The transient behaviour of models of bowed-string motion

R. T. Schumacher
Department of Physics, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213

J. Woodhouse
Department of Engineering, Cambridge University, Trumpington St., Cambridge CB2 1PZ, United Kingdom

(Received 12 July 1993; accepted for publication 4 April 1994)

Theoretical models of the action of a bowed string may be able to shed light on differences of "playability" between different violins. Subjective judgements seem to be concerned, at least in part, with the robustness with which one particular oscillation regime of the string (the "Helmholtz motion") may be obtained under different bowing conditions. In this paper, after a review of bowed-string modelling, systematic simulation is used to obtain plots of the basin of attraction of the Helmholtz motion in a particular subspace of the player's control space. Variations in the size and structure of this basin of attraction are seen when parameters of the problem are varied, and some physical interpretation of these variations is given. Some parallels and contrasts are pointed out between the particular features of the bowed string as a nonlinear system, and the range of more familiar dynamical systems. © *1995 American Institute of Physics.*

I. INTRODUCTION

When a bow is drawn across a violin string, a noise of some kind is invariably produced. Not all such noises are, however, acceptable to the violinist—there is a wide repertoire of whistles, crunches and screeches which can be produced if the bow is not controlled in a suitable way. There is strong anecdotal evidence that some instruments, or some particular notes on a given instrument, are more susceptible than others to such undesirable oscillation regimes. These instruments or notes require more careful control of the bow, and are naturally perceived by players as being "difficult to play." Theoretical modelling of the action of a bowed string has reached a sufficient level that one would expect to be able to shed light on the mechanics of such variations in "playability." A framework for such investigation has been set out in two previous papers,[1,2] and the aim of the present paper is to take up in more detail one of the lines of enquiry initiated there. This concerns the systematic use of numerical simulation to explore the robustness of the desired oscillation regime (the "Helmholtz motion"[3,4]) under variations in the detailed form of initial bowing transient.

A violin string is a rather complicated linear system which is set into self-excited vibration by the action of stick-slip friction between the bow and the string. It falls within the broad class of non-linear dynamical systems, and exhibits many of the phenomena which have become familiar from the extensive investigation of such systems in recent years. However, any model of a bowed string which comes close to realism differs in both style and complexity from the (generally rather simple) systems more commonly studied. In view of the complexity of the system, and especially since the main interest in practice concerns transient, non-autonomous behaviour, it does not seem possible at present to make significant progress other than by numerical simulation. The parameter space of this problem is very large, and in some cases the range of variation of the parameters within which musically acceptable results may be obtained is also very large. It is thus perhaps not surprising that relatively little progress has been made on the problem, even by using simulation studies, although it is now some 15 years since efficient simulation algorithms became available.[5] The reason that some progress is being made now lies in the recent availability of massively-parallel computing engines. By using, typically, 8192 processors of a Connection Machine[6] to run separate simulations simultaneously, it is possible to study the behaviour in a two-dimensional patch of parameter space and see structure in the results which is not readily discernible from serial simulations.

To carry this scheme through it is necessary to be able to classify the outcome of a given transient simulation automatically, since it is clearly not possible to examine the detailed output from each processor manually. The requirement is to recognise whether a periodic regime has been achieved, and if so, which one it is of the wide range of possibilities. The algorithm developed for this will be described in section IIIB. Finally, in section IV some examples are presented to illustrate the results of simulation experiments using models with different parameters. These results are discussed in rather general terms, with the aim of classifying the qualitative types of behaviour observed. The way in which results are displayed can make a significant difference to how much information is conveyed. What seems to be needed is a hierarchy of "pattern recognisers," which distill increasingly high-level, compressed information out of the large volume of individual results. Some suggestions for developments in this direction will be made.

Some interesting inferences can be drawn from the particular results shown here, but the main aim of the present study is to establish a framework for later studies, both in terms of methodology and in terms of the natural groupings of model parameters which reflect different aspects of the physics of the problem. The rather large number of parameters can be divided into those associated with (a) the string; (b) the instrument body; (c) the stick-slip friction constitutive law; and (d) the player's control variables when producing a particular bowing transient. The results to be presented here explore the behaviour in a subset of the space of player's

control variables, category (d). It is hoped that the behaviour in this space reveals something about the robustness or structural stability of the Helmholtz motion under differing bowing transients. The question one would then like to address is "How is this robustness influenced by variations in the parameters (a) and, more particularly, (b)?" This is the formalisation of the question "Why and how does it happen that one violin, or one string, is perceived as easier to play than another?" A series of later papers is envisaged, in which various specific aspects of that question will be investigated using the procedures set out here.

II. MODELLING THE BOWED STRING

A. Summary of basic theory

The modelling of bowed-string motion has been described in detail in previous papers.[1,2,5,7–9] The methods adopted have exploited features specific to this problem, and the models do not relate very directly to the more familiar styles of model of other nonlinear dynamical systems. A brief summary will be given here, introducing the various stages of modelling which incorporate progressively more of the physics of the real system. This relatively lengthy account of modelling is important—our aim is to use some insights of dynamical-systems theory to explore a physical problem, and there is no point in studying models which are over-idealised to the extent of falsifying vital aspects of the physical behaviour. On the other hand, of course one wants the simplest acceptable model. This means that an essential part of the investigation is to be aware of the range of possible physical models, and to develop the model under study iteratively when results show unphysical behaviour. Some individual models may have mathematical interest in themselves, but we are not interested here in following up such suggestions.

All the models to be discussed here share certain idealisations. The most important of these is that the "bow" is assumed to act at a single point on the string, so that any effects of the finite width of the ribbon of hair on a real bow are neglected. With this assumption, the essence of the model is very easy to describe. It is framed in terms of two time-varying quantities relating to behaviour at the bowed point: the friction force $f(t)$, and the string's velocity $v(t)$ at that point. The quantities $f(t)$ and $v(t)$ are connected in two distinct ways, and when these are enforced simultaneously the governing equation for the system is obtained.

The first connection between f and v is through a nonlinear relation describing the tribology of stick-slip friction. We will consider here only models based on the simplest acceptable idealisation of frictional behaviour, in which the force and the velocity are related through an instantaneous functional relation of the kind shown in Fig. 1. (In a more realistic model, the force would depend not only on the current value of the velocity but also on some aspects of the earlier history of the relative motion between the contacting surfaces.[10,11]) The vertical portion of the curve represents the indeterminacy of friction force when there is sticking between the two surfaces. This occurs when the string velocity is equal to the speed v_b of the bow. The curved portion of the

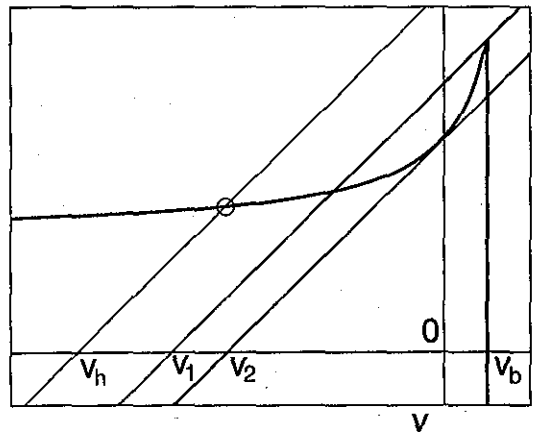

FIG. 1. The heavy curve shows the relation between friction force f and string velocity v at the contact point. The straight line and ringed intersection illustrate the approach to solving $f(v)$ simultaneously with Eq. (3). The other two straight lines demarcate the limits of the ambiguous region for this solution, as discussed in section IIIC.

f–v relation represents the behaviour when there is relative sliding between the surfaces, leading to a reduction in the friction force. The friction force at a given sliding speed is assumed to be directly proportional to the normal force f_b between bow and string, the familiar notion of a "coefficient of friction."

The values of f_b and v_b may vary in time: they are two of the three main "player's control variables" (the third being the position of the bowed point along the length of the string). For the purposes of modelling the string motion, they will be regarded as given functions of time. Particular time histories of f_b and v_b will characterise any particular bowing transient. For example, a martelé transient involves a bow force f_b which starts high and falls rapidly to its eventual steady value, while at the same time the bow speed increases from near-stationary to a final steady speed. By contrast, a string-crossing transient involves an already-moving bow alighting on a string, so that the bow speed will be essentially constant while the bow force increases from zero to a final steady value. It will be seen later that the theoretical models agree with the experience of players in that two such contrasting transients, even if they have the same final values of f_b and v_b, do not necessarily result in the same final regime of oscillation of the string. Probing such behaviour is the main objective of this study.

