9th International Conference on Music Perception and Cognition

Alma Mater Studiorum University of Bologna, August 22-26 2006

Perceptual correlates of violin acoustics

C. Fritz,
Faculty of Music,
Cambridge University
Cambridge, UK
cf291@cam.ac.uk

I. Cross, E. Smith, K.Weaver, U. Petersen
Faculty of Music,
Cambridge University
Cambridge, UK

J. Woodhouse / B.C.J. Moore
Dept. of Engineering / Dept. of Experimental Psychology
Cambridge University
Cambridge, UK

ABSTRACT

This paper presents the results of a preliminary series of experiments exploring the extent to which listeners display consistent patterns of preference and discrimination in respect of violin sounds, an issue which has received little prior attention in the literature but which is of great interest for violinists, violin-makers and psychologists.

The principal characteristics used by violin makers to differentiate between instruments relate to the physical features of the violin body. The present study provides a test of a method which enables the same performance to be replayed on different "virtual violins" and it has yielded preliminary data on the abilities of different groups of listeners to indicate preferences for, and to discriminate between, particular violins on the basis of sound alone.

Recordings of real performances were made using a bridge-mounted force transducer, giving an accurate representation of the signal from the violin string. These were then played through filter sets corresponding to the admittance curves of different violins. A preliminary experiment used three violins which were selected as differentiable on the basis of informal listening tests, using both single notes and phrases to explore listeners' abilities to discriminate between pairs of violins. In the second experiment, one violin was used as a basis on which modifications of different magnitude were applied to resonance peaks in order to determine thresholds of discrimination. Three groups of listeners participated in both experiments; non-musicians, violinists and non-string-playing musicians.

Results of the first experiment indicated that all groups of listeners were able to discriminate between, violins when single notes were used as stimuli. However, the non-musician group performed poorly on when phrases were employed (discriminability index d'<1). Moreover, the violinist and musician groups provided patterns of responses that differed significantly from each other for the phrase stimuli.

Results of the second experiment showed that the threshold depends, not surprisingly, on the type of modification and on the note chosen as input, but was not dependent on the musical training of the subjects.

Keywords

Violin, acoustics, perception

INTRODUCTION

There is an extensive literature on the acoustics of the violin, and an even more extensive literature on human perception of sounds in general, and of musical sounds in particular. However, there is virtually no published research on the combined problem of the human capability for perception, discrimination and judgement of the sounds of violins with particular measurable acoustical properties. This is a very significant gap, since perceptual judgements must define what makes a violin different from other bowed-string instruments, and one violin different from another. A project to begin the process of filling this gap has recently started, and this paper will review the work so far and targets for the near future.

The ultimate aim of this research is to answer the typical question that a violin maker will ask: "What will happen to the sound if I change such-and-such a constructional detail?". There are two stages necessary: to relate the constructional change to an acoustical change, and to evaluate the perceptual effect of that acoustical change. The experiments reported here concentrate on the second stage, that of establishing quantitative links between acoustical parameters of the instrument body and the perceptions of a listener. Two broad types of test are relevant: threshold tests to establish the just-noticeable difference (JND) for changes in each relevant parameter, and descriptive rating tests to quantify the perceptual correlates of these various changes. Both types of test will be employed in this study.

The methodology of the study relies on the large impedance jump between the strings and the bridge of the instrument. The player manipulates the string to vibrate in certain ways, the vibrating string applies a force to the bridge, the body vibrates in response to this force, and thus creates a certain pattern of sound radiation. To a first approximation, the body motion has little backward influence on the string motion. There are exceptions, of course: most obviously the "wolf note"[1,2]. More generally, if the topic of interest was the "playability" of the violin rather than its sound, then it would certainly not be admissible to ignore this back reaction. Similarly, if the study was concerned with the guitar or the piano then string/body coupling would be crucial because it determines the decay rates of the various overtones of the string motion. However, for a bowed string it can be argued that such coupling effects can be ignored in the first instance. If strings of the same type are fitted to two different violins, a skilled player will adjust bowing to coerce the vibration into the standard Helmholtz motion with an acceptably short transient[3]. The force waveforms acting at the bridge in the two cases will be very similar, and one would expect that the major differences in sound between the two instruments could be captured by driving them both with identical forcing.

