Chapter 5

The Virtual Bowed String

5.1 Introduction

So far, this work has examined the potential benefits of adding haptic feedback to an instrument that, in the real case, provides no haptic cues to the player. But what about the case where the real-world instrument has the potential to provide many haptic cues?

The questions that motivate the experiments presented in this chapter are:

1. Is the haptic feedback available via a mechanical coupling between player and instrument crucial to the control of the instrument? and
2. Is this haptic feedback part of a player's internal representation of the behavior of the instrument.

Askenfelt and Jansson (1992) has shown that many musical instruments produce vibrations that are well within the frequency and amplitude range to which mechanoreceptors in the skin are sensitive. Measurements taken for the open G string on the violin played fortissimo (Jansson, 1970), indicate that vibration levels recorded at the top plate for the two lowest partials were above the sensation threshold for mechano-receptors in this range (Low G = 196Hz, which is close to the peak sensitivity of pacinian corpuscles at 250Hz (Bolanowski *et al.*, 1988).) Vibrations at the chin rest, about 15dB lower than those at the top plate, are also above threshold. Furthermore, these vibrations may well be transmitted via the jaw bone to the resonant cavities of the head and to the ear itself. Figure 5.1 shows potential haptic cues available to a violin player.

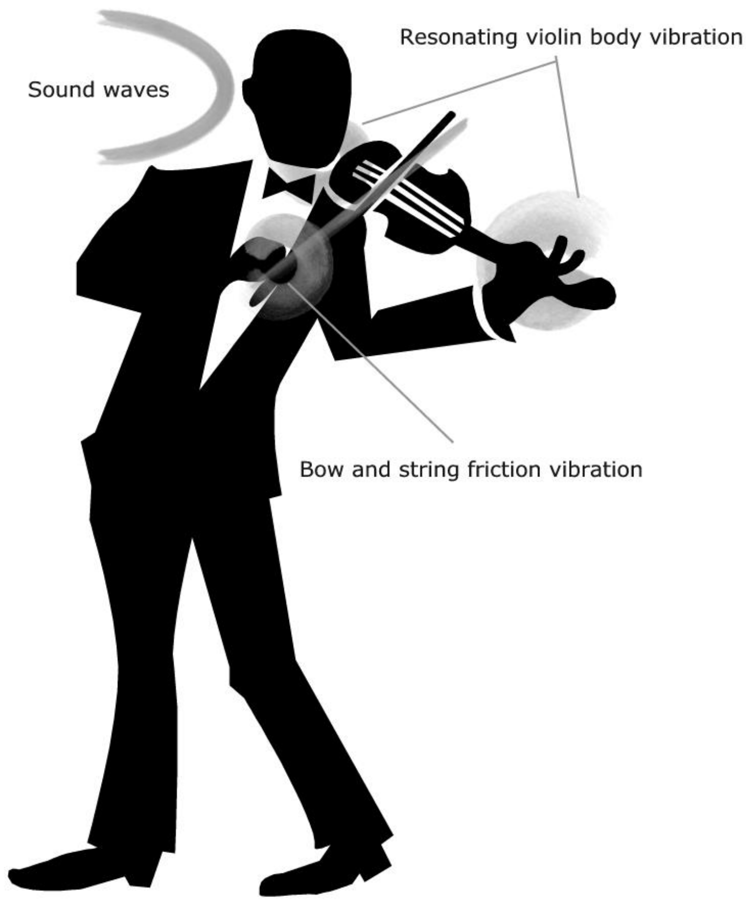

Fig. 5.1. Vibrotactile and force cues potentially available to a violin player.

Given that these cues are potentially available to the player, the question then becomes: Are the reaction forces produced in exciting the strings via the bow, the small vibration cues provided by the strings themselves and the vibration felt at the fingertips of the bowing hand resulting from the friction between bow and strings providing cues that help the player to control the bow? Chafe (1993) has provided evidence that such haptic cues are certainly available to a player. By attaching accelerometers to the bridge of a cello and to the fingernails of the bowing hand, he was able to record the vibrations transmitted from the string to the player's hand and the body of the instrument (see Figure 5.2 taken from Chafe (1993).)

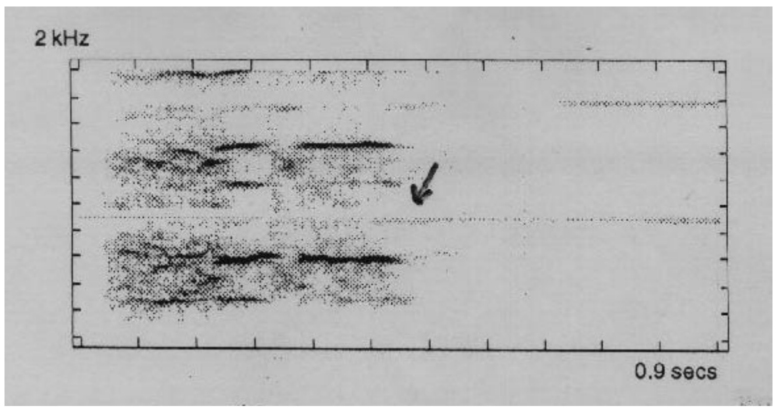

Fig. 5.2. Vibrations recorded from accelerometers on a player's fingertips during normal 'cello bowing. Low frequency components are shown in a spectogram of finger vibration at note onset. These components disappear as the string settles into stable oscillation (seen after the arrow).

