

Design and Analysis of Haptic Interface and Teleoperator Feedback Systems

by

Paul G. Griffiths

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mechanical Engineering)
in The University of Michigan
2008

Doctoral Committee:

Associate Professor R. Brent Gillespie, Chair
Professor James S. Freudenberg
Professor Huei Peng
Professor A. Galip Ulsoy

© Paul G. Griffiths

All Rights Reserved

2008

Table of Contents

List of Figures	iv
Chapter 1 Introduction	1
1.1 Haptic Interface Feedback Systems	1
1.2 Teleoperator Feedback Systems	6
1.3 Outline	9
Chapter 2 Haptic Rendering under Position Feedback	12
2.1 Introduction	12
2.2 Posing Feedback Design as a Model-Matching Problem	14
2.3 Structure of the Haptic Interface controller	16
2.4 Feedback Design Requirements	17
Chapter 3 Algebraic Tradeoffs in Haptic Rendering	20
3.1 Introduction	20
3.2 Three-way Tradeoff between S , T , and Θ	20
3.3 Tradeoff between Performance and Sensitivity	23
3.3.1 Background	25
3.3.2 Application to haptic rendering	25
3.4 Existence of a Proper, Stabilizing Cancellation Controller	29
Chapter 4 Virtual Coupler Design for Haptic Rendering	33
4.1 Introduction	33
4.2 Performance of the Generalized Virtual Coupler	33
4.3 Optimizing for Performance	37
4.4 Experimental Results	38
4.4.1 Spring-Damper Virtual Coupler Design	42
4.4.2 Cancellation Coupler Design	47
4.5 Discussion	51
Chapter 5 Force-Reflecting Teleoperation under Position Feedback	53
5.1 Introduction	53
5.2 Position-Position Teleoperation	55
5.3 Models for Design and Performance Analysis	57

5.3.1	Parameterizing the Free-Free Dynamics	57
5.3.2	Measuring Transparency Performance	60
5.3.3	Measuring Tracking Performance	61
5.4	Analysis of Transparency and Tracking Performance	62
5.4.1	Transparency Performance	63
5.4.2	Tracking Performance	65
5.5	Steer-by-wire Example	66
5.5.1	Shaping the Master and Slave Dynamics & Designing a Virtual Coupler	68
5.5.2	Evaluating Transparency and Tracking	69
5.6	Conclusions	73
Chapter 6	Algebraic Tradeoffs in Teleoperation	74
6.1	Introduction	74
6.2	Tradeoff between Transparency and Sensitivity	74
6.3	Example: Steer-by-wire	77
Chapter 7	Waterbed Tradeoff in Distortion	80
7.1	Introduction	80
7.2	Bode Attenuation Integral in Haptic Rendering	80
Chapter 8	A Fundamental Conflict between Performance and Passivity in Haptic Rendering	86
8.1	Introduction	86
8.2	Bode Gain-Phase Relationship	87
8.3	Conflict between Performance and Passivity	88
8.4	Example	92
8.5	Conclusions	95
Chapter 9	Conclusions	96
9.1	Summary of Results	96
9.2	Near-term Work	98
9.3	Future Work	99
Bibliography	101

List of Figures

Figure		
2.1	Schematic of a rotary haptic device and controller	12
2.2	General control configuration	14
2.3	Model-matching block diagram for haptic rendering	15
2.4	Architecture of the haptic interface controller	16
2.5	Sample Nyquist plot of R (dashed) and sR (solid) for positive frequencies	19
3.1	Algebraic relationships between S , T , and Θ	21
3.2	Design specifications M_Θ , M_S , and M_T	24
3.3	Upper and lower magnitude bounds on the set of feasible virtual environment dynamics	24
3.4	Tradeoff severity Γ for a simple mass device P and a virtual spring environment R_d	29
3.5	Tradeoff severity Γ for a simple mass device P and a virtual mass-spring-damper environment R_d	30
3.6	Open-loop distortion Θ_o for a simple mass device P and a virtual mass-spring-damper environment R_d	31
4.1	Controller partitioned into a generalized virtual coupler V and virtual environment R_d	34
4.2	Frequency responses of R_{d1} and R_{d2} relative to the haptic device P (top) and resulting tradeoff severities Γ (bottom)	40
4.3	Single-axis, rotary haptic device with wooden handwheel (above) and without (below)	41
4.4	Open-loop and closed-loop distortion for the spring-damper virtual coupler rendering R_{d1} (top) and R_{d2} (bottom)	43
4.5	Bode sensitivity function S for the spring-damper virtual coupler rendering R_{d1} (top) and R_{d2} (bottom)	45
4.6	Step responses of virtual environments R_{d1} (top) and R_{d2} (bottom) rendered by the spring-damper virtual coupler	46
4.7	Cancellation coupler design parameters H_{11} , H_{12} and H_{21}	47
4.8	Open-loop and closed-loop distortion for the cancellation coupler rendering R_{d1} (top) and R_{d2} (bottom)	48
4.9	Bode sensitivity function S for the cancellation coupler rendering R_{d1} (top) and R_{d2} (bottom)	49

4.10	Step responses of virtual environments R_{d1} (top) and R_{d2} (bottom) rendered by the cancellation coupler	50
5.1	Block diagram of position-position teleoperation	56
5.2	Free-free dynamics in terms of the controller elements and open-loop master and slave (top) and a reparameterization in terms of design parameters T_m , T_s , A , and B (bottom)	59
5.3	Feedback configuration for transparency	61
5.4	Feedback configuration for tracking	62
5.5	Configuration of a steer-by-wire system	67
5.6	Steer-by-wire master and slave dynamics P_m and P_s	69
5.7	Transparency diagram for design (a) with overlaid contours indicating worst-case $ \Theta_h(j\omega) $ for a given magnitude of P_e	71
5.8	Transparency diagram comparing design (a) and (b)	71
5.9	Steer-by-wire transparency performance (top) and tracking performance (bottom) for design (a) and (b)	72
6.1	Lower bounds on elements of S^* from Prop. 6.2.1	78
6.2	Output sensitivity function S for two steer-by-wire designs	78
7.1	Sample bounds on $\log \Theta(j\omega) $ due to performance and bandwidth requirements	84
7.2	Lower bound on the peak of $\log \Theta(j\omega) $ between ω_0 and ω_1 predicted by (7.14)	85
8.1	Gain-phase weighting factors from (8.3) shown on a logarithmic frequency scale	89
8.2	Right and left sides of the inequality in Proposition 8.3.1 for $M_\Theta = 0.5$. . .	92
8.3	Distortion Θ for designs A, B, C and D	93
8.4	Frequency response of the haptic device P , the virtual environment R_d , and four causal approximations to R_d	94

Chapter 1

Introduction

1.1 Haptic Interface Feedback Systems

The growth of virtual reality systems for medical training [Li and Liu, 2007; Sorid and Moore, 2000; Vlachos et al., 2003], human motor control research [Huang et al., 2006; Patton et al., 2006], and entertainment applications [Laycock and Day, 2003] has driven development of haptic interface systems which are able to synthesize a variety of tactile experiences. A haptic device is typically a hand-held interface, such as a handwheel or pen-shaped stylus. By manipulating the device, the user navigates through a computer-rendered virtual environment, and motors in the device push against the user to synthesize a mechanical interaction with the virtual environment. Realistic haptic portrayal of the virtual environment places significant demands on both the design of the haptic device and the control system, and existing design practice struggles to achieve acceptable feedback properties including performance and stability robustness. Given the challenges faced by designers of haptic interface systems, we undertake a theoretical analysis to relate hardware choices to achievable design goals and to provide controller synthesis techniques that approach the feasible boundaries of the feedback design.

The mechanical design of the haptic device and the quality of sensors and actuators define intrinsic limitations in the feedback design for haptic rendering. This dissertation addresses the most common mechanical design for haptic interface systems, where motor commands and the user force on the haptic device affect the device motion through the same dynamics [Iwata, 1990; Massie and Salisbury, 1994; Buttolo and Hannaford, 1995; Lee et al., 2000]. Devices for which this idealization is appropriate are termed **impedance-type** [Adams et al., 1998; Shen and Goldfarb, 2006; Faulring et al., 2006]. The performance of these devices is limited by mechanical properties including nonlinear friction in bear-

ings and structural compliance between the user and motor. Sensor and actuator properties present further limitations. Impedance-type devices typically use high-quality, low-cogging DC servo-motors for actuation and co-located, high-resolution encoders for position sensing. Intrinsic performance limitations exist within this sensing and actuation scheme as measurement of the haptic device position is subject to noise or quantization effects, and actuation is subject to saturation and bandwidth constraints.

The challenges associated with high-fidelity haptic rendering depend on the nature of the virtual environment. For example, a dental training simulation must synthesize hard contact interactions between a steel tool and tooth enamel [Wang et al., 2003]. High resolution sensors are needed to locate the point of collision, and high-bandwidth actuators are needed to generate a rapid rise in force. On the other hand, important haptic features when simulating a tool interacting with soft tissue include compliance, damping, and inertia [Vuskovic et al., 2000]. Since the objective of haptic rendering is to portray diverse dynamics to the user, an important characteristic of a particular haptic device is the class of virtual environments which it is capable of rendering.

Extensive prior research addresses hardware limitations in the haptic rendering of hard contacts, also called the **virtual wall problem**. Realistic rigid contact requires fast and large force production to arrest device motion when the user reaches the point of contact. While large feedback gains are needed to prevent penetration of the virtual wall, sufficiently large gains give rise to chatter and oscillations as the user pushes into the wall. Previous work has analyzed the limitations imposed by sampling [Colgate et al., 1995; Colgate and Schenkel, 1997; Minsky et al., 1990; Mahvash and Hayward, 2005; Mehling et al., 2005; Diolaiti et al., 2006; Shen and Goldfarb, 2006; Çavuşoğlu et al., 2002; Abbott and Okamura, 2005], quantization [Colgate et al., 1995; Mehling et al., 2005; Diolaiti et al., 2006; Çavuşoğlu et al., 2002; Abbott and Okamura, 2005], hardware damping [Colgate et al., 1995; Colgate and Schenkel, 1997; Minsky et al., 1990; Mehling et al., 2005; Diolaiti et al., 2006; Shen and Goldfarb, 2006; Çavuşoğlu et al., 2002; Abbott and Okamura, 2005], and nonlinearities [Colgate et al., 1995; Colgate and Schenkel, 1997; Mahvash and Hayward, 2005; Mehling et al., 2005; Diolaiti et al., 2006; Shen and Goldfarb, 2006; Çavuşoğlu et al., 2002; Abbott and Okamura, 2005]. This work relates the quality of hardware in terms of mechanical design, sensor resolution, and computing power to its capability in rendering hard contact.

Analysis of the virtual wall problem addresses one, but not all challenges in haptic rendering. Design challenges, limitations, and tradeoffs associated with high-fidelity haptic rendering of linear dynamic systems have received little attention to date. Even without the nonlinearity intrinsic to the virtual wall, accurate rendering of certain linear time-invariant

environments can suffer from undesired oscillations. Prior work has analyzed closed-loop oscillations and chatter of the haptic device through sampled-data models [Colgate and Schenkel, 1997; Miller et al., 2000; Fardad and Bamieh, 2004] and time-domain energy analysis [Gillespie and Cutkosky, 1996; Mahvash and Hayward, 2005; Abbott and Okamura, 2005; Diolaiti et al., 2006]. The analyses to-date address challenges that arise due to sampled-data effects. However, in our experience with haptic interface feedback design, we observe instabilities and poor robustness properties that are not mitigated through increased sample rate.

This dissertation presents the first comprehensive characterization of intrinsic capabilities, limitations, and tradeoffs in haptic rendering of continuous linear time-invariant (LTI) systems. These systems form an important class of virtual environments, since a subset of LTI systems describe idealized models of mechanical elements including masses, springs, dampers, and their interconnections. An analysis of feedback design challenges in rendering LTI systems is a logical starting point given the expressive power of this class of systems and the powerful linear feedback analysis techniques which may be applied to the problem.

Prior work has captured performance in haptic interface systems by the achievable hardness of virtual walls. Larger controller gains improve hardness while excessively large gains induce undesired oscillations when the user interacts with the virtual wall. Measures such as *Z*-width [Colgate et al., 1995] identify the set of proportional and derivative controller gains that can be used without inducing oscillations. However, measuring performance by achievable controller gains has limited utility when the goal is to render LTI environments. High gains are not necessarily appropriate, and furthermore, the hardware dynamics may play an important role in the dynamics rendered to the user, in which case controller gains are not indicative of closed-loop performance.

An important contribution of this dissertation is the introduction of a new performance measure, termed **distortion** which measures the difference between the virtual environment dynamics and the actual closed-loop dynamics provided to the user. This quantity offers two primary benefits over prior assessments of performance in haptic rendering. Unlike the concept of *Z*-width, distortion is a property of the closed-loop system and accounts for both the hardware and controller design. Another important property is that distortion measures error. The quantity termed **transparency ratio** [Lawrence, 1993] has been borrowed from teleoperator literature and applied to haptic interface [McJunkin et al., 2005]; however accurate rendering is achieved as this ratio approaches unity not zero. Measuring error rather than a ratio permits us to apply techniques of magnitude analysis.

Both performance and stability are important goals of the feedback design. Analysis of

closed-loop stability in haptic rendering may be significantly affected by the user's interaction with the haptic device. In addition to standard measures of stability robustness such as phase margin, gain margin, and stability radius, the haptic rendering feedback loop must remain stable even as the user closes an additional feedback path around the device. Modeling the user is not generally practical because musculoskeletal dynamics are nonlinear, vary between users, and depend on grip, posture, and muscle co-contraction (i.e. tensing ones muscles). Coupled stability of the human operator with the haptic interface system is typically solved by guaranteeing that the closed-loop response of the haptic device is energetically passive [Colgate and Hogan, 1989]. (By energetically passive, we mean that the response between power variables such as force and velocity is passive.) One motivation for the passivity criterion is a result from network theory that predicts that the feedback interconnection of two passive systems is closed-loop stable [Khalil, 2002]. Coupled stability is assured if the closed-loop response of the haptic device is passive and the user is assumed to also be passive. In point of fact, the user violates passivity due to volitional feedback; however, in practice it is sufficient to assure that the passive bio-mechanics are stable when coupled to the haptic device.

The process of feedback design necessarily involves compromises between conflicting objectives. The structure of the feedback loop, that is the location of sensors and actuators, dictates a set of tradeoffs between feedback properties such as disturbance attenuation and noise attenuation [Skogestad and Postlethwaite, 1997; Seron et al., 1997; Freudenberg et al., 2003]. Another set of tradeoffs emerge from conditions imposed on transfer functions by physical realizability [Bode, 1945]. In addition to tradeoffs within the feedback design, properties of the available hardware such as sample-rate, delay, actuator bandwidth, and sensor noise impose limitations that cannot be overcome through feedback design. By quantifying tradeoffs and limitations, we can reveal intrinsic relationships between hardware and feedback properties which all controller designs must satisfy.

Design limitations and tradeoffs are well characterized for typical servo-control applications [Skogestad and Postlethwaite, 1997]; however intrinsic differences between servo-control and haptic rendering require us to re-interpret design relationships for haptic rendering. An important objective of feedback in servo-control applications is to reject disturbances that enter at the actuator input. In haptic rendering, the user's applied force also enters at the actuator input; however, here, the control objective is to generate the dynamic response of the virtual environment, not reject the user's input. This design objective may be treated as a model-matching problem rather than a servo-control problem. As a consequence, standard loop-shaping techniques that are appropriate in servo-control systems are not directly applicable to haptic rendering, and additional design tradeoffs exist in haptic

rendering which have no counterpart in servo-control systems.

In our analysis of intrinsic tradeoffs and limitations within haptic rendering we give special attention to an important characteristic of feedback systems which has not been analyzed for haptic rendering: the sensitivity of the closed-loop response to parametric variations in the hardware. Predictions about important system properties including stability or performance become less reliable as the sensitivity to hardware dynamics increases. Although feedback may be used to attenuate sensitivity, feedback may also amplify sensitivity. For certain virtual environments, the hardware dynamics will mask or distort the virtual environment that is to be rendered. Feedback compensation may be used to recover performance as suggested in [Colgate et al., 1995]; however the costs in terms of sensitivity to hardware dynamics have not been quantified. Practical feedback designs must meet bounds on sensitivity as excessive sensitivity jeopardizes performance and stability of the feedback system.

The passivity requirement for coupled stability imposes restrictions on the feedback design that limit performance. Prior work on rendering hard contact identifies ranges of controller parameters which, given sampling and quantization effects, still are passive [Colgate et al., 1995; Diolaiti et al., 2006]. We wish to answer a related, but broader question: what passive LTI virtual environments cannot be rendered passively within given frequency domain specifications on distortion? We note that, without bandwidth constraints, there is no apparent conflict between passivity and rendering a passive virtual environment. However, the bandwidth of closed-loop control is always limited in practice by physics of the sensors and actuators, existence of high-frequency hardware resonances, and computational speed. Subject to bandwidth constraints, the feedback design must satisfy certain analytic conditions that may bring performance requirements in certain frequency ranges into conflict with passivity requirements in other frequency ranges.

Relationships across frequencies follow from the analyticity of transfer functions. Many servo-control applications are subject to an analytic relationship, termed the **waterbed effect**, which predicts that disturbance attenuation at some frequencies is balanced by disturbance amplification at other frequencies. This relationship does not describe a performance tradeoff in haptic interface systems, however, the transfer function that describes disturbance attenuation in servo-control systems is not the same as the distortion transfer function in haptic rendering. We show that, under appropriate hypotheses, the distortion transfer function satisfies a related waterbed tradeoff. Prior work in the field of haptic interface systems has not explored the role of analytic relationships in design tradeoffs.

Chapters 2–4, 8, and 9 of this dissertation analyze feedback design limitations and tradeoffs in haptic rendering. Working with LTI models of the hardware, controller, and

virtual environment, we reveal previously unrecognized fundamental design relationships. Results in Chapter 3 follow from algebraic analysis of the feedback goals in haptic rendering. Insight gained from algebraic analysis is used to develop an improved controller design technique for haptic rendering. Results in Chapters 8 and 9 follow from analytic analysis of closed-loop transfer functions and the restrictions imposed by bandwidth limitations. Section 1.3 provides a guide to the contents of the dissertation. In brief, the novel contributions of this dissertation to the field of haptic interface systems are

- Development of **distortion**, a new performance metric for haptic rendering
- Identification and analysis of a fundamental tradeoff between performance and sensitivity in haptic rendering
- Experimental support of the fundamental tradeoff between performance and sensitivity
- Generalization of the standard **virtual coupler**, a performance-optimized technique for tuning.
- Identification and analysis of a waterbed tradeoff in haptic performance
- Identification and analysis of a bandwidth-induced conflict between performance and passivity

Chapters 5 and 6 analyze design of teleoperator systems under position feedback. Significant parallels exist between haptic rendering and teleoperation, and the work in Chapter 6 is an extension of the algebraic analysis in Chapter 3. The analytic results in Chapters 7 and 8 have not yet been extended to teleoperation.

1.2 Teleoperator Feedback Systems

Whereas a haptic interface system connects the user to a virtual environment, a teleoperator system connects the user to a physical, remote environment. Applications of teleoperation include remote handling of nuclear material [Clement et al., 1985], robot-assisted surgery [Okamura, 2004], and fly-by-wire [Krahe, 1996] and steer-by-wire [Bretz, 2001; Odenthal et al., 2002; Pan et al., 2006]. If the teleoperator is force-reflecting, the haptic feedback provided to the user is representative of the actual forces between the remote robot and environment. Not all teleoperators provide force-feedback; indeed, commonly available fly-by-wire flight controls and surgical robots are not force-reflecting teleoperators [Krahe,

1996; Okamura, 2004]. Demonstrating reliability and fault-tolerance of these real-time, safety-critical applications is daunting even without haptic rendering. Additional feedback pathways due to haptic rendering at the user (or master) interface generate additional sensitivity and stability issues. Better understanding of the intrinsic feedback design limitations and tradeoffs in force-reflecting teleoperation will foster the adoption of this technology.

Certain feedback control challenges in teleoperation mirror those in haptic rendering. One objective of the force-reflecting teleoperation is to approximate the haptic experience of direct mechanical interaction with the environment. As in haptic rendering, the inherent dynamics of the user interface may mask the environment dynamics and require feedback compensation. The feedback design must also assure coupled stability between the user and master as in haptic rendering. The principle additional feedback goal in teleoperation is tracking between the master and slave devices.

Literature on force-reflecting teleoperation recognizes the existence of a tradeoff between performance and stability robustness [Lawrence, 1993; Leung et al., 1995; Kim et al., 1992]. Closed-loop stability in the face of significant delay between the master and slave requires special compensation which necessarily degrades the accurate rendering of the remote environment dynamics [Niemeyer and Slotine, 2004; Stramigioli et al., 2002]. Robust control design has been applied to reduce the cost in terms of performance while guaranteeing stability in the face of a bounded delay [Leung et al., 1995]. The severity of this tradeoff increases as delay increases and disappears for delay-free teleoperation [Lawrence, 1993; Leung et al., 1995; Kim et al., 1992]. The existence of tradeoff between performance and robustness properties in the absence of time-delay has not been previously analyzed.

As in haptic rendering, analysis of feedback design tradeoffs in teleoperation depends on appropriate metrics for performance. One approach is to view teleoperator performance as a dual tracking problem. The motion of the master should track the slave, and the interaction force between the user and the master should track the interaction force between the slave and the environment [Yokokohji and Yoshikawa, 1994; Yan and Salcudean, 1996]. The drawback of this definition of performance is that neither position tracking error nor force tracking error alone describe how well the user can feel the remote environment dynamics. An alternative approach is to compare the dynamic response of the environment to the dynamic response presented to the user. Transparency ratio is defined by the ratio of these two quantities and should ideally have a frequency-response near unity [Lawrence, 1993; Fite et al., 2001]. This definition of performance also has drawbacks. When evaluating deviation of performance from the ideal, both magnitude and phase of the transparency ratio must be monitored. For example, if the closed-loop response rendered to the user is a delayed version of the environment dynamics, the transparency ratio is still precisely

unity in magnitude. We circumvent these challenges by extending the concept of distortion, introduced for haptic rendering, to measure teleoperator transparency. Since distortion describes error, it admits basic magnitude analysis techniques such as application of the triangle inequality.

A comprehensive assessment of teleoperator performance is challenging because transparency depends not only on the control design, but also on inherently uncertain environment dynamics. Performance has been indirectly assessed by examining the rendered response for representative environments such as rigid and free-space dynamics [Hashtrudi-Zaad and Salcudean, 2001]. However, a full assessment of teleoperator performance should capture the class of environments which are rendered to the user within a certain measure of transparency. Prior literature has not provided such a comprehensive characterization of teleoperator performance.

While the robust rendering of environment dynamics is a recognized challenge in teleoperation, prior work has not analyzed the intrinsic tradeoff between environment dynamics and sensitivity to hardware variations. Previous work notes that all solutions to the mixed-sensitivity optimal control problem abide by the identity $S + T = 1$ [Yan and Salcudean, 1996]; however the conventional design interpretation of this identity does not apply to teleoperation. While the performance goals of force-reflecting teleoperation can be expressed as a dual tracking problem, the error dynamics are not described by the sensitivity function S as is typically the case in servo-control systems.

In Chapters 5 and 6, we analyze feedback design of teleoperators under position feedback. We extend distortion and are the first to analyze transparency (as measured by distortion) and tracking in a unified framework. A novel parameterization of the feedback design is developed which provides a useful decomposition of the controller actions. With this parameterization we are able to provide a quantitative characterization of transparency over classes of environment dynamics. We develop a graphical design tool called a **transparency diagram** which relates frequency response of design parameters with distortion. Analysis of performance and transparency reveals an previously unstudied algebraic tradeoff which has certain parallels with the performance/sensitivity tradeoff in haptic rendering. As an overview, the novel contributions of this dissertation in the area of teleoperator design include

- Introduction of **distortion** as an appropriate performance metric for analysis in teleoperation
- Development of a novel parameterization of all position-position LTI teleoperator feedback designs useful for analysis and controller tuning

- Quantitative analysis of teleoperator distortion for classes of environment dynamics
- Development of the **transparency diagram**, a graphical tool for assessing performance and tuning of the teleoperator controller.
- Identification and analysis of a fundamental tradeoff between performance and sensitivity

We now lay out the contents of dissertation, related to both haptic interface systems and teleoperators.

1.3 Outline

This dissertation begins with an analysis of algebraic (frequency-by-frequency) design tradeoffs in haptic rendering. Chapter 2 introduces the general setup of haptic rendering with an impedance-type device under position feedback. Design requirements are then translated into a set of specifications on closed-loop frequency responses. Chapter 3 analyzes conflicts between design goals induced by algebraic relationships between closed-loop transfer functions. We find that haptic rendering is subject to a three-way algebraic tradeoff which has not been not previously identified within haptic interface systems and has no counter-part in servo-feedback control. We isolate one aspect of this tradeoff for study: an inherent tradeoff between performance and sensitivity. Interpreting results from [Freudenberg et al., 2003], we capture the severity of the performance/sensitivity tradeoff in a single frequency-dependent parameter that depends on the virtual environment and haptic device dynamics.

Practical implications of design tradeoffs are best illustrated through a design study. Chapter 4 introduces structural requirements imposed on the control system realization and techniques for controller design. The feedback is partitioned into the virtual environment and a multi-input/multi-output controller called the **generalized virtual coupler** [Adams and Hannaford, 1999]. Analysis of performance leads to a new design technique for the virtual coupler which achieves better performance than standard virtual coupler design. We call this performance-optimized design the **cancellation coupler**. The intrinsic tradeoff between performance and sensitivity is highlighted through an experimental study with both the standard virtual coupler and the performance-optimized design. The analysis, design, and experimental work discussed in Chapters 2 and 3 are also available in Griffiths et al. [2007] and Griffiths et al. [2008b].

Chapter 5 develops a novel parameterization of teleoperator feedback design which is

useful both for controller tuning and for characterizing performance over a class of environment dynamics. The parameterization describes all teleoperator feedback designs using four transfer functions that decompose the controller action into hardware compensation and virtual coupling between the master and slave. We show that these parameters describe a class of environments that are rendered within a certain bound on distortion. Algebraic relationships provide a two-way mapping between controller elements and parameters. Thus we may design the controller by shaping the frequency response of the parameters and then determining the controller realization. We may also use the parameterization to characterize performance of advanced feedback schemes such as wave-variable controllers for which no prior performance bounds exist.

The connection between the parameterization of teleoperator feedback design and performance leads to our development of a graphical tool for design and analysis which we call a **transparency diagram**. An upper bound on the frequency response of distortion may be computed using only the magnitude frequency response of the environment dynamics, discarding the phase information. A contour plot of distortion may then be constructed with frequency along the x -axis and magnitude of the environment dynamics along the y -axis. We show that asymptotes in this plot are functions of the parameterization of the feedback design. By plotting the magnitude frequency response of anticipated environment dynamics against the asymptotes reveals the critical parameters of the feedback design. The design parameterization, performance analysis, and development of the transparency diagram first appeared in Griffiths et al. [2008a].

Chapter 6 analyzes a previously unexplored tradeoff between transparency and sensitivity in force-reflecting teleoperation. Algebraic analysis reveals a lower bound on the sensitivity of the feedback design to achieve a desired level of distortion. As in haptic rendering, feedback compensation of the user's interface or master may be necessary to feel the environment without significant distortion. In contrast to haptic rendering, the lower bound on feedback sensitivity does not approach a finite value as distortion approaches zero.

The final part of this dissertation studies analytic relationships within haptic rendering. Chapter 7 studies integral relationships on the frequency response of the performance measure, distortion. Under appropriate assumptions, distortion satisfies an analogue of the Bode Sensitivity Integral. When subject to bandwidth constraints, the integral relationship induces a waterbed tradeoff between distortion at different frequencies. It follows that an aggressive performance specification in conjunction with limited feedback bandwidth necessarily implies a feedback design with very poor performance near the closed-loop bandwidth.

Chapter 8 reveals a fundamental conflict between performance requirements and coupled stability. The closed-loop response presented to the user should be passive to guarantee coupled stability, which imposes a phase constraint on the frequency response of closed-loop transfer functions. Simultaneously, performance in haptic rendering may be captured by a magnitude constraint on the frequency response of distortion. Bode gain-phase relationships allow us to demonstrate that certain virtual environment dynamics, which are themselves passive, cannot be rendered passively within finite-bandwidth performance bounds. This work will appear in Griffiths and Gillespie [2008].

Chapter 2

Haptic Rendering under Position Feedback

2.1 Introduction

In a standard configuration of haptic rendering, a human operator grasps and applies forces to a motorized, computer-controlled manipulator. Figure 2.1 depicts a direct-drive, single-axis, impedance-type haptic device. For this type of device we make the idealized assumption that the user's torque f applied to the handwheel and the motor torque u affect the handwheel position y through the same dynamics. In general, the haptic device may be linear or rotary. Without loss of generality, we refer to f and u as forces rather than torques.

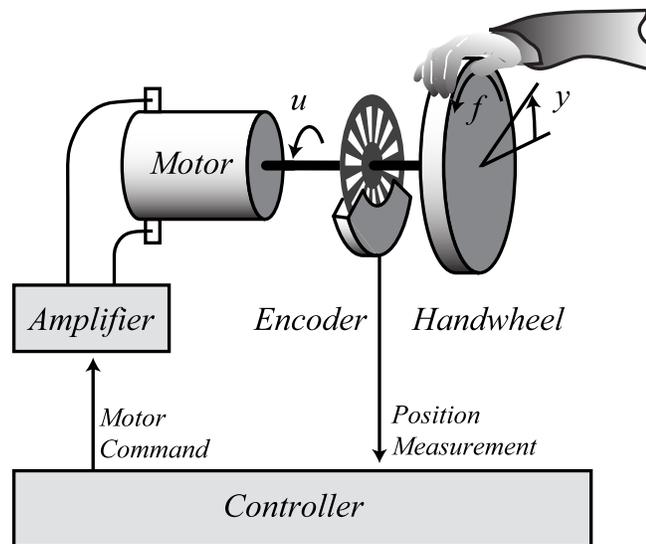


Figure 2.1 Schematic of a rotary haptic device and controller.

