$$\vec{V} = J \vec{\omega},$$

where:

$J =$  the manipulator jacobian ( $\epsilon R^{3 \times 4}$ ),

$\vec{\omega} =$  joint angle rates ( $\epsilon R^4$ ), and,

$\vec{V} =$  Cartesian fingertip velocity ( $\epsilon R^3$ ).

The *manipulability ellipsoid*[28] is defined by examining the singular value decomposition of the Jacobian.

If  $J \in R^{M \times N}$ ,  
 Then  $\exists U \in R^{M \times M}$  and  $V \in R^{N \times N}$ ,  
 Such that,  $J = U \Sigma V^T$ , where:

$$\Sigma = \left[ \begin{array}{cccc|c} \sigma_1 & & & & 0 \\ & \sigma_2 & & & \\ & & \sigma_3 & & \\ & & & \ddots & \\ & & & & \sigma_m \end{array} \right] \in R^{M \times N}, \text{ with,}$$

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_m \geq 0.$$

These  $\sigma_i$  represent the  $m$  singular values of  $J$ . If we constrain the joint space angular rates,

$$|\bar{\omega}| = [\omega_1^2 + \omega_2^2 + \dots + \omega_n^2]^{1/2} \leq 1$$

then we find that the resulting manipulability ellipse delimits the realizable Cartesian space velocities. The subspace ellipsoid is defined by the principally conditioned axes,

$$[\sigma_1 u_1, \sigma_2 u_2, \dots, \sigma_m u_m]$$

where the  $u_i \in R^m$  are the  $m$  column vectors of the  $U$  matrix above. The manipulability ellipsoid describes the conditioning of the transformation from finger joint space to the fingertip Cartesian velocity space.

The Principally Conditioned Axes (PCA's) resulting from the singular value decomposition support the computation of performance indices which reflect the capability of the finger given its joint space configuration. The PCA's themselves represent the *instantaneous velocity conditioning* of the finger configuration. The PCA ellipse can be related to the object's center of rotation to create a twist system which addresses a specified goal.

While the manipulator geometry may be capable of generating instantaneous velocities, it may not be able to do so in the region immediately surrounding its current state. The PCA's are weighted to reflect the physical joint limits of the manipulator. This is accomplished by defining a raised cosine weighting coefficient that is a function of a hypothetical elastic strain energy resulting from the displacement of a joint from its nominal position:

$$W(\vec{\theta}) = \left[ \cos \left\{ \left( \sum_{i=0}^3 \left( \frac{\theta_i - \theta_{nom}}{\theta_{max} - \theta_{min}} \right)^2 \right) \times \pi \right\} + 1.0 \right] / 2.0$$




Figure 3: Weighted Principally Conditioned Axes for the Utah/MIT Finger

The results of applying such a weighting coefficient to varying configurations of the Utah/MIT finger is presented in Figure 3. A finger is viewed from the side and from behind to illustrate the weighted PCA's for a variety of finger geometries.

The index used to discriminate between competing finger configurations is currently proportional to the volume of the weighted manipulability ellipsoid. The manipulability index suggested by Yoshikawa [28] is modified by expressing the joint limits as above and computing the product of the weighted singular values. An inverse kinematic solution based on this metric produces finger configurations well conditioned for the application of forces and velocities in directions defined by the geometry of the finger. Moreover, the index itself describes a vector field for which a well behaved gradient space toward isotropic configurations can be defined.

### The Object Prerogative

Given a set of positions on the object's surface through which interactions may take place, it is possible to produce an incremental improvement in the system state by defining migrations of these sites over the object's surface. It is the object's prerogative to select features of the object and interaction topologies particularly well suited to the task. In this section, we will define the error of a particular set of surface contact positions relative to a task, and compute a trajectory over the surface of the object which reduces this error.

Unions over sets of contact wrench systems span a subspace of the six degrees of freedom of the object. We require as our criterion of stability, that the constraint span all six degrees of freedom. That is, a basis for the union of the wrench subspaces over all contacts must be of rank 6. Given a contact system for which this is not true, it is the object's prerogative to improve the contact system by changing the surface position of the interactions.

