

# Using Haptics to Assist Performers in Making Gestures to a Musical Instrument

**Edgar Berdahl**  
CCRMA  
Stanford University  
Stanford, CA, USA

**Günter Niemeyer**  
Telerobotics Laboratory  
Stanford University  
Stanford, CA, USA

**Julius O. Smith III**  
CCRMA  
Stanford University  
Stanford, CA, USA

## Abstract

Haptic technology, providing force cues and creating a programmable physical instrument interface, can assist musicians in making gestures. The finite reaction time of the human motor control system implies that the execution of a brief musical gesture does not rely on immediate feedback from the senses, rather it is preprogrammed to some degree. Consequently, we suggest designing relatively simple and deterministic interfaces for providing haptic assistance.

In this paper, we consider the specific problem of assisting a musician in selecting pitches from a continuous range. We build on a prior study by O'Modhain of the accuracy of pitches selected by musicians on a Theremin-like haptic interface. To improve the assistance, we augment the interface with programmed detents so that the musician can feel the locations of equal tempered pitches. Nevertheless, the musician can still perform arbitrary pitch inflections such as glissandi, falls, and scoops. We investigate various forms of haptic detents, including fixed detent levels and force-sensitive detent levels. Preliminary results from a subject test confirm improved accuracy in pitch selection brought about by detents.

**Keywords:** Haptic, detent, pitch selection, human motor system, feedback control, response time, gravity well

## 1. Introduction

Haptics is an important application of feedback control to musical instrument design. Kinesthetic haptic feedback can alter the gestures made by musicians during a live performance. In particular, it can assist musicians in making certain types of gestures. For example, Figure 1 depicts a cellist assisted by two SensAble<sup>TM</sup> PHANTOM<sup>®</sup> Desktop haptic devices.<sup>1</sup> Each haptic device would exert forces on one end

<sup>1</sup> PHANTOM, PHANTOM Desktop, SensAble, and SensAble Technologies, Inc. are trademarks or registered trademarks of SensAble Technologies, Inc.



Figure 1. Cellist assisted by two haptic devices

of the bow. By way of feedback control, the physical properties of the device could be changed and even non-physical behaviors could be created, altering the way in which the musician would interact with the bow.

The idealized configuration shown in Figure 1 would be rather involved, so we simplify it in our laboratory. We virtualize the instrument, allowing the musician to interact with it by way of a single haptic device [2]. A *haptic musical instrument* consists of actuators that exert forces on the musician, sensors that measure the response of the musician, and a programmable controller that determines appropriate forces to exert on the musician. Figure 2 illustrates how the musician is included in the feedback loop, allowing the musician's gestures to be controlled. The controller also synthesizes sound signals in response to the way in which the haptic musical instrument is played. We employ the Model T PHANTOM desktop device because it has a relatively large workspace and provides high resolution position measurements. The end of the arm is attached to a thimble, into which a musician can insert his or her index finger (see Figure 3).

Because the haptic device is controlled by a digital feedback system, it can be programmed to behave like any of a wide array of dynamical systems. For instance, it can be programmed to behave like a physical acoustic instrument, such as a bowed string instrument. On the other hand, it can also be programmed to behave like an instrument with an

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.  
NIME09, June 3-6, 2009, Pittsburgh, PA  
Copyright remains with the author(s).




Figure 2. Musician playing a haptic musical instrument




Figure 3. Musician's finger inserted into PHANTOM thimble

internal energy source, or it can be programmed to assist the musician in some way, etc. Clearly haptic musical instrument designers enjoy an amazing amount of freedom, and for this very reason, we investigate some designs for assisting musicians as we expect these designs to be especially useful in practice.

## 2. Assisting the Human Motor System

Both closed-loop and open-loop control models provide distinct perspectives from which the human motor system can be studied. Humans use *closed-loop control* for completing fine tasks such as threading a needle, where the control loop is closed around decision-making centers in the brain. The sequence of actions is as follows: the human attempts to move the end of the thread, waits for feedback from the senses about the new position of the thread, attempts to refine the position of the thread, waits for sensory feedback again, etc. The speed of closed-loop control of the human motor system is affected by the reaction time (RT) of the human motor system. The RT describes the time interval from the sudden presentation of an unanticipated force signal to the beginning of the human motor response, as directed by a signal returning from the brain [10]. Due to the relatively long motor system RT of roughly 120ms-180ms, the RT feedback mechanism is apparently not used for controlling brief tasks, such as for playing trills at up to 20Hz [7]. The reaction times for other modalities (auditory, visual, etc.) have similar durations, so they cannot be used to significantly speed up motor system responses either.

