AN ANALYSIS OF CONTACT INSTABILITY IN TERMS OF PASSIVE PHYSICAL EQUIVALENTS

Ed Colgate
Northwestern University
Department of Mechanical Engineering
Evanston IL, 60208

Neville Hogan
Massachusetts Institute of Technology
Department of Mechanical Engineering
Cambridge MA, 02139

Abstract—This paper explores the sources of and solutions to robot contact instability associated with force feedback. The behavior of several linear robot models is analyzed in terms of the properties of their admittances. Using the novel technique of "passive physical equivalents," an explanation for the often-observed instability of a force-controlled robot contacting a stiff surface is offered, and it is shown that a fundamental limit exists to the efficacy of any force feedback controller implemented on a robot with non-colocated actuators and sensors. Suggestions for improved force control involving both mechanical design and compensator design are also presented.

1 Introduction

The last several years have witnessed explosive growth in the study of robot force control. The development of successful strategies and implementations for force control is seen as a crucial step in enabling robots to perform tasks, such as deburring, drilling, and parts assembly, which require significant interaction with the environment. The implementation of high-bandwidth, high-accuracy force control, however, has proven to be quite difficult, largely due to stability problems that occur upon contact with a rigid surface. It is "contact instability" which is the subject of this paper.

Whitney [20] was the first to provide a stability analysis of a force controlled manipulator. He modeled a manipulator as a velocity-input integrator, and assumed that proportional position, velocity, and force feedback were implemented in discrete time. He modeled the environment as a spring and derived the following stability result:

$$0 < TGK_e < 1$$

where T is the sampling rate, G is the force feedback gain, and K_e is the combined stiffness of the sensor and the environment. For fixed T, this indicates that a tradeoff exists between G and K_e. In other words, high bandwidth force control requires a compliant sensor or environment.

In the years since Whitney published this result, the essential stability tradeoff between force feedback gain and stiffness of the environment[1] has been substantiated many times [6, 7, 9, 12, 14, 21]. It is now known that the problem transcends the digital implementation, that even for analog control the tradeoff exists [21].

[1]The stiffness of the sensor is generally included as part of that of the environment.

In a recent review [13], Paul indicated that the contact problem with a rigid manipulator, rigid sensor, and rigid environment is still unsolved.

Various explanations for contact instability have been offered. Kazerooni [12] and An [1] have both shown that the tradeoff may be attributed to unmodeled dynamics. Eppinger and Seering [5, 6, 7] have carried out extensive stability analyses based on a series of increasingly sophisticated single-axis models. They have shown that stability problems arise due to the *non-colocation*[2] of sensors and actuators. There are two important consequences of their observations: first, that it is specifically those dynamic elements (modeled or not) separating the actuator and sensor that are most seriously implicated in contact instability; second, that, even if the robot is exquisitely modeled, high gain force control is a challenging proposition.

The analyses in this paper support the major conclusions of Eppinger and Seering; however the analytical approach is quite different. Our approach has been to focus on the properties of the robot impedance, as motivated in a number of previous publications [4, 10]. Included in this approach is the use of "passive physical equivalents," a novel technique that allows, even encourages, the simultaneous consideration of mechanical design and controller design. An important consequence of this approach is that it facilitates the generation of new approaches to improved force control. To begin, however, the next two sections will be used to clarify the stability and performance criteria used in this paper.

2 Coupled Stability

In this section several theoretical results will be presented regarding the stability of linear feedback controlled systems interacting with passive environments. The restriction to linear systems makes the mathematics tractable, allowing the development of useful analytical tools. Some restriction must also be placed on the class of allowable environments. It is felt that the restriction to passive environments is the most natural choice, as a great number of the environments with which robots interact *are* passive. Furthermore, it can also be shown that many of the results described here apply to a broad class of active environments [3, 4].

All proofs and derivations are omitted due to space limitations; they may be found in [3]. The following definition of linear 1-

[2]The importance of non-colocation in the stability of feedback systems was first noted by Gevarter [8] in the context of flexible space vehicles.

port[3] passivity will prove useful:

Passivity. A linear, time-invariant 1-port is passive iff:

1. $Z(s)$ has no poles in the right half plane.
2. $Z(s)$ has a Nyquist plot which lies wholly within the closed left half plane.