Impedance-type haptic devices are widely used for desktop-scale applications and a number of such devices are commercially available today. On these devices, position sensing alone is typical. A force sensor may be installed between the motor and handwheel in Fig. 2.1 to monitor u and f . Force sensing becomes a necessity in very large applications such as industrial devices where, due to large friction and device dynamics, the user can no longer move the unpowered interface. (Such devices are termed admittance-type devices because the user can generate a force, not a position input.) In smaller haptic interface applications, however, the use of force sensing is atypical given the price of high-quality force sensors and issues of noise, calibration and physical robustness. We focus on the feedback design under position sensing as this is the commonly used hardware configuration.

We model the dynamics of the haptic device and the controller as transfer functions of the Laplace variable s . Let P denote the haptic device dynamics such that

$$y = P(f + u). \quad (2.1)$$

The haptic device position y is then measured by the encoder, and a quantized version of y is available to the controller. For simplicity, we model sensor quantization with an additive noise n . Letting C describe the controller, we have that

$$u = -C(y + n). \quad (2.2)$$

The primary goal of the controller is to shape the dynamic behavior of the haptic device from f and y to match desired dynamics, termed the **virtual environment**. We denote a linear time-invariant virtual environment by the transfer function R_d . The desired output y_d in response to the user's input f is

$$y_d = R_d f. \quad (2.3)$$

We denote the actual closed-loop dynamics between f and y rendered to the user by R . These dynamics are defined by the feedback interconnection of P with C :

$$R \triangleq \frac{P}{1 + PC}. \quad (2.4)$$

The accuracy of haptic rendering may be measured by **distortion**, defined as the error between the actual and desired closed-loop dynamics [Griffiths et al., 2008b]

$$\Theta \triangleq \frac{R - R_d}{R_d}. \quad (2.5)$$

The error is normalized by R_d to provide a notion of relative error. The dynamic response presented to the user closely matches the desired response if Θ is small along the $j\omega$ -axis.

2.2 Posing Feedback Design as a Model-Matching Problem

The feedback design problem of haptic rendering may be expressed in the standard form of the **general control configuration** [Skogestad and Postlethwaite, 1997]. As shown in Fig. 2.2, the standard form consists of a generalized multivariable plant G in feedback with a generalized controller K . The generalized plant G describes the input/output responses from disturbance inputs w and control inputs u to performance outputs z and measured outputs y :

$$\begin{bmatrix} z \\ y \end{bmatrix} = \begin{bmatrix} G_{zw} & G_{zu} \\ G_{yw} & G_{yu} \end{bmatrix} \begin{bmatrix} w \\ u \end{bmatrix}. \quad (2.6)$$

We denote the closed-loop disturbance response from w to z by T_{zw} . For a scalar system, the disturbance response is given by

$$T_{zw} = G_{zw} + G_{zu}K(1 - G_{yu}K)^{-1}G_{yw}. \quad (2.7)$$

Performance goals are achieved by designing the generalized controller K to attenuate the disturbance response T_{zw} .

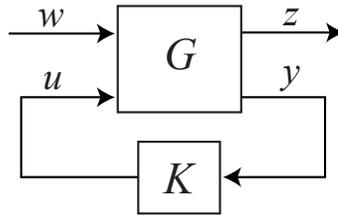


Figure 2.2 General control configuration.

In terms of the general control configuration, the human operator force f provides the exogenous input w . Let the generalized controller K be C , and define the performance output as

$$z \triangleq R_d^{-1}(y - y_d) \quad (2.8)$$

Figure. 2.3 shows the resulting general control configuration. The elements of the generalized plant G , contained within the dashed box, are given by

$$G = \begin{bmatrix} P/R_d - 1 & P/R_d \\ P & P \end{bmatrix}. \quad (2.9)$$

The disturbance response T_{zw} is found by substituting (2.9) into (2.7)

$$T_{zw} = \frac{1}{R_d} \left(\frac{P}{1+PC} - R_d \right). \quad (2.10)$$

Recall from (2.4) that $R \triangleq P/(1+PC)$. It is then evident that, by construction

$$T_{zw} = \Theta. \quad (2.11)$$

It follows that the disturbance attenuation problem (that is attenuating the response from f to z) captures the model-matching goal of making R match R_d .

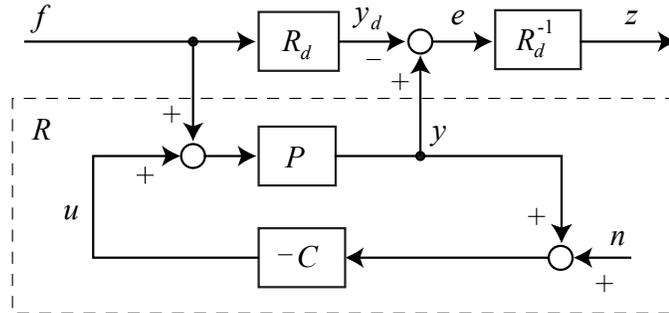


Figure 2.3 Model-matching block diagram for haptic rendering. The human operator feels the rendered virtual environment R , the transfer function y/f ; whereas the desired response is given by the virtual environment R_d . The performance variable z measures distortion, the normalized error between the actual and desired responses.

While both the performance goals of servo-design and haptic rendering can be described as disturbance attenuation problems, note that a typical servo-control application is not also a model-matching problem. For haptic rendering, the desired response of y to the exogenous input f is R_d ; whereas in servo-control problems, the disturbance rejection problem implies ideally a null response from f to y , or $R_d = 0$. The model-matching block diagram then reduces to a simple feedback loop. The model-matching problem presented by haptic rendering, on the other hand, yields an inherently multivariable feedback system.

2.3 Structure of the Haptic Interface controller

The haptic interface controller is typically partitioned into two parts: a simulation of the virtual environment and a virtual coupler [Adams and Hannaford, 1999]. This latter element connects the hardware with the virtual environment. As shown in Fig. 2.4, the virtual coupler produces an input f_e to the virtual environment and receives the desired position of the haptic device y'_d . (Note that $y'_d \neq y_d$ as f_e is not generally equal to f .) Partitioned in this way, design of the controller is split into two problems: creating an accurate simulation of the virtual environment dynamics, and designing a generalized virtual coupler to render that virtual environment accurately. With this structure, various virtual environments can, in theory, be interchanged without redesigning the virtual coupler. Note that for haptic applications involving scaling between the haptic device and the virtual environment, we assume without loss of generality in our analysis that these scaling factors are internal to the virtual environment.

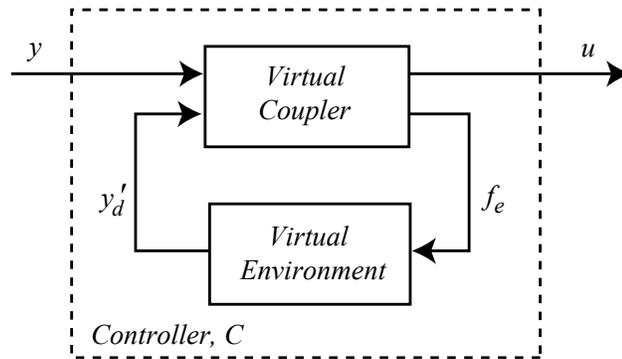


Figure 2.4 Architecture of the haptic interface controller.

The virtual coupler is often modeled after mechanical elements such as springs and dampers, but can more generally be described by four transfer functions relating the two inputs and two outputs. In Chapter 4 we introduce a generalized virtual coupler design that out-performs the standard virtual coupler design by exploiting the full flexibility of this multi-input/multi-output design. However, let us defer further discussion of the virtual coupler. Our primary focus is to analyze feedback design tradeoffs in haptic rendering that govern all designs regardless of the structure imposed on C .

2.4 Feedback Design Requirements

Performance requirements demand that the rendered virtual environment R match the desired dynamics R_d over some finite bandwidth. We may express this requirement as a bound on the magnitude of distortion evaluated along the imaginary axis

$$|\Theta(j\omega)| \leq M_\Theta(\omega). \quad (2.12)$$

The feedback design must also achieve certain stability margins and provide adequate robustness to parameter variations in the hardware. The Bode sensitivity function

$$S \triangleq \frac{1}{1+PC}. \quad (2.13)$$

describes several important feedback properties including stability robustness and the sensitivity of closed-loop transfer functions to variations in hardware dynamics [Skogestad and Postlethwaite, 1997]. Consider, for example, the sensitivity of the rendered virtual environment R to the haptic device dynamics P . Differentiating R by P reveals that

$$\frac{P}{R} \frac{dR}{dP} = S. \quad (2.14)$$

Thus S describes the differential change in the rendered environment dR/R to a differential change in the haptic device model dP/P . As a first-order approximation, relative error in the haptic device dynamics results in a relative error between the nominal and actual rendered virtual environment scaled by the Bode sensitivity function S . A design specification of the form

$$|S(j\omega)| \leq M_S(\omega) \quad (2.15)$$

stipulates a degree of robustness to hardware dynamics.

An inherent cost of feedback is the injection of sensor noise into the plant output. Most commonly, haptic interfaces use encoders which introduce finite quantization and may produce a perception of roughness or chatter. The response from n to y is given by the complementary sensitivity function

$$T \triangleq \frac{PC}{1+PC}. \quad (2.16)$$

For high-fidelity haptic rendering, the magnitude frequency response of T must not become too large as the user will perceive effects of sensor noise and quantization.

In addition to noise transmission, the complementary sensitivity function describes closed-loop bandwidth. Due to practical considerations such as unmodeled hardware dynamics, sensor and actuator bandwidths, T must satisfy a bandwidth constraint of the form

$$|T(j\omega)| \leq M_T(\omega), \quad \forall \omega \geq \omega_c, \quad (2.17)$$

where $M_T(\omega) \rightarrow 0$ as $\omega \rightarrow \infty$. Feedback designs that exceed intrinsic bandwidth limits inject high frequency sensor noise in the plant output and tend to excite lightly damped, structural resonances.

While the controller C should be designed to stabilize the haptic device P , stability of this feedback interaction is not sufficient to guarantee well-behaved interaction between the human operator and the controlled device. When in physical contact with the device, the human operator forms an additional feedback path between the haptic device position y and the force f . The dynamics in this path are variable and depend on many factors such as grasp, posture, muscle co-contraction, and volitional responses. As a result, the coupled user and powered haptic device may give rise to undesired oscillations.

A practical method to assure **coupled stability** of the human operator and haptic interface system is to design the dynamic response of the haptic interface system to be passive [Colgate and Hogan, 1989]. A necessary and sufficient condition for a linear time-invariant transfer function between a pair of power variables—such force and velocity—to be passive is that its poles lie in the open left-half plane and its Nyquist plot lies in the closed-right half plane [Slotine and Li, 1991]. Note that the rendered dynamic response R is a transfer function from force to position rather than from force to velocity. Thus we are interested in whether sR is passive rather than R . Assume that the controller C is stabilizing. The user is presented with a passive dynamic response if

$$\operatorname{Re}[j\omega R(j\omega)] \geq 0 \quad \forall \omega. \quad (2.18)$$

In other words, the Nyquist plot of sR must lie entirely in the right-half plane, or equivalently, the positive- ω locus of R must remain below the real-axis.

Passivity imposes a phase requirement that may be equivalently expressed as a magnitude requirement. When the scattering operator [Haykin, 1970]

$$Q(s) \triangleq \frac{sR(s) - 1}{sR(s) + 1}. \quad (2.19)$$

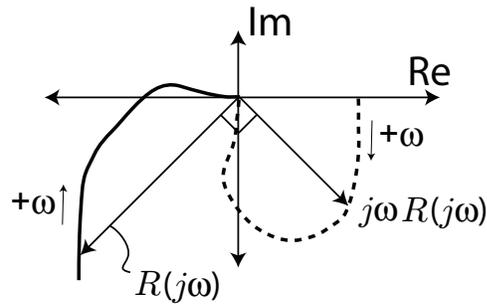


Figure 2.5 Sample Nyquist plot of R (dashed) and sR (solid) for positive frequencies. The response of R and sR at each frequency are separated by 90° degrees. The closed-loop system is passive if the locus of sR lies in the 1st and 4th quadrant, or equivalently, if the locus of R for positive frequencies lies in the 3rd and 4th quadrant. The sample trace does not describe a passive system.

satisfies

$$|Q(j\omega)| \leq 1, \quad \forall \omega \quad (2.20)$$

the transfer function R satisfies (2.18). Either formulation of the passivity requirement may be useful depending on the particular analysis.

Chapter 3

Algebraic Tradeoffs in Haptic Rendering

3.1 Introduction

Specifications on the closed-loop transfer functions S , T , and Θ form a multi-objective design problem. While controller synthesis techniques seek to satisfy these multiple requirements, intrinsic relationships between S , T , and Θ may prohibit certain combinations of feedback properties. We consider algebraic relationships between these transfer functions, that is, relationships that hold at any point s . Of primary concern are the tradeoffs imposed at points $s = j\omega$ as design requirements established in Chapter 2 constrain the closed-loop transfer functions along the imaginary axis.

3.2 Three-way Tradeoff between S , T , and Θ

In servo-feedback systems, performance goals of tracking and disturbance rejection are achieved by attenuating the Bode sensitivity function. In haptic rendering, however, attenuating the Bode sensitivity function $S(j\omega)$ is not equivalent to attenuating distortion $\Theta(j\omega)$ when $R_d(j\omega) \neq 0$. The class of feedback systems for which the Bode sensitivity function does not describe performance are termed **not reducible to a feedback loop**. These systems are inherently multivariable and subject to an algebraic tradeoff between performance and robustness properties which does not exist in servo-feedback systems [Freudenberg et al., 2003].

Algebraic relationships between S , T , and Θ arise due to a shared dependence on C . The diagram in Fig. 3.1 displays definitions of S , T , and Θ on the vertices of a triangle. Along the edges are relationships obtained by eliminating the controller C . Since the relationships do not depend explicitly on the controller, they reveal fundamental relationships

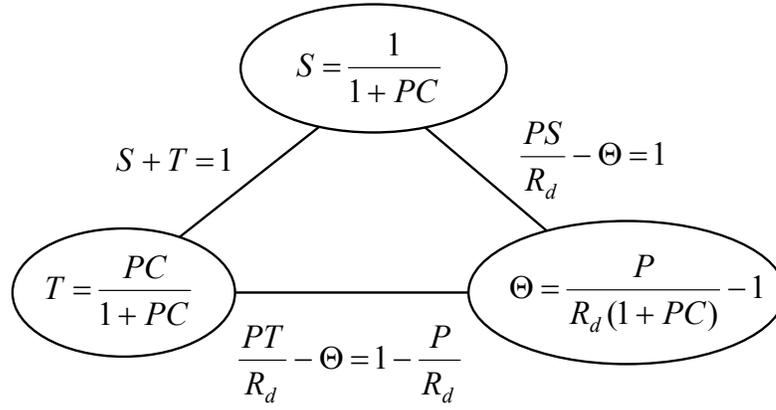


Figure 3.1 Algebraic relationships between S , T , and Θ . Relationships on the edges do not depend explicitly on C and imply intrinsic tradeoffs between feedback goals described by S , T , and Θ .

between performance, robustness properties, and bandwidth constraints that hold for all feedback designs.

The Bode sensitivity function and complementary sensitivity function are related by the identity

$$S + T = 1. \quad (3.1)$$

As a consequence, no controller can achieve arbitrary attenuation of S and T at a frequency. It follows from the triangle inequality that the specifications on S and T given by (2.15) and (2.17) are feasible only if

$$M_S(\omega) + M_T(\omega) \geq 1 \quad \forall \omega. \quad (3.2)$$

The tradeoff between attenuation of S and T is not particular to haptic rendering and is already analyzed for servo feedback design [Skogestad and Postlethwaite, 1997]. However, unlike servo feedback systems where S describes performance, the model-matching performance goal of haptic rendering is described by Θ . Additional tradeoffs are then implied between Θ and S and between Θ and T .

It follows from the definitions of S and Θ that sensitivity and performance are related by

$$\Theta = \frac{PS}{R_d} - 1. \quad (3.3)$$

A feasible set of requirements on S and Θ must satisfy

$$M_{\Theta}(\omega) + \left| \frac{P(j\omega)}{R_d(j\omega)} \right| M_S(\omega) \geq 1. \quad (3.4)$$

As $|P(j\omega)/R_d(j\omega)|$ approaches zero, the performance specification must be relaxed such that $M_{\Theta}(\omega) \geq 1$. The inequality (3.4) implies that reasonable performance and robustness requirements are not achievable when $R_d(j\omega) \gg P(j\omega)$. This condition arises when accurate rendering of the virtual environment dynamics requires significant feedback compensation of the hardware dynamics.

A separate conflict exists between performance and bandwidth constraints as dictated by

$$\Theta = \left(\frac{P}{R_d} - 1 \right) - \left(\frac{P}{R_d} \right) T. \quad (3.5)$$

Applying the triangle inequality, we find that the performance specification (2.12) and the bandwidth constraint (2.17) must satisfy

$$M_{\Theta}(\omega) + \left| \frac{P(j\omega)}{R_d(j\omega)} \right| M_T(\omega) \geq \left| \frac{P(j\omega)}{R_d(j\omega)} - 1 \right|. \quad (3.6)$$

As $M_{\Theta}(\omega)$ and $M_T(\omega)$ become small, the simultaneous constraints can only be satisfied when the difference between the virtual environment response $R_d(j\omega)$ and the open-loop hardware dynamics $P(j\omega)$ is small.

The inequalities (3.4) and (3.6) provide necessary conditions for the existence of a satisfactory feedback design given the haptic device and virtual environment dynamics. The algebraic relationship between Θ , S , and T may also be used to derive limitations on the set of virtual environment dynamics which may be rendered with a particular haptic device. First we consider the restrictions imposed by M_S and M_{Θ} . With some manipulation, (3.3) may be re-written as

$$R_d = \frac{PS}{1 + \Theta}. \quad (3.7)$$

For $M_{\Theta}(j\omega) < 1$, it follows from the triangle inequality that

$$|R_d(j\omega)| \leq |P(j\omega)| \frac{M_S(j\omega)}{1 - M_{\Theta}(j\omega)}. \quad (3.8)$$

For any virtual environment which violates this inequality, there exists no feedback design

that meets the performance and sensitivity specifications. To find restrictions imposed on R_d by M_T and M_Θ , we re-write (3.5) as

$$R_d = P \left(\frac{1+T}{1+\Theta} \right). \quad (3.9)$$

For $M_\Theta(j\omega) < 1$, the triangle inequality applied to both the numerator and denominator yields the inequality

$$|R_d(j\omega)| \leq |P(j\omega)| \left(\frac{1+M_T(j\omega)}{1-M_\Theta(j\omega)} \right). \quad (3.10)$$

For $M_T(j\omega) < 1$, application of the triangle inequality also yields a lower bound

$$|R_d(j\omega)| \geq |P(j\omega)| \left(\frac{1-M_T(j\omega)}{1+M_\Theta(j\omega)} \right). \quad (3.11)$$

The inequalities (3.8), (3.10), and (3.11) restrict, by magnitude frequency response, the class of virtual environment dynamics that may be rendered on a device while meeting specifications performance, sensitivity, and bandwidth requirements.

Example 3.2.1. Let the haptic device dynamics be given by a simple inertia $P = 1/s^2$. Sample design specifications M_Θ , M_S , and M_T are shown in Fig. 3.2. The performance specification M_Θ requires attenuation of distortion at low frequencies, the bandwidth specification M_T requires attenuation of the complementary sensitivity function at high frequencies, and across all frequencies, S must not exceed the bound $M_S = 3$. Figure 3.3 depicts resulting bounds on the magnitude of the virtual environment dynamics R_d . The upper trace is the upper bound on $|R_d|$ given by (3.8), and the lower trace is the lower bound on $|R_d|$ given by (3.11). (The upper bound given by (3.10) has not been included as it is nearly identical to (3.8).) For virtual environment dynamics whose magnitude frequency response enters the regions labeled “infeasible”, no feedback design exists that can achieve the desired combination of performance, sensitivity, and bandwidth.

3.3 Tradeoff between Performance and Sensitivity

The model-matching performance goal of haptic rendering introduces tradeoffs between Θ and S as well as between Θ and T . The tradeoff between Θ and T plays the same role as the tradeoff between S and T in typical servo-feedback design. In this section, we examine the tradeoff between performance and sensitivity in more detail as this tradeoff has no

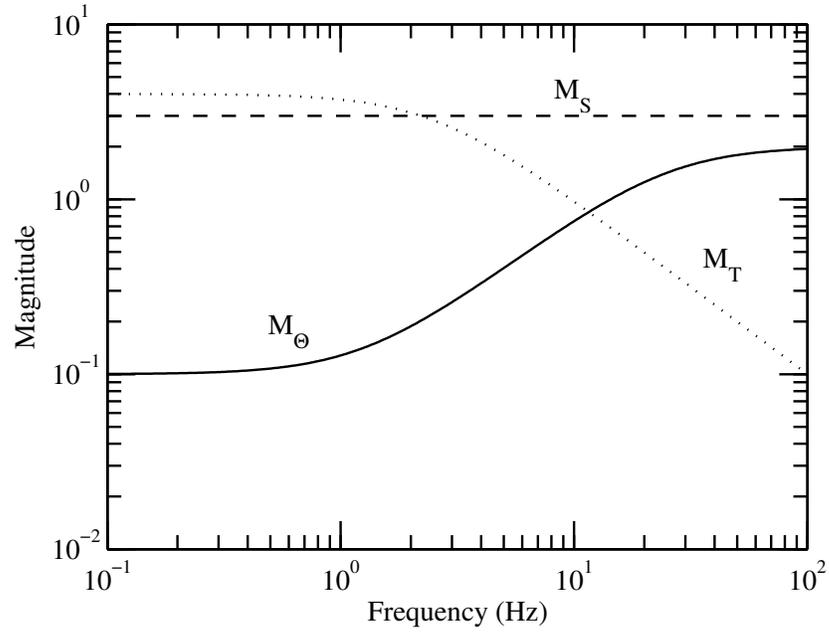


Figure 3.2 Design specifications M_Θ , M_S , and M_T .

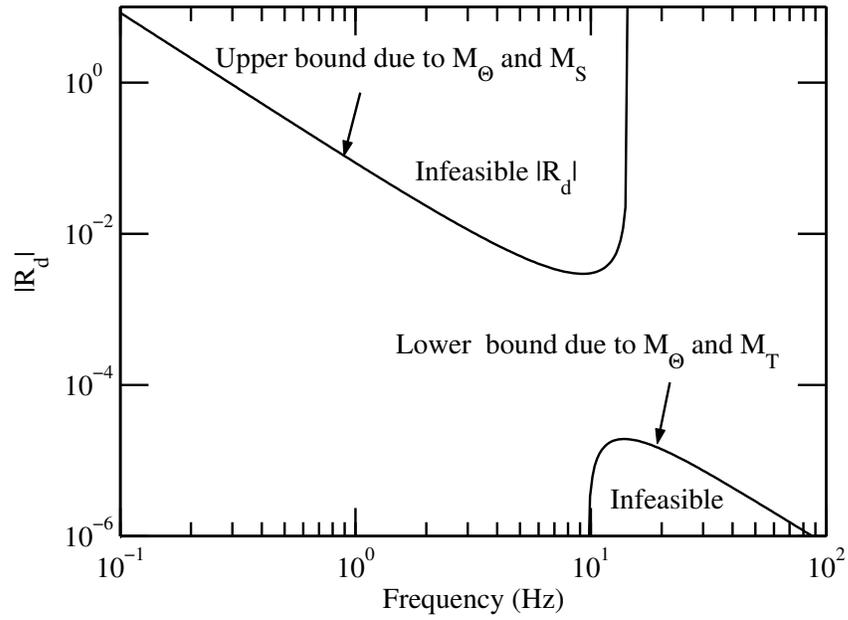


Figure 3.3 Upper and lower magnitude bounds on the set of feasible virtual environment dynamics. The upper trace is given by (3.8) and the lower trace is given by (3.11).

counter-part in servo-feedback design. Our aim is to quantify the intrinsic costs to attenuate performance or sensitivity. We show that the severity of the tradeoff may be captured in a single frequency dependent parameter. Let us begin by relating performance and sensitivity in terms of the general control configuration.

3.3.1 Background

As given by (2.7), the disturbance response T_{zw} of the general control configuration is $G_{zw} + G_{zu}K(1 - G_{yu}K)^{-1}G_{yw}$. The Bode sensitivity function for the general control configuration is $S = 1/(1 - G_{yu}K)$. Combining the expressions for T_{zw} and S , we eliminate the explicit dependence on the controller K . Let us define

$$\Gamma \triangleq G_{zu}G_{yw}/(G_{zw}G_{yu}). \quad (3.12)$$

Then we have the algebraic identity

$$\frac{T_{zw}}{G_{zw}} = 1 + \Gamma(S - 1). \quad (3.13)$$

Only when $\Gamma = 1$ is disturbance attenuation described by the Bode sensitivity function. There are two special cases where $\Gamma \equiv 1$: (a) systems whose performance variable is also the measured output, that is $G_{zu} = G_{yu}$ and $G_{zw} = G_{yw}$, and (b) systems whose control and exogenous input enter the system through the same dynamics, that is $G_{zu} = G_{zw}$ and $G_{yu} = G_{yw}$.

When $\Gamma \neq 1$, performance and robustness properties present competing goals. The severity of the tradeoff is determined by Γ and is generally frequency dependent. At frequencies where $|\Gamma(j\omega)| \ll 1$, the cost of attenuating the closed-loop disturbance response relative to the open-loop disturbance response is large amplification of $S(j\omega)$. On the other hand, at frequencies where $|\Gamma(j\omega)| \gg 1$, the cost of attenuating the Bode sensitivity function is large amplification of the closed-loop disturbance response relative to the open-loop disturbance response. Note that, while there is a limit to how small T_{zw} and S may be made at a frequency, both S and T_{zw} may be large.

3.3.2 Application to haptic rendering

We now interpret the tradeoff implied by (3.13) for haptic rendering with position feedback. In (3.13), the benefit of feedback in improving performance is gauged by the ratio T_{zw}/G_{zw} .

We denote the open-loop disturbance G_{zw} by Θ_o . This is distortion of the open-loop system. Substituting P for R in (2.5), we define Θ_o by

$$\Theta_o \triangleq (P - R_d)/R_d. \quad (3.14)$$

To distinguish open-loop distortion Θ_o from closed-loop distortion, we will refer to distortion Θ by Θ_c . At frequencies where the ratio Θ_c/Θ_o is greater than one, the use of feedback degrades performance compared with not using any feedback control.

Recall that in (3.13) the closed-loop disturbance response T_{zw} is Θ_c and the open-loop disturbance response G_{zw} is Θ_o . The general identity can then be written as

$$\frac{\Theta_c}{\Theta_o} = 1 + \Gamma(S - 1). \quad (3.15)$$

In the absence of feedback, the Bode sensitivity function S and the ratio Θ_c/Θ_o are precisely unity at all frequencies. Performance afforded by feedback control is gauged by attenuation of the ratio $|\Theta_c/\Theta_o|$ evaluated along the $j\omega$ -axis, and the benefit of feedback in terms of reduced sensitivity to the haptic device dynamics is gauged by attenuation of $|S|$ evaluated along the $j\omega$ -axis. The algebraic relationship captured by (3.15) limits the ability to attenuate both quantities simultaneously.

The tradeoff severity Γ is found by substituting (2.9) into (3.12). Further substituting open-loop distortion Θ_o for $(P - R_d)/R_d$, we find that

$$\Gamma = 1 + \frac{1}{\Theta_o}. \quad (3.16)$$

Given a fixed device model P , the term $1/\Theta_o$ approaches 0 as R_d approaches 0. For $R_d \neq 0$, the tradeoff severity $\Gamma \neq 1$ and there exists a tradeoff between attenuating distortion and attenuating the Bode sensitivity function. The tradeoff is most severe at frequencies where $\Theta_o \rightarrow -1$ and frequencies where $\Theta_o \rightarrow 0$.

We first consider the situation where $\Theta_o \rightarrow -1$ and thus $\Gamma \rightarrow 0$. From the definition of Θ_o , we see that $\Theta_o \rightarrow -1$ at frequencies where $|R_d(j\omega)| \gg |P(j\omega)|$. At these frequencies, the magnitude of the virtual environment $|R_d(j\omega)|$ is much greater than the magnitude of the haptic device dynamics $|P(j\omega)|$, and partial cancellation of the device dynamics is required to present R_d accurately. For example, the virtual environment might have less mass than the intrinsic dynamics of the haptic device. The cost of partially cancelling device dynamics, however, is amplification of the Bode sensitivity function.

Let us alternatively consider what happens as $\Theta_o \rightarrow 0$ and thus $\Gamma \rightarrow \infty$. In this situation, the response of the virtual environment $R_d(j\omega)$ is close the open-loop response of the

haptic device $P(j\omega)$. Little or no feedback is required to achieve low distortion since open-loop distortion is already nearly 0. However, the Bode sensitivity function approaches 1 as the feedback gain approaches 0. We may use feedback to attenuate sensitivity, but only by accepting large amplification of $|\Theta_c(j\omega)/\Theta_o(j\omega)|$.

An important consequence of (3.15) is that, at frequencies where $\Gamma \neq 1$, any feedback design that attenuates the ratio Θ_c/Θ_o cannot also attenuate the Bode sensitivity function S . Furthermore, regardless of the controller synthesis technique or controller complexity, $S \rightarrow 1 - 1/\Gamma$ at frequencies where $\Theta_c/\Theta_o \rightarrow 0$. Substituting (3.16) for Γ , we reduce this limit to

$$\lim_{\Theta_c/\Theta_o \rightarrow 0} S = \frac{R_d}{P}. \quad (3.17)$$

At frequencies where $|R_d(j\omega)/P(j\omega)|$ is large, accurate rendering of the virtual environment can only be achieved by accepting very poor robustness to variations in the haptic device dynamics.

We may also establish a bound on the difference between S and R_d/P given closed-loop distortion. It follows from (3.15) and (3.16) that

$$S = \frac{R_d}{P}(1 + \Theta_c). \quad (3.18)$$

Then the difference $S - R_d/P$ is simply $(R_d/P)\Theta_c$, and the relative error between S and R_d/P is given by:

$$\frac{S - R_d/P}{R_d/P} = \Theta_c. \quad (3.19)$$

Thus, at a frequency, closed-loop distortion $|\Theta_c(j\omega)|$ indicates how close $S(j\omega)$ is to $R_d(j\omega)/P(j\omega)$.

