

Contact Transition Control: An Experimental Study

James M. Hyde and Mark R. Cutkosky

Center for Design Research
Bldg. 02-530 Duena St.
Stanford University
Stanford, California 94305-4026

Abstract

Successful control of contact transitions is an important capability of dextrous robotic manipulators. In this paper we examine several methods for controlling the transition from free motion to constrained motion, with an emphasis on minimizing fingertip load oscillations during the transition. A new approach, based on input command shaping, is discussed and compared with several methods developed in prior research. The various techniques were evaluated on a one-axis impact testbed, and we present results from those experiments. The input shaping method was found to be comparable, and in some cases superior, to existing techniques of contact transition control.

Introduction

Contact transitions can be notoriously difficult to control. These events involve all of the well known problems of force control, with the added difficulties of non-zero approach velocities and discontinuous dynamic characteristics. If a robotic manipulator is to interact effectively with its environment, however, it must frequently make and break contact with foreign objects. An event-driven, reflexive capability to manage grasp state transitions is crucial for successful, practical dextrous manipulation.

Compliant fingertips have been used to eliminate some contact transition problems, but many difficult tasks remain. We find that even with soft fingertips, contact transitions tend to excite instability and/or result in undesirably high impact forces. These phenomena can be particularly bothersome when manipulating with sensor-laden fingertips, devices that are sensitive to "glitches" in force control.

Our particular goal in this work is to achieve smooth, stable transitions between motion and force control. We want to avoid instability and large force spikes during the controller transition, while increasing the grasp force from zero to the desired level as rapidly as possible.

Contact problems have not been neglected in prior research. Mills and Lokhorst [5] implemented a discontinuous controller that sought to reduce the problems of contact instability and also managed events of contact loss and trajectory tracking.

Hogan [2] used impedance control in experiments involving contact transitions. His implementation achieved stability against a stiff environment and avoided controller transitions and inverse kinematic computations. Vossoughi and Donath [9] also employed impedance methods in tests with varying environment stiffnesses.

Actively compliant fingertips containing electro-rheological fluids were explored by Akella, Siegart, and Cutkosky [1], who managed to control the damping characteristics of a fluid-filled fingertip to minimize force oscillations following the onset of contact.

Youcef-Toumi and Gutz [9] developed a dimensionless representation of impact behaviour and used integral force compensation with velocity feedback to improve impact response.

Khatib and Burdick [4] proposed a method for dissipating impact oscillations that involved increasing the velocity gains of a proportional-derivative force controller for a limited time following impact. By disabling the high velocity feedback gain after the impact oscillations decayed, response times to subsequent force commands were decreased.

More recently, Qian and De Schutter [6] presented an active nonlinear damping approach. This method examined the force signal derivative, and effectively added a coulomb friction term to the output force command when this derivative exceeded a threshold.

We will examine some of these methods, and compare them to a relatively new technique: input command preshaping. This method has enjoyed success in many applications of position control, as shown in Singer [7] and Hyde [3], but to the authors' knowledge, preshaping has never been applied to force control problems.

Input command preshaping is essentially a feedforward technique that employs linear system theory and a plant's inherent dynamic characteristics to suppress vibration. Shaping an input is straightforward. First, the system frequency targeted for suppression must be identified to within about 10% (similar to the identification of a mode targeted for control by a dominant 2nd order controller). This mode is input into closed-form equations used to

calculate a train of impulses, a process which can occur off-line if necessary.

The impulses are convolved with the input command in real time, imposing only a small computational burden and producing a shaped input which is then fed to the plant. At runtime, the requirements are that the servo rate and actuator authority be sufficient to accurately produce the desired commanded force function -- in practice, servo rates and actuator authorities found adequate for stable force control will be more than sufficient. When applying the shaped input, the system output will be largely devoid of vibrations at the targeted frequency.

As an example, a standard single mode shaper has three impulses. When a step input is convolved with this impulse train, the modified input resembles a three-step staircase, as shown in Figure 1. Any commanded trajectory (step, ramp, or sinusoid; force, position, etc.) can be "shaped" with this method. Preshaping produces only small system response delays, and is relatively insensitive to variations in plant parameters. Since the method operates outside of a standard control loop, preshaping can be used in conjunction with any of the other feedback methods mentioned above. Shapers can also be used to eliminate multiple-mode vibration. For full details on input command preshaping, see Singer [7] and Hyde [3].