Though much of the energy in the steady state portion of this note is too high to be useful to the haptic system (Verrillo, 1992), the burst of energy at the note's start contains frequency components that can be felt. Chafe concluded that these energy bursts, which occur at the beginnings of bow strokes, may well provide important timing cues to a player, particularly in ensemble playing. Since similar broad-band vibration signals are generated whenever there is a large slip between the bow and the string, they may also inform the player of instability in the bow-string contact and potential loss of control of the bow.

In the experiments presented here, we constructed a virtual bowed string with both audio and haptic feedback. We used this model to discover whether the presence or absence of friction in the haptic model would affect a player's ability to imitate a pre-recorded bow stroke.

5.2 Audio and Haptic Models for the Virtual Bowed String

The virtual bowed string model used in these studies was comprised of two components, each running on separate processors in real time.

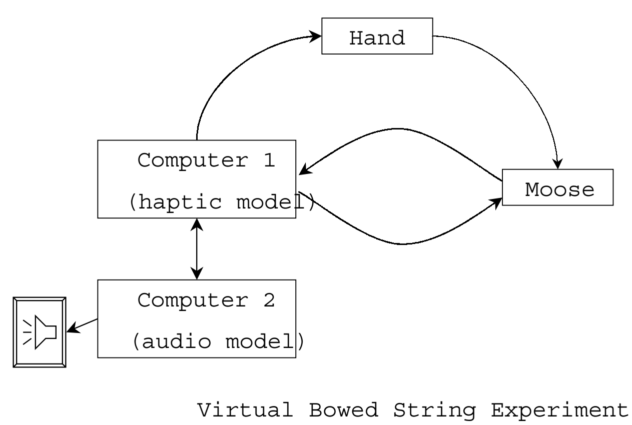

Fig. 5.3. Block diagram showing components of the virtual violin simulation.

Audio output was generated by a computational model of a bowed string (Serafin *et al.*, 1999). The inputs to this model, bow velocity and bow pressure, were derived from the position of the player's hand in the 2D workspace of the haptic display device. These control parameters were transmitted via MIDI to the audio model. Normal and frictional forces simulating the haptic interaction between bow and string were computed locally and fed back to the player's hand.

5.2.1 The Audio Model

The audio model of the virtual bowed string is generated using "digital waveguide synthesis" (Smith, 1998). In this technique, the wave equation is first solved in a general way to obtain traveling waves, which are thereafter simulated using delay lines (in contrast to computationally expensive physical models based on numerical integration of the wave equation.) The outputs of these delay lines are then summed to produce a physical output.

In the case of the bowed string, the bow excites the string at a single point causing two waves to propagate along its length — one toward the bridge and the other toward the nut, where they are each reflected. Thus the bowed string is modeled, in its simplest case, by two delay lines — one simulating propagation of a longitudinal wave in the portion of the string between the bowing point and the nut and the second simulating propagation between the bowing point and the bridge. The position of the bowing point on the string represents the "bow-bridge distance" and is fixed in our simulation at a normalized position of 0.08 where 0 represents the bridge and 1 the nut. The frictional component of bow-string interaction in this simulation is solved as an hyperbolic function depending on the relative velocity between bow and string. Currently no attempt is made to use this solution to drive the haptic friction model because it is not possible to communicate between haptic and audio models at audio sampling rates.

5.2.2 The Haptic Model

The two components of the bow-string interaction that were used as input to the audio model, bow velocity and bow pressure, are associated in a real bow stroke with frictional and normal forces respectively. Therefore, in modeling the haptic feedback present in the interaction between bow and string, both normal (perpendicular) and frictional forces must be present.

Normal force

In our simulation, the string is modeled as a virtual surface and the bow as a single moving point. This is somewhat counter-intuitive as, in reality, the bow is a moving surface and the string is a quasi-static point. The normal force magnitude, Fn, is computed to be proportional to the penetration of the bowing point inside the virtual surface. As a first approximation, this model assumes the following:

1. That the normal force Fn increases monotonically as bow pressure is increased.

2. That the stiffness of the bow is constant along its length.

   These assumptions are valid for the current implementation of the audio model but, as it is extended, the haptic model must be refined to reflect this increased sophistication.

