Integration of Tactile Sensing and Robot Hand Control

A thesis presented

by

Jae S. Son

to

The Division of Applied Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Engineering Sciences

Harvard University

Cambridge, Massachusetts

May 1996

Dedicated to my devoted father.

Suck-Bong Son

(1939 - 1995)

Acknowledgements

This thesis could not have been completed without the numerous ideas and suggestions of Prof. Robert Howe along with his tiredless editing and motivation for the cause of tactile sensing. I would also like to thank him for his understanding and support during my personally challenging times and providing me with the opportunity to work with other researchers. I've learned valuable lessons from Prof. Mark Cutkosky, and enjoyed my semester at Stanford which was an enriching experience. I would also like to thank Prof. Greg Hager for his assistance in helping me finish my last set of experiments at Yale.

My years at Harvard have been enriched by the members of the "sensors lab" due to the wide range of diversified talent. I look forward to following the future successes of Eric Dunn, Aram Hajian, Dimitri Kontarinis, Dianne Pawluk, Bill Peine, and Parris Wellman.

I am greatful for the financial support from the Hughes fellowship program, and the summers at Sensors and Mechanism Department which enabled me to start each academic year refreshed. This thesis has also been supported by ONR grant N00014-92-J-1814 and N00014-92-J-1887.

Abstract

The main contribution of this thesis is the advancement of robot hand dexterity in manipulating objects by using tactile sensing in robot hand control. Many issues in manipulation such as transient events, determining the contact location on a finger, grasping and grasp refinement, and object-to-object interactions are examined in detail. The first two chapters provide useful insights for selecting tactile sensors and algorithms. The following grasp refinement chapter combines the use of visual and tactile feedback for controlling a robot. Finally, Cartesian Object Stiffness Controller is used to study fundamental performance limits of a robot hand to modulate the stiffness of a grasped object. The consequences of these limitations and errors are illustrated by examining insertion forces for a high tolerance peg insertion task. The results of this thesis show that tactile sensing is able to increase robot hand dexterity by providing information about the hand object-relation and the mechanical state of the manipulation process.

Table of Contents

1. Introduction

1.1 Motivation

Robots will become useful in our daily lives when they are able to perform tasks in unstructured human environments. One limiting factor that prevents current robots from achieving this goal is the inability to sense the environment and behave in an intelligent manner. With respect to dextrous manipulation, kinematics, vision, and the sense of touch are critical. Of these sensing modalities, tactile sensing is furthest away from practical use. Therefore, thesis focuses on issues related to the use of tactile sensing for robot hands.

The sensing and actuation system of the human hand is remarkable. It permits wide range of objects to be grasped and manipulated in a dextrous manner. However, multifingered robot hands are far away from the abilities of our human hands. Among the problems encountered are limitations in mechanical robot hand design which result in friction and backlash, difficulties in coordination of many degrees of freedom, and inadequate sensors. Nevertheless, as with human hands, rolling and sliding of the object between the fingers greatly increase the mobility of the object and the dexterity of the robot hand in comparison to simple parallel-jaw robot grippers. This also increases uncertainty of the object location within the hand. Tactile sensing promises to increase robot hand dexterity by providing information about the hand-object relationship and the mechanical state of the manipulation process. The main contribution of this thesis is the advancement of robot hand dexterity in manipulating objects through the use of tactile sensing along with existing theories on multifingered robot hand control.

The sense of touch is an important component of dexterous manipulation. Consider the inevitable clumsiness we experience when our sense of touch is diminished by wearing gloves. Physiological experiments show that when humans are deprived of tactile sensation with anesthetics, manipulation ability is seriously degraded even though musculature is unimpaired [Johansson and Westling 1984a]. Many different parameters may be sensed through touch, including contact forces [Salisbury 1985, Bicchi 1990], pressure distributions [Fearing 1990], and high frequency vibrations [Patterson and Neville 1986, Morrell 1990, Howe and Cutkosky 1989]. Experiments with multifingered robot hands demonstrate that without tactile feedback small unmodeled errors can accumulate, causing failure in complex manipulation tasks [Fearing 1987]. Although many researchers have studied tactile sensing devices, little work has appeared on using tactile information for useful manipulation tasks. Recent research on the use of tactile information in real-time manipulation include Berger and Khosla [1991] and Chen et al. [1995] who used tactile array sensors to track object edges in sliding. Maekawa et al. [1995], Nicolson [1995] and Son and Howe [1996] used tactile information to reorient objects by rolling them between the finger tips. Tactile sensing is also used to sense contact [Son et al. 1994] and determine initial object orientation [Son and Howe 1996, Son et al. 1996].

1.2 A Prototypical Example

Typical manipulation tasks that we perform on a daily basis can be broken down into phases which are separated by events [Hyde 1995, Howe and Cutkosky 1992, Howe et al. 1990]. For example, the simple grasp, lift and replace task can be segmented into phases and events where the first phase of the task requires position control of the finger tips as they approach the object to be grasped. The first event, make contact, signals the transition from position to force control. The grasp force is then increased simultaneously during the loading phase, and the next event is the breaking of contact between the object and the surface on which it rested. This event is classified as remote break contact. Now the object is grasped and the position of the object must be specified so this phase is called the manipulation phase. As the object is replaced and makes contact with the surface, the

remote make contact event signals the transition to unloading phase where the grasp and loading forces are decreased. When the grasp force reaches zero, contact between the finger tips and the object are broken and this event is classified as a break contact. Now the finger tips are retracted in a position control mode. Here, drinking a cup of coffee is used as an example to discuss the various phases of the manipulation task and the sensing requirements. This classification of phases and events is extended to include pre-approach phase and object-to-object interaction phase.

1.2.1 Pre-Approach

Imagine sitting at a desk reading and deciding take a sip from a coffee mug that is sitting on the desk. During the pre-approach phase, we use our sense of vision and kinesthesia. This phase requires that we first locate the mug. The most effective method for determining the location of the mug is to use our sense of vision because large areas can be sensed quickly. If we were reading and didn't want to interrupt this process, groping, which uses the kinesthesia and tactile sensing, could be used with much less effectiveness.

Once the object location has been determined, kinesthesia is used to move our arm towards the mug. Here, vision and tactile information may not be required since the hand is most likely out of the visual field and isn't touching anything. However, as our hand approaches the mug, it comes into field of view, and we use our visual senses to control the hand and finger positions relative to the mug. This visual control scheme cannot be used when the hand occludes the mug. At this point, the hand is so close to the mug, feedforward kinesthesia along with tactile sensing takes precedence over vision.

1.2.2 Acquiring the Object

During object acquisition, the fingers approach the object, and we rely on our sense of touch to determine when contact has been made. This phase is called the approach phase. One mechanical design feature of human fingers is the built-in compliance that minimizes initial contact forces. Finger tip mechanical compliance has been advocated by many researchers in various embodiments [Brockett 1985, Cutkosky 1985, Russell 1987, Akella

and Cutkosky 1989, Nowlin 1991]. However, adding compliance without tactile sensing is analogous to putting on gloves to perform a dextrous manipulation task. The ability to detect small contact forces is also important, especially when grasping objects like a hot cup of coffee.

Once contact has been made with one finger, that finger rests on the object until the others also make contact. Before increasing the grasp force, we use contact location and local object curvature information to determine if the grasp is stable and optimal. Corrections are made using tactile sensing to control sliding and rolling while optimizing the grasp configuration. Now the grasp force can be increased to lift the object in the loading phase. Details of these phases and events are examined in detail by Cutkosky and Hyde [1993].

1.2.3 Object Manipulation and Interaction

Once the mug is grasped and lifted, the important considerations are the grasp force and the position of the object. Grasp force control in humans are studied in detail by Johansson and Westling [1984b, 1988, and 1990]. To control object position, we rely on vision and kinesthetic information. When the position and orientation of the mug needs to be controlled closely, because the mug is filled to the rim with coffee, we use our sense of vision. Often times we find it difficult to perform visual servoing to prevent spills as when we are walking with the cup in our hand.

As we move the mug closer to our mouth, we watch the mug, but eventually as the mug comes closer to the mouth, the lip of the mug is occluded. Again, we rely on kinesthetic information to guide the mug to our mouth and rely on tactile information to determine when contact has been made. Minimizing forces is important since we don't want to generate excessive forces when contacting the mug to our mouth. The natural passive compliance available in our lips help to reduce contact force transients.

To replace the mug, similar analysis can be used, but the issues remain essentially the same. Kinesthesia, vision, and tactile information needs to be used in such a way that

when one sensing scheme doesn't provide the necessary information, others are on hand to fill the gap. Through intelligent combination of these sensing modalities, robust manipulation in unstructured environments can be achieved.

1.3 Thesis Outline

This thesis is separated into four sections which deal with transient events, determining the contact location on a finger, grasping and grasp refinement, and object-to-object interactions during manipulation. The transient events chapter first examines contact detection abilities of force, skin acceleration and the multi-element stress rate sensors. After describing each sensor used in the experiments, analytical and experimental results show why the stress rate sensor is superior for detecting the onset of contact. Next, the multi-element stress rate sensor's ability to detect incipient slip is demonstrated. This ability enables robot hands to control the grasp force such that minimal grasp forces are exerted without dropping the object.

In the contact localization chapter, comparisons are made between tactile array sensing versus force-torque sensing for localizing contact during manipulation. The manipulation tasks involved rotating objects and translating objects using a planar two fingered manipulator at the Stanford Dextrous Manipulation Laboratory. A pin and a box were selected as limiting cases of point and line contact against a cylindrical robot finger tip. Practical implementations of the two sensor types are compared theoretically and experimentally, and three different localization algorithms for the tactile array sensor are considered. Force-torque contact sensing results suffered from difficulties in calibration and transients. Tactile array sensing was immune to these problems and the effects of shear loading was only noticeable for the simple centroid algorithm. The results show that with care, both of these sensing schemes can determine the contact location within a millimeter during real manipulation tasks.

The grasp refinement experiments were conducted with a vision feedback system and a Zebra-Zero manipulator at the Vision and Robotics Laboratory at Yale University. This

chapter presents results in integrating touch with vision for grasping tasks. Since vision provides position and shape information at a distance, while tactile provides small-scale geometric and force information, the complimentary roles of vision and touch are investigated. Experimental results demonstrate that visual feedback can perform the rough positioning required for tactile sensor feedback, and that grasp force and object orientation angles can be sensed and controlled with tactile sensing. A force sensor based approach provided a comparison measure, and the use of tactile sensing results in a more gentle grasp.

The first part of the object interaction chapter examines fundamental performance limits of robot hands to modulate the stiffness and the center of compliance (CC) of a grasped object. Cartesian object stiffness control is implemented on a two fingered planer manipulator, and the limits to object stiffness are measured experimentally. The lower limit to attainable stiffness may be set by friction in the robot hand, which increases "effective stiffness" until the frictional forces are exceeded. One upper bound to object stiffness is finger tip compliance. Maximum stiffness is also related to the geometry of the grasp configuration: due to controller stability limits, the maximum usable stiffness decreases as the CC is moved away from the fingers, and increases with the distance between the fingers.

The second part of the object interaction chapter examines the role of tactile sensing in object-to-object interactions. In performing an assembly task where objects must interact with each other, object position, orientation, and stiffness must be controlled. Uncertainty in the initial grasp pose, sliding, and rolling can produce errors in the calculated object location if the proper contact kinematics are not considered. Since this controller produces restoring forces proportional to the calculated object position, these errors cannot be corrected with an increase in object stiffness gain. This also leads to inaccuracies in the stiffness of the object, since the restoring forces are in error.

Experimental execution of a peg insertion task demonstrates that tactile sensing can effectively determine the initial orientation of an object. This orientation requirement is necessary to successfully initiate the peg insertion. Once the peg has made contact, object

impedance determines the resulting contact forces. Ideally, a very compliant object with the center of compliance placed at the tip of the peg would minimized interaction forces. In reality, manipulator friction dictates the controlled stiffness requirement for a given task and limits the lowest "effective stiffness" obtainable. Controller stability also limit how far the CC can be placed away from the robot finger tips.

The final chapter starts with conclusions and discussion of the significant contributions. Next, suggestions are made for future work to improve robot hand dexterity and general usefulness. Finally, the chapter closes with personal notes on robotics research and problem solving approaches.

2. Transient Events

2.1 Introduction

The motivation behind this chapter is to clarify how touch sensors can provide useful transient information during manipulation. Consequently, three different touch sensors are used for the analytical and experimental comparison: stress rate sensor, skin acceleration sensor, and the force sensor. In particular, the advantages of the multi-element stress rate sensor, which consists of piezoelectric polymer strips molded into the surface of the rubber "skin" covering the robot finger tip, are highlighted. Experimental results confirm the sensor's ability to detect localized transient information which is important for manipulation control.

Transient vibrations are important indicators of "events" such as contact and slip. Detection of these events is particularly useful in unstructured environments. A simple example is the task of grasping and lifting a delicate object. Knowledge of the contact state is important for good control of this operation because before the fingers make contact with the object the positions of the fingers in space must be controlled, while contact forces must be considered afterwards. In an unstructured environment, the exact location and shape of the object may not be known, so the instant of contact cannot be predicted. Vision may provide rough object location information, but as the fingers come to grasp the object, the contact location becomes occluded, making the visual scheme ineffective. To avoid destabilizing or damaging the object, it is essential to detect the contact at the earliest possible moment so that the controller may respond appropriately.

This chapter starts with a description of each sensor. Next, analysis on a simple lumped mass model along with experimental results compare the three sensor's performance for detecting contact. The multi-element stress rate sensor has also been used to detect incipient slip, which is generated by small micro slips near the contact edges just prior to gross slip. Finally, these results and their implications for sensor design and control of robot hands are discussed in the conclusion.

2.2 Contact Sensors

The tactile sensor that has received the most attention in robotic research is the tactile array, which typically consists of individual pressure-sensitive elements arranged in a rectangular array over the contact surface of the finger tip. Although it's possible to use these sensors for detecting contact, the time response is limited since array sensors are usually multiplexed to reduce the number of wires between the transducer elements and the electronics. This impairs the ability to detect high frequency information. In addition, sensor elements are usually covered with a relatively thick layer of elastomer to reduce aliasing and improve durability [Fearing 1990]; this also limits high frequency sensitivity. Other uses for the array sensor are covered extensively in the following chapters.

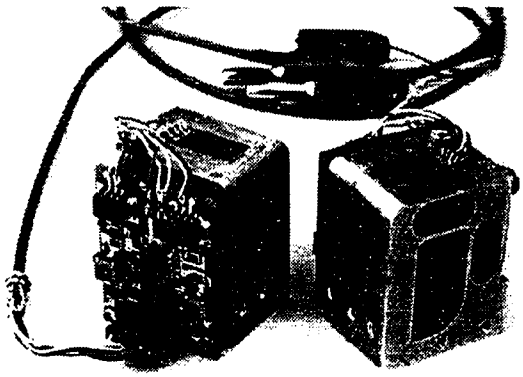

Figure 2.1 The two axis force sensor is designed to measure decoupled forces on the planar manipulator at the Harvard Robotics Lab.

2.2.1 Force Sensor

Many robot researchers have used wrist force sensors, for improving contact force control. By combining force and kinematic information, object shape can be determined, provided that the object is sufficiently large and rigid. One of the sensor used for comparing contact is a two-axis force sensor especially developed here at the Harvard Robotics Lab (Figure 2.1). The flexure design consists of dual parallel beams for each force axis. This reduces torque sensitivity and provides decoupled force signal from each flexture. Surface mount electronics are located next to the flexure to minimize electrical noise pickup. A cylindrical finger tip, which holds the stress rate and skin acceleration sensors, is mounted in front of this force sensor.

2.2.2 Skin Acceleration Sensor

One high frequency sensor of particular interest is the skin acceleration sensor developed by Howe and Cutkosky [1989]. This sensor consists of a miniature accelerometer bonded to the inner surface of the rubber skin, approximately 1.1 mm thick, covering a robot finger tip (Figure 2.2). The 0.5 g accelerometers (Picotron model 8616A500, Kistler Instrument Corporation, Amherst, NY) have the amplifying electronics packaged inside the small casing

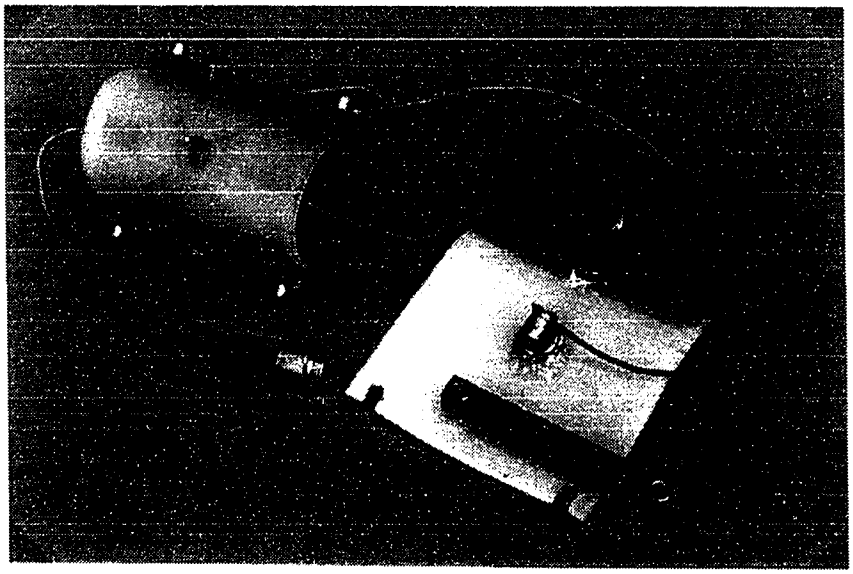

Figure 2.2 Skin acceleration sensor: A small accelerometer is bonded to the inner skin surface which is then wrapped around a cylinderical foam.

(5 mm diameter, 6 mm height) with the output signal and the input power multiplexed on a single coax cable. The accelerometer forms a spring mass system with the local skin stiffness 2000 N/m, and the resonant frequency for this sensor is around 300 Hz, but resonant frequency of 900 Hz was obtained with a thicker and stiffer rubber skin by Howe and Cutkosky [1989]. Beneath the skin is a compliant cylindrical foam structure with a cavity for the accelerometer that is beneficial for grasp stability. The measured stiffness of the foam and rubber skin structure was between 980 - 1270 N/m for a flat half inch diameter indentor. The principal advantage of this sensor is that it intrinsically responds to changes in acceleration, thus providing exceptional sensitivity to high frequency vibrations.

The skin acceleration sensor has been successfully used to detect object slip along with grasp force control to determine the friction coefficient between the finger and the object [Trembley and Cutkosky 1993]. However, application to manipulation tasks with robot hands has proved difficult because the sensor's sensitivity is not spatially localized. Vibrations appearing essentially anywhere in the finger tip are detected by the sensor, including vibrations from sources other than incipient slip such as skin movement and robot mechanism noise. Trembley and Cutkosky [1993] have addressed this problem and shown that by using two accelerometers at different locations on the skin, some measure of disturbance rejection can be achieved. Unfortunately, the large active area of the sensor still poses fundamental limits to its use in slip detection.

2.2.3 Multi-Element Stress Rate Sensor

The multi-element stress rate sensor combines the spatial localization ability of an array sensor with the high frequency derivative sensing advantages of the skin acceleration sensor. Small strips of a piezoelectric polymer are molded into the skin just beneath the surface. These transducers respond to stresses in the skin caused by contact forces and skin shape changes in their immediate vicinity. Signal processing electronics are configured to measure changes in stress, so the sensor is referred to as a "stress rate sensor". An earlier version of this sensor, using a single piezoelectric strip, was designed to measure extremely fine surface features down to a few microns in size by stroking the sensor over the object's surface

[Howe and Cutkosky 1993]. However, this new sensor uses multiple piezoelectric strips and protruding "nibs" on the skin to localize manipulation events.

The prototype version of the multi-element stress rate sensor described here is part of a semi-cylindrical "finger tip" 25 mm in diameter and about 45 mm long (Figure 2.3). There are three primary components: strips of piezoelectric film, a silicone rubber skin with molded surface texture, and a foam core within the skin. Contact with objects in the environment occurs along the side of the cylinder. The skin is 0.8 mm thick and a surface texture consisting of small protruding cylindrical "nibs" has been molded onto the surface. The dimensions of the nibs were, 0.8 mm in length, 0.6 mm in diameter with spherical ends. The nib axes formed a hexagonal pattern with 0.8 mm spacing that minimized the tendency of grasped objects to "squirm" against the nibs. This design was selected to optimize the sensor performance for detecting incipient slip. The skin is then wrapped around a semi-cylindrical polyester foam core and secured at the ends using a clamping mechanism. The stiffness of the nibbed skin was nonlinear following a quadric form with a coefficient of $25 \times 10^6 \, N/m^2$.

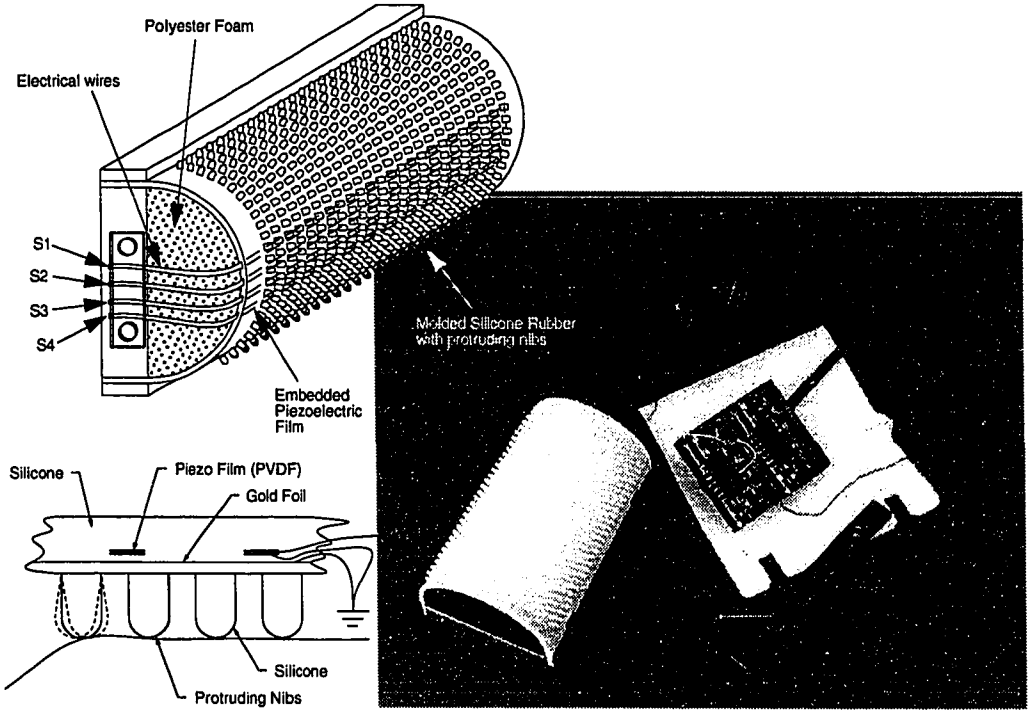

Figure 2.3 The multi-element stress rate sensor is designed to measure localized transient events between the finger and the grasped object.

However, the effective stiffness of the finger tip structure was similar to the skin acceleration sensor.

The piezo film is made of 0.028 mm thick polled polyvinylidene fluoride (PVDF) film (AMP Flexible Film Sensors, Valley Forge, PA), a flexible polymer with very good piezoelectric response. The piezoelectric film acts as a capacitor on which charge is produced in proportion to the strain applied to the film. A thin layer of silver ink on the surface of the film collects the charge generated and provides electrical pathways for signal amplification and processing. The transduction constant between stress and charge density in the directions of the thickness of the film is 33×10^{-12} $(C/m^2)/(N/m^2)$. Along the transverse and parallel directions of the PVDF gain, the transduction coefficients are respectively 3×10^{-12} $(C/m^2)/(N/m^2)$ and 23×10^{-12} $(C/m^2)/(N/m^2)$. The PVDF film was oriented parallel to the skin surface, and the strips were cut such that the direction of higher sensitivity was oriented along the direction of expected slip.

There are four 0.8 mm wide piezo film strips molded into the contact surface of the skin at 3.2 mm spacing. The strips run through the entire width of the finger tip and extend out from one edge of the skin for connecting to the signal processing electronics. To prevent response from the exposed film, this portion of the piezo film was desensitized by heating it to a temperature just above the Curie temperature. In later versions of the sensor, conductive tape was used to attach very thin wires to the silver ink within the rubber skin.

The piezo film strips are connected to a FET-input op-amp configured as a current-to-voltage converter [Howe and Cutkosky 1993]. The piezo film generates a charge q proportional to the stress applied to the film σ. The current i at the input to the amplifier is the time derivative of q, so the output voltage v is proportional to the time rate of change of the stress, or to the stress rate:

$$v = \frac{dq}{dt} R_f \propto \frac{d\sigma}{dt} \tag{2.1}$$

Here R_f is the feedback resistance and the op-amp effectively servos the voltage developed across the film to zero. Since the voltage across the film is minimized, inaccuracies

due to charge leakage through the internal resistance of the film are negligible, and the required amplifier impedance is much lower than a conventional charge amplifier. Furthermore, the sensor has no DC response so it is immune to saturation due to drift and pyroelectric effects which are common problems with piezoelectric tactile sensors [Buttazzo et al. 1986]. The result is a transducer optimized for measuring changes in stress within the rubber skin.

The key to this sensor's success in localizing events on the skin is the limited receptive area for each piezo film element. From solid mechanics we know that in the thin skin material the near-surface stress will be high in regions of high curvature [Fearing1990]. Furthermore, these stresses are far higher than the stress produced by the overall tension in the skin. Since the piezo film strips are located just beneath the surface of the skin, the predominant response is to changes in local curvature. This is in contrast to the situation for the skin acceleration sensor, where small, fast displacements anywhere on the skin are readily detected by the sensor.

The texture on the outer surface of the skin (the pattern of small-scale shapes) is extremely important for determining how local curvature changes when the skin is in contact with an object. A texture consisting of small protruding "nibs" is particularly effective; as the nibs make contact with an object surface they may be bent and even pushed over onto their sides. This generates large curvature changes and stresses at the base of the nib, and the piezo film elements respond very strongly. In defining the skin texture, numerous variations on nib size, shape, aspect ratio and spacing were tested. The size of the nibs and the surface roughness on which slip was occurring were correlated in such a way that larger nibs required a rougher surface to excite or pluck the nibs. However, changing from a flat ended nib to a hemispherical ended nib improved the detection of slip on smooth surfaces. This is probably due to the fact that it is easier to pluck a cylindrical nib rather than a flat ended nib. Furthermore, longer nibs were more easily excited during slip, but shorter nibs produced a more localized signal.

Because of the derivative nature of this sensing modality, there is a great deal of high frequency information present in the signal. Although the nibs which amplify micro slips have a mechanical resonance of 5 kHz, low pass filtering was required to obtain a clean signal which could be used in closed loop control. Thus, following current-to-voltage conversion the sensor signal is amplified and low-pass filtered. Surprisingly, the optimal 3 dB frequency for detecting incipient slip was found to be about 30 Hz for the 2 pole low pass filter; this minimized high frequency noise while leaving the slip signal undiminished.

2.3 Detecting Contact

The importance of the contact transition is illustrated in numerous research studies [e.g. Hyde 1995, Hogan 1985]. Contact events can be detected using virtually any type of sensor, including joint angles or joint torque measurement, force sensors, tactile array sensors, multi-element stress rate sensors, and skin acceleration sensors (Figure 2.4). How well these sensing modalities work is another question. The purpose of this section is to determine the relative sensor performance for detecting contact.

Some of the sensing methods are inherently slow or insensitive. For example, looking at joint angles to see if the finger stopped moving may displace all but the most massive objects. Joint torque sensing is furthest from the event phenomena so it is filtered by, friction, mass, stiffness, and damping of the links, and the finger tip. The tactile array

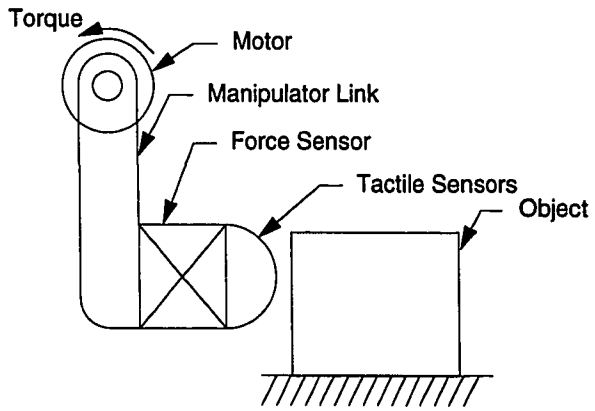

Figure 2.4 Schematic drawing of a one-link finger with a variety of sensors for detecting contact with an object include: stress rate, skin acceleration, tactile array, motor torque, or joint angle position sensing.

sensor time response is dramatically reduced by the required multiplexing. Three sensor choices remain: force sensor, skin acceleration sensor, and the multi-element stress rate sensor.

2.3.1 Physical Model

Sensor performance comparisons of this nature are difficult to generalize since the parameter being sensed is different for each sensor and a particular sensor sensitivity cannot not be decoupled from the experimental conditions. As an example, there is no absolute conversion factor which relates stress rate to acceleration. In addition, the dynamic range and signal to noise ratios are dependent on sensor design and construction techniques. Therefore, rather than looking at absolute gain sensitivity, sensor gain variations in the frequency domain are compared.

Figure 2.5(a) shows how all three sensors were physically incorporated into one finger tip, and figure 2.5(b) shows the simplified lumped mass model of the finger tip. The model assumes that the force sensor is grounded, and that the stiffness and damping coefficients of the force sensor are given by k_1, and b_1. Finger tip mass, m_f, is lumped at the back of the finger tip since most of the weight comes from the mounting hardware and the electronics, and the effective stiffness and damping coefficients for the foam and the rubber skin structure are given by k_f, and b_f. The accelerometer mounted on the skin is modeled as an independent system since the acceleration signal oscillates during contact while the skin

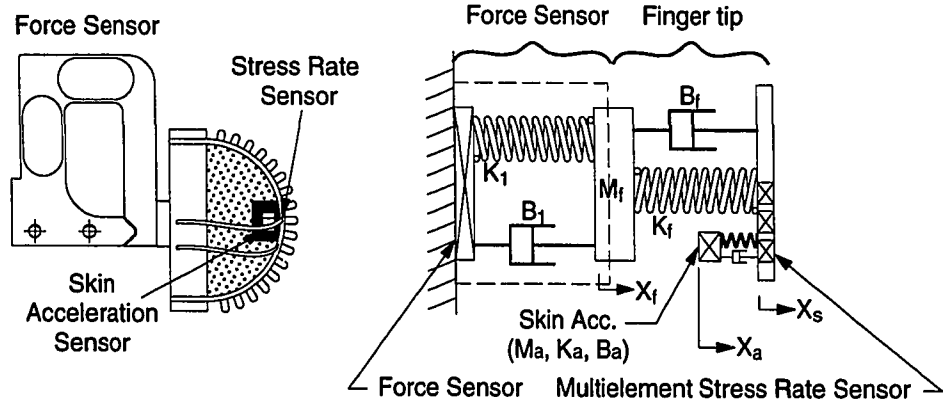

Figure 2.5 Left: Physical structure of the finger tip with force, skin acceleration, and the stress rate sensors. Right: Lumped-element model of the physical structure.

remains in contact with the object and does not bounce. The parameters for the skin acceleration sensor are mass, m_a, stiffness, k_a, and damper, b_a. Since the mass of the skin around the multi-element stress rate sensor is small and distributed, the effective skin mass is assumed to be negligible while the position of the skin is defined as x_s.

This model has three degrees of freedom, one of which is independent of the other two. For simplicity, equations (2.2) and (2.3) describe the motions for the mechanical model without the damping terms since the analysis is aimed at obtaining qualitative comparision results.

$$m_a\ddot{x}_a = k_a(x_s - x_a) \tag{2.2}$$

$$m_f\ddot{x}_f = k_f(x_s - x_f) - k_1 x_f \tag{2.3}$$

Acceleration and position of the accelerometer is defined as $\ddot{x}_a$ and x_a, and the finger tip acceleration and position are given by $\ddot{x}_f$ and x_f. The position of the skin, x_s, is where contact occurs, and is taken as the input to the system. When a robot contacts a rigid object, the object to resting-surface frictional force must be exceeded before the object starts to slide. Ideally, sensors should detect this contact before the object starts to move. Therefore, a displacement of the finger tip skin models the situation more appropriately.