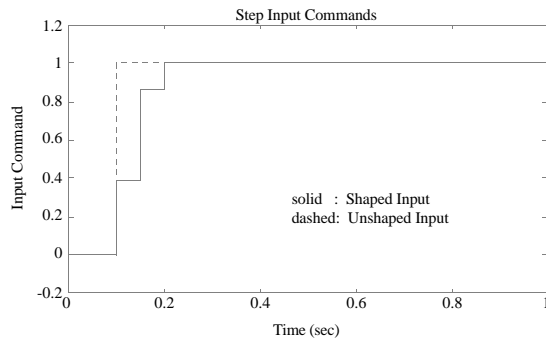


Figure 1: A "shaped" step input.

Control Approaches

The focus of this paper is the experimental comparison of several contact transition control approaches. All of the approaches were evaluated on a simple impact testbed, described below, and the experiments had a common structure of three phases: free motion, impact, and constrained motion. In this section, we provide a brief description of the tested approaches, including a formulaic representation of the control law used, where appropriate.

(I) Baseline: Discontinuous Control

The baseline controller for the following experiments was a simple discontinuous control law. We chose this law as a baseline because it achieved the basic goals of contact transition, but its performance was rather poor. All of the tested controllers, therefore, were able to demonstrate

some improvement over the baseline performance. In this control law, proportional velocity control was exerted in the non-contact phase, switching to proportional force control upon impact. The contact state was triggered when the applied load exceeded a threshold value. This approach resulted in a step input in commanded force and a discontinuous change in closed-loop system properties. The method suffered from serious oscillations due to the discontinuity at contact, even when the force control gains were tuned for stability under steady-state conditions. The laws used to implement this controller in the pre- and post-contact phases are given by equations (1a) and (1b), respectively:

$$f_c = k_{v_1} (v_{des} - v_{act}) \quad (1a)$$

$$f_c = k_f (f_{des} - f_{act}) - k_{v_2} v_{act} \quad (1b)$$

where f_c denotes the commanded force, f_{act} and v_{act} are the actual force and velocity readings, and f_{des} and v_{des} are the desired force and velocity values. The non-contact proportional gain is k_{v_1} , and k_f is the contact force gain. The velocity gain k_{v_2} is applied as needed to maintain stability after contact.

(II) Impedance Control

The second method tested was impedance control, which had the advantage of providing a uniform control approach for both the unconstrained and constrained phases of the task. We implemented a controller paraphrased from Hogan [2]. In contrast with the baseline approach, force and velocity feedback gains remained active throughout the experiment, resulting in less oscillation upon contact. More generally, a trade-off had to be made among rise time, oscillation and peak impact force. The impedance control law tested is given by equation (2).

$$f_c = k_f (f_{des} - f_{act}) - k_v (v_{des} - v_{act}) \quad (2)$$

(III) Active Impact Damping

A more aggressive approach is to apply a large velocity feedback term for a short time immediately following impact, as proposed by Khatib and Burdick [4]. This method was implemented using laws similar to (1a) and (1b), with an additional term added to k_{v_2} for 0.1 seconds following impact.

(IV) Active Nonlinear Damping

We also tested a non-linear damping approach developed by Qian and De Schutter [6]. We employed a low pass filter to calculate relatively clean force derivatives, and implemented the control law in a discontinuous manner. During pre-contact motion, equation (1a) governed the system, and upon impact, the law switched to equations (3a) and (3b).

$$d_n = \text{sgn}(\dot{f}_{act}) T(\dot{f}_{act}) k_c f_{des} \quad (3a)$$

$$f_c = k_f (f_{des} - f_{act}) - d_n \quad (3b)$$

In the previous equations, $T(\dot{f}_{act})$ is a threshold function; it equals zero or unity, depending on the magnitude of the force derivative. The feedback gain k_c determines the strength of the applied damping.

(V) Input Preshaping

In contrast to the previous feedback-based approaches, input command preshaping modifies the feedforward command to suppress post-maneuver vibration (see the introduction for more detail). We conducted frequency identification tests using the baseline controller and some soft robotic fingertips, and determined that the fingertips had one dominant mode (the plant + controller natural frequency), and one or more secondary modes (the "impact" or "ringing" frequencies), which decayed quickly after impact.

Using the equations in Hyde [3], we first generated shaping sequences for the dominant low frequency mode. When used in conjunction with the baseline controller, the shaper removed most of the residual vibration, but some oscillations remained. We then turned our attention to the secondary modes, and by shaping for the lowest secondary frequency, we achieved the best response.