Frictional force

Classical Helmholtz steady-state theory predicts the string displacement, and thereby the driving force on the bridge to be controlled by both bow velocity and bow-bridge distance. In other words, the displacement of the string is only secondarily a function of bow pressure. Bow force needs only to be kept between a maximum and minimum limit to maintain Helmholtz motion. Below a certain minimum force, the bow fails to keep hold of the string during the sticking part of the cycle and proper Helmholtz motion does not develop. Above a maximum force, the circulating Helmholtz corner is insufficient to initiate the slipping phase of the cycle, and oscillations break down (Schelleng, 1973). At the beginning of a bow stroke, the string is stationary. Then it is displaced, driven by the bow - this is the stick phase. When bow force exceeds some maximum value, the bow and string separate and the string slips back relative to the bow. Traditionally, this is thought to be caused by Helmholtz motion in the string — as the peak passes the bow-string interaction point, it causes the bow and string to part. Each cycle (at whatever frequency the string is tuned to) therefore has a long sticking portion and a short slipping portion. The relative lengths of these two parts of the cycle are a function of both bow velocity and bow pressure and their ratio defines a 2-dimensional timbral space described at one corner by light, fast bow strokes and at the opposite by slow heavy strokes. The friction between bow and string therefore plays a key part in allowing a player to negotiate this timbral space (McIntyre *et al.*, 1983). The friction model used in simulating the bow-string interaction is based on Dahl's model of pre-sliding displacement. In this model, the frictional force is proportional to a tension, z, which can be thought of as a spring connecting two points — a point on the moving object, X, and an adhesion point on the fixed object, w. In the case of the bowed string, x is the point on the bow that is currently in contact with the string, and w is an infinitely small cross-section of the string. During adhesion, w is attached to the fixed object so $z = x - w$. This signed quantity describes micro movements between the two objects. The absolute value of z, the spring tension, is capped at z_{max}, beyond which w relocates so that at all times, $|z| <= z_{max}$. While the contact is fully tense, $\dot{x} = \dot{w}$ and $\dot{z} = 0$ (i.e., the model simulates the sliding phase of friction).

The more general form proposed by Dahl is as follows (Dahl, 1976):

$$\dot{f} = \sigma_0 v|1 - \frac{f}{f_c} sgn(v) z|^i sgn(1 - \frac{f}{f_c} sgn(v)), \tag{5.1}$$

Where $v = \dot{x} = \frac{dx}{dt}$, f is the friction force, f_c is the Coulomb force and σ_0 the assumed stiffness relating force to tension.

The friction model implemented here is basically Dahl's model with one modification proposed by Hayward and Armstrong (1997). Hayward proposed that, for the purpose of displaying friction using a haptic display, frictional forces should depend on displacement, not velocity. Because of the compliance in the contact, rapidly varying external applied forces will result in reversals of velocity, i.e. microscopic motion. In the case where these external forces are being applied by a human, such rapidly varying forces are inevitable since they will result from involuntary hand tremor.

Two further modifications to the basic friction model were made to achieve a more

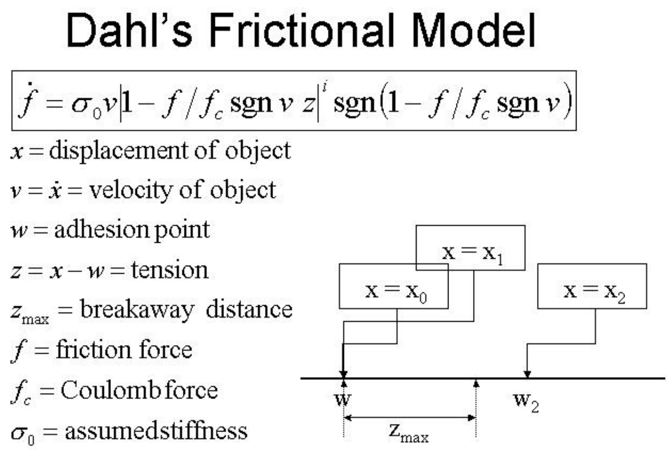

Fig. 5.4. Block diagram of Dahl's pre-sliding displacement model.

realistic simulation.

1. F_C, the coefficient of friction, increased as normal force increased, i.e., the more pressure was applied to the bow, the harder it was to move across the string.

2. A small amount of noise was introduced into the calculation of the frictional force (Green, 1999) to simulate non-uniform behavior of bow hair *.

Tuning the Haptic Model

Before beginning formal testing, we conducted a pilot study with experienced string players to assess the goodness of the "feel" of our virtual bowed string. Players were presented with each of three different variations of feedback, presented in random order, and were asked to indicate which they preferred:

1. normal force alone

*This model assumes that the magnitude of friction forces varies monotonically with normal forces which is unlikely to be the case.

2. normal force + coulomb friction and

3. normal force + our friction model.

All preferred the model we had developed, and further suggested some small modifications which were then implemented.

5.3 Experiment III: Friction

This experiment tested the hypothesis that the presence of friction in a haptic model of a virtual bowed string would affect the novice player's ability to establish and maintain good Helmholtz motion. If the presence of friction helped the novice player to monotor the stability of the contact between bow and string then those players who felt both normal and frictional forces should perform better than those who felt normal forces alone.

5.3.1 Experiment

Participants

20 participants were recruited for this study from graduate and undergraduate music classes at Stanford. All had some basic musical training and some were experienced musicians. None had previously played bowed instruments. All received $5 gift certificates for taking part in the experiment.

Apparatus

To realize the virtual bowed string in real time, we connected our 2 degree-of-freedom haptic display, the Moose (O'Modhrain and Gillespie, 1996), to an audio model of the bowed string via MIDI. The haptic display was oriented so that it provided normal forces in the vertical plane when the player pushed down on the virtual string and frictional forces in the horizontal plane as the bow was moved across the string. Bow pressure and bow velocity were derived from encoder readings in the vertical and horizontal planes respectively. These values were scaled to fall within the range 0-127 and transmitted to the audio model as continuous MIDI control parameters. These values were also used locally to compute normal and frictional forces relayed to the player's hand (see Fig. 5.1).

