A Nonlinear PD Controller for Force and Contact Transient Control

Y. Xu, J.M. Hollerbach, and D. Ma

This article describes a nonlinear proportional and derivative (NPD) controller and its use for both non-contact transient force control (NCTFC) (zero velocity when contact occurs) and contact transient force control (CTFC) (non-zero velocity when contact occurs). The key advantages of NPD control are its high disturbance rejection and robustness to time delay. We present a gain design method for NCTFC. Simulations and experiments on the Sarcos Dextrous Arm show that PD control becomes unstable for CTFC while NPD control is stable and reaches the steady state quickly. Experiments with human subjects on NCTFC and CTFC are also presented.

Despite performance limitations, PD control is most widely used for robot position control because of its simplicity. More advanced position controllers often incorporate PD control, such as computed torque or resolved acceleration control [7]. For robot force control, most controllers are simple PD or PID [1, 2, 13, 15] as well. With a rigid environment, PD based force control often induces instability and large oscillations because the loop gain is too high [1] and there is not enough damping. The performance is even worse than that of open loop force control, which is always stable, but open loop force control has a low bandwidth and is sensitive to force disturbances.

In any real force control task, a robot needs to frequently make and break contact with the environment. The robot must successfully manage two states—free motion and contact—and state transitions. The state transition adds discontinuity and nonzero velocity. These transitions may be event-driven or simply caused by large disturbances. All these problems make the transient controller design a challenging issue.

A number of force control strategies have been proposed: impedance control [3], stiffness control [10], passive and active damping controls [6, 9, 14], and command preshaping control

[4]. Hyde and Cutkosky showed experimentally that all of the methods can result in stable force control and improved response if the gains are judiciously tuned and the environment is compliant. It has been identified that force control with a rigid object is the worst case [1, 2]. An and Hollerbach [1] showed that a dominant pole plus heavy lowpass filtering of the force error stabilizes force control, but they did not address state transitions. Qian and Schutter [9] showed that active damping plus a dominant pole can give stable contact with a rigid object; state transitions also were not tested. Hogan [3] achieved stability with impedance control under contact transitions, but impedance control regulates the interaction port impedance, not the force explicitly. Volpe and Khosla proposed an impact controller with negative feedback gains. Their method can maintain stability but cannot track force commands, which is the goal of the explicit force control.

In this paper, we describe an NPD controller [16, 17], analyze its gain design, and experimentally investigate NCTFC and CTFC against a rigid environment. The goal is to avoid instability and large force spikes during controller transition, while increasing the contact force from zero to the desired level as rapidly as possible.

NPD Theory

Assume a general system equation:

$$\ddot{x} + c\,(\dot{x}, x) = u \qquad (1)$$

where x and $\dot{x}$ represent the system states, and $c(\dot{x}, x)$ is a nonlinear, state-dependent term. For a setpoint task, PD control derives the input proportionately to the output error and current states.
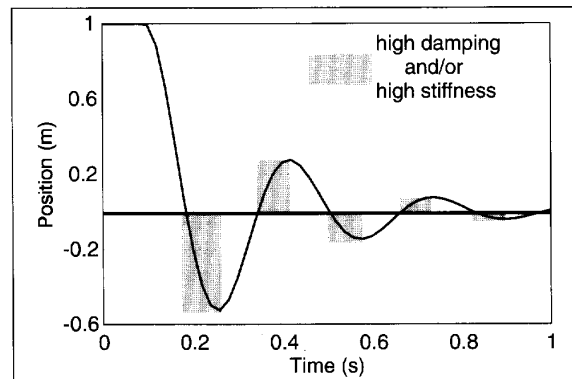
Fig. 1. Illustration of NPD control.

$$u = Ke_f + B\dot{e}_f + f_{des} + \hat{c}(\dot{x}, x) \qquad (2)$$

where K and B are constant gains, $e_f = f_{des} - f$ is the force error, f_{des} is the desired force, $\hat{c}(\dot{x}, x)$ is the estimate of $c(\dot{x}, x)$, which makes (1) close to a double integrator, and f is the actual force. NPD control's action depends on whether the system output is moving toward or away from the desired setpoint:

- Movement away from setpoint: K and B increase to stop the system.
- Movement toward setpoint: K and B decrease to minimize the residual energy at the goal.

A graphical explanation is given in Fig. 1. The position and velocity errors are from an underdamped second order system. At the shaded areas the system is moving away from the desired point and the gains should be higher. At the nonshaded areas the system is moving toward the desired point and lower gains can be used. NPD control enjoys advantages over standard PD control in terms of disturbance rejection, because disturbances create unfavorable motions. This is particularly important in CTFC because a contact transition may be viewed as a large disturbance. NPD control is different from sliding mode control [11], which is essentially high gain position control and uses maximum input to the system whether or not the system is moving toward the setpoint. Furthermore, NPD control is less sensitive to delays than sliding mode and simple PD control because shifting the shaded area slightly to the right would not adversely affect the system stability.

