Chapter 1

Introduction

This thesis addresses the incorporation of computer-generated haptic feedback into interfaces for computer-based musical instruments. In doing so, it draws upon a wide range of topics including haptic perception, haptic simulation, music synthesis and control, human motor control, and human motor skill acquisition.

The term "haptic," derived from the Greek word "haptesta" (to touch), refers to combined feedback from tactile sensors in the skin and kinesthetic sensors in muscles and joints. Though spread throughout our bodies, tactile and kinesthetic sensors are most concentrated in our hands and lips. It is no accident, therefore, that musicians are acutely aware of an instrument's "feel," since the actions of blowing, bowing, plucking, pressing, and tapping used to play most instruments are carried out by hands and lips.

The fundamental principle explored in this work is that haptic feedback can support auditory feedback to inform performers of the consequences of their actions, playing a crucial role in the performer/instrument interaction loop. Where it exists, a mechanical coupling between player and instrument provides a secondary sensory feedback channel through which much information about an instrument's state can be monitored.

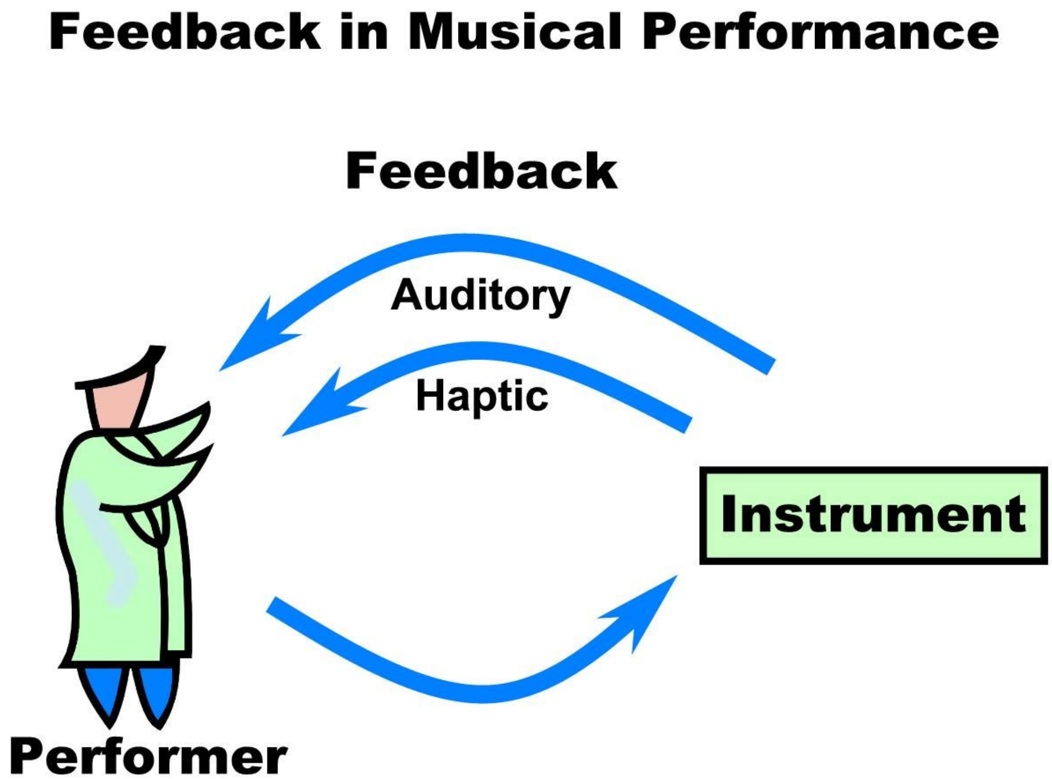

Fig. 1.1. Haptic and auditory feedback paths in the musician/instrument interaction loop.

Leveraging the musician's existing sensitivity to the relationship between an instrument's feel and its sound, this work explores the implications of incorporating haptic feedback into interfaces and controllers for computer-based musical instruments, in other words, the implications of closing a feedback loop between instrument and performer via a second sensory modality, touch. This fundamental principle is the motivating force behind all the theoretical and experimental work presented here.