Where $Z(s)$ is impedance or admittance at the port. A function $Z(s)$ which satisfies these conditions is termed *positive real.* With this definition and an application of the Nyquist stability criterion, the following result may be found [3]:

Coupled Stability 1. A necessary and sufficient condition to ensure the stability of a linear, time-invariant (LTI), stable plant coupled at a single interaction port to an arbitrary stable, passive environment (linear or nonlinear), is that the impedance (or admittance) of the plant be positive real.

In other words, the plant, although feedback controlled, should have the interaction port behavior of a passive system. In fact, it can be shown that, given a positive real impedance, some passive system always exists which has that impedance [2]. Furthermore, such systems can be derived from the impedance using the techniques of network synthesis [19]. Therefore, a *passive physical equivalent* [3] can be defined to be a passive system which has the same port behavior as a feedback controlled system with a positive real impedance.

Finally we note an alternative coupled stability criterion [3, 16] which is useful in identifying certain types of environments that lead to instability:

Coupled Stability 2. A necessary and sufficient condition to guarantee the stability of a stable LTI plant coupled at a single interaction port to an arbitrary stable, passive environment, is that the plant be stable when coupled to all linear masses and all linear springs.

It is also shown in [3] that this criterion can be examined by generating the root loci of $sZ(s)$ and $Z(s)/s$, where $Z(s)$ is the impedance or admittance of the plant. One root locus[4] is parametrized by the mass of the environment, the other by the stiffness. If all branches of both root loci remain in the closed left half plant, coupled stability 2 is satisfied. Thus, pure springs and masses, although they are not the only environments which can cause instability, are termed the "worst environments".

3 Coupled Performance

Specifications such as bandwidth and steady state error are typically used to evaluate the performance of servo systems. Such specifications, however, make no sense in a force control problem unless a specific model of the environment is chosen so that force can be measured as an output. An alternative is to express performance specifications in terms of the properties of the robot's admittance, which is well-defined with or without an environment.

For instance, an ideal force control robot would act as an effort source, exhibiting infinite admittance. A real robot, however, cannot be free of dynamics; it must at least have some mass! Consider a linear mass—its admittance is $Y(s) = 1/Ms$, which is finite (except at d.c.). The admittance, however, is scaled by M—the smaller M, the larger the admittance at a given frequency and the higher the frequency corresponding to a given magnitude of $Y(s)$.

But mass is not the only determinant of force control performance. There are likely to be other dynamic effects such as joint friction, and nonlinearities, such as motor cogging, which cause steady state error (that mass alone cannot) and poor force resolution. Such effects tend to be dominant at low frequencies. Like the mass, however, these effects are manifest in the robot admittance and the solution is generally to increase the admittance as much as possible (perhaps with an emphasis on low frequencies) without jeopardizing coupled stability.

In summary, it is achieving the optimal tradeoff between admittance amplification and robustness to interaction that characterizes the force control design process.

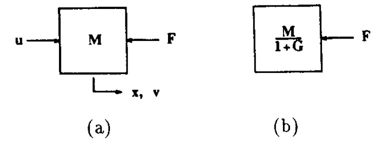

Figure 1: (a) Rigid body robot model. (b) Passive physical equivalent for $u = -GF$.

4 Analysis of Lumped Models

In this section the technique of passive physical equivalents is used to analyze the coupled stability properties of several simple robot models. Due to space limitations, derivation of the equivalents is not presented, however, an example may be found in [3].

The first robot model to be considered is shown in Figure 1(a). It is a mass M acted on by an actuator force u and an environmental force F. This model is extremely simple, yet it captures the essential dynamic behavior of a robot (in particular, a direct drive robot), at least at low frequency. A basic tenet of "impedance control" is that dynamic interaction with the world is a fundamental feature of manipulation. In keeping with this idea, we will attempt to use in this paper the ***simplest models competent to represent the essential dynamic behavior.*** In this case, if $u = -GF$ is selected as the control law, the passive physical equivalent is a mass $M/(1+G)$ acted on by the environmental force F (Figure 1(b)).