The virtual environment R_d and haptic device model P are transfer functions from force to motion; thus a large magnitude of either transfer function corresponds to a small mechanical impedance. It is then not surprising that, as given by (3.17), poor sensitivity results when we accurately render a virtual environment with a small mechanical impedance relative to the mechanical impedance of the haptic device. We note, however, that this sensitivity does not arise from cancellation achieved through feed-forward control. The controller C in Fig 2.3 is in feedback around the haptic device P . Furthermore, the sensitivity to hardware dynamics induced by feedback control may be much greater than the unity sensitivity of feed-forward control. Perhaps less intuitive is the frequency dependent

nature of the tradeoff. For instance, rendering a pure spring with low distortion using a haptic device with inertia induces sensitivity that increases with frequency.

Example 3.3.1. Let us consider the tradeoff severity Γ for two numerical examples. For both examples we use a simple, ideal-mass device model

$$P = \frac{1}{s^2}. \quad (3.20)$$

First we consider a virtual spring environment $R_d = 1$. Figure 3.4 shows the Bode diagram for the device dynamics P , the virtual environment dynamics R_d , and the resulting tradeoff severity Γ . At low frequencies where $|P(j\omega)| \gg |R_d(j\omega)|$, Γ approaches unity, indicating no tradeoff between performance and sensitivity. In other words, by rendering the virtual spring at low frequencies, we also attenuate the Bode sensitivity function. On the other hand, at high frequencies $|R_d(j\omega)|$ is much greater than $|P(j\omega)|$ and thus Γ approaches zero. Thus the higher the frequency at which we wish to attenuate Θ_c/Θ_o , the greater amplification of S we must accept.

As a second example, consider the virtual environment

$$R_d = \frac{1}{s^2 + 0.1s + 1}. \quad (3.21)$$

This ideal mass-spring-damper model has a resonant frequency of 1.05 rad/s, is lightly damped ($\zeta = 0.053$), and has the same mass parameter as the device P . Figure 3.5 shows the Bode diagram for P , R_d , and the resulting tradeoff severity Γ . At low frequencies where Γ is nearly unity, there is not a significant tradeoff between performance and sensitivity. At the resonant peak, Γ approaches zero indicating that large amplification of S is required to render the resonant peak in R_d with the given device dynamics. At high-frequencies, Γ becomes large, indicating that attenuation of sensitivity requires amplification of Θ_c/Θ_o . Inspection of Θ_o shown in Fig. 3.6 reveals that the open-loop distortion approaches zero at high frequencies. This is a consequence of the fact that the virtual environment and the device dynamics have the same high frequency behavior. Since Θ_o is already small at high frequencies, using feedback to achieve robustness to hardware variations causes large amplification of Θ_c/Θ_o .

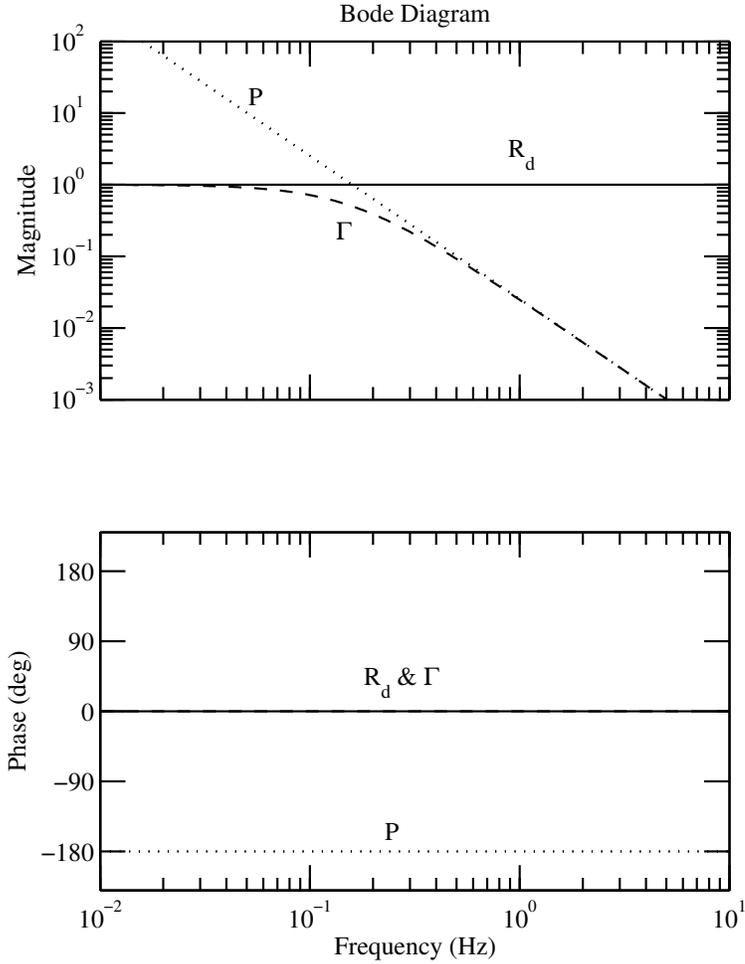


Figure 3.4 Tradeoff severity Γ for a simple mass device P and a virtual spring environment R_d .

3.4 Existence of a Proper, Stabilizing Cancellation Controller

The limit (3.17) predicts the cost to achieve perfect attenuation of distortion; however stability requirements may limit achievable performance. For the model-matching problem of haptic rendering, there exists an exact finite gain controller that achieves perfect performance. Recall that for haptic rendering

$$\begin{bmatrix} G_{zw} & G_{zu} \\ G_{yw} & G_{yu} \end{bmatrix} = \begin{bmatrix} \frac{P}{R_d} - 1 & \frac{P}{R_d} \\ P & P \end{bmatrix}. \quad (3.22)$$

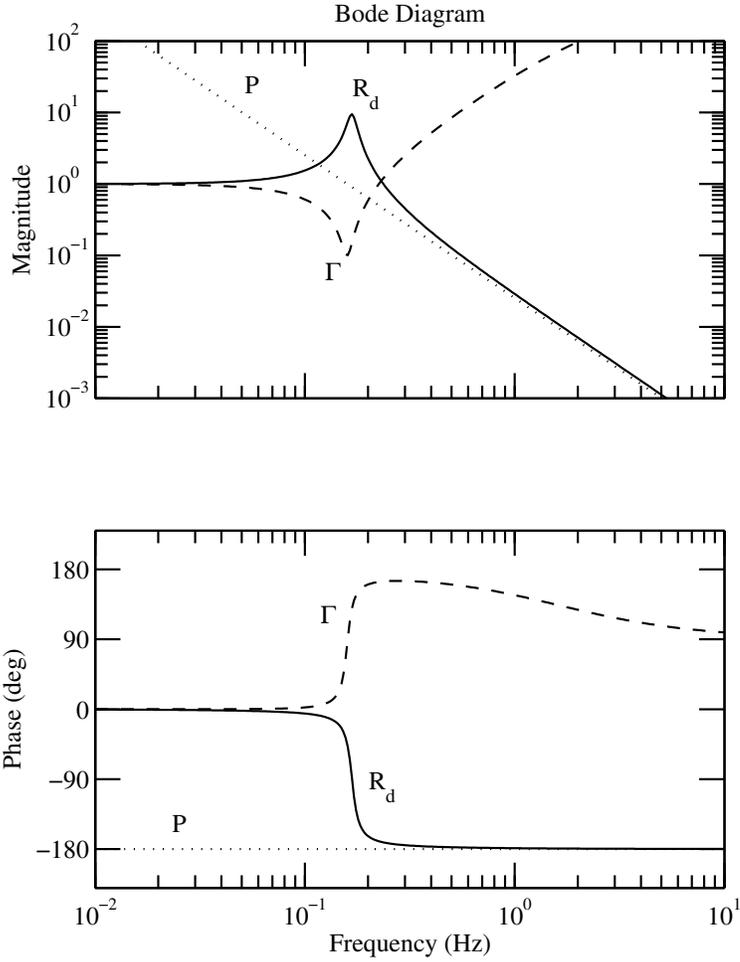


Figure 3.5 Tradeoff severity Γ for a simple mass device P and a virtual mass-spring-damper environment R_d .

and

$$T_{zw} \triangleq G_{zw} + G_{zu}K(1 - G_{yu}K)^{-1}G_{yw} \quad (3.23)$$

$$= \Theta. \quad (3.24)$$

If $P \neq R_d$, the controller that achieves $\Theta = 0$ is

$$K^C = \frac{G_{zw}}{G_{zw}G_{yu} - G_{zu}G_{yw}} \quad (3.25)$$

$$= \frac{1}{P} - \frac{1}{R_d}. \quad (3.26)$$

This controller is termed the **cancellation controller** as it completely blocks transmission from exogenous inputs w to the performance outputs z . However, (3.26) may not be sta-

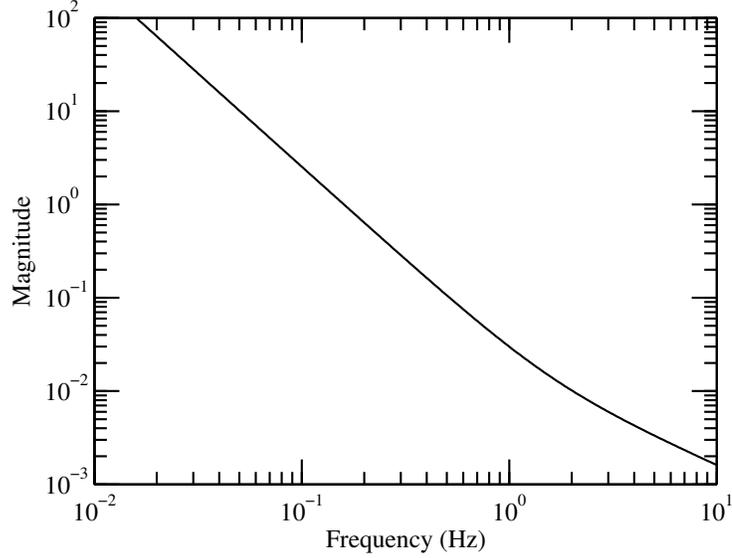


Figure 3.6 Open-loop distortion Θ_o for a simple mass device P and a virtual mass-spring-damper environment R_d . Distortion approaches zero at high frequency because P and R_d have the same mass parameter.

bilizing. The feedback interconnection of the haptic device P with controller $C = K^C$ is bounded-input/bounded-output stable if and only if the four closed-loop transfer functions S , PS , CS , and PCS have no closed right-half plane poles (CRHP) and the transfer functions P and K^C are proper. We begin by finding conditions under which the cancellation controller (3.26) yields closed-loop transfer functions S , PS , CS , and PCS whose poles are contained entirely within the open left-half plane (OLHP).

Proposition 3.4.1. *If P has no CRHP zeros, R_d has no CRHP poles, and $C = K^C$, then the poles of S , PS , CS , and PCS lie in the OLHP.*

Proof. From (2.4) and (2.13) we can show that $PS = R$. Since K^C achieves zero distortion, R equals R_d and PS has no CRHP poles. It follows that any CRHP poles of S must be CRHP zeros of P . However, by hypothesis, P has no CRHP zeros, so S also contains no CRHP poles.

It remains to be shown that P and C have no unstable pole-zero cancellations. By hypotheses, P has no CRHP zeros and R_d has no CRHP poles. At any CRHP pole of P , K^C is given by $-1/R_d$ which is non-zero since R_d has no CRHP poles. Since there are no unstable pole-zero cancellations between P and C and since S has no CRHP poles, $S = 1/(1 + PC)$ has CRHP zeros at the CRHP poles of P and C . It follows that CS and PCS have no CRHP poles. \square

Suppose that the hypotheses of Proposition 3.4.1 are satisfied. If P is proper and K^C is not strictly proper, Proposition IV.7 in [Freudenberg et al., 2003] provides a procedure to obtain a strictly proper, stabilizing approximation to K^C over any finite bandwidth.

For many haptic applications, the controller is separated from the virtual environment dynamics as shown in Fig. 2.4. This architecture is particularly useful for complex virtual environments. For example, in a medical training application, the virtual environment may include a real-time simulation of deformable surfaces. The cancellation controller (3.26) then has limited utility as it integrates the virtual environment dynamics and hardware compensation. A virtual coupler serves as an interface between virtual environment dynamics and control of the physical device [Adams et al., 1998]. In the next Chapter, we introduce a novel virtual coupler design that provides the required separation between hardware compensation and virtual environment dynamics and also recovers the performance of the cancellation controller.

Chapter 4

Virtual Coupler Design for Haptic Rendering

4.1 Introduction

We now discuss design of the generalized virtual coupler introduced in Chapter 2. Our approach is to find a parameterization of the virtual coupler design which provides useful terms for tuning the closed-loop distortion. We then generate design directives for optimizing performance, and use such a design in Section 4 to demonstrate experimentally the sensitivity induced as performance improves. We compare this virtual coupler design with the standard design practice. The algebraic analysis in Chapter 3 permits us to evaluate the relative quality of the designs with respect to the tradeoff between performance and sensitivity.

4.2 Performance of the Generalized Virtual Coupler

The haptic interface feedback controller C is typically partitioned into the virtual environment and virtual coupler. The virtual coupler is generally fixed and should accommodate a range of virtual environments. Referring to Fig. 2.4, the virtual coupler describes the response from the haptic device position y and desired position y'_d to the motor command u and a virtual environment force f_e . This set of input/output responses is often modeled after a physical system such as a spring [Adams and Hannaford, 1999]. Let us simply express

these input/output relationships by the matrix of transfer functions

$$V \triangleq \begin{bmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{bmatrix} \quad (4.1)$$

such that

$$\begin{bmatrix} u \\ f_e \end{bmatrix} = V \begin{bmatrix} y \\ y'_d \end{bmatrix}. \quad (4.2)$$

Figure 4.1 shows the block diagram of the haptic interface system with V interposed between the virtual environment and haptic device. We note that V is more general than the virtual coupler described in [Colgate et al., 1995] or extended in [Adams and Hannaford, 1999]. In contrast to prior work, each element of V may be any transfer function rather a particular structure such as a proportional-derivative term; further, we do not assume any particular relationship between the elements such as making all terms equal. We refer to V as a generalized virtual coupler.

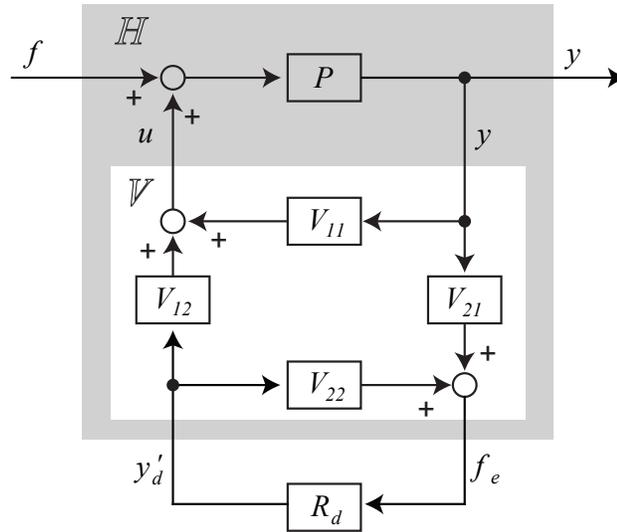


Figure 4.1 Controller partitioned into a generalized virtual coupler V and virtual environment R_d . Design directives to optimize performance of the generalized virtual coupler are found by considering the multivariable response of H .

To focus on the design of V , we find it useful to remove the virtual environment from the problem and just consider the feedback interconnection of the haptic device P with the generalized virtual coupler V . Referring to Fig. 4.1, the input/output response of P in

feedback with V is

$$\begin{bmatrix} y \\ f_e \end{bmatrix} = \begin{bmatrix} \frac{P}{1 - V_{11}P} & \frac{V_{12}P}{1 - V_{11}P} \\ \frac{V_{21}P}{1 - V_{11}P} & V_{22} + \frac{V_{21}V_{12}P}{1 - V_{11}P} \end{bmatrix} \begin{bmatrix} f \\ y'_d \end{bmatrix}. \quad (4.3)$$

As we will show, the terms of the four input-output responses of (4.3) are easily related to closed-loop distortion when R_d is reconnected to the virtual coupler.

Let us denote the matrix of four transfer functions in (4.3) by H in the form of the hybrid matrix (where our mapping is between force and position rather than force and velocity):

$$\begin{bmatrix} y \\ f_e \end{bmatrix} = \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} \begin{bmatrix} f \\ y'_d \end{bmatrix}. \quad (4.4)$$

The multivariable responses of H are indicated by the dark box in Fig. 4.1. For $H_{11} \neq 0$, the elements of H uniquely determine the generalized virtual coupler V according to

$$\begin{bmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{bmatrix} = \begin{bmatrix} \frac{H_{11} - P}{H_{11}P} & \frac{H_{12}}{H_{11}} \\ \frac{H_{21}}{H_{11}} & H_{22} - \frac{H_{12}H_{21}}{H_{11}} \end{bmatrix}. \quad (4.5)$$

We then use H as a re-parameterization of the generalized virtual coupler V .

Before we connect the virtual environment to H , let us remark on the role of the elements of H . Referring to (4.4), the response from the desired position of the haptic device y'_d to the actual position y is described by H_{12} . Clearly, for small error between y and y'_d , H_{12} must be nearly unity. The virtual environment describes the desired response of y'_d to the human operator force f ; however, referring to Fig. 4.1, we see that the virtual environment generates y'_d in response to f_e not f . Thus, to generate the correct desired position y'_d , the term H_{21} , which describes the response from f to f_e , must be also nearly unity. The remaining terms H_{11} , which describes the feed-through from f to y , and H_{22} , which describes the feed-through from y'_d to f_e , should be attenuated. It follows that perfect performance is achieved by [Adams et al., 1998]

$$\begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}. \quad (4.6)$$

We remark that this condition is analogous to ideal kinesthetic coupling in teleoperation [Yokokohji and Yoshikawa, 1994].

Inherent limitations of the feedback structure in Fig. 4.1 prevent (4.6) from being realized. The term H_{11} describes the feed-through from f to y , and from (4.3) and (4.4)

$$H_{11} = \frac{P}{1 - V_{11}P}. \quad (4.7)$$

This term cannot be made identically zero with finite feedback V_{11} . Further, suppose that V_{11} is proper and P is strictly proper. Then H_{11} must approach the open-loop dynamics P at high frequencies. From (4.3) and (4.4)

$$H_{12} = V_{12}H_{11} \quad (4.8)$$

$$H_{21} = V_{21}H_{11}. \quad (4.9)$$

If the controller elements V_{12} and V_{21} are proper, it follows that H_{12} and H_{21} are strictly proper and can only be close to unity over a limited bandwidth. The term H_{22} describes the response from y'_d to f_e and is given by

$$H_{22} = V_{22} + V_{21}V_{12}H_{11}. \quad (4.10)$$

The assumptions on H_{11} , V_{21} , and V_{12} imply that $V_{21}V_{12}H_{11}$ is strictly proper. We will take advantage of this fact later to make H_{22} identically zero with a strictly proper controller element V_{22} .

Let us now re-introduce the virtual environment in the feedback loop and compute distortion. Computing the response from f to y in Fig. 4.1, we find that

$$R = H_{11} + \frac{H_{12}H_{21}R_d}{1 - H_{22}R_d}. \quad (4.11)$$

Recall that closed-loop distortion Θ_c is given by $(R - R_d)/R_d$. Substituting (4.11) into Θ_c , we find that

$$\Theta_c = \frac{H_{11}}{R_d} + \frac{H_{12}H_{21}}{1 - H_{22}R_d} - 1. \quad (4.12)$$

If, as in (4.6), H_{11} and H_{22} are 0, and H_{12} , and H_{21} are 1, then Θ_c is 0 regardless of the virtual environment. However, practical virtual coupler designs will fall short of achieving zero distortion since the requirements of (4.6) violate inherent limitations of the feedback structure.

4.3 Optimizing for Performance

Ideally, the virtual coupler would render any virtual environment dynamics R_d with low distortion. However, a practical performance objective is to achieve low distortion over a finite bandwidth for a limited class of virtual environments. Let us consider guidelines for attenuating distortion for classes of virtual environments R_d defined by their magnitude frequency response. To reduce (4.12) without exploiting phase information of R_d , we must attenuate both H_{11} and $(H_{12}H_{21})/(1 - H_{22}R_d) - 1$. To attenuate the latter term, we can select H_{22} such that $|H_{22}R_d| \ll 1$ and $H_{12}H_{21} \approx 1$. Referring to (4.3), we see that H_{22} can in fact be made identically zero if $V_{22} = -V_{21}V_{12}P/(1 - V_{11}P)$. As noted in the previous section, this cancellation, accomplished by summation rather than inversion, may be realized using only proper elements in V . We call a generalized virtual coupler for which $H_{22} \equiv 0$ a **cancellation coupler**.

Closed-loop distortion achieved by the cancellation coupler is tuned by the free parameters H_{11} , H_{12} , and H_{21} . Through our choice of parameters, we can guarantee that Θ_c satisfies an upper bound for a range of virtual environments. Distortion for the cancellation coupler is found by setting $H_{22} = 0$ in (4.12):

$$\Theta_c = \frac{H_{11}}{R_d} + H_{12}H_{21} - 1. \quad (4.13)$$

Then distortion may be bounded for a class of environment dynamics R_d by application of the triangle inequality. The term $H_{12}H_{21} - 1$ depends only on the design of the virtual coupler, and the term H_{11}/R_d can be upper bounded for a set of environment dynamics whose magnitude satisfies a lower bound. Let $\varepsilon(\omega)$ be a positive function and consider the class of virtual environment dynamics that satisfy

$$|R_d(j\omega)| \geq \varepsilon(\omega), \quad \forall \omega. \quad (4.14)$$

It follows from (4.13) that for this class of virtual environments

$$|\Theta_c(j\omega)| \leq \frac{1}{\varepsilon(\omega)} |H_{11}(j\omega)| + |H_{12}(j\omega)H_{21}(j\omega) - 1|. \quad (4.15)$$

Examining (4.13), we note that distortion grows as the magnitude of the virtual environment approaches zero. This is an intrinsic limitation of the finite element V_{11} in the virtual coupler which can only achieve finite attenuation of H_{11} .

The following is a practical procedure to design a cancellation coupler for a strictly proper P :

1. Design a proper, stable, and stabilizing controller V_{11} around the haptic device dynamics P .
2. Choose stable filters H_{12} and H_{21} that have relative degree greater than or equal to the relative degree of P and that attenuate $|H_{12}(j\omega)H_{21}(j\omega) - 1|$ at low frequencies.
3. Compute H_{11} from V_{11} and use (4.15) to determine whether performance is acceptable for virtual environments to be rendered.
4. Redesign V_{11} , H_{12} , and H_{21} to achieve greater attenuation of H_{11} and $H_{12}H_{21} - 1$ if necessary.
5. Set $H_{22} = 0$ and solve for the controller elements V_{12} , V_{21} , and V_{22} with (4.5).

The procedure results in proper terms V_{12} , V_{21} , and V_{22} . Referring to (4.11), we note that, with $H_{22} = 0$, the poles of closed-loop rendered virtual environment R include the poles of the virtual environment R_d and the stable poles of the filters H_{11} , H_{12} , and H_{21} . While maintaining the poles of the virtual environment is desirable, we caution that there are desirable properties which the above design procedure may not achieve. For instance, although the virtual environment dynamics R_d may model a passive mechanical system, passivity may not be preserved in R .

The design directives for the cancellation coupler minimize distortion at the cost of other feedback goals such as sensitivity. However one may reasonably choose a virtual coupler design that, compared with the cancellation coupler, sacrifices performance but improves sensitivity. In the next section, we investigate how two virtual coupler designs, one common virtual coupler design and one cancellation coupler, strike different balances between performance and sensitivity.

4.4 Experimental Results

We now demonstrate experimentally the consequences of the algebraic tradeoff between performance and sensitivity given by (3.15). To determine the sensitivity of a particular feedback design to the haptic device dynamics, we vary the haptic device dynamics and observe the variations in the rendered response. We compare the performance and sensitivity of two virtual coupler designs rendering two different virtual environments. One feedback design uses a cancellation coupler while the other is a **spring-damper** virtual coupler, a typical design modeled after the response of a parallel spring and damper. As we

will show, the cancellation coupler is capable of achieving better performance than the virtual coupler; however we do not compare these designs on performance alone. We instead consider the compromise between the competing objectives of performance and sensitivity inherent to both designs.

We render two mass-spring-damper virtual environments on a single-axis rotary interface shown in Fig. 4.3 and described in [Gillespie et al., 2003]. The experimentally determined device dynamics are

$$P = \frac{1550}{s^2 + 0.775s} \text{ (rad/N-m)}. \quad (4.16)$$

Both virtual environments have a natural frequency of 2 Hz and a damping ratio of 0.2, but the gain of the second system is five times the gain of the first. Let R_{d1} be the first virtual environment, given by

$$R_{d1} = \frac{474}{s^2 + 5.03s + 158} \text{ (rad/N-m)}. \quad (4.17)$$

Then let R_{d2} be the second virtual environment, given by

$$R_{d2} = 5R_{d1} \text{ (rad/N-m)}. \quad (4.18)$$

In mechanical terms, R_{d2} has 1/5 the mass, damping, and spring stiffness of R_{d1} . Figure 4.2 (left) shows the frequency response of R_{d1} and R_{d2} plotted alongside the haptic wheel dynamics P . For R_{d1} and R_{d2} rendered on the haptic wheel P , we have the tradeoff severity Γ shown in Fig. 4.2 (right).

Let us highlight an important difference between rendering R_{d1} and R_{d2} given the haptic device dynamics P shown in Figure 4.2. While $|R_{d1}(j\omega)|$ is less than $|P(j\omega)|$ at all frequencies, $|R_{d2}(j\omega)|$ is greater than $|P(j\omega)|$ at frequencies above 1 Hz (or $\omega = 2\pi$). In terms of model parameters, the inertia of R_{d1} is greater than the device inertia, but the inertia of R_{d2} is less than the device inertia. Partial cancellation of the device dynamics is necessary to render R_{d2} accurately.

The tradeoff severity Γ shown in Fig. 4.2 indicates that there is little to no algebraic tradeoff between performance and sensitivity at low frequencies. However, the dip in Γ for R_{d2} near 2 Hz reflects the partial cancellation of hardware dynamics required to render the resonant peak. At this frequency the cost of accurately rendering R_{d2} is significant amplification of sensitivity.

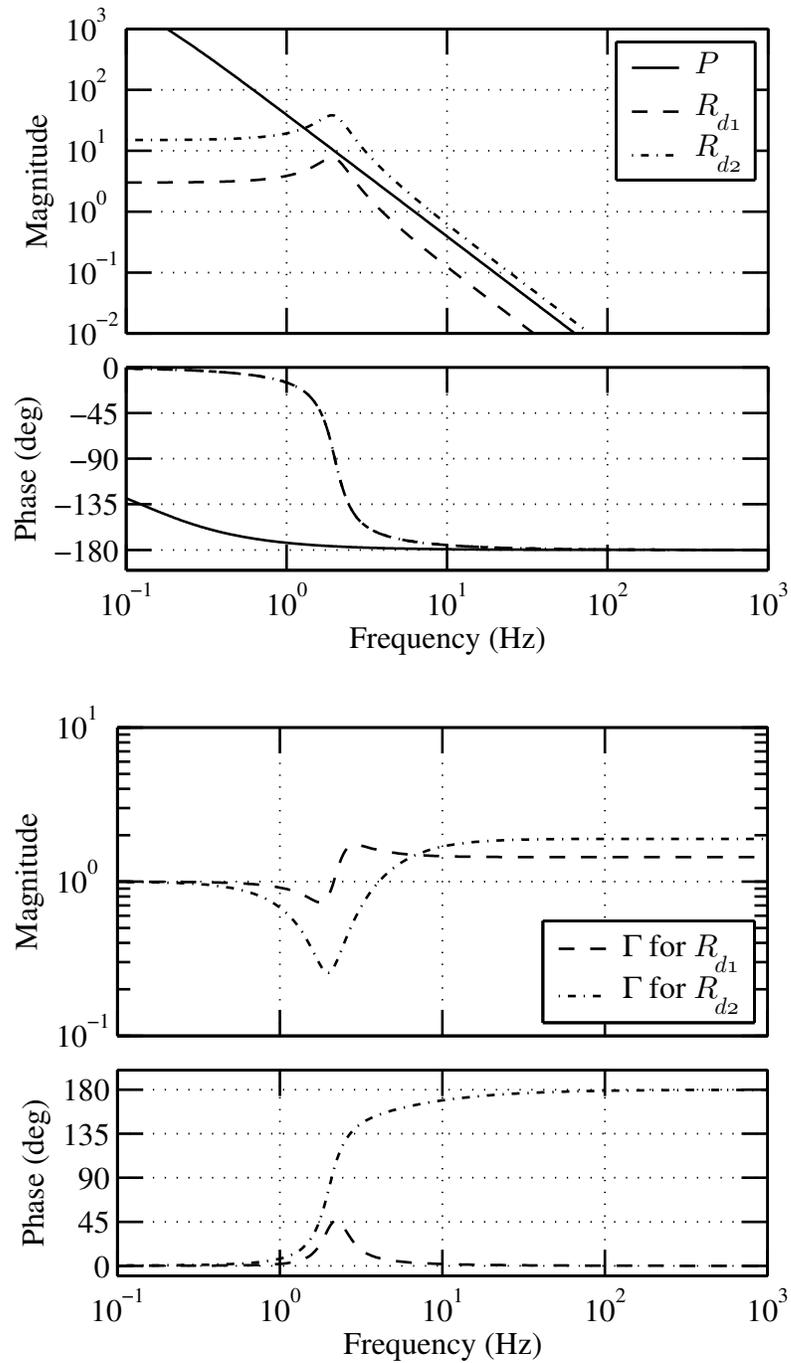


Figure 4.2 Frequency responses of R_{d1} and R_{d2} relative to the haptic device P (top) and resulting tradeoff severities Γ (bottom). As indicated by $\Gamma \approx 1$, no tradeoff exists at low frequencies where $|P|$ is much larger than $|R_{d1}|$ or $|R_{d2}|$. The most severe tradeoff between performance and sensitivity occurs for R_{d2} near 2 Hz. At this frequency, Γ approaches 0 as $|R_{d2}|$ significantly exceeds $|P|$. Note that phase also provides important information; where $|\Gamma|$ crosses 1 near 4–5 Hz, the phase plot indicates that the complex value of Γ is not in fact close to 1, and thus a significant tradeoff still exists.