Finally, audio output from the bowed string simulation was sampled at 22kHz and recorded to disc for later analysis.

Stimuli

In this experiment, participants were randomly assigned to one of two groups, the friction and no friction conditions. In all cases, their task was to imitate as closely as possible a pre-recorded bow stroke that fell well within the playability region for the bowed string model. The sample bow stroke that players were to imitate was recorded onto Compact Disk. Each recorded bow stroke was followed by a 4-second silence. The player's task was to attempt to imitate the stimulus bow stroke each time it was played.

Procedure

Before beginning the experiment, all participants completed a short questionnaire (see Appendix A). They were then shown the virtual bowed string and given approximately 2 minutes to become accustomed to playing. They were shown how to produce various timbres by varying both the velocity and pressure of the bow. They were given several practice trials in which they imitated the recorded bow stroke. When they were ready, they recorded 20 bow strokes, each time imitating the recorded sample.

5.3.2 Results

Unlike the Theremin, the bowed string is a highly complex system in which many input parameters interact to produce good Helmholtz motion. Using analysis techniques, it is possible to observe the force, velocity and position trajectories of the bow during playing. But such techniques cannot yet fully describe what it is in the quality of a tone that appeals to us as musicians. Therefore, in assessing the goodness of the bow strokes obtained as data here, it was important to take both empirical analysis and musical judgements into account. Two experienced string players ranked data bow strokes. Their scores were then compared against the output of an algorithm that detected the presence of Helmholtz motion based on position, velocity and force data recorded from the haptic display device.

Qualitative Measurements of Bow Strokes

In order to assess, musically, the relative success of friction and no-friction feedback conditions, we asked two independent scorers, both of whom were professional string players, to rank individual bow strokes according to how similar they were to the pre-recorded sample. Bow strokes were ranked on a seven point scale with a score of 1 for data bow strokes that most closely matched the sample.

Fig. 5.5 shows the scores obtained by two players, one in the friction feedback condition and the other from the no friction group.

Agreement between the two raters, so-called "inter-rater reliability," was measured using the "κ Coefficient of Agreement"

$$\kappa = (p(A) - p(E))/(1 - p(E)). \tag{5.2}$$

Where $p(A)$ is the observed proportion of agreement and $p(E)$ is the expected proportion of agreement by chance (Siegal, 1988). The κ coefficient of agreement is the ratio of agreement observed in excess of chance, to the maximum possible agreement

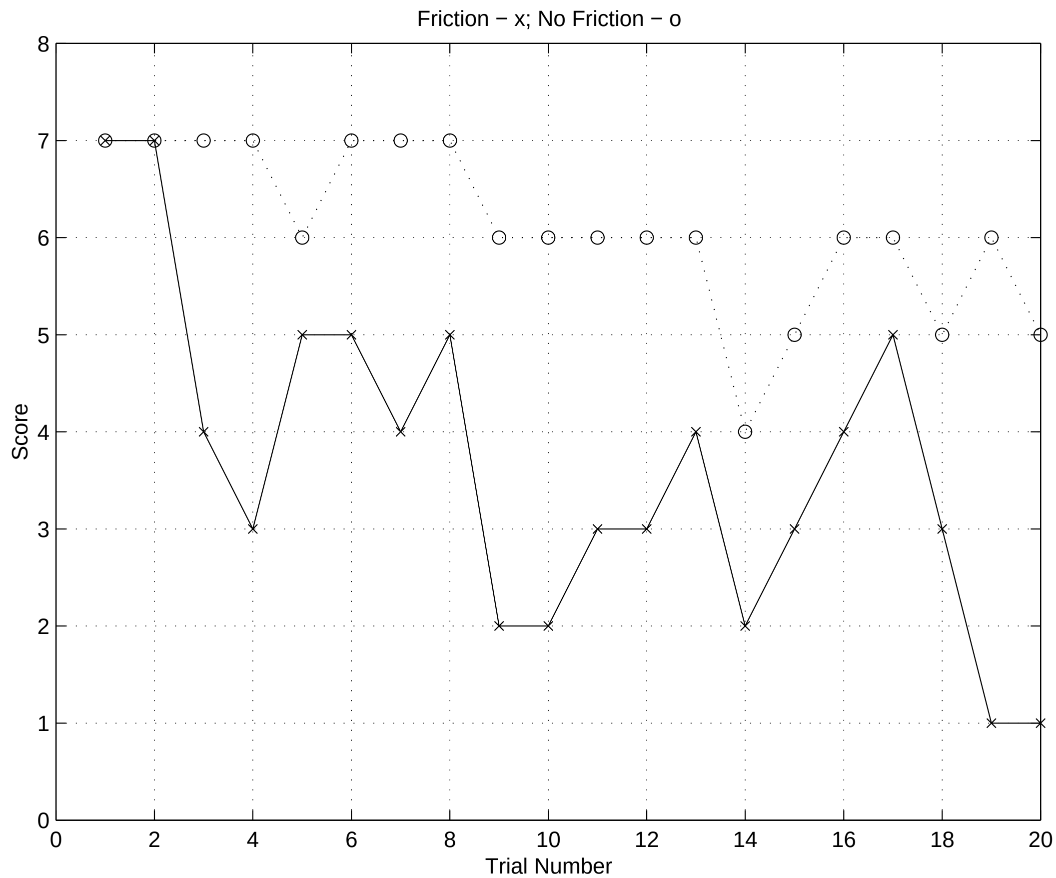

Fig. 5.5. Scores by trial for two players, one player from the friction and one player from the no-friction experimental group. Scores are from one judge and are on a scale from 1-7, where 1 is most similar to the sample bow stroke. Scores are given for each trial ordered in time from left to right.

in excess of chance.