With this in view, representative force waveforms can be recorded using normal playing on a violin whose bridge is instrumented with piezoelectric force sensors. These predetermined force functions can then be applied to different violins, so that sound differences can be compared with no complications arising from variations in playing. Such a test could be carried out using different physical violins, applying the force at the bridge with a vibration shaker of some kind. However, for this study a different approach is taken. The frequency response function of the violin is mimicked using a digital filter of sufficiently high order, and the output signal for listening tests is generated by convolution with the recorded bridge force signal. This filtering can be done offline, using Matlab, or it can be done using a real-time system[4]. Once the violin response is represented in digital filter form, it becomes very easy to make controlled variations of a kind which would be virtually impossible to achieve by physical changes to a violin.

Langhoff et al.[5] performed such experiments in which violin performances were filtered digitally. The frequency response curve used as a filter was modified in several ways, to give enhancement of the Helmholtz resonance, enhancement of mid-range frequencies (around 1.7 kHz) and creation of a smoother decay towards higher frequencies. This experiment did show that it is possible to compare violin spectra by listening to convolved signals, but it did not address the question of how people perceived the different sounds created. No participants were involved; the paper only reports the 'subjective impressions' of one of the authors and that is why further investigation is needed with psychoacoustic tests.

PRELIMINARY EXPERIMENT – PROOF OF METHOD

The experimental methodology was tested, and some preliminary results obtained, in a series of undergraduate projects. The key results of these projects will be described here. First, informal listening tests were used to select three violins which were judged to be clearly distinguishable by blindfolded listeners during live performance. The input admittances of these three violins were measured, and are shown in Figure 1. The violins are labelled A,B,C. Violin A has been played professionally. It is judged to be powerful and flexible, with a soloist character, but perhaps rather crude-sounding. Violin B is owned by an amateur chamber musician, and is successful in that context. Violin C is a student violin of indifferent quality.

The input admittances of these three violins were processed by modal identification techniques, and resynthesised from the fitted parameters in the frequency range up to 4000 Hz to remove the effects of measurement noise and to produce a limited bandwidth without a need for filtering. These resynthesised versions were used to construct digital filters. A short musical fragment was recorded via a force sensor: the chosen passage consisted of the first six notes of the third theme from the Glazunov Concerto for violin in A minor, op. 82, starting on A*b* and played entirely on the G string. The recorded bridge-force signal was used with the three digital filters to create sound files for three "virtual violins", which were then used for listening tests. The sound files were normalised in amplitude to the same peak level.

Note that, in this case, no attempt was made to represent sound radiation behaviour: the synthesised signal corresponds to the body velocity at the bridge. This will obviously not give the usual "sound of the violin", but it has the virtue that the measurement is clear, unambiguous and repeatable, and not subject to any vagaries of microphone placement, room acoustics and so on. These factors will be considered in due course, but it seems plausible that the perception of differences between one virtual violin and another might not depend very much on such details. This will be an interesting hypothesis to test later in the study when more data become available.