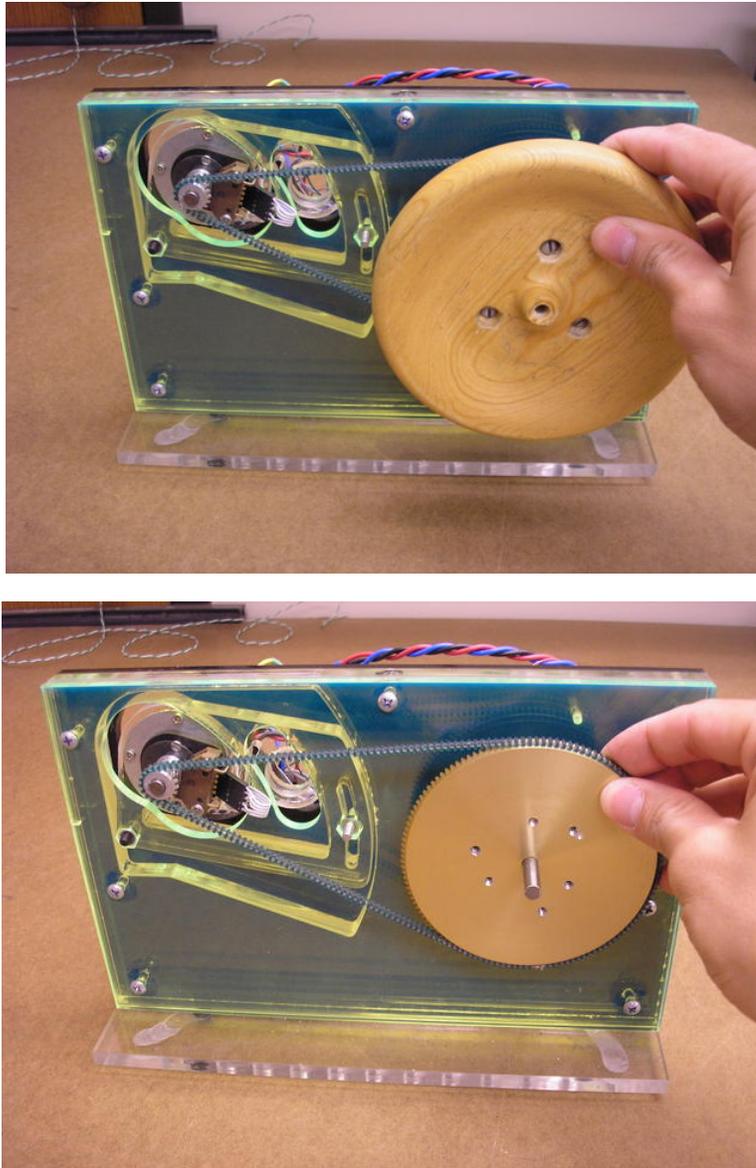


Figure 4.3 Single-axis, rotary haptic device with wooden handwheel (above) and without (below). Controller design was performed on a reduced-order model of the haptic device with the wooden handwheel. System identification on the device without the handwheel showed a 30% reduction in rotational inertia.

4.4.1 Spring-Damper Virtual Coupler Design

A common technique for designing a virtual coupler is to model the dynamics of B after the dynamics of a mechanical coupling such as parallel spring-damper [Colgate et al., 1995; Adams and Hannaford, 1999; Miller et al., 2000]. A proper implementation of the spring-damper virtual coupler is a network of lead filters B given by

$$\begin{bmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{bmatrix} = \begin{bmatrix} \frac{-bs-k}{\tau s+1} & \frac{bs+k}{\tau s+1} \\ \frac{bs+k}{\tau s+1} & \frac{-bs-k}{\tau s+1} \end{bmatrix}. \quad (4.19)$$

The scalar parameters k and b adjust the stiffness and damping of the virtual coupler, and the pole at $s = -1/\tau$ makes the elements of V proper.

A recommended technique for tuning the spring-damper virtual coupler is to first set $R_d = 0$. In Fig. 4.1, the virtual coupler and virtual environment reduce to simply the lead filter V_{11} . This term is then tuned using loop-shaping methods appropriate for servo-control problems [Lawrence et al., 2000]. As is often the case when tuning servo controllers, the presence of high-frequency unmodeled dynamics limits the closed-loop bandwidth. For our experimental setup, we select $k = 17.1$ (N-m/rad) and $b = 0.203$ (N-m-s/rad) which results in an open-loop transfer function PV_{11} with 60° of phase margin and a gain cross-over frequency of 50 Hz. Higher bandwidth is not possible with our hardware as it excites a resonance in the chain drive that connects our motor and handwheel. The pole at $-1/\tau$ was selected to have a break-frequency of 185 Hz, significantly above the 50 Hz crossover frequency and well below the 1 kHz sampling-frequency. A Tustin approximation of the spring-damper virtual coupler provided the digital implementation of V .

The predicted performance of the spring-damper virtual coupler rendering R_{d1} and R_{d2} is shown in Fig. 4.4. The solid traces are closed-loop distortion Θ_c given by (2.5), and the dashed traces are open-loop distortion Θ_o given by (3.14). Attenuation of the ratio Θ_c/Θ_o indicates performance gain achieved through feedback. Closed-loop distortion Θ_c approaches open-loop distortion Θ_o at high-frequencies due to the limited closed-loop bandwidth. For the spring-damper virtual coupler, significant attenuation of Θ_c/Θ_o is achieved below 1 Hz; however above 1 Hz performance results are mixed. For R_{d1} , attenuation of Θ_c is only slightly smaller than Θ_o near 2 Hz, and for R_{d2} , closed-loop distortion is actually worse than open-loop distortion above 2 Hz.

Conventional wisdom directs us to increase the stiffness k and damping b of the virtual coupler to improve performance; however the poor rendering of R_{d2} is not due to

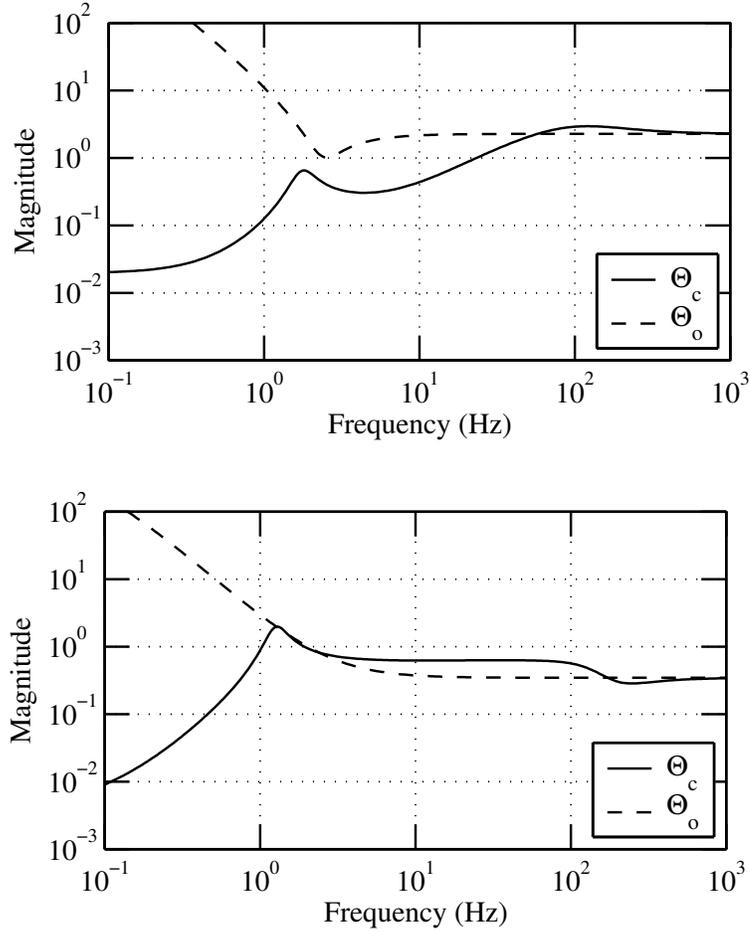


Figure 4.4 Open-loop and closed-loop distortion for the spring-damper virtual coupler rendering R_{d1} (top) and R_{d2} (bottom).

insufficient gains. Let $F \triangleq (bs + k)/(\tau s + 1)$. It follows from (4.5) and (4.19) that

$$H = \begin{bmatrix} \frac{P}{1+PF} & \frac{PF}{1+PF} \\ \frac{PF}{1+PF} & \frac{-F}{1+PF} \end{bmatrix}. \quad (4.20)$$

Closed-loop distortion Θ_c for the virtual coupler may be computed by substituting (4.20) into (4.12). With some algebraic manipulation, distortion for the virtual coupler reduces to

$$\Theta_c = \frac{P - R_d - R_d^2 F}{R_d [1 + F(P + R_d)]}. \quad (4.21)$$

We may strengthen the virtual coupler by increasing k and b . However, this strategy does

not drive Θ_c to 0. This is apparent by considering $|\Theta_c(j\omega)|$ as $|F(j\omega)| \rightarrow \infty$

$$\lim_{|F(j\omega)| \rightarrow \infty} |\Theta_c(j\omega)| = \left| \frac{R_d(j\omega)}{P(j\omega) + R_d(j\omega)} \right|. \quad (4.22)$$

Distortion is small only if $|R_d(j\omega)|$ is small relative to $|P(j\omega)|$. Mechanically speaking, the high-gain virtual coupler rigidly attaches the virtual environment to the haptic device, but does not cancel device dynamics. At frequencies where the virtual environment dynamics are already similar to the haptic device dynamics (i.e. at frequencies where the magnitude of Γ is much greater than unity), open-loop distortion is small and the virtual coupler actually has worse performance in closed-loop than open-loop.

Let us now consider the predicted sensitivity of the spring damper virtual coupler design to hardware variations. The predicted Bode sensitivity function to render R_{d1} and R_{d2} is given by the solid traces in Fig. 4.5. Since S is identically 1 in the absence of feedback, any attenuation of $|S(j\omega)|$ below 1 indicates the beneficial effect of feedback in reducing sensitivity to hardware dynamics. The dashed trace shows the limit for S as the ratio of closed-loop distortion to open-loop distortion Θ_c/Θ_o approaches 0. As predicted by the tradeoff severity Γ shown in Fig. 4.2, driving Θ_c/Θ_o to zero near 2 Hz will induce a peak in sensitivity when rendering R_{d2} but not when rendering R_{d1} .

Due to the tradeoff relationship (3.15), there are algebraic limits on the ability of any feedback design to simultaneously attenuate $\Theta_c(j\omega)/\Theta_o(j\omega)$ and $S(j\omega)$. Let us first focus on frequencies below 1 Hz where Θ_c/Θ_o is small for both R_{d1} and R_{d2} as shown in Fig. 4.4. In this frequency band, the Bode sensitivity function S , indicated by the solid traces in Fig. 4.5, must approach the limit R_d/P as dictated by (3.19). This limit becomes large at frequencies where Γ shown in Fig. 4.2 becomes small. At certain frequencies, the virtual coupler design is suboptimal with respect to the algebraic tradeoff relationship (3.15). Between 1 and 2 Hz, we see that $|\Theta_c/\Theta_o|$ is approximately 1 for R_{d2} and that $|S|$ is greater than 1. Thus, in this frequency range, the virtual coupler amplifies sensitivity to hardware dynamics while not reducing distortion Θ_c relative to open-loop Θ_o .

Experimental step-responses of R_{d1} and R_{d2} rendered by the spring-damper virtual coupler are shown in Fig. 4.6. The applied step torque is selected such that the DC response of both virtual environments is a half rotation of the hand wheel. The solid trace is the virtual environment (the desired response) and the dashed trace labeled ‘‘R (nominal P)’’ is the actual response. The predicted closed-loop distortion Θ_c (shown in Fig. 4.4) is better for R_{d1} than for R_{d2} which is reflected experimentally by a smaller mismatch between the actual and desired traces for R_{d1} than for R_{d2} .

To evaluate sensitivity of the feedback design to hardware variations, we removed part

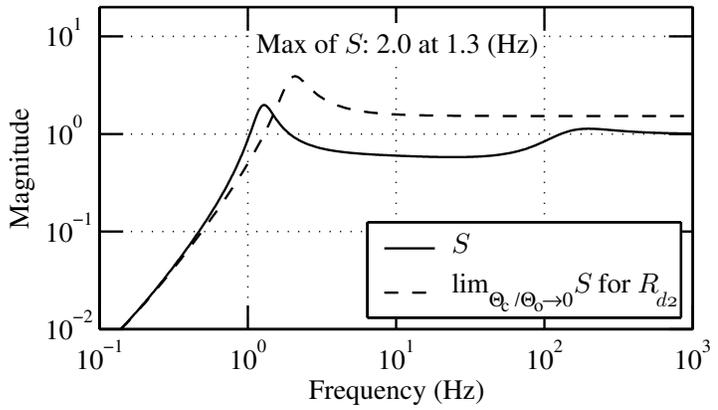
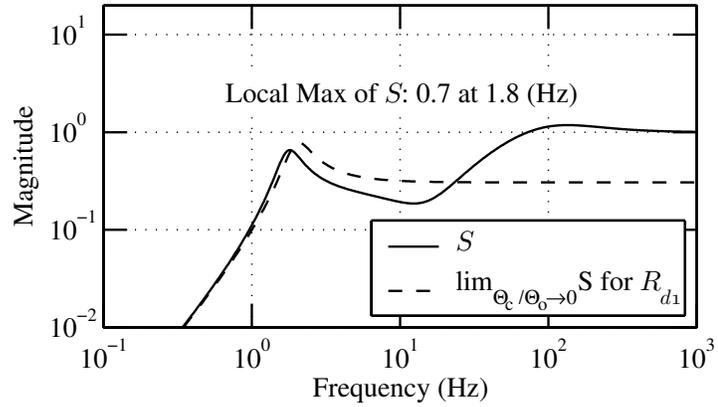


Figure 4.5 Bode sensitivity function S for the spring-damper virtual coupler rendering R_{d1} (top) and R_{d2} (bottom). The dashed traces indicate S in the limit as Θ_c/Θ_o approaches 0. Below 0.5 Hz, where distortion is small, S is close to the predicted limit.

EXPERIMENTAL RESULTS

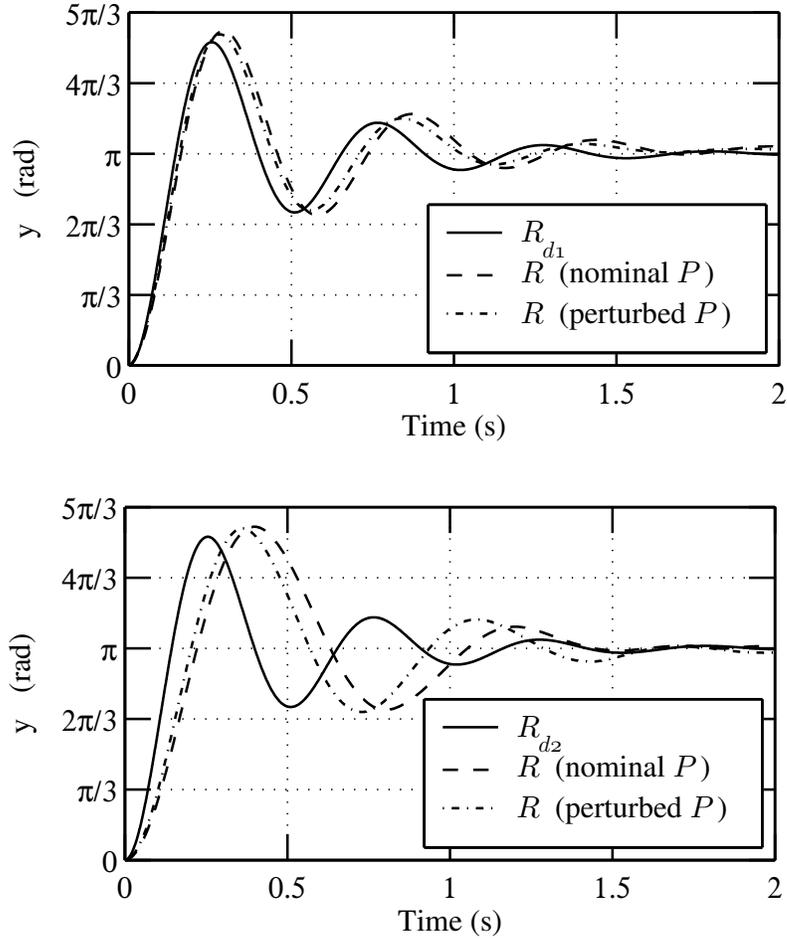


Figure 4.6 Step responses of virtual environments R_{d1} (top) and R_{d2} (bottom) rendered by the spring-damper virtual coupler. The solid traces indicate the desired response. Dashed traces were obtained using the nominal haptic device, whose dynamics are given by (4.16). Dash-dot traces were obtained on a modified device with 30% less rotational inertia.

of the handwheel from our device. System identification showed that this modification reduces the rotational inertia by approximately 30%. The dash-dot traces in Fig. 4.6 are the step responses obtained using the modified device. Sensitivity to hardware variations is evident in the mismatch between “ R (nominal P)” and “ R (perturbed P)”, and as predicted by the larger peak value of S for R_{d2} , the rendered response of R_{d2} exhibits greater sensitivity to hardware variations than the rendered response of R_{d1} .

4.4.2 Cancellation Coupler Design

Parameters of the cancellation coupler H_{11} , H_{12} , and H_{21} are designed according to the design directives given in Section 4. We begin by designing V_{11} . To attenuate H_{11} , we increase the magnitude of V_{11} ; however, as with the spring-damper virtual coupler, the gain of V_{11} is limited by high-frequency modes of the haptic device. Subject to the same hardware limitations, we select the same lead filter V_{11} for the cancellation coupler that we used in the spring-damper virtual coupler. The resulting frequency response for H_{11} is shown in Fig. 4.7. For H_{12} and H_{21} , we select the filter

$$\frac{8.085 \times 10^5 s + 1.27 \times 10^8}{s^3 + 1608s^2 + 8.085 \times 10^5 s + 1.27 \times 10^8}. \quad (4.23)$$

To improve low-frequency performance, we have selected the coefficients of the numerator such that $1 - H_{12}H_{21}$ has two zeros at the origin. Frequency responses of H_{12} , H_{21} , and $1 - H_{12}H_{21}$ are shown in Fig. 4.7.

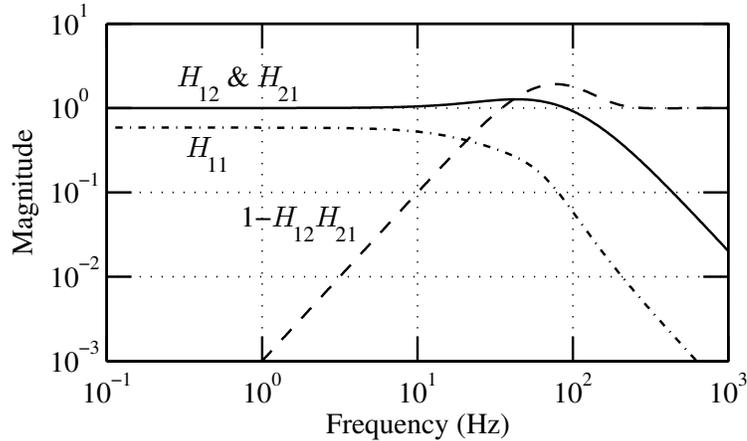


Figure 4.7 Cancellation coupler design parameters H_{11} , H_{12} and H_{21} .

The predicted closed-loop distortion for the cancellation coupler design is shown in Fig. 4.8. Comparing closed-loop distortion of the cancellation coupler with the virtual

coupler, we see that low frequency performance is similar, but the cancellation coupler maintains low distortion over a wider bandwidth than the spring-damper virtual coupler. As a consequence, the Bode sensitivity function for the cancellation coupler design, shown in Fig. 4.9 by the solid trace, approximates the limit R_d/P over a wider bandwidth than the Bode sensitivity function for the spring-damper virtual coupler design. Thus, according to the tradeoff, the peaks in S near 2 Hz for the cancellation coupler cannot be attenuated without compromising performance.

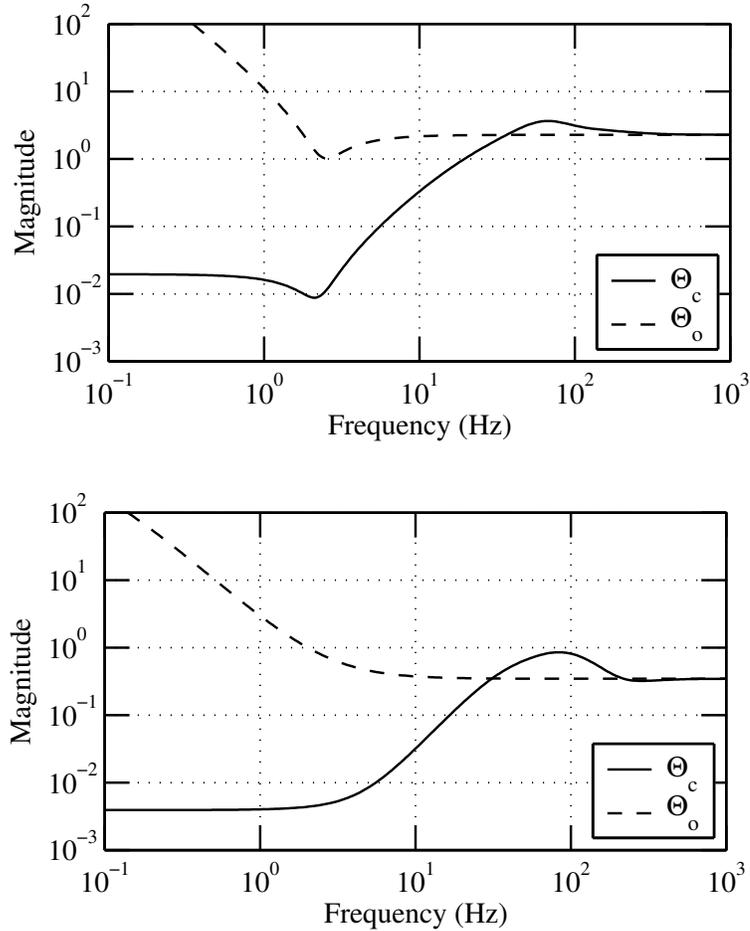


Figure 4.8 Open-loop and closed-loop distortion for the cancellation coupler rendering R_{d1} (top) and R_{d2} (bottom).

In the frequency range from 30 Hz to 150 Hz, the cancellation coupler amplifies rather than attenuates distortion. Furthermore, at some frequencies in this range, sensitivity is also amplified. At these frequencies where both Θ_c/Θ_o and S are amplified, the cancellation coupler design is clearly suboptimal with respect to the algebraic tradeoff between performance and sensitivity. We caution, however, against more general conclusions about the optimality of either cancellation coupler or the spring-damper virtual coupler (which

also exhibits a suboptimal combination of performance and sensitivity at some frequencies.) Other tradeoffs may be imposed by integral relationships across frequency which are beyond the scope of the present discussion.

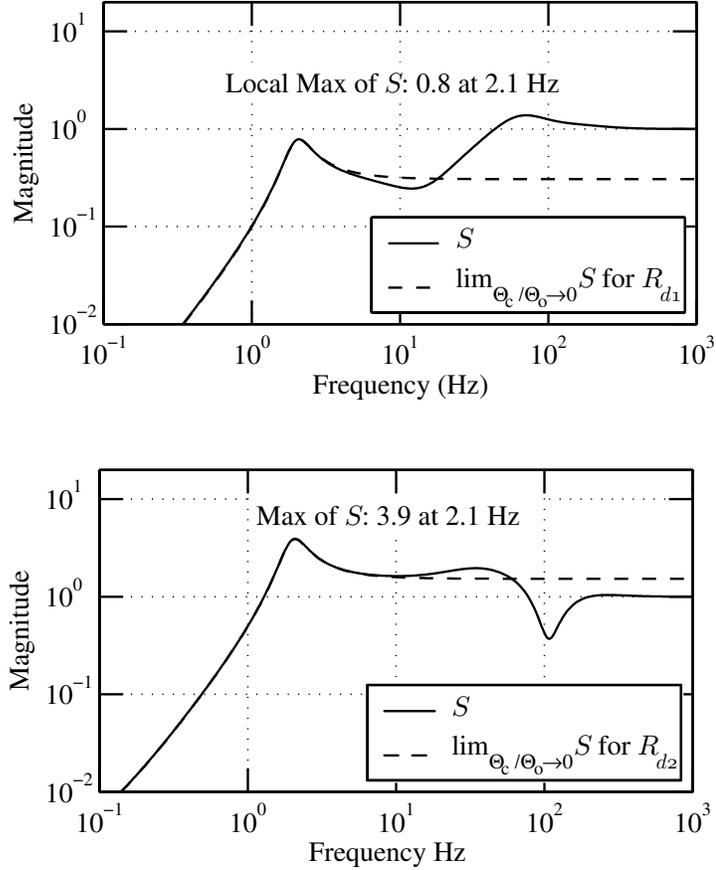


Figure 4.9 Bode sensitivity function S for the cancellation coupler rendering R_{d1} (top) and R_{d2} (bottom). The dashed traces indicate S in the limit as Θ_c/Θ_o approaches 0 for each virtual environment.

Experimental step responses of the rendered virtual environments are shown in Fig. 4.10 along with the desired step responses of R_{d1} and R_{d2} . The predicted performance of the cancellation coupler (Fig. 4.8) is superior to the virtual coupler (Fig. 4.4). This theoretical prediction is confirmed experimentally by the rendered step-responses of the cancellation coupler, given by the dashed traces in Fig. 4.10, which track the desired responses of R_{d1} and R_{d2} better than rendered step-responses of the spring-damper coupler, shown in Fig. 4.6.

As with the spring-damper virtual coupler design, we evaluate sensitivity of the rendered step-response to hardware variations by removing part of the handwheel from our device. This modification reduces the rotational inertia of the device by approximately

EXPERIMENTAL RESULTS

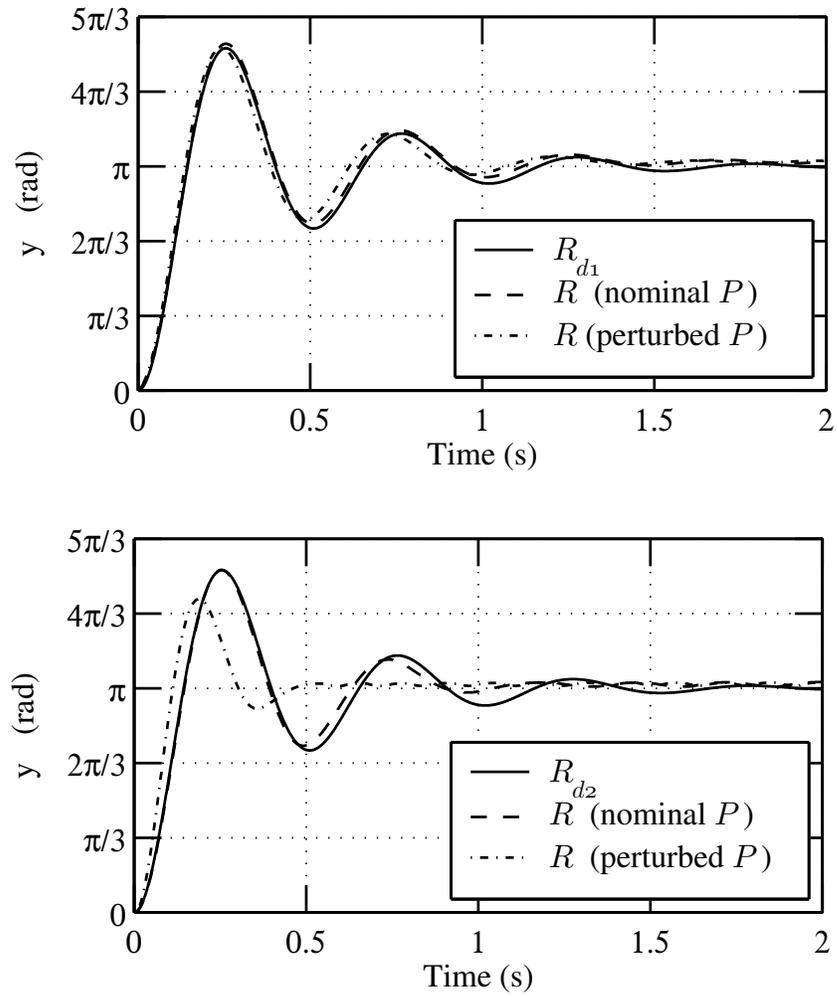


Figure 4.10 Step responses of virtual environments R_{d1} (top) and R_{d2} (bottom) rendered by the cancellation coupler. The solid traces indicate the desired response. Dashed traces were obtained using the nominal haptic device, whose dynamics are given by (4.16). Dash-dot traces were obtained on a modified device with 30% less rotational inertia.

30%. The dash-dot traces in Fig. 4.10 are the rendered step responses of the virtual environments using the modified hardware. While the step-responses of R_{d1} and R_{d2} rendered on the nominal device closely match the desired step-responses of R_{d1} and R_{d2} , the step-response of R_{d2} on the modified haptic device reveals a much greater sensitivity to the hardware dynamics than the step-response of R_{d1} . Thus a design with good predicted performance but large sensitivity may actually suffer from poor performance due to small variations in the haptic device dynamics.

4.5 Discussion

As a practical matter, the set of virtual environments that can be rendered well with a particular haptic device is limited. As the magnitude of the virtual environment dynamics R_d increase relative to the haptic device dynamics P , error in the model of the haptic device P must be reduced to maintain low distortion in the face of high sensitivity. However, the accuracy of available models for a haptic device is typically limited and the hardware dynamics are themselves subject to some variation over time.

While we have not yet addressed coupled stability of the human operator with the haptic interface system, the tradeoff between performance and sensitivity holds important implications for passivity. To guarantee stable interaction with the human operator, we may design sR to be passive, which implies that the Nyquist plot of sR lies in the closed right-half plane. (Note that R is multiplied by s to obtain a transfer function between force and velocity.) As given by (2.14), the Bode sensitivity function describes the sensitivity of the rendered virtual environment R to parameter variations in the haptic device and model uncertainty. Thus robustness of passivity to hardware variations and uncertainty is also determined by S .

The design example demonstrates that the cancellation coupler design provides superior performance compared with the spring-damper virtual coupler at frequencies where the virtual environment dynamics require cancellation of the haptic device dynamics. However, one may reasonably choose not to use the cancellation coupler to improve sensitivity or guarantee other properties of the closed-loop. The cancellation coupler, for instance, does not necessarily generate a passive rendered response to the human operator even when the virtual environment is passive. The limitations passivity would impose on the cancellation coupler is an open research question.