The second connection between $f(t)$ and $v(t)$ is through the linear-system dynamics of the string and violin body. The velocity response at the bowed point of the string can be calculated from the past history of the force applied there, via a convolution integral with the appropriate impulse response function, $g(t)$ say. This impulse response function can be decomposed into several components.[1,5,7,11] It is convenient first to separate an initial delta-function contribution, representing the velocity response of an infinite string to an applied force impulse: write

$$g(t)=\frac{Y_0}{2}\,\delta(t)+\tilde{g}(t) \tag{1}$$

where Y_0 is the characteristic admittance of the string at the bowed point [equal to $(Tm)^{-1/2}$ for transverse motion of an ideal string of tension T and mass per unit length m]. The function $\tilde{g}(t)$ represents the effects of reflections from the ends of the string.

Putting the two relations together, the governing equation for the system is now

$$v(t)=\int_{-\infty}^{t} g(t-\tau)f(v(\tau))d\tau \tag{2}$$

where $f(v)$ is the friction function. Using Eq. (1), this may be written

$$v(t)=\frac{Y_0}{2}f(t)+v_h(t) \tag{3}$$

where

$$v_h(t)=\int_{-\infty}^{t} \tilde{g}(t-\tau)f(v(\tau))d\tau. \tag{4}$$

The quantity $v_h(t)$ depends only on the past history of the motion, so that in a time-stepping simulation it is a known quantity at a given time t. The way that such a simulation proceeds is that at a given time step $v_h(t)$ is calculated, and then the new values of f and v are found by solving Eq. (3) simultaneously with the friction relation $f=f(v)$. This process is most easily visualised as the graphical construction illustrated in Fig. 1.[7,12,13]

If we begin with a model which only allows for transverse motion of the string, $\tilde{g}(t)$ can be written in terms of two "reflection functions" representing the combined effects of propagation, dispersion, dissipation and boundary reflection on the two separate sections of string to the left and the right of the bowed point.[1,5] Denoting these two functions $h_1(t)$ and $h_2(t)$, we may write

$$\tilde{g}(t)=\frac{Y_0}{2}[h_1(t)+h_2(t)+2h_1*h_2+h_1*h_2*h_1 + h_2*h_1*h_2+\ldots] \tag{5}$$

where * denotes the operation of convolution. This decomposition considerably improves the efficiency of a simulation. The calculation of $v_h(t)$ may be done by storing the outgoing velocity waves on the string travelling to the left and to the right from the bow, and performing separate convolutions of these with $h_1(t)$ and $h_2(t)$ respectively to generate the reflected waves arriving at the bow from the two sides.[5] The sum of these two contributions gives the total incident velocity at the bowed point, which is $v_h(t)$. The significance of this approach to the computation is that the reflection functions may be of quite short duration, whereas $\tilde{g}(t)$ typically continues for some hundreds of period-lengths. Thus the time spent calculating convolution integrals is greatly reduced by using the separate reflection functions.

B. A hierarchy of physical models

To obtain a particular model within the class under discussion here, it is necessary to specify the friction relation $f(v)$ and the function $\tilde{g}(t)$. For a model involving only transverse string motion, the latter is best done via the two reflection functions $h_1(t)$ and $h_2(t)$. A variety of models may be obtained in this way, depending on just which aspects of the physics of real strings and bows are taken into consideration. Even the simpler of these models have not yielded very much to analytic treatment. A certain amount can be learned about the rich set of possible periodic solutions, both their morphology and the regions of parameter space within which they are possible and stable,[9,12,14] but that is about all. The governing equation for the system takes the form of neither a set of nonlinear differential equations nor (except in a very special case, see Appendix A of Ref. 5) an iterated map, and such progress as has been made makes explicit use of the particular features of this system, and seems to have little in common with the main body of dynamical-systems theory.

We obtain what is mathematically the simplest model for a bowed string by considering a classical text-book string, rigidly anchored at both ends. In that case, both reflection functions consist of unit delta functions, delayed by the respective travel times from the bowed point to one end of the string and back. This model was first discussed in detail by Friedlander[12] and Keller.[13] Friedlander proved that this model has a strikingly unphysical characteristic. The condition that

$$df/dv>0 \tag{6}$$

must be imposed on the friction relation, to guarantee that the non-oscillatory solution f=constant, $v=0$ is unstable, since it is never (or at least almost never[15]) observed in practice. This turns out to be a sufficient condition for all periodic solutions to be unstable as well, an unsatisfactory state of affairs. The reason is that Friedlander's model has no mechanism of energy dissipation, whereas the sliding bow can act as an energy source when (6) is satisfied. Clearly, any model which is to be considered as a candidate to describe real bowed strings must be capable of supporting stable periodic solutions, and so must have energy dissipation accounted for in one way or another.

There are two physical mechanisms of dissipation. Most obviously, there is some loss of energy during the propagation and reflection of waves from the bowed point to the ends of the string and back. This can arise from internal dissipation in the string itself, or from losses into the violin soundbox (which is after all there in order to draw energy from the string vibration and radiate some of it as sound). These effects can all be allowed for by modifying the reflection functions. The simplest such modification produces "Raman's model,"[9,16] which retains the delta-function reflections of Friedlander's model but allows some energy loss by having a reflection coefficient less than unity. However, Raman's model produces results which are still quite unrepresentative of real string behaviour:[9,11] in particular, the waveforms for $f(t)$ and $v(t)$ take the form of piecewise-constant functions, with sudden jumps between successive constant values.

Instead, we must allow reflection functions of finite width, so that the waveforms $f(t)$ and $v(t)$ turn int' smoother functions. If the reflection functions are given a width which is finite but still narrow compared with the natu-

ral oscillation period of the string, they provide energy loss without significantly upsetting the precise harmonic relation between the mode frequencies of the string.[1,2] A further stage is to use models which provide for perturbations in these mode frequencies as well as energy loss: the main effects contributing to such perturbations are coupling to vibration modes of the violin body, and wave dispersion due to bending stiffness in the string itself. Reflection functions which can account for both these effects have been described previously.[1,7]

The second mechanism of energy loss is via torsional motion of the string. So far, only transverse string motion has been considered. However, the friction force between bow and string acts tangentially at the string's surface, so that it will excite both transverse and torsional motion. The torsional waves are presumably not responsible for significant sound production, but what little data there is suggests that they are much more highly damped than the transverse waves,[17] and they play an important role in the overall energy budget of the system. Torsional energy loss may be allowed for in two ways. The simpler makes the rather gross assumption that torsional waves are generated at the bow, but that their reflections from the ends of the string may be ignored as if the waves were totally absorbed there. In that case, a simple modification to the friction-velocity curve involving a linear shear is all that is necessary.[5,18]

However, in reality the torsional waves are reflected. The model described above can be readily extended to allow for this, with only a slight increase in complexity. The function $\tilde{g}(t)$ must now contain contributions from both wave-types on the string. To achieve this, a second pair of reflection functions to describe torsional reflection can be introduced, in exactly the same way as those used to describe transverse reflection.[5] (It might also be necessary for complete realism to introduce "cross-reflection functions," to allow for transfer between transverse and torsional motion during propagation and reflection.) The outgoing torsional waves to the left and right of the bowed point are stored, and incoming torsional waves calculated from them by convolution with these new reflection functions. Once this step has been taken, then a wide range of linear-system behaviour of the torsional waves can be represented by choosing suitable reflection functions, just as in the case of the transverse waves.

The governing equation is made only slightly more complicated by including torsional waves in this way. The string surface velocity $v(t)$ is now made up of two components:

$$v(t)=v_t(t)+v_r(t) \tag{7}$$

where $v_t(t)$ is the transverse velocity at the string centre, and $v_r(t)$ is the transverse velocity at the string's surface caused by rotation alone. The total incoming velocity $v_h(t)$ is now a sum of four terms, calculated by performing the appropriate convolutions of the outgoing waves with the reflection functions for transverse and torsional waves on each side of the bow. The friction force $f(t)$ still obeys Eq. (3), except that the value of the constant Y_0 is now the sum of the transverse and torsional wave admittances. This force generates outgoing transverse and torsional velocity waves, the amplitude of each being governed by the corresponding wave admittance. (This torsional wave admittance is in the same units as the transverse wave admittance, since it has been defined in terms of the tangential linear velocity at the string's surface in response to a tangential force.)

It seems likely that a fairly good representation of bowed-string behaviour might be given by a suitable model embodying the features described, and the study described here will be confined to such models. However, one should be aware that many aspects of the physical system have been idealised, and relaxing these idealisations might prove necessary at a later stage in the study. Examples might be a more sophisticated constitutive law for friction, or a better treatment of the contact mechanics near the bow, allowing for the finite width of bow hair, local bending stiffness of the string, bow-hair elasticity and so on.

III. SIMULATION AND REGIME CLASSIFICATION

A. The delta/Cremer and stiff/Cremer models

We will present results in this paper for two models which retain reasonable simplicity while representing at least some aspects of the behaviour of real strings. Both models share the same reflection function for the short section of string between bow and violin bridge. (A violin string is always bowed at a point fairly close to the bridge.) This is derived from a model due to Cremer,[19] in which the free decay times for the first 13 overtones of an open violin A string (440 Hz) were fitted to the frequency dependence given by a simple spring/dashpot model for the termination at the bridge. The reflection function consists of a delta function followed by a small exponential "tail."[1] For one of our two models, the "delta/Cremer model," the reflection function for the other end of the string (the long section from the bow to the nut or player's finger) is simply a unit delta function, representing perfect reflection on an ideal string. When these ingredients are combined, the first few period-lengths of the impulse response function $g(t)$ for transverse motion are as plotted in Fig. 2(a). The exponential tails following the delta functions are only barely visible here, but nevertheless they introduce the desired level of dissipation.

