PERCEPTUAL TESTS WITH VIRTUAL VIOLINS

J Woodhouse University of Cambridge, Department of Engineering
I Cross University of Cambridge, Department of Music
BCJ Moore University of Cambridge, Department of Experimental Psychology
C Fritz University of Cambridge, Department of Music

1 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 be able 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. This project is concentrating on the second stage: to establish 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.

2 KEY CONTROL VARIABLES

The first stage is to identify a suitable set of variables to characterise the acoustical behaviour of a given violin. In this experimental methodology, the parameters of the strings (tension, mass per unit length, bending stiffness etc.) are not directly relevant: their influence is contained within the recorded bridge force waveforms. The recordings for the tests described here all used Thomastik “Dominant” strings, “mittel” gauge.

The body behaviour is most naturally described in terms of the normal modes. Each mode has a natural frequency, a damping factor, a spatial mode shape, and a sound radiation characteristic embodying variation in space and frequency. A typical violin has perhaps 100–200 modes in the audible frequency range. However, it is not realistic to expect all of these to have effects which are separately audible. Furthermore, it is certainly not reasonable for a violin maker to expect to control them individually.

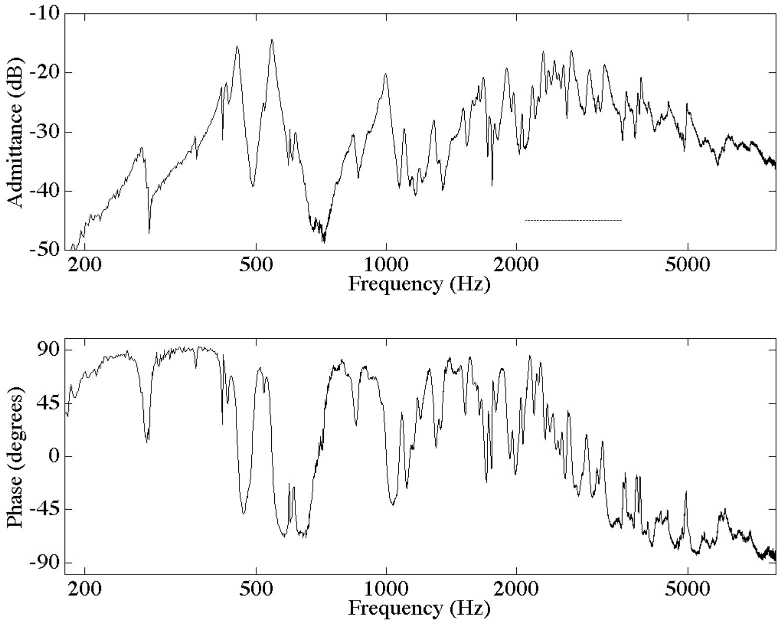

Figure 1. Input admittance of a typical violin, in amplitude and phase.

Figure 1 shows a typical measured input admittance of a violin. Careful study of data of this kind reveals that up to about 1000 ~~k~~Hz the modes show a low to moderate modal overlap factor, and also moderate statistical overlap factor when different violins are compared[5]. At higher frequencies, both overlap factors increase so that they exceed unity. This suggests that individual modes have a clear identity at low frequencies~~y~~, so that it seems sensible to explore the perceptual effects of variation of individual modal parameters. However, at higher frequencies it is likely that only statistical features of the response will prove to have a reliable perceptual identity. Relevant quantities might be: (1) the broad distribution of energy across the spectrum, smoothed in some appropriate way; (2) the typical peak-to-valley height fluctuation of the frequency response, or equivalently the standard deviation of the frequency response; (3) the typical frequency spacing of adjacent major peaks in the response. Some understanding of the behaviour of these quantities can be gained from statistical room acoustics[6] and recent developments in statistical theories of vibration of complex structures[7].