A common challenge in haptic rendering is the **virtual wall** environment [Colgate et al., 1995; Mehling et al., 2005; Diolaiti et al., 2006; Abbott and Okamura, 2005]. Although

our theoretical framework does not address the switching inherent to the virtual wall problem, we can comment on operation within the virtual wall where $R_d \rightarrow 0$ and outside the virtual wall where $|R_d| \rightarrow \infty$. Operating inside the virtual wall, the virtual environment is very stiff and the algebraic tradeoff between performance and sensitivity is small. The free-space motion outside the virtual wall, however, involves a tradeoff between performance and sensitivity. The sensitivity required to over-come the haptic device dynamics to render the desired free-space dynamics may be large in certain frequency ranges. A well-known relationship predicts the minimum physical damping of the haptic device to render a passive virtual wall using a sampled proportional-derivative control [Colgate et al., 1995]. It is suggested that physical damping added to maintain passivity be actively cancelled to some extent to recover free-space performance. Such a strategy, however, must be considered in light of the algebraic tradeoff which predicts that the cost of compensating for hardware dynamics is amplification of sensitivity. Furthermore, as we have noted, this sort of feedback compensation may exhibit much greater sensitivity than a feedforward control scheme.

The algebraic tradeoff between performance and sensitivity to hardware dynamics assumes a control architecture with position feedback. Additional sensors can mitigate the tradeoffs inherent to the position feedback architecture. A logical choice is the addition of a force sensor measuring the human operator's force f on the haptic device. With this measurement we may compute the desired position y_d by $R_d f$. Then typical high-gain control techniques, not subject to the algebraic tradeoff between performance and sensitivity, could be applied to make the haptic device position y track y_d .

Chapter 5

Force-Reflecting Teleoperation under Position Feedback

5.1 Introduction

The relationship between open and closed-loop properties of feedback systems is well understood for typical scalar servo-control systems, and rules of thumb provide guidance in shaping the loop-gain transfer function to obtain desired closed-loop tracking performance and stability margins. The situation is more complicated for teleoperator feedback systems since there are multiple feedback paths and, further, the human operator in the feedback loop presents special performance and stability requirements. Rules of thumb for teleoperator feedback design, were they available, would allow the designer to relate the open-loop responses to key characteristics of the closed-loop behavior such as transparency, tracking, and passivity.

Insight into the effects of certain design parameters can be a valuable tool for tuning a feedback controller. Many feedback laws for teleoperation are based in part on proportional-derivative (PD) feedback of the error between the master and slave positions [Alvarez-Gallegos et al., 1997; Anderson and Spong, 1989; Lawrence, 1993; Yokokohji and Yoshikawa, 1994; Lee and Spong, 2006]. Intuitively, a PD controller behaves like a tunable, parallel spring-damper connecting the master and slave. As described in [Leung et al., 1995; Yan and Salcudean, 1996], optimal synthesis tools can be applied to the teleoperator feedback design problem to address multiple design objectives. These tools exploit a more general controller structure which may afford better performance than the more restrictive structure of a PD controller; however an intuitive understanding of the more complicated controller is not immediately apparent.

The dependence of transparency and tracking performance on inherently variable en-

environment and human operator dynamics presents a special challenge in the analysis of teleoperator feedback design. A full assessment of performance should consider tracking and transparency for a class of environment and human operator dynamics. One approach to simplify the problem is to analyze the teleoperator behavior for a select set of environment dynamics. Typical choices include the zero impedance environment, where the slave may move freely, and the zero admittance environment, where the slave is held immobile [Hashtrudi-Zaad and Salcudean, 2001]. Further analysis, however, is required to generalize from particular environment dynamics to a class of environments.

We reveal design relationships between the controller and measures of performance by introducing a parameterization of all linear teleoperator feedback controllers with position sensing at the master and slave. This parameterization allows us to interpret any controller as a combination of local hardware compensation and a virtual coupler that provides position tracking between the master and slave. Two parameters characterize the compensated master and slave dynamics, and a further two parameters describe the gain and frequency dependent scaling of the virtual coupler. (Note that while there exist parallels between the teleoperator virtual coupler and the haptic interface virtual coupler discussed in Chapter 4, they are distinct concepts which require separate design and analysis.) We are able to find simplified relationships between controller design, transparency, and tracking by expressing measures of transparency and tracking performance in terms of the teleoperator design parameters. Further analysis allows us to prove bounds on transparency for a class of environment dynamics. To obtain this result, we measure transparency by **distortion** [Griffiths et al., 2008b] rather than the standard transparency ratio [Lawrence, 1993; Fite et al., 2001; Hashtrudi-Zaad and Salcudean, 2001]. The principle advantage of this choice in our analysis is that the magnitude of distortion is defined by the difference of transfer functions, and by analyzing expressions describing the magnitude of distortion, we may apply the triangle inequality to upper bound the magnitude of relative error.

Special coupled stability issues arise in teleoperation when the human operator and environment dynamics are in contact with the master and slave hardware. For passive human operator and environment dynamics, coupled stability is assured if the teleoperator is a passive two-port network, or less conservatively, if the teleoperator two-port network satisfies Llewellyn's criteria [Llewellyn, 1952; Haykin, 1970]. We do not address the problem of designing to satisfy these requirements as techniques exist to assure these conditions. However, if one has found a passive feedback design, our analysis provides insight into the controller's structure, proves bounds on performance, and suggests directives for the hardware and controller to improve performance.

The teleoperator feedback structure we analyze assumes that the master and slave are

ideal impedance-type devices equipped with position sensors. Force sensing at the slave is required in certain large robotic applications where the slave is better characterized as an admittance-type device. (Hydraulic actuators or lead screw mechanisms used to achieve power requirements may prevent the slave from being driven by the environment.) However, in practical applications such as automotive steer-by-wire, the master and slave hardware may be reasonably modeled as an impedance-type devices for which position feedback alone is sufficient [Odenthal et al., 2002; Yao, 2006]. Furthermore, price and physical durability of force sensors are significant drawbacks in many applications while the benefits of combined force and position sensing are unclear. A prerequisite for an analysis of the costs and benefits of adding force sensing is a thorough, quantitative analysis of limitations and design tradeoffs in teleoperator feedback design with only position sensing. The work presented here establishes an appropriate framework for this analysis and begins, but does not close, the analysis of teleoperation under position feedback.

5.2 Position-Position Teleoperation

The block diagram in Fig. 5.1 captures the general feedback configuration of a teleoperator with position sensing at the master and slave. The position of the master device x_m is the response of the master dynamics P_m to the sum of the motor force u_m and human operator force f_h :

$$x_m = P_m(u_m + f_h). \quad (5.1)$$

Note that all signals and transfer functions are functions of the Laplace variable s . Symmetric with the master model, the position of the slave manipulator x_s is the response of the slave dynamics P_s to the sum of the motor force u_s and environment force f_e :

$$x_s = P_s(u_s + f_e). \quad (5.2)$$

The master and slave models (5.1) and (5.2) assume that the master motor affects the master device through the same dynamics as the human operator and the slave motor affects the slave manipulator through the same dynamics as the environment. We assume without loss of generality that the master and slave models are normalized such that perfect tracking is achieved when $x_m = x_s$.

We model the linear dependence of the human operator force f_h on the master position x_m by P_h^{-1} and the linear dependence of the environment force f_e on the slave position x_s by

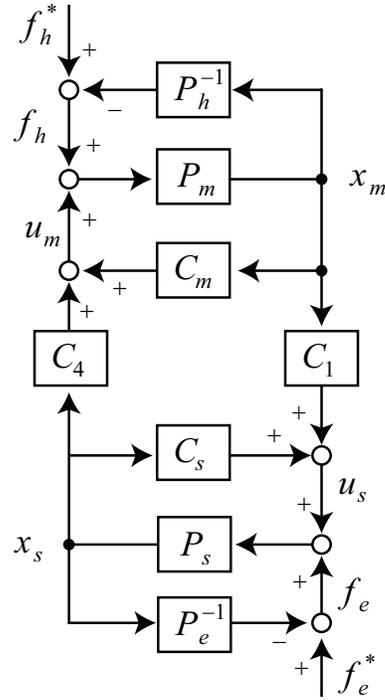


Figure 5.1 Block diagram of position-position teleoperation.

P_e^{-1} . We express these models in inverse form for consistency of notation (P_m , P_s , P_h , and P_e are all in forward-dynamics form, expressing transfer functions from force to motion). Then the human operator applies

$$f_h = -P_h^{-1}x_m + f_h^*, \quad (5.3)$$

where volitional input of the human operator enters through the exogenous input f_h^* . Similarly, the environment applies

$$f_e = -P_e^{-1}x_s + f_e^*, \quad (5.4)$$

where f_e^* is the exogenous input from the environment.

A multivariable controller computes motor commands u_m for the master and u_s for the slave based on the sensed master position x_m and sensed slave position x_s . As shown in Fig. 5.1, the controller elements C_m , C_s , C_1 , and C_4 define the four input/output relation-

ships:

$$\begin{bmatrix} u_m \\ u_s \end{bmatrix} = \begin{bmatrix} C_m & C_4 \\ C_1 & C_s \end{bmatrix} \begin{bmatrix} x_m \\ x_s \end{bmatrix}. \quad (5.5)$$

Controller elements C_1 and C_4 , called **channel controllers** [Lawrence, 1993], provide feedback between the master device and slave manipulator, while controller elements C_m and C_s operate locally. Note that our notation is consistent with the standard four-channel teleoperator framework introduced in [Lawrence, 1993] except that these four controller terms assume position rather than velocity feedback.

5.3 Models for Design and Performance Analysis

Teleoperator feedback design must meet two performance goals: transparency, which requires that the dynamics of the environment be accurately presented to the human operator through the master, and tracking, which requires that the error between the master and slave positions remain small. For many feedback systems, multiple performance goals can be captured with multiple performance variable outputs in a single model of the feedback system. However, the goals of transparency and tracking require two models, one model that includes only the environment dynamics, and one model that includes both the human operator and the environment dynamics.

Since the human operator and environment dynamics are variable, it is useful to characterize the teleoperator behavior by just the closed-loop response of the master and slave in feedback with the controller. The dynamics of the teleoperator without the human operator or environment dynamics are termed the **free-free** dynamics. Through a loop-transformation, we find a parameterization of the free-free dynamics that admits an intuitive interpretation. Models of the teleoperator including the human operator and environment dynamics may then be expressed as functions of these parameters.

5.3.1 Parameterizing the Free-Free Dynamics

Let us begin by introducing the free-free dynamics, defined by the feedback interconnection of the master, slave, and controller elements shown in Fig 5.2(a). In this configuration, the master and the slave are not in contact with the human operator or the environment dynamics. The controller C shapes the free-free dynamics; however the role of each controller element is not apparent. The confounding factor is that local feedback terms C_m

and C_s serve dual roles: they both compensate for the hardware dynamics and contribute to tracking between the master and slave.

To isolate the functions of the controller, we seek a decomposition of the controller into a part that provides hardware compensation and a part that only affects tracking. Since compensation of the hardware dynamics is a local, that portion of the controller should be diagonal (responding exclusively to x_m or x_s). The portion of the controller responsible for tracking should respond to a differences between x_m and x_s . In particular we consider the tracking portion of the controller to be the portion of u_m and u_s that responds to the generalized tracking error $x_m - Ax_s$, where the transfer function A is finite and not identically zero. Suppose that each element of C is finite except at isolated points s in the complex plane and that the term $C_1(s)C_4(s)$ is not identically zero. The controller C may then be decomposed into

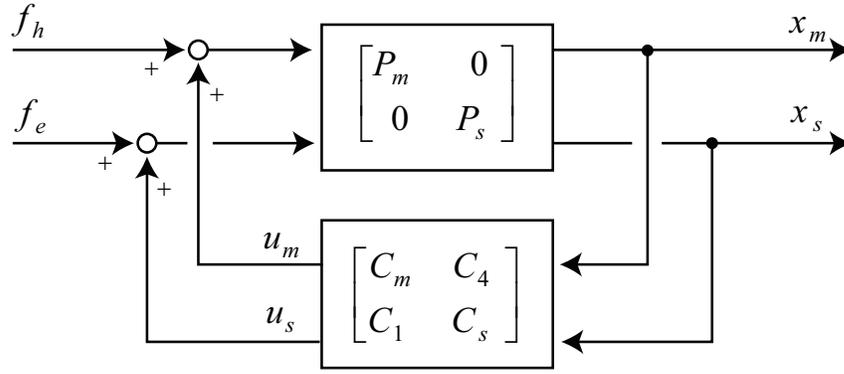
$$\begin{bmatrix} C_m & C_4 \\ C_1 & C_s \end{bmatrix} = \begin{bmatrix} -C'_m & 0 \\ 0 & -C'_s \end{bmatrix} + B \begin{bmatrix} -1 & A \\ 1/A & -1 \end{bmatrix}. \quad (5.6)$$

The diagonal term acts to shape the master and slave dynamics. The remaining term, which we refer to as a virtual coupler, describes the control action proportional to the tracking error $x_m - Ax_s$. The additional transfer function B describes the gain of virtual coupler—in other words, the gain on the generalized tracking error. An important observation is that the compensation of the hardware dynamics is not given by C_m or C_s alone; instead it is a function of all four controller elements.

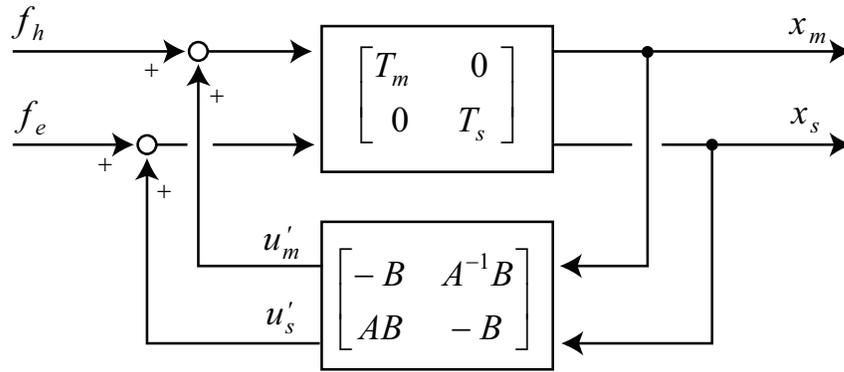
The virtual coupler defined by transfer functions A and B is distinct from the haptic interface virtual coupler introduced in Chapter 4. Whereas the haptic interface virtual coupler given by (4.2) includes four independent elements and incorporates compensation of the haptic device dynamics, the virtual coupler as we define it for teleoperation has only two independent elements A and B and does not incorporate hardware compensation.

The terms C'_m and C'_s provide local feedback that shapes the closed-loop dynamics of the master and slave. These terms may be used to compensate for undesirable hardware dynamics such as excessive damping or inertia, or can be used to introduce behavior such as stiffness which may not be inherent to the master or slave dynamics. We denote the shaped master dynamics by

$$T_m \triangleq \frac{P_m}{1 + P_m C'_m} \quad (5.7)$$



(a) Free-free dynamics



(b) Loop-transformation of the free-free dynamics

Figure 5.2 Free-free dynamics in terms of the controller elements and open-loop master and slave (top) and a reparameterization in terms of design parameters T_m , T_s , A , and B (bottom).

and the shaped slave dynamics by

$$T_s \triangleq \frac{P_s}{1 + P_s C'_s}. \quad (5.8)$$

Having incorporated a portion of the local feedback into the master and slave dynamics, we define new control inputs u'_m and u'_s which act on the shaped master and slave dynamics. These new control inputs are given by

$$\begin{bmatrix} u'_m \\ u'_s \end{bmatrix} = \begin{bmatrix} -B & AB \\ A^{-1}B & -B \end{bmatrix} \begin{bmatrix} x_m \\ x_s \end{bmatrix}. \quad (5.9)$$

The transformed feedback system, in terms of the parameters T_m , T_s , A , and B , is shown in Fig. 5.2(b). Note that the feedback remaining in u'_m and u'_s after removing the contribution due to $C'_m x_m$ and $C'_s x_s$ is proportional to the generalized tracking error $x_m - Ax_s$.

The parameters A and B describe a virtual coupler between the shaped master dynam-

ics T_m and the shaped slave dynamics T_s . When $A = 1$ at all frequencies, the network of transfer functions (5.9) can describe, for example, the response of a parallel spring-damper. While B may be selected such that (5.9) mimics the response of a mechanical element such as a spring, we do not impose any particular structure on B . For our analysis, the key characteristic of B is its magnitude frequency response which determines the feedback gain on the error $x_m - Ax_s$.

The transfer function A accounts for asymmetry of the teleoperator. For $A \equiv 1$ the virtual coupler is a fully symmetric; however as A approaches zero, the master position x_m becomes a reference signal for the slave to track, and as A approaches infinity, the slave position x_s becomes a reference signal to be tracked by the master. In other words, when the magnitude of A is much less than one, the controller provides assist at the master, and when the magnitude of A is much greater than one, the controller provides assist at the slave. In our subsequent analysis of transparency and tracking, we will analyze the fully symmetric teleoperator where A is unity across all frequencies.

5.3.2 Measuring Transparency Performance

To evaluate transparency, we place the environment dynamics in contact with the free-free dynamics and evaluate the response of the master position x_m to f_h and f_e^* . In terms of Fig. 5.2, contact between the environment and slave closes a feedback path from x_s to f_e according to (5.4). As shown in Fig. 5.3, the virtual coupler described by A and B then acts on the shaped master dynamics T_m and the combined shaped slave and environment dynamics $T_s P_e / (T_s + P_e)$. We denote these combined dynamics by $T_{se} \triangleq T_s P_e / (T_s + P_e)$. The response of the master and slave positions to the force inputs f_h and f_e^* is then

$$\begin{bmatrix} x_m \\ x_s \end{bmatrix} = \frac{1}{1 + B(T_m + T_{se})} \begin{bmatrix} T_m + T_m T_{se} B & T_m T_{se} A^{-1} B \\ T_m T_{se} A B & T_{se} + T_m T_{se} B \end{bmatrix} \begin{bmatrix} f_h \\ f_e^* \end{bmatrix}. \quad (5.10)$$

For perfect transparency, the response of the master position x_m to both the human operator force f_h and the external environment force f_e^* should be given by P_e . Thus the desired response for x_m , which we denote by $x_m^{(d)}$, is given by $P_e(f_h + f_e^*)$.

We measure transparency by the relative error between the actual and desired closed-loop dynamics. In haptic rendering, this error is captured by **distortion** Θ [Griffiths et al., 2008b]. The multi-input/multi-output nature of teleoperation requires us to extend this definition. Let us denote the closed-loop response from f_h to x_m in Fig. 5.3 by x_m/f_h , and the closed-loop response from f_e^* to x_m by x_m/f_e^* . Distortion then consists of two components

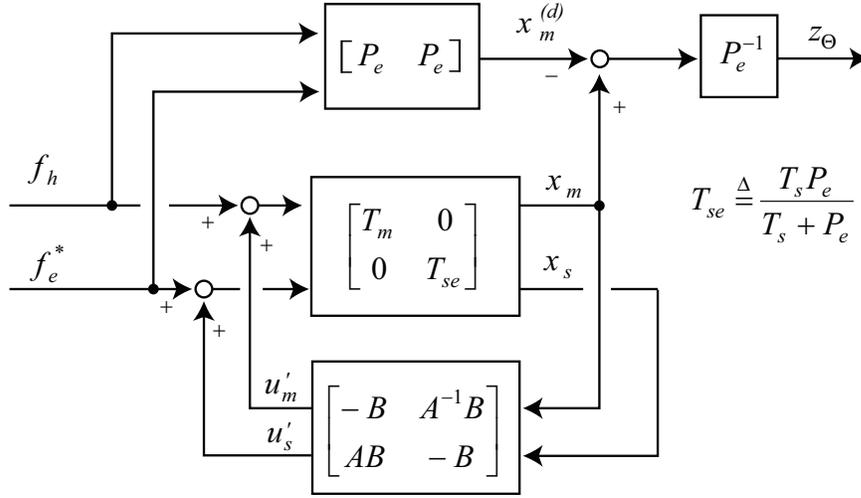


Figure 5.3 Feedback configuration for transparency. The coupled dynamics of the shaped slave T_s and environment dynamics P_e are captured by T_{se} . The closed-loop dynamics of the transparency performance variable z_Θ to the exogenous inputs f_h and f_e^* are denoted by Θ_h and Θ_e .

given by

$$\Theta_h \triangleq \frac{x_m/f_h - P_e}{P_e} \quad (5.11)$$

$$\Theta_e \triangleq \frac{x_m/f_e^* - P_e}{P_e}. \quad (5.12)$$

Based on the difference between the actual response x_m and the desired response $x_m^{(d)}$, we may define a performance variable $z_\Theta \triangleq \frac{1}{P_e}(x_m - x_m^{(d)})$. (Note that $x_m = (x_m/f_h)f_h + (x_m/f_e^*)f_e^*$.) Thus the response of z_Θ in Fig. 5.3 is given by

$$z_\Theta = \Theta_h f_h + \Theta_e f_e^*. \quad (5.13)$$

Attenuating the response of z_Θ is achieved by reducing the magnitude of Θ_h and Θ_e evaluated along the $j\omega$ -axis. Through the construction of the performance variable z_Θ , we have expressed the problem of achieving transparency in the standard form of a disturbance attenuation problem, where f_h and f_e^* are the exogenous (or disturbance) inputs.

5.3.3 Measuring Tracking Performance

To address tracking performance, we include both the human operator dynamics P_h and environment dynamics P_e in our system model. Now in Fig. 5.2, the human operator closes

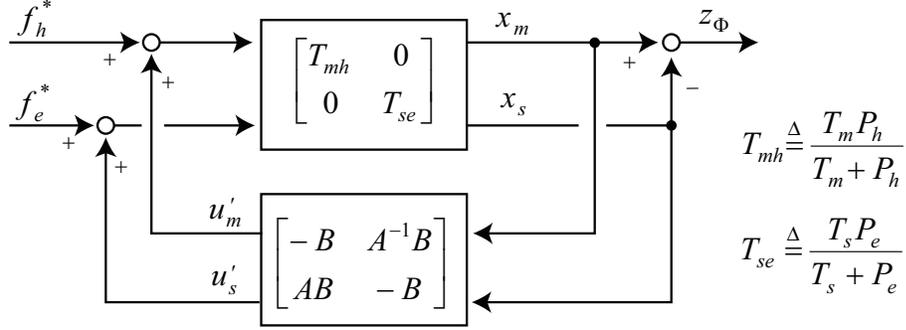


Figure 5.4 Feedback configuration for tracking. The performance variable z_Φ measures the error between the master and slave positions. The coupled dynamics of the shaped master T_m and human dynamics P_h are denoted by T_{mh} , and similarly, the coupled dynamics of the shaped slave T_s and environment dynamics P_e are denoted by T_{se} .

a feedback path from x_m to f_h according to (5.3) and the environment closes a feedback path from x_s to f_e according to (5.4). The block diagram in Fig. 5.4 shows the combined dynamics of the master and human $T_{mh} \triangleq T_m P_h / (T_m + P_h)$, and the combined dynamics of the slave and environment $T_{se} \triangleq P_s P_e / (P_s + P_e)$ in feedback with the virtual coupler. The response of the master and slave positions to the force inputs f_h^* and f_e^* is then

$$\begin{bmatrix} x_m \\ x_s \end{bmatrix} = \frac{1}{1 + B(T_{mh} + T_{se})} \begin{bmatrix} T_{mh} + T_{mh}T_{se}B & T_{mh}T_{se}A^{-1}B \\ T_{mh}T_{se}AB & T_{se} + T_{mh}T_{se}B \end{bmatrix} \begin{bmatrix} f_h^* \\ f_e^* \end{bmatrix}. \quad (5.14)$$

To measure tracking, we form the performance variable $z_\Phi \triangleq x_m - x_s$. Let $\Phi_h \triangleq z_\Phi / f_h^*$ and $\Phi_e \triangleq z_\Phi / f_e^*$ such that

$$z_\Phi = \Phi_h f_h^* + \Phi_e f_e^*. \quad (5.15)$$

For position tracking in the face of the exogenous inputs f_h^* and f_e^* , the feedback design must attenuate the frequency response of Φ_h and Φ_e .

5.4 Analysis of Transparency and Tracking Performance

A primary goal of feedback design in teleoperation is to attenuate the frequency response of Θ_h , Θ_e , Φ_h , and Φ_e which describe transparency and tracking performance. These dynamics depend on the environment and human operator dynamics, and the challenge of feedback design is to find a controller that makes these terms sufficiently small for a suitable class of human operator and environment dynamics. We compute these performance

dynamics as functions of the design parameters T_m , T_s , A , and B , and subsequently interpret these expressions to obtain useful rules for design.

5.4.1 Transparency Performance

Substituting expressions for x_m/f_h and x_m/f_e^* from (5.10) into the definitions for Θ_h and Θ_e , we find that

$$\Theta_h = \left(\frac{T_m + T_m T_{se} B}{1 + B(T_m + T_{se})} - P_e \right) \frac{1}{P_e} \quad (5.16)$$

$$\Theta_e = \left(\frac{T_m T_{se} A^{-1} B}{1 + B(T_m + T_{se})} - P_e \right) \frac{1}{P_e}. \quad (5.17)$$

To improve transparency, we must design the controller and its associated parameters T_m , T_s , A , and B such that Θ_h and Θ_e are attenuated along the $j\omega$ -axis. Furthermore, as much as feasible, the design should attenuate Θ_h and Θ_e irrespective of the particular environment dynamics P_e .

To highlight important relationships between the design parameters and the environment dynamics we consider two limiting cases. In the first case, the gain of the virtual coupler B approaches infinity while T_m and T_s are finite, and in the second case, B is finite but the magnitude of T_m and T_s approach infinity. We do not explore the role of the frequency dependent motion scaling A in the present discussion and simply let $A = 1$. From (5.16) and (5.17), as $|B(j\omega)| \rightarrow \infty$

$$\Theta_h(j\omega) \rightarrow \Theta^*(j\omega) \quad (5.18)$$

$$\Theta_e(j\omega) \rightarrow \Theta^*(j\omega) \quad (5.19)$$

where

$$\Theta^* \triangleq \frac{T_m T_{se}}{P_e(T_m + T_{se})} - 1. \quad (5.20)$$

It follows from the definitions of T_{mh} and T_{se} that

$$\Theta^* = \frac{-P_e}{\frac{T_m T_s}{T_m + T_s} + P_e}. \quad (5.21)$$

Note that the ratio of P_e to $T_m T_s / (T_m + T_s)$ plays a critical role in the expression for trans-

parency. We denote the rigidly coupled dynamics of the master and slave by

$$T_{ms} \triangleq \frac{T_m T_s}{T_m + T_s}. \quad (5.22)$$

Then, at a frequency ω , let

$$a \triangleq \left| \frac{P_e(j\omega)}{T_{ms}(j\omega)} \right|. \quad (5.23)$$

For T_{ms} fixed, we consider the class of environment dynamics for which a lies between $a_l > 0$ and $a_u < 1$. Let $E \triangleq \{P_e \mid a_l \leq a \leq a_u\}$.

Proposition 5.4.1. *For any $P_e \in E$*

$$\frac{a_l}{1 + a_l} \leq |\Theta^*(j\omega)| \leq \frac{a_u}{1 - a_u}. \quad (5.24)$$

Proof. Applying the triangle inequality to (5.21) yields

$$\frac{a}{1 + a} \leq |\Theta^*(j\omega)| \leq \frac{a}{|1 - a|}. \quad (5.25)$$

By hypothesis, $a_l \leq a \leq a_u < 1$ and the result follows. \square

Proposition 5.4.1 predicts bounds on distortion based on the magnitude of the environment dynamics. Note that the lower bound is non-zero despite letting the gain of the virtual coupler approach infinity. From the upper bound we see that $|P_e(j\omega)|$ must be small relative to $|T_{ms}(j\omega)|$ for good transparency.

To investigate the role of finite gain B , we let $|T_{ms}|$ approach infinity, which implies that both $|T_m|$ and $|T_s|$ approach infinity. We note that $T_{mh} \rightarrow P_h$ as $|T_m| \rightarrow \infty$ and that $T_{se} \rightarrow P_e$ as $|T_s| \rightarrow \infty$. From these facts together with (5.16) and (5.17), we find that as $|T_{ms}(j\omega)| \rightarrow \infty$

$$\Theta_h \rightarrow \frac{1}{BP_e} \quad (5.26)$$

$$\Theta_e \rightarrow 0. \quad (5.27)$$

We see that the $|P_e|$ must be significantly greater than $|1/B|$ for distortion Θ_h to be small.

Proposition 5.4.2. *For any P_e such that $|P_e(j\omega)| \geq m$*

$$\lim_{|T_{ms}(j\omega)| \rightarrow \infty} |\Theta_h(j\omega)| \leq \frac{1}{m|B(j\omega)|} \quad (5.28)$$

and for any P_e such that $|P_e(j\omega)| \leq m$

$$\lim_{|T_{ms}(j\omega)| \rightarrow \infty} |\Theta_h(j\omega)| \geq \frac{1}{m|B(j\omega)|}. \quad (5.29)$$

Proof. The result follows directly from (5.26). \square

The results of Proposition 5.4.1 and 5.4.2 provide rules of thumb for the design of the free-free parameters. Key quantities for transparency are the inverse of the virtual coupler gain $|B(j\omega)^{-1}|$ and the rigidly coupled dynamics of the shaped master and slave $|T_{ms}|$. Proposition 5.4.1 and 5.4.2 predict that increasing the magnitude of these quantities at a frequency improves distortion for a given environment and expands the set of environments for which particular bounds on distortion are satisfied. However, as we will discuss in Chapter 6, making T_{ms} large necessarily induces poor stability robustness and sensitivity.