The nonlinear gain function for NPD control is not unique. One of the implementations [16] is shown here:

$$K_n = \frac{K_1}{1 + \beta \exp[\alpha sign(\dot{e}_f)e_f]} + K_0$$

$$B_n = \frac{B_1}{1 + \beta \exp[\alpha sign(\dot{e}_f)e_f]} + B_0 \qquad (3)$$

where K_n and B_n are the nonlinear proportional and derivative gains, K_1 and B_1 are the top values of the gains, and K_0 and B_0 are the bottom values of the gains. The constant β determines whether the function's value at the setpoint is closer to its maximum (smaller β) or minimum (larger β); typically $\beta = 1$. The constant α determines the width of the transient region from low to high gain; typically it is around 100. Fig. 2 shows a 3-D view of the nonlinear function. The actuator input is:

$$u = K_n e_f + B_n \dot{e}_f + f_{des} + \hat{c}(\dot{x}, x) \qquad (4)$$

Stability and Gain Design for Force Control

Stability for a general system (1) under NPD control is not much different from standard PD control [16]: it depends on the accuracy of the estimated model. If the model is accurate, the closed loop system becomes linear and mimics a stable mass-damper-spring. The advantage of NPD control is its nonlinear gains which reduce the disturbance effect more than simple PD control.

For position control, a rule of thumb [16] is: (1) choose the bottom values K_0 and B_0 of the nonlinear gains based on PD
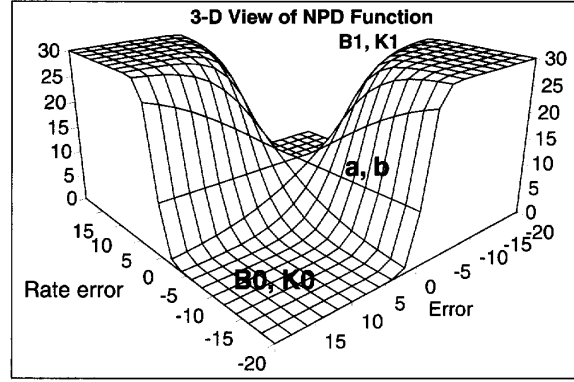

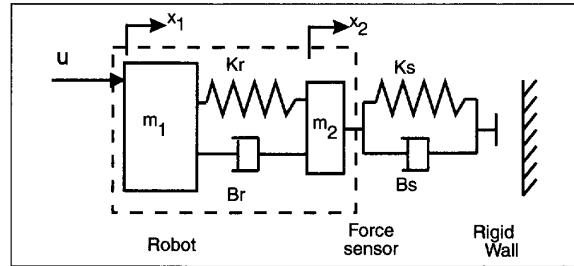Fig. 2. A 3-D view of the nonlinear gain function.

Fig. 3. Model of flexible robot and rigid environment.

control, and (2) choose the top values K_1 and B_1 as high as possible up to the point that sensor noise causes instability. Under NCTFC, we can be more specific about the gain design based on the nature of the instability, and use classic control theory to do an analysis. This analysis will assume that the robot does not lose contact with the environment. Under CTFC, the stability and performance will be assessed in simulation and experiment.

Linear analysis has shown that there are three important modes affecting force control: a rigid-body robot mode, an environmental mode, and a noncolocated actuator sensor mode due to link flexibility [2]. Even though the first two modes cause force control to be underdamped, they do not actually cause instability. It is the third mode that induces instability [1, 2]. The NPD gain design for force control will assume a rigid environment (the worst case) and a flexible robot. We will consider a one dimensional system only for simplicity (Fig. 3).

The NPD design procedure is closely related to classical PD gain design:

Step 1. Estimate the flexible robot model. This can be done using spectrum analysis.

Step 2. Find two sets of stable PD gains, conservatively and non-conservatively chosen. The conservative gains are chosen deliberately lower than those stability allows. With this set of the gains the system will be guaranteed to be stable under PD control despite the modeling error from Step 1. The non-conservative gains are chosen close to the stability boundary. Another way to find the conservative gains is to divide the non-conservative gains by a factor to insure stability of the conservative gains, using standard techniques such as root locus.

Step 3. Construct NPD gains. Use the conservative gains for the bottom values (K_0 and B_0) and the difference between the con-

Table 1
System parameters for the simulation model

m_1	119.4 kg	m_2	13.24 kg
K_r	110100 N/m	K_s	11010 N/m
B_r	10 N/ms^{-1}	B_s	10 N/m

Table II
Gains and poles from root locus design

	Conservative Gains	Non-Conservative Gains
Force Gain (N/N)	0.3	1.3
Force Rate Gain (N/Ns^{-1})	0.01	0.01
pole 1	-125.7	-125.9
pole 2,3	$-0.322\pm9.96i$	$-0.020\pm13.32i$
pole 4,5	$-0.433\pm99.9i$	$-0.625æ99.7i$

Fig. 4. Root locus plot without force rate feedback.