The second model has a more complicated reflection function for the long section of string. In order to simulate the dispersive effects of bending stiffness on the string, a function is used whose derivation and functional form were given in Ref. 1, the precise function to be used here being plotted as Fig. A2 in that reference. Dispersion would be negligible on the short section of string to the bridge, so it does not seem necessary to modify the Cremer-model reflection function. When these two functions are combined according to Eq. (5), the first few cycles of the transverse impulse response $g(t)$ are as plotted in Fig. 2(b). We will refer to this as the "stiff/Cremer" model.

The final ingredient of both models concerns the treatment of torsional motion. We use a simple treatment which comes reasonably close to mirroring real behaviour. Torsional reflections are taken into account by using a model with identical Gaussian functions for both reflection functions. The width and area of the Gaussians were chosen to produce reasonable damping factors for the first few tor-

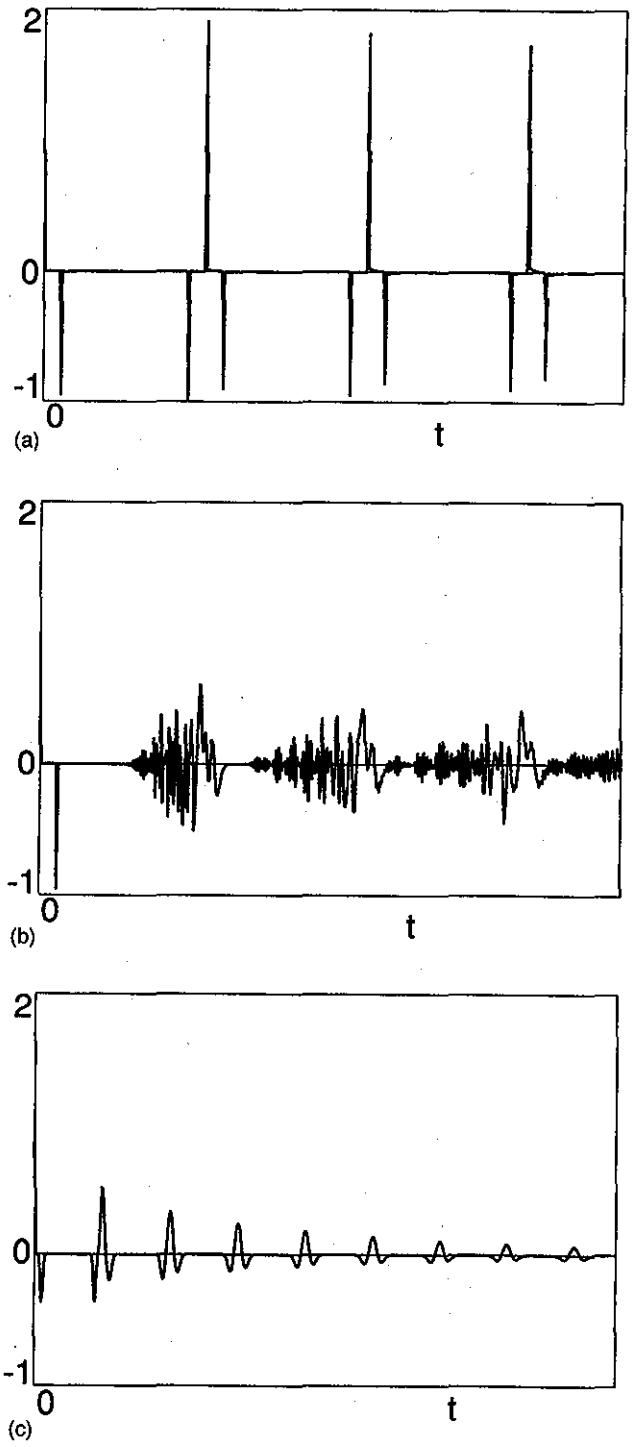

FIG. 2. (a) The start of the time history of the impulse response function $g(t)$ for the delta/Cremer model; (b) the corresponding plot for the stiff/Cremer model; (c) the corresponding plot for the impulse response function for torsional motion as assumed in this study.

sional modes (see the Appendix for details). Of course, it is necessary to give them different time delays appropriate to the two lengths of string, and to allow for the fact that the torsional wave speed is different from the transverse wave speed. Values for the ratio of torsional-to-transverse wave speeds and admittances were taken from measurements on a particular violin D string (Dominant aluminium-

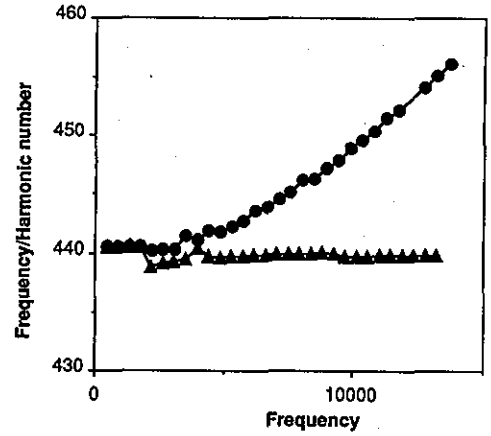

FIG. 3. Mode frequency over harmonic number for the first 30 overtones of the string transverse motion, for the delta/Cremer model (triangles) and the stiff/Cremer model (circles).

on-perlon):[20] the values are 2.4 and 0.617 respectively. To represent the appropriate time delays in the discrete form of the reflection functions needed for numerical implementation, use was made of a result (proved in section 3 of Ref. 1) showing that the effective centre of a narrow reflection function may be calculated by requiring the first moment of the function to vanish. The first few period-lengths of the resulting torsional impulse response function are plotted in Fig. 2(c), on the same time scale as Figs. 2(a) and 2(b). The higher speed of the torsional waves is immediately apparent. There is assumed to be no interaction between torsional and transverse motion (except that induced by the friction force at the bow).

To show what linear-system behaviour is represented by these models, we may look at the corresponding modal frequencies and damping factors. The frequencies of the first 30 transverse modes of the delta/Cremer and stiff/Cremer models, with the parameters used in the simulations to be shown later, are plotted in Fig. 3. (What is actually plotted is the ratio of the nth mode frequency to the harmonic number n, in order to reveal clearly any departures from harmonicity.) The corresponding Q factors are shown in Fig. 4. For the

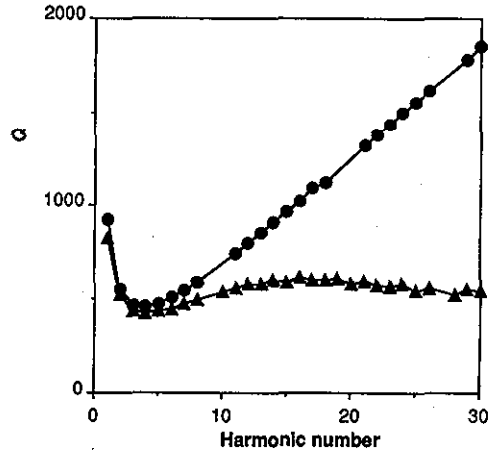

FIG. 4. Q factors for the first 30 overtones of the string transverse motion, for the delta/Cremer model (circles) and the stiff/Cremer model (triangles).

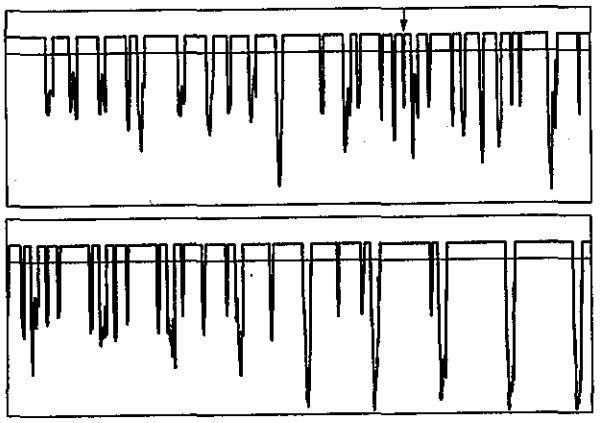

FIG. 5. The first 2400 time-samples of the string-surface velocity for a transient solution of the delta/Cremer model. The two boxes show sequential segments: the upper box shows samples 1-1200, the lower box samples 1201-2400. The horizontal segments of waveform correspond to sticking, at the bow speed v_b. The bow force had the constant value 5.5 throughout, and $v_{\text{mid}}=0.7$ (units and other details are described in the Appendix).

torsional motion, the damping is much higher so that fewer modes could be measured. The first few frequencies are quite accurate harmonics of $2.4\times440=1056$ Hz, and the first few Q factors are 25, 43, 45, 45, 40, 37, 33, and 29 (440 Hz being the assumed fundamental frequency of transverse motion).