5.4.2 Tracking Performance

The position tracking error dynamics Φ_h and Φ_e may be determined from (5.14) which describes the closed loop response from $[f_h^* \ f_e^*]^T$ to $[x_m \ x_s]^T$. Then Φ_h is given by the difference between the (1,1) and (2,1) elements, and Φ_e is given by the difference between the (1,2) and (2,2) elements:

$$\Phi_h = \frac{T_{mh} + T_{mh}T_{se}B(1-A)}{1 + B(T_{mh} + T_{se})} \quad (5.30)$$

$$\Phi_e = \frac{-T_{se} + T_{mh}T_{se}B(A^{-1} - 1)}{1 + B(T_{mh} + T_{se})}. \quad (5.31)$$

We let $A = 1$ which simplifies (5.30) and (5.31) considerably. Let us again consider the limiting case where $|B(j\omega)|$ approaches infinity and the case where $|T_m(j\omega)|$ and $|T_s(j\omega)|$ approach infinity. In the other limiting case, as $|B(j\omega)| \rightarrow \infty$

$$\Phi_h(j\omega) \rightarrow 0 \quad (5.32)$$

$$\Phi_e(j\omega) \rightarrow 0. \quad (5.33)$$

As is typical in servo-control design, high-gain feedback from B drives tracking error to zero. As $|T_m(j\omega)| \rightarrow \infty$ and $|T_s(j\omega)| \rightarrow \infty$

$$\Phi_h(j\omega) \rightarrow \frac{P_h(j\omega)}{1 + B(j\omega)(P_h(j\omega) + P_e(j\omega))} \quad (5.34)$$

$$\Phi_e(j\omega) \rightarrow \frac{-P_e(j\omega)}{1 + B(j\omega)(P_h(j\omega) + P_e(j\omega))}. \quad (5.35)$$

Note that the magnitude of tracking error depends on the interaction of P_h and P_e . As $B(j\omega)$ approaches $-1/(P_h(j\omega) + P_e(j\omega))$, the magnitudes of $\Phi_h(j\omega)$ and $\Phi_e(j\omega)$ approach infinity. Practically speaking this situation implies a resonance between the virtual coupler and the dynamics of the environment and human operator. Due to the interaction of the environment dynamics and the human operator, we do not have bounds on tracking performance analogous to those for transparency. However, as a general rule, one must increase $|B(j\omega)|$ to improve tracking. Further insight into the interaction of B with tracking performance can be found by computing Φ_h and Φ_e for sample environment and human operator dynamics.

5.5 Steer-by-wire Example

Our parameterization of position-position feedback designs allows us to separate the part of the feedback control that compensates for hardware dynamics from the part that forms a virtual coupler between the master and slave. In this simulation example, we demonstrate the utility of compensating for hardware dynamics by comparing two feedback designs with the same virtual coupler properties (i.e. the parameters A and B are the same) but with different shaped master and slave dynamics. We introduce a graphical tool we call a **transparency diagram** that is based on the results of Propositions 5.4.1 and 5.4.2. By overlaying the transparency diagrams of the two designs, we can interpret the effect of hardware compensation on transparency performance.

The models for the master, slave, and environment dynamics are motivated by the application of automotive steer-by-wire. As described in [Odenthal et al., 2002], the feedback design for automotive steer-by-wire may be treated as a teleoperation problem. The general configuration of steer-by-wire depicted in Fig. 5.5 includes a motorized steering-wheel for the master hardware and a motorized rack-and-pinion gear for the slave hardware. The environment dynamics include the vehicle's road-wheels and the interaction between the tire, the road, and the vehicle dynamics. Linear dynamics of the rack and environment are

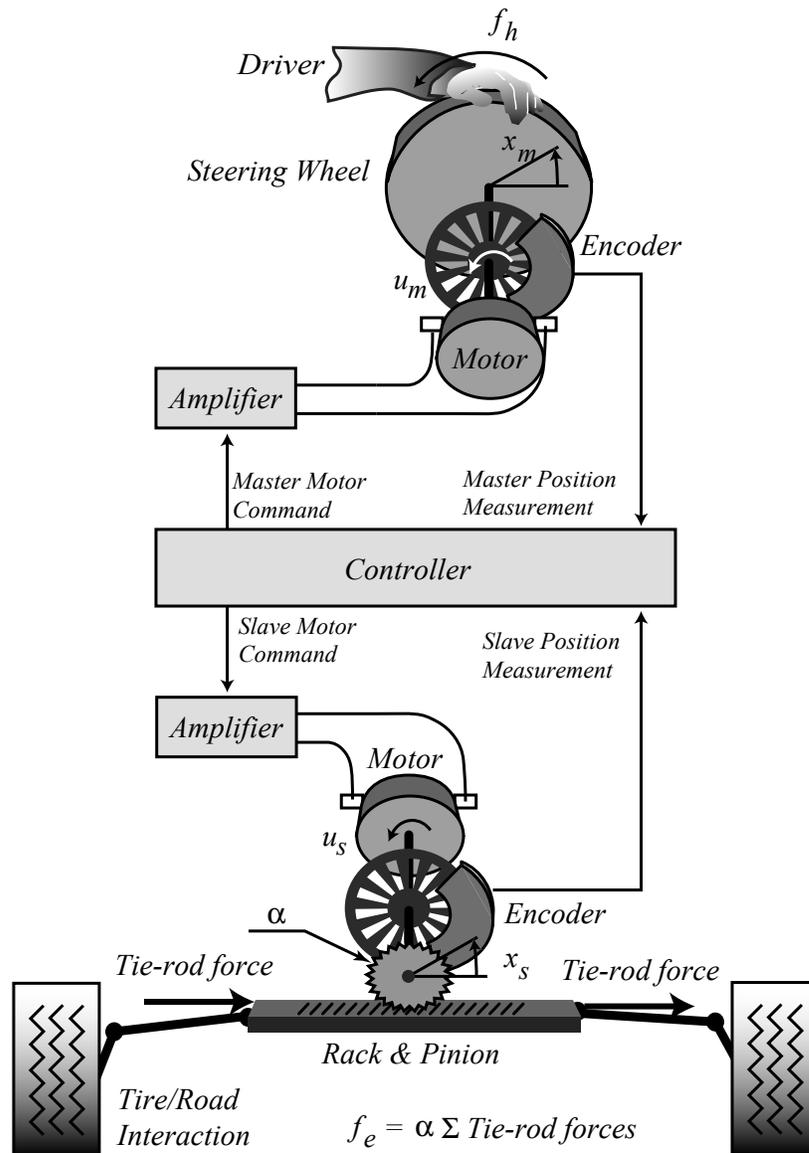


Figure 5.5 Configuration of a steer-by-wire system.

reflected back through the pinion radius α .

Our models for the steering-wheel subsystem, the rack-and-pinion subsystem, and the steering dynamics of the road wheels are representative of a medium-sized passenger vehicle, drawing upon models and parameter values from [Odenthal et al., 2002; Baxter, 1988; Dixon, 1991; Gillespie, 1992]. The steering-wheel and rack dynamics are modeled by simple damped-mass systems. The road-wheel steering dynamics includes rotational inertia and damping of the road wheel about its steering axis and first-order effects of caster, tire aligning moment, and tire sidewall deflection, assuming a constant vehicle speed of 20 m/s. Frequency responses of the master dynamics P_m and slave dynamics P_s are shown in Fig. 5.6, and the frequency response of the environment dynamics P_e is shown in Fig. 5.7.

5.5.1 Shaping the Master and Slave Dynamics & Designing a Virtual Coupler

Damping in the steering-wheel and rack dynamics may cause a sluggish response at low-frequencies, but with active control, we may reduce the effects of hardware damping. For the purpose of forming comparisons, our first design does not shape the master or slave dynamics. We denote the parameters and closed-loop properties of this design with ‘(a)’. Then as indicated in Fig. 5.6, the shaped dynamics for this design are given simply by the open-loop dynamics, $T_m^{(a)} = P_m$ and $T_s^{(a)} = P_s$. Our second design, denoted ‘(b)’, shapes the dynamics of the steering-wheel and rack to cancel a portion of the damping. We use feedback around the steering-wheel subsystem to reduce the steering-wheel damping from 0.6 N-m-s/rad to 0.3 N-m-s/rad, and we use feedback around the rack subsystem to reduce damping from 1.5 N-m-s/rad to 0.3 N-m-s/rad.

Shaping of the master and slave dynamics occurs through the terms C'_m and C'_s . To reduce the effective damping in T_m and T_s to the desired level, we let $C'_m = -0.3s$ and $C'_s = -1.2s$. The negative feedback gain of these derivative controllers mimics the response of negative dampers acting in parallel with the hardware damping. Computation of derivatives in the controller generally presents practical difficulties; however for this example, the controller elements can be made proper by introducing high-frequency poles, and since the open-loop inertia dominates the effects of friction above about 10 Hz, the additional poles placed above this frequency have little effect on the shaped master and slave dynamics T_m and T_s . Frequency responses of $T_m^{(b)}$ and $T_s^{(b)}$ are shown in Fig. 5.6. Partial cancellation of the hardware damping is evident at low frequencies where damping dominates inertial effects.

The virtual coupler is captured in the motion scaling parameter A and the gain B . For

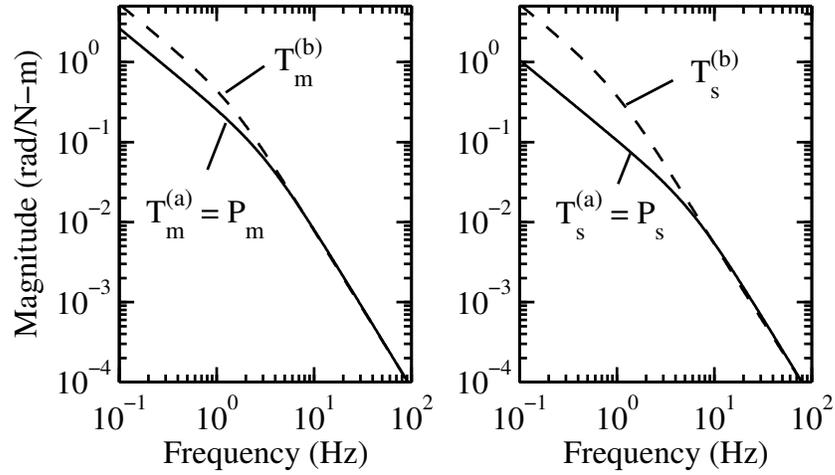


Figure 5.6 Steer-by-wire master and slave dynamics P_m and P_s . For design (a) the closed-loop dynamics of the master and the slave are the same as the open-loop. Design (b) cancels 50% of the master damping and 80% of the slave damping.

both designs we let $A = 1$ and choose $B = 150$ N-m/rad, a stiffness value motivated by conventional steering systems [Odenthal et al., 2002]. Then, for the first design which does not shape the hardware dynamics, the controller is simply

$$C^{(a)} = \begin{bmatrix} -150 & 150 \\ 150 & -150 \end{bmatrix}. \quad (5.36)$$

The second controller also includes the negative derivative elements of C'_m and C'_s . It follows from (5.6) that

$$C^{(b)} = \begin{bmatrix} 0.3s - 150 & 150 \\ 150 & 1.2s - 150 \end{bmatrix}. \quad (5.37)$$

5.5.2 Evaluating Transparency and Tracking

Propositions 5.4.1 and 5.4.2 permit predictions about the magnitude of distortion from the magnitude of P_e and the design parameters T_m , T_s , and B . We graphically compare the frequency response of T_{ms} and B^{-1} to the frequency response of the environment dynamics in what we call a **transparency diagram**. Note that the bounds on distortion provided by Propositions 5.4.1 and 5.4.2 are unity when $|P_e(j\omega)| \leq 1/2|T_{ms}(j\omega)|$ and $|P_e(j\omega)| \geq |B(j\omega)^{-1}|$. (We focus on Θ_h for which both bounds apply.) The transparency diagram in Figure 5.7 depicts the frequency-response of environment dynamics together

with $(1/2)T_{ms}^{(a)}$ and B^{-1} . Also shown for reference are exact contours of the worst-case performance for $|\Theta_h(j\omega)|$ given the magnitude of $P_e(j\omega)$. The worst-case performance is computed numerically by finding the phase of $P_e(j\omega)$ that maximizes $|\Theta_h|$ for a fixed magnitude of $P_e(j\omega)$. Figure 5.7 confirms that $(1/2)|T_{ms}(j\omega)|$ and $|B(j\omega)^{-1}|$ serve as good approximations to the unity distortion curve.

The transparency diagram in Fig. 5.8 contrasts the frequency response of $(1/2)T_{ms}$ for design (a) and (b). Several interesting design features can be determined by inspection of the diagram. Conventional wisdom suggests that improved performance can be achieved by increasing the gains of the virtual coupler (e.g. proportional and derivative gains for a PD controller.) While increasing $|B|$ will expand the region between $|T_{ms}|$ and $|B^{-1}|$, transparency for the particular road-wheel steering dynamics is not limited by the virtual coupler gain B but by the dynamics of the shaped master and slave, T_m and T_s . By compensating for the hardware damping, design (b) increases the frequency range of good transparency in comparison to design (a). The transparency diagram also predicts that further compensation of hardware damping would not significantly improve transparency. With the effective damping in design (b), inertial effects dominate above 1 Hz and so improved transparency can only be achieved by compensating for master and slave inertia.

We can verify the predictions of the transparency diagram by computing the terms of distortion Θ_h and Θ_e . Frequency responses of these terms are shown in Fig. 5.9. As predicted, design (b) achieves low distortion over a wider bandwidth than design (a). The transparency diagram is based on worst-case performance, and we see that for the particular environment dynamics actual performance is considerably better. The transparency diagram predicts a worst-case performance of unity distortion at 0.3 and 0.9 Hz for designs (a) and (b), whereas $|\Theta_h^{(a)}|$ and $|\Theta_e^{(a)}|$ are approximately 0.4 at 0.3 Hz and $|\Theta_h^{(b)}|$ and $|\Theta_e^{(b)}|$ are approximately 0.5 at 0.9 Hz. The peaks and valleys in Θ_h and Θ_e around 11 Hz are coincident with a resonant frequency between the master and slave sprung by the virtual coupler.

Tracking error depends on both the road-wheel dynamics and the particular dynamics imposed by the driver on the steering-wheel. For the purpose of simulation, we assume a very simple model for the driver: $P_h = P_m$. In effect, the driver doubles the mass and damping of the steering-wheel. While no linear model will accurately capture the complex behavior of the operator, our simple model assumes a relatively large effect of the driver on the master dynamics and produces tracking error that differs little from that obtained assuming no driver model. The frequency responses of the tracking error Φ_h and Φ_e are shown in Fig. 5.9. At DC, the response of Φ_f is 1/150 rad/N-m, which is simply the DC stiffness of the virtual coupler. There is zero DC tracking error in response to exogenous

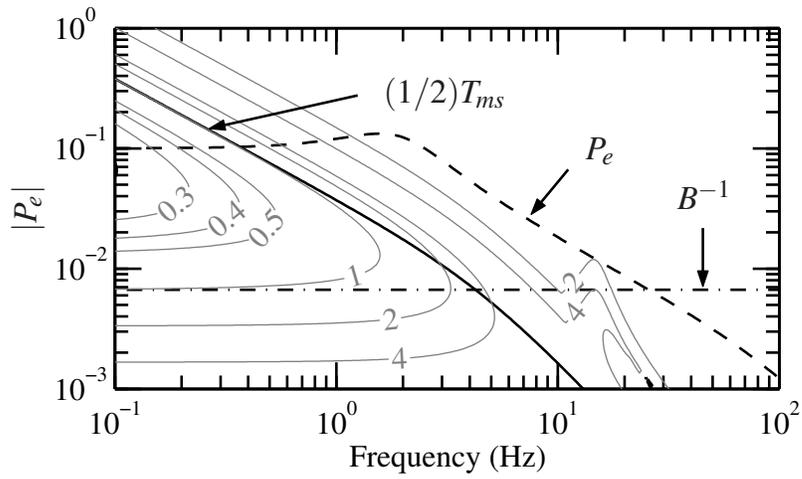


Figure 5.7 Transparency diagram for design (a) with overlaid contours indicating worst-case $|\Theta_h(j\omega)|$ for a given magnitude of P_e . The terms $(1/2)T_{ms}$ and $1/B$ approximate the contour of unity distortion.

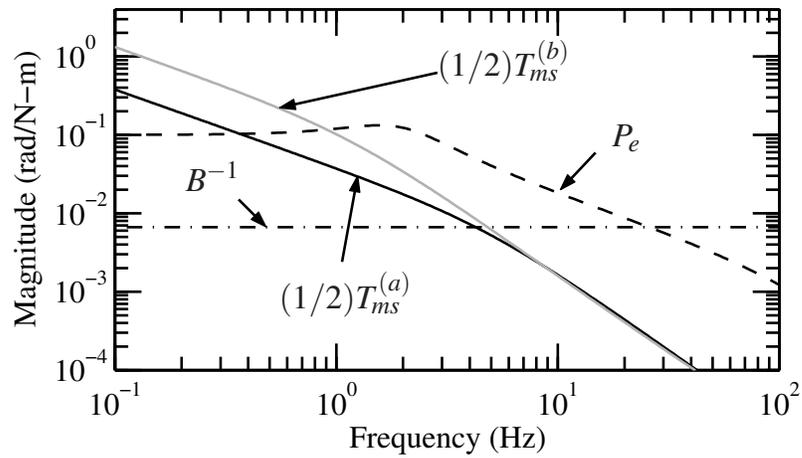


Figure 5.8 Transparency diagram comparing design (a) and (b). At frequencies where $|P_e|$ lies beneath $(1/2)|T_{ms}|$ and above $|B^{-1}|$ distortion is approximately unity or less. Design (b) expands this range by increasing $(1/2)|T_{ms}|$.

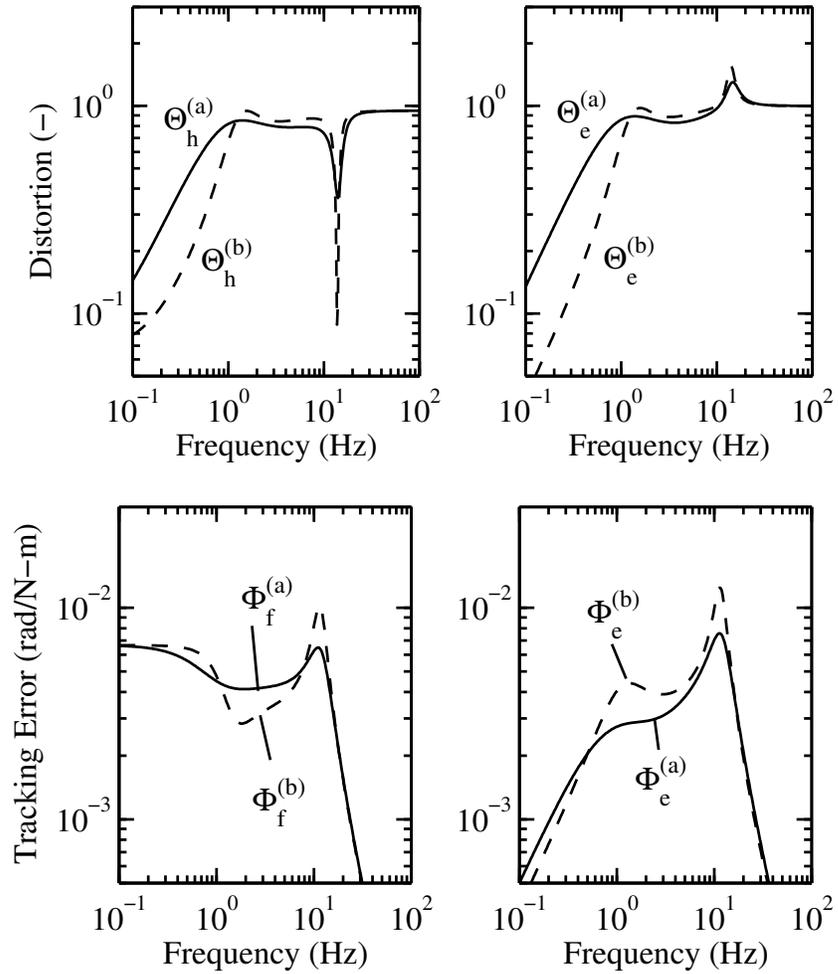


Figure 5.9 Steer-by-wire transparency performance (top) and tracking performance (bottom) for design (a) and (b). Performance is evaluated for particular road-wheel dynamics and driver dynamics.

environment forces f_e^* due to a pole at the origin in both the master and human operator dynamics.

We note that the effect of partially cancelling hardware dynamics in design (b) does not generally result in improved or degraded tracking except at the resonance between the master and slave. Around 11 Hz, the effect of partially cancelling hardware damping is to increase the peak resonant response. This effect could be mitigated by adding damping (or a derivative element) in B . This modification would also increase the transparency range given by the transparency diagram, although no significant gain in transparency would be achieved for the given environment dynamics P_e since the limiting factor in the feedback design is T_{ms} .

5.6 Conclusions

We have shown that all linear feedback designs for teleoperation based on position sensing at the master and slave can be characterized by four transfer function parameters. These parameters describe the shaped master and slave dynamics and the rigidity and scaling factor of a virtual coupler connecting the master and slave. By interpreting the relationship between these parameters and our measures of transparency and tracking performance, we found rules of thumb for the choice of parameters describing the feedback design. For transparency, we showed that a certain class of environment dynamics, defined by their magnitude in relation to the parameters that characterize the feedback design, will be presented transparently through the teleoperator.

The transparency diagram allows one to assess at a glance which environment dynamics a given teleoperator design will present transparently to the human operator. For a particular environment, it indicates whether aspects of the shaped master and slave dynamics or the rigidity of the virtual coupler will be limiting factors for transparency. While the design directive for good tracking is straight forward, that is, increase the rigidity of the virtual link, we have not found a counter-part to the transparency diagram for tracking performance. The situation is more complex because tracking depends both on the environment dynamics and the human operator dynamics.

Chapter 6

Algebraic Tradeoffs in Teleoperation

6.1 Introduction

The feedback design for teleoperation is a problem of balancing multiple design objectives. We have determined how T_m , T_s , A , and B affect transparency and tracking, and generated design directives for improving both. However, transparency and tracking are not the only feedback goals to consider. Another important property is the sensitivity of the teleoperator design to variations in the hardware dynamics. Ideally, through feedback design, we could achieve both good transparency and tracking while also reducing sensitivity to hardware dynamics. We find, however, that as in haptic rendering, there exists a fundamental tradeoff between transparency and sensitivity of the free-free dynamics to hardware variations.

6.2 Tradeoff between Transparency and Sensitivity

To measure the sensitivity of a multivariable feedback control system to variations in the hardware dynamics we examine the output sensitivity function [Skogestad and Postlethwaite, 1997]. For the position-position teleoperator, the output sensitivity associated with the free-free dynamics is

$$S = \left(I - \begin{bmatrix} P_m & 0 \\ 0 & P_s \end{bmatrix} \begin{bmatrix} C_m & C_4 \\ C_1 & C_s \end{bmatrix} \right)^{-1}. \quad (6.1)$$

If the elements of S are large, small variations in the dynamics of the master and slave will cause large changes in the actual free-free dynamics. Note that in the absence of feedback, or as the elements of the controller approach 0, the sensitivity function approaches I .

The tradeoff between transparency and sensitivity follows from an algebraic relationship between the free-free dynamics and the sensitivity function. Let the free-free dynamics shown in Fig. 5.2 be denoted by T_{ff}

$$\begin{bmatrix} x_m & x_s \end{bmatrix}^T = T_{ff} \begin{bmatrix} f_h & f_e \end{bmatrix}^T. \quad (6.2)$$

In terms of the master, slave, and controller, the free-free dynamics are

$$T_{ff} = \left(I - \begin{bmatrix} P_m & 0 \\ 0 & P_s \end{bmatrix} \begin{bmatrix} C_m & C_4 \\ C_1 & C_s \end{bmatrix} \right)^{-1} \begin{bmatrix} P_m & 0 \\ 0 & P_s \end{bmatrix}. \quad (6.3)$$

It follows from (6.1) and (6.3) that

$$T_{ff} = S \begin{bmatrix} P_m & 0 \\ 0 & P_s \end{bmatrix}. \quad (6.4)$$

This identity is a multivariable version of the identity $R = SP$ which holds in haptic rendering. However, while S and the diagonal plant matrix parallel their scalar counter-parts S and P in haptic rendering, the free-free dynamics T_{ff} and the rendered virtual environment R are not entirely analogous. The free-free dynamics exclude the environment dynamics, whereas R implicitly includes the virtual environment dynamics. It follows that the requirements on R and T_{ff} differ, and that the tradeoff between performance and sensitivity in teleoperation differs from the tradeoff in haptic rendering.

Chapter 5 reveals design relationships between transparency and the design parameters T_m , T_s , A , and B . To relate transparency and sensitivity, let us now re-express T_{ff} in terms of these parameters

$$T_{ff} = \frac{1}{1 + B(T_m + T_s)} \begin{bmatrix} T_m + T_m T_s B & T_m T_s A^{-1} B \\ T_m T_s A B & T_s + T_m T_s B \end{bmatrix}. \quad (6.5)$$

Assume that the master and slave are not identically zero and thus are invertible. It follows from (6.4) and (6.5) that

$$S = \frac{1}{1 + B(T_m + T_s)} \begin{bmatrix} \frac{T_m + T_m T_s B}{P_m} & \frac{T_m T_s A^{-1} B}{P_s} \\ \frac{T_m T_s A B}{P_m} & \frac{T_s + T_m T_s B}{P_s} \end{bmatrix}. \quad (6.6)$$

For simplicity we consider the fully symmetric teleoperator ($A = 1$). Let S^* denote S in the

limit as the virtual coupler gain $B \rightarrow \infty$. Then

$$S^* = \begin{bmatrix} \frac{T_{ms}}{P_m} & \frac{T_{ms}}{P_s} \\ \frac{T_{ms}}{P_m} & \frac{T_{ms}}{P_s} \end{bmatrix}. \quad (6.7)$$

We see that the elements of S are then the ratio of the rigidly linked, shaped master and slave dynamics (recall $T_{ms} \triangleq T_m T_s / (T_m + T_s)$) to the open-loop dynamics of the master and slave P_m and P_s . We note that the term $(1/2)T_{ms}$ also forms an asymptote for the unity distortion curve in the transparency diagram. It follows that increasing T_{ms} to achieve lower distortion also increases the sensitivity to hardware dynamics.

We now show that any position-position feedback design which achieves a desired distortion specification also induces a lower bound on the elements of S^* .

Proposition 6.2.1. *If $|\Theta^*(j\omega)| \leq \varepsilon$ for all $|P_e(j\omega)| \leq m$ then*

$$|S^*(j\omega)| \geq \frac{m(1+\varepsilon)}{\varepsilon} \left\| \begin{bmatrix} \frac{1}{P_m(j\omega)} & \frac{1}{P_s(j\omega)} \\ \frac{1}{P_m(j\omega)} & \frac{1}{P_s(j\omega)} \end{bmatrix} \right\| \quad (6.8)$$

where “ $|\cdot|$ ” are “ \geq ” element-wise.

Proof. Let $a_u = \varepsilon/(1+\varepsilon)$. By hypothesis

$$|\Theta^*(j\omega)| \leq \frac{a_u}{1-a_u}. \quad (6.9)$$

It follows from Proposition 5.4.1 that

$$\left| \frac{P_e(j\omega)}{T_{ms}(j\omega)} \right| \leq a_u. \quad (6.10)$$

for all $|P_e(j\omega)| \leq m$. Thus

$$|T_{ms}(j\omega)| \geq \frac{m}{a_u}. \quad (6.11)$$

The inequality (6.8) follows by lowering bounding $|T_{ms}(j\omega)|$ by $m(1+\varepsilon)/\varepsilon$ in each element of $|S^*(j\omega)|$. \square

The inequality (6.8) imposes a tradeoff between attenuation of sensitivity and attenuation of distortion. Note that in contrast to haptic rendering, the sensitivity function does

not approach a finite value as distortion approaches zero. For small ε , every halving of distortion doubles the cost in terms of sensitivity.

6.3 Example: Steer-by-wire

We now return to the application of steer-by-wire and consider the cost and benefits of shaping the master and slave dynamics. Recall that we have examined performance for two feedback designs: (a) a design that provides a virtual coupler between the master and slave but does not compensate for hardware dynamics, and (b) a design that also partially cancels friction in the steering-wheel and rack subsystems. Feedback compensation in design (b) increases the magnitude frequency response of T_{ms} as shown in Fig. 5.8, and thereby improves transparency. Proposition 6.2.1 predicts that improved transparency is necessarily at the cost of increased sensitivity to hardware dynamics. Let us begin with a direct interpretation of the results of Proposition 6.2.1.

Suppose we wish to achieve Θ^* less than $\varepsilon = 0.5$ for the class of road-wheel steering dynamics whose magnitude frequency response lies below the nominal road-wheel steering dynamics P_e shown in Fig. 5.7. At each frequency, we let m in Proposition 6.2.1 be given by $|P_e(j\omega)|$. It follows from Proposition 6.2.1 that any feedback design that guarantees $\Theta^* \leq 0.5$ for all environments with magnitude up to an including $|P_e(j\omega)|$ must induce sensitivity greater than or equal to that shown in Fig. 6.1. The cost to hold Θ^* less than ε increases at high frequencies where the ratios $|P_e/P_m|$ and $|P_e/P_s|$ become large. Practical constraints on the acceptable sensitivity to hardware dynamics imply a fundamental limit on the bandwidth over which we can achieve the performance objective.

We now consider the additional sensitivity induced in design (b) to achieve greater transparency compared with design (a). Figure 6.2 compares the actual sensitivity $S^{(a)}$ for design (a) to the sensitivity $S^{(b)}$ for design (b). At low frequencies, S approximates (6.7). Specifically, the (1,1) and (2,1) terms approximate T_{ms}/P_m and the (1,2) and (2,2) terms approximate by T_{ms}/P_s . Since the hardware is the same in both designs, the increase in sensitivity at low frequencies is given by the ratio of $T_{ms}^{(b)}/T_{ms}^{(a)}$. Damping effects dominate at low frequencies, and taking account, we find that the total damping of the uncompensated steering-wheel and rack is 2.1 N-m-s/rad, and the total damping after compensation is 0.6 N-m-s/rad. Thus feedback compensation of the hardware damping results is a 3.5-fold amplification of T_{ms} with associated benefits in transparency, but at the cost of a 3.5-fold increase in the DC elements of S .