Equations (2.4) and (2.5) are respectively the forces at the finger tip skin and the force measured by the force sensor.

$$f_s = k_f(x_f - x_s) \tag{2.4}$$

$$f_1 = k_1 x_f \tag{2.5}$$

For simplicity, f_s is assumed to be only a function of k_f, since the forces generated by the accelerometer mass are small in comparison to the finger tip stiffness forces.

Equations (2.6), (2.7), and (2.8) are the transfer functions in the Laplace domain for displacement inputs at the skin, x_s, to respective outputs for the force, skin acceleration, and multi-element stress rate sensors.

$$\text{Force Sensor:} \quad \frac{f_1}{x_s} = \frac{k_f k_1}{m_f s^2 + k_f + k_1} \tag{2.6}$$

$$\text{Skin Acceleration Sensor:} \quad \frac{x_a s^2}{x_s} = \frac{k_a s^2}{m_a s^2 + k_a} \tag{2.7}$$

$$\text{Multi-Element Stress Rate Sensor:} \quad \frac{f_s s}{x_s} = \frac{k_f^2 s}{m_f s^2 + k_1 + k_f} - k_f s \tag{2.8}$$

The actual parameters used in the analysis are specific to our sensors, but it provides some realistic numbers for comparison purposes. The force sensor stiffness, k_1, is 720,000 N/m, the finger tip foam stiffness, k_f, is 1100 N/m, and the effective local skin stiffness around the accelerometer, k_a, is 2000 N/m. The mass of the accelerometer, m_a, and the mass of the finger tip, m_f, are 0.0005 and 0.028 kg respectively.

Gains for the transfer functions of equations (2.6) through (2.8) are plotted in figure 2.6. Notice that in equation (2.6), the force sensor transfer function gain decreases with higher frequency due to the s^2 term in the denominator. This says that the force sensor sensitivity drops off with higher frequency at a rate of -12 dB per octave. In practice, the force sensor outputs are low pass filtered to reduce aliasing and resonance effects. The frequency response of the force sensor shown in figure 2.6(a) agrees with this observation. Therefore, the force sensor has higher sensitivity at the lower frequencies.

Equation (2.7) is the transfer function for the skin acceleration sensor. The s^2 term in the nominator says that the skin acceleration sensor DC response is zero, and it rises at a

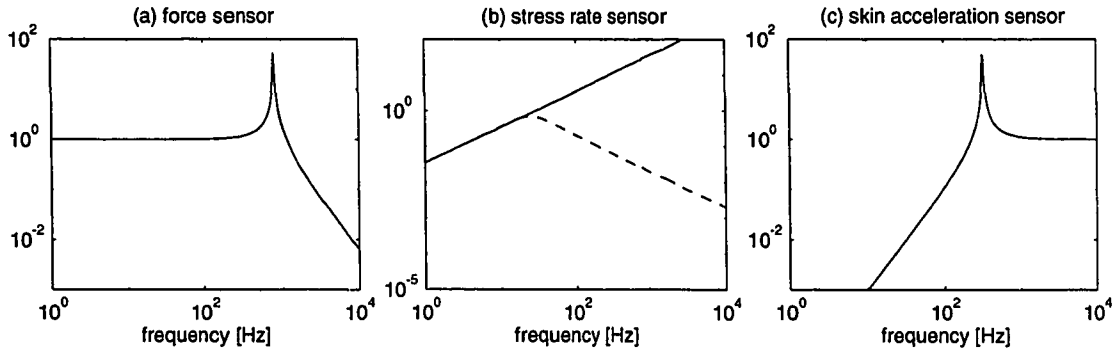

Figure 2.6 Normalized transfer function gains for the force, stress rate, and skin acceleration sensors for displacement at the finger tip skin as the input. Dotted line in (b) indicates response with 2 pole low pass filtering with 3 dB set to 30 Hz.

rate of +12 dB per octave. At higher frequencies, the s^2 term in the denominator level the frequency response. Therefore, the skin acceleration sensor has higher sensitivity at high frequencies, and figure 2.6(c) indicates this.

There are two components of the multi-element stress rate sensor transfer function (2.8). The first term has a resonance at 800 Hz, which rises and falls at a rate of ±6 dB per octave, but the dominant component is the second term, which is increasing with frequency at +6 dB per octave. As a result, the net transfer function shown in figure 2.6(b) rises at +6 dB per octave.

The current design of the multi-element stress rate sensor filters the high frequency signal with a two pole low pass filter with the 3 dB at 30 Hz. Therefore, the actual sensor gain response will have a maxima resulting in a triangular shape. The motivation behind low pass filtering the signal was to optimize incipient slip detection for servo control. Other signal processing could be used for optimizing contact detection such as rectifying the high frequency signal. It seems that the multi-element stress rate sensor falls between the force and skin acceleration sensors where the low frequency response is higher than the skin acceleration sensor and the high frequency response is better than the force sensor. Other advantages of the multi-element stress rate sensor will be evident in the following experiments.

2.3.2 Contact Experiments

To confirm the relative advantages of the sensors on a real manipulator, the multi-element stress rate sensor was mounted on one of the two-fingered planar robot hand shown in figure 2.7. Each finger has 2 degrees of freedom, and the workspace is oriented in the vertical plane. In contrast to most robot hands, this manipulator trades off large workspace and many degrees of freedom for a clean mechanical design with high bandwidth and good dynamic range. The manipulator uses brushless DC motors in a parallel link, direct drive configuration to eliminate backlash and reduce friction. A contactless magneto-resistive potentiometer is used for determining joint angle position, and the joint velocities are electronically generated from the position signal using a

differentiating amplifier. Further details of the manipulator design can be found in [Howe 1992].

First, the multi-element stress rate sensor's ability to detect contact events during manipulation was demonstrated in an autonomous robotic task. The manipulator was programmed such that the finger with the multi-element stress rate sensor approached the other finger at 2 cm/s and touched the other finger until a threshold voltage of 15 mV was detected. This threshold voltage was set based on sensor noise levels to provide reliable contact detection. Using the other finger as the "object" provided an independent verification of the actual contact force. After detection of contact, the finger velocity was reversed, and the whole process was repeated several times. The rectangular velocity profile resulted in peak accelerations of 7 N/m^2.

Figure 2.8 shows that the multi-element stress rate sensor was better able to detect contact in comparison to the force sensor. The actual contact force of 0.05 N shown in figure 2.8(b) is mostly a result of the response time required to reverse the manipulator direction. Even though the threshold value of 15 mV was used for the multi-element stress rate sensor, the peak recorded voltages are close to 100 mV because of the manipulator inertia which does not allow instantaneous changes in direction of motion. As shown in figure 2.8(c), the force sensor responded to both contact and rapid direction reversal. The mass of the finger tip acts as a proof mass thus forming an accelerometer with the force sensor. However, the multi-element stress rate sensor is insensitive to changes in acceleration

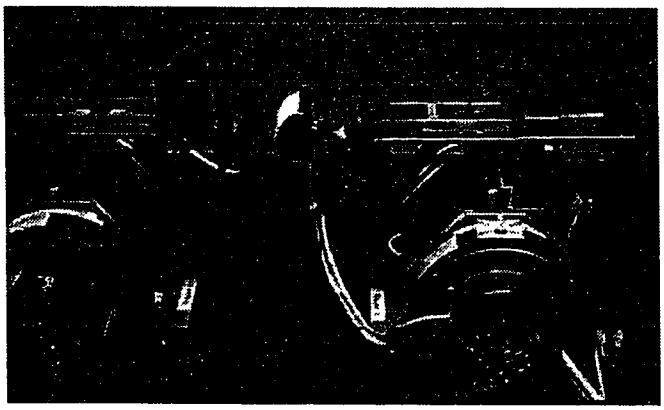

Figure 2.7 The multi-element stress rate sensor and the two axis force sensor is mounted on the right finger during this autonomous contact detection experiment.

caused by the manipulator. This is mainly due to the small effective mass of the finger tip skin which is distributed around the transducer along with a high damping coefficient of the rubber skin.

The skin acceleration sensor is good at detecting contact especially for fast approach velocities, but the large signal obtained at contact is a resonance phenomenon caused by the sensor mass bouncing on the rubber skin rather than the actual skin acceleration. This large signal is useful since it can be used to detect contact at moderate to high speeds. However, at low speeds, the excitation of the resonance is small and inconsistent even though a sensitive instrumentation accelerometer was used.

It should be noted that a great deal of effort has been put forth to make the planar hand smooth by using a direct drive configuration with debrushed motors and trapezoidal acceleration trajectories. Manipulators with higher transmission noise will reduce the

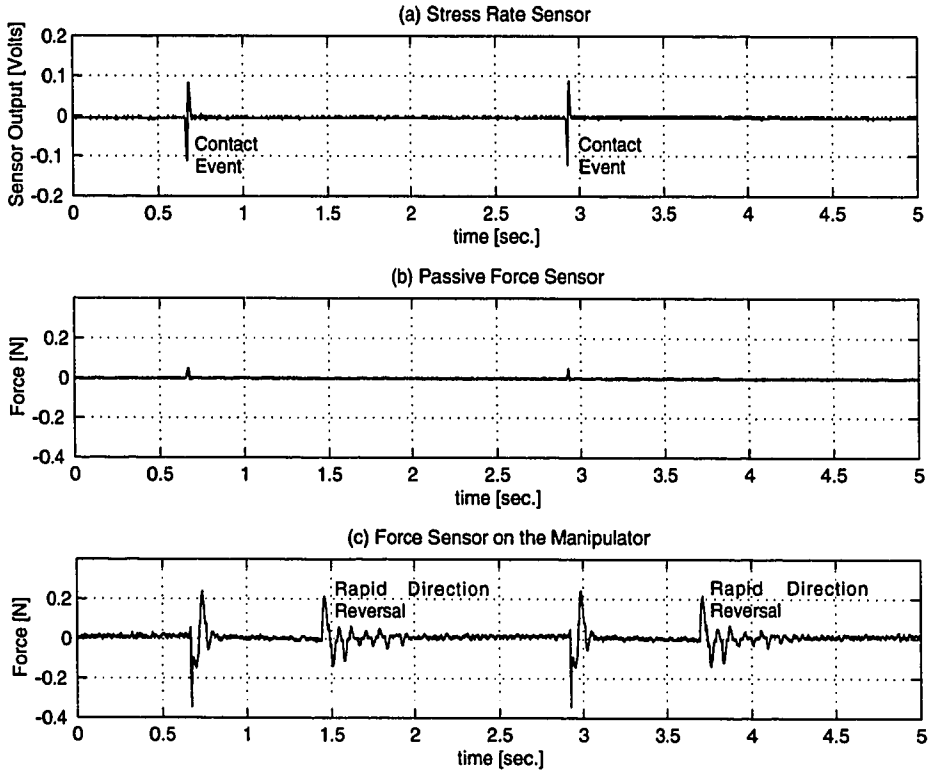

Figure 2.8 (a) Multi-element stress rate sensor responds only to contact events. (b) The actual contact force measured by a separate force sensor shows that the contact forces are less than 0.05 N which is due to the time delay associated with reversing the manipulator direction. (c) The force sensor on the manipulator responds to both contact and rapid accelerations.

performance the skin acceleration sensor since the accelerometer responds whenever the entire finger tip is accelerated. For the skin acceleration sensor, differentiation of contact events from accelerations can be difficult.

To compare sensor performance on a geared manipulator, the finger tip with these three sensors was mounted on an industrial manipulator (American Cimflex, Merlin) shown in figure 2.9. The manipulator brought the finger tip in contact with a fixed steel plate and the sensor outputs were recorded. Figure 2.10 shows the output from the force, skin acceleration, and multi-element stress rate sensors. High frequency resonance is evident in the force sensor and the skin acceleration sensor during free motion due to the vibrations generated by the gears. The signal component of the force sensor shown in figure 2.10(a) is only slightly greater than the high frequency transmission noise.

On the other hand, the skin acceleration sensor output, figure 2.10(b), shows much higher sensitivity to the high frequency transmission noise such that it exceeds the A/D input voltage range of 10 volts. When the finger tip is in contact with the rigid plate, the friction of the finger tip skin eliminate the resonant response. Consequently, the output amplitude of the skin acceleration sensor during contact is smaller than the free motion phase.

As shown in figure 2.10(c), one significant advantage of the multi-element stress rate sensor is its insensitivity to vibrations. The stress rate is related to the changes in contact force. However, the localized high stresses generated by the nibs upon contact provides

Figure 2.9 The multi-element stress rate sensor, skin acceleration sensor, and the two axis force sensor are mounted on this geared industrial manipulator.

higher sensitivity in comparison to the force sensor signal. Thus, even when the approach velocity is small, the stress rate can be large since it is also a function of the environment and finger tip properties.

One might argue that the reason why the multi-element stress rate sensor performs better is due to the low pass filtering. To test this idea, the skin acceleration sensor output was also filtered at 30 Hz with a two pole low pass filter, and the unfiltered multi-element stress rate sensor output was also recorded. The results are plotted in figure 2.11. The low pass filter dramatically reduced the high frequency noise, and the multi-element stress rate sensor does exhibit more high frequency noise in the unfiltered signal. However, even the unfiltered multi-element stress rate sensor signal has a better signal to noise ratio than the filtered skin acceleration signal. Also, vibration characteristics can change with gear ratios, joint configuration, speed, etc., so it will be difficult to build a filter which works for all situations.

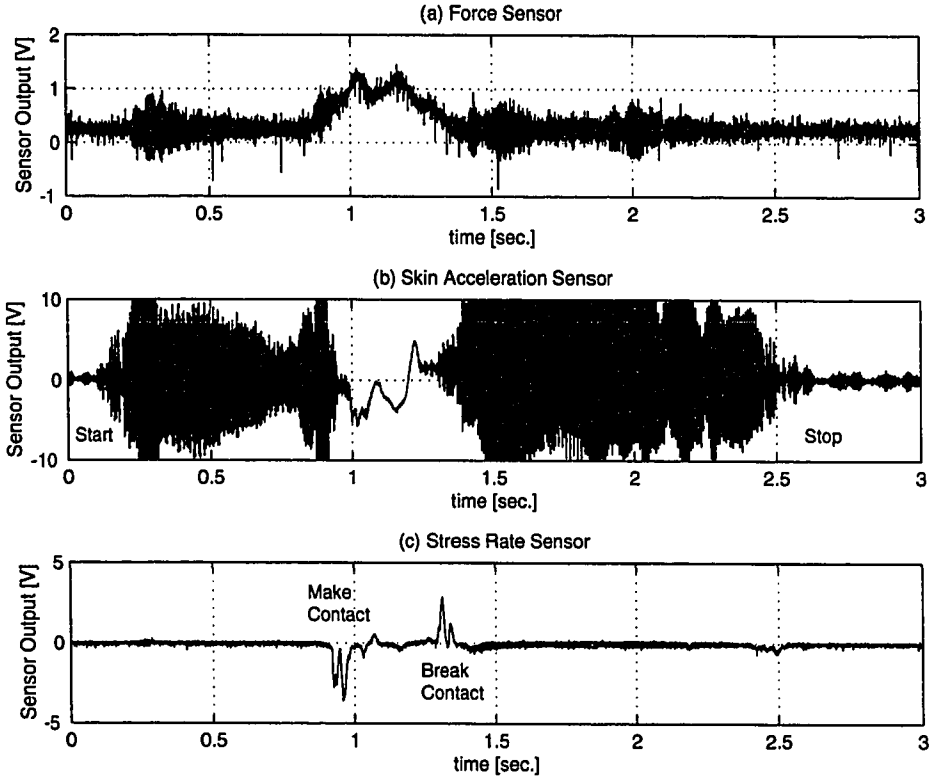

Figure 2.10 The manipulator brought the sensors into contact with a rigid plate and then reversed directions back to the starting position. (a) Force sensor output shows low signal to noise ratio and high frequency noise. (b) The skin acceleration sensor output is so sensitive to high frequency vibrations, it saturated the A/D input voltage range during free motion. (c) The multi-element stress rate sensor is insensitive to transmission noise.

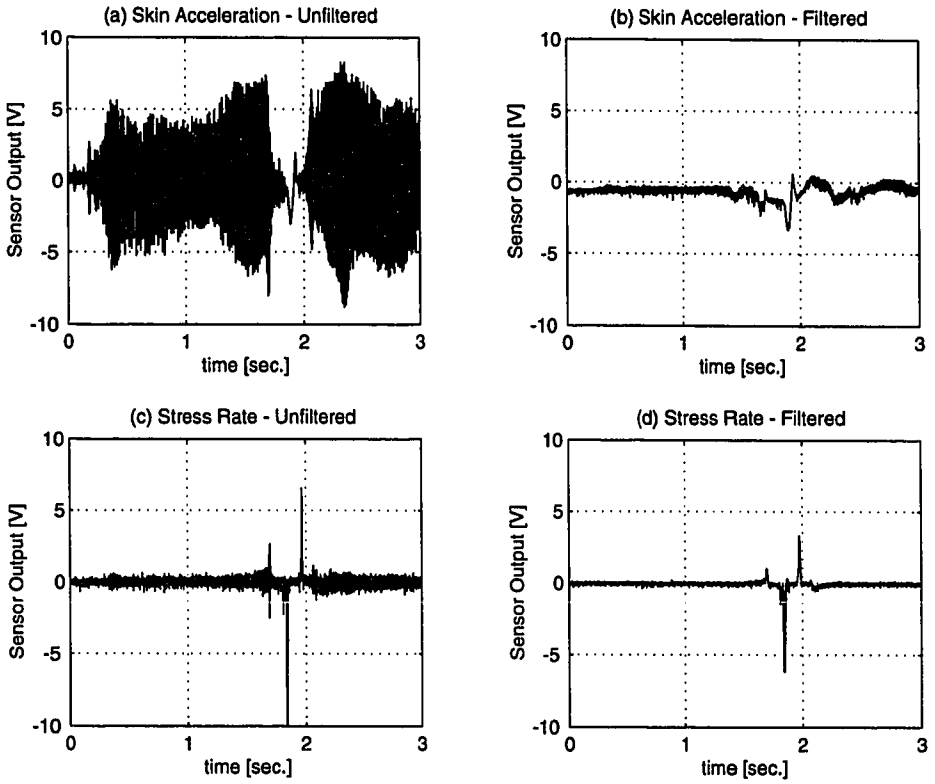

Figure 2.11 (a, b) Unfiltered and filtered skin acceleration response. (c, d) Unfiltered and filtered multi-element stress rate sensor response. The contact response signal to noise of the unfiltered stress rate sensor, (c), is still better than the filtered skin acceleration sensor, (b).

2.4 Incipient Slip Detection

Another important event that must be detected for skillful manipulation is incipient slip, which permits optimal setting of the grasp force between the finger tips. If the grasp force is too low, the object will slip between the fingers; if it is too high, the object may be damaged, particularly for fragile objects such as fruit. Humans show a remarkable ability to use a grasp force that is just slightly higher than the minimum to prevent slipping, even without a priori knowledge of object weight or frictional characteristics [Johansson and Westling 1984b]. This ability is apparently based on sensing small vibrations that indicate the earliest stage of slip before gross sliding of the object has commenced. This permits appropriate adjustment of the grasp force for a wide range of load forces, while minimizing the amount of energy exerted while grasping the object.

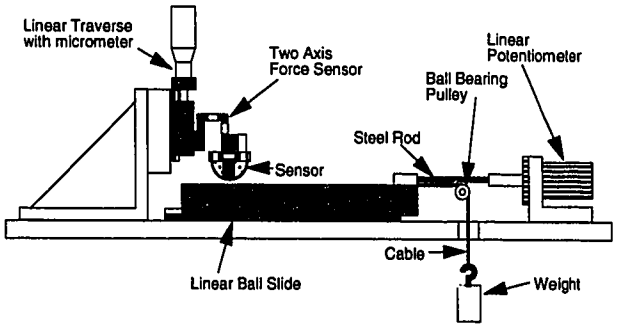

Figure 2.12 Bench top experimental apparatus for measuring the multi-element stress rate sensor's ability to detect incipient slip.

Several sensors have been specifically designed to measure vibrations that indicate incipient slip [Howe and Cutkosky 1989, Morrell 1990]. These sensors function by detecting the small vibrations that precede gross sliding. The mechanism responsible for these vibrations is not known, but is believed to be due to micro slips near the edges of the contact. For a curved finger tip the pressure is lower at the edge of the contact than at the center. When slip starts to occur, regions near the edges will give way before the center starts to slip. This motion around the periphery appears in the form of vibrations, and since the center is still fixed, overall motion of the grasped object does not occur. Since the multi-element stress rate sensor is able to detect localized vibrations, its performance in detecting incipient slip is better than the skin acceleration sensor.

To investigate the multi-element stress rate sensor to detect incipient slip, two bench-top tests were performed using the test apparatus shown in figure 2.12. The multi-element stress rate sensor was mounted on the end of a two-axis force sensor and the force sensor in turn was mounted on a micrometer-driven translation stage oriented in the vertical direction. A linear bearing slide simulated the contact surface of a grasped object below the sensor. The slide was connected to a linear potentiometer for measuring its position,

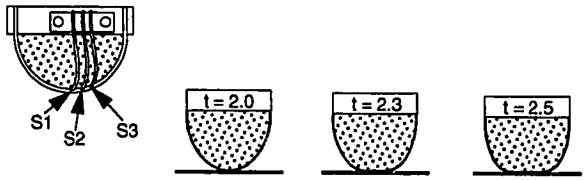

Figure 2.13 Skin curvature profile as the contact force increases. The area with the highest local curvature is at the edges of the contact area.

and a weight suspended by a string across a pulley provided the shear force required to make the contact surface slip beneath the sensor. The forces at the finger tip, the position of the linear slide, and the four elements of the multi-element stress rate sensor were measured at 3 kHz using a 12 bit A/D converter in a 486 PC.

The first test measured the sensor's ability to detect localized curvature changes. One cause of such change is the deformation of the finger tip skin as it is pressed against an object. Because the foam core of the sensor is highly compliant, a region of high curvature proceeds outwards from the center of the contact towards the periphery as shown in figure 2.13. The finger tip was lowered smoothly onto the contact surface. Figure 2.14(a) shows three piezo element responses corresponding to figure 2.13. The S1 trace is the piezo strip response of the center element, which makes contact first, and the S3 trace corresponds to the piezo strip element furthest away from the center. The sensors respond in succession as the local curvature changes near each piezo element along with an increase in normal force. This demonstrates that the sensor does in fact respond to localized

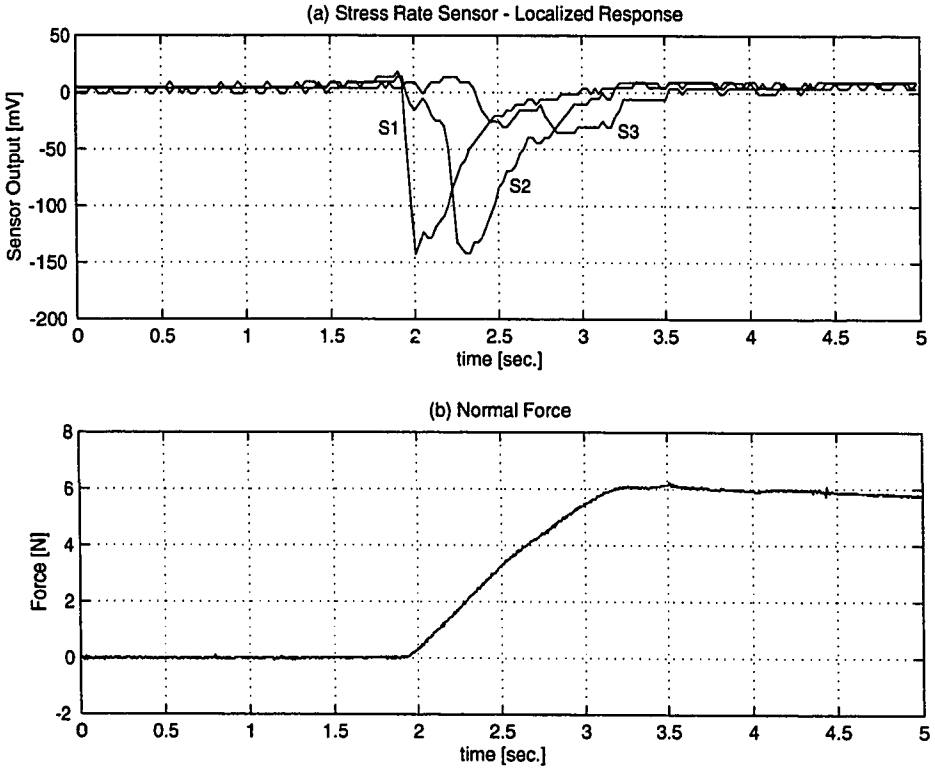

Figure 2.14 (a) The multi-element stress rate sensor response to changes in local skin curvature. (b) Normal contact force measured during the experiment.

curvature, and that it can effectively localize contact on the sensor surface, which is important for detecting incipient slip.

For the incipient slip detection experiment the linear slide was moved to the left and the finger was lowered against it to prevent it from sliding. Afterwards, the finger was slowly raised to reduce the contact force until the linear slide began slipping under the force of the hanging weight. Plot 2.15(a) shows the responses from the multi-element stress rate sensor, and plot 2.15(b) is the normal force and position of the slide. The two sensor signals S3 and S4 starts to change 95 and 47 ms earlier than any gross motion of the slide. This early warning provides sufficient time for a robot hand controller to increase the grasp force and prevent the onset of gross slip. The other two sensing elements do not signal the impending slip because they are not in contact with the slide. This spatial discrimination will permit the controller to robustly discriminate slip signals from other sources of vibration.

Initially, experiments were conducted with a hard rubber sensor over a hard material because of the larger signal amplitudes. However, closer examination revealed that the

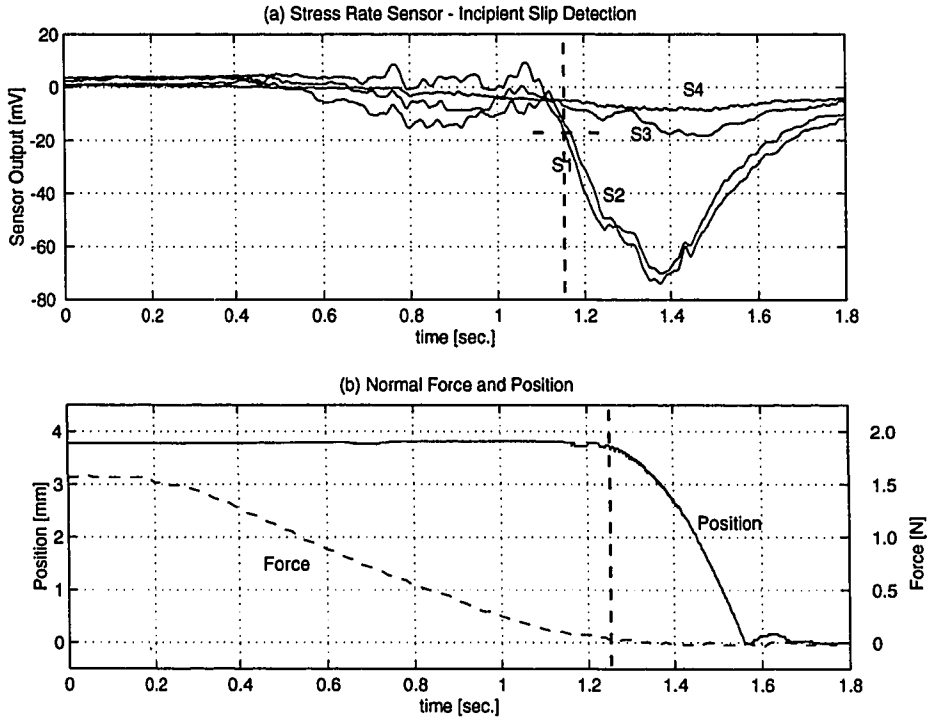

Figure 2.15 Detection of incipient slip. As the normal force is slowly decreased, (b), the multi-element stress rate sensor shows a slight response. The output from S1 and S2 indicate that slip is about to occur 95 and 47 ms before gross slip.

normal force and the incipient slip signal always started to change together. As it turns out, the sensor was measuring the normal stress rate as the skin curvature changed as the finger was released from the surface. By using a compliant foam material, the normal force changes were reduced so that it could be differentiated from the incipient slip signal as shown in figure 2.15(a).

Detecting slip is important in unstructured environments where the friction coefficient could change for reasons such as liquid surface contamination, or during assembly where task forces cause the part to slip within the gripper. Currently, the signals from incipient slip are much smaller than signals from contact events, so detection of slip during autonomous manipulation is difficult to implement. However, signals from the sensors were recorded during teleoperated manipulation to show the current performance of the sensor for detection of slip. The planar manipulator system also includes a kinematically identical "master" manipulator that can be used to direct the "slave" hand described above.

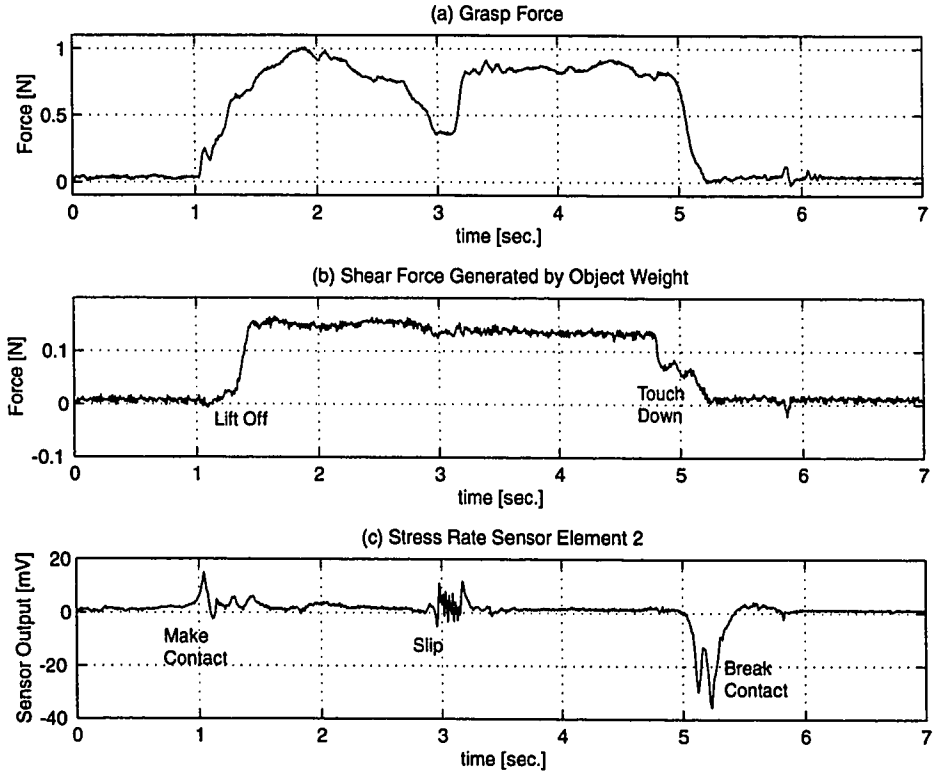

Figure 2.16 Force and multi-element stress rate sensor signals recorded during teleoperation. The task involved grasping and lifting the object, decreasing the grasp force to induce slip, regrasping and replacing the object.

This experiment consisted of picking up an object from the table top, gradually decreasing the grasp force until the object began to slip, followed by regrasping the object, then replacing it on the table. As shown by the force and the multi-element stress rate sensor signals in figure 2.16, the sensor is able to detect the onset of contact, incipient slip, and release of the object. By measuring the force signals just prior to slip, the friction coefficient between the sensor and the manipulated object can be determined and the appropriate grasp force calculated [Bicchi et al 1989].

2.5 Conclusions and Discussion

Detection of transient events in manipulator is important for adapting to the environment and minimizing disturbances to the object. As demonstrated above, the multi-element stress rate sensor is well suited for detecting contact and incipient slips because of the localized vibration sensing ability. The force and skin acceleration sensors both response to sudden accelerations and transmission vibrations found in robot manipulation, although analysis shows that the force sensor gain is larger at lower frequencies while the skin acceleration sensor gain is larger at higher frequencies.

The multi-element stress rate sensor fills a gap between the tactile array sensor and the skin acceleration sensor. Tactile array sensors provide highly localized but low frequency stress information. On the other hand, skin acceleration sensors are sensitive to high frequency information but have poor spatial localization ability. The experiments presented here demonstrate that the multi-element stress rate sensor bridges these capabilities by providing spatially resolved vibratory information.