According to the *open-loop control* theory of the human motor system, once a human makes the decision to carry out a brief movement, a motor program is called up, which sequentially issues commands to the muscles [10]. The way in which the commands are issued may depend on feedback, but the feedback in this loop is limited by the RT. There is evidence that many fast portions of tasks involved in playing a musical instrument, such as rapidly pressing a sequence of keys, are governed primarily by open-loop control. For instance, Schmidt discusses the “example of a skilled pianist playing a piano with a broken key that could not be depressed. As the pianist played a string of notes, the attempts to press the broken key did not interrupt the series of actions at all. In fact, only after the entire sequence was completed did the individual notice and remark that the key was broken” [10]. Since it took the pianist longer to notice the broken key than the duration of a single note in the sequence, it seems likely that the pianist must have been playing the broken key using some form of open-loop control. Skilled typists [11] and telegraphers [1] are also believed to use open-loop control for completing motor tasks rapidly [7].

If an assistive haptic interface ever takes an unexpected action too quickly, then a musician who uses some elements of open-loop control will not be able to respond appropriately, possibly making a mistake. Hence, we argue that *an assistive haptic interface should not take an action that the musician does not expect, unless the interface takes the action so slowly that the musician has time to react*. In practical contexts, we recommend simply that an *assistive haptic interface be deterministic and relatively simple so that no actions surprise the musician*. We now shift gears slightly to focus on a more specific problem.

## 3. Pitch Selection Problem Overview

A musical tone can be described as having a distinct pitch, loudness, and timbre [15]. We study how haptic devices can assist musicians in selecting pitches accurately. Consider the different ways in which pitches are selected while playing the piano versus playing the cello. The piano has 88 strings of fixed length, so a pianist using standard playing techniques can play only a discrete set of pitches. If the strings are tuned consistently, then no matter which keys the pianist presses, the performance will be approximately in tune. In contrast, the length of a cello string can be adjusted directly by the musician's hand, meaning that a cellist can play pitches over a continuous range. This freedom can be troublesome for beginning cellists who might for example wish to initially limit themselves to playing a discrete set of pitches corresponding to a scale. In fact, this freedom can also allow vocalists, trombonists, violinists, viola players, string bass players, etc. to unintentionally play out of tune.

In a broad sense, we seek to combine the way that pitches are selected on a piano and on a cello in order to make it

easy to play pitches accurately while still allowing arbitrary continuous pitch inflections to be created. More specifically, we consider using haptic feedback to improve the accuracy with which a musician can select a pitch over a continuous range. We design the haptic assistance to help guide the musician to desirable pitches. Of course, musicians may disagree over the meaning of the word *desirable*, so for simplicity, we restrict ourselves to more traditional music and the equal temperament tuning system.<sup>2</sup> Then given a scale, only a discrete set of pitches are considered desirable. Despite the haptic assistance, the musician should still be able to play pitch bends, glissandi, etc. to take advantage of the additional ranges of expression afforded by continuous pitch control.

#### 4. Prior Work on Pitch Selection Problem

The Theremin electronic instrument, which was patented in the United States in 1928, produces a harmonic tone as output [12]. The pitch of the tone is controlled by the position of one hand in free space, while the amplitude of the tone is controlled by the position of the other hand. The Theremin instrument provides no haptic feedback, and it is often informally considered to be difficult to play as it is non-trivial for the musician to orient his or her hands in free space.

O’Modhrain studied the accuracy with which musicians select pitches with a Theremin-like interface implemented using a haptic device. She compared controlling pitch given the complete absence of haptic feedback versus several kinesthetic haptic feedback conditions such as a spring force, a viscous damping force, and a constant force. In chapter 4 of her PhD thesis, she draws the conclusion that the “existence of force feedback in a computer-based instrument marginally improves performance of a simple musical [pitch selection] task” (p. 49) [7].