Figure 1(b) is extremely valuable insofar as it represents the relationship between force feedback and robot inertia. As described

[3] A port is a physical location where a system can contact its environment. A port is characterized by power variables, which, in the mechanical realm, are generally taken to be force and velocity. For a 1-port, these variables are scalars.

[4] Which one depends on whether the impedance or admittance causality has been chosen for the plant.

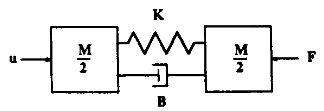

Figure 2: Manipulator with non-colocated actuation and sensing.

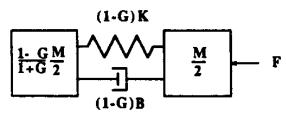

Figure 3: Passive physical equivalent of the two-mass model.

in detail in [11, 21], negative force feedback acts to reduce the robot's inertia; conversely, *positive* force feedback acts to increase the robot's inertia. Note that $G \geq -1$, corresponding to an equivalent mass less than infinity, is necessary for the closed loop system to be passive.

The rigid body model, however, does not predict the contact instability that has been observed so frequently. This is not very surprising, as the actuator and sensor are colocated. Therefore, the non-colocated model shown in Figure 2 is considered next. In this model, K and B generally represent structural dynamics, although they may represent other effects, such as actuator or transmission dynamics. Note that the overall mass of this manipulator is the same as that of the first model.

A passive physical equivalent for this manipulator, given the same control law, $u = -GF$, is shown in Figure 3. Note that the endpoint mass, that acted on by the environment, is unchanged. In the language of passive physical equivalents, this endpoint mass is referred to as an "uncontrollable element," as no causal controller can affect its value [3]. Despite the fact that this mass does not change, one would expect this model to predict the same low-frequency behavior as the rigid body model; thus, the total mass of the realization should be $M/(1+G)$. An examination of Figure 3 shows that this is the case. A consequence of this low frequency behavior is that the the mass near the actuator exhibits a very strong dependence upon G.

Clearly, the force feedback gain is restricted to lie between -1 and 1 for this to be a passive realization. Large negative force feedback gains ($G > 1$) would result in a negative spring, mass, and damper. Negative system elements indicate a violation of the coupled stability criterion and, therefore, the potential for instability when coupled to some passive environments. One way to determine which passive environments cause difficulty is to use the root locus test described in Section 2. These loci, for example model parameter values and two different values of gain, are shown in Figure 4. The result is that, for $G > 1$, stiff environments are destabilizing, a finding that is consistent with observation. Closer inspection reveals that the environmental stiffness must be on the order of the structural stiffness of the robot to create the instability. This does not weaken the result significantly, because many potential workpieces (e.g., an engine block) are at least as stiff as commercially available robots.

To summarize, the behavior of this system may be attributed to

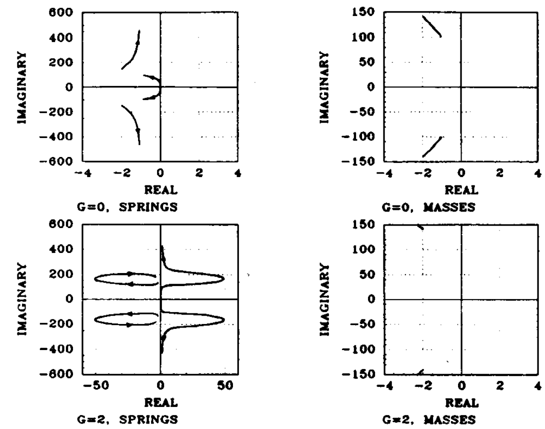

Figure 4: Worst environment root loci, $G = 0$ (open loop) and $G = 2$. Plant parameters are $M = 2$, $B = 2$, and $K = 10,000$.