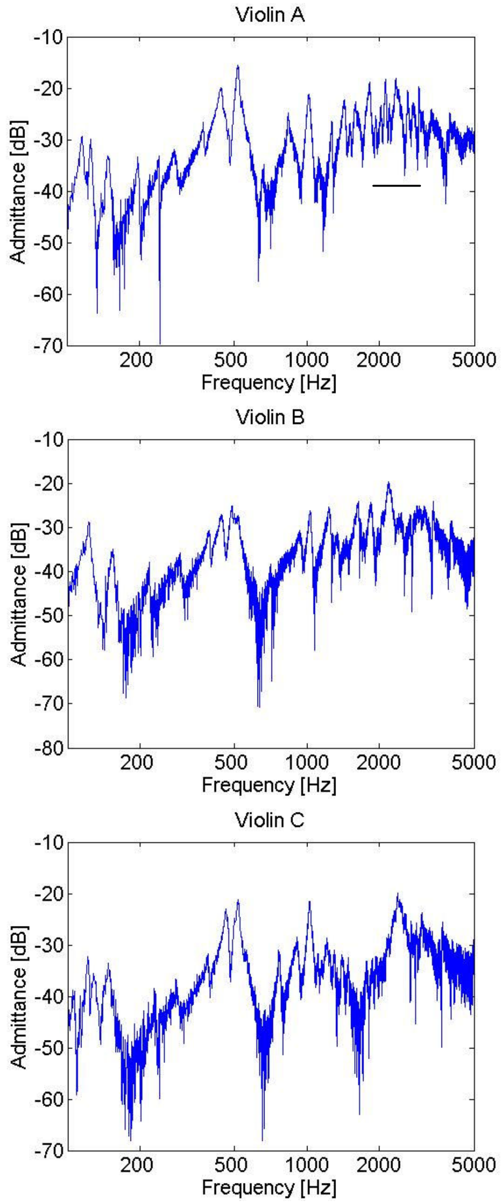

Figure 1: Input admittance measured at the bridge for the three violins used in the listening tests. Levels are plotted in dB re 1 m/s/N.

Two sets of listening tests were carried out using these sound files, one based on the entire musical phrase, and another based on a single note extracted from that phrase. Participants, all of them students at Cambridge University, listened via headphones to groups of three presentations of the note or phrase and were asked to pick the matching pair in an M-X-N paradigm, where M represents one violin, N represent another, and X is the same as either M or N. Each of the three possible pairings of violins A/B/C, in both possible orders, was presented twice, in randomised order. The results were processed using standard signal detection[6] theory to give values of the discriminability index *d*', for which a value exceeding unity implies a reliable ability to discriminate two stimuli.

The participants were divided into three groups: 15 "violinists" (all at or near diploma level), 14 non-string playing "musicians" (all currently studying music) and 12 "non-musicians" (in the rather restricted sense that they were not currently studying music, had in the past studied music for less than five years, and were not involved in music performance). The results for these three groups in the two tests are summarised in Figure 2. It is immediately clear that discrimination was better with the single note test than with the musical phrase. Every group could distinguish each pair of violins with the single note, whereas with the phrase the scores are grouped closer to the "threshold" value of unity. Interestingly, the violinists generally performed better than the other groups based on the musical phrase, whereas with the single notes the non-string musicians tended to outperform them somewhat. The non-musicians consistently showed the lowest performance, as one would have expected.

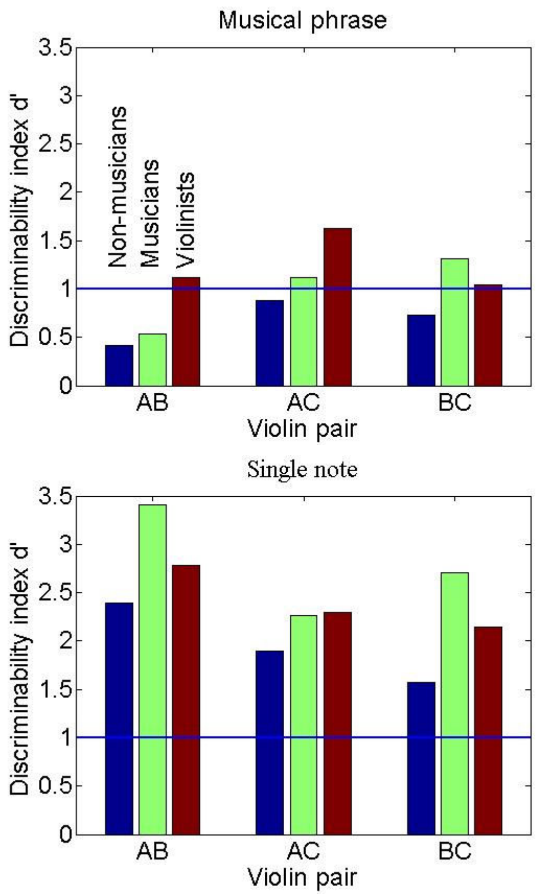

Figure 2. Discriminability index *d*' for the three groups of test subjects and the three pairings of violins, based on a short musical phrase (top) and on a single note (bottom).