Besides Theremin-type interfaces, glove-based and other continuous interfaces may lack significant haptic feedback. We generalize O’Modhrain’s conclusion by hypothesizing that if a musical instrument does not provide any haptic feedback at all, it will probably be more difficult to play accurately. We term this hypothesis the “*Theremin hypothesis*.” As a consequence, we recommend that musical instrument designers incorporate haptic feedback into their instrument designs.

We hypothesize further that specific kinds of active force feedback assisting the musician in selecting desirable pitches may be even more effective. This hypothesis is also suggested by Moss’ and Cunitz’s work in which a specific kind of haptic feedback pushes the musician’s finger toward the notes of the chromatic scale; however, Moss and Cunitz do not consider any other types of haptic feedback [4].

<sup>2</sup> We acknowledge that the equal temperament tuning system is only approximate [15]. Consequently, we choose pitch selection tasks difficult enough that tuning ambiguities are considerably smaller than task errors.




Figure 4. Force profile  $F(y)$  and terrain height profile  $h(y)$

### 5. Simple Pitch-Controlled Instrument

We have strived to make our example instrument as simple and easy to play as possible. Although the PHANTOM Model T can move in three dimensions, we restrict motion of the thimble to a linear axis and simultaneously measure pressure applied normal to this axis in the vertical direction. For simplicity, we allow the musician to adjust only the pitch of the sound. Due to the finite workspace size, the horizontal position  $y$  (see Figure 3) of the musician’s finger can vary over about 20cm and is mapped to the logarithm of the pitch of the musical instrument. This mapping allows the distance between each pair of adjacent notes in the chromatic scale to be about 0.75cm. Higher pitches are further to the right.

### 6. Assistive Haptic Feedback

We consider controllers that exert force in the lateral  $y$ -axis as a function of the current and past  $y$ -positions of the thimble (see Figure 3). Many forms of haptic feedback are imaginable, so we limit these forms here by considering the literature and the guidelines developed in section 2.

#### 6.1. Basic Detent

The detent is simple, deterministic, and can help the musician orient himself or herself. A detent can be created even using 1DOF haptic interfaces. Figure 4 illustrates how to implement a simple piecewise linear detent. Near the center of the detent, the force in the lateral  $y$ -axis behaves like that of a spring, while the force goes to zero when the position  $y$  moves further from the detent’s center [13]. The force profile is consistent with the lateral forces one would experience in the terrain height map  $h(y)$  shown at the bottom of in Figure 4.

Researchers from the human computer interaction (HCI) literature have experimented with piecewise linear as well as piecewise nonlinear detents [4]. However, they report that altering the shape of the detent nonlinearly manifests itself psychophysically as apparently only an intensity difference [14][8]. We have also carried out some informal tests and also believe that the details of the detent shape are not of primary importance, so we believe the piecewise linear detent design to be sufficient for our study.

#### 6.2. Force Feedback Conditions

We now describe the specific haptic force feedback conditions that we incorporated into our subject test.




Figure 5. Force profile for DET condition

### 6.2.1. Multiple Detents (DET)

We extend the basic detent described in section 6.1 to assist the musician in playing notes from a diatonic scale consisting of whole and half steps, which is based at the origin of the haptic device. Figure 5 shows the piecewise linear force profile for the first five notes of the scale. A spring force field is centered around each note, making the  $y$ -position a locally stable equilibrium point, as denoted by each dashed blue circle [9]. In between each pair of notes, the forces must be tapered toward each other to avoid creating any distracting discontinuities. Since there is a half-step instead of a whole-step between *mi* and *fa*, the force profile is warped to retain the same form (see Figure 5). Note that we have made the detent force profiles as wide as possible to ensure that we are providing haptic assistance for the largest possible set of positions  $y$ .