One interesting aspect of this work is that human mechanoreceptors seem to evidence much the same "division of labor" between sensors. There are four types of specialized mechanoreceptor nerve endings in human finger tip skin [Johansson and Vallbo 1983]. Two of these sensors respond to static and slowly changing stimulus, and are believed to be responsible for sensing force and fine-scale pressure and shape information. These are perhaps analogous to the finger tip force and tactile array sensors. Another

mechanoreceptor, the FAII nerve ending, is most sensitive to vibratory stimulus in the 50-500 Hz range. The variation in response with frequency for these receptors suggests that they are sensitive to accelerations. These sensors have large receptive areas (as large as an entire finger surface in some cases), so they are clearly analogous to the robotic skin acceleration sensor. Finally, the FAI nerve endings respond to changing skin shape from roughly 5-60 Hz, and show the spatial selectivity that the multi-element stress rate sensor also possesses.

Chapter 3. Contact Localization

3.1 Introduction

One of the more important parameters to sense for dexterous manipulation is the location of the contact between the object and the robot finger tip. Theoretical work on grasping and manipulation shows that contact location affects many aspects of manipulation, particularly in determining grasp stability and in mapping between finger and object motions and forces [e.g., Kerr and Roth 1986, Salisbury 1985, Cutkosky and Kao 1989]. A number of studies have analyzed how contact locations change as objects pivot and slide against finger tips [Montana 1986, Kao and Cutkosky 1992, Cole et al. 1992, Howe and Cutkosky 1996]. However, many of the factors influencing contact location cannot be predicted in unstructured environments; these factors include detailed object geometry, external object forces, friction, and compliance. This implies that real-time sensing of the contact location is required.

Two methods of sensing contact location have received significant attention in robotics research literature: tactile array "extrinsic" sensing, and force-torque based "intrinsic" contact sensing. The first method uses a tactile array sensor to measure local pressure or displacement across the finger tip. Many transducer technologies have been proposed for this purpose; see Howe [1994] and Nicholls and Lee [1989]. Only a few studies have appeared on the use of array information in real time control. Notable examples include Berger and Khosla [1991] and Sikka et al. [1993], who used contact location information from an array sensor to track object features, and Maekawa et al. [1995], who used contact location sensing to prevent unwanted object rolling during manipulation.

These studies demonstrate that array sensors can localize contacts. However, the precision and generality of the algorithms for extracting location from the sensor data are far from clear. Such algorithms are often chosen ad hoc for a particular application; for example, Berger and Khosla used a simple peak-finding technique to determine the contact location. At the other extreme, Fearing [1990a] and Nicolson and Fearing [1995] have developed elaborate algorithms based on solid mechanics analysis to localize contact with known objects with remarkable accuracy (e.g., 0.03 mm). However, this technique is relatively slow, and experimental testing was based on data acquired under static laboratory conditions where noise levels are low and the mechanical perturbations of manipulation are absent.

One other method for determining contact location is based on a completely different sensing scheme. It uses a multi-axis force-torque sensor located near the finger tip. When the force and torque measurements are combined with known shape of the finger tip, the contact location can be calculated. (A detailed exposition of this technique is provided below.) This method was originally proposed by Salisbury [1984, 1985] and further developed by Bicchi et al. [1993]. The name "intrinsic contact sensing" (ICS) has been proposed to distinguish this approach from tactile array "extrinsic" sensing. ICS represents a fast and potentially advantageous method of determining contact location, particularly since force-torque sensors are useful for many other perceptual and control purposes. While these schemes have been implemented for robotic exploration of object shape and mechanical properties [Brock and Chiu 1985, Tsujimura and Yabuta 1989], the use of this sensing modality in realistic manipulation tasks has not been investigated. As with tactile array sensing, there is little understanding of nonideal effects on performance.

In this chapter, the abilities of these disparate sensing modalities to determine contact location during manipulation are compared. Simple signal processing algorithms that are easily executed in real time are chosen, and their performances are examined along with the effects of noise, vibrations, and other disturbances which are inevitably present in real task execution. The goal here is not to conduct a definitive study of the absolute accuracy

of these sensors, but to provide practical guidelines for immediate use in experiments. This will facilitate the use of tactile sensing and help to improve dexterous manipulation in real experiments.

First, a description of tactile array and intrinsic contact sensors are presented followed by an error source analysis for each sensing technique. Next, results from manipulation experiments, in which a robot finger tip containing both a tactile array and force-torque sensor, is presented. The manipulation tasks require changing contact forces and contact locations. This initial study restricts attention to planar tasks, although the results are immediately applicable to the three-dimensional case. Finally, the results and their implications for practical implementation are discussed.

3.2 Tactile Array Sensing

3.2.1 Description

Tactile array sensors measure the distribution of pressure across the finger tip. This signal can be processed to provide a great deal of information about the hand-object system. Among the parameters that can be extracted are contact location, object shape, and effective width of the pressure distribution (which determines resistance to rotational sliding [Bicchi et al. 1988, Howe et al. 1988, Howe and Cutkosky 1996]). This study focuses on the sensor's ability to localize contacts; several sensors have been developed that directly measure the effective center of the pressure distribution [Nakamura et al. 1986, Maekawa et al. 1995], but these sensors cannot be used to determine the other parameters listed above.

Tactile array sensors usually consist of a regular pattern of sensing elements. A vast number of tactile array sensing devices have been developed, although only a few have gone beyond bench-top testing [Howe 1994]. The experiments reported here use a capacitive tactile array sensor, based on an earlier design by Fearing [1990]. This sensor has the advantage of simple construction and well characterized, essentially linear response. Figure 3.1 is a photograph of the array sensor mounted on a cylindrical finger tip which

was used in for the experiments in this chapter. As shown in figure 3.2, the sensor is composed of two crossed layers of copper strips separated by strips of thin silicone rubber. As force is applied to the surface, the distance between the strips decreases, which causes the capacitance between the strips to increase. By measuring the capacitance at each crossing point, the spatial distribution of pressure across the sensor can be determined.

The sensor has 8 strips at 2 mm spacing in each direction, providing 64 force sensitive elements. The sensor forms a thin, compliant layer which can be easily attached to a variety of finger tip shapes and sizes. By encapsulating the sensor in a layer of elastomer, a range of surface compliance and sensitivities can be obtained. For these experiments, the 0.6 mm thick sensor is wrapped around a semi-cylindrical finger tip 25.4 mm in diameter, and covered by a 2.4 mm thick coating of silicone rubber with a modulus of 7 x 105 N/m2. In contrast to soft rubber, this relatively hard rubber coating minimizes creep and viscoelastic effects that can seriously degrade dynamic response. With averaging, the sensor elements were able to detect a point probe indentation of approximately 0.1 gram. Special-purpose electronics scan the array to measure the capacitance at each element, and the resulting pressure distribution signal for the entire array is sampled by a computer at 30 Hz.

The array sensor is calibrated by applying a uniform pressure to the surface with compressed air. First, the sensor output is measured at atmospheric pressure to obtain baseline voltages. The baseline signals for the elements are not uniform due to capacitance

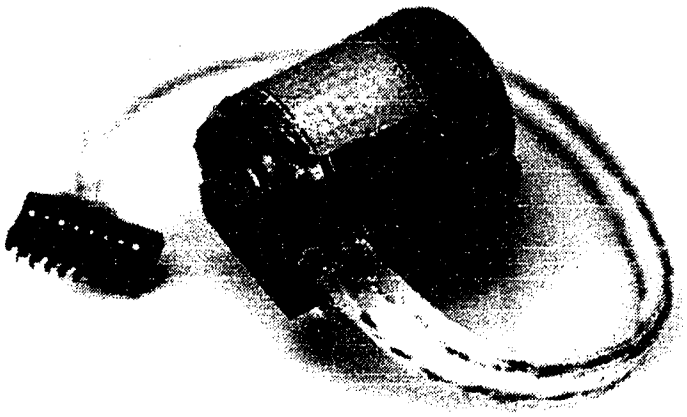

Figure 3.1 Capacitive tactile array sensor, an 8 x 8 matrix at 2 mm spacing.

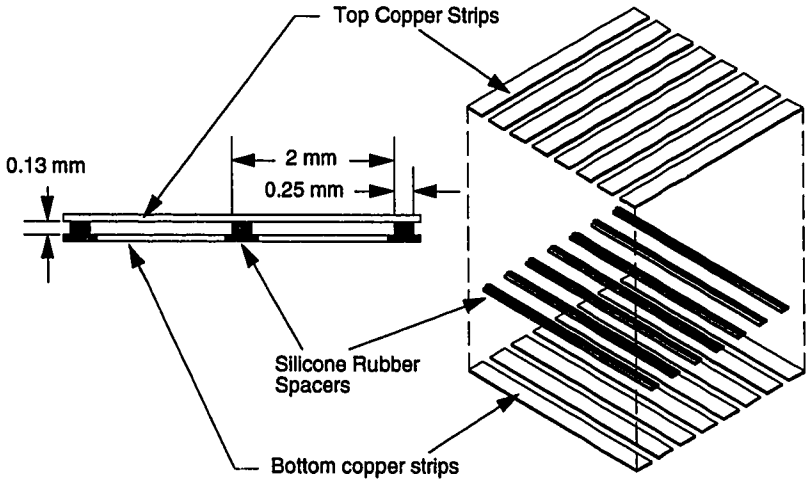

Figure 3.2 Tactile array sensor. Left: side view; contact occurs on upper surface, compressing silicone rubber and forcing top and bottom copper strips closer together. Right: exploded view showing sensor construction.

variations within the sensor and the cables. The sensor is then placed in a compression chamber and the pressure is increased. The array sensor output voltages are measured while the air pressure is measured with a separate pressure gauge. This calibration procedure is possible since the separation between the copper strips is mostly air and that cavity is completely sealed by the outer layer of rubber. Thus, the gain of the ij th element is

$$G_{ij} = \frac{P}{(V_{ij}^{P} - V_{ij}^{baseline})} \tag{3.1}$$

where $V_{ij}^{baseline}$ is the baseline voltage and V_{ij}^{P} is the voltage measured at applied pressure P. Several different pressure settings were used to verify that the gain matrix multiplied by the output voltages resulted in a uniform response. The calibration scheme neglects the compression of air inside the sensor which reduce the sensor output with increasing pressure. However, absolute pressure gains were less important than obtaining a gain matrix which compensates for the differences between the elements.

3.2.2 Contact Localization Algorithms

The goal here is to examine this sensor's ability to quickly determine the contact location of an object. For these experiments, three simple algorithms that are easily implemented and executed in real time are considered. Therefore, centroid and peak detection schemes were selected as shown in figure 3.3. The assumption here is that only one contact area is formed on the cylindrical finger tip. The first scheme is a straightforward method of describing the contact location by the center of pressure, or centroid, as described by

$$x_{centroid} = \frac{\sum_{i=1}^{m}\sum_{j=1}^{n} x_{ij}\delta p_{ij}}{\sum_{i=1}^{m}\sum_{j=1}^{n} \delta p_{ij}} \quad \text{and} \quad y_{centroid} = \frac{\sum_{i=1}^{m}\sum_{j=1}^{n} y_{ij}\delta p_{ij}}{\sum_{i=1}^{m}\sum_{j=1}^{n} \delta p_{ij}}. \tag{3.2}$$

For an $m \times n$ array, the ij th tactile element position is given by x_{ij} and y_{ij}, and the pressure at that element is given by δp_{ij}. The second localization scheme aims to lessen sensitivity to low level noise on elements that are not near the contact area by calculating the centroid with a threshold. In this case, only those elements whose values were larger than 50% of the maximum element value, were included in the summation of (3.2).

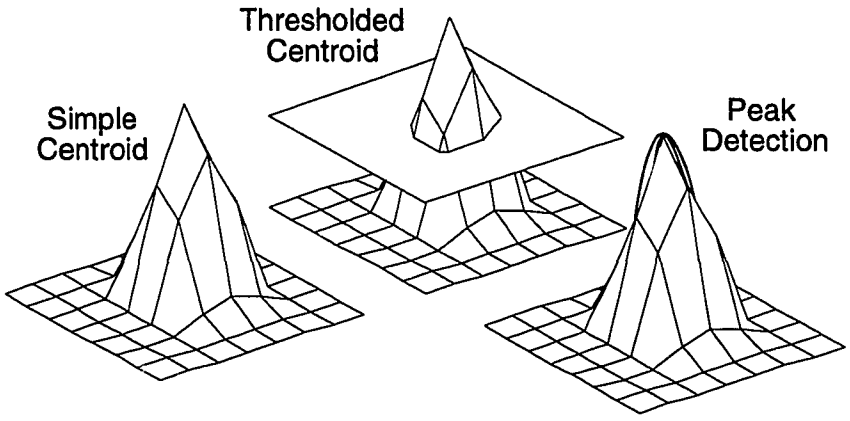

Figure 3.3 Three algotithms for determining contact location for the array sensor. The simple centroid scheme determined the pressure centroid using all the elements of the array sensor. The thresholded centroid algorithm only used tactile elements whose value was larger than 50% of the maximum tactile value to calculate the centroid. The peak detection algorithm used a quadratic fit about the maximum pressure element for each direction along the elements.

Finally, the third localization scheme interpolates the peak strain location. For point contacts, Fearing [1990a] used a sinusoidal curve fitted to the entire array and a gradient search method for determining peak strains, and obtained localization accuracy of 0.01 times the tactile element spacing. To facilitate real time computation in these experiments, a quadratic curve was fitted to three points around the tactile element with the maximum strain. From the quadratic equation, the peak strain can be interpolated with much greater precision than the tactile element spacing.

Analysis suggests that the constraining limit in determining contact location may be the presence of shear forces at contact. Previous work has shown that the response of this type of sensor is predicted with reasonable accuracy by a solid mechanics model [Fearing 1990, Howe and Cutkosky 1993]. Here the finger tip is treated as a linear elastic half space, and each sensor element output provides a normal strain measurement in the vertical direction at the appropriate surface location and depth. This analysis suggests that all three of the localization schemes will be adversely affected by shear forces (i.e., the component of the contact force applied tangential to the surface of the sensor). Although it is possible to separate the even and odd components of the sensor response to better localize contact, this scheme requires that the object shape at the point of contact be symmetric

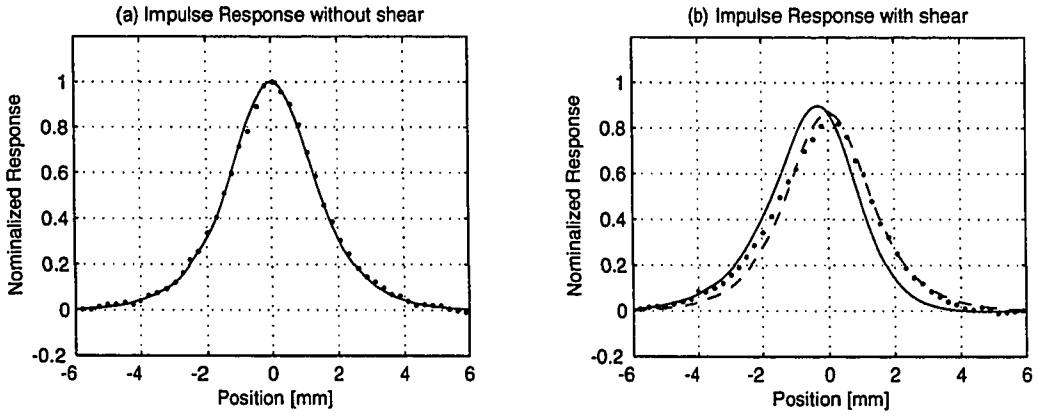

Figure 3.4 Measured and theoretical impulse response of the tactile array sensor. The skin thickness and Poisson's ratio used were 2.4 mm and 0.12, respectively. The solid lines correspond to the expected strain based on a linear elastic model and the points are measured data points. In (b), forces were applied at 30 degrees from the surface normal, and the dashed line corresponds to the predicted strain generated by the normal component of the force. The close agreement with the normal component response shows that the sensor is relatively insensitive to shear forces.

with respect to the array sensor surface normal [Fearing 1990], so the separation algorithm cannot be applied to all object shapes.

However, measurement of the sensor response to shear forces shows that shear sensitivity is less than predicted by the model (Figure 3.4). When the force vector is purely normal to the sensor surface, there is good agreement between the model and the measured response. When the force is inclined at 30 degrees, the model predicts that the peak strain moves about 1 mm in the direction of the shear force, whereas the measured response remains centered at essentially the same location. This effect is more pronounced when the shear force is acting along the top copper strips rather than perpendicular to it.

This discrepancy is due to the sensor construction, which is more appropriately modeled by a thin beam comprised of copper strips and a relatively hard rubber outer layer supported by strips of RTV in between the tactile elements. The air gap between the copper strips decreases the effective Poisson's ratio, and the copper strips distribute shear forces further away from the point of contact. The ramification of the sensor construction is that small but significant measurements are found many cells away from the contact location. Since many tactile array sensors have continuous strips of relatively rigid materials joining the elements, other types of sensors will share this characteristics. Consequently, using a threshold is desirable when computing contact centroids. However, minimal shear sensitivity is a benefit for localizing contact, since the errors generated by shear forces are significantly reduced.

The three localized schemes represent specific points along a continuum. The centroid scheme takes the whole array into consideration, while the peak localization scheme focuses on just a few elements. The threshold centroid scheme is a compromise where only certain elements meeting the criteria are taken into consideration. The peak localization scheme will be less stable than the centroid methods. For example, without contact the peak localization scheme can jump all over the sensor area, but the averaging effect for the centroid scheme provides a more stable output. Both of these outputs are incorrect, so the algorithms require total contact force thresholding.

3.3 Intrinsic Contact Sensing

3.3.1 Analysis

Intrinsic contact sensing (ICS) combines finger tip force-torque sensing with a model of the finger tip shape to estimate contact location. The principle of operation is illustrated in figure 3.5a. The ICS algorithm assumes that the object contacts the finger tip at a single point with no torques transmitted at the contact. Then the torque measured by the sensor is caused exclusively by the measured force at the contact location. The combination of forces and torque determine a line of action, and the intersection of this line of action with the finger tip surface is the point at which the contact occurs.

Although the ICS algorithm as developed by Salisbury [1984] and Bicchi et al. [1993] is fully three dimensional, these experiments are restricted to the planar case. Since the manipulator operates in the plane and the cylindrical finger tip is oriented to form a circular cross section in the plane, the origin of the coordinate frame was located at the center of the finger tip (Figure 3.5b). The force-torque relation is

$$\mathbf{m} = \mathbf{r} \times \mathbf{f} \tag{3.3}$$

where $\mathbf{m}$ is the three-element measured torque vector, $\mathbf{f}$ is the three-element force vector, and $\mathbf{r}$ is the three-element vector from the origin to the point of contact. In the planar case there are only two force components and one torque component, so this reduces to

$$m = f_y x - f_x y \tag{3.4}$$

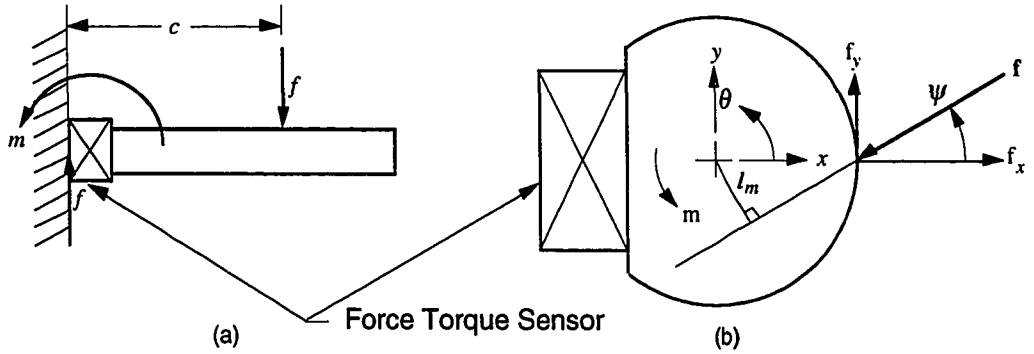

Figure 3.5 (a) Simplest example of ICS for a straight line "fingertip" with perpendicular force f. The contact location c = m/f. (b) Finger tip cross section, coordinates and variables used in the ICS calculations for the experiments.

where m is the moment perpendicular to the plane, f_y, f_x are the planar force components, and x, y are the contact coordinates in the plane. The finger tip shape is described by

$$r^2 = x^2 + y^2 \tag{3.5}$$

where r is the radius of the finger tip. Because the line of action will intersect the finger tip at two points, one must specify that $x > 0$ to ensure that the correct point is found. In general, $\mathbf{r} \cdot \mathbf{f} < 0$ which requires the force vector to point into the fingertip.

Polar coordinates were used for the experiment since this permits specification of the contact location in terms of one parameter, θ, as r is constant. The distance along the finger circumference, measured from the x-axis, is then simply $r\theta$.

The force vector orientation can be determined easily by the ratio of f_y over f_x. By extending the force vector line of action, the lever arm for the moment about the origin, l_m, can be determined. Thus,

$$m = f\, l_m = f\, r \sin(\psi - \theta) \tag{3.6}$$

where $f = \sqrt{f_x^2 + f_y^2}$ is the force magnitude and ψ is the force vector angle. Solving for θ we obtain

$$\theta = \tan^{-1}\left(\frac{f_y}{f_x}\right) - \sin^{-1}\left(\frac{m}{fr}\right). \tag{3.7}$$

This equation imposes a condition that $\left|\frac{m}{f\,r}\right| < 1$. This requirement simply limits the maximum moment allowed for a given force magnitue by the fingertip radius, which means that the force sensor values must be checked before applying this ICS algorithm.

To understand the limits of this scheme, the theoretical limits to localization accuracy is examined. Using equation 3.7, one can take the partial derivative with respect to each parameter and multiply it by the parameter estimation error to obtain an estimate of the noise sensitivity. Since we have two forces, one torque, and a radius in the estimation formula, the maximum estimation error is

$$\delta\theta_{\max} = f(f_x, f_y, m, r) = \left[\left(\frac{\partial\theta}{\partial f_x}\Delta f_x\right)^2 + \left(\frac{\partial\theta}{\partial f_y}\Delta f_y\right)^2 + \left(\frac{\partial\theta}{\partial m}\Delta m\right)^2 + \left(\frac{\partial\theta}{\partial r}\Delta r\right)^2\right]^{1/2}$$

$$= \left[\left(\frac{mf_x\Delta f_x}{rf^3\sqrt{1-\left(\frac{m}{fr}\right)^2}} - \frac{f_y\Delta f_y}{f^2}\right)^2 + \left(\frac{mf_y\Delta f_y}{rf^3\sqrt{1-\left(\frac{m}{fr}\right)^2}} + \frac{f_x\Delta f_y}{f^2}\right)^2 + \frac{\Delta m^2}{fr(fr-m)} + \frac{m^2\Delta r^2}{f^2r^4\left(1-\frac{m}{fr}\right)}\right]^{1/2} \quad (3.8)$$

Figure 3.6 shows the estimation error along the cylindrical surface as a function of each parameter for the fingertip used in the experiments described below. The parameters are varied plus and minus twenty percent of the nominal values for the contact condition shown in figure 3.5b, for the case of contact location $\theta = 0$ and force orientation $\psi = 30°$. Within this range, contact location varies almost linearly with small errors in f_x, f_y, m and r. Similar results are obtained for other contact locations and force orientations. The predicted total errors are on the order of a few millimeters, which correlates well with the experimental results presented below.

The foregoing analysis treats the contact as a single point with no transmitted torques. For compliant surface contacts, it can be shown that intrinsic contact sensing determines the contact centroid or center of pressure as an estimate of the contact location [Bicchi et al. 1993].

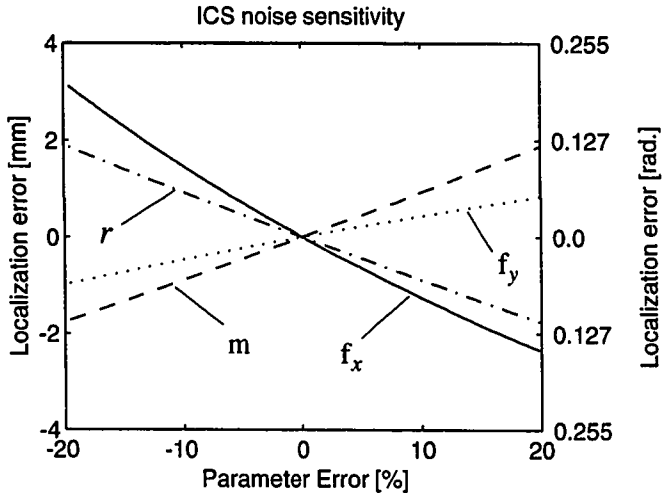

Figure 3.6 Predicted localization errors based on variations in forces, moment and the finger tip radius. Within ± 20% variation, the localization errors are approximately linear.

The measured force acting on a contact centroid is equivalent to the distributed pressure over the contact area. This analysis represents the simplest of three algorithms presented by Bicchi et al. [1993]. The other methods attempt to compensate for the use of compliant finger tips in three dimensions, where contact torques about the surface normal may be generated. In the planar experiments presented here, these contact torques are not pertinent.

3.3.2 Calibration

Multi-axis force-torque sensors typically use a set of coupled single-component strain gauges to make the force and torque measurements. The relationship between the measured strains and the actual forces and torques can be expressed as a linear transformation

$$F = C\,V \tag{3.9}$$

where $F = [f_x f_y m]^T$ and V is the measured strain voltages. The C matrix must be obtained through calibration, and this procedure requires applying controlled forces on multiple locations on the finger tip surface at various orientations. By recording the strain gauge voltages and calculating the moment generated at the finger tip center by the applied forces, a least squares fit calibration matrix can be obtained.

In these experiments, the three-axis sensor is a U-shaped aluminum structure with semiconductor strain gauges mounted at three locations. Despite care in the calibration process, the calibration matrix was not very well-conditioned, so small errors in the measured voltages could produce significant errors in the calculated force components. A major source of this error was due to the compliant finger tips. Because a pointed indenter was used to apply desired forces on the finger tip, the accuracy of the calculated moment was reduced by the uncertainty in contact location due to skin deformation. The finger tip deformation also generated unmodeled shear forces which reduced the accuracy of the applied force. Bicchi and Canepa [1993] provides a thorough discussion of the calibration process and the issues associated with it. The calibration matrix used in these experiments is

$$C = \begin{bmatrix} -126 & -51 & 127 \\ 378 & -6 & -72 \\ 396 & -128 & 22 \end{bmatrix}.$$

Because the force sensor was designed for planar applications with forces only in the manipulation plane, it showed slight cross-sensitivity to out-of-plane forces. The experiments were conducted in the horizontal plane, so the weight of the grasped object resulted in a constant offset of a few millimeters in the ICS localization results. This offset was subtracted from the results presented in the following section.

3.4 Experiments

3.4.1 Experimental Apparatus and Procedures

A simple two-fingered planar hand shown in figure 3.7 was used to conduct these manipulation experiments. As seen in figure 3.7 and 3.8, each finger has direct drive motors and a five bar linkage, providing two degrees of freedom in the horizontal plane. The servo rate was set to 300 Hz and object stiffness control was used to command the desired

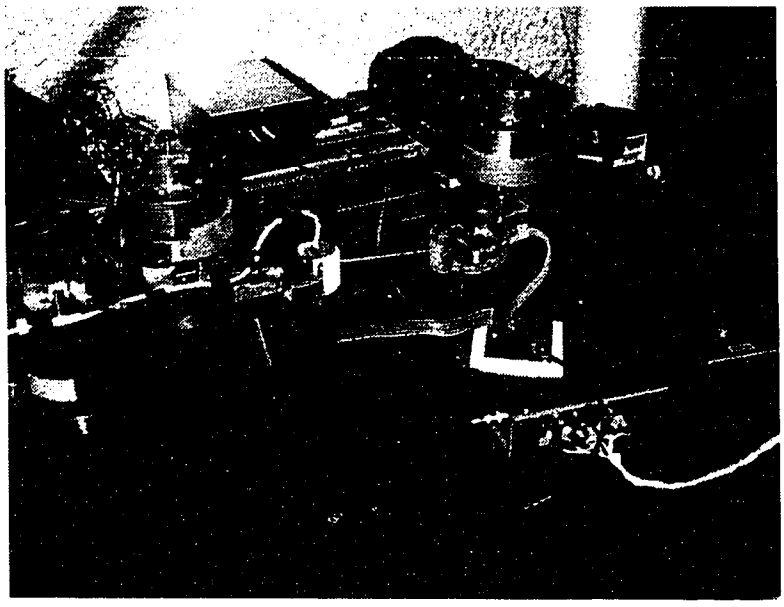

Figure 3.7 Planar manipulator at the Stanford University Dextrous Manipulation Lab.

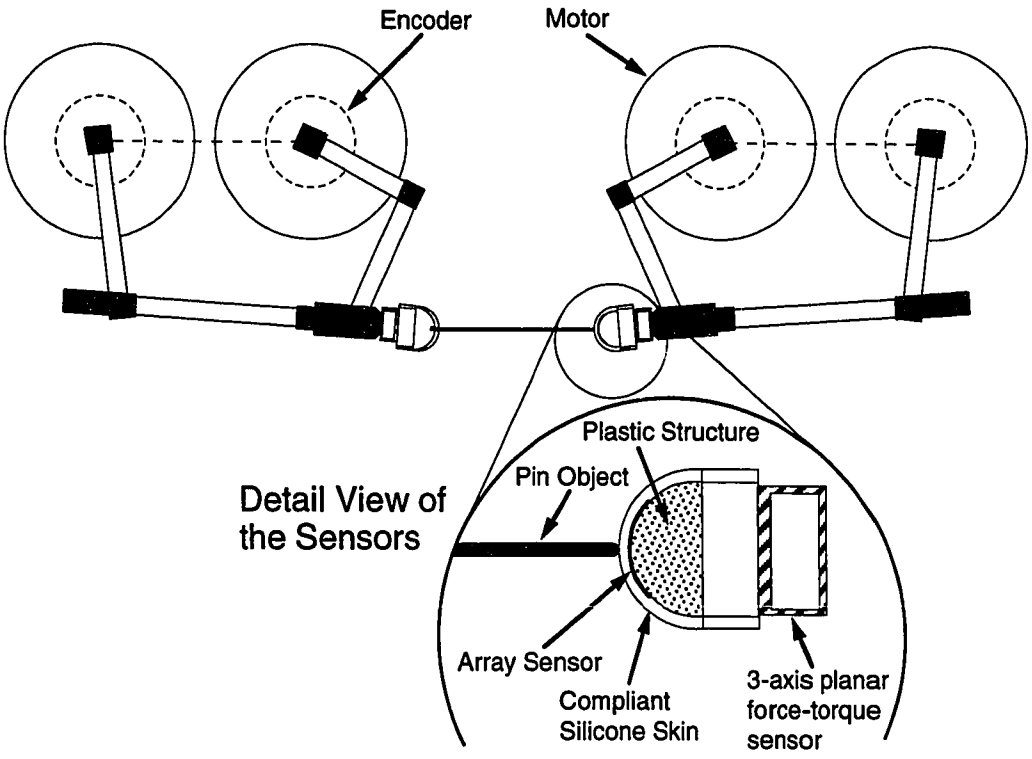

Figure 3.8 Schematic of the manipulator during the pin manipulation experiment. A tactile array sensor and a force sensor are mounted on the right finger tip.

object location and orientation within the workspace. Both finger tips were mounted on the three-axis force sensors described above, and the finger tip with the tactile array sensor weighed 35 grams, with an outer radius of 15.7 mm.

To investigate contact localization during manipulation, two objects and two manipulation tasks were chosen. A 10 cm long pin with spherical ends and a 1.6 mm radius tip, and a 10 cm wide square box were used as the objects. The pin and the box weighed 4.2 and 77 grams respectively. The two manipulation tasks involved rotating the object ±20 degrees about the center of the object, and moving the object in a 2 cm by 2 cm square trajectory at a relatively high velocity of 4 cm/second.

These particular objects and tasks were selected to include many of the conditions that might be encountered in general dexterous manipulation tasks. In particular, the tasks generate significant shear forces and accelerations. When rotating the pin, the force angle with respect to the finger changes. This generates shear forces at the contact location which

may reduce the localization ability of the array sensor. On the other hand, intrinsic contact sensing measures shear forces and uses this information to localize contact, so variations in shear forces are taken into consideration. Since the object is a pin and the contact location remains essentially constant, the actual localization plot as a function of time should be a horizontal line. To measure the ability to track changes in contact location, the pin was replaced with a box having the same dimensions in length. As the object rotates, the finger tip rolls on the object and the expected trace for contact localization is approximately a triangular wave form as a function of time.