In calculating $p(E)$, the expected proportion of chance agreement, we first assumed that all scores from 1 to 7 were equally likely a priori ($p(E) = 1/7$). Using the approximately normal large-sample distribution of κ and the following equation for the variance of the estimate of κ, we calculated a standardized Z-value of 15.2, which exceeded the $\alpha = .01$ significance level.

$$Var(\kappa) \approx \frac{2}{Nk(k-1)} \frac{p(E) - (2k-3)(p(E))^2 + 2(k-2)\sum p_j^3}{(1-p(E))^2}. \quad (5.3)$$

where k is the number of raters and p_j is the proportion of j ratings among all ratings assigned. We therefore concluded that agreement between raters was significantly above chance level.

Overall, players performed equally well in both friction and no-friction feedback conditions. The mean score for all friction trials was 4.50 while that for no-friction trials was 4.40, representing a difference of less than 3 per cent.

Quantitative Measurements of Bow Strokes

The ultimate test of playability for the virtual bowed string is whether the presence of friction enabled players to establish and maintain good Helmholtz motion, and hence good tone, over the duration of a bow stroke. As discussed in (Serafin *et al.*, 1999), the region of playability for the bowed string is clearly bounded by maxima and minima for the three principal input parameters, bow velocity, bow pressure and bow-bridge distance. Given this, we were able to derive from stored force and position data the envelopes for bow velocity and bow pressure for one player's bow strokes. We then compared these values to the envelope for the sample bow stroke (see Fig. 5.6) to obtain objective scores for the goodness of data bow strokes. Helmholtz motion was detected using the algorithm described in (Serafin *et al.*, 1999). In these experiments, bow-bridge distance was fixed at a normalized position of 0.08. The corresponding playability region obtained by varying bow velocity and bow force is shown in Fig. 5.7.

Fig. 5.8 shows the envelopes for a stroke rated very similar to the sample bow stroke. Note how the envelopes are close to those of the sample stroke, and how the velocity and force values fall into the playability region of Fig. 5.7 for almost the whole duration of the stroke. Fig. 5.9 shows the envelopes for a stroke rated very dissimilar to the sample. Note how the envelopes are also dissimilar, and how the values almost

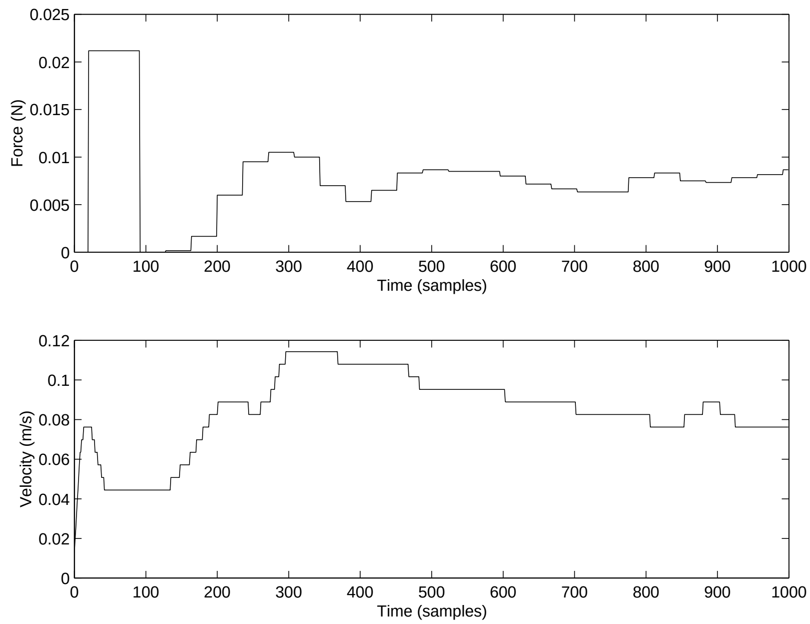

Fig. 5.6. Force and velocity envelopes for the sample bow stroke plot.