A system is quantified by computing the error of the current state with respect to the

task. The task specification is the sum of two terms as shown below.

$$\vec{T} = [\Omega_{R1} \ \Omega_{R2} \ \Omega_{R3} \ \Omega_{R4} \ \Omega_{R5} \ \Omega_{R6}] + [\Omega_{T1} \ \Omega_{T2} \ \Omega_{T3} \ \Omega_{T4} \ \Omega_{T5} \ \Omega_{T6}]$$

where:

$\Omega_{Ri}$  = the task specified force robustness threshold,

$\Omega_{Ti}$  = terms used to anticipate the task, inertial or external loads.

These task descriptions represent the thresholds of contact impedances required for the six degrees of freedom of the object. The system is, therefore, not permitted to be stable only with respect to infinitesimal perturbations, but is compelled to seek contact topologies which are predictably robust with respect to arbitrary perturbations. We view this approach as producing a more robust and predictable stability than is possible with analytic methods (such as Liapunov or Hessian criteria).

If we represent the wrench space spanned by a contact system using a set of  $n$  orthonormal 6D basis vectors, then we may express the error of this system relative to the task as follows.

$$\vec{E} = \vec{T} - \sum_1^n mag_i \times \hat{b}_i$$

where:

$\vec{E}$  = Contact system error relative to the task

$\vec{T}$  = task vector

$\hat{b}_i$  = an orthonormal basis vector for the contact wrench space

$mag_i$  =  $\min(\text{the contact system capability along } \hat{b}_i, \text{ the projection of } \vec{T} \text{ onto } \hat{b}_i)$

The procedure above removes the components of the task which project onto the contact wrench space and are within the magnitude limitations of the contact system. The residual 6D vector is, therefore, the deficiency of the current contact system relative to the task.

In order to reduce this object state error, the planner interrogates the object model to determine the value of the derivative wrench systems with respect to orthogonal surface migrations. If the error is expressed as a linear combination of the derivative wrenches, a migration of the interaction sites in surface coordinates can be defined which incrementally improves the state of the system. The procedure is illustrated in Figure 4 for a system consisting of an infinite cylinder and a single contact point. The Object Prerogative attempts to eliminate stability deficiencies and effectively condition the system for a force in the  $-\hat{y}$  direction by moving a single contact over the surface of the cylinder. Released from a non-optimal position, the system quickly converges to the site for which the wrench space error is minimized as shown in Figure 4.

A second demonstration is presented in Figure 5. Here, a migrating interaction site is directed to improve the net wrench space defined by it and a fixed interaction site. The task consists of a uniform stability robustness. The migrating interaction site seeks




Figure 4: The Directed Migration of a Single Interaction Site

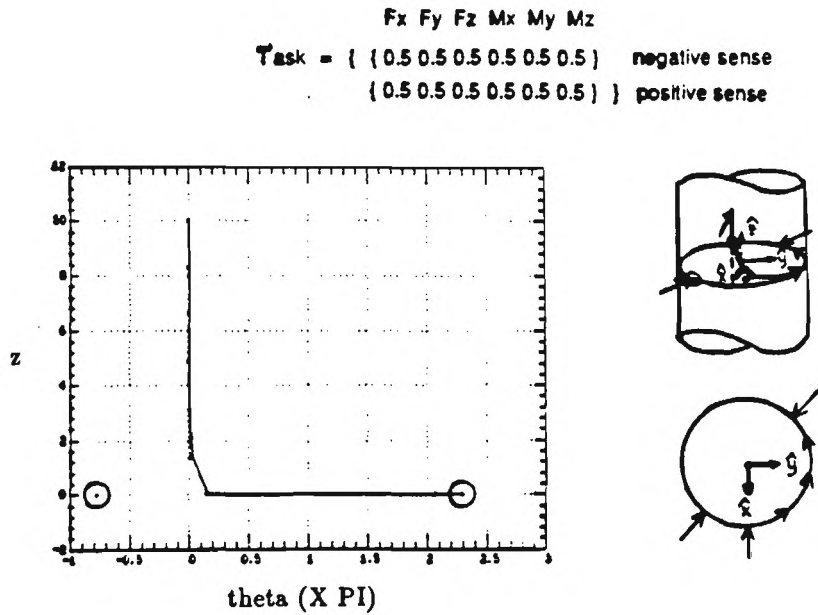


Figure 5: The Directed Migration of a Two Interaction Site Contact System

a position on the object's surface where a wrench subspace complementary to that of the fixed interaction site can be generated. It is in essence a two dimensional gradient following across a six dimensional manifold. The task is never completely satisfied by this contact system, but the state error with respect to the task is minimized.