For a structural response such as the input admittance shown in Figure 1 it is a fairly straightforward matter to represent the behaviour in terms of modal parameters. There is a standard formula for any structural transfer function in terms of normalised mode shapes, natural frequencies and damping factors, which can be used as a basis for parameter estimation techniques to allow the modal parameters to be determined from measured data[8]. Unfortunately, the kind of transfer function really needed for this study relates a force applied at the violin bridge to the pressure measured by a~~–~~ remote microphone. While there are computational techniques to predict such transfer functions for particular structures (e.g., the boundary element method), there is no universal formula that covers the entire frequency range of interest and allows measured data to be fitted directly by a standard set of modal parameters. This issue is a matter of ongoing research, and it will be sidestepped for the purpose of the present study, at least in its early stages.

3 EXISTING PROPOSALS TO TEST

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.

3.1 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.

3.2 “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 of response in a frequency range around 2000–3000 ~~k~~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].

3.3 Trend data in modal parameters

A more complicated proposal for acoustical quantities linked to the perception~~s~~ 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 are testable hypotheses, which could readily be probed using the digital filter methodology.

3.4 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 ~~by~~ 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 interaction 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: one might expect in this case that the key parameters ~~might~~ would involve the statistical quantities at higher frequencies discussed in section 2.

3.5 "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~~ions~~ 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.

4 PRELIMINARY EXPERIMENTS

The experimental methodology has been tested, and some preliminary results obtained, in a series of undergraduate projects. Some key results of these projects will be described briefly 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 2. 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 ~~k~~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 Ab 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 is 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 ~~may be guessed~~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~~s~~ available.

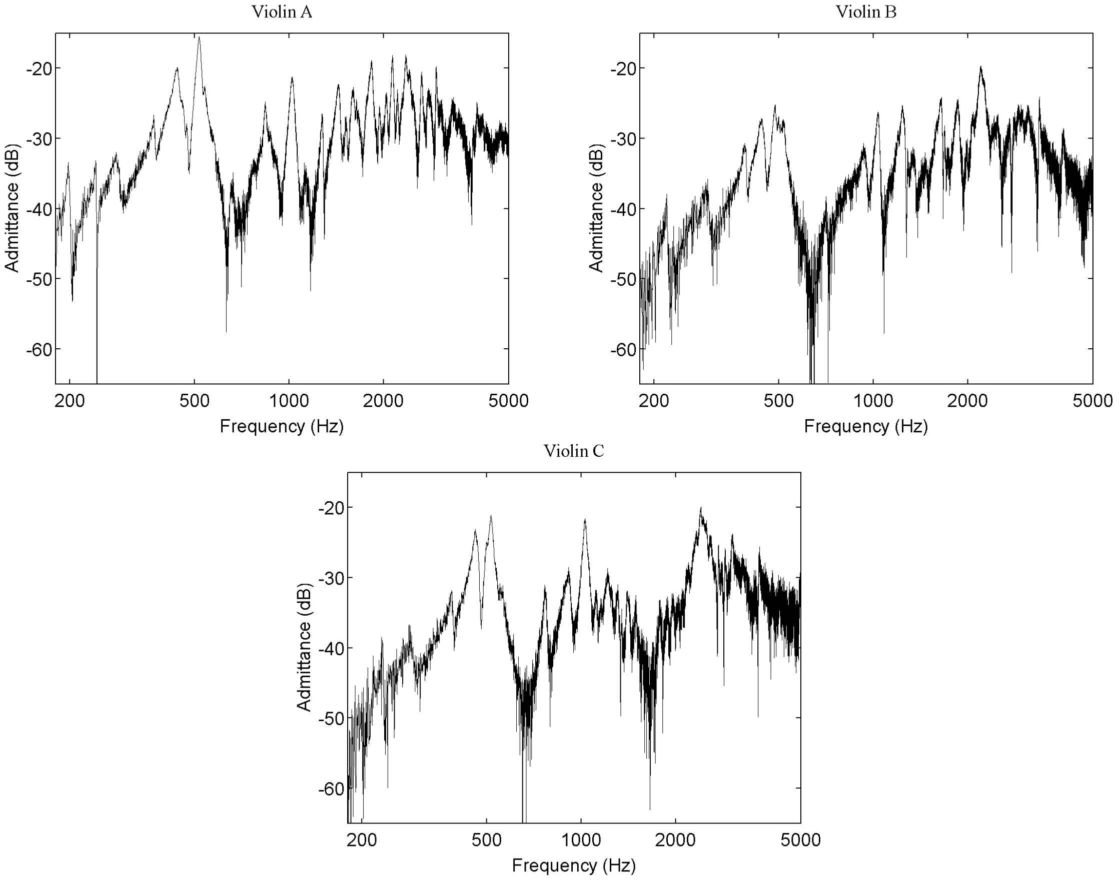

Figure 2: Input admittance measured at the bridge for the three violins used in the listening tests. Levels are plotted in dB re 1 m/s/N.