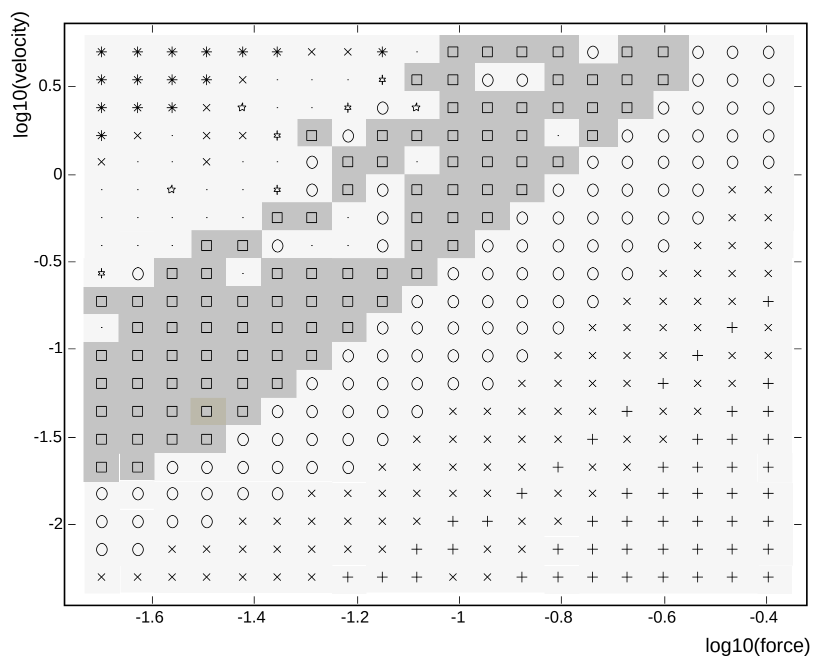

Fig. 5.7. Playability chart for a fixed normalized bow position of 0.08. x-axis=bow velocity, y-axis=bow force.

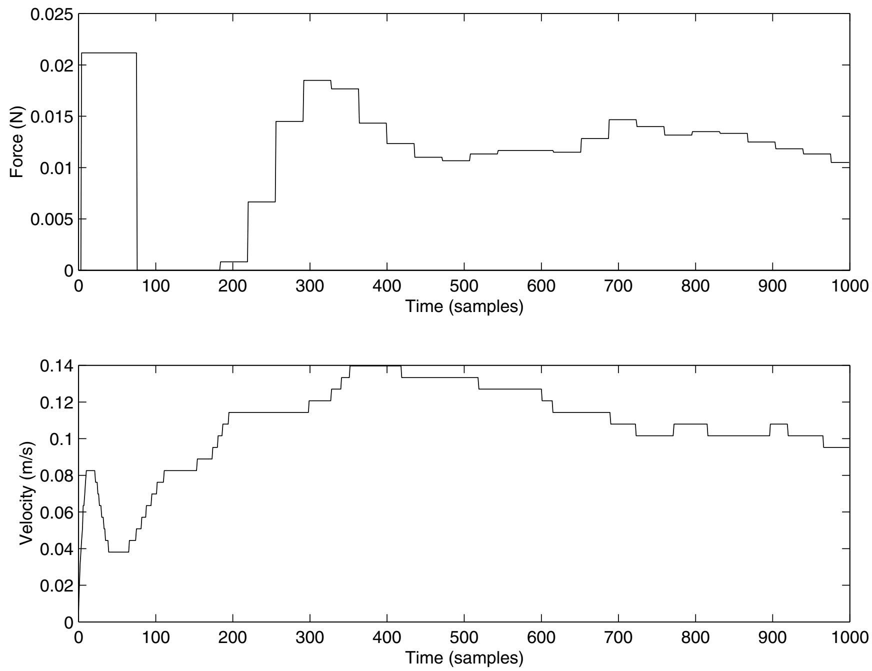

Fig. 5.8. Force and velocity envelopes for a bow stroke rated as 1.

never fall inside the playability region.

The next step in data analysis was to use the detection algorithm to process all data bow strokes. Since the algorithm's output is a series of parameters whose evolution determines the existence of Helmholtz motion in the audio signal, it was necessary to develop a second algorithm that could recognize patterns in these parameters and compute a measure of "goodness," a quantitative score for each data bow stroke. However, developing the statistical pattern matching algorithm which it turned out would be necessary to perform this analysis was considered to be beyond the scope of this present dissertation.

We next turned to the velocity, position and force data recorded from the haptic display during the experiment. Schelleng (1973) determined that the boundaries of the region of playability for the bowed string could be described in terms of a maximum value for bow force. Adapted for a fixed bow position of 0.08 in our simulation, 'f_{max}', maximum allowable bow force, is defined as:

$$f_{max} = Z * v_b/(0.5) * \beta \tag{5.4}$$

In our case: $Z = 1.1$ and $\beta = 0.1$. Therefore:

$$f_{max} = 1.1 * v_b/(0.5) * 0.1 \tag{5.5}$$

So

$$f_{max} = 2.2 * v_b \tag{5.6}$$

In other words, we reject bow force values greater than 2.2 * bow velocity. Having determined boundaries for the region of playability, we developed a simple algorithm to calculate the amount of time recorded force and velocity values remained with in the region of playability. However, when we compared the output from this algorithm with that of (Serafin *et al.*, 1999) this new measure of performance proved to be unreliable. We determined that velocity and force data, recorded at the servo rate of the haptic display device, were too coarse to provide meaningful results on which to base a quantitative analysis of "goodness" of data bow strokes.

5.3.3 Discussion

The results of this experiment indicate that, for novice players, the presence of friction in the haptic response of our virtual bowed string was not providing cues about the stability of the contact between bow and string, since the presence of friction had no effect on performance.

This finding appears to suggest that the ability to use friction as a cue to maintain stability within the context of bowing could be based solely on experience intrinsic to the bowing task. A second possibility is that our simulation was not close enough in its feel to a real instrument for players to be aware of the connection between instability in bow-string contact and sound quality.