### 6.2.2. Force-Sensitive Detents (FRC)

While the HCI literature is in agreement that a single detent centered around a single target, such as a menu item, improves performance, it may be harder for interface designers to improve performance when faced with multiple possible targets, each having its own detent, especially when the user must traverse through distracting detents in order to reach a specific target [6][5]. We hypothesize that by considering notes to be the analog of menu targets, we might take advantage of this feature to improve upon the DET condition. We design a novel force condition that has the advantage that the musician can actively adjust the strength of the haptic assistance. In other words, the musician may control the assistance so that he or she only receives it when he or she requests it by pressing downward with force  $p > 0$ . If the musician instead pulls the thimble even slightly upward (i.e.  $p < 0$ ), then no haptic assistance will be provided at all, allowing the musician to breeze over any distracting detents. We define  $F_{FRC}(y, p)$  as follows:

$$F_{FRC}(y, p) = F_{DET}(y) \cdot \min(lp, M_0) \cdot (p > 0), \quad (1)$$

for some constants  $l > 0, M_0 > 0$ . The parameter  $M_0$  serves as the maximum detent gain, which prevents the detents from becoming unreasonably strong if the musician ever presses down especially hard.

### 6.2.3. Spring Force (SPR)

For the sake of comparison with O’Modhrain’s experiment, we introduce also the SPR condition, where  $F_{SPR}(y)$  de-

scribes a spring. The virtual spring center is placed near the edge of the workspace so that the spring always exerts a force to the left.

### 6.2.4. No Feedback (NOFB)

We finally introduce a control condition, in which the force in the  $y$ -dimension is always zero:

$$F_{NOFB}(y) = 0. \quad (2)$$

Due to O’Modhrain’s work and the Theremin hypothesis, we expect this condition to be the most frustrating for musicians [7].

## 7. Subject Test

### 7.1. Experiment Design

To verify the efficacy of haptic assistance, we are conducting a subject test to compare how accurately subjects can select pitches under the DET, FRC, SPR, and NOFB force conditions. The experiment design was motivated by the desire to obtain data that could be compared across subjects, even if some learning effects could possibly be seen in the data. We recruited seven subjects from the Stanford Symphonic Orchestra and six subjects from the M.S. program at the Center for Computer Research in Music and Acoustics (CCRMA) at Stanford University. Of the thirteen subjects total, one was left-handed and four were female.

The total time commitment required of each subject was about one hour on average, and each subject was compensated with \$20 for his or her efforts. Each subject was trained and tested as follows:

- At the beginning, the subject was presented with a sheet containing seven simple melody excerpts. The first excerpt was for training, while the remaining excerpts were for testing. Each excerpt contained an average of about 9 notes from the C major scale and was followed by a glissando. To help the subject learn the excerpts before being tested, the subject was asked to play them once on a standard piano keyboard-based instrument. Then the subject was asked to play them again to reinforce learning of the excerpts.
- The subject was introduced to the musical instrument described in Section 5. The subject was asked to feel each force condition using his or her dominant hand, while the operation of the force condition was explained. The subject was instructed to try to use the following strategy when using the FRC condition: the subject should press down slightly when playing small intervals and lift up slightly or maintain a neutral hand weight when playing larger intervals or glissandi.
- The process for recording performances was explained to the subject. This process modeled the procedure for recording a part of a song in a music studio and consisted of the steps 1) listen, 2) practice, 3) perform,

4) consider whether to move on. In the third step, the subject performed an excerpt for a given force condition along with a metronome track. In the fourth step, the subject could elect to re-record a performance if he or she were unsatisfied with how he or she played the given excerpt.

- Next the subject practiced using the recording process and manipulating the PHANTOM by recording his or her performance of the single training melody excerpt for each of the force conditions. Any additional questions were then answered.
- Finally the subject recorded himself or herself performing according to the four force conditions. In order to minimize the ordering effects of the subjects' being tested on one force condition before another, the force conditions were presented across subjects according to a balanced Latin square [3]. However, during the testing of each force condition, the six test melodies were presented in always the same order. We chose to do so because we believed that since the subjects learned the (very simple) melodies before performing the test, the ordering of the force conditions would affect the results much more than the ordering of the melodies. Moreover, in this configuration, it would be more straightforward to compare the performances of the melodies for a given force condition across subjects.
- The subject was asked to fill out an exit survey.

## 7.2. Data Analysis

Let  $P_s(m, c, t)$  be the MIDI note contour played as a function of time  $t$  by subject  $s$  for the  $m$ th melody excerpt and the  $c$ th force condition. For example, the fairly typical measurement  $P_5(5, c, t)$  is shown in Figure 6 for all four conditions. It can be seen that the DET and FRC force conditions tend to push the subject's finger toward notes of the C major scale such as 71, 72, and 74. In contrast, the NOFB and SPR conditions allow the subject to drift further away.