In practice, the object prerogative is constrained by the ability of the hand to accommodate the surface trajectory. Therefore, only that portion of the trajectory that is reachable by the hand in its current configuration is used. The object prerogative is, therefore, constrained by the amount of the workspace available along the trajectory for the finger associated with the contact site. The contact position which most effectively addresses the system deficiencies can be identified and a trajectory of this site toward the stability robustness goal can be computed.

### **The Hand Prerogative**

We remarked earlier, that the hand is required to comply to the task and to the geometry of the object. To accommodate the trajectories over the surface of the object defined by the Object Prerogative, the hand coordinate frame must move to positions for which all fingers are well conditioned. The responsibility of the hand to comply is supported by the workspace model for the fingers. This inverse kinematic model contains the weighted manipulability index at every position within the workspace. This index defines a smooth continuous scalar field with a single maximum which can be used to define a gradient space toward well conditioned configurations. It is critical to the conditioning process that the scalar field has these properties, they constrain the range of candidate weighting functions we may use to express the joint limits of the manipulator. During the phase of the conditioning process when the Object Prerogative is not yet satisfied, this manipulability gradient is used to direct the hand into configurations which are generally well suited to arbitrary displacement tasks.

Following the convergence of the Object Prerogative, the hand must comply to the requirements of the task. The Principally Conditioned Axes, weighted to reflect joint limits, were developed in the Finger Model section to reflect the conditioning of the finger anywhere within its workspace. When several fingers contribute to a contact system, we may use the PCA's to compute the wrench space capabilities of this configuration relative to the object's center of rotation. In this way, the wrench subspace spanned by the contact system may be used to define an error with respect to the task as was described in the discussion of the object's prerogative. This error is resolved by expressing the *hand's prerogative*, suggesting a migration of the hand frame which optimally applies these principally conditioned axes to the object<sup>2</sup> in light of the task.

To accomplish this objective, the transform representing the position and orientation of the hand frame relative to the object is changed by a virtual displacement. The error resulting from this hypothetical state is compared with that resulting from other virtual displacements and the current state. The trajectory that reduces the error by the greatest amount is selected. The process continues until no further improvement in the hand state is possible with the current set of contact sites. Once again, tactile feedback must eventually be incorporated into this planning structure to acknowledge additional

---

<sup>2</sup>specifically considering the local normal and the center of rotation




Figure 6: The Conditioning of One Finger for an Isotropic Wrench

contacts that occur as the hand moves relative to the object.

### 2.3 Behavior of the Conditioning Agent

Figure 6 illustrates the result of submitting a task to the Conditioning Agent. The figure illustrates the top view of a system consisting of a cylindrical object four inches in diameter and a two fingered Utah/MIT hand. The actual fingertip position is displaced slightly from the correct surface position. This is due to the coarseness of the inverse kinematic solution<sup>3</sup>. In practice, the planning must use a virtual object which is slightly larger than the real object to ensure that no contact actually takes place while the grasp is evolving. The Conditioning agent integrates the hand's prerogative and the object's prerogative. The task illustrated is to generate an isotropic wrench space<sup>4</sup> using only the index finger of the hand. The state of the system cannot be improved from the object's perspective, but the hand frame migrates to a position relative to the object for which the index finger is more isotropically conditioned.

Figure 7 presents the results of modifying the task presented to the Conditioning Agent. In addition to the isotropic wrench space, a preferred force in the negative  $\hat{y}$  direction is requested. The object prerogative is not immediately satisfied in this case, it directs the migration illustrated in the figure. The hand frame complies to the object's prerogative producing a markedly different behavior in the system. Prior to the time that the object prerogative is satisfied, the Hand Prerogative seeks to improve the conditioning of the fingers isotropically. This directs the hand to seek well conditioned states, from which it will be better prepared to accommodate subsequent state changes. A wrench/twist space task is submitted to the agent representing the hand's prerogative only after the

<sup>3</sup>the characteristic separation between bins in the workspace models is 0.25 inches

<sup>4</sup>A uniform, robust stability task




Figure 7: The Conditioning of One Finger for an Isotropic Wrench +  $F_{-y}$

object prerogative has converged, therefore, the hand frame retreats slightly from its most advanced position.

The behavior of the Conditioning Agent is demonstrated once again in Figure 8. The system now consists of a two fingered contact given the same task as the previous example. The system is directed to improve its state by altering the position of the contact created by the index finger; the contact created by the thumb remains fixed.