The next task of translating the object was designed to test the ICS scheme. At the corners of the square trajectory, the grasped object and the finger tips experience high accelerations, so the mass of the object and the finger tips will generate inertial forces to which the force sensor will respond. These forces will appear as interference on the force sensor signal, but the tactile array sensor, with its low mass and distributed sensing elements, should not respond. The purposes of these experiments were to observe the effects of these inertial forces on the ICS algorithm and to verify that the array sensor is immune to them. When translating the box, the finger tip must roll slightly against the box to obtain the desired trajectory, while the expected localization trace for the pin is approximately constant as a function of time.

3.4.2 Initial Contact Results

In figures 3.9 through 3.14, typical localization results are plotted for intrinsic contact sensing and the three localization schemes for the array sensor. For all the array sensor plots, the heavy solid dots are the peak localization algorithm, the dashed lines are the simple centroid algorithm, and the dotted lines are the thresholded centroid algorithm. The coordinate origin for contact localization is set at the location corresponding to $\theta = 0$ and the distance is along the circumference of the finger tip.

First, the tactile array and ICS localization algorithms are compared during the contact transition. Manipulator fingers were positioned close to the object and the controller commanded a 3 N feed forward grasp force to grasp the object, and data acquisition started

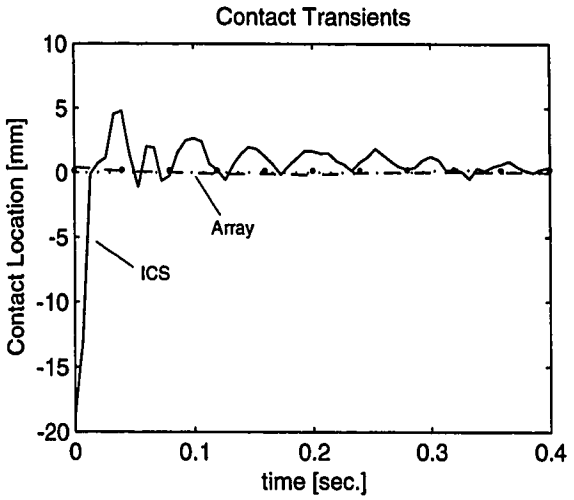

Figure 3.9 Comparison of transient responses of ICS and array sensing schemes. Large negative spike in ICS output is due to low force levels at initial contact.

immediately. Initial impact with the object caused a force spike of approximately 1 N in magnitude. This, along with the inertial forces before contact, explains the large localization error associated with the intrinsic contact sensing scheme shown in figure 3.9. The array sensor is insensitive to this contact transient.

3.4.3 Pin Manipulation Results

Results from translating the pin in a square trajectory are shown in figures 3.10 and 3.11. In plots 3.10a and 3.10b, the pin is located at the center of the finger and shear forces were negligible, while for plots 3.10c and 3.10d, the contact location is shifted by approximately 3 mm and shear forces with respect to the surface normal were about 19% of normal forces. Since the object velocity was 4 cm/sec and the total length of the trajectory was 8 cm, the trajectory was completed in 2 seconds, but it was repeated several times. The measured peak accelerations at the corners of the trajectory were between 80 and 130 cm/sec^2.

For plots 3.10a and 3.10b the actual trajectory of the contact location is approximately a horizontal line at position 0 mm, and for plots 3.10c and 3.10d the actual trajectory is a horizontal line at position -3 mm. The random noise associated with ICS is larger in comparison to the tactile array sensor. There are two reasons for this. First, the random noise level of the force sensor was higher than the array sensor. Second, the finger tip mass

acts along with the force sensor to mechanically amplify any vibrations and manipulator actuation noise.

Although the contact location does not change against the finger tip, localization spikes are seen every 0.5 seconds in plots 3.10b and 3.10d where the trajectory goes through a corner and inertial forces are generated. A clear correlation between object accelerations and localization error for the ICS is shown by the periodicity of the trace corresponding to object accelerations.

Plot 3.10a shows that the array sensor is immune to accelerations, but plot 3.10c shows that when the contact location is shifted approximately 3 mm, shear forces influence the centroid algorithm. Since the simple centroid algorithm uses all the tactile elements of the

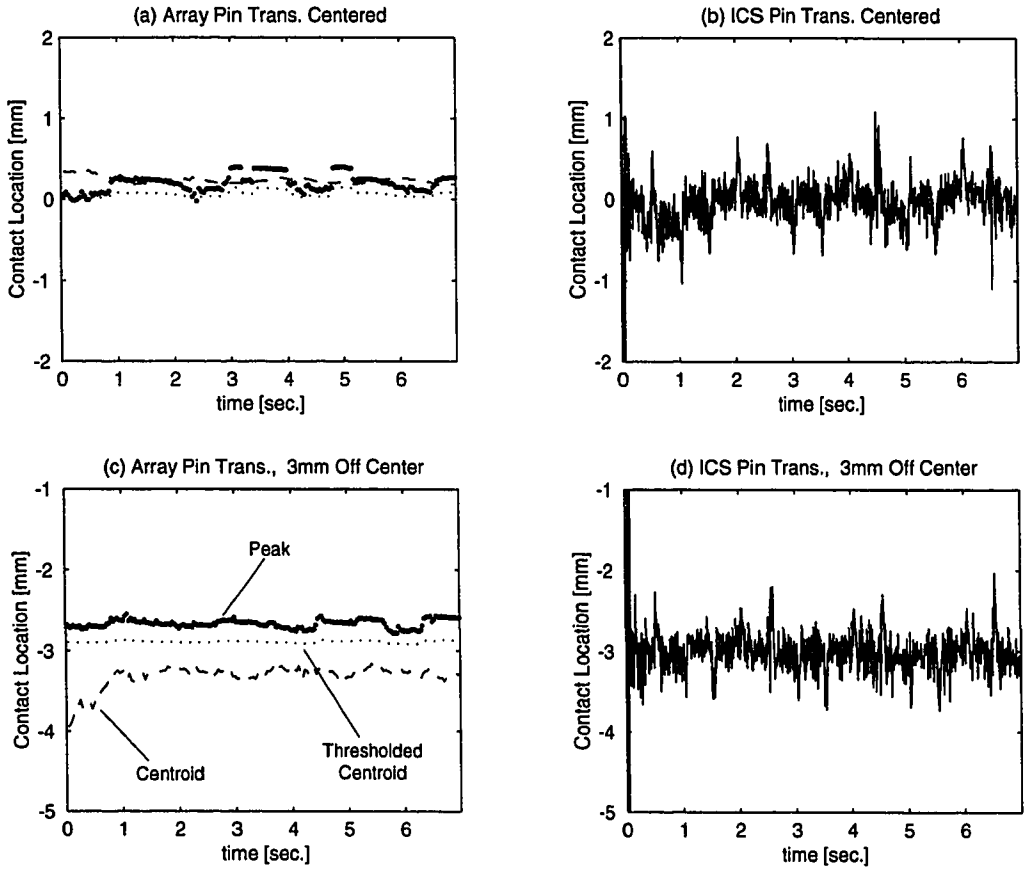

Figure 3.10 Contact localization for pin translation experiment. (a), (c) Array sensor; (b), (d) ICS. (a), (b) pin centered on finger tip ($\theta = 0$, no shear forces); (c), (d) pin displaced 3 mm ($\theta = 0.19$ rad, shear force present).

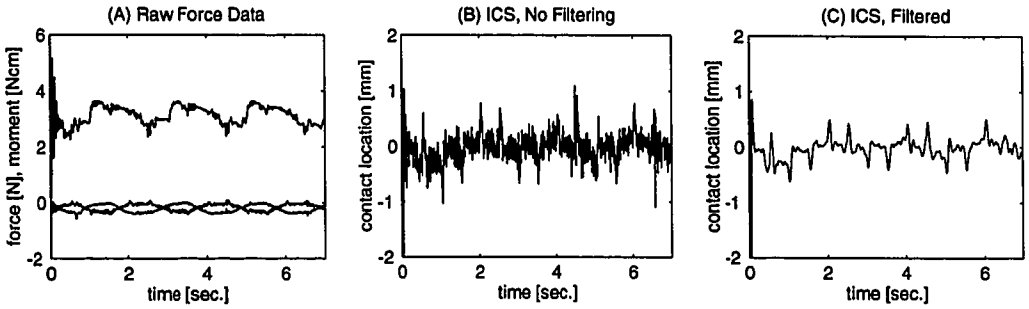

Figure 3.11 (A) Forces and moment measured by the force sensor. (B) Contact location extracted by the ICS algorithm. (C) Filtered contact location using a two-pole low pass Butterworth filter with the corner frequency set to 10 Hz.

sensor, shear force generated deflections far from the contact location are included in the calculations, thus biasing the results.

The force sensor signal is shown in plot 3.11a, which reveals the level of sensor noise present in the data for the pin centered translation experiment. Plot 3.11b is the contact location generated using the force data shown in 3.11a. Plot 3.11c shows the results when contact location data was filtered using a two-pole low pass filter with the 3 dB set to 10 Hz. The random noise associated with ICS shown in figure 3.11b is eliminated, but the accelerational transients remain essentially unchanged.

Results from the pin rotation experiments are shown in plots 3.12a through 3.12d. In plots 3.12a and 3.12b, the pin is located at the center of the finger while for plots 3.12c and 3.12d, the location is shifted again by 3 mm. While holding a pin and rotating it, the contact location does not change against the finger tip, but the force vector changes along with the rotation angle. Since the contact location is fixed, the actual trajectory for these plots is also a horizontal line. Shear forces generated during the pin rotation experiment reduces the tactile array sensor's ability to localize contact since shear forces create normal strains. However, this localization error is only pronounced for the simple centroid scheme at very sharp angles as shown in plot 3.12c.

Although the noise associated with the ICS signal hasn't increased, the localization trajectory is significantly different from the expected and the array sensor results. Theoretically, the ICS scheme should be insensitive to shear forces since it uses this

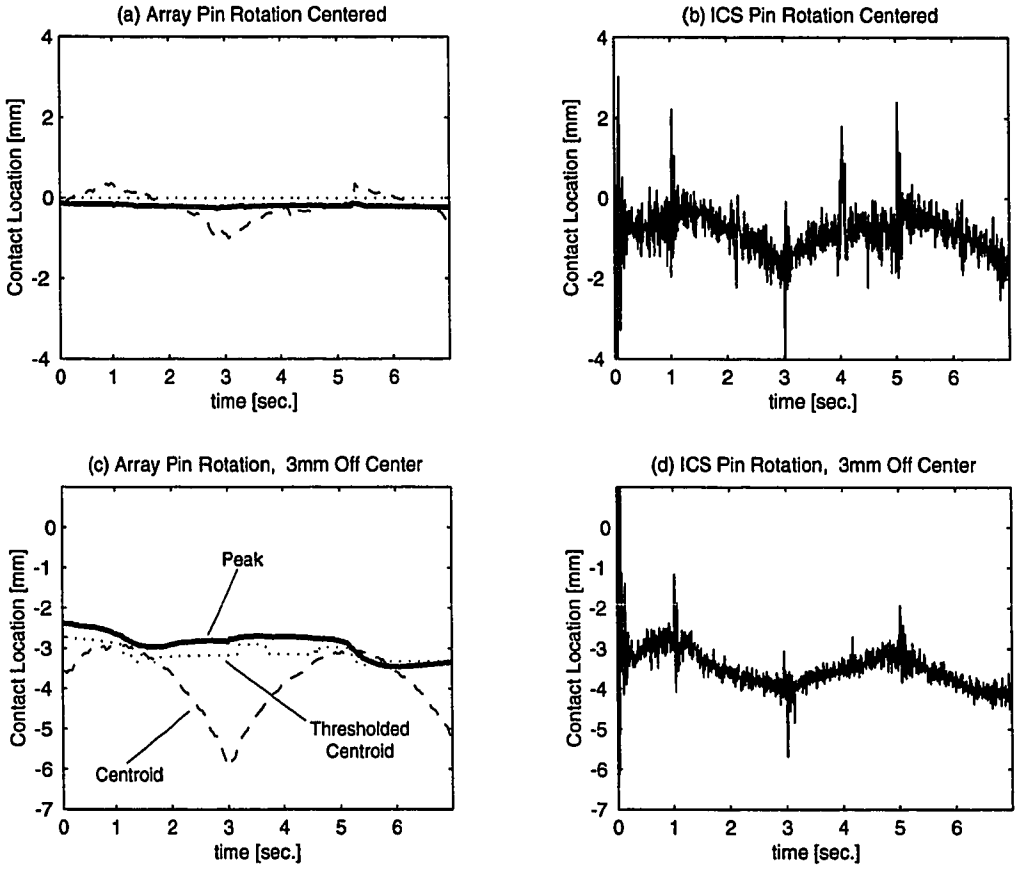

Figure 3.12 Contact localization for pin rotation experiment. (a), (c) Array sensor; (b), (d) ICS. (a), (b) pin centered on finger tip (θ = 0, no shear forces); (c), (d) pin displaced 3 mm (θ = 0.19 rad, shear force present).

information, but calibration difficulties result in errors, as seen by the "triangular wave" deviations from the actual location.

As seen in plots 3.12b and 3.12d, localization error spikes are observed for ICS due to rapid acceleration when the direction of rotation changes. Additional localization spikes are seen when $\theta = 0$ for plot 3.12b. Force transients can be generated when the object goes though $\psi = 0$ due to the sign reversal of the shear force on the compliant rubber skin.

Table 3.1 compares the localization errors associated with each of the sensing schemes for the pin rotation and translation experiments. First, transient peak errors for ICS are several times larger than the array sensor. Second, shear forces introduced in the pin rotation experiment only affects the centroid algorithm of the tactile array sensor. This error is most clearly seen during off centered pin rotation.

	Centroid	Threshold	Peak	ICS
Pin Translation Centered - RMS Error [mm]	0.05	0.10	0.09	0.18
Pin Translation Centered - Peak Error [mm]	0.12	0.17	0.23	1.10
Pin Rotation Off Center - RMS Error [mm]	0.78	0.18	0.25	0.34
Pin Rotation Off Center - Peak Error [mm]	2.94	0.40	0.52	2.42

Table 3.1 RMS and peak localization errors for the pin rotation and translation experiments. The first three columns are the tactile array sensor errors and in the fourth column is the ICS errors less the initial contact transient errors.

3.4.4 Box Manipulation Results

For figures 3.13 and 3.14, the experiments involved translating and rotating a box, similar to the pin experiments. Contact locations calculated based on the kinematics of the robot and the object are plotted as a solid line for comparison purposes. This calculation requires a priori knowledge of object properties which would not be available in unstructured environments, but represents a reasonable approximation to the actual contact location.

The expected trace for the box translation experiment in plots 3.13a and 3.13b is a trapezoidal wave since the fingers roll mostly during the vertical traverse phase. Again, the ICS trace has more high frequency noise than the array sensor traces, even though the frequency content of the ICS data is mostly under 20 Hz, and thus comparable to the bandwidth of the array sensor signal. The peak localization scheme is significantly different

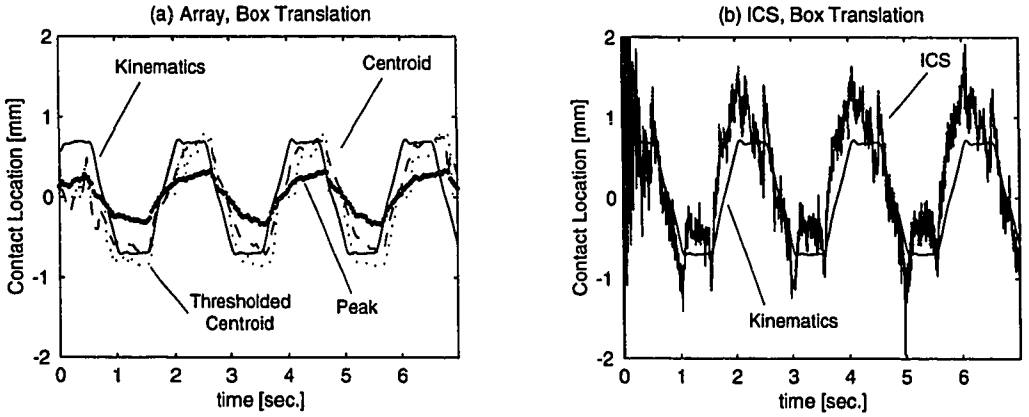

Figure 3.13 Contact localization for box translation experiment. (a) Array sensor; (b) ICS. Solid line represents calculated results using robot kinematics and finger tip and object models.

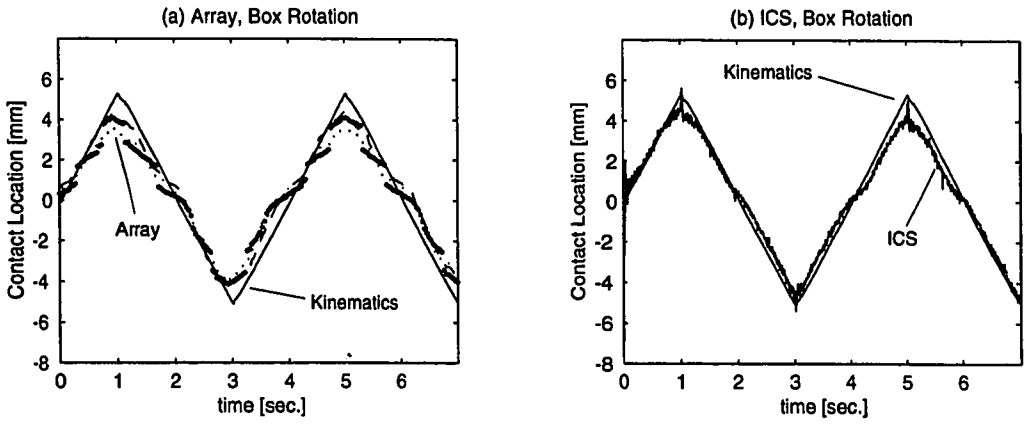

Figure 3.14 Contact localization for box rotation experiment. (a) Array sensor; (b) ICS. Solid line represents calculated results using robot kinematics and finger tip and object models.

from the centroid schemes. This suggests that violating the initial assumption about single point contact affects the peak localization algorithm the most.

For the case of rotating a box shown in plots 3.14a and 3.14b, the contact location changes as the finger tip rolls against the box. Actual peak to peak object rotation is slightly less than 40 degrees since the fingers roll against the object, and this effect is not taken into consideration in the controller calculations. Because of slow manipulation speed and the absence of shear forces during rolling, both sensing schemes show about the same accuracy in tracking the contact location.

Careful examination of the peak localization scheme in plot 3.14a reveals discontinuities in the localization trace. Ideally, the finger and the box should form a uniform line contact, but during manipulation, the peak strain jumps positions within this line of contact due to edge effects and non uniform loading associated with manipulation. Since all the tactile elements were used for the peak detection algorithm, these effects reveal themselves as discontinuities in the trace. This problem could easily be fixed by averaging the rows perpindicular to the line of contact.

3.5 Discussion

The goal of this chapter is to examine fast, practical methods which can be implemented easily on real manipulators. From these calculations and experiments one can derive several

important lessons. Both the tactile array and the ICS schemes were able to localize contact within a few millimeters in the presence of manipulation noise such as shear and inertial forces. Although the hardware for the ICS is much simpler, there are serious difficulties in the calibration procedure, especially with soft finger tips. Many of the errors associated with the intrinsic sensing method can be traced to difficulties in calibration. When calibrating a force sensor for this application, it is important to precisely control the contact locations as well as the magnitudes and orientations of the applied point-load forces. In addition, it is desirable to do so with the finger tip attached. However, a point load indents the rubber finger tip, making contact location imprecise and generating shear forces. This is a generic problem for all force sensors with compliant finger tips.

The ICS algorithm should be thresholded to obtain reliable contact location during initial contact. Before firm contact is established, the signal levels are on the order of the noise level so the localization errors are large. In addition, inertial forces can affect the ICS method, so rapid accelerations generated during manipulation should be minimized or ICS derived location should be ignored during that period. One way to reduce inertial forces is to use a smooth trajectory which minimizes accelerations. Another method involves minimizing the finger tip weight. Since the calibration is conducted on the finger tip surface, any forces generated on the surface, including object inertial forces, will be taken into consideration in the calculations. Therefore, only the inertial forces generated by the finger tip mass affect the localization scheme. This phenomena can be observed by the fact that despite an object weight ratio of 18:1 for the box to the pin, localization errors associated with the translation experiments remain essentially constant.

The tactile array sensor, on the other hand, is a more complex device. It requires a far greater effort to fabricate the sensor and the associated electronics. The calibration procedure required for the array sensor would have been similar, but using compressed air simplified the calibration procedure tremendously. In addition, the spatial characteristics of the sensor elements are uniform, so spatial calibration is not required. The lower mass associated with the array sensor makes it less susceptible to sudden accelerations for those elements not in contact with the object.

The three schemes for localizing contact provided interesting comparisons. The primary difference between the schemes seemed to be their sensitivity to shear forces, with the simple centroid scheme being most susceptible. For point contacts, the peak and thresholded algorithms worked well in rejecting shear induced strains. However, the centroid schemes performed better for larger contact areas. This suggests that the type of algorithm used should be tailored to the expected contact conditions.

Examination of the array sensor pressure distribution during box manipulation reveals that the pressure distribution is not uniform as expected. There are edge effects and shifting of peak locations. This creates some difficulty for the peak localization scheme which jumps from one peak to another in a discrete fashion. This suggests that for manipulating flat objects, a spherical tactile sensor or a softer rubber surface might provide more consistent information.

3.6 Conclusions

This chapter presents calculations and measurements of localization errors in characteristic manipulation tasks using intrinsic contact sensing and tactile array sensing schemes. With the exception of ICS transients, experimental results show that the RMS localization error for all the sensing schemes are less than 1 mm. The ICS scheme, while potentially fast and accurate, requires a great deal of care in the calibration process, and it is affected by transient forces (e.g., initial contact, object accelerations) which can produce errors of several millimeters. It also suffers from a poor signal-to-noise ratio at low contact forces. Therefore, to effectively use ICS, finger tip mass should be minimized, contact forces should be thresholded before applying the algorithm, and either accelerations should be minimized or inertial transients must be taken into account.

The tactile array sensor works well for contact localization. However, it is inherently slower than ICS and much more difficult to construct and integrate with the rest of the manipulator and controller. The tactile array sensor's sensitivity to shear loads was minimal and only noticeable at very sharp angles using a simple centroid scheme. For manipulation purposes, the shear load effects can be minimized using a peak detection or a thresholded

centroid algorithm. Since peak localization schemes work well for point contacts, and centroid localization schemes work better for larger contact areas, algorithm selection should be based on manipulation object types. If this information is not available, a thresholded centroid algorithm should be a good compromise.

With care, both of these sensing schemes can provide contact localization accuracy of less than a millimeter in real manipulation tasks. This should be sufficient for a large number of application domains. Regarding the general question of which type of sensor should be used, or which sensor performs better, other benefits of each sensor should be considered as well. The force sensor has many other uses such as directly measuring contact forces and implementing stiffness controllers. The array sensor can provide shape information as well as pressure distribution across the contact area. Naturally, both sensors can be used to provide redundant information to increase robustness of the entire system.

Chapter 4. Grasping and Grasp Refinement

4.1 Introduction

Human performance of delicate manipulation tasks is governed by a task-dependent combination of visual, kinesthetic and tactile cues. Delicate manipulation in particular demands minimization of the forces transmitted to the object during grasping. This suggests tactile or force sensing should play a primary role in controlling interaction with objects. In unstructured environments, vision is essential for determining a location to place the robot hand for grasping, as well as a source of positional feedback once the object is grasped.

In this sense, vision and touch play largely complimentary roles: vision provides position and shape information at a distance, while touch provides small-scale geometric and force information. By quickly locating and recognizing the object, vision enables a fast and direct approach before a manipulator contacts an object. Thus, vision is particularly effective for free-space motions. Once the fingers make contact with the object, however, touch sensing becomes preeminent, since it bypasses the obstruction and resolution limits of vision. Tactile sensing reveals the exact locations of the contacts between the object and the fingers, as well as the object's local surface shape, which permits optimization of the grasp. In addition, touch sensing can be used to minimize contact forces, thus preventing unwanted disturbance of the object before the grasp is secured.

Most previous work combining vision and touch has concerned perceptual issues. One focus has been the use of touch sensing to supplement visual information for object

recognition or exploration [Allen 1987, Stansfield 1988, Boshra and Zhang 1995]. Other work has developed representations for the integration of visual and tactile spatial information [Rafla and Merat 1990]. Rucci and Bajcsy [1995] have examined coordination of active visual and tactile perception, and Ono et al. [1995], have begun to apply these sensing modalities to manipulation by combining vision with a simple form of tactile sensing for handling flexible materials.

This chapter begins by outlining conceptual frameworks for combining visual and tactile sensory information in control of manipulation, along with descriptions of the visual and tactile primitives used in this chapter. The subsequent section describes the experimental hardware and procedures. Experimental results are then presented, and performance is evaluated in terms of final object orientation angles and net forces applied to the object. This chapter concludes with a discussion of the experimental results and implications for future research.

4.2 Combining Vision and Touch Primitives

Effective combination of vision and touch requires consideration of many factors, especially the capabilities and limitations of each sensing modality and the task context. The unique strength of vision, which fundamentally measures distances and directions, is its ability to sense motion over a large space. Because of the large sensing region associated with vision, the smallest resolvable feature is limited for typical manipulation tasks. Vision also requires an unobstructed view. Touch, particularly tactile arrays, senses local geometry and pressures, providing both length and force measurements. Sensing takes place at contact locations which are hidden from view by the robot hand. Existing tactile array sensor resolutions are coarse, but because of the small sensing region, it may provide better spatial resolution than vision. Since tactile sensing only provides local contact information, without vision, the use of tactile sensing is limited to groping during the initial acquire phase of a task.

These considerations suggest that a variety of combined sensing modalities will prove useful. In some situations both visual and tactile information provide spatial information, which can be "fused" to enhance resolution. One example is finding finger tip contact locations for grasping a complex object: vision can provide global shape information at limited resolution, while tactile sensing can measure local curvature and surface normals. Finger placement can be optimized by using both measurements to ensure full kinematic restraint while minimizing the chance of slips at each finger tip with tactile sensing.

In other situations, the sensed information will be complementary, and the robot controller may be configured to use the both sensing modalities. For example, while an object is grasped, the tactile sensor could provide the grasp force information while the vision system provides the object position information. The experiments described below involve grasping a target object at an arbitrary location in the workspace with a two-fingered gripper using minimal forces. At the beginning of the task, vision provides the only source of information, and the robot must be visually servoed. When the fingers of the gripper make contact with the object, contact location and force information are provided only by touch, and the controller must use contact information to minimize forces and servo the contact location.

This chapter focuses on the complimentary nature of visual and tactile information. The approach taken is based on sets of visual and tactile "primitives," which are basic sensing and control modes appropriate to various segments of a task. These primitives are then combined for task execution. In the following sections we describe these vision and touch primitives, and frameworks for their integration.

4.2.1 Framework for Vision "Feedback Primitives"

Hager [1995a], proposes a conceptual framework for static hand-eye coordination tasks. The central notion in this framework is that of a task-space kinematic constraint. Any robot task is specified in terms of the kinematic constraints that should be achieved between the robot and a target object. For example, if the goal is to place a screwdriver onto a screw,

the desired kinematic constraints are the alignment of the axis of the screwdriver with the axis of the screw, and the tip of the screwdriver touching the head of the screw.

The central idea in Hager's visual feedback primitives is to translate task-space kinematic constraints into sensor-level constraints. Given sensor-level constraints, an image-based error function is defined such that the error function is zero if, and only if, the task-space constraints are satisfied. Standard linearization techniques and PID control methods are then used to drive the manipulator into the desired configuration. Details on this process can be found in [Hager 1995b]. The robot is treated as a velocity-controlled device which acts as a pure integrator. This approach neglects many potentially important dynamic effects, but for low-velocity tasks it has been experimentally validated as a reasonable model of the system [Hager et al. 1995]. This transforms the control problem into one of generating velocity commands from sensor data.

Here, the essential ideas behind the theory using simple point-to-point motion "primitive" is reviewed. Let R and T be points in space attached to the robot manipulator and a target object, respectively expressed in robot base coordinates. These points are imaged into a camera using the usual perspective map which is denoted by $\pi : \Re^3 \rightarrow \Re^2$. Suppose r and t are the projections of these points in the camera image, and the robot translates with some velocity V. Then the projection of R will move with some velocity v. The relationship between these velocities is given by the so-called image Jacobian $J = \partial\pi / \partial R$. Suppose now that we have two cameras observing R and T. It is possible to show that R is coincident with T if and only if their projections are coincident in both cameras. Let r_i, t_i, and J_i denote the projections of R and T and the image Jacobian for cameras $i = 1,2$. Define the image error, e, by the errors in the camera image

$$e = \begin{bmatrix} r_1 - t_1 \\ r_2 - t_2 \end{bmatrix}, \text{ and } J = \begin{bmatrix} J_1 \\ J_2 \end{bmatrix}. \tag{4.1}$$

Let J^τ denote the generalized inverse of a matrix. It follows that applying feedback

$$s = -J^\tau e \tag{4.2}$$

to the robot will stabilize it so that $R \rightarrow T$, as $time \rightarrow \infty$ [Hager et al. 1995]. In general, a manipulator is controlled by a velocity screw s which controls six degrees of freedom of motion. Thus, visual kinematic constraints up to and including six degrees of freedom can be specified and carried out by this type of feedback arrangement. Note that the Jacobian will generally be $m \times 6$, where m is the degree of the visual constraint.

By defining other features such as a line which is formed using two points, and a corner which is formed using two intersecting lines, the number of features used in the visual "feedback primitives" can be increased. These features can then be used to define other visual "feedback primitives" such as point-to-line and line-to-line primitives. In point-to-line primitive, the shortest distance from the specified point to the line is minimized, and in line-to-line primitive the orientation and positional error between the two lines must be reduced.

4.2.2 Tactile Primitives: Task-Space Approach

This section describes an task-space approach for creating tactile primitives which are low level control scheme based on tactile information. As described in section 1.2, manipulation tasks can be broken down into phases separated by events. Tactile sensing can provide information for continuous feedback control during a phase of a manipulation task in addition to detecting events, the discrete transitions between manipulation phases. Common events include making or breaking of contact between the fingers and the object, and the onset of sliding between the finger and the object. A number of researchers have proposed higher level control algorithms which require tactile event information to increase the robustness of task execution [e.g., Cutkosky and Hyde 1993, Brock 1993]. The primitives described here are representative of the range of tactile based behaviors that can be constructed.

Sensing first contact between the hand and an object is important for minimizing forces on the object. A wrist force sensor's ability to detect small forces is limited due to dynamic

range issues, and it is affected by gear noise due to the gripper mass in front of the sensor as shown in chapters 2 and 3. The array sensor used here has a relatively slow response because of the multiplexing required, but it can be used to determine many other parameters including contact location. A more suitable sensor for detecting contact, such as the stress rate sensor, could provide more sensitive and timely response. The "make-contact primitive" begins with the hand approximately centered on the object by visual control, and the fingers closing in on the object. When the sum of the readings from all elements of an array exceeds a threshold, the gripper is commanded to stop. Typically, contact occurs on one finger first followed by others.

Before closing the other fingers, a "reorient primitive" senses the object's orientation with respect to the gripper and reorients the hand to permit the desired grasp. Determining object pose has been an important area of research for tactile sensing. Fearing [1990a] has analyzed this problem in the general 3–D case using contact ellipsoid models, and Chen et al. [1995] have determined orientations using moment area methods. Here the focus is on the rotational axis parallel to the cylindrical sensor axis, where the contact centroid location is directly proportional to the object orientation angle.

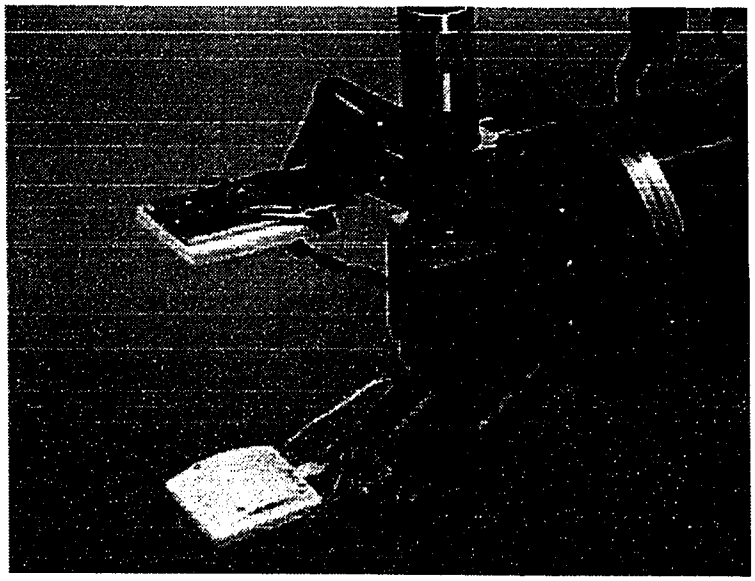

Figure 4.1 Tactile array sensors mounted on the Zerbra-Zero gripper.