The benefit of using feedback to partially cancel hardware dynamics is evident in the

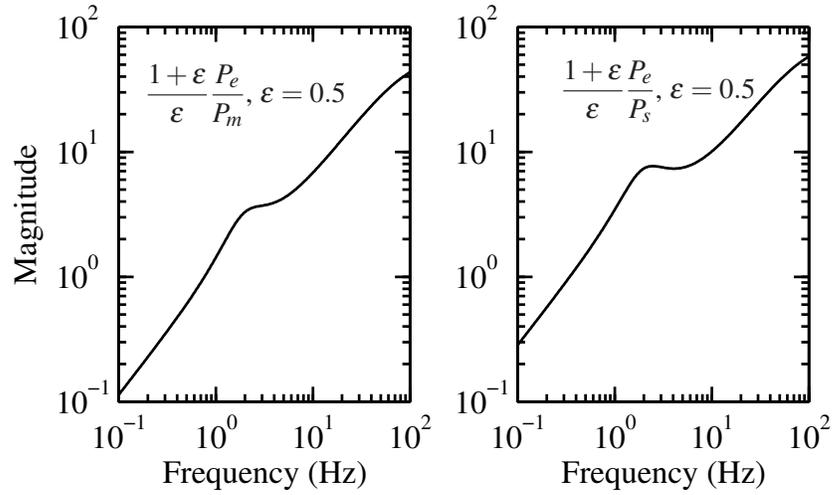


Figure 6.1 Lower bounds on elements of S^* from Prop. 6.2.1. In (6.8), m is given by $|P_e(j\omega)|$. The left subplot is the lower bound for the (1,1) and (2,1) elements of S^* , and the right subplot is the lower bound for the (1,2) and (2,2) elements.

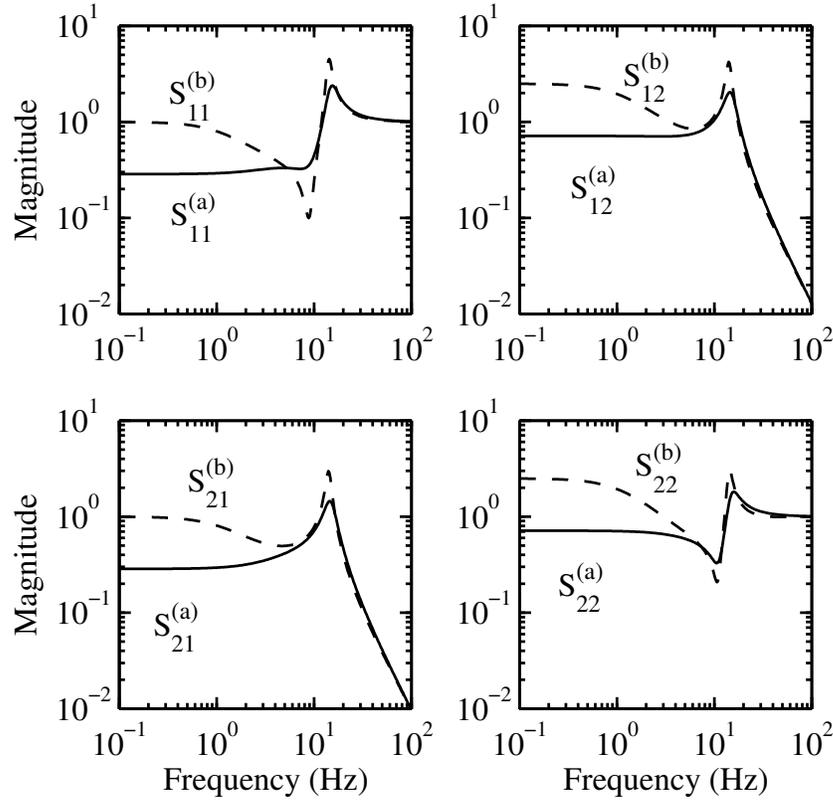


Figure 6.2 Output sensitivity function S for two steer-by-wire designs.

improved low-frequency transparency (shown in Fig. 5.9). However, the peak sensitivity at the resonant frequency of 11 Hz, while induced by the feedback controller, does not provide clear benefits in performance. We see that while improving transparency necessarily imposes a penalty in terms of sensitivity, poor sensitivity does not necessarily yield improvement in performance. Both performance and sensitivity may be poor at a frequency and striking a favorable compromise must be a goal of our feedback design. We have held the virtual coupler gain B constant to highlight the effect of cancelling hardware dynamics. It would appear, however, to be good design practice to add damping in the virtual coupler gain B when damping is actively cancelled from the hardware.

Chapter 7

Waterbed Tradeoff in Distortion

7.1 Introduction

Algebraic tradeoffs, such as the tradeoff between performance and sensitivity in haptic rendering, predict a cost to achieving performance objectives at a frequency. In addition, tradeoffs may exist within performance objectives across frequencies. For many servo-feedback systems, attenuation of the disturbance response at certain frequencies is balanced by amplification of disturbance response at other frequencies. This analytic tradeoff, termed the waterbed effect, is a consequence of the Bode Attenuation Integral applied to the Bode sensitivity function [Seron et al., 1997]. In haptic rendering, however, model-matching performance is described by distortion Θ rather than the Bode sensitivity function S . In this chapter, we study analytic relationships in haptic rendering to reveal analogous waterbed effects in distortion.

7.2 Bode Attenuation Integral in Haptic Rendering

We begin by stating Bode’s Attenuation Integral in its general form. Let $H(s)$ denote any rational transfer function with relative degree zero. Then, in the limit as $s \rightarrow \infty$,

$$H(s) \rightarrow H(\infty) \tag{7.1}$$

where $H(\infty)$ is finite and non-zero. Let z_i for $i = 1, \dots, n_z$ denote the zeros of $H(s)$ in the open right-half plane, and let p_i for $i = 1, \dots, n_p$ denote the poles of $H(s)$ in the open right-half plane. Also, let “log” denote the natural logarithm (base e).

Proposition 7.2.1 (Bode's Attenuation Integral). *Let $H(s)$ be a rational transfer function with relative degree zero. Then $H(s)$ satisfies*

$$\int_0^\infty \log \left| \frac{H(j\omega)}{H(\infty)} \right| d\omega = \frac{\pi}{2} \lim_{s \rightarrow \infty} \frac{s(H(s) - H(\infty))}{H(\infty)} + \sum_{i=1}^{n_z} \operatorname{Re} z_i - \sum_{i=1}^{n_p} \operatorname{Re} p_i. \quad (7.2)$$

Proof. See [Seron et al., 1997]. □

Under appropriate assumptions on the haptic device dynamics P and the virtual environment dynamics R_d , the distortion transfer function Θ satisfies the hypotheses of (7.2).

Proposition 7.2.2. *Assume that P is strictly proper and has no open-right half-plane (ORHP) poles, that R_d has no ORHP zeros, and that P/R_d is strictly proper. Assume further that C is proper and stabilizes P . Then*

$$\int_0^\infty \log |\Theta(j\omega)| d\omega \geq -\frac{\pi}{2} A, \quad (7.3)$$

where

$$A \triangleq \lim_{s \rightarrow \infty} \frac{sP(s)}{R_d(s)}. \quad (7.4)$$

Proof. The hypotheses imply that Θ is proper and stable, and that $\Theta(\infty) = -1$. Denote the ORHP zeros of Θ by $\zeta_i : i = 1, \dots, N_\zeta$. Then Θ satisfies *Bode's Attenuation Integral*

$$\int_0^\infty \log \left| \frac{\Theta(j\omega)}{\Theta(\infty)} \right| d\omega = \frac{\pi}{2} \lim_{s \rightarrow \infty} \frac{s(\Theta(s) - \Theta(\infty))}{\Theta(\infty)} + \sum_{i=1}^{N_\zeta} \operatorname{Re} \zeta_i, \quad (7.5)$$

and thus

$$\int_0^\infty \log |\Theta(j\omega)| d\omega = -\frac{\pi}{2} \lim_{s \rightarrow \infty} s(\Theta(s) + 1) + \sum_{i=1}^{N_\zeta} \operatorname{Re} \zeta_i. \quad (7.6)$$

The hypotheses imply that T is strictly proper. The result then follows by replacing Θ with (3.5) and by noting that $\operatorname{Re} \zeta_i > 0$. □

The presence of bandwidth constraints generates a nontrivial waterbed tradeoff between the level of distortion over a finite bandwidth. Suppose there exists specification on distortion over a performance bandwidth ω_0

$$|\Theta(j\omega)| \leq M_\Theta, \quad \forall \omega \leq \omega_0. \quad (7.7)$$

Further suppose there exists a limit on closed-loop bandwidth above a cut-off frequency ω_c

$$|T(j\omega)| \leq M_T(\omega), \quad \forall \omega \geq \omega_c > \omega_0 \quad (7.8)$$

where $M_T(\omega) \rightarrow 0$ and $\omega M_T(\omega)$ is bounded as $\omega \rightarrow \infty$. The area under $\log |\Theta(j\omega)|$ up to ω_0 is upper bounded by $\omega_0 \log M_\Theta$. To bound the integral of $\log |\Theta(j\omega)|$ above ω_c , let us factor Θ by

$$\Theta = \left(\frac{P}{R_d} - 1 \right) \left(1 + \frac{P/R_d}{1 - P/R_d} T \right) \quad (7.9)$$

Thus $\log |\Theta(j\omega)|$ can be expressed as the sum of a term that depends only on the haptic device dynamics P and the virtual environment dynamics R_d , and a term that depends on the feedback design through T . However, we can bound this latter term in a manner that is independent of the particular feedback design.

Assume that the hypotheses of Proposition 7.2.2 are satisfied together with the specification (7.8). It follows that P/R_d is strictly proper and, for any positive $\delta < 1/2$, there exists a frequency ω_a such that

$$\left| \frac{P(j\omega)}{R_d(j\omega)} \right| < \frac{\delta}{2} \left(\frac{\omega_a}{\omega} \right)^l. \quad (7.10)$$

Note that by the triangle inequality, $\left| \frac{z}{1-z} \right| < \delta$ for any complex z such that $\frac{|z|}{2} < \delta < 1$. Thus

$$\left| \frac{P(j\omega)/R_d(j\omega)}{1 - P(j\omega)/R_d(j\omega)} \right| < \delta \left(\frac{\omega_a}{\omega} \right)^l, \quad \forall \omega > \omega_a. \quad (7.11)$$

Due to the bandwidth constraint (7.8), there exists $\omega_b \geq \omega_c$ such that

$$|T(j\omega)| < \left(\frac{\omega_b}{\omega} \right), \quad \forall \omega \geq \omega_b. \quad (7.12)$$

If we let ω_1 be the greater of ω_a and ω_b , then there exists a frequency $\omega_1 \geq \omega_c$ and $\delta < 1/2$, $l > 0$ such that

$$\left| \frac{P(j\omega)/R_d(j\omega)}{1 - P(j\omega)/R_d(j\omega)} T(j\omega) \right| < \delta \left(\frac{\omega_1}{\omega} \right)^{1+l}, \quad \forall \omega > \omega_1. \quad (7.13)$$

Corollary 7.2.1. *Assume that the hypotheses of Proposition 7.2.2 are satisfied along with*

the specifications (7.7) and (7.8). Then

$$\sup_{\omega \in (\omega_0, \omega_1)} \log |\Theta(j\omega)| \geq \frac{1}{\omega_1 - \omega_0} \left(-\frac{\pi}{2}A - \omega_0 \log M_\Theta - B \right), \quad (7.14)$$

where A is given by (7.4), and B is finite and defined by

$$B \triangleq \int_{\omega_1}^{\infty} \log \left| \frac{P(j\omega)}{R_d(j\omega)} - 1 \right| d\omega + \frac{3\delta\omega_1}{2l}. \quad (7.15)$$

Proof. The performance requirement (7.7) implies that

$$\int_0^{\omega_0} \log |\Theta(j\omega)| d\omega \leq \omega_0 \log M_\Theta. \quad (7.16)$$

We use the fact that $|\log(1+z)| \leq \frac{3|z|}{2}$ for $|z| < 1/2$ to show that

$$\begin{aligned} \int_{\omega_1}^{\infty} \log |\Theta(j\omega)| d\omega &\leq \int_{\omega_1}^{\infty} \log \left| \frac{P(j\omega)}{R_d(j\omega)} - 1 \right| d\omega + \frac{3\delta}{2} \int_{\omega_1}^{\infty} \left(\frac{\omega_1}{\omega} \right)^{1+l} d\omega \\ &= B. \end{aligned} \quad (7.17)$$

By hypothesis, P/R_d is strictly proper, and thus B is finite. We can lower bound the maximum of $\log |\Theta(j\omega)|$ between ω_0 and ω_1 by

$$\sup_{\omega \in (\omega_0, \omega_1)} \log |\Theta(j\omega)| \geq \frac{1}{\omega_1 - \omega_0} \int_{\omega_0}^{\omega_1} \log |\Theta(j\omega)| d\omega \quad (7.18)$$

It follows from (7.3), (7.16), and (7.17) that

$$\int_{\omega_0}^{\omega_1} \log |\Theta(j\omega)| d\omega \geq -\frac{\pi}{2}A - \omega_0 \log M_\Theta - B_1. \quad (7.19)$$

The result (7.14) follows by substituting (7.19) into (7.18). \square

Example 7.2.1. Suppose that

$$P = \frac{1}{s(s+1)}, \quad \text{and} \quad R_d = \frac{1}{s+10}. \quad (7.20)$$

The device model P captures the effects of hardware damping and inertia, and the virtual environment is an idealized spring-damper. Note that P and P/R_d are strictly proper and that R_d has no ORHP zeros. Thus the hypotheses of Proposition 7.2.2 are satisfied. It

follows that for any proper stabilizing controller

$$\int_0^\infty \log |\Theta(j\omega)| d\omega \geq -\frac{\pi}{2} \lim_{s \rightarrow \infty} \frac{sP(s)}{R_d(s)} = -\frac{\pi}{2}. \quad (7.21)$$

Further suppose that the performance objective is described by (7.7) and the bandwidth constraint is captured by (7.13). We apply Corollary 7.2.1 to lower bound the peak of $\log |\Theta(j\omega)|$ between ω_0 and ω_1 . Figure 7.1 depicts the bounds on $\log |\Theta(j\omega)|$ where $M_\Theta = 0.1$, $\omega_0 = 6$ (rad/s), and the bandwidth constraint (7.13) has parameters $\omega_1 = 10$ (rad/s), $\delta = 0.3$, and $l = 2$. Below 6 rad/s, $\log |\Theta(j\omega)|$ is upper bounded by the rectangular area. Above 10 rad/s, $\log |\Theta(j\omega)|$ is bounded above and below by converging bounds which are symmetric around $\log |P(j\omega)/R_d(j\omega) - 1|$. The area under the upper bound corresponds to B_1 . All feedback designs that satisfy both performance objective and the bandwidth constraint must have a peak value of $\log |\Theta(j\omega)|$ that reaches or exceeds approximately 2.64 (or an absolute magnitude of $e^{2.64} = 14.0$), shown by the dotted line between ω_0 and ω_1 . Fig. 7.2 shows how the bound on the peak of $\log |\Theta(j\omega)|$ varies with ω_0 . The cost of increased performance bandwidth ω_0 is increased distortion amplification between ω_0 and ω_1 .

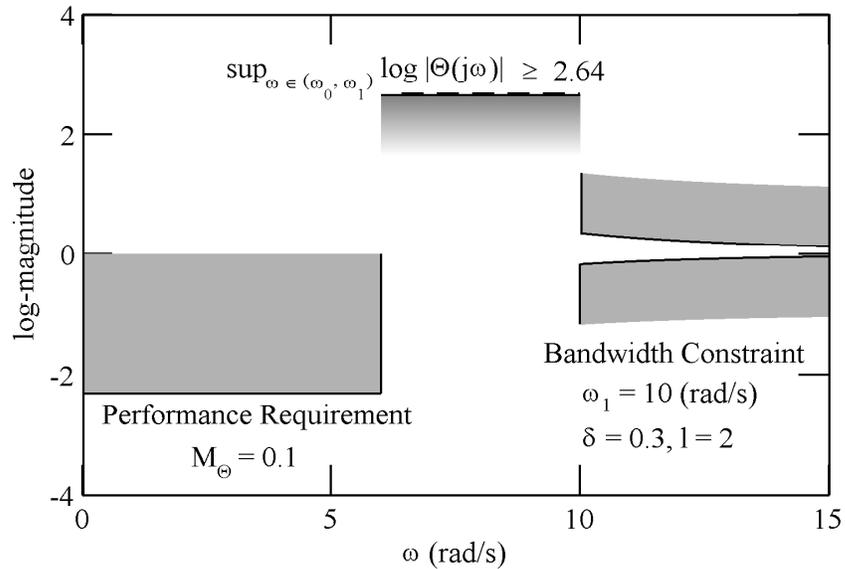


Figure 7.1 Sample bounds on $\log |\Theta(j\omega)|$ due to performance and bandwidth requirements. Note that “log-magnitude” is base- e .

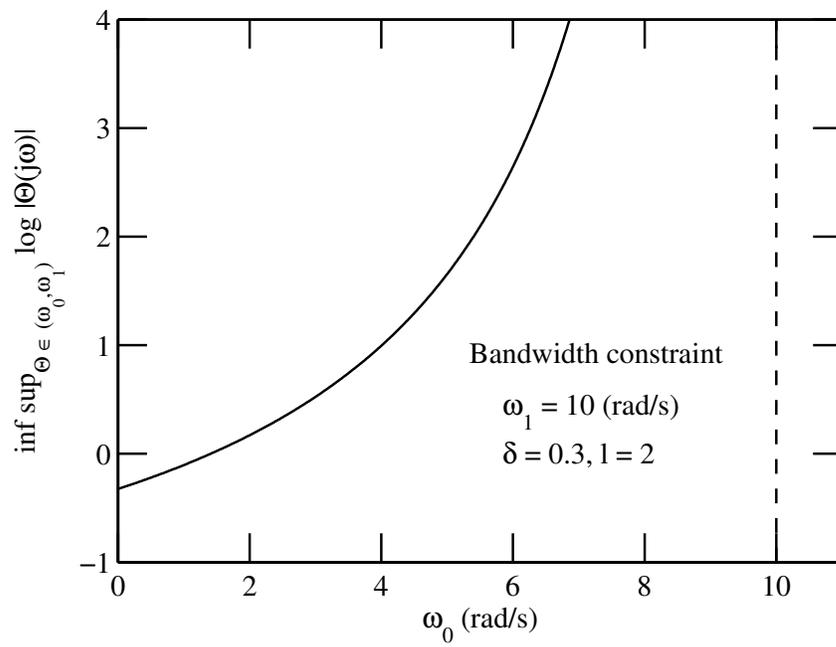


Figure 7.2 Lower bound on the peak of $\log |\Theta(j\omega)|$ between ω_0 and ω_1 predicted by (7.14).

Chapter 8

A Fundamental Conflict between Performance and Passivity in Haptic Rendering

8.1 Introduction

Depending on the intrinsic dynamics of the haptic device and the desired dynamic response, compensation for the haptic device dynamics may be required for the actual closed-loop response to closely match the desired dynamic response. In general, the mechanical properties of the device such as damping and inertia are minimized during hardware design to avert the need for compensation in the feedback design. However, minimizing these hardware dynamics is in conflict with other design considerations such as structural strength or component cost. Feedback control offers the ability to compensate for hardware dynamics; however the designer must be simultaneously aware of inherent costs and limitations.

The human-in-the-loop aspect of haptic rendering raises special stability concerns. Typical frequency-domain techniques for assessing stability such as phase margin and gain margin are not appropriate for haptic rendering since the mechanical interaction with the human operator introduces significant nonlinear, time-varying dynamics. This particular stability problem is termed *coupled stability* [Colgate and Hogan, 1989], and a common solution is to design the controller such that the closed-loop response presented to the user through the haptic device has a passive dynamic response. Coupled stability is assured if the human operator behaves passively. It is then important to determine the class of dynamic responses that can be rendered by the haptic device while maintaining the passivity condition for coupled stability with the user.

Previous analysis of haptic rendering identifies limitations on the feedback design due

to sampling effects and quantization error. When rendering a *virtual wall*, passivity imposes limitations on the gains of the controller parameters [Minsky et al., 1990; Colgate et al., 1995; Mahvash and Hayward, 2005; Mehling et al., 2005; Diolaiti et al., 2006; Shen and Goldfarb, 2006]. Extensions to the work on the virtual wall problem predict limits for more general nonlinear virtual environments [Colgate and Schenkel, 1997; Miller et al., 2000]. Prior literature on haptic rendering has not analyzed the challenges of accurately rendering passive linear time-invariant dynamics without regard to sampling or quantization issues. In this respect, our analysis is similar to analysis of limitations in end-point impedance control for robots [Colgate and Hogan, 1989]. However, unlike typical impedance control which uses force sensing, we consider haptic rendering using position sensing as the feedback signal.

We show that certain linear time-invariant passive dynamic systems cannot be approximated passively over a given finite bandwidth through feedback design. This fundamental conflict between performance and passivity holds for all linear time-invariant controllers. In contrast to previous work, the limitation we reveal is not mitigated by faster sampling or finer sensor quantization. A necessary condition is determined for the existence of a feasible feedback design, and interpretation of this condition reveals that certain compensation of hardware dynamics cannot be achieved while maintaining a passive closed-loop response.

8.2 Bode Gain-Phase Relationship

Analyticity of transfer functions imposes relationships between the magnitude and phase that may conflict with feedback design goals. The Bode gain-phase integral predicts that the phase of a proper, stable, minimum phase transfer function along the $j\omega$ -axis is completely determined by the magnitude of the transfer function along the $j\omega$ -axis [Bode, 1945; Seron et al., 1997]. We use a related integral expression to relate performance and passivity requirements in haptic rendering. We state the result in terms of S ; however we note that the result is not particular to the Bode sensitivity function.

Let p_i for $i = 1, \dots, n_p$ be the poles of S and z_i for $i = 1, \dots, n_z$ be the zeros of S such that

$$S(s) = \frac{k \prod_{i=1}^{n_z} (s - z_i)}{\prod_{i=1}^{n_p} (s - p_i)}. \quad (8.1)$$

We assume that S is normalized such that k is positive. Let $\arg(s)$ assume values from $-\pi$

to π and define

$$\arg S(s) = \sum_{i=1}^{n_z} \arg(s - z_i) - \sum_{i=1}^{n_p} \arg(s - p_i). \quad (8.2)$$

Then $\log S(s)$ is given by $\log |S(s)| + j \arg S(s)$.

Bode gain-phase relationships follow from Cauchy's integral theorem applied to contour integrals that enclose the right-half plane. We consider the contour integral of $\log S/\sqrt{1+s^2/\omega_0^2}$. Assume that all p_i 's lie in the open left-half plane and that all z_i 's lie in the closed left-half plane. Then the integrand $\log S/\sqrt{1+s^2/\omega_0^2}$ is analytic on and inside a contour defined by a large semicircle around the right-half plane and the $j\omega$ -axis. Small semi-circle indentations at $\pm j\omega_0$ and at any zeros of S on the imaginary axis are added to avoid singularities. We take $\sqrt{1-\omega^2/\omega_0^2}$ to be positive for $-\omega_0 < \omega < \omega_0$, positive imaginary for $\omega > \omega_0$, and negative imaginary for $\omega < -\omega_0$. By the hypotheses on P and C , it follows that $S(j\omega) \rightarrow 1$ as $\omega \rightarrow \infty$. Then S satisfies (cf. [Bode, 1945, Eqn. 13-36])

$$\int_0^{\omega_0} \frac{\log |S(j\omega)|}{\sqrt{1-\omega^2/\omega_0^2}} d\omega = - \int_{\omega_0}^{\infty} \frac{\arg S(j\omega) d\omega}{\sqrt{\omega^2/\omega_0^2 - 1}}. \quad (8.3)$$

This integral equality implies that amplification of $S(j\omega)$ below ω_0 must be balanced with negative phase area above ω_0 . If we change the variable of integration to $v \triangleq \log \omega/\omega_0$, the weighting factors $1/\sqrt{1-\omega^2/\omega_0^2}$ and $1/\sqrt{\omega^2/\omega_0^2 - 1}$ become $e^v/\sqrt{1-e^{2v}}$ and $e^v/\sqrt{e^{2v}-1}$. Inspection of these weighting factors (shown in Fig. 8.1) reveals that the left-hand side of (8.3) is strongly influenced by $\log |S(j\omega)|$ within a decade below ω_0 . The contribution of the phase of $S(j\omega)$ to the right-hand side of (8.3) is more evenly distributed with bias towards values near ω_0 .

8.3 Conflict between Performance and Passivity

Performance and passivity requirements generate constraints on the magnitude and phase of closed-loop transfer functions Θ and R . The Bode gain-phase relationship (8.3) provides a means to relate magnitude and phase requirements if they are expressed in terms of a single closed-loop transfer function. To do this, we express performance and passivity requirements in terms of the Bode sensitivity function S . It follows from the definitions of R

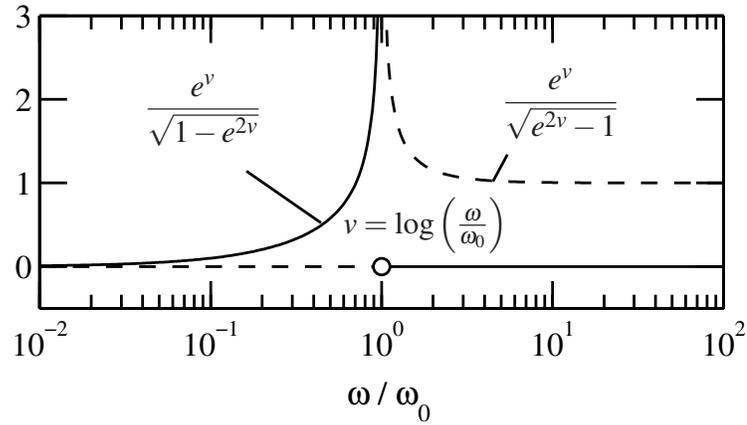


Figure 8.1 Gain-phase weighting factors from (8.3) shown on a logarithmic frequency scale.

and S that

$$R = PS. \quad (8.4)$$

Thus the magnitude and phase of S are given by

$$|S| = |R|/|P| \quad (8.5)$$

$$\arg S = \arg R - \arg P. \quad (8.6)$$

From (2.5), distortion may be written as

$$R = (1 + \Theta)R_d. \quad (8.7)$$

It follows from (8.5) and (8.7) that

$$\log |S| = \log \left| \frac{R_d}{P} \right| + \log |1 + \Theta|. \quad (8.8)$$

Assume that the performance specification is of the form

$$|\Theta(j\omega)| \leq M_\Theta, \quad \text{for } 0 \leq \omega \leq \omega_c. \quad (8.9)$$

where $M_\Theta < 1$. It follows from the triangle inequality that, for any feedback design that achieves the performance specification, the inequality

$$\log |S(j\omega)| \geq \log \left| \frac{R_d(j\omega)}{P(j\omega)} \right| + \log |1 - M_\Theta|. \quad (8.10)$$

holds for $0 \leq \omega \leq \omega_c$. This inequality provides a lower bound on the magnitude of $S(j\omega)$ that depends only on the hardware P , the desired rendered dynamics R_d and the performance specification (8.9). Notably, the lower bound on the magnitude of $S(j\omega)$ does not depend on the feedback design.

The constraint on the phase of $R(j\omega)$ given by (2.18) can be expressed as a constraint on the phase of $S(j\omega)$ using (8.6). For sR to be passive, we require that

$$-\pi - \arg P(j\omega) \leq \arg S(j\omega) \leq -\arg P(j\omega). \quad (8.11)$$

For our present discussion, we focus on the implications of the lower bound on the phase of $S(j\omega)$ when the haptic device dynamics are dominated by inertial effects at high frequencies. For such devices, the phase of $P(j\omega)$ approaches $-\pi$ at high frequencies. It follows that any phase lag in $S(j\omega)$ must be small at high frequencies so as to not violate the lower bound in (2.18).

An inherent conflict emerges between performance and passivity when (8.10) and (8.11) are related through the Bode gain-phase relationship (8.3).

Proposition 8.3.1. *Assume that P has relative degree 2, that P has no open right-half plane poles, and that $0 \leq M_\Theta < 1$. A necessary condition for the existence of a proper, stabilizing controller C that meets the performance specification (8.9) and passivity requirement (2.18) is*

$$\int_0^{\omega_0} \frac{\log \left| \frac{R_d(j\omega)}{P(j\omega)} \right| + \log |1 - M_\Theta|}{\sqrt{1 - \omega^2/\omega_0^2}} d\omega \leq \int_{\omega_0}^{\infty} \frac{\pi + \arg P(j\omega)}{\sqrt{\omega^2/\omega_0^2 - 1}} d\omega. \quad (8.12)$$

for all $0 < \omega_0 \leq \omega_c$.

Proof. Suppose that there exists a proper stabilizing controller C that achieves the performance specification and satisfies the passivity requirement on sR . Since open right-half plane poles of C are open right-half plane zeros of R (recalling that $R = P/(1 + PC)$), it follows that C has no open right-half plane poles. Together with the hypotheses on P , it follows that S is stable and has no open right-half plane zeros. Then S must satisfy the Bode gain-phase relationship (8.3).

For $0 \leq \omega \leq \omega_c$, the left-hand side of (8.3) is lower bounded by (8.10) which captures the performance requirement. We use the passivity condition (8.11) to upper bound $-\arg S(j\omega)$ by $\pi + \arg P(j\omega)$ for all positive frequencies. Thus the resulting inequality given by (8.12) is a necessary condition for the existence of a controller that satisfies the hypotheses. \square

We can demonstrate the existence of passive virtual environment dynamics that violate (8.12) by choosing any non-zero passive virtual environment R_d and multiplying it by a sufficiently large scalar. Suppose that ω_0 and M_Θ are fixed. The left-hand side of (8.12) grows with the magnitude of R_d , but the right-hand side depends only on the phase of P . A sufficiently large scaling of any passive virtual environment will violate (8.12). It follows that no proper, stabilizing controller C exists that meets the performance specification while presenting passive dynamics to the user.

Proposition 8.3.1 implies a limitation on the ability to compensate for hardware dynamics while presenting a passive response to the user. We say that the feedback design *partially cancels* hardware dynamics at a frequency if $|R(j\omega)| > |P(j\omega)|$. At these frequencies the magnitude of the closed-loop response y to the human operator's force f exceeds the open-loop response. This definition generalizes the intuition that shaping the dynamic response of a device with inertia to make it respond like a device with less inertia is partial cancellation of hardware dynamics. It follows that, at frequencies where $|R_d(j\omega)| > |P(j\omega)|$, accurate rendering of the virtual environment requires a feedback design that partially cancels the hardware dynamics. Examining (8.12), we note that only desired closed-loop dynamics R_d that require partial cancellation can cause the inequality to be violated. The unpowered hardware dynamics sP are passive, so $\pi + \arg P(j\omega)$ is positive. It follows that the right-hand side of (8.12) is always positive. Furthermore, the term $\log|1 - M_\Theta|$ is always negative, so (8.12) will always hold if $|R_d(j\omega)| < |P(j\omega)|$ for $0 < \omega < \omega_c$.