1.1 Control of Computer-Based Musical Instruments

Music performance can be described at its most basic level as the evolution in time of four parameters — time, pitch, dynamics (or amplitude), and timbre. For the performer, the manipulation of these dimensions is embodied in physical actions such as striking a piano key or bowing a string. In real instruments these "instrumental gestures" are determined by the sound-producing mechanisms of the instrument. The sound produced, in turn, carries the characteristics of the movement that gave rise to it (Cadoz, 1988). A by-product of sound synthesis techniques such as physical modeling (Smith, 1998) is the ability to decouple the synthesis of an instrument's sound from the physics of the instrument's sound-producing mechanism. Thus the affordances of a synthetic music controller can be very different from those of the instrument being controlled. A piano keyboard, for example, might be used to control a physical model of a trumpet. The advantage of such a modular approach is that a player can perform on instruments with no knowledge of specific playing technique. This is not to say that the player lacks skill, for to be successful it is necessary to maximize the capabilities of the controller being used. More crucial to the success of the paradigm than the player's skill, however, is the mapping between physical gesture and sound synthesis parameters, the relationship between the player's actions and the instrument's response.

It is an unfortunate accident of history that early music controllers were predominantly based on the piano keyboard. For the piano, more than for any other instrument (excepting percussion instruments), there exists a one-to-one mapping between a note and the movement that produced it. Thus the subsequent development of MIDI (Musical Instrument Digital Interface), the serial protocol by which controllers and synthesizers communicate, is predicated on the assumption that each note is an isolated event with controls for its pitch, duration, timbre and amplitude, that cannot interact. For most musical instruments, this one-to-one mapping is the exception rather than the rule. For bowed instruments, for example, many notes of a single slurred phrase will be executed with one bow stroke and hence a single arm movement. Not only does this single arm movement cause notes of a phrase to be linked by a common bow stroke, but its trajectory also embodies expressive nuance, shaping the dynamic and timbral arc of the phrase. Parameters for dynamics and tone color are thereby correlated and co-vary in response to the trajectory of the movement of the arm. To capture such nuance in performance and convey its subtleties to a synthesis model requires a control architecture where the gesture, not the note, is the primary unit of musical time.

Several researchers (most notably Cadoz, 1988; Chafe, 1988; McMillen, 1994) have proposed that music controllers and the protocol that supports their communication with synthesis algorithms should be founded on an asynchronous hierarchical structure with the performance gesture, not the score-based note list, as its unit of currency. As McMillen (1994) points out, players of non-keyboard instruments have been reluctant to embrace the digital revolution in music. Woodwinds, bowed strings, and brass instruments all place the player in direct physical contact with the vibrating element — reeds, strings, or columns of air — providing the player with fine control of a single note or a small group of notes. Most commercially available control devices provide limited control of multiple notes and are inappropriate for most melodic instruments. Faced with trading fine control of a real instrument for the infinite timbral possibilities but coarse control of today's synthesis models, most players opt to remain in the real world. Even in those cases where real instruments are adapted to transmit MIDI messages, or so-called hybrid controllers, the limitations of MIDI still present a significant bottleneck.

In order to better understand the demands of real-time control of non-keyboard instruments, Chafe (1988) developed a control language and supporting architecture to control a physical model of a bowed string using simulated performance gestures. The system's input consisted of lists of descriptors for physical actions such as finger placement and bowing gesture (staccato, martello, etc). Insofar as the input "score" describes musical events in terms of the position of the player's hands on the instrument, rather than as abstract notes on a stave, the method more closely resembles tablature notation than the staff notation of a conventional musical score. Rather than representing gestures in actual physical terms, the system represents gestures in terms of their effect on the string. There is therefore no notion of the extent or rate of motion of the bow or the movements of the fingers. Each hand is represented separately in the score, and a supra-hand scheduler organizes their interaction, correlating parameter trajectories that are coupled by the same physical gesture. This hierarchical structure that can simultaneously represent both low-level movement trajectories and higher-level musical articulation trajectories provides an ideal model on which to build a control architecture for gesture-driven interaction.