Fig. 5. Root locus plot with force rate feedback.

Fig. 6. Simulated NCTFC time responses: (a) contact force (b) actuator input.

servative and non-conservative gains for the top values (K_1 and B_1) in the NPD function (3). β should be chosen to be greater than one initially so that the NPD effect is reduced to insure system stability; it should be fine-tuned later to get better performance. The system stability is not very sensitive to α, which can be fixed at around 100.

If the model is accurate, then stability can be ensured because the NPD gains do not exceed the boundary of the PD gains. Since, like the PD control, the NPD control does not explicitly use a system model, the control is relatively robust to modeling errors if enough gain margin is considered in the gain design.

Because the intention of NPD control is to damp out unfavorable movements, it deviates from PD control only at those moments when the unfavorable movements occur (Fig. 1). There is no reason why the non-conservative gain values should stay in the stability boundary of the PD control. The simulation in the next section shows that even when the non-conservative gain values exceed the stability boundary of PD control, the system under NPD control is still stable and gives better performance. It remains a research challenge to find better ways to design the NPD gains beyond the method presented here, which is conservative.

Simulation

Simulations are performed on the system in Fig. 3 with parameters in Table 1 (modified from [2]).

Gain Design

To make the simulations more realistic, we assume the actuator input is prefiltered by a 20 Hz, first order, lowpass filter (dominant pole). Previous research has shown that a dominant pole is necessary for filtering out the high frequency oscillations in the system. The cutoff frequency of the filter based on our rule-of-thumb should be 5-10 times smaller than the first resonance frequency of the system (see the section on the Hardware
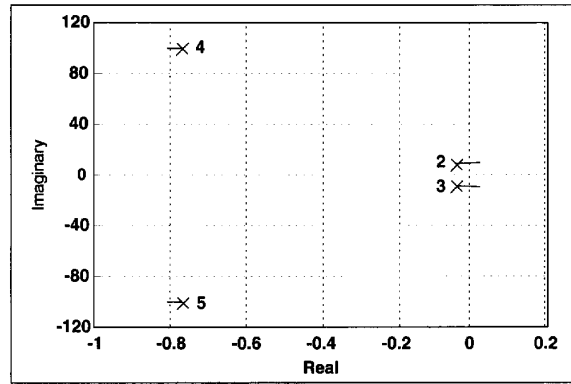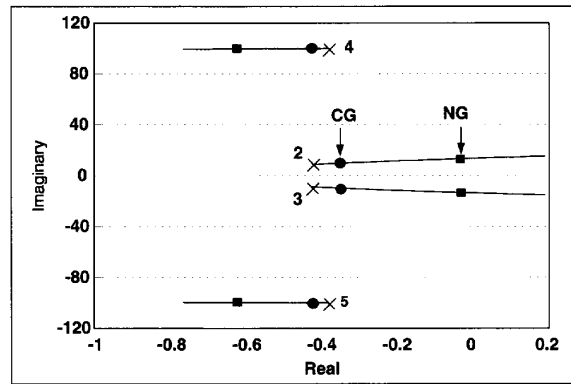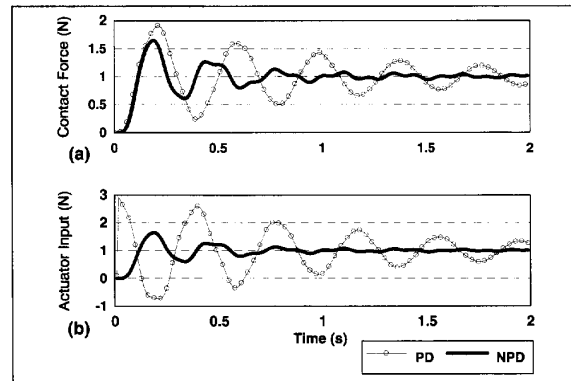
Description) or the lowest noise frequency, whichever is smaller. It will affect the feedback gain values for a given closed loop bandwidth.

Fig. 4 shows the root locus plot without force rate feedback (determined by using Matlab). Pole 1 is not visible because it is quite far to the left. Fig. 5 shows the root locus plot with force rate feedback ($0.01N/Ns^{-1}$), whose effect is to move poles 2 and 3 to the left, and poles 4 and 5 to the right. The conservative and

non-conservative gains are marked in Fig. 5. Table 2 lists the numerical values for the conservative and non-conservative gains and the pole locations. Gain B_1 is set to 0.01, which still results in stable poles based on the root locus analysis.