Another demonstration of the behavior of the Conditioning Agent is presented in Figure 9. Here, we have presented the top and side views of the evolving grasp. The task is once again a uniform wrench representing robust stability. The example illustrates a 3D behavior when both contact positions may vary to improve the state of the system. The system required 5.4 seconds of CPU time on an Vax 750 to produce the behavior illustrated in Figure 9.

Finally, we have applied the Conditioning Agent to the geometry of the Utah/MIT hand. The resulting manipulator can reproduce the results presented earlier by defining a contact system consisting of the thumb and the index finger. Moreover, we may request three and four fingertip contacts and use arbitrary combinations of fingers. A simple, four fingered grasp of the cylinder is presented in Figure 10. Figure 10(a) presents the top view and side view of the initial hand/object configuration. The initial configuration of the fingers is identical, so that the top view in Figure 10(a) appears to show only the thumb and a single finger. Figure 10(b) illustrates the intermediate hand coordinate frame positions and the final configuration of the hand. These results demonstrate that the system can be applied to multiple finger contacts, but it required 18.6 seconds of CPU time on the Vax 750 to reach the final state. Managing four independent contacts is pushing the real time capabilities of the system. It is conceivable that the concept of virtual fingers [3], identified by experience with a class of objects and tasks, may help alleviate the computational burden on the planner while still allowing the system to




Figure 8: The Conditioning of Two Fingers for an Isotropic Wrench +  $F_y$







Figure 9: The Conditioning of Two Fingers for an Isotropic Wrench



(a) initial geometry



(b) final geometry




Figure 10: The synthesis of a four fingered grasp for the Utah/MIT hand

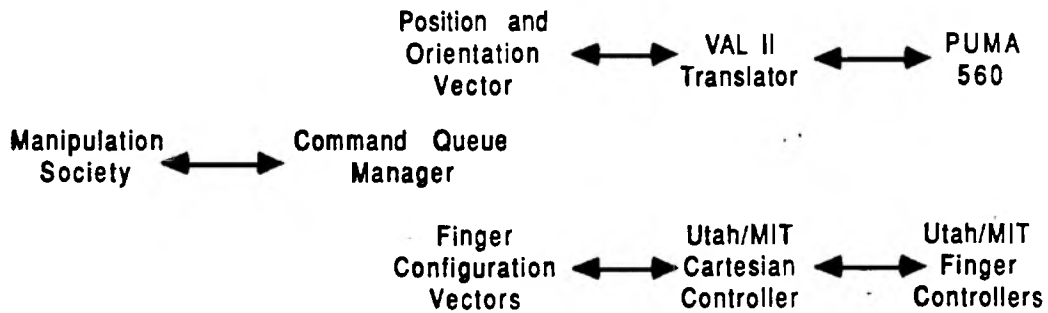


Figure 11: The Control Structure for a Manipulation Test Bed

respond to new or unexpected circumstances.

### 3 Discussion and Future Work

The development of the system will be supported by a graphical simulator and by a robotic manipulation testbed constructed using the Utah/MIT hand. This capability will allow the user to learn how to properly express tasks in wrench space, and will provide insight into methods of learning in manipulation.

The system illustrated in Figure 11 is being constructed to conduct experiments. The manipulation society will run on a VAX 750 and create a child process which maintains a command queue. The planner can then proceed at its characteristic rate while the child process submits subtasks to the mechanical systems and waits for the completion of those tasks asynchronously.

The task submitted to the mechanical arm is a homogeneous transform representing the hand frame position and orientation. The transform is translated into a sequence of VAL II<sup>5</sup> expressions which deliver the hand frame to the correct position and orientation.

The tasks submitted to the Utah/MIT hand consist of the position of the contact site, the orientation of the last phalange and stiffness matrix for every finger, in hand frame coordinates. This task is distributed over four 68020's which act as a Cartesian controller for the hand.

The task may also require the system to partition the contact set into cooperating subsets with independent tasks. This allows a system which has become ill conditioned in the process of executing a task to stably constrain the object with a subset of the contact system while reconditioning the complementary subset. The organization of the system lends itself to the partitioning of contact elements into sets which may respond to independent tasks. The ability of the system to accomodate tasks which require finger replacement, or which require multiple tasks within a contact system will be developed.