FURTHER EXPERIMENTS WITH ONE VIOLIN MODIFIED IN DIFFERENT WAYS

The preliminary study clearly demonstrated the efficacy of the method. However, it raised perhaps too many questions (the violins used were extremely heterogeneous), and it was decided to narrow the range of potential parameters in the subsequent study. Rather than using completely different violins, we selected a single violin (made by David Rubio) as the original model and "designed" other virtual violins differing from the orginal by slight precisely controlled changes involving manipulation of the three most important modes (found in every violin frequency response curve, see Figure 3). These are:

- Two ‘plate modes’ which arise primarily from the bending and stretching of the front and back plates; B1+ (usually found between 530-570 Hz) and B1- (usually in the range 470-490 Hz).
- The modified Helmholtz resonance (‘air mode’) A0, which is associated with the air-pressure variation inside the violin and is usually around 280 Hz.

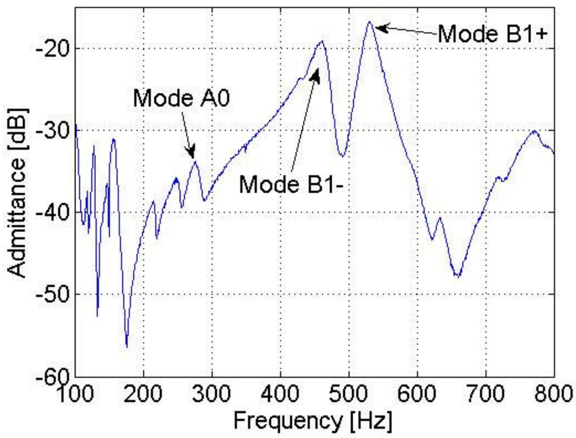

Figure 3. A violin admittance curve showing the three modes that were modified.

The following modifications were made:

1. The frequency of the air mode A0 was altered.
2. The frequency of the plate mode B1- was altered.
3. The frequency of the plate mode B1+ was altered (see Figure 4).
4. The frequencies of all three modes were altered (in the same direction and by the same percentage).

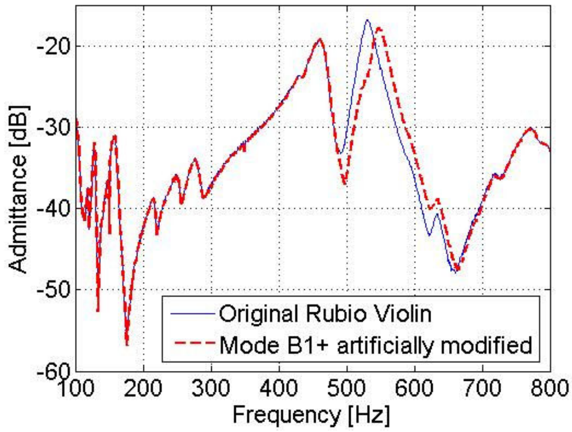

Figure 4. Graph showing the original admittance curve and a modified version with the frequency of mode B1+ altered.

As in the previous study, no sound radiation model was included, but the mode A0 was multiplied by five to account empirically for sound radiation, which plays a much more important role for this mode than for the two others.