The two models show reasonably realistic behaviour. The delta/Cremer model has modal damping factors for both transverse and torsional behaviour which are of the right order of magnitude, the torsional modes being much more highly damped. Because this model uses narrow reflection functions, both the transverse and torsional mode frequencies form quite accurate harmonic series. The dispersive effect of bending stiffness in the stiff/Cremer model perturbs the harmonic series for the transverse modes. The characteristic sharpening of the higher modes, following approximately a quadratic function of mode number,[21] is clear in Fig. 3. The Q factors for the lower modes are approximately the same for both models, since the reflection function is intended to represent dispersion alone, without dissipation. In practice, as has been explained previously,[1] a filter must be applied to the reflection function to avoid aliasing problems in the discrete realisation, and the effect of this filter is to produce additional damping of the higher modes. The Q factors are still of the right order of magnitude for real strings—if anything, the stiff/Cremer model gives rather more realistic Q factors for the higher modes. This version of the stiff/Cremer model seems a good candidate for a benchmark case, since it includes a good first approximation to the main physical effects.

B. Classification and automatic recognition of results

An example of a simulated initial transient, using the delta/Cremer model, is shown in Fig. 5. The string was at rest in its undisturbed position at time $t=0$, and constant values of f_b and v_b were then imposed. The model equation, in discrete form with approximately 139 time samples per nominal period length, was then iterated forward in time as described above. (Further details of this and all other simulations are given in the Appendix.) The particular parameter values were chosen such that the desired "Helmholtz motion" oscillation regime was achieved after a relatively short transient, of some 16 period-lengths. The periodic solution is just becoming established by the end of the time-window plotted in Fig. 5, which shows the first 2400 samples of the string surface velocity waveform $v(t)$. Little apparent structure can be discerned in the transient, just a rather complicated sequence of intervals of sticking [when $v(t)=v_b$] and slipping [when $v(t)<v_b$]. Before further discussion of transients, we consider the simpler question of classifying the periodic oscillations which arise in such simulations.

Figure 6(a) shows approximately three cycles of $v(t)$ for the stable periodic solution which eventually became established. Figure 6(b) shows the same three cycles of the velocity waveform $v_t(t)$ of the string's centre. These two differ because of the effect of torsional motion. The final waveform plotted, in Fig. 6(c), is that of the outgoing transverse velocity wave from the bow towards the bridge. This will be an important quantity for determining which of the possible oscillation regimes has been achieved, and it is also, to a first approximation, the quantity responsible for the sound of the violin. The transverse force exerted by the vibrating string on the violin bridge will have this waveform. This force in turn drives the vibration of the structure of the violin body and thence sound waves in the surrounding air.

Figure 6(a) shows very clearly an alternation of sticking and slipping once per cycle in the periodic motion. The particular oscillation regime achieved here is the "Helmholtz motion," the only one of the possible regimes which is normally acceptable to violinists. The motion is characterised by a single sharp "corner," or velocity jump, which travels back and forth between the ends of the string and triggers transitions between sticking and slipping each time it passes the bow. This is seen most clearly in Fig. 6(c), which we will refer to as the "bridge-force waveform." This waveform has an obvious sawtooth shape, and the sharp "flyback" of the sawtooth represents the passage of the Helmholtz corner.

In a pioneering study of bowed-string motion, Raman[16] argued that all periodic motions of a bowed string might be expected to exhibit travelling "corners" of this type. His argument may be summarised as follows: (i) for any periodic motion with a period close to the string's natural period, the frictional force must be approximately constant, since the damping is light so that significant force fluctuations would evoke a very large resonant response; (ii) with a friction function like that of Fig. 1, it follows that the velocity $v(t)$ must alternate between v_b and a single sliding velocity determined by the value of the constant friction force; (iii) from the D'Alembert solution for waves on an ideal string, the individual travelling velocity waves must thus have a (more or less complicated) "sawtooth" form, with a ramp of fixed slope and occasional jumps of fixed magnitude, in other words "corners." Raman developed a classification of the possible periodic oscillation regimes in terms of the number of corners on the string which they entail. The Helmholtz motion is in these terms the simplest oscillation, with only one corner. Other regimes, called by Raman "higher types," involve larger numbers of corners. Raman's scheme pro-

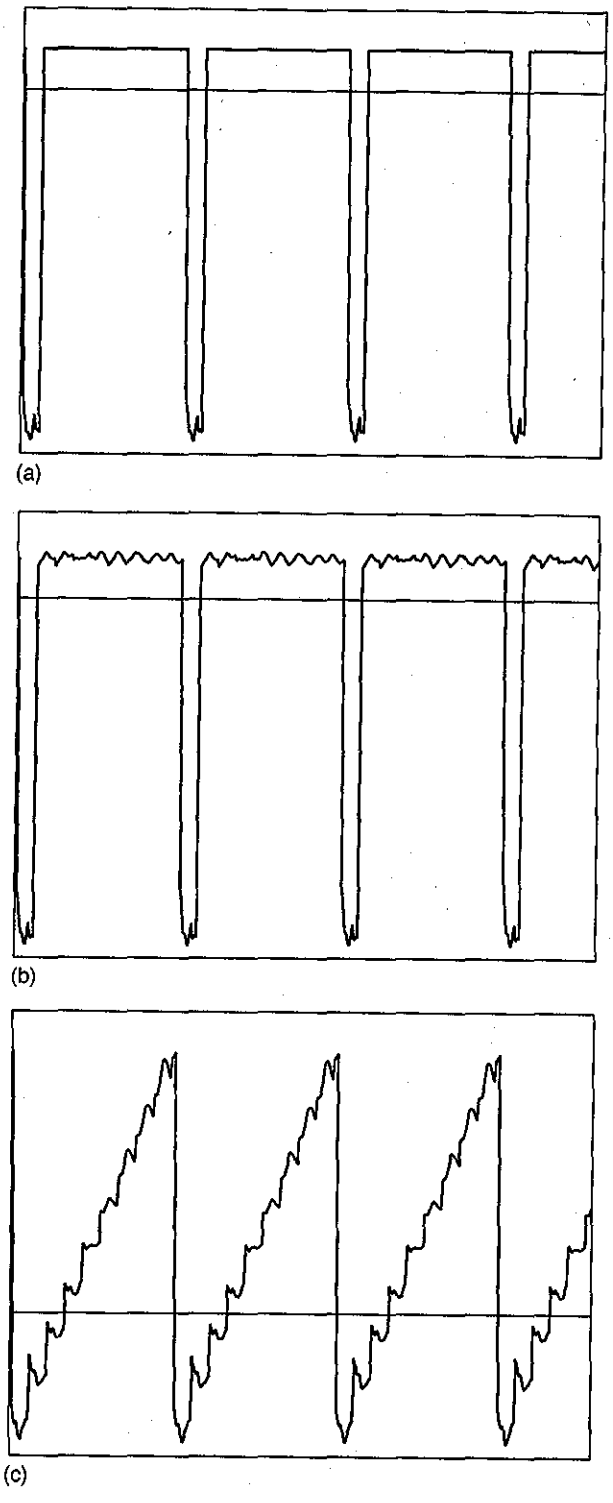

FIG. 6. A portion of the eventual stable periodic solution to the transient shown in Fig. 5, showing (a) the string surface velocity; (b) the string centre velocity; and (c) the outgoing velocity wave towards the bridge, or "bridge force."

duced good agreement with experimental results available in those days,[16,22] and we would expect to be able to use it as a basis for automatic classification of the oscillation regimes arising from simulations.

When this simulation program is run with a variety of parameter values, a wide range of behaviour is seen. As well as the Helmholtz motion, many other periodic regimes can occur. In the great majority of cases, these belong quite recognisably within Raman's scheme. The exceptions of which we are aware are associated with particular values of the torsional-to-transverse wave speed ratio, a complication not considered by Raman, and we will not discuss them in this paper. However, not all runs of the program yield periodic regimes. Particularly with high values of bow force f_b, persistently aperiodic solutions are found. Such solutions are at least in qualitative agreement with the experience of playing the violin—by pressing too hard with the bow, it is only too easy to produce an unmusical noise of some kind. Whether such behaviour represents "chaos" in any strict sense is not known, but it seems likely.

For the purposes of the study reported here, an algorithm is needed which can recognise an acceptable Helmholtz motion by processing the output from a given simulation run. There are three defining characteristics which we can seek to detect: (1) the motion should be periodic; (2) there should be just one episode of slipping per period; and (3) there should be just one Raman travelling corner involved. Separate tests are made for these three conditions. Each poses problems in its detailed implementation, and we address them in turn.