In all previous work on music controllers, only Gillespie (1996) and Cadoz (1988) have considered the possibility that the instrument's "feel" plays a key role in a player's ability to control it. Do performers gravitate toward hybrid controllers because they provide a tight coupling between their actions or gestures and the instrument's response, because they simply "feel" right? Certainly for Cadoz, the instrumental gesture, with its associated haptic feedback, is the musician's way of apprehending and being fully conscious of the sound-producing object. This consciousness is an integral part of the musician's mastery of the instrument.

If, as Cadoz suggests, the musician's understanding of the instrument's behavior is tightly coupled to its haptic response, then adding appropriate haptic feedback to computer-based musical instrument controllers should greatly enrich performer/instrument interaction and restore some level of support for the instrumental gesture and hence finer control. However, for haptic feedback to be useful and meaningful, it is necessary to understand its role in the musician's mental representation of an instrument's behavior. This work presents a new perspective on the design of music controllers, by focusing on movement control and its associated sensory feedback as the primary factors in the design of a series of new computer-based musical instruments.

1.2 Scope of Thesis

The implications of closing the loop by providing haptic feedback about an instrument's state directly to a player's hands or lips are far-reaching. Firstly, we need to understand how the human sensorimotor system gathers knowledge about the affordances of its environment, i.e., how a player learns to play a new musical instrument. Secondly, we need to design haptic interactions that make sense within the context of a given task.

Currently, no formal methodology exists to investigate the processes by which a player builds an internal representation of the behavior of an instrument. Sensor technology is not sufficiently advanced to allow for unobtrusive monitoring of a player's movements and the vibrations and reaction forces generated in response to these movements. In studying the potential role for haptic feedback in music controllers, therefore, this work borrows concepts of experimental design and performance evaluation from the fields of manual skill acquisition, telepresence and virtual reality research. In doing so, this dissertation attempts to formulate a theoretical foundation for the design of haptic feedback for digital musical instruments.

1.3 Research Contributions

The primary contribution of this work to the field of computer music research is a new approach to the design of computer-based musical instruments which takes as its starting point the human sensorimotor system. This approach differs from previous research in that the movements required to execute notes and sequences of notes, in conjunction with their associated sensory feedback, are considered to be the vehicles of musical expression.

Because performance gestures are inextricably bound up with the dynamics of the instrument being played, this work explores the processes by which the musician builds an internal representation of an instrument's behavior. Experimental data are presented that support the hypothesis that haptic feedback is a part of this internal representation. Further, this work provides evidence that computer-generated haptic feedback can potentially support the transfer of skill from real to virtual environments. Finally, this work contributes techniques for constructing haptic feedback for two classes of computer-based musical instruments: those that are entirely novel, and those that simulate the feel of real-world instruments. In addition to detailing the design of these instruments, this work suggests some goals for hardware and software music control protocols that must be met if haptic feedback is to become an integral part of controllers for computer-based musical instruments.

1.4 Outline of this Dissertation

The remaining chapters of this dissertation are organized as follows: Chapter 2 reviews previous research in haptic feedback for music applications and presents findings from the literature of haptic interaction design relevant to the present work. Studies on the role of motor learning in acquisition of musical skill are also reviewed. Chapter 3 draws upon the body of research in human manual skill acquisition to posit a role for haptic feedback in the musician/instrument interaction loop. In particular, this chapter suggests a changing role for sensory feedback in the process of learning to play a new musical instrument. Chapters 4 and 5 present the four experimental studies that are the core of this research. Experiments I and II explore mappings between audio and haptic responses for a new musical instrument, while experiments III and IV simulate the feel of an existing instrument to examine the role haptic cues play in the learning process. Chapter 6 summarizes the results of these studies and discusses their relevance to the design of computer-based musical instrument controllers. Chapter 6 also outlines some of the hardware and software goals that need to be met in order to provide performers, composers, and musical instrument designers with the tools they need to make haptic feedback in computer-based musical instruments a reality.