As participants in the informal tests reported that they were almost exclusively using the first and last note of each phrase to tell the violins apart, and in the light of the results of the previous study, it was decided that the listening tests would involve only single notes. G3 and E4 were chosen because the second harmonic of G3 is near the mode B1- and the third harmonic is near B1+, whereas the harmonics of E4 are in between the modes we modified. For this reason, we expected to find that the modifications would have a greater effect on the note G3, and that the threshold for discrimination between modified versions of this note would therefore be lower than that for E4.

Procedure

A three-alternative forced-choice (3AFC) test was used in all four tests. Three sounds – two of which were the same, and one different – were played, and each participant was asked to identify that which was different. To allow echoic memory to be effective[7], the sounds were shortened to 300 ms so that all three sounds were heard within one second. Initial values were set for the modification to be explored and if the participant answered wrongly the initial percentage modification was increased, while if three correct answers were given in a row the percentage modification decreased. The position of the mode was moved around a central value; for example, for a modification of 14% of the change was –7% for one sound and +7% for the other sound. The difference between the sounds was at first divided or multiplied by √2 at each stage, but after two ‘reversals’, the difference was divided or multiplied by √(√2). After six reversals the test ended and the mean and standard deviation of the modification values at the last six reversals was calculated. Participants also provided written comments on their experience of the tests.

Results

Table 1 gives the average mean and standard deviation for each test for string players, other musicians and non-musicians. Most participants produced a result for every test, but a small number found one or two of the tests too difficult to complete.

Table 1. Mean threshold in % and standard deviation for each modification, for both notes and for the three groups of subjects.

Note	Modification	19 string players mean	19 string players std	17 musicians mean	17 musicians std	10 non-musicians mean	10 non-musicians std
G	all	7.3	5.6	3.9	3.2	7.4	5.0
E	all	9.2	4.1	10.2	9.3	15.8	14.7
G	B1+	6.7	4.1	4.9	3.4	5.4	2.6
E	B1+	19.6	5.1	15.9	6.2	18.7	4.5
G	B1-	9.0	4.6	5.8	2.7	7.6	4.1
E	B1-	22.0	6.6	19.9	5.3	17.7	6.7
G	A0	27.0	10.4	27.4	8.3	29.5	8.1
E	A0	22.5	3.9	22.8	4.2	22.9	6.3

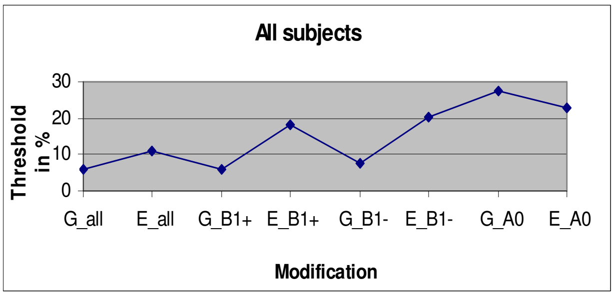

Figure 5. Average thresholds (modification in %) obtained for all subjects, for the different modifications, as a function of the input note.

A repeated measures ANOVA showed that the effect of the note used (G or E) was highly significant ($F_{(1,38)}=79$, $p<0.0005$). For the first three modifications (all three modes, B1+ and B1-), the threshold for G was significantly lower than for E. This was not true, however for the final modification (A0), where the threshold for G was significantly higher than for E. This is not surprising though, as the harmonics of G are near B1- and B1+ but not A0, so the modifications to this latter mode would not be expected to lower the threshold for G.

As expected, the effect of the type of modification was also significant ($F_{(3,114)}=127$, $p<0.0005$) in respect of the thresholds obtained. It is unsurprising to find that the modification of the mode A0 is hard to detect, while a simultaneous modification of the three modes gives a relatively low threshold.

It was very interesting to note that although most participants were more successful in the tests involving G, a small number commented that they thought the tests on E were easier. An oboist suggested that she found the 'E' tests easier because she was more used to listening to notes in that pitch range; she plays the oboe regularly (the range of an oboe includes E4 but not G3), does not play the piano and sings soprano. Also a male pianist and singer suggested that he found the 'G' tests easier because the G was within his vocal range whereas the E was not, and he was more used to making fine judgments about pitches in this lower range.