Although the "reorient primitive" described here is specialized to the one degree-of-freedom two-fingered cylindrical gripper used in the experiments (Figure 4.1), it illustrates the issues for the general multifingered grasping task. The primitive starts with one finger in contact with the object. From a cylindrical sensor contacting a parallel flat sided object, the orientation of the object relative to the finger is simply $\phi = \alpha / \rho$, where α is the contact pressure centroid location on the sensor surface and ρ is the sensor radius. In this task, the goal is to achieve contacts at symmetric locations on opposite sides of the object, which will minimize the net applied torque and maximize grasp stability. This pertains when $\phi = 0$.

Once the object's orientation has been determined, there are many ways to reorient the hand to ensure that the fingers reach the desired final contact locations. With a multijointed, multifingered hand, it may be possible to move only the finger joints to achieve the goal configuration. Since the gripper used in the experiments has only one degree-of-freedom, the joints of the robot arm must also be used. As a start, the gripper is rotated through an angle opposite the measured object orientation angle. Because the manipulator rotates about a tool coordinate frame located at the midpoint between the gripper finger tips, the sensor temporarily breaks contact with the object.

After the correct orientation is obtained, the gripper must close to bring both fingers into contact. This "make-two-finger-contact primitive" requires coordinated motion of the arm and gripper to avoid applying large forces at the contact. Tactile array information can be used in the feedback loop to minimize disturbances during the grasping process. The array sensors' readings are summed to find, $F_i^{act.}$, the total contact force on the finger tip i, and this is then used in a simple proportional gain control law:

$$v_i = K_i(F_i^{des.} - F_i^{act.}) \tag{4.3}$$

Transforming these velocities into tool frame, v_t, and gripper, v_g, velocities,

$$v_t = (v_1 + v_2)/2, \text{ and } v_g = (v_1 - v_2)/2. \tag{4.4}$$

Once stable two-finger contact has been made, the tactile force information permits grasp force control. Instead of commanding velocities at each finger to generate a desired

grasp force, only the gripper velocity is required for the "increase-grasp-force primitive." If the "make-two-finger-contact primitive" was used to increase the desired grasp force, any differences in total force calculated by the sensors will generate tool frame velocities which will cause the gripper to drift. Therefore, it is advantageous to send gripper velocity commands based on the average grip force rather than defining each finger tip velocities.

$$v_g = K(F^{des.} - F^{avg.}) \tag{4.5}$$

4.2.3 Tactile Primitives: Sensor-Space Approach

An alternative approach to generating tactile primitives is analogous to vision primitives described above in section 4.2.1. Instead of focusing on the different segments of the manipulation task, task-space kinematic constraints are mapped to sensor-level constraints and tactile feedback is used to obtain the desired results. More specifically, feedback control is implemented on the force and contact centroid information provided by the tactile array sensor. The advantage of this scheme, which uses a feedback loop to maintain contact, is that if the object moves during the grasping process, the controller can compensate for these changes.

The "sensor-space primitive" compensates for object orientation angle and the contact force by commanding manipulator velocities which reduce these errors. It uses equations (4.3) and (4.4) in the "make-two-finger-contact primitive," while at the same time controlling the tool frame rotational velocity based on contact centroid information. Here, conflicts between goals for reorienting and contact force control is not an issue since they control different velocity axes, but in general the "feedback primitives" must be constructed to avoid conflicting goals.

Again, a simple proportional control law is implemented,

$$\omega_z = K_\omega(\phi^{des.} - \phi^{act.}). \tag{4.6}$$

where, ω_z is the rotational axis controlled by the tactile sensor, and $\phi^{des.}$ and $\phi^{act.}$ are the desired and measured object orientation angles from the contact centroid information. For a parallel sided object, servoing the contact centroid to the center of the sensor corresponds

to $\phi = 0$. This primitive finishes when the desired contact forces are within the threshold on both fingers. Since the gripper does not have additional degrees of freedom, the orientation of the object with respect to the gripper cannot be corrected after the object has been grasped.

4.2.4 Integration for Task Control

The immediate issue that arises is how vision and tactile feedback should be combined. Here, two simple and natural combinations, serial and prioritized parallel combination, are discussed. Serial combination can be described as "vision then touch." This is a useful mode of interaction when vision can be used to roughly place the robot end-effector near the desired goal configuration, and then tactile sensing is used to determine the object to finger interactions. The advantage of serial composition is that the two modalities operating independently makes their combination extremely simple. The disadvantage is that feedback from tactile sensing may undo constraints that have been achieved by visual sensing.

Prioritized parallel composition provides for more direct combination of tactile and visual feedback by superimposing feedback from one modality onto degrees of freedom which the other modality does not constrain. For example, suppose that a 6 degrees of freedom manipulator is performing point-to-point motion as described above. Then the image Jacobian will be 3×6. Suppose that s_t is a motion commanded by tactile sensing. If the goal is to maintain the visual constraint, but superimpose tactile-based motion, then a feedback of the following form can be used:

$$s = -J^{\dagger} e + (I - J^{\dagger} J) s_t \tag{4.7}$$

where as before J is the image Jacobian for the visual constraint, and e is the image error for that constraint. The term on the right projects the tactile motion into the nullspace of the image Jacobian. Physically, this causes the robot to move under tactile control only in directions which are uncontrolled by visual sensing. Note that to the extent to which tactile sensing can be written as a feedback system of this form, it is also possible to give

tactile priority over visual feedback. A simple case of this feedback law is orthogonal combination, in which tactile and visual feedback control independent degrees of freedom. For example, the vision system positions the gripper using a point-to-point primitive, and the tactile sensing system is used to orient the gripper.

4.3 Experiments

Most of the tactile sensing experiments presented in this thesis were performed using a planar manipulator which limited manipulation tasks. Greg Hager's group at Yale was using a Zebra-Zero manipulator (Integrated Motions, Inc., Berkeley CA), and providing visual feedback to control it. Their ability to grasp objects was limited to feedforward approached because of visual occlusions. Therefore, the purpose of the experiments presented in this chapter was to combine the sense of vision with touch while grasping objects.

4.3.1 Hardware Description

The entire system architecture required the use of four separate processing platforms: dedicated motor controllers, a host PC for the Robot, another PC for the tactile signal processing, and a workstation for the vision signal processing (Figure 4.2). The Zebra-

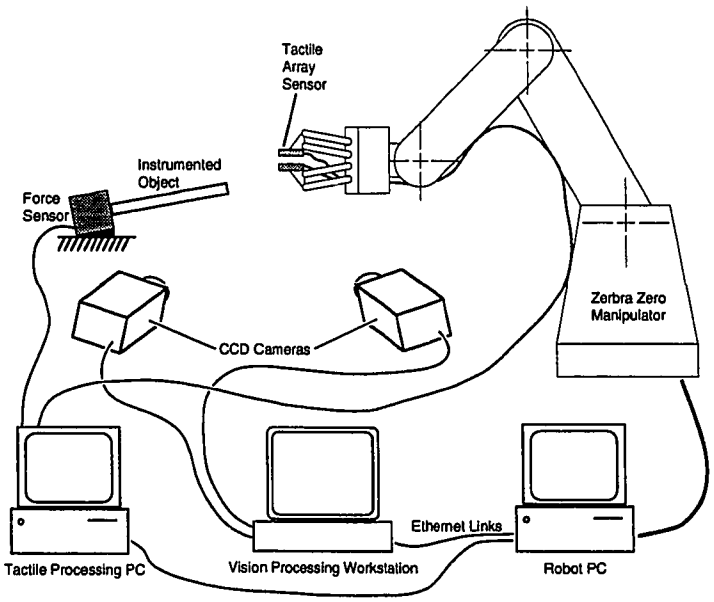

Figure 4.2 Schematic of the experimental setup.

Zero manipulator is a medium sized 6-axis manipulator especially designed for the research market. The gripper is actuated using an electric motor with a screw gear. At the lowest level, each joint is controlled by a dedicated motor control IC. These controllers can be configured for velocity or position control which leaves the host computer to focus on coordination issues such as calculating the kinematics, Jacobian, coordinate transformations, etc. The signal processing computers send tool frame velocities to the robot host PC via ethernet connections.

The advantage of this type of distributed system is that it is possible to segment the problems into manageable pieces such as vision, tactile, kinematics, and motor control. This allows independent process development and simplifies debugging the entire system. In addition, computational power is increased through the use of multiple processors. The disadvantage is that communication between the subsystems is limited. In our implementation, we are only able to send desired velocity commands to the robot host PC.

The vision system consists of two Sony XC-77 CCD cameras with 12.5 mm lenses connected to a Sun 10/42 workstation via two Imaging Technologies FG101 frame grabbers. Visual tracking for both cameras and visual feedback calculations are performed on the Sun host using the XVision tracking system [Hager 1995b] and the Servomatic hand-eye coordination system [Hager 1995a].

In the experiments reported here, initial positioning of the gripper is performed using the point-to-point motion primitive described in section 4.2.1. The tracking system follows the motion of the gripper using a simple pattern matching method, which tracks by locally optimizing the sum of the squared differences between an area of the live image and a stored reference image [Hager et al. 1995]. A setpoint is computed in each camera, and then the distance from this setpoint to the object grasp location (which is also tracked) is used to implement point-to-point positioning. Relative positions between the setpoint and the desired grasp location can be set by specifying offset displacements.

The tactile sensors used in the experiment use the same capacitive tactile array technology presented in chapter 2. Each sensor contains 64 elements in an 8 x 8 matrix

equally spaced at 2 mm, and the entire sensor is encapsulated in silicone rubber to protect the sensor and increase friction and compliance. The sensors are mounted on a 2.54 cm radius cylinder which permits better contact localization and increase object mobility for multifingered robot hands. Since the Zebra-Zero manipulator has a much larger workspace in comparison to the planar fingers, extra long cables from the sensor to the electronics were required. This reduced the sensor gain which is approximately a function of sensor capacitance divided by the line capacitance as explained in [Fearing 1990]. Figure 4.1 shows a close-up of the tactile array sensors mounted on the gripper.

Several objects were used in these experiments. A 10 mm diameter aluminum tube mounted on a force-torque sensor measured two force components and one torque component in the plane of the gripper and object. This sensor was mounted on an aluminum frame which was securely clamped in place. Other objects used in the experiments were a pen and a floppy disk which were readily available and a heavy machinist's right angle plate which was stationary.

4.3.2 Procedures

The first experiment was conducted to compare different control algorithms for correcting initial object orientation angle. Here, no feedback, force sensor feedback, feed forward tactile sensing and tactile feedback are compared. The back end of the machinist right angle plate is used as the object, and it is presented in two orientations, parallel to the

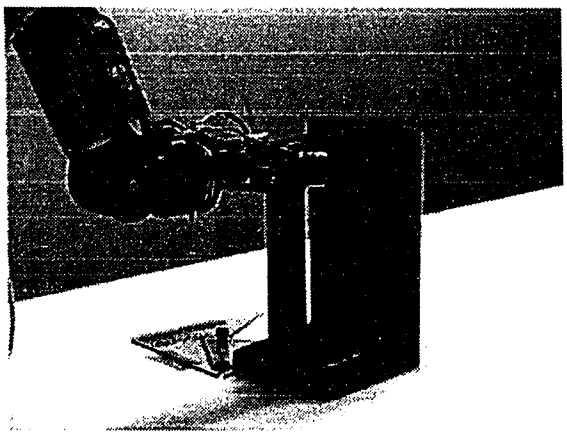

Figure 4.3 Experimental setup for the grasp orientation experiment which compared different control algorithms.

gripper and 10° off parallel, as shown in figure 4.3. The final orientation angle of the gripper with respect to the object is used as a comparison measure for the different control algorithms. Orientation angles were measured using the tactile array sensor. Calibration from contact centroid to object orientation angle was achieved using an instrumented object and an optical tracker which provided accurate object orientation information. Four hundred orientation angles from the optical tracker and the centroid data was fitted using a least squares approach with a standard deviation of 0.51 degrees.

The second experimental task required the robot to gently grasp an aluminum rod, which was fixed within the camera visual field and the manipulator workspace. To simulate an unstructured environment, the controller had no model based information about the object position or orientation aside from the assumption that it was grasping a parallel sided object. Task execution followed a serial combination of the primitives described in sections 4.2.1 and 4.2.2: once the vision system moved the gripper into position, tactile feedback minimized disturbances to the object while grasping it. Although numerous methods for combining vision with tactile sensing could be used, initial efforts were focused on integrating the entire system to complete the task successfully.

In this experiment, minimization of forces and torques applied to the object was used to evaluate performance of the system. Alternatively, displacement of the object could have been measured if the object was compliantly mounted. Minimization of disturbance forces, however, is a realistic constraint in handling delicate objects. It also emphasizes the differences between vision and touch, since vision cannot provide force information to the controller. Tactile sensing is used to detect contact, compensate for the difference between the orientation of the object and the gripper, make contact on both fingers, and finally increase the grasp force to the desired value. Not compensating for the orientation angle would impose a torque on the object through the difference in the contact locations where the grip force was applied. Grasp stability would also be reduced.

As a comparison benchmark, the task was executed using an alternative technique to minimize disturbance forces. One of the built-in force control modes of the Zebra-Zero

robot is "sfloat" or "servo float," which controls the robot joints to minimize the forces and torques measured at the wrist force sensor. In this mode, the robot attempts to compensate for the orientation angle difference and the translational difference of the object through sensed forces at the wrist. This wrist-based compliant motion scheme provides a simple and immediate means of verifying the benefits of tactile sensing.

In the third set of experiments, "sensor-space primitives" are use to perform the human hand off task where the robot grasps an object held by an individual. Since the object is hand held, it can easily move during grasping. Rather than having the human adapt to the robot, the goal is to have the robot adapt to the human just as in handing the object to another person. A pen and a floppy disk are held by an individual and moved during the grasping process.

When the object orientation changes out of the camera view plane, visual determination of the object orientation is difficult because of a singularity in the stereo mapping function. However, rough positioning and orientation in the remaining axes is possible. Hence, this is a situation where tactile and vision provide some redundant and some complementary

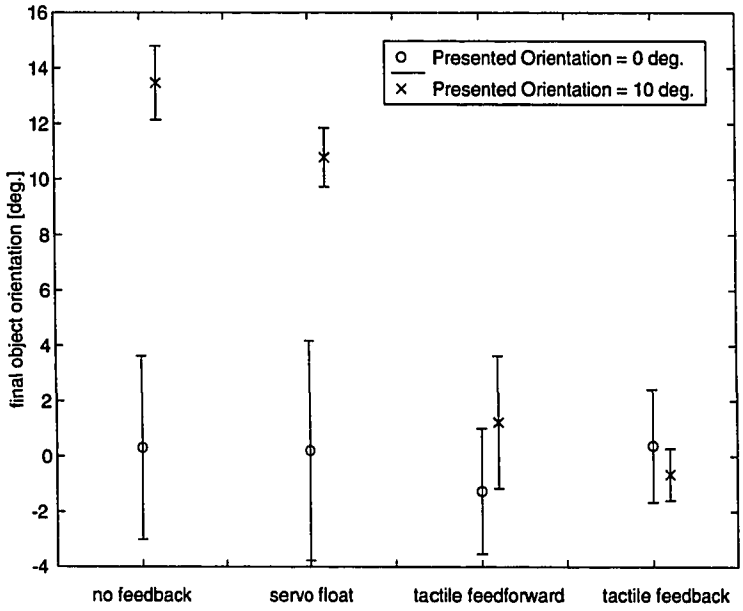

Figure 4.4 Final orientation of the gripper with respect to the machinist's square handle for different control algorithms. The handle was presented at zero and ten degrees from the gripper. Symbols are mean values and the bars represent standard deviations from for 4 trials.

information. In this implementation, the vision system controlled the robot motion in the camera view plane and tactile sensing controlled translation in the camera axis and orientations out of the camera view plane.

4.3.3 Results

The first experiment compares different control strategies for reorienting the gripper parallel to the object orientation angle. Figure 4.4 is a compilation of orientation angles from multiple trials. When the object is parallel to the gripper, the orientation angles measured using all of the approaches (no feedback, force feedback, tactile feedforward and tactile feedback) all resulted in a mean final orientation angle near 0°. The error bars correspond to one standard deviation. The variations in measured orientations are large due to friction and backlash of the robot, which limit the ability to control it precisely. In a less significant way, tactile sensor noise and calibration also added to the orientation errors.

When the object is presented at 10°, the expected final orientation angle is 10° for the no feedback trial and 0° for all the other trials. The handle was not centered exactly between the fingers so disturbance forces were generated during the grasping process, which explains the larger than 10° orientation angle for the no feedback case. The force feedback or "servo float" approach was not able to compensate for the orientation angle due to lack of sensitivity, and the moment generated by the gripper was relatively small. The backlash of the gripper mechanism and the compliance between the contact point and the force sensor also limit this approach. Both of the tactile sensing approaches did however control the manipulator such that the final orientation angle was close to the desired 0°. The feedback approach performed a little better than the feedforward approach described in the "reorient primitive."

In the second experiment, integration of vision and tactile sensing was serial. The vision system positioned the gripper to the end of the rod and tactile sensing was used to perform the tactile primitives generated with the task-space approach. The top three images in figure 4.5 show how the vision system is able to move the gripper, and allowing it to

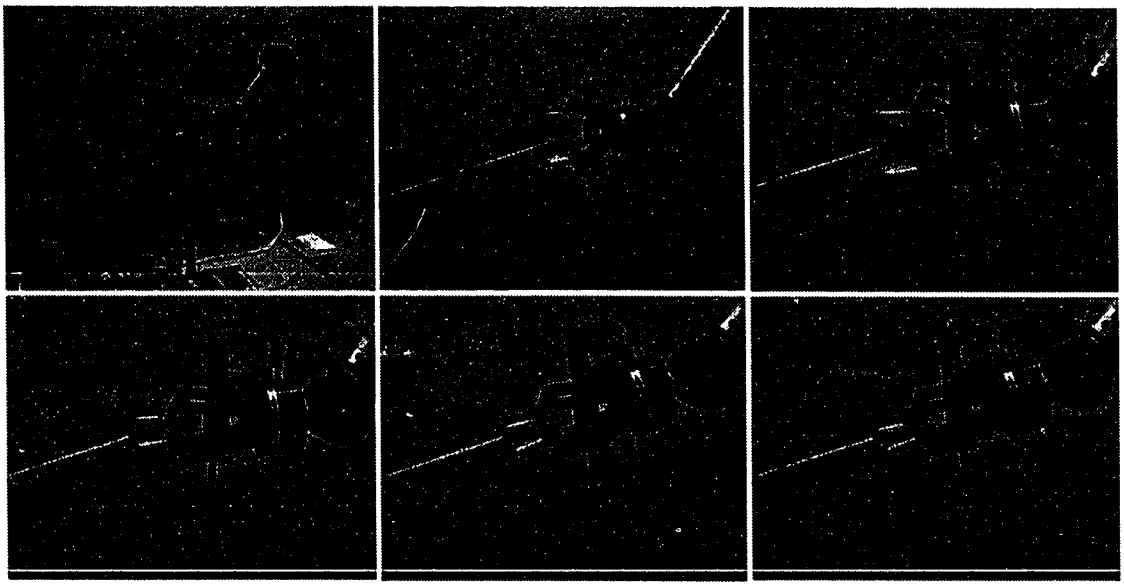

Figure 4.5. In the top three photos, the vision system brings the gripper to the object using a point-to-point primitive. In the bottom three photos, the tactile sensors are used to grasp the object by detecting contact, compensating for orientation angle, and controlling the grasp force.

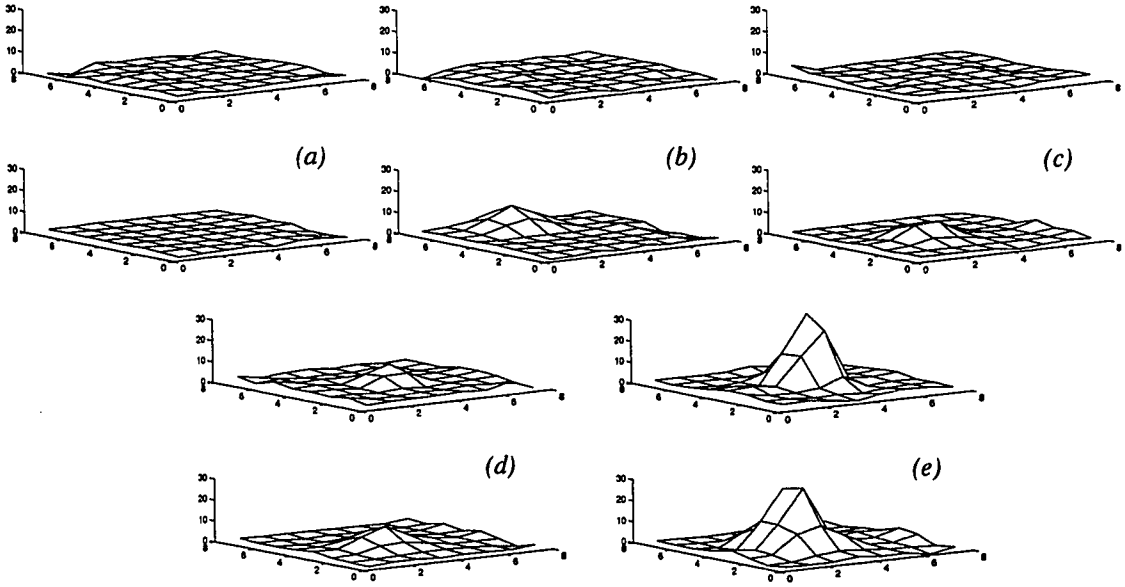

Figure 4.6 Tactile array sensor response during the grasping process. The vertical axis units are in psi. The response is flat (a) until contact occurs at the edge of the bottom sensor (b). As the manipulator compensates for this orientation and starts closing the gripper, only the bottom sensor is in contact (c). Contact forces are minimized until both sensors make contact (d), where upon the grasp force is increased to the desired value (e).

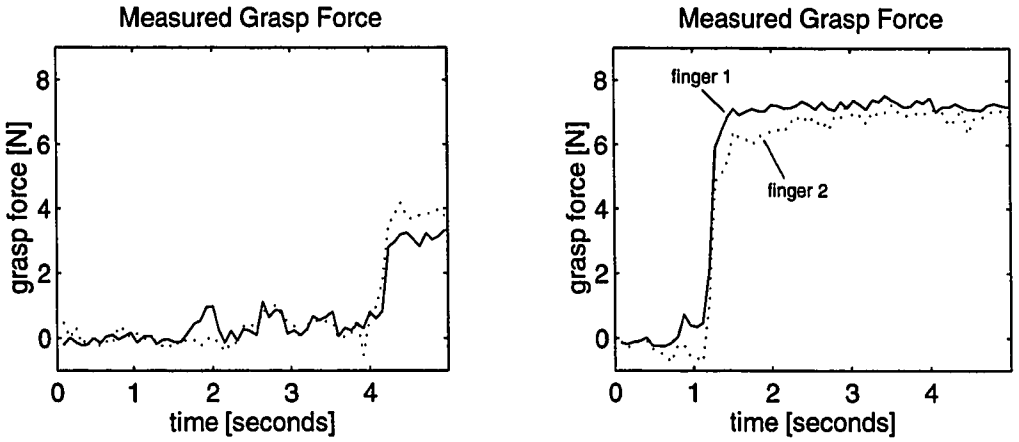

Figure 4.7 Left: grasp force when tactile sensor is used. Right: grasp force without tactile sensing.

grasp the end of a rod. Feedback computed from the stereo camera images provides sufficient information for a complete three-axis positioning relative to the tip of the rod. Once the visual feedback velocities to the robot endpoint dropped below 0.5 mm/sec in all axes, tactile feedback was invoked with the "task-space primitives" described above. The bottom three images in figure 4.5 show how the tactile array sensor is able to detect contact, and compensate for object orientation before grasping the object. Figure 4.6 shows measurements from the array sensor during each of these phases.

The final orientation angle reached using the "reorient primitive" was also observed to be much closer to the desired value than that achieved by "servo float" mode. Even when the cylindrical sensor was replace with flat parallel jawed gripper, it was not sufficient to overcome friction and completely reorient the gripper.

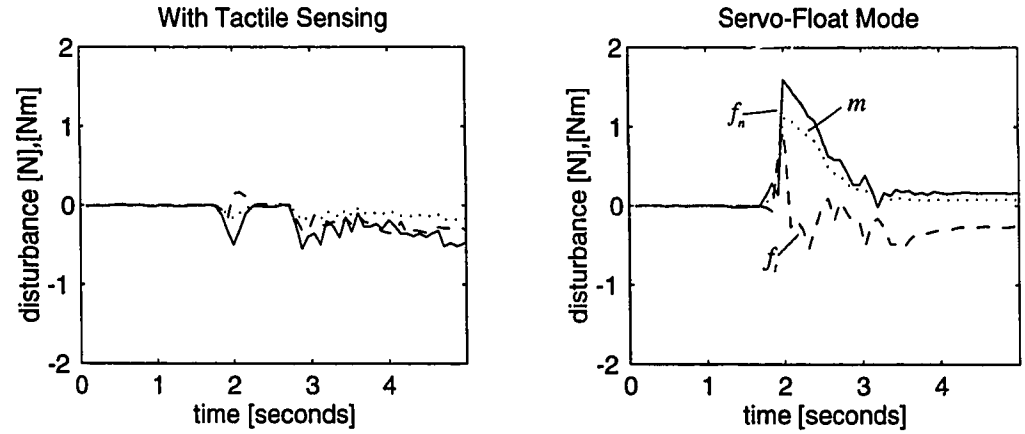

Figure 4.8 Disturbance forces and moment measured on the instrumented object. Left: tactile array sensor is used to minimize disturbances. Right: wrist force sensor is used to minimize disturbances.

Without tactile sensing or finger tip force sensing, the grasp force on the object is difficult to obtain. Some crude measure of the grasp force can be estimated though the motor current, but errors associated with this approach are large due to friction. Figure 4.7 shows the total grasp force computed by adding all the array elements. By using this information, the grasp force on the object can be controlled. In figure 4.7a, the final desired grasp force was set to 2.2 N, and the higher than desired result is due to the feedforward approach combined with the motor inertia which slowed the response of the gripper. Without tactile sensing, the grasp force was over 6 N, due to small positional changes which can generate large forces and poor motor torque control which is limited by friction.

Next, the forces and moment imparted to the object with tactile sensing and "servo float" mode are compared. Figure 4.8 shows that with tactile sensing, the peak contact forces are reduced from 1.6 N to 0.5 N and the peak torques applied to the object was reduced from 1 Nm to 0.1 Nm. The sensitivity of the array sensor was reduced due to the extra long cables required for the experiment. The tactile sensing scheme also exhibited larger time delays than the force sensor. This is due to several factors, most notably the 0.04 ms required to scan the array information and compute velocity commands, and the time consumed in transmitting the commands to the host PC. Despite these limiting factors, all of which are unfavorable for the tactile array sensor, the array sensor is able to improve the performance of the grasping process.

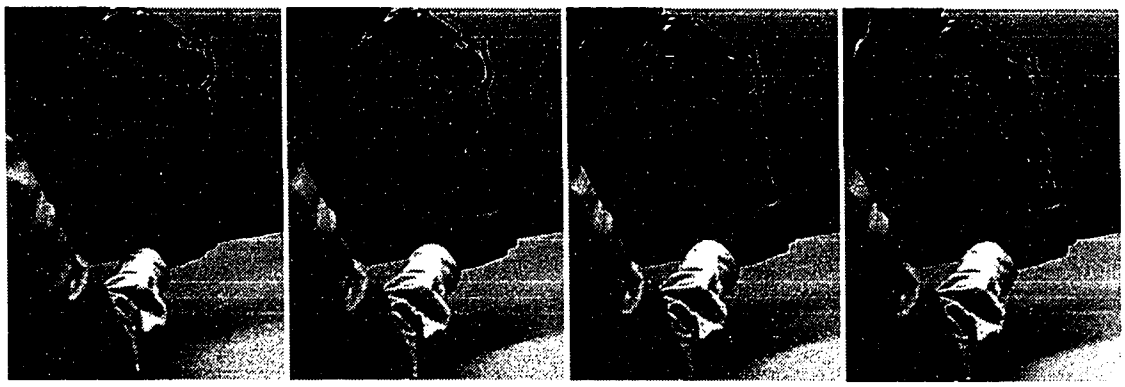

Figure 4.9 Here, the robot controls the contact force and the gripper-pen orientation angle while the gripper is closing. The "feedback primitive" compensates for changes in object position and orientation during the grasp by using tactile feedback.

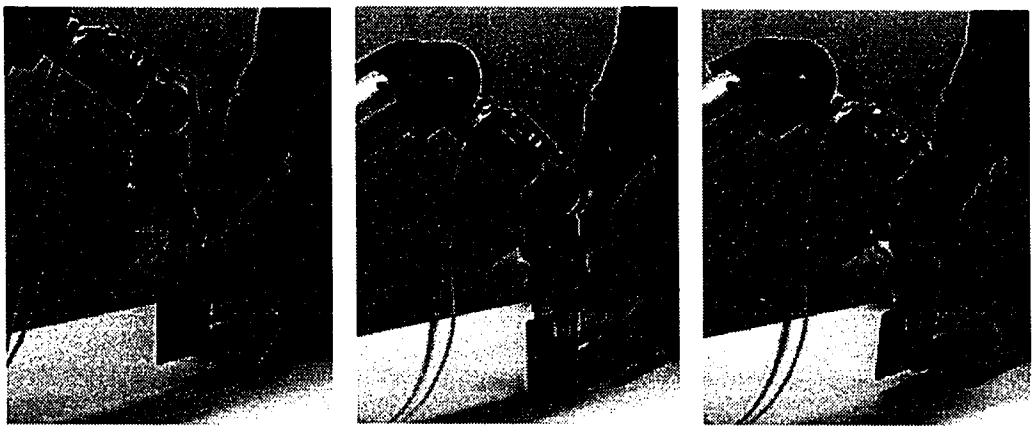

Figure 4.10. Orthogonal combination of vision and tactile sensing. The vision system controlled the robot in the camera image plane while the tactile sensing controlled the rest of the axes. The photos show that the vision system is able to track the moving disk and place the gripper around the floppy disk. When the disk is rotated by the person holding the disk, tactile feedback compensates for the changes in orientation angle.

Finally, human hand-off is demonstrated where the pen is moved and rotated during the grasping process in figure 4.9. Forces required to move the manipulator were small and the manipulator was able to compensate for the changes in pen orientation quite well. Orthogonal combination of vision and tactile primitives is demonstrated in figure 4.10, where a person moved the floppy disk during visual servoing, but the vision system was able to move the manipulator to the final desired location. During grasping, the person also rotated the floppy disk in the orientation difficult for the vision system to determine, and the tactile feedback adapted to the new orientation angle. Since both sensing modalities were used, the robot was able to compensate for the disk moving during the grasping process and minimized forces on the object.

4.4 Conclusions and Discussion

As noted above, these preliminary experiments have focused on integrating the tactile and visual system, and developing a set of tactile feedback strategies. Two simple methods of combining vision with touch, serial combination and orthogonal combination, were demonstrated. Both approaches worked reasonably well with the proper initial conditions.

The experiments presented here also demonstrate that vision is sufficiently accurate to provide an initial set of conditions in which tactile feedback operates correctly. Tactile sensors can compensate for uncertainties in object orientation angles and minimize forces on the object, and tactile sensing resulted in a more gentle grasp in comparison to a wrist force sensor approach.

Two approaches were used to generate tactile primitives. For adapting to changes in the object position and orientation "sensor-space primitives" proved to be more robust than the "task-space primitives." This is not surprising, since the "task-space primitives" were merely strung together without a supervisory controller.

The real advantage of the "sensor-space primitive" approach is that the desired actions are carried out automatically without the need for a higher level controller which must decide when to initiate a particular phase execution. The "sensor-space primitive" also can compensate for changes in object orientation right up to the point where both fingers make contact. The "task-space primitive" requires a higher level planner which determines that the object has moved and responds accordingly. This is not to say that the "sensor-space primitives" does not require a higher level planner. It just means that the high level planner needs to intervene less frequently.