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. Tests were done for preference and also for discrimination: only the latter will be described here. The test subjects, 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 ~~A~~M-X-~~B~~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 theory to give values of the discriminability index *d*', for which a value exceeding unity implies a reliable ability to discriminate two stimuli.

The test subjects were divided into three groups: 15 "violinists" (all at or near diploma level), 14 non-string playing "musicians" (all currently studying ~~the ~~music~~ tripos~~) 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 3. It is immediately clear that discrimination was better based on 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 performances are grouped closer to the “threshold” value of unity. Interestingly the violinists generally perform better than the other groups based on the musical phrase, whereas with the single notes the non-string musicians tend to outperform them somewhat. The non-musicians consistently show the lowest performance, as one would have ~~guessed~~expected.

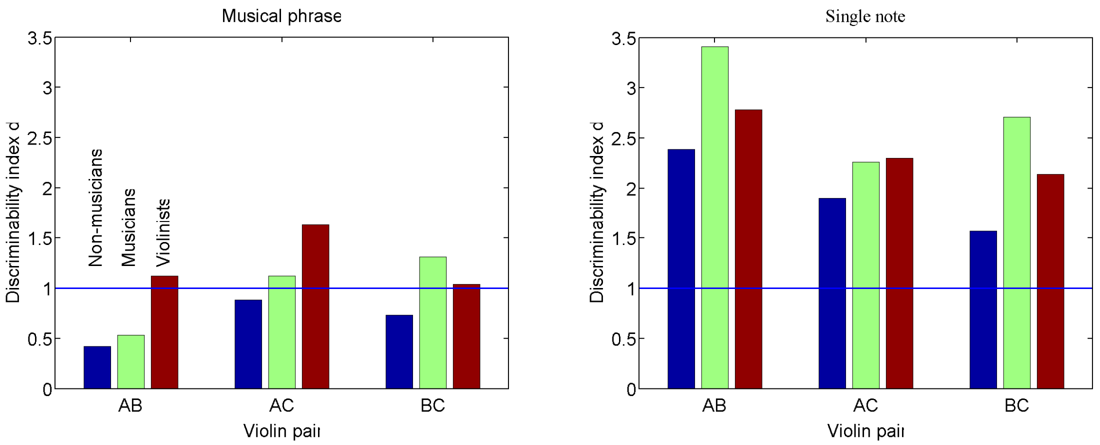

Figure 3. Discriminability index *d*’ for the three groups of test subjects and the three pairings of violins, based on a short musical phrase (left) and on a single note (right).

5 FUTURE PROGRAMME

The results of these preliminary experiment 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 ~~would presumably~~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. 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. This last approach is currently being tested in another undergraduate project, and the initial results are very encouraging. As one might guess, a player can be much more sensitive to changes in a virtual instrument than a listener who is outside the feedback loop involved in performance.

It is envisaged that exploratory tests will be made in which the various control parameters described in sections 2 and 3 are varied systematically. The initial goal will be to establish ~~JND information~~the just-noticeable difference for the various kinds of parameter, to give a first impression of the relative sensitivity of a listener or player. That will guide the priorities of later work, to explore first the most sensitive regions of the parameter space. The first set of tests will involve varying the modal parameters of key low-frequency modes. The natural candidates for a first systematic study are the air mode A0 and the two “main body modes” B1- and B1+[5,14], widely regarded as “signature modes” for the general behaviour of a violin.