To obtain an estimate of the discrete note sequence that a subject intended to play, we started by simply quantizing the measured MIDI note contour to the nearest MIDI notes from the MIDI scale. However, some spurious notes were detected in the note sequences, especially for large note intervals. To greatly reduce the number of these artifacts, we developed an improved quantizer called *Estimated Intended Note Quantizer*, or *EINQ*, which eliminated quantized notes with durations shorter than 0.2 sec. As each short note was eliminated, it was replaced by its neighbors, where the transition time was chosen to minimize the error measure described in the following section.

To evaluate the quality of the tuning of each performance, we calculated the norm between the estimated intended note

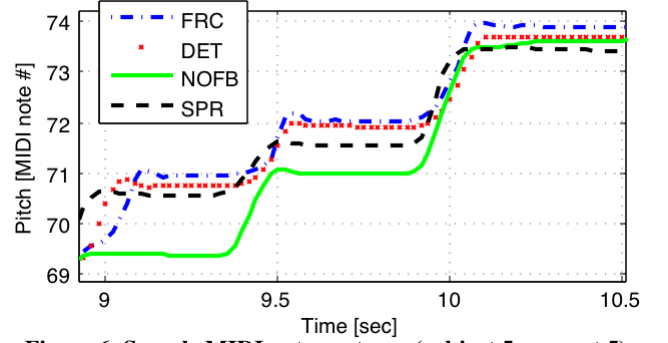


Figure 6. Sample MIDI note contours (subject 5, excerpt 5)




Figure 7. Mean absolute error for subject 5

sequence and the actual performed note sequence. We chose to use the  $L_1$ -norm instead of the  $L_2$ -norm so that the error measure did not emphasize the inevitably large error contributions stemming from note transitions, but rather focused more on the typically more constant error contributions from within notes (see the contours in Figure 6). The  $L_1$ -norm of the error, also known as the mean absolute error, from time  $T_1$  to time  $T_2$  is

$$\frac{1}{T_2 - T_1} \int_{T_1}^{T_2} |P_s(m, c, t) - \text{EINQ}(P_s(m, c, t))| dt. \quad (3)$$

$T_1$  and  $T_2$  were chosen by hand for each performance so that the time segment included only the notes prior to the glissando even if the subject made any small errors in timing. Figure 7 shows a typical set of mean absolute errors.

## 7.3. Preliminary Results

### 7.3.1. Mean Absolute Pitch Error

The clearest trend is that the DET and FRC conditions generally resulted in smaller mean absolute error than the SPR and NOFB conditions (see Figure 7 for an example). In other words, detents seem to cause the average error to be smaller. In particular, for 89% of the melody-subject pairs, the DET error was smaller than both the SPR and NOFB errors. The FRC error was also smaller than both the SPR and NOFB errors for 89% of the melody-subject pairs.

All subjects were able to perform the glissandi for all of the force feedback conditions, including the detent con-

ditions FRC and DET. In addition, some subjects learned to play glissandi smoothly with FRC by lifting upward slightly with the PHANTOM thimble.

### 7.3.2. Surveys

The thirteen subjects were asked to order the force feedback conditions from the most to the least preferable for playing glissandi accurately as well as for playing notes accurately. In agreement with section 7.3.1, twelve subjects most preferred either FRC or DET for playing notes accurately. Of these subjects, nine most preferred FRC for playing notes accurately.

Conversely, ten subjects most preferred either SPR or NOFB for playing glissandi accurately. The remaining three subjects most preferred FRC for playing glissandi accurately, presumably because these subjects had learned sufficiently well to lift the thimble up slightly while playing glissandi, so as not to get distracted by detents.

### 7.3.3. Conclusions

Haptic detents improve the accuracy with which musicians select pitches over a linear range in comparison with a simple spring force or no haptic feedback at all. In addition, more musicians seem to prefer the novel detent type FRC to DET. Indeed, musicians can even learn to play glissandi and large melodic intervals when using FRC by lifting upward slightly. Testing more subjects and carrying out a more complete statistical analysis will help us solidify our conclusions from the subject test.