The success of this approach can be explained by the interaction between the impact dynamics and the applied control law. The fingertip under baseline control was stable when gently placed against the environment (by hand); the small oscillations produced by an impact tended to disturb the system, helping to excite the dominant controller + plant frequency. By removing the small impact vibrations, we eliminated part of the source of the dominant low frequency oscillations.

Thus it appears to be useful to preshape for the dominant frequency of the destabilizing impact event. For the present experiments, this frequency was the lowest "ringing" mode in the fingertip structure. The control law for this case (V) used in the subsequent experiments was simply the baseline controller whose output was shaped to suppress vibration at the fingertip's lowest "ringing" mode.

Experimental Apparatus

The experiments used to evaluate the various control strategies were conducted on a linear one-axis impact testbed, consisting of a large voice coil linear motor, fingertip mounting surface, compliant fingertip, and a solid wall (a firmer fingertip striking a softer wall produces similar results). The fingertip mounting surface was attached to the motor armature, which traversed a hardened track on low friction ball bearings. Sensors measured armature position and velocity, and fingertip load. The fingertip load cell, the most important sensor in the tests, provided readings up to 25 lbf., with a resolution of 0.04 lbf. A microcomputer equipped with a Digital Signal Processing (DSP) board served as the controller for the experiments. The system was capable of 10 kHz servo rates while managing 8 channels of input and 2

channels of output. Most of the experiments described here were conducted at 1 kHz. The maximum obtainable open-loop system bandwidth was 175 Hz for a 2 mm peak-to-peak position oscillation. A schematic of the experimental apparatus is shown in Figure 2.




Figure 2: One-Axis Impact Testbed.

Experiments and Results

In the following tests, we were mainly interested in controller performance with a type of fingertip suited for dextrous manipulation with tactile sensing. This fingertip consists of a rubber skin over a supporting layer of foam rubber. However, for purposes of comparison we also conducted experiments with a more easily characterized "fingertip" consisting of a coil spring. Each contact transition control approach was tested on both finger types, under varying conditions. The first set of experiments were performed under idealized conditions including clean velocity and force feedback signals (< 2% noise, RMS), and a 1 kHz servo rate. A second set of experiments was conducted with the rubber/foam fingertip under less ideal conditions in which the velocity signal was corrupted with 15% RMS white noise. In a third set of tests on the rubber/foam fingertip, the sensor readings were clean, but the servo rate was reduced to 500 Hz.

The performance of the baseline discontinuous control law (I) with the compression spring fingertip is shown in Figure 3. The fingertip force signal remained near zero during the constant velocity approach phase, rose above the contact threshold, triggering the control law switch, and then oscillated as the proportional control law servoed toward the desired force. The controller gains were tuned to achieve stable steady-state force control, but only marginally stable impact response. Figure 4 shows the Fourier transform of the data in Figure 3, immediately following the impact. The dominant low-frequency peak is off the scale; this view shows the secondary frequencies at approximately 50, 110, and 150 Hz. The presence of multiple secondary modes can be explained by the "ringing" of the steel coils of the spring. The frequency spectrum, as mentioned earlier, becomes important when applying the preshaping method.

Figure 5 shows the response of the "spring" fingertip to the active nonlinear damping method (IV) and to impedance control (II). When compared to the baseline

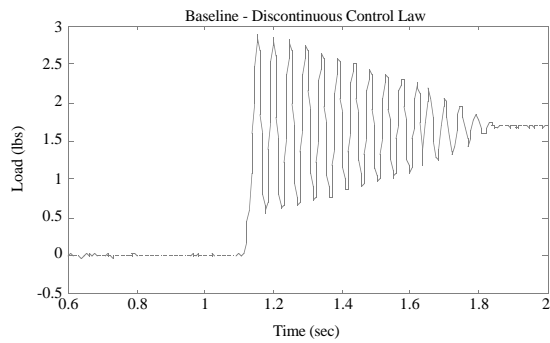


Figure 3: Contact response of the "spring" fingertip.