The following graph shows the results of each group of participants.

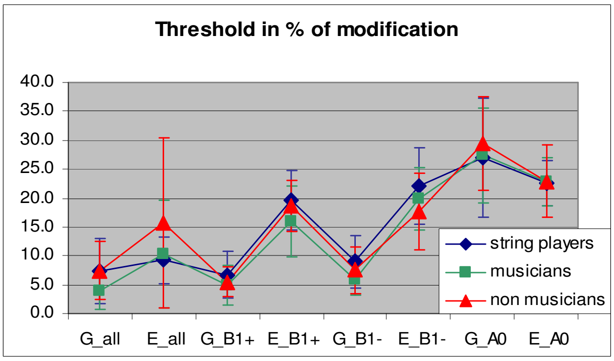

Figure 6. Average thresholds in % of modification obtained for the different groups of subjects, for the different modifications, in function of the input note.

A mixed, repeated measures ANOVA showed that the participants' level of musical training did not have a significant effect on their performance ($F_{(2,38)}=0.7$, $p = 0.491$).

The questionnaires provided a number of comments suggesting that the differences between the sounds seemed to change in the last two tests. Participants said that the penultimate test 'took some getting used to' and that 'the difference between sounds changed in the last test which caught me out', while another said that in the last two tests it 'sounded like the one that had to be chosen had fewer high frequency harmonics whereas the previous ones sounded like they had more on the odd one out'. This can perhaps be explained by the fact that A0 is lower than the modes B0- and B1+, and was therefore affecting the note in a different way; rather than amplifying certain harmonics to differing degrees, the modification of the mode affected the amplitude of the fundamental frequency. This would have 'caught out' anyone who was listening specifically to the overtones of the note, and may explain why some people with otherwise very low thresholds suddenly found it very difficult to discriminate between the sounds in the last two tests.

The questionnaire also provided some interesting information about how people discriminated between the sounds. Even though the participants were not told how the different sounds had been created, many mentioned that they had been listening for 'overtones' or 'harmonics'. For example one participant said, 'I heard two with a lower 'harmonic' and one with a higher one and worked according to that.' Some participants heard the different sounds as differing in pitch, while others mentioned 'loudness' and 'tone quality' as discriminating factors. Two participants described the difference between the sounds as similar to different vowel sounds; one said 'if you think of the sounds as vowel sounds it makes them easier to tell apart', while the other compared the sounds to the different vowel sounds of his Indian language[8].

FUTURE PROGRAMME

Choice of the input signal

The results of these preliminary experiments illustrate some of the challenges to be faced in formulating a systematic programme of tests to map out the full set of musically relevant perceptual attributes of a violin. The differences between the single-note and phrase tests show that the choice of input signal needs to be carefully considered. Three levels can be distinguished. First, the single-note approach: this could be extended, but it might yield different results with different choices of note, and different bowings of that note. To cover all combinations of variables will require a very large number of tests. Second, the musical phrase method: again, results might be different with different choices of passage, string and bowing.

Furthermore, the literature of violin acoustics contains a number of suggestions for acoustical attributes which may correlate with quality judgments by listeners. These provide a "shopping list" of predictions to explore and test in the present project.

Individual modes at low frequency

Many authors have written about the individual modes of a violin body in the low-frequency range. Some of them have explicitly considered the practical issues of adjusting plate geometry to control the parameters of these modes: for example Hutchins[9] and Schleske[10]. Certain of these low modes are sometimes called "signature modes", and the clear implication is that these authors expect the control of these individual modes to have significant perceptual effects. This gives a first and obvious target for study. It is very easy in the digital-filter context to vary the frequency, amplitude and damping of any individual mode, or of groups of modes.