In the second experiment presented here, our participants were experienced string players whose knowledge of both the task domain and the feel of bowed instruments provided us with the opportunity to critically assess the playability of the haptic model.

5.4 Experiment IV: Transfer of Skill

The three experiments presented in this work to date have looked at the role of haptic feedback in the very early stages of learning a new musical instrument. As noted in Chapter 3, the early stages of learning are the most fruitful ground for exploring the influence of different forms of sensory feedback on performance. At this stage, the learner depends on cues from the environment to estimate the success of actions, having no earlier experience of the task on which to draw.

For the experienced player, on the other hand, feedback from the environment plays a much smaller role (see Chapter 3). Only when the response of the instrument changes suddenly is the player again conscious of its "feel." Switching a string player's bow, for example, will disorient them temporarily until the player adapts to the dynamics of the new bow.

If the goal of controllers for computer-based instruments is to enable the translation of movement into expressive musical gesture, then one possibility is to leverage instrumental technique built up over years of practice. For this to be effective, computer-based instruments and their controllers must resemble their real-world counterparts sufficiently to enable the transfer of skill from the real to the virtual domain.

In this last experiment, experienced string players were recruited to play the virtual bowed string. Again, the "feel" of the model was manipulated in order to determine whether the presence of friction would facilitate transfer of technique from a real to a virtual bowed instrument. Specifically, would the presence of friction affect in any way the player's ability to control the interaction between bow and string.

5.4.1 Experiment

This experiment tested the hypothesis that the presence or absence of friction in a virtual bowed string model would affect experienced player's performance of a simple bowing task.

Participants

6 string players, four violinists, one viola player and one cellist, were recruited for this study. All were experienced players having an average of 10 years training on their instrument. All received $5 gift certificates for taking part in the experiment.

Apparatus

The experimental apparatus was the same as that used for Experiment III.

Stimuli

In this experiment, all players recorded trials in both friction and no-friction conditions, though the order of presentation of these conditions was random.

Again, the players' task was to imitate as closely as possible a pre-recorded bow stroke that fell well within the playability region for the bowed string model.

Procedure

After a short training period (approximately 2 minutes) players recorded 20 trials in their first friction condition. They then rested, before being presented with the second friction condition. Again they were given a short practice period before recording 20 further trials.

At the end of the experiment, players was presented with both the no-friction and friction feedback and asked to indicate which they preferred.

5.4.2 Results

As with Experiment III, we asked 2 independent scorers who were experienced string players to rank data bow strokes according to their similarity to the sample bow stroke. Again, the κ Coefficient of Agreement was used to determine inter-rater reliability and agreement was found to be significant ($z = 13.9668$) which exceeds the $\alpha = 0.01$ significance level [†].

The mean score for all trials where friction was present was 5.4 and that for no-friction trials was 4.5, representing approximately a 10 percent improvement in performance for trials where no friction was present.It should be noted that the performance in no friction trials was very similar to that for both friction and no-friction trials in Experiment III (friction = 4.50 and no-friction = 4.40.)

Fig. 5.10 shows mean scores by trial for one player in both friction and no-friction feedback conditions[‡].

5.4.3 Discussion

Skill transfer is defined as the effect of practicing one skill on the subsequent performance of another. In examining the "transfer" of a motor skill from one environment to another, a transfer study will typically measure performance of a task B, comparing groups of subjects who practiced another task A, against those who just practiced task B.

In the present study, however, all participants were skilled string players, bringing to

[†] An analysis of the mean scores for each friction condition revealed no significance for order of condition presentation.

[‡] Rankings were on a scale of 1-7, with 1 representing a bow stroke most similar to the sample and 7 least similar.

the task years of prior experience. Ideally, performance of this task would be measured against base-line performance of bowing on a real instrument. In the absence of accurate techniques for measuring the position, velocity and force at the point of bow-string interaction, however, it was necessary to develop a paradigm that could measure skill transfer assuming prior knowledge of the task. One possibility that we considered was to measure performance on hybrid instruments such as the 'celletto' (see Fig. 5.11) that could drive the same audio model.

This option was rejected because the midi transducers on the 'celletto' provided coarse control of the audio model compared with the experimental set-up.

The experimental hypothesis, that the presence or absence of friction in virtual bowed string model would affect performance of a simple bowing task, was tested under the assumption that, like a strange bow, players would quickly adapt to the virtual bowed instrument. In terms of skill transfer, if there were no differences in performance between friction and no friction trials, then friction had no impact on performance of the task. If performance in friction trials was better, then players were tapping into some element of learned playing technique that relied on friction. If, on the other hand, performance on friction trials was worse than that on no-friction trials, the friction feedback was simply confusing players.

The results, which indicate poorer performance in the presence of our friction model i.e. negative transfer, suggests that we had tapped into some component of playing technique where friction is used, but our simulation was not good enough to leverage skilled performance on the real instrument. Moreover, though their performance with friction feedback was worse than that without friction, most players indicated a preference for the virtual bowed string model with friction.