A practical implication of Proposition 8.3.1 is that even a small amount of compensation for hardware inertia may be impossible without violating the passivity requirement on sR . Consider haptic device dynamics P with only inertia. Since $\arg P(j\omega)$ is $-\pi$, the right-hand side of (8.12) is exactly zero. The left-hand side of (8.12) must remain less than or equal to zero. This requirement implies that $|R_d(j\omega)/P(j\omega)|$ must not significantly exceed unity over a range of frequencies below ω_c .

8.4 Example

We now consider the conflict between performance and passivity for a simple example problem. Let the haptic device model be a damped mass system $P = \frac{1}{s^2+s}$ and let the virtual environment dynamics be an ideal parallel spring-damper $R_d = \frac{1}{s+10}$. For the performance specification, we let $M_\Theta = 0.5$ and we consider several different values for ω_c . Figure 8.2 shows the left-hand and right-hand sides of (8.12) using four values for ω_c indicated by lines A, B, C, and D. Proposition 8.3.1 predicts that there is no proper stabilizing feedback design that achieves $|\Theta(j\omega)| \leq 0.5$ up to 10 (rad/s) and that makes sR passive. Thus the performance bandwidth indicated by A cannot be achieved passively.

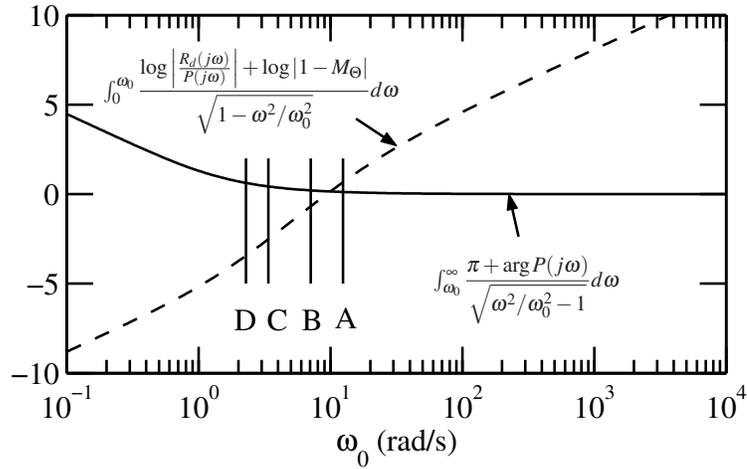


Figure 8.2 Right and left sides of the inequality in Proposition 8.3.1 for $M_\Theta = 0.5$. Lines A, B, C, and D indicate four specification for ω_c .

Let us examine the conflict between performance and passivity through example controller designs. Since the rendered dynamics R are given by $P/1 + PC$, we can algebraically solve for the controller that makes $R \equiv R_d$. However, the solution $C = R_d^{-1} - P^{-1}$ is not necessarily proper or stabilizing. In the present example, the exact algebraic solution for C is $-s^2 + 10$. This improper controller is not practical, and furthermore, multiplying this expression by a low-pass filter to obtain a proper controller does not necessarily result in closed-loop stability. Note that classical design techniques for loop-shaping do not provide a direct method of shaping the closed-loop dynamics, and optimal synthesis tools such LQG and H_∞ do not necessarily produce stable controllers, a prerequisite for sR to be passive. To generate feedback designs which approximate the desired closed-loop response using a stable, proper controller, we use controllers of the form

$$C = \left(\frac{(\tau s + 1)^2}{R_d} - \frac{1}{P} \right) \left(\frac{1}{\gamma s + 1} \right)^3. \quad (8.13)$$

For the values of τ and γ given in Table 8, the controller given by (8.13) yields closed-loop stability and the desired performance. The frequency responses of distortion Θ are shown in Fig. 8.3.

Design	$(2\pi\tau)^{-1}$ Hz	$(2\pi\gamma)^{-1}$ Hz
A	8	400
B	4	100
C	1.6	15
D	0.6	2

Table 8.1 Parameters τ and γ of the four feedback designs.

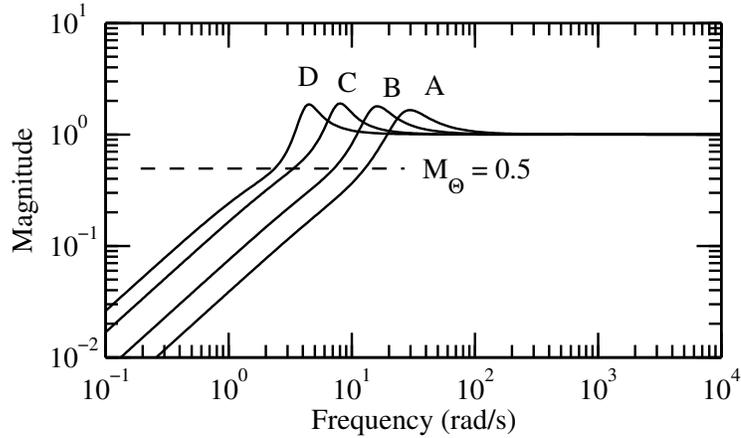
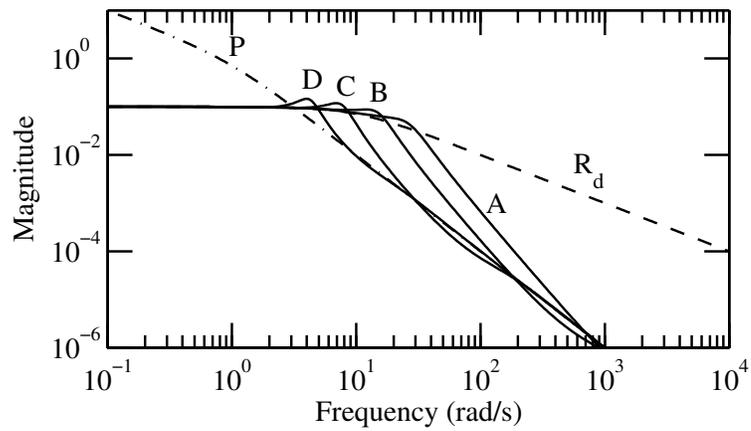
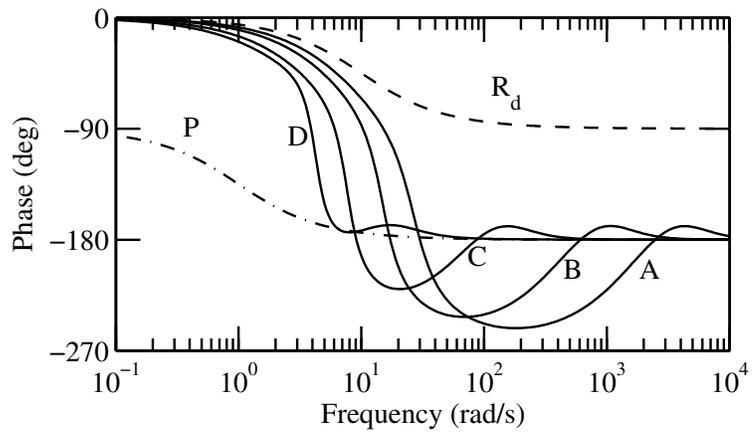


Figure 8.3 Distortion Θ for designs A, B, C and D.

Frequency responses of the rendered virtual environment R for each design, labeled A through D, are shown in Fig. 8.4. Inspection of the phase plot reveals that designs A, B, and C violate the lower bound of -180 degrees and thus do not present the user with passive dynamics. Proposition 8.3.1 proves that no feasible design exists that meets the performance specification for design A and satisfies the phase criteria for passivity. Progressive relaxations of the bandwidth incur smaller excursions below -180 degrees. Only design D both meets the performance specification and the passivity requirement. For designs B and C, inequality (8.12) is not violated so Proposition 8.3.1 provides no determination on whether some other feedback design exists that achieves the bandwidth of C without violating passivity. There may then exist a higher-order controller than (8.13) which satisfies both performance and passivity requirements.



(a) Bode magnitude plot



(b) Bode phase plot

Figure 8.4 Frequency response of the haptic device P , the virtual environment R_d , and four causal approximations to R_d . Designs A, B, and C violate passivity since their phase drops below -180 degrees.

8.5 Conclusions

We have shown that, for given haptic device hardware and performance specification for the feedback design, a certain class of passive transfer functions cannot be approximated while presenting a passive response to the human operator. This class of transfer functions is characterized by compensation of hardware dynamics. In Proposition 8.3.1 we have provided a necessary condition for the existence of a stabilizing, proper feedback design that can passively approximate a desired closed-loop response. If the inequality (8.12) is violated at any frequency, we can conclude that no feasible feedback design exists that satisfies both performance and passivity requirements.

Passivity is a strong requirement which may be overly restrictive. Less conservative coupled stability criteria may be available through a robust stability analysis assuming a set of possible user dynamics. Although relaxing passivity would mitigate the conflict between performance and passivity, other tradeoffs may exist such as between performance bandwidth and the closed-loop bandwidth. As seen in the example, small increases in the performance bandwidth cause large growth in the controller bandwidth. While a high-order controller may be used to increase the rate of roll-off, such a design will also increase the amount by which the phase requirement of passivity is violated.

Chapter 9

Conclusions

9.1 Summary of Results

Feedback design for haptic interface and teleoperator systems must abide by limitations imposed by hardware and inherently involves tradeoffs between mutually incompatible goals. Knowledge of design limitations and tradeoffs inform both hardware and control design. In this dissertation, we have derived previously unrecognized design limitations and tradeoffs in feedback design for haptic rendering and teleoperation. Additionally, analysis of these feedback systems has led to our development of new tools and techniques for controller design.

Key to our analysis is the introduction of **distortion** as a measure of performance in haptic rendering and teleoperation. By analyzing error in rendering of environment dynamics, we are able to apply the triangle inequality to derive tradeoff relationships as well as prove bounds on performance. The inequality (4.15) provides an upper bound on distortion over a class of virtual environments rendered with the cancellation coupler. Similarly, Proposition 5.4.1 bounds distortion in teleoperation over a class of environment dynamics defined by their magnitude frequency response. These results were not previously available in literature where performance was gauged by expressions other than distortion.

In contrast to servo-control systems where the performance goals are achieved through attenuation of the Bode sensitivity, accurate rendering of environment dynamics in haptics and teleoperation is achieved by attenuation of the distortion transfer function. As such, haptic rendering and teleoperation are subject to design tradeoffs not present in typical servo-control systems. Through algebraic analysis we have shown that attenuation of S , T , and Θ are mutually incompatible goals and imply a three-way tradeoff in haptic rendering. For certain combinations of virtual environment dynamics and hardware, desired specifi-

cations on S , T , and Θ may be unachievable. The inequalities (3.2), (3.4), and (3.6) form a set of feasibility requirements on design specifications.

In the three-way algebraic tradeoff, the tradeoff between performance and sensitivity properties in haptic rendering has no analogue in typical servo-control design. We have shown that the severity of this tradeoff is captured by the transfer function Γ given by (3.16). Practically speaking the tradeoff reflects the fact that feedback compensation of hardware dynamics is required to render certain virtual environments but also induces sensitivity to variations in the hardware. If the sensitivity required is too large to be practical, worse performance must be accepted or the hardware can be re-designed to reduce intrinsic dynamics.

Both teleoperation and haptic rendering are subject to an algebraic tradeoff between performance and sensitivity; however we have shown that the two tradeoffs are not entirely analogous. In haptic rendering, the sensitivity function approaches a finite value as distortion approaches zero. In contrast, the elements of the multivariable teleoperator sensitivity function are unbounded as distortion approaches zero. The primary difference is that the environment dynamics are internal to the controller in haptic rendering but extrinsic for teleoperation.

Insight gained from the feedback analysis of haptic rendering and teleoperation has led us to refine current design practice. Typical practice in haptic rendering is to select a virtual coupler that mimics the response of a mechanical coupler such as a spring. However, this approach unnecessarily constrains the multi-input/multi-output response of the virtual coupler. An analysis of distortion provides design directives for all elements of the virtual coupler which leads us to propose the **cancellation coupler**. Experimental results demonstrate the performance improvements over standard practice.

Our analysis of feedback design tradeoffs in teleoperation has also contributed techniques for controller design and tuning. Parameterization of teleoperator feedback designs was essential preliminary work that enabled our analysis of the tradeoff between performance and sensitivity. However, the parameterization on its own is a useful tool for design. Current practice assumes particular structures for the controller such as a simple PD feedback on tracking error or the more sophisticated structure of wave variable control. The parameterization we introduce describes all LTI feedback designs and permits design in the frequency domain. The **transparency diagram** provides design directives for frequency shaping the design parameters, and using (5.6)–(5.8), we can map design parameters back to controller elements.

Prior analysis of haptic interface and teleoperator control design has not taken into account the limitations imposed by bandwidth constraints. By interpreting the Bode atten-

uation integral for distortion, we are the first to analyze a **waterbed tradeoff** in the haptic rendering. As a consequence, improved accuracy in rendering virtual environments at some frequencies is necessarily at the cost of worse accuracy at other frequencies.

The relevance of bandwidth constraints in feedback design is evident in our analysis of passivity in haptic rendering. In the absence of bandwidth constraints, there is no conflict between passivity and performance when the virtual environment is itself passive. However, when the closed-loop bandwidth is finite, interpretation of Bode gain-phase relationship 8.3 reveals that accurate rendering of certain passive virtual environments may necessarily violate passivity.

9.2 Near-term Work

The teleoperator design tools and tradeoff relationships developed in this dissertation have been demonstrated through worked examples, but still need be supported by experiments. Certain idealizations are made in our theoretical treatment which are violated in practice. For example, we assume linear hardware models and impedance-type device dynamics (that is hardware where the user’s input and the control input affect the plant output through the same dynamics). The merit of our theoretical predictions is only in their ability to predict significant features of design practice.

The use of distortion as a performance metric in teleoperation is entirely new and has not yet been measured experimentally. To determine this quantity we need a calibrated dynamic system such as a mass, spring, or damper to serve as the known environment dynamics. Experimental estimates of Θ_h and Θ_e then require measurement of two dynamic responses: the response of the master x_m to excitations from f_h and the response of x_m to excitations from f_e^* . The signal x_m is easily obtained from the control system; however providing known excitations f_h and f_e^* is more challenging. One approach is to superimpose these forces on the master and slave control outputs u_m and u_s . The accuracy of this method depends on the quality of the motor calibration and has limited bandwidth due to compliance between the motor and user interface. Alternatively independent actuators may be used to drive the master and slave with inline force sensors to measure f_h and f_e^* .

Experimental support for the algebraic tradeoff between transparency and sensitivity may be obtained by comparing feedback designs which employ varying degrees of feedback compensation. We can empirically determine the output sensitivity by injecting virtual noise into the sensed signals x_m and x_s and measuring the closed-loop response of x_m and x_s . This dynamic response is the output complementary sensitivity function T_o , and the

output sensitivity function S_o is given by $I - T_o$. The magnitude of the virtual noise must, of course, be much larger than sensor quantization.

The parameterization of all LTI teleoperator feedback includes an asymmetry parameter A which we have assumed to be unity to obtain the results obtained in Chapters 5 and 6. When the master and slave are asymmetric, some frequency shaping of A may help reduce tracking error. However, transparency degrades as A approaches zero or infinity. Further analysis of the role of asymmetry in teleoperator design is merited and may uncover a fundamental tradeoff between tracking and transparency performance.

Prior literature on haptic rendering has addressed violation of passivity due to sampling and quantization, but has not to date experimentally studied or demonstrated the coupled instability which may arise when two stable LTI systems are connected in feedback. We note that prior analyses of passivity in haptic rendering have not imposed bandwidth restrictions which are faced in practical design and which give rise to the conflict studied in Chapter 8. Further study of coupled instability would be valuable as the passivity criterion we have imposed in haptic rendering is a strong condition that may be overly conservative. With a rough frequency-domain characterization of possible user dynamics, passivity requirements could be relaxed. Experimental study of coupled instability which arises in a continuous LTI framework would fill a significant gap in the study of human-in-the-loop feedback systems.

9.3 Future Work

Parallels between the virtual coupler and teleoperator feedback design suggest that the results developed originally for teleoperation may apply to haptic rendering with the virtual coupler structure. The novel parameterization of teleoperator feedback designs introduced in Chapter 5 was key to proving bounds on performance and providing valuable design interpretation to the transparency diagram. The actions of the virtual coupler might be succinctly captured by a gain parameter, a compensated haptic device parameter, a virtual slave parameter, and an asymmetry parameter. Parameterized in this fashion, the transparency diagram could be used as a tool for tuning the virtual coupler and predicting performance. However, the practical utility of viewing the virtual coupler design through a teleoperation framework requires further study.

The design and analysis of the teleoperator control system presented in Chapter 5 and 6 does not address the closed-loop stability of the free-free dynamics nor coupled stability when the environment and human operator are in the loop. A reasonable approach to ad-

dress closed-loop stability is to view position-position teleoperator feedback system as an inner/outer loop structure, where the terms C'_m and C'_s close inner loops around the master and slave. For coupled stability, we may require that the teleoperator's free-free dynamics be passive. Performance limitations and tradeoffs imposed by stability and passivity can now be studied quantitatively using distortion.

Force sensing has not been included in our analysis of haptic rendering and teleoperation as they are not frequently used on impedance-type devices and the force sensors have drawbacks including cost, calibration, and fragility. However, if intrinsic hardware dynamics are already minimized (for example due to strength requirements) force sensing may overcome certain limitations of position-based feedback design. For instance, the tradeoff between performance and sensitivity in haptic rendering may be circumvented with force sensing as the sensitivity function we have analyzed is associated with the position feedback loop. It is interesting to note, however, that the conflict between performance and passivity in haptic rendering is apparently intrinsic to impedance-type devices and not the sensor suite. Thus compensating for hardware dynamics with a force sensor may necessarily come at the cost of coupled stability. The strong parallels between haptic rendering and shaping master and slave dynamics suggests that similar conclusions likely apply to feedback compensation in teleoperator feedback design.

Bibliography

- Abbott, J. J. and Okamura, A. M. (2005). Effects of position quantization and sampling rate on virtual-wall passivity. *IEEE Trans. Robot.*, 21(5):952–64.
- Adams, R. J. and Hannaford, B. (1999). Stable haptic interaction with virtual environments. *IEEE Trans. Robot. & Automat.*, 15(3):465–74.
- Adams, R. J., Moreyra, M. R., and Hannaford, B. (1998). Stability and performance of haptic displays: Theory and experiments. In *Proc. ASME Int'l Mechan. Engin. Congr. and Expos., DSCD*, volume 64, pages 227–34.
- Alvarez-Gallegos, J., Rodriguez, D. C., and Spong, M. W. (1997). A stable control scheme for teleoperators with time delay. *IEEE Trans. Robot. Automat.*, 12(3).
- Anderson, R. J. and Spong, M. W. (1989). Bilateral control of teleoperators with time delay. *IEEE Trans. Automat. Contr.*, 34(5):494–501.
- Baxter, J. (1988). Analysis of stiffness and feel for a power-assisted rack and pinion steering gear. Technical Report SAE 880706.
- Bode, H. W. (1945). *Network Analysis and Feedback Amplifier Design*. Nostrand, New York.
- Bretz, E. A. (2001). By-wire cars turn the corner. *IEEE Spectrum*, 38(4):68–73.
- Buttolo, P. and Hannaford, B. (1995). Pen-based force display for precision manipulation in virtual environments. In *Proc. Virtual Reality Annual Inter. Symp.*, pages 217–24, Research Triangle Park, NC.
- Çavuşoğlu, M. C., Feygin, D., and Tendick, F. (2002). A critical study of the mechanical and electrical properties of the PHANToM haptic interface and improvements for high-performance control. *Presence*, 11(6):555–68.
- Clement, G., Vertut, J., and Fournier, R. (1985). An overview of CAT control in nuclear services. In *Proc. IEEE Int'l Conf. on Robot. and Automat.*, pages 713–18, St. Louis, MO. IEEE Computer Society.

- Colgate, E. and Hogan, N. (1989). An analysis of contact instability in terms of passive physical equivalents. In *Proc. IEEE Int'l Conf. Robot. & Automat.*, volume 1, pages 404–9, Scottsdale, AZ. IEEE Comp. Soc.
- Colgate, J. E. and Schenkel, G. G. (1997). Passivity of a class of sampled-data systems: application to haptic interfaces. *J. Robotic Sys.*, 14(1):37–47.
- Colgate, J. E., Stanley, M. C., and Brown, J. M. (1995). Issues in the haptic display of tool use. In *Proc. IEEE/RSJ Int'l Conf. on Intell. Robots and Contr.*, pages 140–45, Pittsburgh, PA.
- Diolaiti, N., Niemeyer, G., Barbagli, F., and Salisbury, J. K. J. (2006). Stability of haptic rendering: Discretization, quantization, time delay, and Coulomb effects. *IEEE Tran. Robot.*, 22(2):256–68.
- Dixon, J. C. (1991). *Tyres, Suspension and Handling*. Cambridge University Press, Cambridge, UK.
- Fardad, M. and Bamieh, B. (2004). A frequency domain analysis and synthesis of the passivity of sampled-data systems. In *Proc. 43rd IEEE Conf. on Decision and Contr.*, volume 3, pages 2358–63, Nassau, Bahamas.
- Faulring, E. L., Colgate, J. E., and Peshkin, M. A. (2006). The cobotic hand controller: Design, control and performance of a novel haptic display. *Int'l J. Robot. Research*, 25(11):1099–119.
- Fite, K. B., Speich, J. E., and Goldfarb, M. (2001). Transparency and stability robustness in two-channel bilateral telemanipulation. *JDSMC*, 123:400–7.
- Freudenberg, J. S., Hollot, C. V., Middleton, R. H., and Tsochinda, V. (2003). Fundamental design limitations of the general control configuration. *IEEE Trans. Automat. Contr.*, 48(8):1355–70.
- Gillespie, R. B. and Cutkosky, M. R. (1996). Stable user-specific haptic rendering of the virtual wall. In *Proc. ASME DSCD*, volume 58, pages 397–406, Atlanta, GA.
- Gillespie, R. B., Hoffman, M. B., and Freudenberg, J. (2003). Haptic interface for hands-on instruction in system dynamics and embedded control. In *Proc. 11th Symp. on Haptic Interfaces for Virtual Environ. and Teleop. Sys.*, pages 410–15, Los Angeles, CA.
- Gillespie, T. D. (1992). *Fundamentals of Vehicle Dynamics*. SAE, Warrendale, PA.
- Griffiths, P. G. and Gillespie (2008). A fundamental conflict between performance and passivity in haptic rendering. to be presented at the *American Control Conference*.
- Griffiths, P. G., Gillespie, R. B., and Freudenberg, J. S. (2007). A fundamental trade-off between performance and sensitivity in haptic rendering. In *Robotics: Science and Systems*, Atlanta, GA.

- Griffiths, P. G., Gillespie, R. B., and Freudenberg, J. S. (2008a). Characterizing teleoperator behavior for feedback design and performance analysis. In *Symp. on Haptic Interf. for Virtual Environ. and Teleop. Sys.*, Reno, NV.
- Griffiths, P. G., Gillespie, R. B., and Freudenberg, J. S. (in press, 2008b). A fundamental tradeoff between performance and sensitivity in haptic rendering. *IEEE Trans. Robot.*
- Hashtrudi-Zaad, K. and Salcudean, S. E. (2001). Analysis of control architectures for teleoperation systems with impedance/admittance master and slave manipulators. *Int'l J. Robot. Research*, 20(6):419–45.
- Haykin, S. S. (1970). *Active Network Theory*. Addison-Wesley, London.
- Huang, F. C., Gillespie, R. B., and Kuo, A. D. (2006). Human adaptation to interaction forces in visuo-motor coordination. *IEEE Trans. Neural Sys. and Rehab. Engineer.*, 14(3):390–97.
- Iwata, H. (1990). Artificial reality with force-feedback: Development of desktop virtual space with compact master manipulator. *Computer Graphics*, 24(4):165–70.
- Khalil, H. K. (2002). *Nonlinear Systems*. Prentice Hall, Upper Saddle River, NJ, 3rd edition.
- Kim, W. S., Hannaford, B., and Bejczy, A. K. (1992). Force-reflection and shared compliant control in operating telemanipulators with time delay. *IEEE Trans. Robot. Automat.*, 8(2).
- Krahe, C. (1996). Airbus fly-by-wire aircraft at a glance. *FAST Airbus Technical Digest*, 20:2–9.
- Lawrence, D. A. (1993). Stability and transparency in bilateral teleoperation. *IEEE Trans. Robot. & Automat.*, 9(5):624–37.
- Lawrence, D. A., Pao, L. Y., Dougherty, A. M., Salada, M. A., and Pavlou, Y. (2000). Rate-hardness: A new performance metric for haptic interfaces. *IEEE Trans. Robot.*, 16(4):357–71.
- Laycock, S. D. and Day, A. M. (2003). Recent developments and applications of haptic devices. *Comp. Graphics Forum*, 22(2):117–32.
- Lee, C. D., Lawrence, D. A., and Pao, L. Y. (2000). A high-bandwidth force-controlled haptic interface. In *Proc. ASME Inter. Congr. and Expos., Symp. on Haptic Inter. for Virt. Environ. and Teleop. Systems*, Orlando, FL.
- Lee, D. and Spong, M. W. (2006). Passive bilateral teleoperation with constant time delay. *IEEE Trans. Robot.*, 22(2):269–81.
- Leung, G. M. H., Francis, B. A., and Apkarian, J. (1995). Bilateral controller for teleoperators with time delay via μ -synthesis. *IEEE Trans. Robot. Automat.*, 11(1):105–16.

- Li, M. and Liu, Y. H. (2007). Dynamic modeling and experimental validation for interactive endodontic simulation. *IEEE Trans. Robot.*, 23(3):443–58.
- Llewellyn, F. B. (1952). Some fundamental properties of transmissions systems. *Proc. IRE*, 40(3):271–83.
- Mahvash, M. and Hayward, V. (2005). High-fidelity passive force-reflecting virtual environments. *IEEE Tran. Robot.*, 21(1):38–46.
- Massie, T. H. and Salisbury, J. K. (1994). The PHANTOM haptic interface: A device for probing virtual objects. In *Proc. ASME Winter Annual Meeting, Symp. on Haptic Inter. for Virt. Environ. and Teleop. Systems*, Chicago, IL.
- McJunkin, S. T., O’Malley, M. K., and Speich, J. E. (2005). Transparency of a Phantom Premium haptic interface for active and passive human interaction. In *Proc. American Contr. Conf.*, volume 5, pages 3060–65, Portland, OR.
- Mehling, J. S., Colgate, J. E., and Peshkin, M. A. (2005). Increasing the impedance range of haptic display by adding electrical damping. In *Proc. First Joint Eurohaptics Conf. and Symp. on Haptic Inter. for Virtual Environ. and Teleoperator Systems*, pages 257–62, Pisa, Italy.
- Miller, B. E., Colgate, J. E., and Freeman, R. A. (2000). Guaranteed stability of haptic systems with nonlinear virtual environments. *IEEE Trans. Robot. & Automat.*, 16(6):712–19.
- Minsky, M., Ouh-Young, M., Steele, O., Brooks, F. P., and Behensky, M. (1990). Feeling and seeing: Issues in force display. In *Proc. Symp. on Interactive 3D Graphics*, volume 24, pages 235–41, Snowbird, UT.
- Niemeyer, G. and Slotine, J. E. (2004). Telemanipulation with time delays. *Int’l J. Robot. Research*, 23(9):873–90.
- Odenthal, D., Bünte, T., Heitzer, H., and Eicker, C. (2002). How to make steer-by-wire feel like power steering. In *Proc. 15th IFAC World Congr. on Automatic Control*, Barcelona, Spain.
- Okamura, A. M. (2004). Methods for haptic feedback in teleoperated robot-assisted surgery. *Industrial Robot*, 31(6):499–508.
- Pan, Y., Canudas-de Wit, C., and Sename, O. (2006). A new predictive approach for bilateral teleoperation with applications to drive-by-wire systems. *IEEE Trans. Robot.*, 22(6):1146–61.
- Patton, J. L., Kovic, M., and Mussa-Ivaldi, F. A. (2006). Custom-designed haptic training for restoring reaching ability to individuals with poststroke hemiparesis. *Journal of Rehabilitation*, 43(5):643–55.

- Seron, M. M., Braslavsky, J. H., and Goodwin, G. C. (1997). *Fundamental Limitations in Filtering and Control*. Springer-Verlag, London.
- Shen, X. and Goldfarb, M. (2006). On the enhanced passivity of pneumatically actuated impedance-type haptic interfaces. *IEEE Trans. Robot.*, 22(3):470–80.
- Skogestad, S. and Postlethwaite, I. (1997). *Multivariable Feedback Control: Analysis and Design*. Wiley, New York.
- Slotine, J. E. and Li, W. (1991). *Applied Nonlinear Control*. Prentice Hall, Englewood Cliffs, NJ.
- Sorid, D. and Moore, S. K. (2000). The virtual surgeon. *IEEE Spectrum*, 37(7):26–31.
- Stramigioli, S., van der Schaft, A., Maschke, B., and Melchiorri, C. (2002). Geometric scattering in robotic telemanipulation. *IEEE Trans. Robot. and Automat.*, 18(4):588–96.
- Vlachos, K., Papadopoulos, E., and Mitropoulos, D. (2003). Design and implementation of a haptic device for training in urological operations. *IEEE Trans. Robot. and Automat.*, 19(5):801–9.
- Vuskovic, V., Kauer, M., Szekely, G., and Reidy, M. (2000). Realistic force feedback for virtual reality based diagnostic surgery simulators. In *Proc. IEEE Int'l Conf. Robot. and Automat.*, volume 2, pages 1592–8, San Francisco, CA. IEEE.
- Wang, D. X., Zhang, Y., and Wang, Y. (2003). Development of dental training system with haptic display. In *Proc. 12th IEEE Int'l Workshop on Robot and Human Interact. Comm.*, pages 159–64, Millbrae, CA. IEEE.
- Yan, J. and Salcudean, S. E. (1996). Teleoperation controller design using H_∞ -optimization with application to motion-scaling. *IEEE Trans. Contr. Syst. Technol.*, 4(3):244–58.
- Yao, Y. (2006). Vehicle steer-by-wire system control. Technical Report SAE 2006-01-1175.
- Yokokohji, Y. and Yoshikawa, T. (1994). Bilateral control of master-slave manipulators for ideal kinesthetic coupling—formulation and experiment. *IEEE Trans. Robot. Automat.*, 10(5):605–20.