## 8. Final Words

The possibility of providing musicians with haptic assistance is indeed intriguing. After taking the finite reaction time of the human motor system into account, we suggested designing assistive haptic interfaces that are deterministic and relatively simple.

Our subject test shows that detents can be used effectively to assist musicians in accurately selecting pitches over a modest linear range. The novel FRC condition appears to be the most promising. At first our conclusions might seem to be at odds with O’Modhrain’s experiment. Indeed while O’Modhrain suggested simply that some form of haptic feedback was better than none at all [7], we suggest that detents are even more helpful in assisting musicians in selecting pitches over a modest linear range. However, we think that the real lesson to learn is that *the best type of haptic assistance depends on the task*. In fact, this was the motivation behind the development of the novel force-sensitive FRC condition, where the degree of haptic assistance depends directly on user input. We believe that future haptic assistance devices should consider taking advantage of real-time user input in a similar manner, allowing users to constantly regulate the assistance that they receive.

## 9. Acknowledgments

We gratefully acknowledge Sile O’Modhrain, Carr Wilkerson, Hiroko Terasawa, Javier Sánchez, Nelson Lee, Rob Hamilton, Fernando Lopez-Lezcano, Bill Verplank, Chris Chafe, and Jonathan Berger for making this work possible.

## References

- [1] Bryan and Harter. Studies in the physiology and psychology of the telegraphic language. *Psychological Review*, 4:27–53, 1897.
- [2] C. Cadoz, A. Luciana, and J.-L. Florens. Synthèse musicale par simulation des mécanismes instrumentaux. *Revue d’acoustique*, 59:279–292, 1981.
- [3] G. Grindlay. Haptic guidance benefits musical motor learning. In *Proc. of Symposium on Haptic interfaces for Virtual Environment and Teleoperator Systems*, pages 397–404, Reno, NV, March 2008.
- [4] W. Moss and B. Cunitz. Haptic theremin: Developing a haptic musical controller using the sensible phantom omni. In *Proceedings of the International Computer Music Conference*, Barcelona, Spain, Sept. 2005.
- [5] I. Oakley. *Haptic Augmentation of the Cursor: Transforming Virtual Actions into Physical Actions*. PhD thesis, University of Glasgow, Glasgow, Scotland, May 2003.
- [6] I. Oakley, A. Adams, S. Brewster, and P. Gray. Guidelines for the design of haptic widgets. In *Proc. 16th British HCI Group Annual Conference*, pages 195–211, London, England, Sept. 2002.
- [7] S. O’Modhrain. *Playing By Feel: Incorporating Haptic Feedback into Computer-Based Musical Instruments*. PhD thesis, Stanford University, Stanford, CA, USA, 2000.
- [8] F. Picon, M. Ammi, and P. Bourdot. Force model for CAD selection. In *Proc. ACM symposium Virtual reality software and technology*, pages 283–284, Bordeaux, France, 2008.
- [9] L. Rosenberg and S. Brave. Using force feedback to enhance human performance in graphical user interfaces. In *Proc. Conference on Human Factors in Computing Systems*, pages 291–292, Vancouver, British Columbia, Canada, April 1996.
- [10] R. Schmidt. *Motor Control and Learning: A Behavioral Emphasis*. Human Kinetics Publishers, Champaign, Illinois, 1982.
- [11] Shaffer. Cognition and motor programming. *Tutorials in motor neuroscience*, pages 371–383, 1991.
- [12] L. Theremin. U.S. Pat. No. 1,661,058, 1928.
- [13] B. Verplank. Haptic music exercises. In *Proc. of the Int. Conf. on New Interfaces for Musical Expression*, pages 256–257, Vancouver, BC, Canada, 2005.
- [14] T. Yamada, D. Tsubouchi, T. Ogi, and M. Hirose. Desk-sized immersive workplace using force feedback grid interface. In *Proc. IEEE Virtual Reality Conference*, pages 135–142, Washington, DC, USA, 2002.
- [15] E. Zwicker and H. Fastl. *Psychoacoustics: Facts and Models*. Springer, New York, NY, 2nd edition, 1999.