Non-Contact Transient Force Control

Fig. 6 compares the NCTFC time responses of the PD and NPD controllers. The contact force under PD control has a large oscillation and does not die out quickly. The NPD controller gives a better time response, as the oscillation dies out in two cycles. The actuator inputs show similar behaviors. Further tuning the PD parameters did not result in significant improvement. This is not surprising: with the non-colocated actuator and sensor, the root locus in the previous section shows that for any combination of force and force rate gains there will always be underdamped poles. However, both controllers are stable.

Contact Transient Force Control

In CTFC the robot is positioned away from the environment initially. When the contact force is zero the robot is under position control. When the force rises above a threshold (0.5 N) the control switches to force control [4]. This will cause jumps in the actuator input. We ignore this jump because the low level command filter will smooth it out. Note the initial velocity will be non-zero when the contact occurs.

Fig. 7 shows the comparison of the CTFC time responses of the PD and NPD controllers. PD control cannot establish stable contact within the time limit. NPD control establishes the stable contact after one bounce and reaches steady state. The actuator input with NPD control has a couple of spikes around the times when the contacts occur and break. These are clearly the effect of the nonlinear proportional derivative gains. For NPD control the top gain is higher than that listed in Table 2.
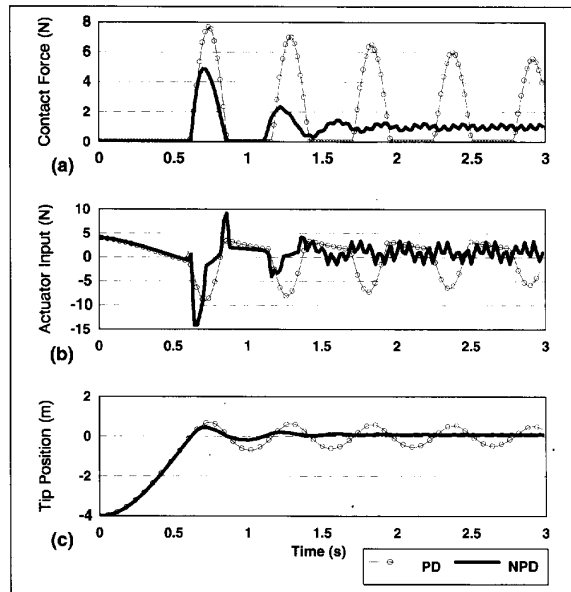

Fig. 7. Simulated CTFC time responses: (a) contact force (b) actuator input (c) robot tip position.

Experiments

PD and NPD control have been implemented on the thumb of the Sarcos Dextrous Arm. For comparison, we also tested how well humans perform in the NCTFC and CTFC tasks.

Hardware Description

The Sarcos Dextrous Arm is hydraulically powered, with high bandwidth actuators utilizing custom suspension jet pipe valves [5]. The hand has a thumb with two degrees of freedom (DOF), called the thumb lateral and vertical DOFs (Fig. 8). The lateral DOF is used in the current experiment. The thumb joint position is measured by a precision potentiometer. A beam force sensor with strain gauges is integrated in the thumb structure, approximately halfway between the joint center and the tip.

Each DOF of the arm has an analog controller, called the Advanced Joint Controller (AJC), which includes analog position and force controllers (Fig. 9). The position error and force cross gains can be set digitally. By setting one of them to zero we can obtain position control only or force control only. The valve current is filtered by an onboard, one pole lowpass filter (1.0 Hz). The open loop dynamics (relation between the current input to

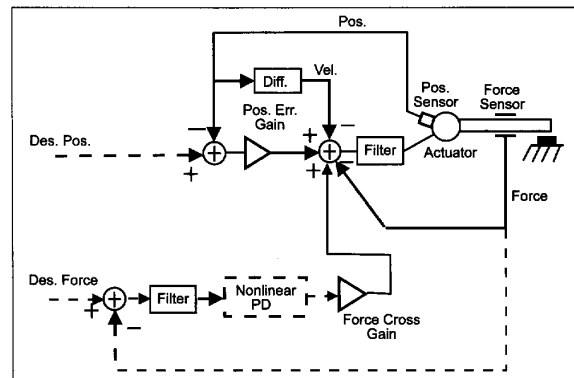Fig. 8. Sarcos Dextrous Arm hand and the environment fixed to the palm.

Fig. 9. Schematic diagram of the control for the Sarcos thumb joint. The solid lines stand for the analog control (implemented in AJC) and dashed lines for the digital control.

the valve and force output) is highly nonlinear [8], but qualitatively close to a first-order integrator. With force feedback the closed loop actuator system is quite linear and can assume very high bandwidth (on the order of 1 kHz) [8]. This inherent lowpass nature of the actuator eliminates the need to implement a dominant pole digitally for force control.