The natural way to test for periodicity is by correlation analysis. If the correlation coefficient between one period and the previous one, based on a suitable test waveform such as bridge force, exceeds a threshold which is close to unity, then presumably a more-or-less periodic motion has been achieved. (Of course, the value of the threshold must be fine-tuned using many simulation runs whose outcome is inspected in detail, so that the criterion for acceptance is in line with what seems intuitively correct.) This correlation analysis can be implemented without difficulty if the period is known and is a whole number of samples. However, neither of those conditions is met in general. With the delta/Cremer model, with its narrow reflection functions, the nominal period is quite easy to predict. There is a frequency-independent "end correction," given by the vanishing of the first moment of the reflection function.[1] However, there is no equivalent simple result for the stiff/Cremer model. The effective reflection time is far from obvious [see Fig. 2(b)], and in any case the inharmonicity induced by the bending stiffness leads one to expect the pitch to vary somewhat depending on the values of parameters such as bow force, just as happens with woodwind instruments.[23] There is also another effect which makes pitch hard to predict in advance: it has been shown that frictional hysteresis (to be discussed shortly) causes a systematic lowering of pitch as the bow force is increased.[5,7,18,24]

These effects all mean that a correlation analysis is required which can detect "periodic" signals whose period is not known in advance, and which will not in general be an integer number of time samples. The approach adopted was as follows. A nominal period is assigned, then the correlation coefficient is calculated over a range of one sample above and below this nominal value using the quadratic interpolating polynomial on triads of adjacent data points. The maximum value of this coefficient as a function of the offset from the nominal period can be found analytically. If a maximum

occurs within the allowed range, on either side of the nominal period, then the value of that maximum is taken as "the" correlation to be compared with the threshold value. The algorithm can be extended to encompass a wider range away from the nominal period, simply by repeating the calculation with the nominal period increased or decreased. The overall maximum value found within the range searched can then be taken as the best approximation to the desired answer. As a side effect, this analysis gives an estimate of the period length for non-integer periods.

The second test for Helmholtz motion involves counting the number of episodes of slipping per nominal period length. This is essentially a trivial operation, since in order to implement the piecewise-defined friction curve (and the hysteresis rule) the program keeps explicit track of whether the current state is sticking or slipping. However, a little care is needed if the result of the automatic algorithm is to agree with a human judgement based on watching the simulation proceed. It is common to find solutions which are "almost Helmholtz," but with very brief slips during the nominal sticking phase. These often appear to be of little dynamical consequence. The bridge force signal looks very much like the normal sawtooth, and one suspects that if such a regime occurred on a real violin, then a listener would judge it an acceptable Helmholtz motion (although of course there may well be audible consequences of the extra slips). So the criterion adopted was to count the number of "significant" slips in a five-period interval, where "significant" was defined by setting a threshold level for the integrated velocity during one slip interval. The threshold was tuned so that the acceptance criterion gave subjectively reasonable results.

The hardest of the three tests to implement robustly was to count Raman travelling corners. The precise form of the Helmholtz motion, and the other periodic regimes, varies with the parameters of the problem, in a way which has been discussed extensively in previous works. There are two main effects. First, any reflection functions of finite width cause the Helmholtz corner to have a finite extent (i.e. to be "rounded," so that these models are collectively known as "rounded-corner models"). The degree of roundedness varies systematically with bow force f_b: higher forces produce a sharper corner, lower forces a more rounded corner.[7,25] The second departure from the idealised Helmholtz motion is evident in Fig. 6(c): on the rising ramp of the sawtooth wave, there is a superimposed oscillation. This arises from velocity perturbations which are to a large extent trapped in the short section of string between bow and violin bridge, so that the period of the oscillation is governed by the distance of the bow from the bridge. The phenomenon is usually known as "Schelleng's ripples."[26] The magnitude of these ripples is also sensitive to bow force f_b: larger forces produce greater velocity perturbations (when the slightly rounded Helmholtz corner interacts with the bow[8]).

These effects make it rather tricky to implement a robust algorithm to count the number of Raman travelling corners in a given solution. The chosen method is based on the slope of the bridge-force waveform. A simple threshold test on the slope is certainly not adequate: the Schelleng ripples at high bow force produce larger values than does the Helmholtz corner itself at low bow force. However, the effect of ripples can be reduced by averaging, taking advantage of the known period of the ripples. A simple moving average of the derivative of the bridge force waveform, using a top-hat weighting function with a width equal to the ripple period, produces a signal which can be tested against a threshold value to detect Raman travelling corners. As with the other two tests, the value of the threshold must be fine-tuned to give satisfactory results.

When these three tests are combined, and their respective threshold values carefully adjusted, the discrimination of Helmholtz motion from other outcomes of a given transient simulation was found to be sufficiently reliable. The combined test is certainly not perfect, but we will argue in section IV that it discriminates the types of behaviour that we are most concerned about in this study.

C. Transients and divergence of solutions

The motivation for this entire study is the idea, stemming from the experience of string players and from watching simulations, that the outcome of a given bowing transient might depend quite sensitively on the details of the imposed bow force and bow speed time-histories. Furthermore, it is thought that the *degree* of this sensitivity might be influenced by some aspects of the linear-system behaviour of the string or instrument body, leading to perceived differences in "playability." Before showing any systematic computational results, it is useful to enquire whether such sensitivity might be anticipated from the general description of the models which we have already given. If a possible source of sensitivity can be identified, that would provide guidance on where to search in the multi-dimensional parameter space to maximise the chance of finding a significant effect.

It is commonplace in dynamical systems theory that such sensitivity would stem from some mechanism which causes "neighbouring" solutions to diverge from one another. The stronger the divergence, the more sensitivity can be anticipated. Two mechanisms of divergence can be readily identified, both associated with the friction function. The first of these effects stems from the fact that the gradient df/dv is positive on the portion of the curve corresponding to slipping. When carrying out a linearised stability analysis of any given solution, periodic or not, this positive gradient acts as a "negative resistance,"[12,26,27] an energy source which can amplify a small velocity perturbation. Linearised stability analysis for periodic solutions to a range of bowed-string models has been discussed in some detail elsewhere.[14] In all cases, the threshold of stability of any particular solution is determined by a condition of balance between energy loss (via the reflection functions or into torsional motion) and energy gain from this negative resistance at the slipping bow. The possibility of instability, in this linearised sense, obviously implies the possibility of divergence of nearby solutions, although with realistic parameter values this is usually a rather slow process.

A much stronger mechanism of divergence of solutions arises from an aspect of the graphical construction shown in Fig. 1, whereby the new values of $f(t)$ and $v(t)$ are found from a knowledge of $v_h(t)$. It is plain that if $v_h(t)$ falls

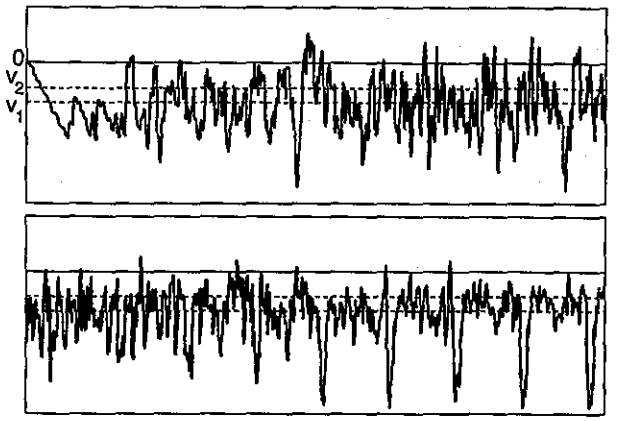

FIG. 7. The waveform of incident velocity $v_h(t)$ for the transient shown in Fig. 5, plotted over the same time range and in the same format. The two dashed lines show the limits of the ambiguous range within which hysteresis operates.

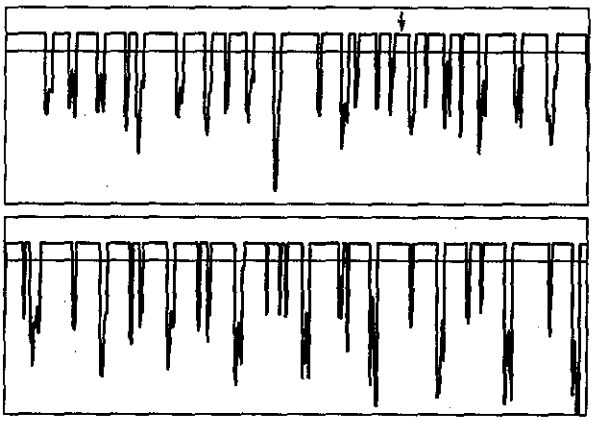

FIG. 8. String-surface velocity in a transient differing from that of Fig. 5 only in that the (constant) value of the bow force was 5.3. The two plots have identical scales and formats. The waveforms diverge strongly at the time indicated by an arrow.

between the two values v_1 and v_2 (shown on Fig. 1), there is an ambiguity of intersections of the straight line with the friction curve. The resolution of this has been shown to be a hysteresis rule, as one might guess.[7] Slipping continues until v_h reaches v_2, and when it passes through this value there is a jump to a state of sticking. Sticking then continues until v_h reduces as far as v_1, and there is then a larger jump to a state of slipping just outside the ambiguous region. This hysteresis has been shown to be responsible for a fall in pitch of a bowed note when the bow force is increased.[5,7,18,24]

More important for the present purpose, hysteresis can obviously cause a strong divergence of solutions. Suppose an occasion arises during a transient when v_h rises to approach v_2 (during slipping), and then reduces again. The solution will continue in a state of slipping. But a small perturbation might cause v_h to go just beyond v_2. Then there will be a jump to a state of sticking, from which there is no escape until v_h decreases all the way to v_1. By then, the stored outgoing velocity waves will have been changed sufficiently that the entire subsequent development of the solution is likely to be different. A similar drastic effect of a small perturbation can, of course, be envisaged whereby v_h is carried just below v_1 so that it slips, when v_2 the unperturbed case continued to stick. Again, hysteresis means that there is no easy recovery from this change, since when v_h rises again, it will not jump back to the sticking state until it reaches v_2.