Once again, we stress that the results presented in this paper represent only the

<sup>5</sup>VAL II is the control language for the PUMA 560 robot used in these experiments

conditioning phase of the hand/object interaction. The mechanical system may be idle or performing an entirely different task while the Conditioning Agent defines a contact geometry for a new object and task. Furthermore, the mechanical system need not respond to all the intermediate states produced by the Conditioning Agent. Alternatively, we may submit only the final state to the mechanical system.

The Conditioning Agent concludes by defining the contact geometry, the command to the arm and hand will be relative to a virtual object. The actual contact must be achieved by implementing a compliant guarded move. The Utah/MIT hand facilitates this action by combining Cartesian stiffness control with low-level reflexive movements, such as proximal stiffening and distal curling [14]. An unexpected contact on a proximal phalange causes the distal kinematic chain to curl, while more proximal joints stiffen. Taction is not required for this behavior, contacts are sensed by measuring tendon tensions. Reflexive movements can be applied at low levels to improve the application of the fingers to the object while these contact sites are managed by the planner.

A characteristic of the society control structure unified by a common task, is that each component of the grasp complements the others. That is, two fingers responding to the same wrench space task will cooperate rather than conflict. In the simple experiments presented earlier, the interactions sites never converge to the same position on the surface of the object since this geometry spans the smallest volume of the wrench space task.

Another result observed in the development of this control society structure is that the use of potential functions, or gradient following, to seek solutions can be useful (despite well known problems with local minima) if the complexity of the gradient space is minimized. The objectives of complex gradient spaces can be distilled into separate *prerogatives* which represent independent goals. Moreover, we may use the results of independent, but sequentially prioritized prerogatives to further constrain the gradient following: the hand frame migrates to a position which conditions the fingers for displacements along the surface of the object, the positions through which interactions occur traverse the object's surface to improve the state of the object with respect to the task constrained by the finger's ability to comply, and finally, when this state can no longer improve, the hand agent addresses the application of task defined forces and velocities.

The Actuation Agent presented in Figure 1 must control the hand/object geometry designed by the Conditioning Agent. The problem facing the Actuation Agent is constrained by the results of the Conditioning Agent. The Equilibrium and Displacement Prerogatives illustrated in Figure 1 generate joint torques and velocities given knowledge of the surface character, the contact types, and the ability of the fingers to generate forces and velocities.

A great deal of specific knowledge is recorded in the models of the hand and object; therefore, modeling is a primary concern. The knowledge necessary to express the object's prerogative is entirely contained in the object models. The demonstrations presented earlier represent the cylinder analytically, since such an approach works well for primitive shapes. In general, it will be necessary to tag surface models with the surface's wrench space, and to build derivative wrench spaces which highlight graspable surface elements. The surface model will instantiate a *realm of influence* for the derivative wrench spaces, thus accenting surface elements which are especially useful<sup>6</sup>. The exact nature of the

---

<sup>6</sup>Such as *finger sized* concavities in the surface

models use to represent general classes of objects and the automatic generation of such models from a CAD representation will be examined further.

We also wish to investigate the extent to which the object model can be used to support learning in manipulation. It is possible to learn which initial hand/object interactions permit the system to migrate to a solution. A manipulation strategy can be discovered, therefore, by cataloging the portion of the objects surface which is upstream of a solution along the object models gradient space for a particular class of tasks. The state error integrated over a succession of intermediate states can be used to identify initial approach vectors and orientations for the hand and contact sets which provide effective manipulation strategies. A tessellated sphere representing the object can be used to delineate hand approach vectors which have in the past proved useful for a particular class of tasks. In addition, we would like to examine the automatic classification of virtual fingers, to exploit the reduction in complexity warranted in certain tasks. Skill refinement behaviors based on these ideas will be investigated further.