In summary, degradation in performance in the presence of this friction model in the haptic feedback for a virtual bowed string suggests that players were confused by the "feel" of the model. This confusion, together with a marked preference for friction over no friction feedback is encouraging because it suggest that a more finely tuned model has the potential to promote positive skill transfer.

5.5 Summary and Conclusions

The experiments presented here have sought to explore the role of haptic feedback in playing a bowed string instrument. By endeavouring to simulate the "feel" of frictional and normal forces present during bowing, it was hoped to discover whether feedback from friction between bow and string played a role in informing the player about the stability of the bow-string contact. However, results of an experiment with experienced string players lead us to conclude that our simulation was close enough to confuse player's internal representation of the "feel" of a bowed string, but not close enough to promote transfer of skill from the real task.

As Loomis (1992) has pointed out, only simulated environments that provide accurate cues to the perceptual system will promote transfer of skill from real to virtual task domains. When simulations are less than realistic, they are so confusing that they are of no use in training and misleading to experienced operaters. In Loomis's terms, players are unable to model the linkage between their actions and the reactions of the objects they control.

It is interesting to note that even though the presence of friction for experienced players resulted in poorer performance, most indicated that they preferred even our unrealistic friction model. In the same way that people were found to prefer color monotors to monochrome monotors even though color made no difference to their performance (Christ, 1975; Kellogg *et al.*, 1984), designers of musical instruments might simply find that people prefer computer-based musical instruments with simulated haptic feedback to those controllers such as MIDI keyboards that have no instrument-specific feedback.

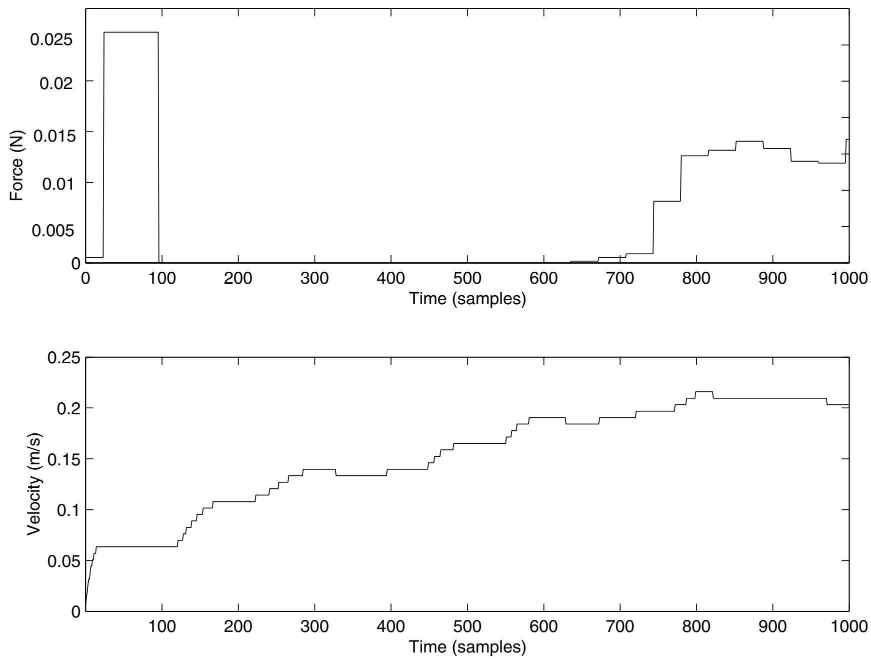

Fig. 5.9. Force and velocity envelopes for a bow stroke rated as 5.

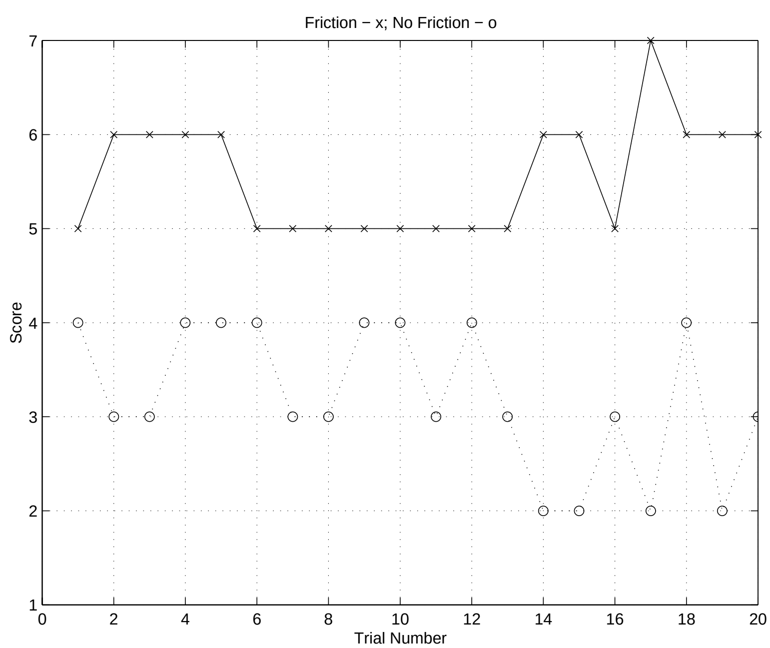

Fig. 5.10. Mean scores for one player in both friction and no-friction feedback conditions.

Fig. 5.11. Celletto.