This behaviour can be illustrated with the transient shown in Fig. 5. Figure 7 shows a plot of $v_h(t)$ for that transient, with the two critical levels v_1 and v_2 indicated. It is easy to see that the variation of $v_h(t)$ is sufficiently strong, and apparently random, that one or other of the situations envisaged in the previous paragraph might arise, so that a solution with slightly different parameters might follow this one closely for a while, then be diverted away from it. Figure 8 shows an example of this occurring: it shows a simulation identical to that of Fig. 5 except that the value of f_b is reduced by approximately 4%. Careful comparison of the two plots shows that they remain very similar until the time indicated by an arrow, when Fig. 5 shows a slip while Fig. 8 continues to stick. After that the two transients rapidly diverge. In fact, this particular pair both lead to a stable Helmholtz motion in due course, after transients of different lengths.

It appears that frictional hysteresis might play a role in this system somewhat analogous to that of saddle points in dynamical systems based on differential equations. If qualitatively different solutions are separated by a transition of the kind just described, involving an extra stick/slip transition, then the boundary separating the basins of attraction of these two solutions will be marked out by an unstable solution in which $v_h(t)$ just reaches the critical value but does not cross through it. This is reminiscent of the solution which marks the unstable manifold of a saddle point. Such basin boundaries will be of particular interest in this study, and if a way could be found of charting these unstable solutions it might offer a more efficient approach than the exhaustive simulation to be used here.

Another question raised by dynamical systems theory is whether the bowed-string problem has a natural state space, in which the divergent behaviour could be represented graphically in some way. A phase space based on derivatives is, of course, a nonsense for this system. But one can certainly define a state-space dimension of some kind as the minimum number of parameters necessary to specify a given state of the system completely. At first sight this number is infinite, since the string is thought of as a continuum. However, this is rather misleading. Certainly for the Friedlander and Raman models there is a well-defined dimension, equal to the number of past values of v and f which must be stored to iterate Eq. (8) or (11). (It is necessary to add one more parameter, to specify whether the current frictional state is sticking or slipping, to cope with hysteresis.) Rounded-corner models also have a finite dimension, at least when they are implemented in discrete form for simulation. Again, a definite number of past values (of outgoing velocity waves) is stored by the program to allow the convolutions to be carried out, and the dimension of the system cannot be greater than this.

However, even this finite dimension is usually uncomfortably large, and no simple graphical presentation of a sub-

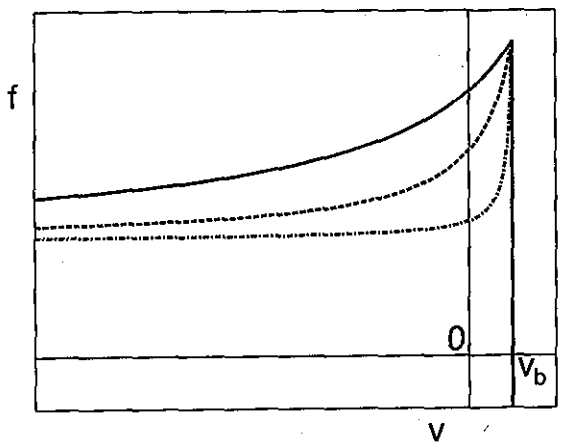

FIG. 9. Friction curves corresponding to $v_{mid}=-2$ (solid curve); $v_{mid}=0.2$ (dashed curve); and $v_{mid}=0.9$ (dash-dot curve).

set of it (for example via Poincaré maps) has yet been shown to be useful. The spaces which have been found to show very interesting structure are subspaces of the player's *control space*, and to this we now turn. As mentioned in the introduction, the many parameters of the bowed-string models can be divided into four categories. First is a group which describes the properties of the string: this would include the tension and mass per unit length, the torsional stiffness, admittance and wave speed, the bending stiffness, and the internal damping behaviour. The second group concerns the violin body. In the particular models discussed here, this includes only the parameters of Cremer's model (a spring constant and a dashpot rate) and the width and area of the torsional reflection functions. In a more complete model it would include the natural frequencies, damping factors and admittances (measured at the string terminations) of the various vibration modes of the instrument body. Since these are the mechanical quantities which contain the information about differences of material and construction between different violins, they will be of prime importance in later stages of this study.

The third group of parameters specifies the frictional constitutive law. In the present model the slipping portion of the friction curve is represented by a rectangular hyperbola, so that we need just three parameters: the coefficient of limiting sticking friction, μ_s, the asymptotic coefficient of friction at high slip speeds, μ_d, and one more parameter to specify the "tightness" of the curve joining these two values. The particular parameter we use for this, denoted v_{mid}, is the value of the slip speed at which the coefficient of friction is mid-way between μ_s and μ_d. Friction curves corresponding to three different values of v_{mid} are shown in Fig. 9.

The fourth group of parameters consists of the player's control variables. These specify the time-histories of the bow speed, the bow force, and the position of the bowed point along the length of the string. For the present purpose, only the bow force will be allowed to be time-varying. A family of exponential bow force transients will be studied, so that an individual member is completely specified by an initial bow force f_{init}, an eventual asymptotic bow force f_{asym}, and an exponential time scale τ_f such that the initial offset from f_{asym} decays as $\exp(-t/\tau_f)$. Results of systematic simulation will be shown in a plane whose axes are f_{asym} and $f_{\mathrm{init}}/f_{\mathrm{asym}}$. The aim is to find the basin of attraction of the Helmholtz motion in this subspace, and then to see how its qualitative features change when parameters of the system are varied. This is, of course, by no means the only possible subspace of control space which might be studied, but it will be seen to be quite illuminating for an initial investigation.

IV. RESULTS AND DISCUSSION

A. The influence of friction-curve shape

We first show a set of results for one particular run, to illustrate the detailed workings of the automatic classification algorithm discussed in section IIIB. The architecture of the Connection Machine makes it convenient to use 8192 processors in parallel. Each processor is given its own set of parameter values, and they all run independently. In a single run, a grid of 128×64 points in the chosen parameter plane can thus be studied, and all pictures will be shown at this resolution. A preliminary report on this work[2] showed a picture with higher resolution, achieved by splicing together the results of several runs (and using 16384 processors in parallel), but we have found that sufficient information can generally be gained at the lower resolution. The horizontal axis in each picture shows f_{asym}, on a logarithmic scale. The vertical axis shows $f_{\mathrm{init}}/f_{\mathrm{asym}}$, varying on a linear scale in the range 0-2. Thus points lying on a horizontal line through the centre of a picture represent "switch-on" transients in which the bow force remains constant throughout (as in the examples shown in Figs. 5 and 8). Below this line, the bow force starts low and builds up exponentially to f_{asym}, while above the line it starts high and falls exponentially to f_{asym}. The exponential time scale is constant throughout, at about three period-lengths.

Figure 10 shows four plots obtained after running the stiff/Cremer model for 40 period lengths, and carrying out the three tests on the last five periods. The results of the three separate tests are shown, and then that of the composite "Helmholtz" check. Figure 10(a) shows the correlation coefficient between successive period-lengths, averaged over the last five periods. More-or-less periodic motion is found in most of the region studied, the exception being at high bow force as one would expect. To find out the nature of the (nearly) periodic motions, we must examine the results of the other two tests. Figures 10(b) and 10(c) show respectively the number of significant slips in the last five period-lengths, and the number of Raman travelling corners counted in the same time. Only where these are both equal to five, and when the correlation coefficient is high enough, do we have Helmholtz motion. The points which pass this composite test are shown white in Fig. 10(d) while all others are shown black.

Thus Figure 10(d) shows a snapshot of which points in the chosen portion of parameter space lead to acceptable Helmholtz motion after 40 periods. However, there is nothing special about the choice of 40 periods—some points may have reached Helmholtz motion much earlier, while others which have not yet done so may do so after a longer time.