## References

- [1] P.K. Allen. Surface Descriptions from Vision and Touch. In *Proceedings of the IEEE Conference on Robotics*, pages 394-397, March 1984.
- [2] P.K. Allen. Sensing and Describing 3-D Structure. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 126-131, April 1986.
- [3] M. Arbib, T. Iberall, and D. Lyons. *Coordinated Control Programs for Movements of the Hand*. Technical Report COINS 83-25, University of Massachusetts, 1983.
- [4] R.M. Bolle and D.B. Cooper. On Optimally Combining Pieces of Information, with Application to Estimating 3-D Complex Object Position from Range Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(5):619-638, 1986.
- [5] R.C. Bolles and R.A. Cain. Recognizing and Locating Partially Visible Objects: The Local-Feature-Focus Method. *Robotics Research*, 1(3):57-82, 1982.
- [6] R.C. Bolles and P. Houraud. 3DPO: A Three-Dimensional Part Orientation System. *Robotics Research*, 5(3):3-26, 1986.
- [7] M.R. Cutkosky and P.K. Wright. Modeling Manufacturing Grips and Correlations with the Design of Robotic Hands. In *Proc. of IEEE Conference on Robotics and Automation*, pages 1533-1539, 1986.
- [8] H.A. Ernst. *MH-1, A Computer-Operated Mechanical Hand*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Mass, December 1961.
- [9] C.G. Geschke. A System for Programming and Controlling Sensor-Based Robot Manipulators. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-5(1), January 1983.
- [10] W.E.L. Grimson. Disambiguating Sensory Interpretations Using Minimal Sets of Sensory Data. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 286-292, April 1986.
- [11] H. Hanafusa and H. Asada. A Robot Hand with Elastic Fingers and its Application to Assembly Process. In *IFAC Symposium on Information and Control Problems in Manufacturing Technology*, Tokyo, 1977.
- [12] C. Hansen and T. Henderson. CAGD-Based Computer Vision. In *Proceedings of the IEEE Workshop on Computer Vision*, page to appear, Miami Beach, Florida, November 1987.
- [13] S.C. Jacobsen, J.E. Wood, D.F. Knutti, and K.B. Biggers. The Utah/MIT Dextrous Hand. In *MIT/SDF IRR Symposium*, August 1983.
- [14] S.C. Jacobsen, J.E. Wood, D.F. Knutti, and K.B. Biggers. The Utah/MIT Dextrous Hand Work in Progress. *Journal of Robotics Research*, 3(4):21-50, Winter 1984.

- [15] C. Klein and B. Blaho. Dexterity Measures for the Design and Control of Kinetically Redundant Manipulators. *Journal of Robotics Research*, 6(2), Summer 1987.
- [16] H. Kobayashi. Control and Geometrical Considerations for an Articulated Robot Hand. *Journal of Robotics Research*, 4(1), Spring 1985.
- [17] H. Kobayashi. Grasping and Manipulation of Objects by Articulated Hands. In *Proc. of the IEEE Conference on Robotics and Automation*, pages 1514-1519, April 1986.
- [18] Z. Li and S. Sastry. Task Oriented Optimal Grasping by Multifingered Robot Hands. In *Proceedings of IEEE Conference on Robotics and Automation*, pages 389-394, Raleigh, NC, April 1987.
- [19] D. Lian, S. Peterson, and M. Donath. A Three-Fingered, Articulated Hand. In *Proc. SME 13th International Symposium on Industrial Robots*, 1983.
- [20] M.T. Mason. Compliance and Force Control for Computer Controlled Manipulators. *IEEE Systems, Man, and Cybernetics*, 11(6):418-432, June 1981.
- [21] M. Minsky. *The Society of Mind*. Simon and Schuster, New York, NY, 1986.
- [22] T. Okada. Object-Handling System for Manual Industry. *IEEE Systems, Man, and Cybernetics*, 9(2), February 1979.
- [23] R. Paul and B. Shimano. Compliance and Control. In *Proc. of the Joint Automatic Control Conference*, pages 694-699, Purdue University, July 1976.
- [24] J.K. Salisbury. *Kinematic and Force Analysis of Articulated Hands*. PhD thesis, Stanford University, May 1982.
- [25] J.K. Salisbury and J.J. Craig. Articulated Hands: Force Control and Kinematic Issues. *Journal of Robotics Research*, 1(1), Spring 1982.
- [26] R. Tomovic, G. Bekey, and W. Karplus. A Strategy for Grasp Synthesis with Multifingered Robot Hands. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 83-89, April 1987.
- [27] D.E. Whitney. Force Feedback Control of Manipulator Fine Motions. *Journal of Dynamic Systems, Measurement, and Control*, 91-97, June 1977.
- [28] T. Yoshikawa. Analysis and Control of Robot Manipulators with Redundancy. In *Robotics Research: The First International Symposium*, pages 735-747, 1984.
- [29] T. Yoshikawa. Manipulability of Robotic Mechanisms. *Journal of Robotics Research*, 4(2), Summer 1985.