"Graphic equaliser" effects

Another theme which runs through the literature is that important aspects of the sound quality of a violin might be captured by the pattern of sound energy in various quite broad frequency bands. Any description of this kind can be thought of in terms of formant-like characteristics, or as a "graphic equaliser effect", since these are precisely the kind of changes which can routinely be made on a domestic hi-fi system. It is an interesting question how far one go in creating "the sound of a Stradivarius" simply by such broad-brush changes to the frequency spectrum. The most thorough study of this kind is by Dünnwald[11], who measured the frequency response of a large number of violins and made very explicit proposals about the correlation of "quality" with the relative levels in certain frequency bands. These proposals are ripe for psychoacoustical testing, and the present methodology offers an easy way to do so.

A particular effect which falls in this category has been studied in some detail: the so-called "bridge hill"[12]. Many violins show a broad maximum response in a frequency range around 2000–3000 Hz: it can be seen clearly in Figure 1, in the range indicated by the horizontal line in the upper plot. This feature is thought to derive from an in-plane resonance of the normal violin bridge, modified by the coupling through the feet to the vibration characteristics of the violin body[13].

Trend data in modal parameters

A more complicated proposal for acoustical quantities linked to the perception of quality comes from the work of Bissinger[14]. He has carried out very detailed measurements of a number of violins, using both a fine grid of test points on the violin body and a microphone array to measure the radiated sound field. All this information has been processed by standard modal analysis methods. In parallel, each of his tested violins was subjected to a standardised "quality rating test" by a professional player[15]. Bissinger has noted correlations between his "quality" results and certain features which show up in a trend analysis of his modal and radiation results. These lead to hypotheses, which could readily be tested using the digital filter methodology.

Vibrato sensitivity

A different kind of prediction from the earlier literature of violin acoustics concerns sensitivity to vibrato. The idea goes back to the pioneering studies of Mathews and Kohut[16] and Gorrill[17], who in the 1970s experimented with electronic filters to do a similar job to that proposed here with digital filters. One aspect of their results was interpreted by McIntyre and Woodhouse[18] in terms of the inter-

ISBN 88-7395-155-4

action of a "spiky" frequency response function with vibrato, to produce the sense of "liveliness" or "richness" often associated with violin tone (and conspicuously absent from the vibrato effect on most keyboard synthesisers). Similar ideas have recently been explored by Gough[19]. The digital filter methodology gives a simple way to explore such effects.

"Directional tone colour"

A final proposal from the existing literature comes from the work of Weinreich[20]. He suggested that the complex directional character of the radiated sound field from a vibrating body like a violin may, after interaction with the acoustics of the room, be responsible for some of the important perceptual qualities of live violin performance. Loudspeakers are generally designed to have quite different directional characteristics (except perhaps for "distributed-mode" loudspeakers), and this might explain why it is so hard to reproduce the sound of a recorded violin performance well enough to fool a listener that they are hearing a live performance. This is a fascinating and inherently plausible suggestion, which in principle could be explored by the digital filter methodology by generating stereo signals with two different filters, to be listened to via headphones. However, this would be a very challenging task, and it will be deferred to a later stage of the research.

Another approach with an electric violin

Finally, a player-based approach can be used: instead of listening to predefined sound files, a violinist can be given a short time to play whatever they choose on a mute electric violin, listening to the results via a real-time filter system. Preliminary studies are already under way employing this method.

CONCLUSIONS

A methodology has been proposed to perform systematic psychoacoustical evaluations of the perception of controlled variations in the vibration behaviour of a violin body. The method employs input recorded from real playing, via a force transducer in the violin bridge. This input is fed through a digital-filter realisation of the desired "virtual violin". This will typically be based on a measurement of a real instrument, modified to change one parameter at a time. A set of likely parameters to explore has been identified: some of these are deterministic modal properties, while others involve statistical information about the vibration behaviour at higher frequencies. These parameters map quite well onto a number of proposals found in the existing literature of violin acoustics for acoustical quantities showing a correlation with judgements of "quality".