Note that some points which appear in Fig. 10(a) as "more-or-less periodic" after 40 periods actually correspond to regimes which are slowly evolving, and which may in time give way to different regimes. We can encapsulate information about a range of possible transient lengths by using colour coding in an enhanced version of Fig. 10(d). The same set of simulations can be run, and the set of tests for Helmholtz motion carried out every five periods. Then a composite picture can be plotted in which points which never reach Helmholtz motion are shown black as before, but points which do get there at some time are coloured according to the length of the transient before the first occasion when that point passes the tests. A set of such pictures is shown in Fig. 11, in which the simulations were run for 100 period-lengths. In particular Fig. 11(d) shows the same case of the stiff/Cremer model as shown in Fig. 10(d). Specifically, the white region in Fig. 10(d) encompasses all points in Fig. 11(d) coloured to indicate a transient of 40 periods or shorter (except for certain pixels which were showing acceptable Helmholtz motion after 40 period-lengths, but which later "regressed" to something else).

The set of runs whose results are depicted in Fig. 11 was designed to test the suggestion from the previous section, that sensitivity of behaviour may relate to the magnitude of frictional hysteresis. The two columns of the figure show the delta/Cremer model (on the left) and the stiff/Cremer model (on the right). The three rows show runs of these models which differ only in the shape of the friction curve: three different values of v_{mid} have been used, corresponding to the values used in Fig. 9. The main effect of this variation is to change the magnitude of hysteresis, from rather small (solid curve) to quite large (dash-dot curve).

It is immediately apparent that v_{mid} has a strong effect on the behaviour of the system. The two models give similar behaviour, but differ in one very important respect: the stiff/Cremer model appears to be "easier to play" than the delta/Cremer model, in the sense that it generally has larger contiguous areas of colour, less interspersed with black. The pattern of behaviour seen in Fig. 11, and in others not reproduced here showing intermediate values of v_{mid}, may be summarised as follows. When v_{mid} is close to v_b, in Figs. 11(a) and 11(b), both models have a rather restricted range within which Helmholtz motion is found. The allowed region takes the form of a vertical stripe, suggesting that the behaviour is mostly governed by the value of f_{asym}, and that the precise form of the bowing transient leading to this value does not influence very strongly whether the eventual outcome is acceptable Helmholtz motion. There are quite well-defined limits on f_{asym}, but note that these do not equate to the values of minimum and maximum bow force which have been calculated by Schelleng[26] and others.[9,28] Both limits are set, somehow, by the process of pattern selection in the initial transient: more than one periodic regime is possible over most of the range of f_{asym} considered here, and of these, the Helmholtz motion is being chosen only when f_{asym} lies between relatively narrow limits.

When v_{mid} moves farther away from v_b, in Figs. 11(c) and 11(d), the two models show rather different behaviour. In the stiff/Cremer model the range of allowed values becomes wider. There is a sprinkling of black spots within the coloured region, but most transients lead to Helmholtz motion sooner or later. For the delta/Cremer model, the points where Helmholtz motion may occur are confined within the same general region, but a great deal of this region is covered by a pattern of curving black patches. In those regions, periodic regimes are found which differ from the Helmholtz motion in having more travelling Raman corners. We have christened these particular regimes "multiple flyback motion," and examples have been shown previously.[2] They are all characterised by a bridge-force waveform which looks roughly like the Helmholtz sawtooth, except that in place of the single "flyback" in each cycle, there is a cluster of 3, 5, 7 or more closely-spaced flybacks.

It is very encouraging to encounter this regime in the results of simulations, since observations of real bowed strings have indicated that multiple-flyback motion is indeed a particularly common source of undesirable noise. When a competent player produces such a noise, unless the musical context has dictated a very low bow force, multiple-flyback motion is the most likely culprit. The results shown here indicate that, at least with the present very simplistic model of the violin body (Cremer's model), a stiffer string is easier to play than a perfectly flexible string in the sense of being less prone to multiple-flyback motion.

This makes some physical sense: the rounding of the travelling Raman corners necessitated by string stiffness may well make it more difficult to accommodate a group of closely-spaced corners on the string. However, it is a little surprising that the trend of better playability (in this particular sense) with stiffer strings continues beyond the level of string stiffness which has been suggested previously as the maximum acceptable value.[26] This fact has been verified with runs not reproduced here, using stiff-string reflection functions representing a range of values of string stiffness. Perhaps other effects make very stiff strings unacceptable in practice, either deterioration of tone quality or influences on playability from sources not included in this first investigation.

When v_{mid} moves even farther away from v_b, in Figs. 11(e) and 11(f), the difference between two models becomes rather less drastic. At low bow forces, both models now show a large area in which Helmholtz motion arises, although often with rather long transients. The stiff/Cremer model has a fuzzy vertical line at a maximum value of bow force, which is at a rather lower value than in Fig. 11(d). The delta/Cremer model again shows a significant area of black in the characteristic curving pattern which indicates multiple-flyback motion.

It is worth noting that the differences in friction curve between the three rows of Fig. 11 have some immediate physical significance to a violin player. Of course, the underlying constitutive law describing rosin friction cannot be changed by the player (except perhaps by choosing different rosin). However, for a given constitutive law, changes in the bow speed produce a somewhat similar effect. To see this, it is sufficient to consider the ideal Helmholtz motion. At a slow bow speed, the whole friction curve will be shifted to the left. The Helmholtz slipping speed will be low, since it is

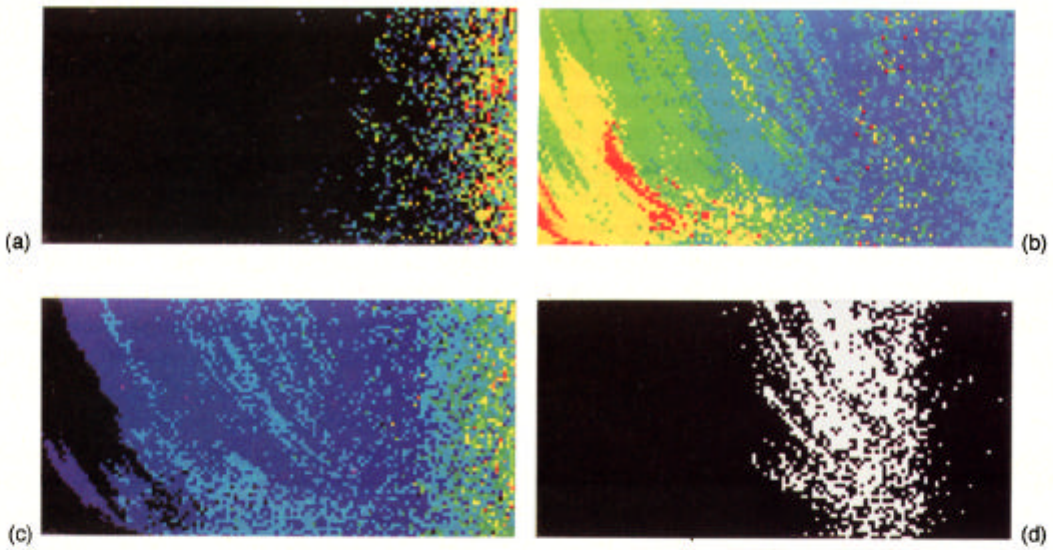

FIG. 10. Plots of (a) the correlation coefficient; (b) the number of significant slips in five periods; (c) the number of travelling corners in five periods; and (d) the cases passing the test for Helmholtz motion, for a set of simulation runs for 40 period-lengths as described in the text and the Appendix. The horizontal axis shows asymptotic bow force between 0.5 and 20, on a logarithmic scale. The vertical axis shows initial bow force as a fraction of asymptotic bow force, on a linear scale from zero (at the bottom) to 2 (at the top). The colour scales have the same sequence as is shown in the colour bar of Fig. 11.