This "sensor-space primitive" maintains contact with the object, but introduces sliding between the object and the sensor since the rotation is about the tool frame between the fingers. The disadvantage is that sliding can introduce disturbances, and sliding repeatability will depend on friction, which can vary significantly. Alternatively, to avoid sliding and breaking contact during reorientation, the robot can be commanded to roll the finger tip against the object, by rotating about the instantaneous contact centroid location. The rotational velocity about the contact centroid frame must be transformed to the tool frame coordinates. The transformation can be calculated from the contact location and gripper opening information. This transformation requires constant update since the contact location and the gripper position changes during manipulation.

It is interesting to note that one of the concerns before the experiment was manipulator controllability. In the force control literature, it is well known that transition to contact and maintaining contact both require high bandwidth feedback. For tactile sensing, it was unclear how much positional accuracy was available, and how much accuracy was needed. Backlash and manipulator stiffness were also thought to be issues. A very stiff manipulator could generate high forces, which could damage the sensor. However, none of these concerns proved to be problems during the experimental execution. The tactile sensor combined with the compliance of the manipulator and the ability to actively control grasp forces produced a reliable feedback mechanism.

The next step is to demonstrate a closer coupling between vision and touch, and to develop an intelligent sensor manager which controls the character and interaction of sensor feedback as a task develops. In particular, manipulation robustness will be improved when it is not possible for either vision or touch to achieve the task alone. For example, a sensor manager could intervene when visual obstruction prevents the use of visual servoing. One approach would use vision until the occlusion occurs, and then reduce the velocity while scanning the tactile or force sensors for contact, just as we put our hands forward and walk cautiously when walking into a dark room. Since the choice of which sensor to use and how to use it varies as the task progresses, sensor management is crucial.

Within the sensor management function, there are accuracy consideration issues similar to those discussed by Nelson and Khosla [1996] in the context of vision and force control. In particular, before transitioning to tactile based control, the visual system needs to bring the gripper into a position such that when the gripper is closed, the object makes contact with some active portion of the array sensor. This establishes the visual servoing accuracy requirements. One interesting aspect to explore in this vein is to move to an active camera so that visual servoing accuracy can be improved at certain points in the task execution.

There are several architectural issues that need to be resolved. A micro-controller interface to the tactile array sensor will reduce tactile feedback delays since the array sensor values can be quickly read from memory at a convenient time from the tactile interface PC.

Also, the current network architecture only provides one-way flow of information from the various sensor processors to the robot host computer. A more reasonable architecture is to use one management computer that monitors feedback from all three subsystems, and determines the appropriate actions.

Chapter 5. Object Interactions

5.1 Introduction

This chapter investigates how fundamental limits robot hand control affect tasks which require object interactions. In performing object interaction tasks such as assembly, the positional error between the two objects needs to be minimized. Object position error caused by robot hand friction can be minimized by increasing the controller gains, but errors in determining the object position cannot be compensated without tactile sensing, and higher stiffness gains also lead to larger interaction forces. For a peg insertion task where the objects are in contact, controlled object stiffness and specified center of compliance affect the interaction forces. However, the gain margin for modulating these parameters are fixed for a given robot hand grasp. Therefore, reducing the required object stiffness, through accurately determining object position with tactile sensing and compensating for manipulator friction, can provide lower interaction forces and improved ability to specify the object center of compliance.

A number of studies have demonstrated that mechanical impedance control permits effective robotic execution of certain tasks [Hanafusa and Asada 1983, Whitney 1982, Kazerooni 1988], and Hogan [1985] argues that impedance is a fundamental way to characterize mechanical interactions. In addition, many algorithms have been proposed for controlling a robot hand to achieve a desired impedance for a grasped object [Salisbury 1985, Khatib 1987, Cutkosky and Kao 1989, Li et al. 1989].

For simplicity, Salisbury's [1985] Cartesian object stiffness control is chosen to study the quasi-static component of mechanical impedance. This control algorithm attempts to make the grasped object act as if it was suspended by a set of springs and dampers. The controller senses displacement of the object from its desired location and orientation, and uses the robot fingers to generate the appropriate restoring forces and torques. Starr [1992] has implemented this controller on a Salisbury hand, but basic performance questions are not address.

Many of the multifingered control algorithms mentioned above assume fixed point contact conditions for simplified analysis. However, real fingers with compliance will form contact areas, and the contact location can move on the finger surface. Without tactile sensing, the stiffness controller can only sense object displacements through motions at the finger joints. Therefore, tactile sensing can enhance the performance of stiffness control by providing object-finger contact information.

To provide a focus, peg-in-hole insertion task is used as the prototype for a range of dexterous manipulation tasks involving contact between the environment and the grasped object. Whitney [1982] has provided a thorough analysis of the mechanical interactions involved in the insertion process, and the stiffness control algorithm permits the object center of compliance to be relocated enabling an active RCC behavior.

This chapter investigates three issues related to effectively modulating the impedance of a grasped object. First, Cartesian stiffness control for multifingered hands is reviewed, and fundamental performance limits of a robot hand to modulate the stiffness of a grasped object is experimentally validated. Second, maximum initial object orientation uncertainty values are calculated, and tactile sensing is used to compensate for this uncertainty thereby improving the performance of the stiffness controller. Finally, a high tolerance peg-insertion task is performed to determine the effectiveness of tactile sensor based rolling, and object stiffness modulation for improving mechanical interactions.

5.2 Cartesian Object Stiffness Control

5.2.1 Controller

From the many proposed schemes for coordinating the control of multifingered robot hands, Salisbury's Cartesian object stiffness controller [Salisbury 1985] was selected for its simplicity and immediate applicability to the peg insertion task. The Cartesian object stiffness controller enables us to specify a center of compliance (CC) and modulate the stiffness of the object about that CC in Cartesian coordinates. Figure 5.1 shows the concept behind this controller where the CC location and the spring constants can be arbitrarily chosen and varied. This controller acts to make the object behave as if linear springs are attached to the specified CC and the ground. If the stiffness matrix is diagonal, the object stiffness is decoupled at the CC, so forces acting through the CC cause only translations, while pure moments only cause rotations about the CC.

By measuring the difference between the desired and actual CC locations, a force on the object is calculated though the stiffness gain,

$$\mathbf{f}_{obj} = \mathbf{K}_{obj}\delta\mathbf{x}_{obj}, \tag{5.1}$$

where $\mathbf{f}_{obj}$ is the net force on the object, which is augmented to include grasp force and object weight, $\mathbf{K}_{obj}$ is the desired stiffness matrix at the CC, and $\delta\mathbf{x}_{obj}$ is the difference between

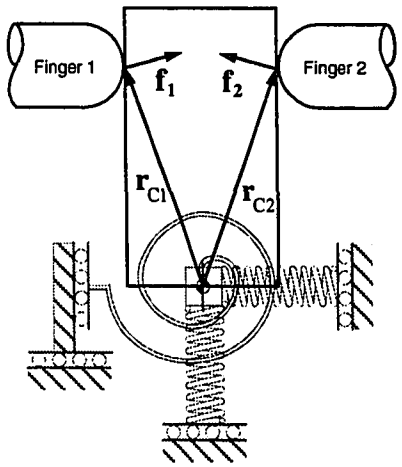

Figure 5.1 The Cartesian object stiffness controller allows the Center of Compliance (CC) to be relocated and the stiffness about that CC modulated.

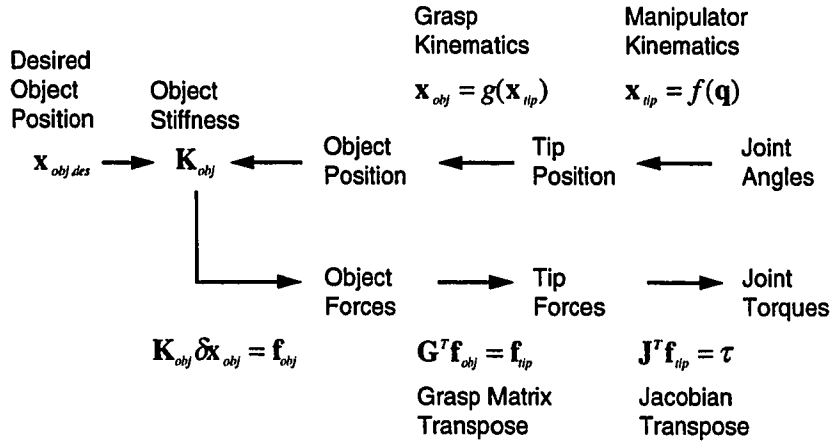

Figure 5.2 In Cartesian object stiffness control, the two components of object position sensing and joint torque generation are coupled through the object stiffness matrix.

the desired and actual CC locations. This object force is mapped to the finger tip forces, $\mathbf{f}_{tip}$, with transpose of the grasp matrix $\mathbf{G}^T$.

$$\mathbf{f}_{tip} = \mathbf{G}^T\mathbf{f}_{obj} \tag{5.2}$$

The finger tip forces are then mapped to joint torques, τ, with the transpose of the Jacobian matrix, $\mathbf{J}^T$, and substituting (5.1) into (5.2) gives us,

$$\tau = \mathbf{J}^T\mathbf{f}_{tip} = \mathbf{J}^T\mathbf{G}^T\mathbf{K}_{obj}\delta\mathbf{x}_{obj}. \tag{5.3}$$

To determine the object position error, we start with measured joint angles, $\mathbf{q}$, then calculate the manipulator forward kinematics, f, and finally the finger tip to object grasp kinematics, g.

$$\mathbf{x}_{obj} = g(f(\mathbf{q})) \tag{5.4}$$

The resulting equation computes motor torques based on object stiffness and position error.

$$\tau = \mathbf{J}^T\mathbf{G}^T\mathbf{K}_{obj}(\mathbf{x}_{obj,des} - g(f(\mathbf{q}))) \tag{5.5}$$

Figure 5.2 shows the above transformation equations in a graphical form.

5.2.2 Limits to Force Generation

One component of the stiffness controller involves generating desired forces on the object. The force that is actually generated in response to object displacements may contain several sources of error. We can explicitly enumerate these errors by combining (5.2) and (5.3) above.

$$\mathbf{f}_{obj,\,act} = \mathbf{G}_{act}^{-T}\mathbf{f}_{tip,\,act} = \mathbf{G}_{act}^{-T}\mathbf{J}_{act}^{-T}\tau_{act} \tag{5.6}$$

where the subscript "act" refers to the actual physical situation, as opposed to the calculated quantities based on joint angle measurements. We can describe the actual torque, τ_{act}, as composed of the desired torque, τ_{calc}, calculated from (5.3) above, plus a disturbance term, $\tau_{disturbance}$, due to joint-level friction, motor torque ripple, D/A quantization, amplifier noise, etc. Combining (5.3) with (5.6) yields,

$$\begin{aligned}\mathbf{f}_{obj,\,act} &= \mathbf{G}_{act}^{-T}\mathbf{J}_{act}^{-T}(\tau_{disturbance} + \tau_{calc}) \\ &= \mathbf{G}_{act}^{-T}\mathbf{J}_{act}^{-T}\tau_{disturbance} + \mathbf{G}_{act}^{-T}\mathbf{J}_{act}^{-T}\mathbf{J}_{calc}^{T}\mathbf{G}_{calc}^{T}\mathbf{K}_{obj}\delta\mathbf{x}_{obj}.\end{aligned} \tag{5.7}$$

If there are no kinematic errors, then $\mathbf{J}_{act} = \mathbf{J}_{calc}$ and $\mathbf{G}_{act} = \mathbf{G}_{calc}$, and the joint torques will be accurately mapped to the object force. The magnitudes of these errors will vary greatly with the details of the grasp and joint configuration; explicit values may be estimated with the above relation.

The most serious source of disturbance torques with present robotic hands is friction, which is a particular problem in this context because it is highly nonlinear and difficult to model especially for complex tendon drive systems used in many robot hands. Proper gains must be used for the torque constants and amplifier gains, and torque ripple and amplifier offset current must be accounted for. Alternatively, joint torques or finger tip forces can be directly sensed, and a controller can compensate, at least partially, for these disturbance torques. Finally, even with a very low friction and properly modeled actuation system, there is a limit to the maximum amplitude and resolution in which the desired torques can be controlled due motor coil heating and to DA resolution.

For the planar manipulator described in section 2.3.2, the peak force generated at the home position is 3.5 N with theoretical control resolution of 0.0017 N, although noise and filtering limit the control resolution. Maximum force error for the planar manipulator is approximately 0.3 N at the finger tip, 0.1 N of this is due to joint friction and the rest are modeling errors.

5.3 Stiffness Limits

5.3.1 Planar Hand Hardware Description

Cartesian object stiffness controller was implemented on the two-fingered planar robot hand shown in figure 5.3, and a full description of the manipulator was given in section 2.3.2. One unique important design feature for these experiments is that the finger tip orientation remain constant due to the manipulator's parallel linkages.

For these stiffness limit experiments, the manipulator was fitted with special finger tips to provide approximate "point contact". The pointed finger tips were made from a 25.4 mm square L shaped aluminum bracket 1.5 mm thick. The aluminum edges were wrapped with a 0.4 mm thick silicone rubber with a radius of 2 mm to provide a high frictional contact surface which was required for reliable manipulation. The finger tip-object contact was modeled as a point neglecting the finger tip radius and the compliance.

The object used in these experiments was a rectangular block fitted with infrared LEDs. An optical tracking device was used to measure the object position. Each LED was turned on in sequence, and the optical tracker gave the position of each LED individually. Electrical noise resulted in 0.1 mm of peak to peak position uncertainty, but multiple samples were averaged to reduce this error even further. The object was perturbed by a probe attached to a two-axis force sensor described in section 2.2.1, and the applied forces along with the resulting object displacements were measured using a separate data acquisition computer.

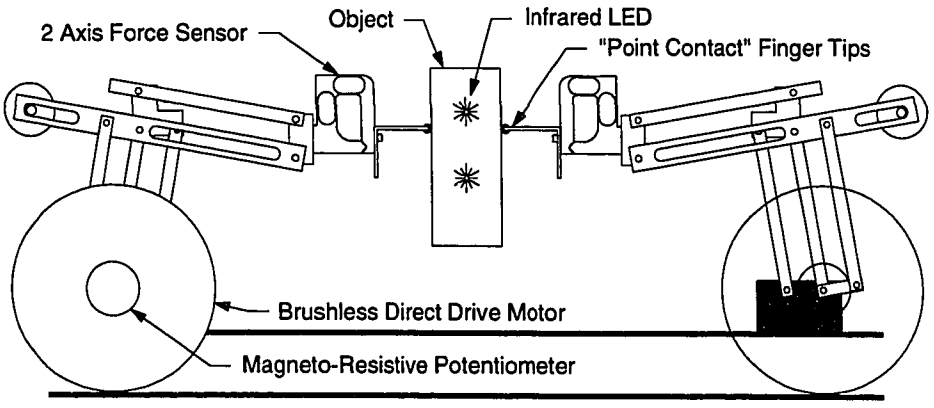

Figure 5.3 Cartesian object stiffness controller was implemented on this planar hand, and object position was measured using an optical tracking device.

5.3.2 Minimum Object Stiffness

The first set of experiments involved measuring the object stiffness and comparing it to the desired stiffness. The purpose of the measurements was to observe the magnitudes of the stiffness errors, and observe the effects of friction on object stiffness. A low object stiffness was commanded and the actual stiffness was measured with two levels of friction. Friction was added to the motor joints by wedging wooden splints between the motor axle and the support frame.

The manipulator was commanded to pick up an object and hold it in space. Forces were applied to the object in the Y direction at the CC using the probe described above, and object displacement was recorded with the optical tracker. The probe was lowered until 1 N was measured on the force sensor, then back up again until 1 N was reached in the other direction followed by moving down again until the force sensor reading was back to zero. Each force-displacement data point was taken after moving the probe and allowing the forces to stabilize to ensure quasi-static stiffness measurements.

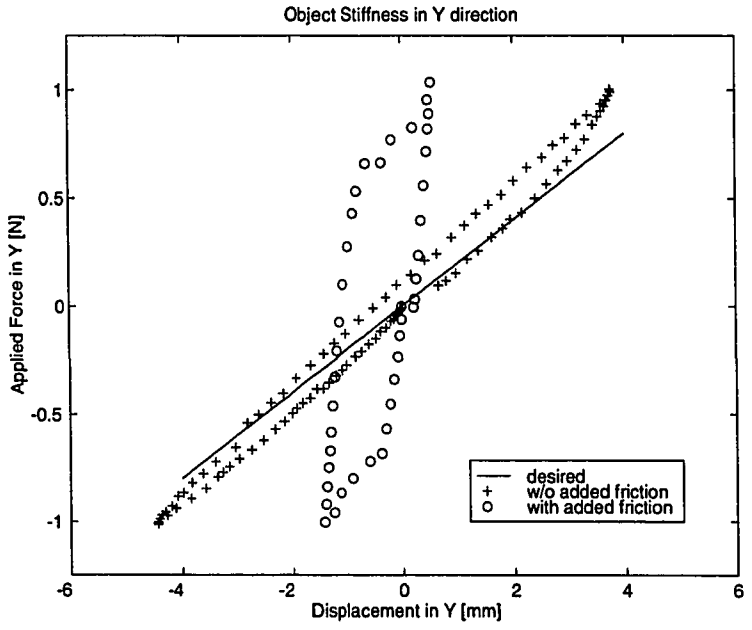

Figure 5.4 Desired and measured object stiffness in the Y direction shows that friction forces must be exceeded before the desired stiffness behavior can be obtained.

The force displacement measurements for the y (vertical) direction of the object are shown in figure 5.4. The object exhibits substantial hysteresis due to the friction associated with the manipulator joints. Without adding friction, the frictional force is approximately 0.1 N, while the wooden splints in the joints increased this friction to 0.8 N. Frictional force was calculated as half the difference between the two force values at a give position.

Since friction prevents motion at the joints until the frictional force is exceeded, the addition of friction effectively increases the stiffness. This high stiffness is observed, rather than a vertical line, because of the compliance in the manipulator finger tips, force sensor and the linkages. The desired stiffness was set to 200 N/m; with and without added friction, the measured stiffness beyond the friction dead bands were 320 N/m and 240 N/m respectively. The stiffness associated with the dead band for the former case is on average 2430 N/m. As shown in the peg insertion experiments below, this high friction results in much higher insertion forces.

In a sense, the lowest obtainable stiffness is limited by the frictional force. Although the desired stiffness of the object can be set to zero, any disturbances which are smaller than the friction force will not result in a restoring force from the controller. Instead, the restoring force will be generated by the finger tip and the manipulator structure.

5.3.3 Maximum Stiffness Gain

To demonstrate that one upper bound on stiffness is set by the finger tips, the stiffness of the object was measured at high stiffness gains. The stiffness measurements indicate that the actual stiffness was much lower than the desired stiffness. For the x axis, the stiffness gain could be set up to 20,000 N/m and remain stable. However, the highest stiffness measured was 6300 N/m for the case with friction. Even using a feed forward approach of commanding object forces opposite and equivalent to the measured force yielded only slightly higher 7300 N/m. With the friction at the joints removed, the stiffness were 1200 and 2100 N/m for the stiffness and feed forward cases respectively.

This is due to the compliance of the rubber covering the finger tips. Although finger tip compliance would decrease at higher forces, this can be a problem for tasks which

require small forces. In general, soft finger tips can be advantageous, as they have high friction, and their passive compliance can increase grasp stability. Trying to use stiffer finger tips caused problems due to low friction coefficient causing the object to slip. Therefore, some tradeoff considerations needs to be made about the finger tip surface material [Cutkosky et al. 1987, Shimoga and Goldenberg 1992].

Next, the stiffness limits due to CC location were considered. With the CC located between the finger tips, the stiffness gain could be adjusted all the way from zero to k_x= 800, k_y=200 N/cm, and k_θ=250 N cm/rad. The damping gain was increased along with the higher stiffness gains, but after reaching a certain threshold, increasing the damping gains reduced the total system stability. The maximum damping values which did not generate oscillations were k_{vx}=30, k_{vy}=2.5 N sec/cm, and $k_{v\theta}$=2.4 N cm sec/rad. This ability to freely adjust the gains became increasingly difficult as the CC moved away from the finger tips.

This observation can be easily understood by considering the example shown in figure 5.5. If an external disturbance force, f_x, is applied through the CC, the object should only translate since the controller makes the grasped object behave as if it was attached to virtual springs at the center of compliance. In order for this to happen, the object displacement must be measured and appropriate forces must be applied to the object. Therefore, finger 2 must increase the grasp force to counteract the disturbance force to obtain an equilibrium position. However, since the object is solely supported by the finger

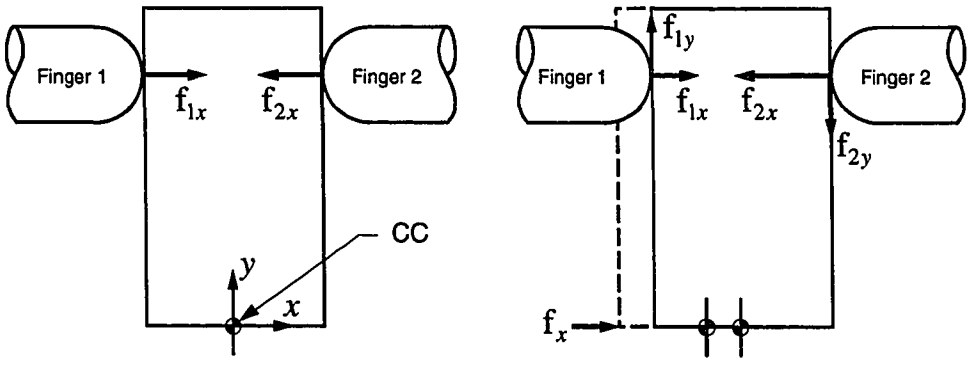

Figure 5.5 This figure demonstrate how the fingers must generate a moment through f_{1y}, f_{2y}, and the object width to balance the moment generated by the disturbance force at the CC and the distance between the CC and the finger tips.

tips, a moment will be generated by the disturbance force and the reaction force at finger 2. To equilibrate this object moment, the controller applies equal and opposite vertical forces at the finger tips such that $-f_{1y} = f_{2y}$. In this particular example, the forces are determined so that a "finger tip gain," k_{tip}, which relate finger-tip-force to finger-tip-displacement can be calculated

$$k_{tip} \equiv \frac{f_{1y}}{dx_1} = -k_{xx}\frac{r_{c1y}}{r_{12x}}. \tag{5.8}$$

Here, dx_1 is the displacement of finger tip 1 in the x direction, k_{xx} is the stiffness matrix component which maps object displacement to forces in the x direction, r_{c1y} is vertical distance from the CC to finger tip 1, and r_{12x} is the distance from finger tip 1 to finger tip 2. This shows that the effective finger tip gain at the finger tip increases in proportion to the distance from the finger tips to the CC and decreases inversely with the distance between the fingers.

To verify the limits of maximum stiffness gain as a function of $\mathbf{r}_{C1}$ and $\mathbf{r}_{1i}$, the manipulator was commanded to hold objects with various width and the stiffness gain, k_x, was increased

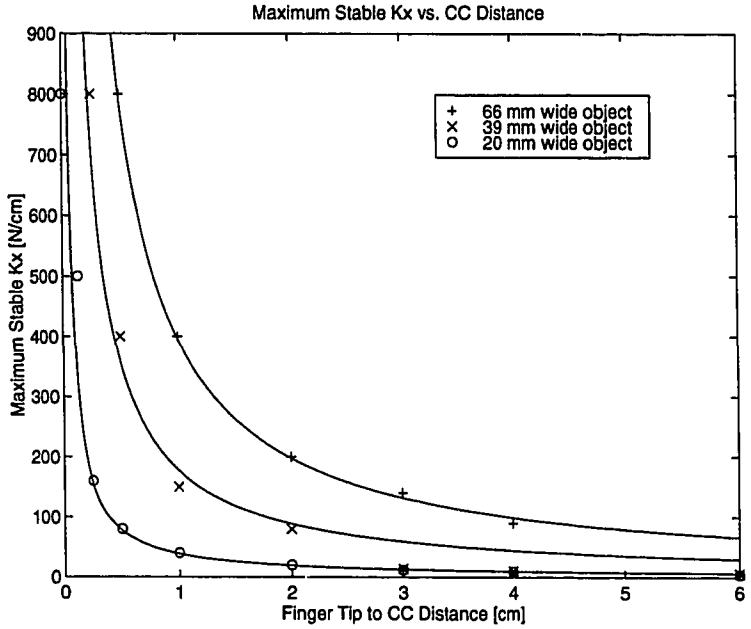

Figure 5.6 This plot of $\mathbf{r}_{c1}$ versus $k_{x,max}$ is for the grasp configuration shown in figure 5.5. It shows that the product must remain constant for a given object width.

until the system went unstable for a range of r_{C1}. The grasp configuration and disturbance force used in the experiments were identical to figure 5.5. Although small high frequency vibrations were allowed, the criteria used for determining stability of a gain was that the measured finger tip velocities must fall below 1 mm per second when the object was held steady. In addition, any small disturbances to the object had to damp out quickly. The maximum gains which met these criteria were recorded. Afterwards, the CC was moved to a new location and the procedure was repeated. Figure 5.6 shows the maximum gains permissible for k_x when r_{C1} is varied in the y direction. Below the maximum stiffness obtainable with the CC between the fingers, the product of k_x and r_{C1} is nearly constant for a particular object width.

This relationship breaks down for large r_{C1}; perhaps due to some other limitations in the controller such as the noise associated with object position values. For example, object position and velocity error in the x direction is now a function of translational error plus the rotational error multiplied by the lever arm r_{C1}.

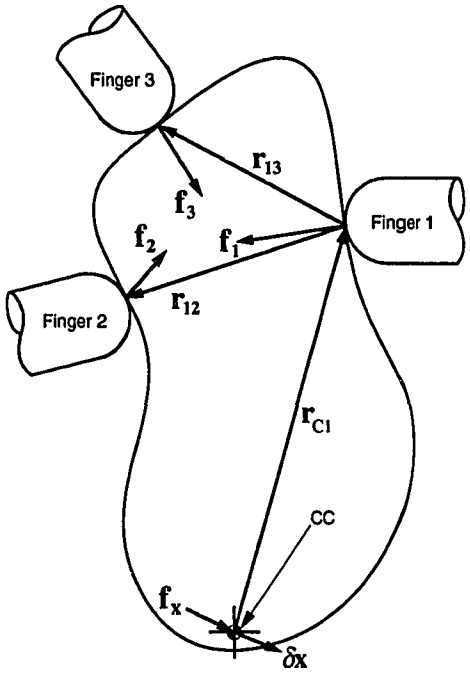

Figure 5.7 Because of the large lever arm associated with the CC in comparison to the finger tip distances, small disturbance at the CC requires large finger tip forces to balance the resulting torque.

To obtain quasi-static equilibrium, object forces and moments needs to equal the forces and moments generated by the finger tip forces. For frictional grasps with more than two fingers, this "finger tip gain" relation cannot be found without specifying additional constraints, since the number of degrees of freedom tends to be larger than the force and moment equilibrium constraints. The location for the moment balance can be chosen arbitrarily, but in figure 5.7 (a generic example of an object being grasped by a multifingered hand), the contact location at finger 1 is selected. The advantage of this selection is that the contact force at finger 1 need not be considered for the moment summation since the lever arm associated with that force is zero. Another way to write the moment balance requires that the moment generated by the force at the CC must equal the moment generated by the finger tip forces. The required force at the CC is determined by the stiffness matrix and the CC displacement, which is related to the finger tip displacement by the grasp matrix. By substituting equation (5.1) into the moment balance, we obtain

$$\mathbf{r}_{C1} \times \mathbf{k}_x \mathbf{G} \delta \mathbf{x}_{tip} = \sum_{i}^{n} \mathbf{r}_{1i} \times \mathbf{f}_i , \tag{5.9}$$

where $\mathbf{r}_{C1}$ is the position vector from the CC to finger 1, $\mathbf{k}_x$ is the translational object stiffness matrix at the CC, $\mathbf{G}$ is the grasp matrix which maps small finger tip displacement to object CC displacement, $\delta \mathbf{x}_{tip}$ is positional displacement vector for the finger tips, $\mathbf{r}_{1i}$ is the position vector from selected finger 1 to the other grasping fingers, $\mathbf{f}_i$ is the force vector exerted by finger i, and finally n is the number of fingers used to grasp the object. Looking at (5.9) we can see that the "effective finger tip gain", defined as the force required at the fingers for a given finger tip displacement $\delta \mathbf{x}_{tip}$ is proportional to $\mathbf{k}_x$, $\mathbf{r}_{C1}$ and inversely proportional to $\mathbf{r}_{1i}$ for a given grasp $\mathbf{G}$.

Similar analysis for the rotational stiffness requirement shows that,

$$\mathbf{k}_q \delta \mathbf{q} = \sum_{i}^{n} \mathbf{r}_{1i} \times \mathbf{f}_i \tag{5.10}$$

where $\mathbf{k}_q$ is the rotational stiffness vector and $\delta \mathbf{q}$ is the rotational displacement vector. In this case, only the distance between the finger tips amplifies the stiffness gains.

5.4 Contact Location Uncertainty

For the previous stiffness experiments, the contact localization problem was minimized through the use of finger tips which provided near ideal "point contact." This is unrealistic for practical hands which have significant finger tip radii. Furthermore, with large finger tips, the object being manipulated can roll against the finger tip surface and the contact location can change. Errors in sensing object finger tip contact location results in incorrect object orientation estimates for the peg insertion task. This in turn can result in significant displacement discrepancy at the peg tip. Use of a tactile array sensors mounted on the larger radius finger tips can provide object orientation information, in addition to improved grasp stability (Figure 5.8).

In this section, contact information is used to improve the performance of object stiffness control with multifingered hands. As shown by equation (5.4), the object location can be determined by measuring joint angles of the multifingered hand and applying several transformations that ultimately provide the object location. Determination of finger tip position based on joint angles is straight forward and accurate. For this particular manipulator, finger tip position noise was about 0.03 mm peak to peak.

However, going from the finger tip location to the contact location is less accurate because of rolling, sliding, and compliance of finger tips, which adds additional uncertainty. This can be a major source of error for extended objects like a peg. Solutions to this problem include modeling the finger-object interaction [Montana 1986; Cole et al. 1989], and using tactile sensing to determine the contact location [Maekawa et al. 1995; Son et al. 1996].

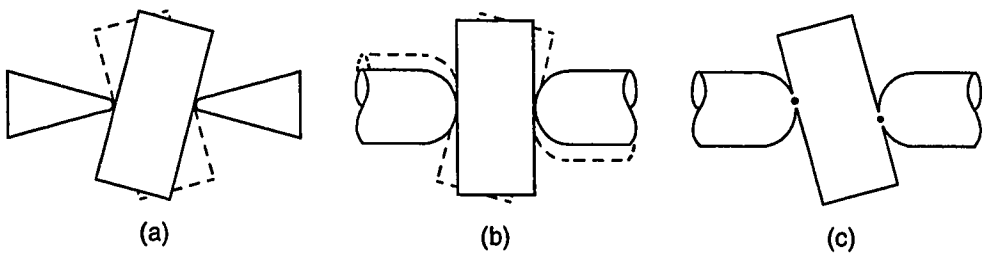

Figure 5.8 (a) With pointed finger tips, the object orientation is unknown within the friction cone. (b) With rounded finger tips, the contact location changes during manipulation. (c) Tactile sensing can provide contact location information to better determine the object position and orientation.

It should be noted that, although force generated errors can be reduced with higher stiffness gains, the controller cannot compensate for contact kinematic errors by increasing object stiffness. If the controller thinks that the desired position has been obtained due to an incorrect mapping of the hand-object relation, no compensating forces will be generated. This implies that the object finger tip interaction needs to be taken into consideration either through modeling or through sensing to improve manipulation dexterity.

5.4.1 Orientation Uncertainty from Friction

Uncertainty in object orientation after grasping can arise from friction between the fingers and the object. To illustrate this uncertainty, consider an unconstrained object with parallel flat sides at a skewed orientation as two fingers come together to grasp it. One assumption here is that during this initial approach phase, the finger tips are controlled with high stiffness in the vertical direction which minimize deflections. If there is no friction, the object will slide against the fingers and reorient so that its surfaces are perpendicular to the line between the finger tip centers. The object orientation may then be inferred from the positions of the finger tips, i.e. from joint angle measurements.