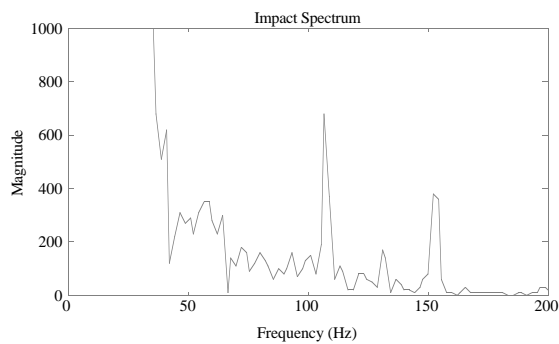


Figure 4: Frequency spectrum of the post-impact force signal.




Figure 5: Response of the "spring" fingertip to Impedance (II) and Active Nonlinear Damping (IV) controllers.

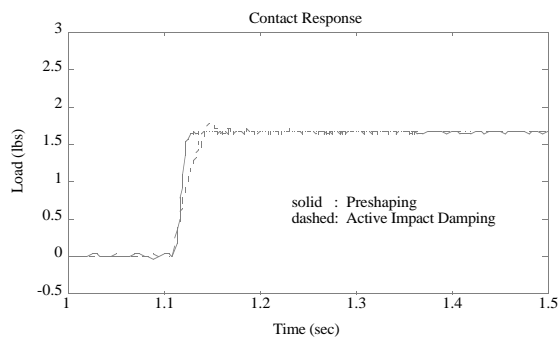


Figure 6: Response of the "spring" fingertip to Active Impact Damping (III) and Preshaping (IV) controllers.

response in Figure 3, impedance control produced excellent stability and vibration reduction, at the expense of a modest increase in rise time. The nonlinear damping method effectively suppressed the vibration shown in the baseline tests, but some residual vibration remained. In implementing this method, we took care to tune the force derivative threshold levels and gains to improve the response. We found that higher gains or lower thresholds than those used to generate Figure 5 amplified variations in the force derivative calculations to the point of destabilizing the system.

In Figure 6, we compare the active impact damping method (III) to the preshaping technique (V). In the dashed trace, the impact damping was turned on for 0.1 seconds following the onset of contact. Extending the impact damping "life" improved the response, but tended to delay post-contact force commands. The solid trace shows the response to the input command as shaped for the lowest secondary frequency of Figure 4. The shaper performed well, with a rise time and vibration reduction comparable to the impedance control response of Figure 5.

The "spring fingertip" experiments were repeated with the rubber/foam fingertip. The baseline controller gains were once again tuned to achieve marginal stability, but in this case, the force control law involved proportional action only. Any system damping was due to the natural damping characteristics of the foam. We took this action to minimize the effects of noisy velocity feedback on the baseline controller, a point discussed in detail below. The rubber/foam fingertip had a higher fundamental frequency than the spring fingertip, and the spectrum of the rubber/foam impact vibration revealed only one secondary frequency, at approximately 70 Hz, as shown in Figure 7.

As Figure 8 shows, the impedance control (II) of impacts with the foam/rubber fingertip generated a response nearly identical, except for a small increase in rise time, to the corresponding spring finger response. The active nonlinear damping approach (IV) also performed nearly as well with the foam/rubber fingertip; the only difference between the solid traces of Figures 8 and 5 being a slightly larger overshoot and higher-frequency residual vibration.

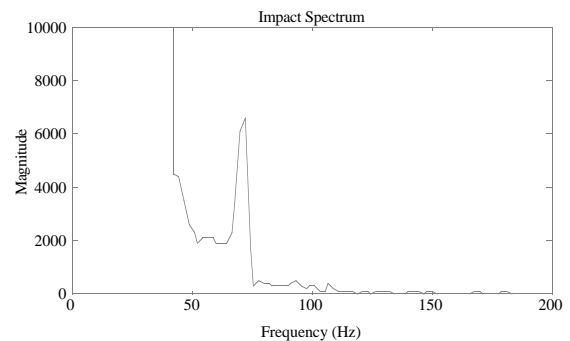


Figure 7: Frequency spectrum of the post-impact force signal.

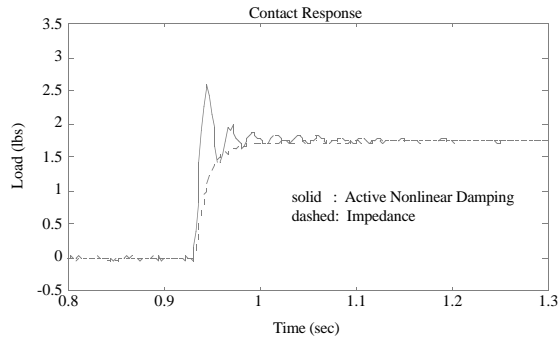


Figure 8: Response of the rubber/foam fingertip to Impedance(II) and Active Nonlinear Damping(IV) controllers.