All of the computation is done in a C40 microprocessor, one of whose comports is interfaced to the AJC's digital IO. The C40 resides in a PC 486 50MHz host computer. The position and force signals are sampled through a 12 bit-ADC on the AJC. The digitally implemented servo runs at 1200 Hz.

Both PD and NPD controllers were tested. In CTFC the position control phase relies on the AJC's position controller; we only feed in desired positions. After contact occurs and the contact force rises over 0.5 N, the position controller is disabled and the digital force controller is enabled. The analog force control loop is always active. The force error is filtered by a lowpass filter with a cutoff frequency at 30 Hz, which attenuates the force error rate due to noise by 8 times. The force rate error is estimated by differentiating the filtered force error signal. The computational need is negligible (less than 1 μs with the C40). The servo rate is limited by the interface with the AJC controller.

To estimate the first bending mode of the thumb in the lateral direction, we mounted a high-precision accelerometer (Bruel and Kaejr type 8001, 20 kHz resonance) to the tip of the thumb. A light knock was applied to the thumb. The acceleration signal was amplified by a charge amplifier (Bruel and Kaejr type 2535). The data were collected and analyzed by using a Dynamic Signal Analyzer (Hewlett Packard type 35665A). Fig. 10 shows the normalized spectrum plot. There is a clear peak at 144 Hz. Because the force sensor is integrated in the thumb we assume that it has the same resonance.
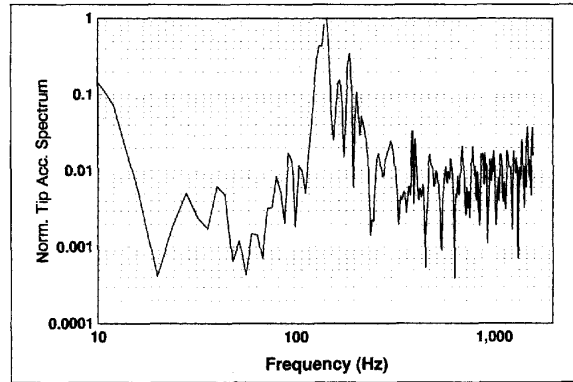

Fig. 10. Experiment: normalized tip acceleration spectrum.

Using a two mass lumped parameter model as in Fig. 3 and a 0.01 damping ratio, and assuming the force sensor has the same stiffness, we followed the steps of the gain design method described in the previous section to calculate the gains. The gains obtained are not optimal because we did not do an exhaustive search. For the conservative gain, the force gain is 1.005 N/N and the force rate gain is 1.001 N/Ns^{-1}. The digital control loop gains for PD control are 0.005 N/N and 0.001 N/Ns^{-1}. For NPD control, K_1 and K_0 are both set to 0.005 N/N. B_0 is 0.001 N/Ns^{-1} while B_1 is set to 0.002 N/Ns^{-1}. With the maximum NPD gains the PD control was stable based on the estimated model.

Non-Contact Transient Force Control

Fig. 11 compares the NCTFC time responses for PD and NPD control. In all of the experiments the desired force is 4 N. The contact forces are similar except that the NPD response settles a little faster. The force rate response is similar to the force. The actuator input shows clearly that initially the NPD control has spikes which indicates the action of the nonlinear gains. In the experiments we tuned up the gains for both PD and NPD. The force gain can go to 0.04 N/N and the force rate gain to $0.003N/Ns^{-1}$ without instability.

Contact Transient Force Control

Fig. 12 compares the CTFC time responses for PD and NPD control. The contact force shows that PD control is not stable: the robot bounces on the environmental surface. The force, force rate, tip velocity, and actuator input all show a limit cycle behavior. In comparison, NPD control is stable. The contact force response shows that after the first two bounces, stable contact is established. Initially there are several spikes in the actuator input for the NPD control, which are due to the nonlinear gains. The force rate and tip velocity are smaller than those of PD control, as NPD control dampens the force rate more.