However, if friction is present, sliding will stop before this condition is reached. This static equilibrium occurs when the grasp force, which is directed between the finger-object contact points, falls within the friction cones centered about the object's surface normals. For "point contact" finger tips whose radius is small compared with the size of the object,

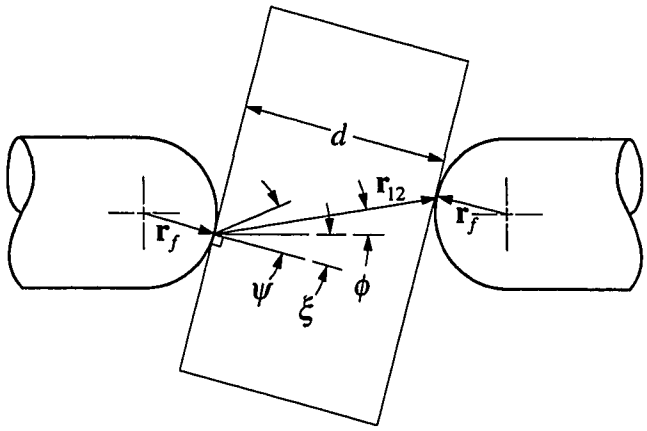

Figure 5.9 Variables used in calculating the object rotational uncertainty for cylinderical finger tips and a parallel-sided object.

the range of stable object orientations is determined by the friction cone angle, $\psi = \tan^{-1}\mu$, where μ is the coefficient of friction. For fingers with finite radius, the limiting equilibrium angle of the object also depends on the ratio of object-width to the finger-tip-radius along with the friction coefficient.

For cylindrical finger tips grasping a parallel, flat sided object, figure 5.9 shows the situation in the plane containing the finger tip centers and the finger-object contact points. For grasp stability

$$\psi \geq \phi + \xi \tag{5.11}$$

where ξ is the object orientation angle and ϕ is the grasp force angle, both defined with respect to the horizontal. The relationship between object orientation and grasp force orientation is

$$\phi = \sin^{-1}\left(\frac{2r_f \sin\xi}{r_{12}}\right) \tag{5.12}$$

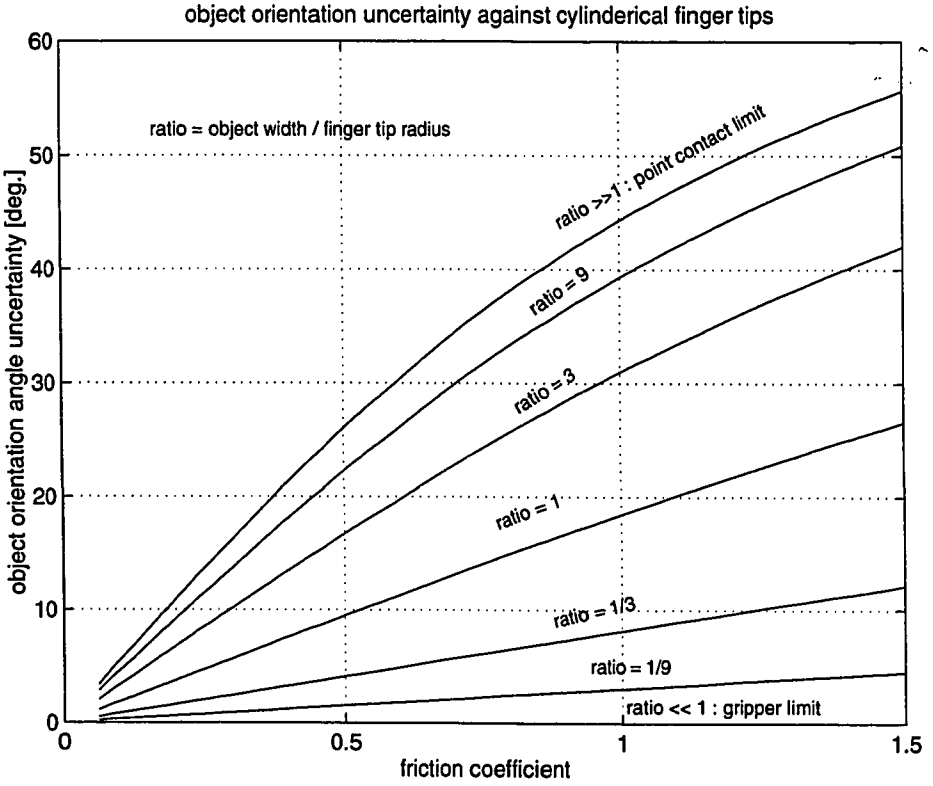

Figure 5.10 Object orientation uncertainty as a function of the friction coefficient for a variety of object-width to the finger-tip-radius ratios.

where the distance between the contact points $r_{12} = \left[\left((d + 2r_f)\tan\xi\right)^2 + d^2\right]^{1/2}$ is in the plane of the grasp vector, r_f is the finger tip radius, and d is the object width. This expression is immediately applicable to the peg insertion experiments described below; extension to objects with nonparallel sides is accomplished by considering when the grasp force vectors at each contact point pass into the friction cones.

Figure 5.10 shows the maximum object orientation angles calculated from these relationships as a function of the friction coefficient for various values of the object-width to finger-tip-radius ratio. For pointed finger tips ($\mathbf{r}_f << d$) the orientation uncertainty is simply the friction cone angle, and for parallel-jaw grippers with flat fingers ($\mathbf{r}_f >> d$), the uncertainty is zero. In a typical dexterous manipulation task, a robot hand might manipulate objects about as wide as the finger tips ($\mathbf{r}_f \approx d$) with a coefficient of friction of about 1, resulting in maximum object orientation uncertainty of ±15–20 degrees.

5.4.2 Kinematic Errors and Tactile Sensing

Tactile sensing can limit kinematic errors in stiffness control in two ways. The first is by eliminating uncertainty in initial object orientation as demonstrated in chapter 4. By detecting the contact location on the finger tips as the object is grasped, the object orientation can be determined. In some cases it may be possible to infer the global object shape and orientation without an a priori model, by using tactile sensing to find the local curvature of the object and then extrapolating the overall object shape [Fearing 1990a].

However, in many tasks which require object stiffness control, some knowledge of the object will be available; we need to know something about the grasped object in order to know what to do with it. This is particularly true for tasks that involve tool use or parts mating, where stiffness control is most useful. In the experimental peg insertion task presented below, the controller implicitly assumes that the object is parallel-sided, and so for the cylindrical finger tips, the sensed contact locations provide the object orientation information.

The second way that tactile sensing can limit kinematic errors in stiffness control is by tracking changes in the location of contacts between the finger tips and the grasped object. Tactile sensing is not the only means of tracking these changes. If a mathematical model of the finger tip and the object is available, finger joint angle measurements can be used to infer the object velocities [Montana 1986]. This approach is reasonable for simple objects, but determining a mathematical model of an irregular object may pose computational challenges. Another requirement which limits the application of Montana's approach is that the object surfaces must be smooth and differentiable. In addition, these model-based schemes require initial values for the contact locations and cannot detect object slips during manipulation.

Maekawa et al. [1995] introduced a scheme that eliminates the need for a mathematical model of the object by using tactile sensor-based contact localization and integrating the sensed displacements of contact location. The advantage of this approach is that it is straightforward to manipulate around sharp corners, and a mathematical model of the object is not required. Once again, however, some model of object shape must be available for many dexterous tasks. Next, the effects of contact location uncertainty on object stiffness control is examined.

5.4.3 Object Stiffness Errors

Here, a relation that describes the object stiffness errors caused by inaccurate determination of hand-object relationship is presented. In particular, contact location errors will degrade the restoring force and object velocity transformations, and incorrect object stiffness will result. Disturbance torques such as friction, motor torque ripple, and amplifier gain variation may sometimes be modeled or directly sensed and corrected, and improved manipulator designs can help limit these disturbances. However, incorrect determination of contact locations cannot be corrected in this manner for control algorithms which use joint angles to close the feedback loop.

If we neglect disturbance torques in (5.7), the force on the object reduces to

$$\mathbf{f}_{obj,\,act} = \mathbf{G}_{act}^{-T}\mathbf{J}_{act}^{-T}\mathbf{J}_{cal}^{T}\mathbf{G}_{cal}^{T}\mathbf{K}_{obj}\delta\mathbf{x}_{obj}. \tag{5.13}$$

However, equation (5.13) does not take into consideration that the actual object displacement calculated will also be in error due to errors in determining the hand-object relation. As described in figure 5.2 and equation 5.4, direct kinematic calculations were used in the controller implementation. However, if the displacements are small, it is possible to formulate the controller in terms of infinitesimal changes in the joint angles, $\delta\theta \equiv \theta_{desired} - \theta_{act}$, [Salisbury 1985]. If this approach is used to map small displacements in joints to object displacements, equation (5.13) can easily be written to include the difference between the calculated object displacement, $\delta\mathbf{x}_{cal} = \mathbf{G}_{cal}\mathbf{J}_{cal}\delta\theta$, to the actual displacement, $\delta\mathbf{x}_{act} = \mathbf{G}_{act}\mathbf{J}_{act}\delta\theta$, such that,

$$\mathbf{f}_{obj,\,act} = \mathbf{G}_{act}^{-T}\mathbf{J}_{act}^{-T}\mathbf{J}_{cal}^{T}\mathbf{G}_{cal}^{T}\mathbf{K}_{obj}\mathbf{G}_{cal}\mathbf{J}_{cal}\mathbf{J}_{act}^{-1}\mathbf{G}_{act}^{-1}\delta\mathbf{x}_{act} \tag{5.14}$$

where the subscript "act" refers to the "actual" and subscript "cal" refers to the "calculated" transformations. We can define a stiffness error matrix, $\mathbf{E}$, which map the actual stiffness matrix, $\mathbf{K}_{act}$, from the desired stiffness matrix, $\mathbf{K}_{obj}$, as

$$\mathbf{E} = \mathbf{G}_{act}^{-T}\mathbf{J}_{act}^{-T}\mathbf{J}_{cal}^{T}\mathbf{G}_{cal}^{T}. \tag{5.15}$$

Substituting equation (5.15) into equation (5.14), the actual object stiffness, $\mathbf{K}_{act}$, is given by,

$$\mathbf{K}_{act} = \mathbf{E}\mathbf{K}_{obj}\mathbf{E}^{T} \tag{5.16}$$

The $\mathbf{E}$ matrix provides a measure of capturing object stiffness deviations generated by improper mapping of the hand-object relation. If the kinematics are correctly described by the $\mathbf{G}_{calc}^{T}$ and $\mathbf{J}_{calc}^{T}$ matrices, then $\mathbf{E}$ will be an identity matrix and the actual stiffness matrix will agree with the desired one. The divergence from an identity matrix describes the effects of the kinematic errors, independent of the particular $\mathbf{K}_{obj}$ values. The actual stiffness matrix can then be found by pre-multiplying any specified desired stiffness matrix $\mathbf{K}_{obj}$ by $\mathbf{E}$ and post multiplying by $\mathbf{E}^{T}$.

The calculation of $\mathbf{E}$ for a particular example will demonstrate the nature of the resulting stiffness errors. Since this manipulator works in a plane, the object forces augmented by the grasp force, f_g, and displacements are $\mathbf{f}_{obj} = [f_x, f_y, \tau, f_g]^T$ and $\delta\mathbf{x}_{obj} = [x, y, \theta, 1]^T$ respectively, and $\mathbf{K}_{obj}, \mathbf{E} \in \Re^{4\times4}$. Two planar fingers, with radius of 1.5 cm, are grasping a parallel flat sided object of width 2 cm, as shown in figure 5.9. The object stiffness controller is commanded to place the CC symmetrically between the finger tips. In the absence of tactile information about the actual orientation of the object, the controller must assume $\xi = \phi = 0$. If the actual object orientation is 5 degrees, then the difference between this assumption and the actual contact locations results in a stiffness error matrix which is given by,

$$\mathbf{E} = \begin{bmatrix} 1.000 & 0 & 0 & 0 \\ 0 & 1.000 & 0 & 0 \\ 0.131 & -0.006 & 1.000 & 0 \\ 0 & 0.044 & -0.044 & 0.996 \end{bmatrix}.$$

Therefore, if the desired object stiffness matrix was an identity matrix, then the actual stiffness matrix for this particular case is,

$$\mathbf{K}_{actual} = \begin{bmatrix} 1.000 & 0 & 0.131 & 0 \\ 0 & 1.000 & -0.006 & 0.044 \\ 0.131 & -0.006 & 1.000 & -0.044 \\ 0 & 0.044 & -0.044 & 0.996 \end{bmatrix}.$$

The off diagonal terms results in coupling of stiffness between the different axes. This coupling phenomena is mostly a result of inaccurately modeling the contact kinematics. If the manipulator exhibited backlash and the linkages were compliant, then the stiffness error matrix would also capture these errors as well since the Jacobian errors are also included. However, manipulator kinematics is a well understood problem and most robots have great positioning capabilities in comparison to determining the hand-object relation which is difficult to obtain without tactile sensing.

5.4.4 Compensating for Object Orientation

To demonstrate the benefits of tactile sensing for stiffness control, tactile array sensors were mounted on cylindrical finger tips for the planar manipulator as shown in figure 5.11. The tactile array sensors used here were very similar to the ones described in chapter 3.

The first experiment demonstrates the challenges in determining initial object orientation of an object which is disturbed during the grasping process. The robot is commanded to grasp the peg, lift it, and rotate it in the vertical plane symmetric about the vertical axis. Because the peg is arbitrarily placed between the fingers, one finger makes contact with the peg before the other. The base width of the pin is small compared with the height of the contact, so it can be tipped with less than 0.2 N of force. Although the finger tip force sensors are used to monitor contact forces in order to stop the finger motion until the other finger makes contact, the manipulator cannot reverse directions fast enough to prevent the peg from being tipped over.

The task is first performed without tactile sensing. Figure 5.12 shows the desired peg orientation, the controller's calculated peg orientation, and the actual peg orientation as measured by an optical tracker. Since tactile sensing was not used for this trial, the

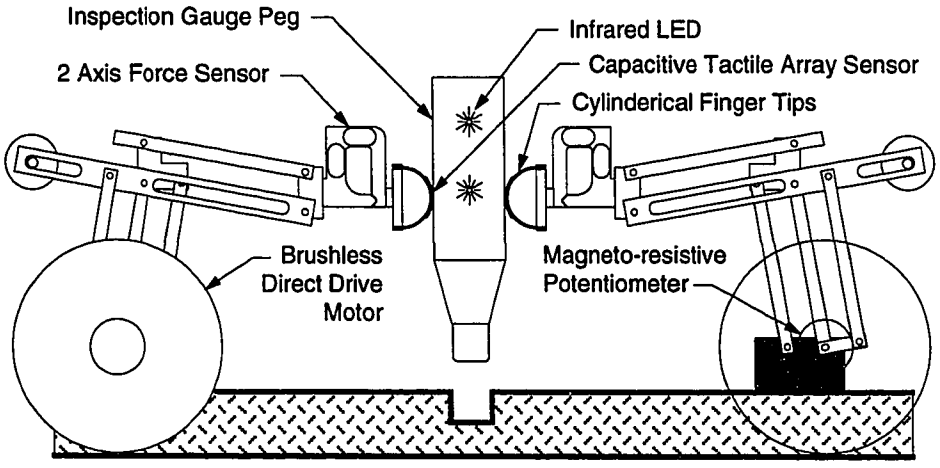

Figure 5.11 Two-fingered planar hand is now fitted with cylinderical finger tips with tactile array sensors. As before, the peg position and orientation is measured by tracking LEDs embedded in the peg with an optical position sensing device.

controller's peg orientation was calculated based on the assumptions that the object has flat parallel sides with a known object width, and that contact occurs at $\phi = 0$.

The peg is tipped against the other finger, as indicated by event 1 in the figure. The initial peg orientation of 14 degrees from vertical is reduced to about 6.5 degrees as the fingers close and the grasp force increases. This object orientation is stable after the grasp force enters the friction cones between the fingers and the object. The observed error is within the maximum orientation angle error of 10° predicted by figure 5.10. However, due to the absence of tactile feedback, the controller is unaware of this error, and it persists throughout the remainder of the task as the object is lifted and rotated. Increasing the controller gain does not reduce this kinematic error since the object orientation error calculated by the controller is small.

Figure 5.13 shows the same task, but here the orientation of the peg is determined from the tactile array sensor. Again, the peg is tipped before it is grasped, resulting in a significant deviation from the desired orientation. However, since the correct contact locations are determined through tactile sensing, the desired object orientation is set to the initial orientation of the object while the grasp force is increased. This permits gentle grasping

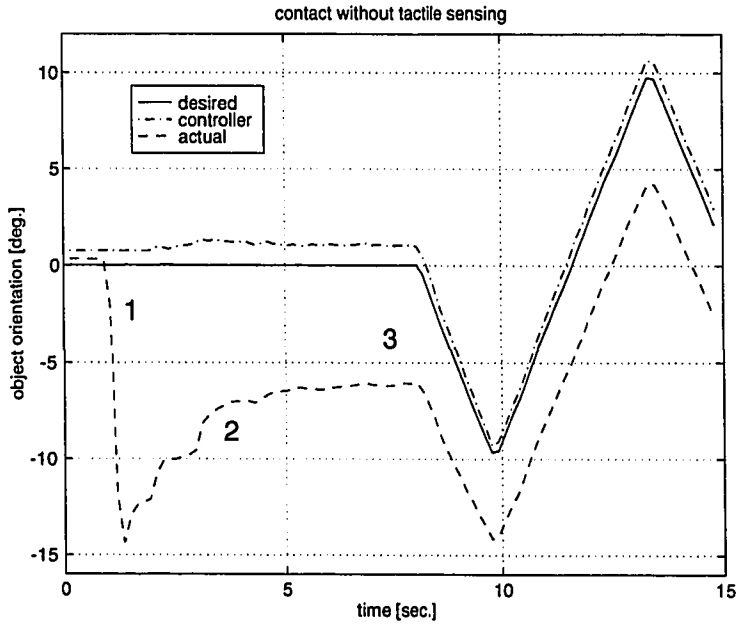

Figure 5.12 Peg orientation during grasp-lift-rotate task without tactile sensing. Events: (1) peg is tipped when one finger makes contact first; (2) actual and controller calculated peg orientation differ after lifting; (3) error persists due to lack of tactile sensing.

that minimizes disturbances to the object. Now through tactile sensing, the controller is able to compensate for this orientation error before performing the desired rotational trajectory.

An alternative approach is to set the desired object orientation angle to zero when switching from finger tip stiffness control to Cartesian object stiffness control. This causes the object to rotate immediately when the controller is switched because of the large difference between the desired and the actual object orientation angle. This approach may not be desired when handling fragile or delicate objects, but it might be more desirable if manipulation speed is an important criteria.

Since the correct object orientation angle has been determined through tactile sensing, the difference between the controller and the actual peg orientation angle is minimized. As a result, object orientation can be accurately determined, and the object stiffness error is also reduced. One tradeoff, which might be a general issue for using tactile array information, is the difference between the manipulator servo rate and the tactile array sample speed. In general, high resolution tactile array sensors will require more time to sample. However, integrating slow sensory data to a fast controller can reduce the system

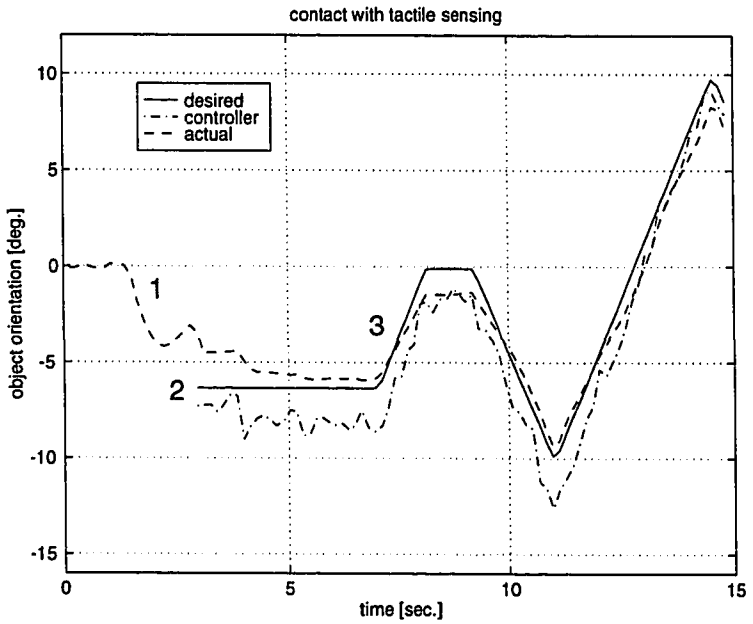

Figure 5.13 Peg orientation during grasp-lift-rotate task with tactile sensing. Events: (1,2) Peg is tipped again, but (3) tactile signals permit controller to roll the peg near the desired orientation.

gain margin. In this particular implementation, the manipulator servo rate was 500 Hz while the array scanning rate was 25 Hz. This reduced the maximum stiffness gain obtainable while maintaining controller stability. In the previous chapter, controller stability was not a problem since the robot used fast dedicated servo motor controllers with a slower higher level controller which accepted desired trajectories.

5.5 Peg Insertion

Here, analysis on the minimum stiffness required for a peg insertion task is presented followed by expected maximum insertion forces. Experimental execution of a peg insertion task examines object interaction forces with different levels of manipulator friction. Finally, tactile array sensor's ability to compensate for unstructured conditions is used to successfully perform a peg insertion task after a human subject hands the peg to the planar manipulator.

The variables used in this section follows [Whitney 1982] since the analysis uses his results and extend it. Whitney provides a planar analysis of peg insertion task which segments the problem into four phases of approach, chamfer crossing, one point contact, and two point contact. He examines the required positions and forces, and develop conditions which must be satisfied to prevent jamming and wedging.

5.5.1 Stiffness Requirements

To successfully execute the task of inserting a peg into a hole, the position of the peg must be controlled within the clearance and chamfer width during the approach phase of the manipulation. The amount of positional error at the tip of the peg is due to lateral translation error plus the rotation error multiplied by the distance between the object center of compliance and the peg tip. This relationship is given by Whitney [1982] as,

$$cR + w > L_g\theta_o + \varepsilon_o \tag{5.17}$$

where w is the chamfer width, cR is the clearance between the peg and the hole, L_g is the distance between the object center of compliance and the tip of the peg, and ε_o and θ_o is the lateral and rotational errors at the center of compliance (Figure 5.14).

There are two components to the position error, ε_o and θ_o. The first component is the accuracy limit of the sensing scheme in determining the object position and orientation. The second component is limited by the manipulator's ability to position the object at the desired position. The latter is dependent on manipulator friction and controller stiffness. Although mapping manipulator friction to object friction is dependent on the grasp configuration, if we define friction observed at the object as f_μ and τ_μ, the initial position error can be expressed as,

$$\varepsilon_o = \varepsilon_a + \frac{f_\mu}{k_x}, \text{ and } \theta_o = \theta_a + \frac{\tau_\mu}{k_\theta} \tag{5.18}$$

where ε_a and θ_a are errors in determining the object position and orientation, and k_x and k_θ are desired object stiffness used in the controller.

Positional accuracy of a grasped object can be improved by increasing the controller stiffness gains, but this results in larger interaction forces that inhibit task execution. Furthermore, increasing the controller stiffness gains do not compensate for errors caused by inaccurately determining the object position. However, decreasing the error in determining object position allows a lower controller stiffness to be used for a given task requirement. Tactile sensing can improve the ability to determine object position by providing accurate contact kinematic information.

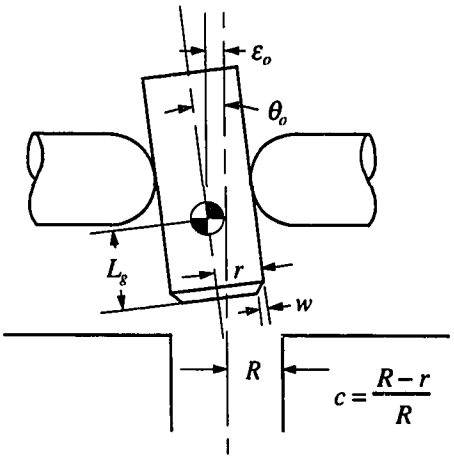

Figure 5.14 Variables used in calculating the required object stiffness and expected maximum insertion forces is the same notation used by Whitney [1982].

If the object position was controlled perfectly, performing assembly tasks such as precision insertion would not pose any problems. For any given clearance between the peg and the hole, insertion without contact is possible if the position of the peg could be controlled within the tolerance of the peg and hole clearance. However, positional errors are bound to be present in real systems which results in interaction forces. Interaction forces are dependent on the mechanical impedance of the objects. Section 5.3.2 presented results which showed that the lower limit to actual object stiffness, the quasi-static component of impedance, is set by manipulator friction. Therefore, reducing and compensating for manipulator friction are important for reducing actual object stiffness and minimizing interaction forces.

These observations can be translated to required controller stiffness by substituting equations (5.18) into (5.17). This gives us the minimum object stiffness required to successfully start the insertion task based on the criteria that the chamfer must hit the hole edge.

$$k_x > \frac{f_\mu}{cR + w - L_g\theta_o - \varepsilon_a} \tag{5.19}$$

$$k_\theta > \frac{\tau_\mu L_g}{cR + w - L_g\theta_a - \varepsilon_o} \tag{5.20}$$

This tells us that the required object stiffness increases linearly with manipulator friction as defined by f_μ and τ_μ. Errors in object position also increase the required object stiffness by making the denominator terms approach zero. If the errors are larger than the specified clearance and the chamfer width, the denominator terms become negative and the insertion will fail. Therefore, the accuracy in determining object position needs to be increased and manipulator friction needs to be reduced.

5.5.2 Insertion Forces

To prevent jamming or wedging the peg, Whitney [1982] has calculated limits to position and orientation requirements in addition to applied forces and moments for the given task parameters such as peg hole clearance and friction coefficient. This is based on physical

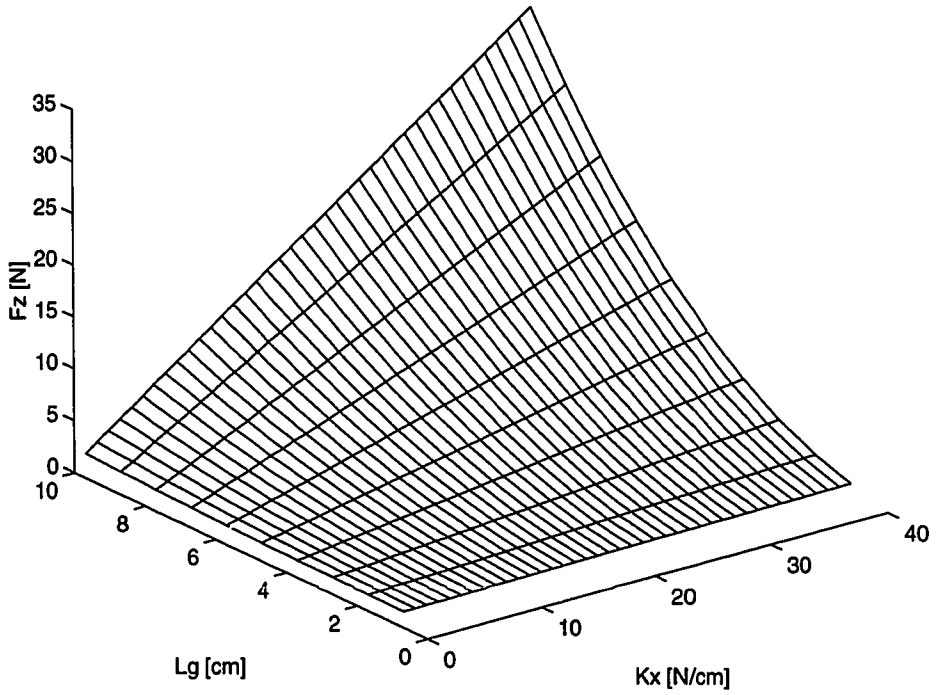

Figure 5.15 Maximum peg insertion force during two point contact peg insertion. The object stiffness and the object center of compliance were varied for the following condition. Rotational stiffness was 100 N cm/rad, initial angular error was 1°, initial positional error was 0.01 cm, friction coefficient was 0.1 and the peg and hole radius were 0.6325 and 0.6350 cm respectively.

analysis of the mechanics involved in the task. Whitney shows that insertion force increases with lateral stiffness, position error, fiction coefficient and the distance between the center of compliance and the peg-hole contact point.

Whitney's analysis of the peg insertion mechanics reveals that the insertion force during the two point contact phase of the insertion task is

$$F_z = \frac{2\mu}{l}\left[\left(k_x L_g^2 + k_\theta\right)G + E\right] + \mu\left(\frac{1+\mu 2r}{l}\right)\left[-k_x L_g G - \frac{E}{L_g}\right]. \tag{5.21}$$

Variable definitions used here are, μ which is the friction coefficient, r and R which are the radius of the peg and the hole respectively, $E = k_x L_g \varepsilon_o''$, $G = \theta_o - 2(R-r)/l$, and $\varepsilon_o'' = \varepsilon_o + R - r$. Whitney also explains how passive remote center of compliance (RCC) devices can perform critical assembly tasks with reduced forces. An RCC device attempts to minimize the distance between object CC and the peg-hole contact point to maintain one point contact

and reduce the insertion force. By placing the CC of the object near the tip of the peg through stiffness control, similar benefits can be obtained.

Differentiating equation (5.21) with respect to the insertion depth and solving for zero, we can obtain the maximum insertion force for a given initial condition. By varying the lateral stiffness, k_x, and the CC distance i.e., L_g, we can plot maximum insertion force during two point contact as a function of those two parameters. Figure 5.15 shows that the insertion force decreases with smaller L_g, but also with stiffness k_x. However, previous results show that the actual minimum object stiffness is set by the manipulator friction. If it was possible for the object to have zero stiffness, the insertion forces would approach zero for the quasi-static insertion. Furthermore, the ability to specify the CC location is limited by the "finger tip gain" as described in section 5.3.3.

Alternatively, insertion forces can be reduced by reducing object stiffness provided it doesn't increase the positional errors, this can be achieved with tactile sensing. Alternatively, object stiffness could be set to a high value for the initial approach phase and lowered gradually during insertion. This method of modulating the controller gains is suggested by Brockett [1988].

Now, a close-tolerance peg insertion task is performed with the planar hand. Since the manipulator motions are constrained in the plane, a slot oriented out of the plane is used instead of a hole. The peg is a precision gauge pin used for inspecting machined parts and the slot is constructed from machinist's parallels. To permit independent measurement of the peg position and orientation, the optical tracker described above was used to monitor the position and orientation of the peg handle.

To demonstrate the usefulness of controlling the CC and the actual object stiffness, peg insertion task is executed by varying both the CC and the manipulator friction at the joints. During this experiment, the near ideal "point contact" finger tips were used. The 12.7 mm diameter peg is chamfered 0.5 mm at 45°, and it is mounted in a 2 cm wide handle with flat parallel sides. The slot is assembled using differing height machinist parallels which provide high accuracy and uniformity; slot to peg clearance is less than 0.01 mm.

The CC was located between the finger tips and at the tip of the peg which was 45 mm away. Friction forces at the finger tips measured before and after adding friction at the joints were 0.1 and 0.4 N respectively. Since the chamfer width is only 0.5 mm for the peg, the stiffness gain required, for the case with added friction, was relatively high to obtain the positional accuracy required to make it into the slot. The gains used on all trials were $k_x = 4.0$, $k_y = 10$ N/cm, $k_\theta = 100$ N cm/rad, $k_{vx} = 0.8$, $k_{vy} = 1$ N s/cm, and $k_{v\theta} = 2$ N cm s/rad.

Since the "point contact" finger tips were used, the manipulator was commanded to grasp the object while the peg was already in the hole. After making contact with both fingers, the actual object position was set as the desired object position in the x direction. After lifting the peg out of the hole, the manipulator moved the peg 5 mm to one direction then back to the slot before attempting to insert the peg. The purpose of this additional motion was to stay on one side of the friction dead band on a consistent basis. Object position and tip forces were recorded at 100 Hz during the insertion task.

Without the added friction, insertion forces were minimal for both cases with and without active RCC as shown by figure 5.16. Here, the actual object stiffness was close to the specified lateral stiffness due to the low levels of manipulator friction. Even when tilting the slot over 4.6 deg., the insertion was successful for both cases with minimal forces.