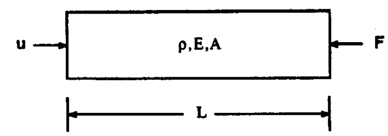

Figure 5: Uniform distributed model of a robot.

one, the "uncontrollability" of the endpoint mass, and two, the preservation of the rigid body low frequency behavior. In light of the apparent importance of the structural dynamics, however, it is reasonable to question whether this intuitive explanation and the quantitative limits on G will hold for higher-order robot models. Passive physical equivalents for three-mass and four-mass uniform[5] models are found to support the qualitative and quantitative behavior shown by the two-mass model. They further indicate that springs, masses, and dampers are all affected by the same scaling factors near the ends of the structure. The logical culmination of this line of investigation is a distributed beam model, which is treated next.

5 Analysis of a Distributed Model

In this section, the uniform distributed model of a robot shown in Figure 5 is analyzed. The compression/tension beam has uniform density ρ, modulus of elasticity E, area A, and length L. As before, the control law is $u = -GF$. This is a better robot model than any of the lumped parameter versions previously analyzed, but is not as realistic as a cantilever beam. It will be seen, however, that the primary value of the distributed model is in elucidating the role of transmission delays—toward such an end, the compression/tension beam is sufficient. Further, it is much easier to analyze than a cantilever beam, and is the natural extension to the lumped parameter models of the last section.

The network synthesis procedures become decreasingly tractable as the order of the system increases; using them to analyze an

[5]Uniform means that all springs have the same value, all masses do, and all dampers do. Transmission lines are frequently modeled as uniform dynamic chains.

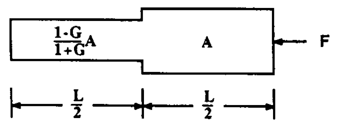

Figure 6: Passive physical equivalent of the uniform beam model.

infinite dimensional system would be asking too much. However, the insight that was created by the analyses of the last section may be used to guess a solution, which may then be checked by assuming a sinusoidal environmental input, $F = f\sin(\omega t)$, and showing that the endpoint behaviors of the controlled system and the realization match for all f and ω. The realization is shown in Figure 6.

Note that, whereas the element values of the lumped parameter models were scaled by feedback, the *area* of this model is scaled. In fact, this is done simply for the purposes of illustration, as both the mass per unit length and stiffness per unit length of a distributed beam are proportional to cross-sectional area. We could just as well have kept the area constant and have indicated that ρ and E were scaled by $(1-G)/(1+G)$.

This realization has the expected low frequency behavior—that the total mass is scaled by a factor of $1/(1+G)$. In addition, this realization has what should be the expected high frequency result—that the half of the beam close to the endpoint is unaffected by feedback. That this should be the expected result is explained by considering transmission delays. If the transmission time over one length of the uniform beam (Figure 5) is T, and if the sensing and actuation are instantaneous (as has been assumed throughout), then there will be no effect at the endpoint of the beam due to the control until time T has elapsed since the first effect due to the environment (assumed to occur at $t = 0$). If the realization is to be consistent with this, which it must be to have the correct high frequency behavior, then there can be no change in the response to an endpoint disturbance until time T has passed. This requires that the realization be uniform over half its length, and that at its midpoint, a finite change occur. In this way, there will be a reflection from the midpoint, and a response will be felt at the endpoint after a delay of $T/2 + T/2 = T$.

This argument based on time delays is independent of the particular control law that is used, so long as the control law is causal. Thus we have a fundamental result—the best that can be done without violating the positive real condition, given a uniform robot model, is to reduce the mass (or amplify the admittance) by a factor of two. In fact, this is rather optimistic, as typical digital controllers have lengthy computational delays of a few milliseconds, and analog controllers exhibit finite delays as well.

6 Dynamic Compensation

In the last section it was demonstrated that, for the uniform robot, no controller could do better than mask half the mass without violating the positive real condition. But we may ask if it is possible to violate this condition in such a way (e.g., only at low frequency) as to cause no difficulty with stiff environments, while improving low frequency admittance amplification.