The tip position measurements are not accurate as can be seen from Fig. 12(d). There are only 4 bits accuracy there. This indicates that in force control with a stiff environment, the position as well as the velocity estimated from the position are not suitable for control. An and Hollerbach [1] noticed this, and concluded that the limited sensor resolution prohibits the online identification of the environment.
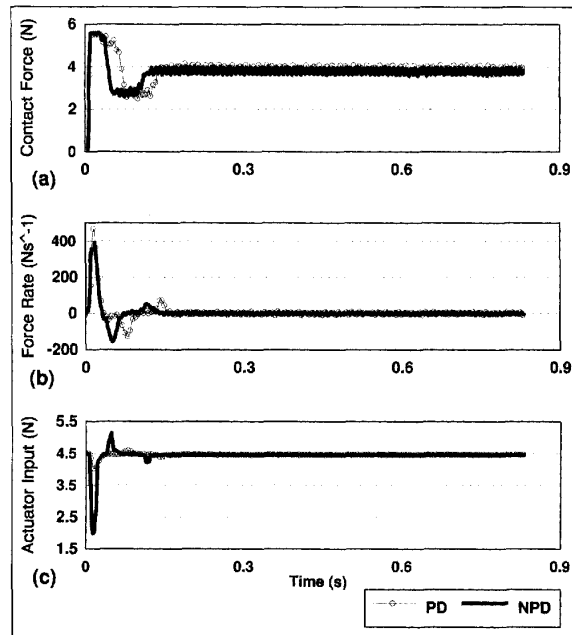

Fig. 11. Experimental NCTFC time responses: (a) force steps (b) force rates and (c) actuator inputs.
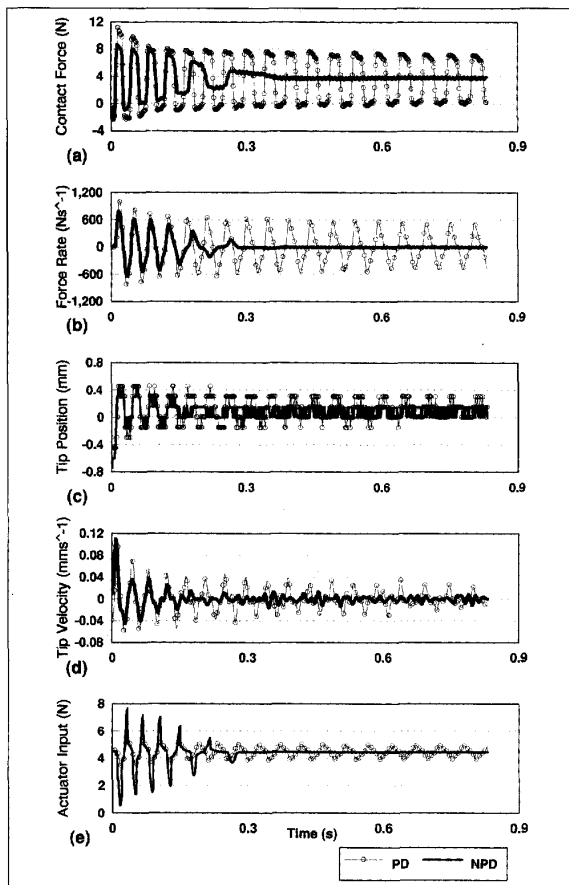
Fig. 12. Experimental CTFC time responses: (a) force steps (b) force rates (c) tip positions (d) tip velocities (e) actuator inputs

Contact and Non-Contact Transient Force Control by Humans

For interest, we did NCTFC and CTFC experiments on humans. Each subject was asked to hold one end of an aluminum rod 12 mm in diameter and 200 mm long, and to use the other end to apply force to a load cell (OMEGA, LCC-200) also with an aluminum surface. The tip position of the rod was measured by an Opotrak system (Northern Digital Inc., Waterloo, Ontario) with a sampling frequency of 500 Hz and accuracy of 0.05 mm. The force signal was sampled at 2000 Hz and later filtered digitally with a 100 Hz cutoff frequency. In the experiment the force signal was displayed to a subject by a scope, and the subject was allowed to observe the impact. The task was to apply a step force to the load cell with or without contact transition as accurately as possible. Furthermore, subjects were instructed to use the wrist only. Four subjects participated in the study. Since results from all the subjects were similar, we show the results from one subject only.

Fig. 13 shows the human NCTFC time response. The position is not changed at all, as it should not be. The force response has a clear overshoot. Both the rise and the settling times are larger than for the robot force control in Fig. 11. Fig. 14 shows the human CTFC time responses. The tip of the rod was initially a half meter from the surface of the load cell. Fig. 14 only shows

a segment of the data in order to see the details when the contact occurs. There is only one bounce. The velocity right before the contact is higher (1.2 mms^{-1}) than that of the robot (0.12 mm^{-1}, Fig. 12), yet stable contact is established quickly.

Discussion

In this paper we described an NPD control method. Even though we emphasized force control, the method is equally applicable to position control [16]. The method has the advantages of higher disturbance rejection and less sensitivity to delays than simple PD control. We proposed a gain design method for the NPD controller for force control by using the root locus technique. The design finds two stable sets of gains for PD control; one set is conservative and another not. They are then used as the bottom and top gains in the nonlinear PD gain function. The method is in general conservative, especially for the force rate gain. Other possible approaches for the gain design include frequency methods, such as Bode and Nyquist plots, and the energy method. The frequency methods may not be better than the root locus method because they all design gains based on the steady-state behaviors of a system. The energy method may lead to optimal gain design because we can examine the total energy loss in an oscillation cycle and minimize it accordingly; this is a current research topic.