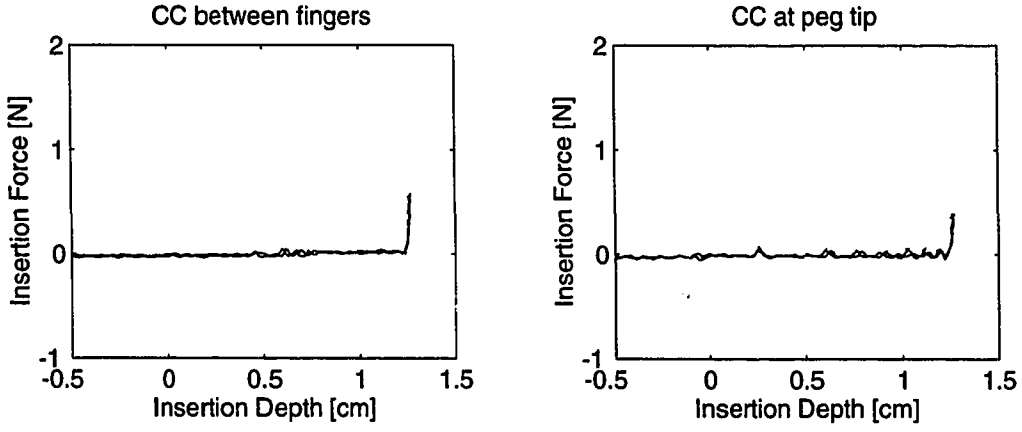

Figure 5.16 Peg insertion forces with the center of compliance (CC) placed between the fingers and at the tip of the peg. Insertion forces are low for both cases because low manipulator friction provided actual object stiffness close to the desired low stiffness and high positional accuracy.

Since the stiffness in the x direction was so low, the insertion force reduction due to placing the CC at the tip of the peg could not be measured in the presence of disturbance noises.

Figure 5.17 shows a typical trial with and without friction both placing the CC at the tip of the peg. Adding friction to the joints increased the insertion force by an order of magnitude. Part of the increase in force is due to an increase in positional error, but this increase was only three times as large compared to the case without friction added. Therefore, the additional increase in insertion force is due to the higher actual stiffness of the object caused by the manipulator friction.

5.5.3 Human Interaction Uncertainty

In this experiment, one of the experimenter hands the peg to the manipulator shown in figure 5.11, and it inserts the peg into the slot. The orientation angle is not controlled as the peg is handed off, so the robot first needs to sense and compensate for this orientation angle before attempting the insertion. For the experimental setup, six degrees of object orientation error results in a tip positional error of 4.7 mm which far exceeds the chamfer width of 0.5 mm. Thus, compensating for this orientation uncertainty is critical to the success of a peg insertion task. Figure 5.18 shows a photograph series of the manipulator accepting a peg from a human hand, determining the orientation of the peg, compensating for the initial orientation error, moving the peg to the desired location, and successfully inserting the peg into the slot.

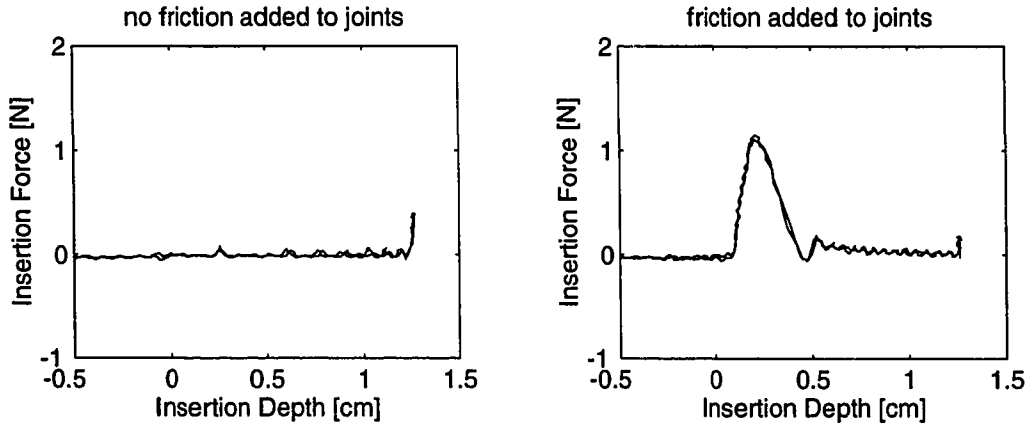

Figure 5.17 Peg insertion force with and without the addition of friction at the joints. Both cases placed the CC at the tip of the peg.

Due to limitations in the tactile array scanning electronics and the time required for sensor signal processing, the rate for determining the contact locations was limited to 25 Hz. When this contact centroid information was directly incorporated into the controller, the time delay in the centroid information reduced the system bandwidth, so the stiffness gains could not be increased to the level possible without tactile feedback.

It is possible to increase the array sensor sampling rate and decrease the signal processing time; an improved version of the sensor system should attain 200 Hz update rates. However, there is a fundamental tradeoff between increasing spatial resolution (i.e. number of array elements) and increasing update rates, which implies that techniques for incorporating slow sensory data into real-time controllers will be required for hands with large, high resolution tactile sensors.

5.6 Conclusions

This chapter looked at limitations in controlling object stiffness and its impact on a peg insertion assembly task. From the analysis and the experiments conducted in this chapter, the following conclusions are evident.

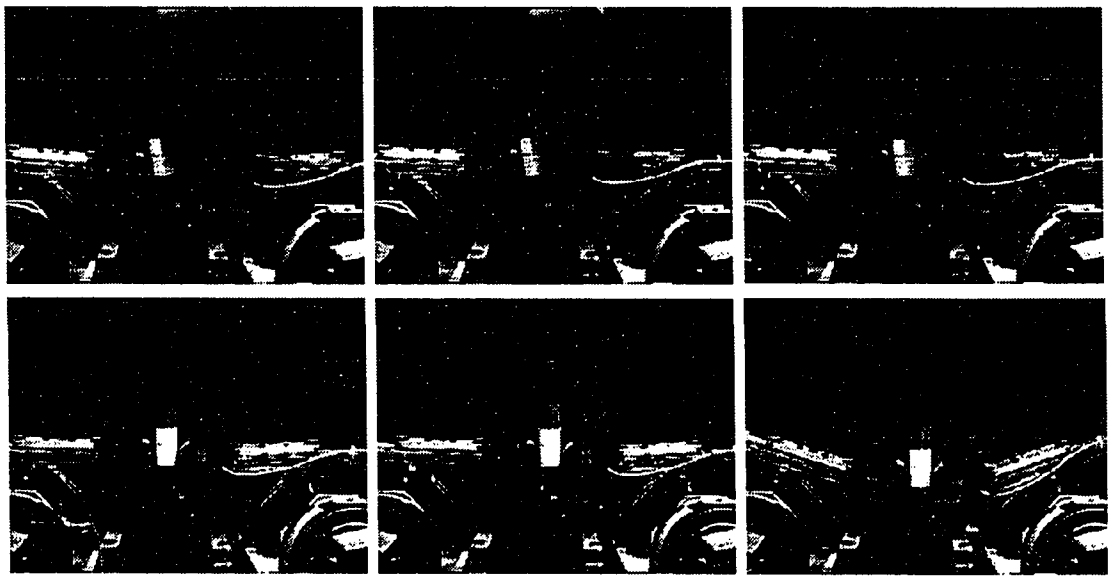

Figure 5.18 A sequence of photographs shows that a peg can be handed to the manipulator at an arbitrary angle, and the manipulator is able to sense the object orientation angle, compensate for it and perform the peg insertion task successfully.

Because stiffness control works by measuring object displacement and generating restoring forces, performance can be separated into position sensing and force generation processes. Position sensing is limited by the ability to determine the contact location since the object can roll and slide against the finger tip. Force sensing is also limited by this uncertainty, and by disturbance torques such as friction at the joints of the fingers. Increasing stiffness gains can compensate for errors in the force generation components of the stiffness controller, but hand-object relationship errors cannot be corrected in this manner. However, tactile sensing can provide the additional information needed for correcting this uncertainty, and allow the use of lower object stiffness.

The lower limit on controlling object stiffness with current hands is friction; this increases "effective stiffness" until the frictional forces are exceeded. The maximum stiffness is determined by finger tip compliance, which is needed for reliable manipulation, and the manipulator structural stiffness. Another factor which limit maximum stiffness is the "effective finger tip gain." For static equilibrium, the summation of moments generated by disturbance and finger tip forces should equal to zero. However, as $\mathbf{r}_{ci}$ becomes large, the torque needed to maintain static equilibrium increases in proportion. In addition, if the distances between the fingers are small, the ability to generate torques is reduced. Thus, "effective finger tip gain", as defined by the required finger tip forces for a given finger tip displacement, increases in proportion to $\mathbf{r}_{ci}$, and is inversely proportional to $\mathbf{r}_{1i}$. This limits how far the CC can be placed for a given grasp.

In an unstructured environment, the orientation of an object after grasping can have considerable uncertainty. Under realistic manipulation conditions, analysis and experimental measurements show that these errors can pose significant limits to task execution. Imprecise information about the location of the finger-object contact also creates errors in the object stiffness. The stiffness error matrix, (5.15), permits calculation of object stiffness errors caused by grasp kinematic errors. The stiffness error matrix results in coupling terms that generate inappropriate forces.

For typical dextrous hands, determining contact location is important since large contact areas provide more stable grasps despite the added uncertainty in determining the contact location. As the object-width to finger-tip-radius ratio increase, the object gains mobility within a grasp at the expense of additional uncertainty. Tactile sensing can provide contact location and object orientation information and improve the performance of object stiffness control for multifingered hands.

Experimental execution of the peg insertion task confirm that lowering lateral object stiffness reduces insertion forces, and the limitation placed on obtaining the lowest object stiffness is set by manipulator friction. The second peg insertion experiment demonstrated that tactile sensing can effectively determine the initial orientation of the object after grasping and be used to track object rolling in real-time control.

6. Conclusions and Discussion

6.1 Summary of Results

The analysis and experiments presented in this thesis are intended to advance robot hand dexterity in manipulating objects through the use of tactile sensing. Many issues in manipulation, from acquisition of the object to interactions between the grasped object and the environment, are studied in detail. In addition, many practical issues are revealed through actual hardware implementation of realistic manipulation tasks. Here, the important conclusions from the previous chapters are summarized.

Analytical and experimental comparison of force, skin acceleration, and the multi-element stress rate sensors for detecting contact show that the multi-element stress rate sensor has significant advantages in rejecting other manipulation transients. As a result, it provides the most reliable contact information in comparison to the other sensors. The multi-element stress rate sensor is also able to detect incipient slip significantly prior to gross motion of the object. Detecting incipient slip is important in minimizing grasp force of fragile objects in unstructured environments.

Comparison of contact localization abilities of tactile array and ICS revealed that ICS sensing results suffered from difficulties in calibration and manipulation transients. Tactile array sensing was immune to these problems, and the effects of shear loading can be minimized using a peak detection or a thresholded centroid algorithm. In the experiments, both of these sensing schemes provided contact localization accuracy of less than a millimeter during realistic manipulation tasks.

Experiments in coordinated grasping with vision and touch showed that vision can provide the necessary information to get the hand to the object, and that tactile sensing can provide the mechanical interaction information necessary for gentle grasping and grasp refinement. Experimental results confirm that serial and orthogonal combinations of vision and touch are effective in unstructured environments. Initial results also indicate that tactile primitives generated using sensor-space approach might provide a more direct and robust approach to tactile sensory integration than the task-phase primitives.

When an object is grasped by a multifingered robot hand, the lower limit to attainable object stiffness is set by the robot hand friction, which increases "effective stiffness" until the frictional forces are exceeded. Finger tip compliance which increase grasp stability also set an upper bound to attainable object stiffness. Maximum stiffness is also related to the geometry of the grasp configuration: due to controller stability limits, the maximum usable stiffness decreases as the specified center of compliance is moved away from the fingers, and increases with the distance between the fingers.

In performing an assembly task where objects must interact with one another, object position, orientation, and stiffness must be controlled. Uncertainty in the initial grasp pose can produce errors in the object orientation if the proper contact kinematics are not considered. Since the stiffness controller produces restoring forces proportional to the calculated object position, these errors cannot be compensated with an increase in object stiffness gain. Contact location errors also lead to inaccuracies in the stiffness of the object since the mapping between the finger-object is in error.

Experimental execution of a peg insertion task demonstrates that tactile sensing can effectively determine the initial orientation of an object which is necessary to successfully initiate the peg insertion. Once the peg has made contact with the hole, object impedance determines the resulting contact forces. A compliantly grasped object with the center of compliance placed at the tip of the peg will minimized interaction forces. However, manipulator friction dictates the controlled stiffness requirement for a given task and the

lowest "effective stiffness" obtainable. Controller stability also limits how far the center of compliance can be placed away from the robot finger tips.

These conclusions are particularly important for future researchers looking to extend robot hand usefulness by performing tasks which we humans take for granted. We live in unstructured environments, and our sensory feedback system is rich with wide variety of information. For manipulation, the important senses are kinesthesia, vision, and touch. We also possess a remarkable ability to selectively extract the important information from the different senses. In robot manipulation, one type of sensors cannot provide all the information required for successful manipulation. Especially in unstructured environments, the ability to use many complimentary sensing modes and the ability to fill in voids caused by a particular sensor, is crucial for robust manipulation.

This is not surprising when we examine our abilities to manipulate objects without one of the senses. When we loose our sense of touch, either through wearing a glove or anesthesia, our manipulation abilities are impaired. If we close our eyes during a manipulation task, we have to resort to groping and our manipulation abilities are likewise reduced. Current robots typically do not use vision nor tactile sensing. Imagine how dextrous we would be if we could not see or feel the object that we are manipulating.

6.2 Suggestions for Future Work

Tactile sensing provides a rich set of information about the manipulation process, and the unexplored areas of tactile sensing research are numerous. To advance tactile sensing in robot hand manipulation, improvements are needed in two areas. One area is related the hardware and hardware integration, and the other is intelligent use of sensory information.

Since the design of a smooth dextrous robot hand is a research topic in itself, this thesis used two fingered planar hands to investigate the benefits of tactile sensing. These systems might not look like robot hands, but their ability to control positions and forces were much better than typical robot hands. The next logical phase in robot hand design is to add

additional degrees of freedom while keeping the friction low or at least compensated such that the desired forces at the finger tips are obtained. Since coordinating large degrees of freedom is a difficult task, perhaps the next generation of hands should fall between a planar hand and a multifingered hand such as Salisbury or Utah/MIT hand.

Determining the hand-object relationship is critical for successful manipulation of objects. This is where tactile sensing is needed. Therefore, a tactile system needs to be integrated into the robot hand design at the initial design stage. Designing robot hands with surfaces to mount tactile sensors is easy in comparison to providing connections and routing space required for the electrical signal and power lines. The problem is an integration issue where the subsystems like joint angle sensing, tactile sensing, and joint actuation subsystems needs to work together in a very limited space. Imagine the number of nerve fibers which must run through our hands; it is an incredible thought. However, one idea to keep in mind here is to start simple, make it work and learn from it, only then move on to the next level of challenges.

Currently, there are too many different types of tactile sensors and not enough research being done with it. One of the reasons why many researchers have not used tactile sensing for manipulation is the scarcity of robot hands and tactile sensors along with the difficulty of integrating the two systems. Although integration problems my not be considered academically challenging, seemingly mundane issues such as packaging, connectors and cabling limit practical uses for tactile sensing.

Some researchers have developed VLSI tactile sensors with built-in electronics which resolve the cabling problem, but the sensors themselves are limited to relatively small, flat, and rigid sensors. A more reasonable approach is to use conventional tactile sensing technology which can be wrapped around robot fingers, and amplify the signals using custom designed integrated circuit with micro-packaging. These integrated circuits should provide the multiplexing and signal processing required for the sensors along with a common interface bus structure such that adding addition sensors can be achieved without increasing the number of electrical wires. Signal processing near the sensor also has the

added advantage of electrical noise reduction which is especially important in noisy environments where electrical noise can reduce the dynamic range and sensitivity of the sensors. Odeberg [1996] suggests a similar framework for distributed sensor processing.

Just as the current field of vision research relies on reliable and inexpensive CCD cameras, tactile sensing research will be widespread when the hardware becomes readily available. Researchers will be able to focus on research issues such as tactile sensor signal processing, sensor fusion of tactile information with other sensory information, robust robot hand manipulation, etc.

The most immediate and fruitful extensions of this work is in the area of combining vision with touch. There are many interesting experiments and tasks which can be performed using the vision and tactile system described in chapter 4. The area of research should focus more on sensor fusion and robust task planning. More specifically, when to use kinematic, vision or tactile information and how to optimize the information available from the various sensors are important issues which needs to be addressed. Suggested improvements to the hardware include a simple two fingered robot hand which could be mounted on the manipulator, faster signal processing computers, and a system architecture with one central computer which can easily exchange information between the robot and the sensory subsystems.

Furthermore, ideas like the "event manager" proposed by Trembley and Cutkosky [1995] could be extended to include vision, and it can be validated on a new platform. Whether this approach to manipulation is more robust than a sensor-space generated tactile primitives is another interesting research question. There might be advantages and limitations with both approach. The task-space based approach could segment a task more than necessary and make it unnecessarily complicated. Sensor-based primitives might not allow the use of different controllers since feedback control is performed in the sensor domain. These lower level tactile primitives needs to be intelligently integrated with vision using context and planning. Here, comparing different approaches

by Cutkosky and Hyde [1993], Brock [1993], McCarragher and Asada [1993] should result in a impressive demonstration of robot manipulation ability.

6.3 Closing Remarks

In comparison to Hollywood's depiction of robot technology, the current state of robotics research is embarrassingly primitive. Part of the reason is due to the fact that robotics is a new area of research which requires many different disciplines. To make useful and intelligent robotics as depicted in the movies, all disciplines in robotics needs to contribute together.

Currently, many of the researchers have focused on specific areas within robotics and tried to advance the knowledge in that specific area. Consequently, more people have specialized expertise in small segments of the robotics research rather than a stronger knowledge of robotics as a whole. If you can visualize a graph, with research topics within robotics along the horizontal axis, and the depth of knowledge along the vertical axis, current state of robotics research would look similar to the Boston sky line.

Since collaboration from different areas within robotics is rare, the specialized knowledge does not spread to the whole robotics community. In addition, the number of people directly using the results of other people's research is very small since the emphasis in research is on developing new ideas. I urge researchers to take initiative and collaborate within the robotics field. Collaboration between vision and manipulation, low and high level control, VLSI and tactile sensing, are a few examples which could benefit immediately.

Even if collaboration is not chosen, there are ways to advance the robotic field as a whole. Rather than thinking that the products of research are publications, contribute hardware and algorithms that other researchers can use along with the publications. These results needs to be robust and useful such that other researchers will want to use it as part of their research. By having different researchers combining these advancements, we connect the advancements of particular areas and the whole research area of robotics will advance.

I would like to close this thesis with a few personal notes. Problem solving skills are valuable, but it should not be applied to developing new ideas for the sake of it being new.

Make good ideas work well.

Conceptually, most systems are relatively simple. To make the concept work however, every little detail needs to be worked out and solved before the entire system functions properly. I've learned on many occasions to first understand the fundamentals. Having the wrong assumptions about a problem can be costly. Verify what you believe to be true by testing it extensively. Only continue onto the next level, when the problem has been understood and resolved. Systematically build on top of each other. It is a slow and tedious process, but it is the only process which enabled me to solve challenging problems.

References

P. K. Allen, Robotic object recognition using vision and touch. Kluwer Academic Publishers, Boston, 1987.

P. Akella and M. Cutkosky. Manipulating with soft fingers: modeling contacts and dynamics. Proceedings of the IEEE International Conference on Robotics and Automation, pages 764-769, Scottsdale, AZ, May 1989.

A. D. Berger and P. K. Khosla. Using tactile data for real-time feedback. International Journal of Robotics Research, 10(2):88-102, April 1991.

A. Bicchi, J. K. Salisbury, and P. Dario. Augmentation of grasp robustness using intrinsic tactile sensing. Proceedings of the IEEE International Conference on Robotics and Automation, pages 302–307, Scottsdale, AZ, May 1989.

A. Bicchi. Intrinsic contact sensing for soft fingers. Proceedings of the IEEE International Conference on Robotics and Automation, pages 968-973, Cincinnati, OH, May 1990.

A. Bicchi and G. Canepa. Optimal design of multivariate sensors. Measurement Science & Technology, 5(4):319-32, 1993.

A. Bicchi, J. K. Salisbury, and D. L. Brock. Contact sensing from force measurements. International Journal of Robotics Research, 12(3):249-262, June 1993.

M. Boshra and H. Zhang, A constraint-satisfaction approach for 3D vision/touch-based object recognition. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, vol. 3, pp. 368-73. 1995.

D. L. Brock. A Task-level Control System for Automatic Robotic Grasping. MIT AI Lab Number 1469. January 1993.

D. Brock and S. Chiu. Environment perception of an articulated robot hand using contact sensors. Robotics and Manufacturing Automation, volume 15 of PED, pages 89-96, ASME Winter Annual Meeting, Miami, November 1985.

R. W. Brockett. On the Computer Control of Movement. Proceedings of the IEEE International Conference on Robotics and Automation, pages 534-540, Philladelphia, April 1988.

R. W. Brockett. Robotic hands with rheological surfaces. Proceedings of the IEEE International Conference on Robotics and Automation, pages 942-946. St. Louis, March 1985.

G. Buttazzo, P. Dario, and R. Bajcsy. Finger based explorations. In David Casasent, editor, Intelligent Robots and Computer Vision: Fifth in a Series, pages 338–345, Cambridge, MA, October 28-31, 1986. Proceedings of SPIE Volume 726.

N. Chen, H. Zhang, and R. Rink. Edge tracking using tactile servo. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS '95. vol. 2, pages 84-89, Pittsburgh, PA, August 1995.

A. Cole, P. Hsu, and S. Sastry. Dynamic regrasping by coordinated control of sliding for a multifingered hand. Proceedings of the IEEE International Conference on Robotics and Automation, pages 781-786, Scottsdale, AZ, May 1989.

A. Cole, P. Hsu, and S. Sastry. Dynamic control of sliding by robot hands for regrasping. IEEE Transactions on Robotics and Automation, 8(1):42-52, 1992.

M. R. Cutkosky. Robotic Grasping and Fine Manipualation. Kluwer Academic Publishers, Boston, 1985.

M. R. Cutkosky, J. M. Jourdain, and P. K. Wright. Skin materials for robotic fingers. Proceedings of the IEEE International Conference on Robotics and Automation, pages 1649-54 vol. 3, Raleigh, NC, March 31 - April 3, 1987.

M. R. Cutkosky and I. Kao. Computing and controlling compliance of a robotic hand. IEEE Transactions on Robotics and Automation, vol. 5, pages 151-165, April 1989.

M. R. Cutkosky and J. M. Hyde. Manipulation Control with Dynamic Tactile Sensing. 6th International Symposium on Robotics Rresearch, Hidden Valley, PA, October 1993.

R. S. Fearing. Some experiments with tactile sensing during grasping. Proceedings of the IEEE International Conference on Robotics and Automation, pages 1637-1643, Raleigh, NC, March 31-April 3, 1987.

R. S. Fearing. Tactile sensing mechanisms. International Journal of Robotics Research, 9(3):3–23, June 1990.

R. S. Fearing. Tactile sensing for shape interpretation. In S. T. Venkataraman and T. Iberall, editors, Dextrous Robot Hands. Springer-Verlag, pages 209-238, 1990a.

G. D. Hager. A modular system for robust hand-eye coordination using feedback from stereo vision. To appear in the IEEE Transactions on Robotics and Automation, submitted Jan. 1995a.

G. D. Hager. The "XVision" system: a general purpose substrate for real-time vision-based robotics. Proceedings of the Workshop on Vision for Robotics. pages 56-63, 1995b.

G. D. Hager, W. Chang, and A. S. Morse. Robot feedback control based on stereo vision: towards calibration-free hand-eye coordination. IEEE Cont. Sys. Mag. 15(1):30-39, 1995.

H. Hanafusa and H. Asada. A robot hand with elastic fingers and its application to assembly process. In Brady et al., editors, Robot Motion: Planning and Control. Cambridge, MA, MIT Press, 1983.

N. Hogan. Impedance control: An approach to manipulation: Part 1, 2, and 3. Transactions of the ASME Journal of Dynamic Systems, Measurement, and Control, 107:1–24, March 1985.

R. D. Howe. A force reflecting teleoperated hand system for the study of tactile sensing in precision manipulation. Proceedings of the IEEE International Conference on Robotics and Automation, pages 1321-1326, Nice, France, May 1992.

R.D. Howe. Tactile sensing and control of robotic manipulation. Journal of Advanced Robotics, 8(3):245-261, 1994.

R. D. Howe and I. Kao and M. R. Cutkosky. The sliding of robot fingers under torsion and shear loading, Proceedings of the IEEE International Conference on Robotics and Automation, pages 103-108, Philadelphia, PA, 1988.

R. D. Howe and M. R. Cutkosky. Sensing skin acceleration for texture and slip perception. Proceedings of the IEEE International Conference on Robotics and Automation, Scottsdale, AZ, May 1989.

R. D. Howe, N. Popp, P. Akella, I. Kao, and M. Cutkosky. Grasping, manipulation, and control with tactile sensing. Proceedings of the IEEE International Conference on Robotics and Automation, pages 1258–1263, Cincinati, OH, May 1990.

R. D. Howe and M. R. Cutkosky. Touch sensing for robotic manipulation and recognition. O. Khatib, et al., editors, Robotics Review 2, pages 55–112, MIT Press, Cambridge, MA, 1992.

R. D. Howe and M. R. Cutkosky. Dynamic tactile sensing: Perception of fine surface features with stress rate sensing, IEEE Transactions on Robotics and Automation, 9(2):140-151, April 1993.

R. D. Howe and M. R. Cutkosky. Practical force-motion models for sliding manipulation. International Journal of Robotics Research, in press. 1996.

J. Hyde. Ph.D. Thesis, Stanford University. July 1995.

R.S. Johansson and Å.B. Vallbo. Tactile sensory coding in the glabrous skin of the human hand. Trends in Neuroscience, 6(1):27-32, 1983.

R. S. Johansson and G. Westling. Factors influencing the force control during precision grip. Experimental Brain Research, 53:277-284, 1984a.

R. S. Johansson and G. Westling. Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. Experimental Brain Research, 56:550-564, 1984b.

R. S. Johansson and G. Westling. Programmed and triggered actions to rapid load changes during precision grip. Experimental Brain Research, 71:72-86, 1988.

R. S. Johansson and G. Westling. Tactile afferent signals in the control of precision grip. In M. Jeannerod, editor, Attention and Performance, vol XIII, pages 677-713. Hilldale, NJ, 1990.

I. Kao and M. R. Cutkosky. Quasistatic manipulation with compliance and sliding, International Journal of Robotics Research, 11(1):20-40, Feb. 1992.

H. Kazerooni. Direct-drive active compliant end effector (active RCC). IEEE Journal of Robotics and Automation, 4(3):324-333, June 1988.

J. Kerr and B. Roth. Analysis of multifingered hands. The International Journal of Robotics Research, 4(4):3-17, 1986.

O. Khatib. A unified approach for motion and force control of robot manipulators: the operational space formulation. IEEE Journal of Robotics and Automation, 3(1):43-53, February 1987.

Z. Li, P. Hsu, and S. Sastry. Grasping and coordinated manipulation by a multifingered robot hand. International Journal of Robotics Research, 8(2):33-50, August 1989.

H. Maekawa, K. Tanie, and K. Komoriya. Tactile sensor based manipulation of an unknown object by a multifingered hand with rolling contact. Proceedings of the IEEE International Conference on Robotics and Automation, pages 743-750, Nagoya, Aichi, Japan, May 1995.

B. J. McCarragher and H. Asada. A discrete event approach to the control of robot assembly tasks. Proceedings of the IEEE International Conference on Robotics and Automation, Atlanta, GA, May 1993.

D. J. Montana. Tactile Sensing and the Kinematics of Contact. Ph.D. thesis, Harvard University. August 1986.

J. B. Morrell. Force modulation of a robot gripper using slip detection. Master's thesis, University of Washington, June 1990.

Y. Nakamura and H. Hanafusa and N. Ueno. A piezoelectric film sensor for robotic end-effectors. In A. Pugh, editor, Robot sensors, Volume 2: Tactile and Non-Vision, pages 247-57, IFS Publications/Springer-Verlag, Berlin, 1986.

B. J. Nelson and P. K. Khosla. Force and vision resolvability for assimilating disparate sensory feedback, To Appear in IEEE Transactions on Robotics and Automation, Sept. 1996.

H. R. Nicholls and M. H. Lee. A survey of robot tactile sensing technology. International Journal of Robotics Research, 8(3):3–30, June 1989.

E. J. Nicolson. Ph.D. thesis, University of California Berkley. November 1994.

W. Nowlin. Tactile Sensing with Compliant Manipulators. PhD thesis, Harvard University. February 1991.

H. Odeberg. A tactile sensor data-processing system. Sensors and Actuators A (Physical), vol. A49, no. 3, p.173-80, 1996.

E. Ono, N. Kita, and S. Sakane. Strategy for unfolding a fabric piece by cooperative sensing of touch and vision. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems. vol. 3, pages 441-5, 1995.

R. W. Patterson and G. E. Nevill, Jr. The induced vibration touch sensor - a new dynamic touch sensing concept. Robotica, 4:27-31, 1986.

N. I. Rafla and F. L. Merat. Vision-taction integration for surface representation. Proceedings of the IEEE International Conference on Systems Engineering, pages 519-22, 1990.

M. Rucci and R. Bajcsy. Learning visuo-tactile coordination in robotic systems. Proceedings of the IEEE International Conference on Robotics and Automation, vol. 3, pages 2678-83, 1995.

R. Russell. Compliant-skin tactile sensor. Proceedings of the IEEE International Conference on Robotics and Automation, pages 1645-1648. Raleigh, NC, March 31-April 3 1987.

J. K. Salisbury. Interpreting contact geometries from force measurements. Proceedings of the 1st International Symposium on Robotics Research. Bretton Woods, NH, September 1984.

J. K. Salisbury. Kinematic and force analysis of articulated hands. In M. T. Mason and J. K. Salisbury, editors, Robot Hands and the Mechanics of Manipulation. MIT Press, Cambridge, MA, 1985.

K. B. Shimoga and A. A. Goldenberg. Soft Materials for Robotic Fingers. Proceedings of the IEEE International Conference on Robotics and Automation, pages 1300-5, Nice, France, 1992.

P. Sikka, H. Zhang and S. Sutphen. Tactile servo: control of touch-driven robot motion. Experimental Robotics III: The 3rd International Symposium. Springer-Verlag, Lecture Notes in Control and Information Sciences, vol. 200 pages 221-233. Oct. 1993.

J. S. Son, E. A. Monteverde, and R. Howe. A tactile sensor for localizing transient events in manipulation, Proceedings of the IEEE International Conference on Robotics and Automation, vol. 1, pages 471-476, San Diego, CA, May 1994.

J. S. Son and R. D. Howe. Tactile sensing and stiffness control with multifingered hands. Proceedings of the IEEE International Conference on Robotics and Automation, pages 3328-3234, Minneapolis, April 1996.

J. S. Son, M. R. Cutkosky, and R. D. Howe. A comparison of object localization by contact sensors. Journal of Robotics and Autonomous Systems, Vol. 17, No. 4, June/July 1996.

J. S. Son, R. D. Howe, J. Wang, and G. D. Hager. Preliminary results on grasping with vision and touch. Submitted to IEEE/RSJ International Conference on Intelligent Robots and Systems. November 1996.

S. A. Stansfield. A robotic perceptual system utilizing passive vision and active touch. International Journal of Robotics Research, 7(6):138-61, 1988.

G. P. Starr. An experimental investigation of object-stiffness control using a multifingered hand. Robotics and Autonomous Systems, pages 33-42, vol.10, no.1, 1992.

M. Trembley and M. Cutkosky. Estimating friction using incipient slip sensing during manipulation tasks. Proceedings of the IEEE International Conference on Robotics and Automation, Atlanta, GA, May 1993.

M. Trembley and M. Cutkosky. Using sensor fusion and contextual information to perform event detection during a phase-based manipulation task. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems. vol. 3, pages 262-267. Pittsburgh, PA, August 1995.

T. Tsujimura and T. Yabuta. Object detection by tactile sensing method employing force/torque information. IEEE Transactions on Robotics and Automation, 5(4):444-50, August 1989.

D. E. Whitney. Quasi-static assembly of compliantly supported rigid parts. Journal of Dynamic Systems, Measurement, and Control, 104:65-77, March 1982.