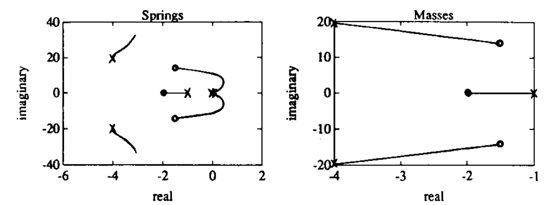

Figure 7: Worst environment root loci for first-order lag force feedback compensator and two mass plant. Parameters are $M = 1$, $B = 2$, $K = 100$, $G = 1$, and $\tau = 1$.

A reasonable place to start would be with the non-colocated two-mass model (Figure 2), and a first-order lag compensator:

$$u = -\frac{G}{\tau s + 1}F$$

Unfortunately, there is no "good" passive physical equivalent, meaning that no version of the equivalent makes the closed loop structure and the dependence upon G and τ apparent. It should be understood that this is because we are seeking a *symbolic* realization. If plant parameters, a range of values for G and τ, and appropriate software were available, passive physical equivalents would continue to be used. For the current purposes, however, a straightforward analysis of the closed loop admittance will suffice. The following points summarize this analysis:

- At sufficiently low frequency, the admittance violates the coupled stability criterion for any combination of $G > 0$ and $\tau > 0$. The consequence is that the closed loop system will always exhibit instability upon coupling to soft environments, as illustrated by the worst environment root loci in Figure 7.
- The centroid of the root locus for coupling to springs always lies in the left half plane ($\sigma_{centroid} = -B/2M$), exactly where it lies in the case of no control, regardless of the values of τ and G. Because there are only two asymptotes in this case (see Figure 7), the closed loop system is always stable when coupled to very stiff environments, so long as the admittance contains no nonminimum phase zeros. This seems contrary to the results of Eppinger and Seering [7].
- There are gains, G, for which there are no nonminimum phase zeros. This is a possible explanation for the success that An [1] and Youcef-Toumi [23] report with this compensator.
- If the simplifying assumptions $2B^2 \ll KM$ (low structural damping) and $1/\tau \ll \sqrt{2K/M}$ (control bandwidth well below first structural resonance) are made, then the following conditions guarantee minimum phase zeros:

  $$\tau > 0, \quad -1 \leq G \leq 2B\tau/M$$

  Furthermore, $2B\tau/M > 1$. Thus, for sufficiently low bandwidth control, significant low frequency admittance amplification is possible.

In conclusion, this form of compensation shows promise for stable control of the force exerted on a rigid surface, but will not ex-

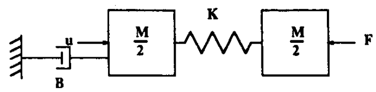

Figure 8: Two-mass model for the analysis of integral force control.

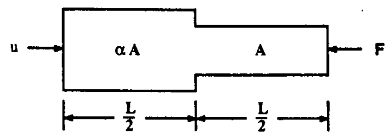

Figure 9: Non-uniform distributed beam model.

hibit high bandwidth or robustness to interaction with all passive environments.

Integral control ($u = -(G_I/s)F$) is another form of compensation which has received some attention in the past, particularly because of the promise of zero steady state error. Of course, since integral action should mask *all* inertia at low frequency, it is not reasonable to expect a passive physical equivalent. In fact, it is fairly easy to show that integral control applied to any of the models treated so far will lead to instability when coupled to environments of any stiffness.

In order for integral control to satisfy the positive real condition at all, the robot model must include some damping to ground. One such model is shown in Figure 8. Structural damping is left out of this model for simplicity's sake. Again, there is no good realization; nevertheless, the bottom line is passivity, and analytic limits on the values of G_I which will satisfy the positive real condition may be found. The limits are:

$$0 \leq G_I \leq \frac{B}{M}$$

At low frequencies, the admittance is $Y(s) = G_I/Bs$; thus, at best, the apparent mass of the robot is M. In other words, integral control results in no mass reduction at all without jeopardizing coupled stability, but it does result in high admittances at low frequencies. Note that the gain G_I may also be increased by increasing the viscous damping to ground, B.

7 Non-Uniform Manipulators

With a fundamental limit on what control can achieve in terms of masking a manipulator's inertia, an alternative approach to the force control problem is redesign of the manipulator. In this section we consider one type of redesign that fits neatly into the context of the models considered so far.