Simulations and experiments both confirm stable control with the designed gains. We discovered experimentally the instability of the PD controller in contact transient control to a rigid environment even though the closed loop system is stable without contact transition. The NPD control is stable with contact transition; it settles down to a steady state quickly after one or two bounces.

In these force control experiments we observed five difficulties associated with force control: (1) the dynamics of the system before and after contact can change dramatically, especially with a stiff environment, (2) the transition can cause instability, even though the controller designed for force control is stable, (3) the position sensor resolution limits its use for force control when the robot is handling a stiff environment, (4) the contact force rate information is very important, even though the force signal is noisy, and (5) input command prefiltering is essential, as it slows down the overall system so that the actuator has time to react to force change; furthermore it filters out the high-frequency oscillations in the system. These observations are subject to further test. Some of them have been reported before, such as the importance of the dominant pole [1, 9].

Because of limitations in sensor resolution and the discontinuous change of the system dynamics, traditional control design may have difficulties. Nonlinear control can play an important role. In the world of manipulation a robot controller must not just handle only one type of object, but any kind like the human hand. The design of the NPD controller presented in this paper is just a beginning. The gain design is still conservative; the exact boundary of the stability is unknown for a given system, and our design is only one possibility. For the future research, we intend to consider control delay, friction, and high order systems for the CTFC with the NPD controller.

Experiments on human subjects clearly show that the biological system performs better than the robot in force control. Instability is nowhere seen in the contact transition with the human motor control system. Vision may play a role in overall stability,
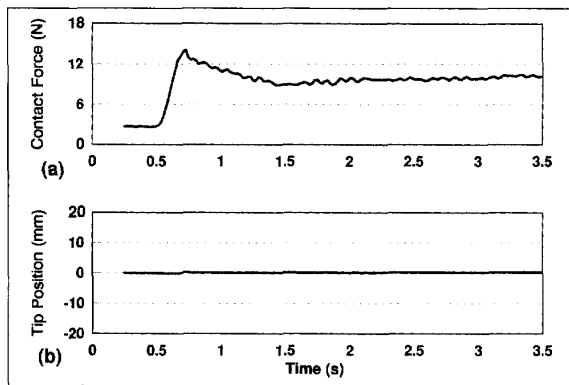
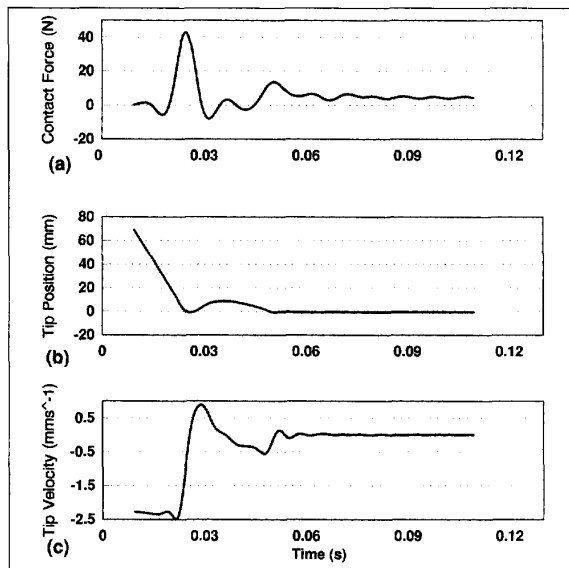Fig. 13. Experimental NCTFC time responses for a human subject.

Fig. 14. Experimental CTFC time responses for a human subject.

but its role must be very limited during initial transition because there is a significant delay in the visual pathway of over 100ms. Even though we do not propose that NPD control is used, it may be concluded that simple PD control is not used by humans. Understanding human force control is certainly an interesting research direction.

References

[1] C.H. An, C.G. Atkeson, and J.M. Hollerbach, *Model-Based Control of a Robot Manipulator*. MIT Press, Cambridge, MA, 1988.

[2] S.D. Eppinger and W.P. Seering, "Understanding Bandwidth Limitations in Robot Force Control," in *Proc. IEEE Int. Conf. Robotics and Automation*, Raleigh, pp. 904-909, March 31-April 3, 1987.

[3] N. Hogan, "Stable Execution of Contact Tasks Using Impedance Control," in *Proc. IEEE Int. Conf. Robotics and Automation*, Raleigh, pp. 1047-1054, March 31-April 3, 1987.

[4] J.M. Hyde and M.R. Cutkosky, "Contact Transition Control: An Experimental Study," in *Proc. IEEE Int. Conf. Robotics and Automation*, Atlanta, pp. 363-368, May 2-6, 1993.