The issue of sound radiation must not be forgotten, and for this initial study there is a very simple way to deal with it. A transfer function has been measured on a “typical violin” (violin A from the previous tests), from force at the bridge to a single microphone in a position which has been recommended to us by experienced recording engineers (0.5 m directly above the bridge), in an anechoic chamber. This contains at least some of the information about sound radiation in a welldefined way. The low modes A0, B1- and B1+ are all very clear in this transfer function. Using standard modal-fitting methods, these three peaks can be fitted and then subtracted from the transfer function, leaving everything else intact. Then a series of filters can be made in which these modes are added back in with the desired variable properties: frequency, amplitude and damping. It is encouraging that filters obtained in this way produce a much more convincing "violin sound" than was the case for the previous tests using the input admittance.

6 CONCLUDING REMARKS

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 ~~to~~ with judgements of "quality".

The methodology has been tested in several preliminary small projects, a sample of which have been described here. It appears that the method is robust, and capable of giving quantitative information on this important subject. There now opens a vista of many tests to be carried out, which in time will map out the perceptual landscape of the violin in an unprecedented manner.

7 ACKNOWLEDGEMENTS

The authors thank the Leverhulme ~~Foundation~~ Trust for financial support, and Professor Patrick Gaydecki for making available to us the real-time digital filter system and associated software. The undergraduate projects described in section 4 were carried out by Laura Kaye in the Department of Engineering, Cambridge University in 2003, and by Ulrike Petersen and Kevin Weaver in the Department of Music, Cambridge University during the academic year 2004/5. The musical passage recorded for input in the tests was performed by Paul Galluzzo.

8 REFERENCES

1. L. Cremer. The physics of the violin: MIT Press, (1985).
2. J. Woodhouse, 'On the playability of violins: Part 2 Minimum bow force and transients', Acustica 78, 137-153 (1993).
3. K. Guettler and A. Askenfelt, 'Acceptance limits for the duration of pre-Helmholtz transients in bowed string attacks', J. Acoust. Soc. Amer. 101, 2903-2913 (1997).
4. P. Gaydecki and B. Fernandes, '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 (2003).
5. J. Woodhouse, 'Body vibration of the violin - what can a maker expect to control?' J. Catgut Acoust. Soc. Series2, 4(5) 43-49 (2002).
6. L. Cremer, H.A. Müller and T.J. Schultz, Principles and applications of room acoustics, Applied Science (1982).
7. R.S. Langley and V. Cotoni. 'Response variance prediction in the statistical energy analysis of built-up systems', J. Acoust. Soc. Amer. 115, 706-718 (2004).
8. D. J. Ewins, Modal testing: theory and practice. Research Studies Press, (1984).
9. C.M. Hutchins, 'A 30-year experiment in the acoustical and musical development of violin-family instruments', J. Acoust. Soc. Amer. 92,639-650 (1992).

10. M. Schleske, 'On making "tonal copies" of a violin', J. Catgut Acoust. Soc. Series 2, 3(2) 18-28. (1996).
11. H. Dünnwald, 'Deduction of objective quality parameters on old and new violins', J. Catgut Acoust. Soc. Series 2, 1(7) 1-5. (1991).
12. E. V. Jansson, 'Admittance measurements of 25 high quality violins', Acustica - Acta Acustica 83, 337–341. (1997).
13. J. Woodhouse, 'On the "bridge hill" of the violin', Acustica - Acta Acustica 91, 155-165 (2005).
14. G. Bissinger, 'Sound radiation by violins radically modal", J. Acoust. Soc. Amer. 114, 2438 (2003).
15. G. Bissinger and F. Gearhart, 'A standardized qualitative violin evaluation procedure?', J. Catgut Acoust. Soc. Series 2, 3(6) 44-45 (1998).
16. M.V. Mathews and J. Kohut, 'Electronic simulation of violin resonances', J. Acoust. Soc. Amer. 53, 1620-1626 (1973).
17. S. Gorrill, 'A viola with electronically synthesised resonances', Catgut Acoust. Soc. Newsletter 24, 11-13 (1975).
18. M. E. McIntyre and J. Woodhouse, 'Toward a psychoacoustically realistic violin physics', Catgut Acoust. Soc. Newsletter 22, 18-19 (1974).
19. C. Gough, 'Physical aspects of the perception of violin tone', Proceedings ISMA 2001, Perugia, 117-122 (2001).
20. G. Weinreich, 'Directional tone color', J. Acoust. Soc. Amer. 101, 2338-2346 (1997).