The methodology has now been tested in several studies as described above. It appears that the method is robust, and capable of giving quantitative information on this important subject. Some of the thresholds already measured are in line with expectations: differences of around 5-6%, corresponding to a semitone, are very much in line with what violin-makers might predict However, some thresholds are surprisingly large, being in the order of 15%, and further investigation is required to understand the bases for such divergent patterns of judgment. Nevertheless, the success of the method has opened out a vista of many tests to be carried out, which in time should map out the perceptual landscape of the violin in an unprecedented manner.

ACKNOWLEDGMENTS

The authors thank the Leverhulme Trust for financial support, and Professor Patrick Gaydecki for making available to us the real-time digital filter system and associated software.

REFERENCES

1. Cremer, L. (1985). The physics of the violin: MIT Press.
2. Woodhouse, J. (1993).On the playability of violins: Part 2 Minimum bow force and transients, *Acustica 78*, 137-153.
3. Guettler, K. and Askenfelt, A. (1997). Acceptance limits for the duration of pre-Helmholtz transients in bowed string attacks, *J. Acoust. Soc. Amer. 101*, 2903-2913.
4. Gaydecki, P. and Fernandes, B. (2003). An advanced real-time digital signal processing system for linear systems emulation, with special emphasis on network and acoustic response characterization, *Meas. Sci. Technol. 14*, 1944-1954.
5. Langhoff, A., Farina, A. and Tronchin, L. (1995). Comparison of Violin Impulse Responses by Listening to Convoluted Signals, ISMA95, International Symposium on Musical Acoustics, Paris.
6. Green, D. M., and Swets, J. A. (1974). Signal Detection Theory and Psychophysics, Krieger, New York.
7. Darwin, C. J., Turvey, M. T., and Crowder, R. G. (1972). An auditory analogue of the Sperling partial report procedure: evidence for brief auditory storage, *Cog. Psych. 3*, 255-267.
8. Traube, C., & Depalle, P. (2004). Phonetic gestures underlying guitar timbre description. In S. D. Lipscomb & R. Ashley & R. O. Gjerdingen & P. Webster (Eds.), Proceedings of the 8th ICMPC (pp. 658-661). Evanston, Ill.: Causal Productions.
9. Hutchins, C.M. (1992). A 30-year experiment in the acoustical and musical development of violin-family instruments, *J. Acoust. Soc. Amer. 92*, 639-650.
10. Schleske, M. (1996). On making "tonal copies" of a violin, *J. Catgut Acoust. Soc. Series 2*, 3(2) 18-28.
11. Dünnwald, H. (1991). Deduction of objective quality parameters on old and new violins, *J. Catgut Acoust. Soc. Series 2*, 1(7) 1-5.

12. Jansson, E.V. (1997). Admittance measurements of 25 high quality violins, *Acustica - Acta Acustica 83*, 337–341.
13. Woodhouse, J. (2005). On the “bridge hill” of the violin, *Acustica - Acta Acustica 91*, 155-165.
14. Bissinger, G. (2003). Sound radiation by violins radically modal, *J. Acoust. Soc. Amer. 114*, 2438.
15. Bissinger, G. and Gearhart, F. (1998). A standardized qualitative violin evaluation procedure?, *J. Catgut Acoust. Soc. Series 2*, 3(6) 44-45.
16. Mathews, M.V. and Kohut, J. (1973). Electronic simulation of violin resonances, *J. Acoust. Soc. Amer. 53*, 1620-1626.
17. Gorrill, S. (1975). A viola with electronically synthesised resonances, *Catgut Acoust. Soc. Newsletter 24*, 11-13.
18. McIntyre, M.E. and Woodhouse, J. (1974). Toward a psychoacoustically realistic violin physics, *Catgut Acoust. Soc. Newsletter 22*, 18-19.
19. Gough, C. (2001). Physical aspects of the perception of violin tone, *Proceedings ISMA 2001*, Perugia, 117-122.
20. Weinreich, G. (1997). Directional tone color, *J. Acoust. Soc. Amer. 101*, 2338-2346.