[5] S.C. Jacobsen, J.E. Wood, D.F. Knutti, and K.B. Biggers, "The Utah/MIT Dextrous Hand: Work in Progress," *International Journal of Robotics Research*, 3, no. 4, pp. 21-50, 1984.

[6] O. Khatib and J. Burdick, "Motion and Force Control of Robot Manipulators," in *Proc. IEEE Int. Conf. Robotics and Automation*, San Francisco, pp. 1381-1386, April 7-10, 1986.

[7] J.Y.S. Luh, M.W. Walker, and R.P. Paul, "Resolved-Acceleration Control of Mechanical Manipulators," *IEEE Transactions on Automatic Control*, AC-25, pp. 195-200, 1980.

[8] T.W. McLain, E.K. Iversen, C.C. Davis, and S.C. Jacobsen, "Development, Simulation and Validation of a Highly Nonlinear Hydraulic Servosystem Model," in *Proc. 1989 American Control Conference*, Pittsburgh, pp. 385-391, June 1989.

[9] H.P. Qian and J. De. Schutter, "Introducing Active Linear and Nonlinear Damping to Enable Stable High Gain Force Control in Case of Stiff Contact," in *Proc. IEEE Int. Conf. Robotics and Automation*, Nice, pp. 1374-1379, May 12-14, 1992.

[10] J.K. Salisbury, "Active Stiffness Control of a Manipulator in Cartesian Coordinates," in *Proc. 19th IEEE Conf. Decision and Control*, Albuquerque, pp. 95-100, Dec. 1980.

[11] J.-J.E. Slotine and S.S. Sastry, "Tracking Control of Nonlinear Systems Using Sliding Surfaces, With Applications to Robot Manipulators," *Intl. J. Control*, S8, pp. 465-492, 1983.

[12] R. Volpe and P. Khosla, "A Theoretical and Experimental Investigation of Impact Control for Manipulators," *Intl. J. Robotics Research*. 12, no. 4, pp. 351-365, 1993.

[13] K. Youcef-Toumi and D. Li, "Force Control of Direct-Drive Manipulators for Surface Following," in *Proc. IEEE Int. Conf. Robotics and Automation*, Raleigh, NC, pp. 2055-2060, March 31-April 3, 1987.

[14] K. Youcef-Toumi and D.A. Gutz, "Impact and Force Control," in *Proc. IEEE Int. Conf. Robotics and Automation*, Scottsdale, pp. 410-417, May 14-19, 1989.

[15] K. Youcef-Toumi and D.A. Gutz, "Impact and Force Control: Modeling and Experiments," *J. Dynamic Sys. Meas. and Cont. 116*, pp. 89-98, 1994.

[16] Y. Xu, D. Ma, and J.M. Hollerbach, "Nonlinear Proportional and Derivative Control for High Disturbance Rejection and High Gain Force Control," in *Proc. IEEE Int. Conf. Robotics and Automation*, Atlanta, pp. 752-759, May 2-6, 1993.

[17] Y. Xu, J.M. Hollerbach, and D. Ma, "Force and Contact Transient Control Using Nonlinear PD Control," in *Proc. IEEE Int. Conf. Robotics and Automation*, San Diego, pp. 924-930, May 1994.

Yangming Xu received the B.S. degree from Northwestern Polytechnical University, China in 1982, and the S.M. and Ph.D. degrees in aeronautics and astronautics from the Massachusetts Institute of Technology, Cambridge, in 1986 and 1991, respectively. From 1991-1994, he was a postdoctoral fellow in Biorobotics Laboratory at McGill University, Montreal, Quebec, Canada. He is currently with Sarcos Research Corporation, Salt Lake City, Utah.

John M. Hollerbach is professor of computer science at the University of Utah. From 1989-1994 he was the Natural Sciences and Engineering/Canadian Institute for Advanced Research Professor of Robotics at McGill University, jointly in the Departments of Mechanical Engineering and Biomedical Engineering. From 1982-1989 he was on the faculty of the Department of Brain and Cognitive Sciences and a member of the Artificial Intelligence Laboratory at MIT; from 1978-1982 he was a research scientist. He received his BS in chemistry ('68) and MS in mathematics ('69) from the University of Michigan, and SM ('75) and PhD ('78) from MIT in Electrical Engineering and Computer Science. In 1984 he received an NSF Presidential Young Investigator Award; in 1988 he was named a fellow of the Canadian Institute for Advanced Research.

Donghai Ma received his B.Eng. degree in Instrumentation from Tsinghua University, China, in 1990. He is completing an M.Eng. in mechanical engineering at McGill University, Montreal, Canada. His interests include autonomous calibration, nonlinear control, and robotics.