The non-uniform model shown in Figure 9 is considered. The area has been increased by some factor $\alpha > 1$ over the half close to the actuator. The reason for doing this is to offset the very strong dependence on the force feedback gain G of that half. Given such a manipulator and the control law $u = -GF$, the realization is shown in Figure 10.

The limits on G are now $-1 \leq G \leq \alpha$. There are two ways to

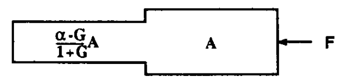

Figure 10: Passive physical equivalent to non-uniform beam model.

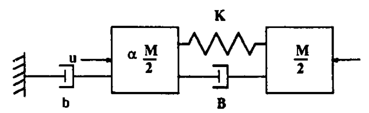

Figure 11: Two-mass model with damping to ground to illustrate the effect of joint friction.

interpret this result in terms of robot design. One is simply that it pays to put the bulk of the mass near the actuator, leaving the distal part of the robot quite light[6]. The second interpretation is that it would be useful to add weight near the actuator, even if nothing is done to the weight near the endpoint. No more mass could be masked than before, but a higher gain G could be used. If effects such as joint friction are present, the larger gain will be useful in low frequency admittance amplification (see Section 3).

To illustrate this, consider the model in Figure 11. Here, a viscous damper (b) plays the role of friction. Given the control law $u = -GF$, the low frequency admittance is $(1 + G)/b$. Thus, increasing G decreases the apparent damping. Although the addition of this damper complicates the analysis somewhat, the result is that the upper limit on G is approximately equal to α.

In conclusion, the act of adding mass near the actuator can improve the steady state response by allowing larger gains, which have the effect of reducing the apparent joint friction. It should also be noted that the integral control described in the last section is theoretically capable of entirely eliminating viscous joint friction, but not while simultaneously reducing the inertia and providing a guarantee of coupled stability. The choice of an approach will depend upon actual joint characteristics as well as the relative importance of low frequency and high frequency admittance amplification.

8 Summary: Improving Force Control

The most important point in this paper has been that force feedback controllers are subject to a fundamental performance limitation: given any manipulator with non-colocated actuation and force sensing, there is a limit to the amount of inertia reduction which can be achieved without risk of coupled instability. But the various analyses have also suggested a number of potential routes to improved force control. The object of this section is to summarize these routes.

If mass reduction is the primary design goal, it may be more fruitful to consider *design of the hardware* rather than design

[6]Note that the robot end-effector is typically beyond the force sensor (usually a wrist sensor), and therefore does not add to the weight near the endpoint; i.e., this strategy does not prohibit end-effectors.

of the controller, because of the fundamental limit described in Sections 4 and 5. To amplifiy the high frequency admittance, it is necessary to decrease physically the mass near the endpoint. One viable way of doing this is to use a macro/mini design—a small robot mounted on the end of a larger robot—and to close the feedback loop around the mini, thus reducing the mass to be masked [15, 17].

If low frequency admittance amplification or steady state error is the primary goal, it is useful to achieve high gain. For instance, integral control, while it does not result in any mass reduction, ensures that the manipulator behave as a perfect force source at d.c., so long as some amount of viscous damping is present at the joints. High gain and low steady state error can also be achieved without sacrificing mass reduction by adding mass near the actuator, as described in Section 7. We make special note of the fact that this is an approach to improved force control which would likely have gone unnoticed without the use of passive physical equivalents. Futhermore, it is an approach which requires redesign of the plant rather than redesign of the controller, a result which is quite natural in the context of passive physical equivalents.

Finally, of course, one need not use force feedback to implement force control. A conventional treatment of joint friction as well as other nonlinearities, such as motor cogging, is to implement a torque control loop around each joint [22]. Although this eliminates structural dynamics as a source of non-colocation, it does not eliminate actuator or transmission dynamics, nor does it mask any link inertia.