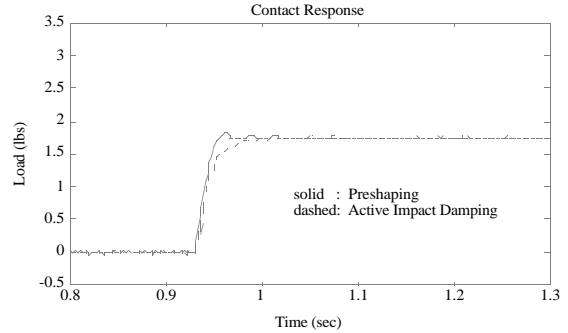


Figure 9: Response of the rubber/foam fingertip to Active Impact Damping(III) and Preshaping(V) controllers.

Figure 9 compares the performance of the active impact damping (III) and preshaping (V) approaches. In these runs, the impact damping still caused a delayed rise time as compared to the preshaping method, but the impact damping appeared to be more effective at suppressing the residual vibration. Recall that the shaper was developed assuming a linear system. The foam rubber exhibits non-linear behaviour, and we believe that this is responsible for the slightly degraded shaper performance shown in Figure 9. Both methods still inhibit a large fraction of the vibration caused by the baseline controller.

In the next set of experiments, we retained the rubber/foam fingertip and corrupted the velocity feedback signal with 15% RMS white noise. As expected, the noise caused performance degradation in those control methods that relied on velocity information to execute vibration suppression. Figure 10 shows the performance of the impedance (II) and active nonlinear damping (IV) algorithms. The impedance control method was clearly affected by the noisy velocity information, but the nonlinear damping approach still performed well once contact was initiated.




Figure 10: Response to Impedance(II) and Active Non-linear Damping(IV) controllers with velocity noise.

Figure 11 compares the active impact damping (III) and preshaping (V) techniques. The impact damping approach, with its reliance on velocity feedback, was adversely affected by the velocity noise. The preshaping method, however, was largely unaffected by the noisy velocity signal. The response to the shaped input shown in Figure 11 is similar to the response in the noise-free tests, shown in Figure 9.

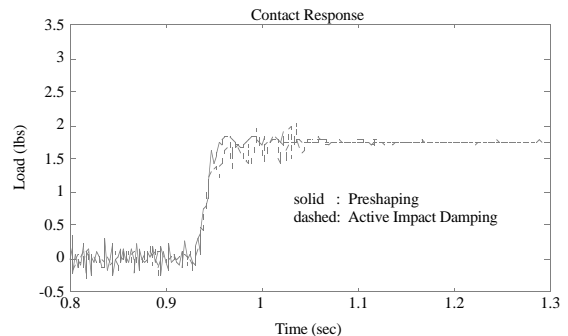


Figure 11: Response to Active Impact Damping(III) and Preshaping(V) controllers with velocity noise.

In the final group of tests, we limited our examination to the active impact damping (III) and preshaping (V) techniques. These experiments involved the rubber/foam fingertip and clean velocity signals, with the servo rate reduced from 1 kHz to 500 Hz. The new servo rate produced two major effects. First, the controller gains required adjusting, since in a digital system the gains are functions of the servo rate. Second, the new sampling rate decreased the frequency resolution available to the shaping method. As a result, the "shaped for" frequency was a rougher approximation of the desired frequency.




Figure 12: Response to Active Impact Damping(III) and Preshaping(V) controllers with reduced servo rate.

After the velocity gains were reduced to achieve stability, the impact damping controller lost some of its effectiveness, as shown in a comparison of the dashed traces of Figures 12 and 9. The preshaping technique, however, was only mildly affected by the drop in servo rate. The amplitude of the residual vibration in Figure 12's solid trace is higher than in Figure 9, a result of the reduced frequency resolution, but the majority of the baseline vibration was still suppressed.