Another approach is to use impedance control to establish a known dynamic relation (impedance) between the position of the end of the robot and the force that it exerts on the environment [11]. Current implementations of impedance control, however, do not ensure that the desired impedance is implemented, so that effects such as joint friction are still liable to cause significant steady state error. A final approach is to design robots that do a better job of controlling force in an open loop fashion [18, 23], thereby avoiding control problems altogether.

9 Acknowledgements

This research was performed in the Newman Laboratory of Biomechanics and Human Rehabilitation at M.I.T. Partial support was supplied by the NSF grants number ECS-8307461 and EET-8613104, NIDRR grant number G000830074, and the Whitaker Foundation.

References

[1] C.H. An. *Trajectory and Force Control of a Direct Drive Arm.* PhD thesis, M.I.T. Department of Electrical Engineering and Computer Science, August 1986.

[2] O. Brune. Synthesis of a finite two-terminal network whose driving-point impedance is a prescribed function of frequency. *J. Math. Phy., MIT*, 10:131–236, August 1931.

[3] J.E. Colgate. *The Control of Dynamically Interacting Systems.* PhD thesis, M.I.T. Department of Mechanical Engineering, August 1988.

[4] J.E. Colgate and N. Hogan. Robust control of dynamically interacting systems. *International J. of Control*, 48(1):65–88, 1988.

[5] S.D. Eppinger and W.P. Seering. Modeling robot flexibility for endpoint force control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 165–170, 1988.

[6] S.D. Eppinger and W.P. Seering. On dynamic models of robot force control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 29–34, 1986.

[7] S.D. Eppinger and W.P. Seering. Understanding bandwidth limitations in robot force control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 904–909, 1987.

[8] W.B. Gevarter. Basic relations for control of flexible vehicles. *AIAA Journal*, 8(4):666–672, April 1970.

[9] G. Hirzinger. Direct digital robot control using a force-torque sensor. In *Proceedings of the IFAC Symposium on Real Time Digital Control Applications*, pages 243–255, 1983.

[10] N. Hogan. Impedance control: an approach to manipulation: part I—theory. *J. Dynamic Systems, Measurement, and Control*, 107:1–7, March 1985.

[11] N. Hogan. Stable execution of contact tasks using impedance control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1987.

[12] H. Kazerooni. Robust, non-linear impedance control for robot manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 741–750, 1987.

[13] R.P. Paul. Problems and research issues associated with the hybrid control of force and displacement. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1966–1971, 1987.

[14] R.K. Roberts, R.P. Paul, and B.M. Hillberry. The effect of wrist force sensor stiffness on the control of robot manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 269–274, 1985.

[15] A. Sharon, N. Hogan, and D. Hardt. More analysis and experimentation on a macro/micro manipulator system. In R. Shoureshi, K. Youcef-Youmi, and H. Kazerooni, editors, *Modeling and Control of Robotic Manipulators and Manufacturing Processes*, pages 417–422, ASME, 1987.

[16] A. Talbot. Some theorems on positive functions. *IEEE Trans. Circuit Theory (Correspondence)*, CT-12:607–608, December 1965.

[17] S.W. Tilley, R.H. Cannon, Jr., and R. Kraft. End point force control of a very flexible manipulator with a fast end effector. In F.W. Paul and K. Youcef-Youmi, editors, *Robotics: Theory and Applications*, ASME, New York, 1986.

[18] W.T. Townsend. *The Effect of Transmission Design on the Performance of Force Controlled Manipulators.* PhD thesis, M.I.T. Department of Mechanical Engineering, May 1988.

[19] M.E. Van Valkenburg. *Introduction to Modern Network Synthesis.* John Wiley and Sons, New York, 1960.

[20] D.E. Whitney. Force feedback control of manipulator fine motions. In *Proceedings of the Joint Automatic Control Conference*, 1976.

[21] J.J. Wlassich. *Nonlinear Force Feedback Impedance Control.* Master's thesis, M.I.T. Department of Mechanical Engineering, February 1986.

[22] C.H. Wu and R.P. Paul. Manipulator compliance based on joint torque control. In *Nineteenth IEEE Conference on Decision and Control*, pages 88–94, 1980.

[23] K. Youcef-Toumi and D. Li. Force control of direct drive manipulators for surface following. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1987.