A potential criticism of preshaping is that the method is merely a complicated force ramping technique, and its performance is obtained only by removing some of the high-frequency, high magnitude energy from the input command. We present Figure 13 as a counter-argument. This plot is an "insensitivity curve." The x-axis represents normalized frequency, i.e. the shaper frequency divided by the actual system frequency. The y-axis represents normalized vibration amplitude. If the vibration amplitude reaches unity, then the shaper had zero effect (properly implemented, input shapers can cause no increase in the natural system vibration). Figure 13 presents an experimental verification of the shaping method's insensitivity to plant parameter variations, and this figure demonstrates that the shaper was most effective at a particular frequency, a rebuttal to the "ramping" argument mentioned above. This insensitivity helped to insure the shaper's success even when the servo rate dictated a marginal frequency resolution.

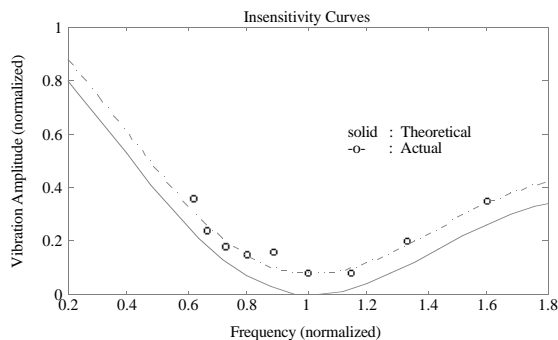


Figure 13: Comparison of theoretical and experimental shaper insensitivity curves.

Conclusions

As the results demonstrate, any of the methods in the referenced literature can result in improved response if gains are judiciously tuned. For the feedback-based methods, a tradeoff among rise time, peak impact force and duration of oscillation typically results. The methods that rely on velocity or force feedback signals are, not surprisingly, sensitive to noise in those signals. In contrast with these feedback based methods, input command preshaping suppresses vibration by modifying feedforward information, and is therefore less susceptible to noise.

Applying the preshaping technique to contact transition control appears to be a relatively new concept, even though the method has enjoyed success in position and

velocity control problems. We found knowledge of the frequency structure of the impact response, if not the exact modes, to be particularly useful in realizing the full potential of the shaping method. Shaping for the dominant mode will remove approximately 85% of the original vibration, but by identifying and shaping for the lowest secondary frequency, roughly 95% of the vibration can be suppressed. The method is not unduly sensitive, however, to modeling errors; as Figure 13 shows, over- or under-estimating the frequencies by as much as 20% results in only a 10% increase in the amplitude of post-impact oscillation.

The experimental results indicate that the benefits of input command preshaping are comparable in magnitude to those of feedback methods previously presented. Where noise in velocity signals is a problem (perhaps because the velocity must be obtained by estimation from digital position information) preshaping may be particularly advantageous.

Acknowledgments

This paper describes work performed at the Stanford University Center for Design Research. Funding for this research was provided in part by the Office of Naval Research under the University Research Initiative contract #N-00014-90-J-4014-P01. The authors would like to thank Marc Tremblay, Dean Chang and Rob Howe for their technical advice and electronics expertise.

References

- [1] Akella, P., Siegwart, R., and Cutkosky, M., "Manipulation with Soft Fingers: Contact Force Control," IEEE International Conference on Robotics and Automation, 1991.
- [2] Hogan, N., "Stable Execution of Contact Tasks Using Impedance Control," IEEE International Conference on Robotics and Automation, 1987.
- [3] Hyde, J. M., "Multiple Mode Vibration Suppression in Controlled Flexible Systems," MIT Artificial Intelligence Laboratory Technical Report #1295, June 1991.
- [4] Khatib, O., and Burdick, J., "Motion and Force Control of Robot Manipulators," IEEE International Conference on Robotics and Automation, 1986.
- [5] Mills, J. K., and Lokhorst, D. M., "Experimental Results in Manipulator Contact Task Control," IEEE International Conference on Robotics and Automation, 1991.
- [6] Qian, H. P., and De Schutter, J., "Introducing Active Linear and Nonlinear Damping to Enable Stable High Gain Force Control in Case of Stiff Contact," IEEE International Conference on Robotics and Automation, 1992.
- [7] Singer, N. C., "Residual Vibration Reduction in Computer Controlled Machines," MIT Artificial Intelligence Laboratory Technical Report #1030, February 1990.
- [8] Vossoughi, R., and Donath, M., "Robot Hand Impedance Control in the Presence of Mechanical Nonlinearities," *Robotics and Manufacturing Automation*, PED Vol. 15, ASME Winter Annual Meeting, Miami, November 17-22, 1985.
- [9] Youcef-Toumi, K., and Gutz, D. A., "Impact and Force Control," IEEE International Conference on Robotics and Automation, 1989.