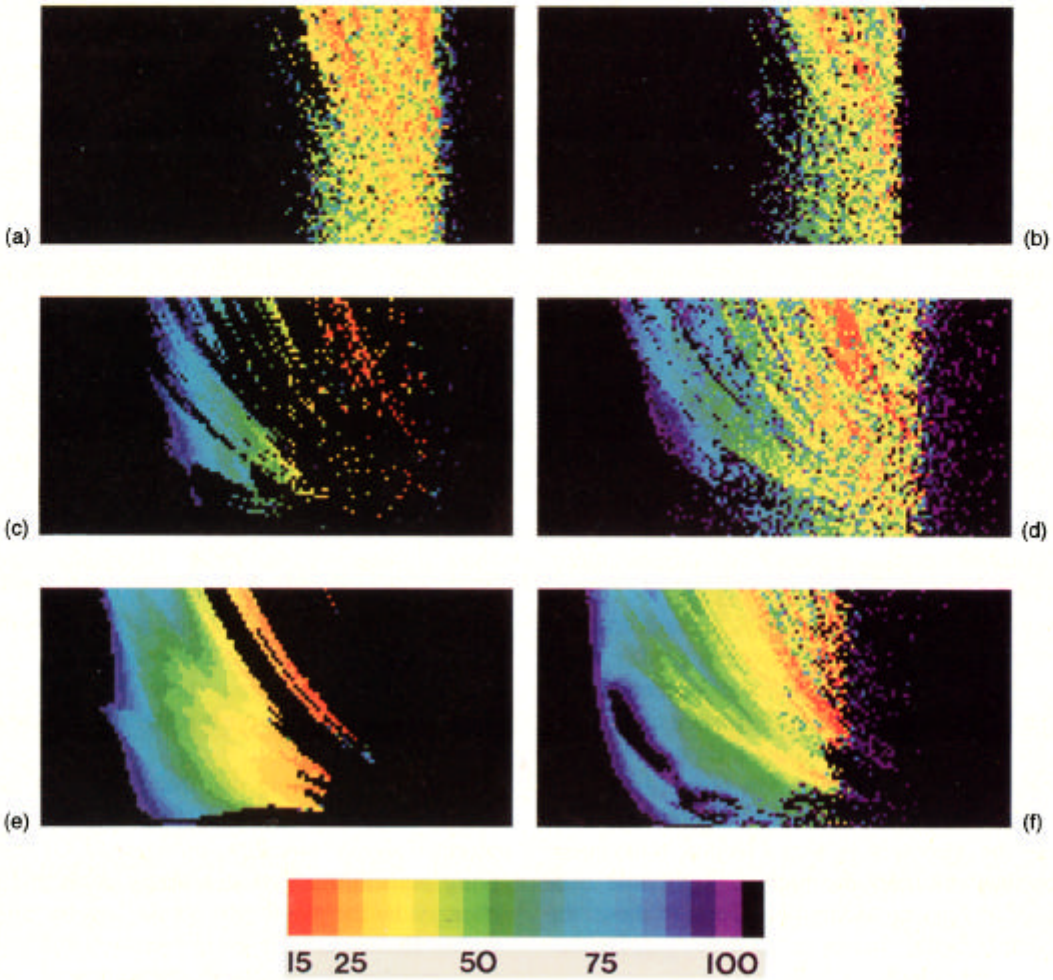

FIG. 11. Results of six different sets of simulations, plotted in the same format as Fig. 10. The colour scale denotes length of transient before an acceptable Helmholtz motion was achieved. The left-hand column shows the delta/Cremer model, the right-hand column the stiff/Cremer model. The three rows show $v_{mid}=0.9$ (top), $v_{mid}=0.2$ (middle), and $v_{mid}=-2$ (bottom).

approximately proportional to the bow speed. Thus all activity is confined to a section of the force/velocity characteristic which is close to the sticking portion—this is a rather similar effect to changing the friction curve from that of the dashed curve in Fig. 9 to that of the solid curve. At a higher bow speed, the opposite argument applies, and the effect is similar to that of changing the friction curve from the dashed to the dash-dot curve in Fig. 9. The variations in playability with changes to v_{mid} may thus manifest themselves as variations with bow speed.

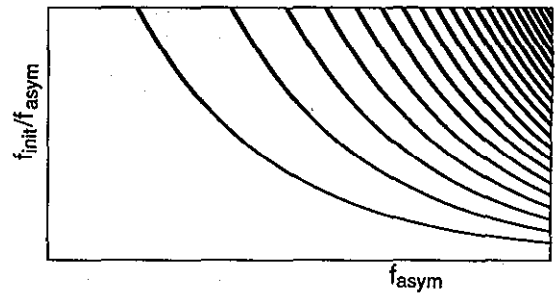

FIG. 12. Contour lines of bow force at time $t=0$, on the scales of Figs. 10 and 11.

B. Classification of terrain in control space

The discussion of the results of Fig. 11 has drawn some conclusions which are specific to the particular models and parameter values used in those simulations. But there is a more general issue which can also be illustrated by these results. The overall objective of this study is to take results like these for a wide range of models and parameter values, and compare them in terms of some high-level measures of "playability." In view of the large parameter space to be explored, a further degree of automation of this process is probably necessary. To move towards defining suitable high-level measures, it is useful to classify the qualitatively different types of "terrain" found in the results shown here.

We may distinguish five such types. First, there are uniform areas of black. Such areas are unambiguous—the model cannot be "played" in a satisfactory way in such a region. Second, there are coloured areas where the transient length is smoothly-varying. Such areas occur to the left of centre in Figs. 11(e) and 11(f). Provided the transient length is not too long, these are regions where the model is "docile" for the player: a small change of bowing gesture may make a small change in transient length, but nothing unexpected occurs. As a first guess, this may be the type of terrain which represents ideal playability.

Next we may identify two rather different types of "speckled" terrain. In Fig. 11(b), for example, the coloured stripe contains only a low density of black spots, but the colours are certainly not smoothly-varying. Under such conditions a Helmholtz motion is obtained fairly reliably, but small changes in bowing gesture may make a significant difference to the precise length and nature of the transient. The transients shown in Figs. 5 and 8 come from such a region of the control space, and illustrate the kind of variation of transient which occurs there. Whether such differences have significant audible consequences is not known, but it is a reasonable guess that they do. Such a region might be experienced by the player as moderately safe, but rather uneven.

The second type of speckly terrain is not well represented in Fig. 11, but an example of sorts occurs towards the right-hand side of Fig. 11(d). Here, coloured pixels are interleaved, apparently randomly, with a sprinkling of black ones. The sensitivity of transients to details of the bowing gesture is such that a small variation may make the difference between the selection of the Helmholtz motion and that of a different, undesirable, regime. Such a region must surely be experienced by the player as unreliable and dangerous. These are the conditions under which unwanted squeaks and whistles may be produced, because to avoid them requires extreme, maybe superhuman, bow control.

The final category of terrain can be described as "streaky." To the left of centre in the top half of Fig. 11(c), coloured pixels are interleaved with black regions. The player's perception of such a region is probably similar to that of the previous category, with random black speckles. But in this case there is obviously some structure in the pattern. This may make little difference to the player, but it suggests some significant difference in the underlying physics. Random speckles suggest sensitive dependence without offering much scope for further theoretical study. The pattern of curving streaks, by contrast, gives a tantalising suggestion of some order in the behaviour which might be amenable to explanation.

Not much progress has yet been made in seeking such an explanation. One thing can be said, though: the pattern of streaks is very reminiscent of the shape of the lines in Fig. 12, which shows contours of bow force at time $t=0$. (At later times, as the bow force tends towards f_{asym}, these contours tend towards vertical lines.) Since the magnitude of hysteresis is governed by the bow force, for a given friction curve, these lines are also contours of the difference v_2-v_1 at time $t=0$. This suggests that perhaps some critical event occurs early in a starting transient, triggered by the precise hysteresis level, and that this eventually determines the choice of periodic oscillation regime, for example between Helmholtz and multiple-flyback motion. We have examined many starting transients in such regions to look for some such critical event, but so far without success. Continuing effort in this area is surely justified.

A target for a future stage of the study of playability by systematic simulation might be to use image-processing techniques (such as fractal compression) to classify these five qualitative types of terrain in a given set of output. One could then construct measures of playability from the areas of the different types. That would give a quick basis for comparison of a wide range of cases. It would not be necessary for such measures to be perfectly in accordance with the judgements of players, provided they capture something of the right flavour of the results. When regions of parameter space were found in which these simple measures indicated strong variation, then one would revert to looking at the more detailed results to see what was going on.

Finally, it is important to address the question of whether the discussion in this section is based on a solid foundation. Terrain has been divided into categories and physical significance attached to these, without enquiring whether at least some of the behaviour might not be an artifact of the algorithm for automatic classification of oscillation regimes. To address this question seriously requires the examination of a wide range of waveforms, to see whether the automatic recognition algorithm agrees with human judgement. Unfortunately, it does not seem possible to convey very much by reproducing just a few examples here. Instead, we will simply describe the qualitative conclusion reached after many studies. When the terrain is smoothly varying, we have invariably found that the recognition algorithm works quite robustly. When the terrain is "speckly" or "streaky," things are not quite so clear. If a line of individual runs through such terrain is studied (corresponding to a line of pixels in the plots shown above), there is a complex intermixture of Helmholtz and non-Helmholtz motion. We then generally find that most individual pixels are described by the automatic algorithm in accordance with a judgement by eye, but that occasional ones are not. However, the qualitative nature of the *terrain* is never in doubt: in terms of the five-fold classification described above, the output of the program has always been found to be in accordance with human judgement. Since the terrain is more interesting than the individual points for the study of variations in playability, this degree of success of the algorithm seems adequate.

V. CONCLUSIONS

It has been shown, by systematic simulation studies, that a carefully-chosen "benchmark" model of bowed-string motion displays many of the qualitative features of the motion of real strings. A variety of periodic and non-periodic oscillation regimes are found, in good general agreement with playing experience. Two different bowing transients with the same final value of such parameters as the normal bow force do not necessarily produce the same regime of oscillation of the string. In particular, one transient may produce the desired Helmholtz motion while a rather similar transient may produce a different, undesirable, regime. It has been shown that at least one parameter of the problem (concerning the form of the friction function) has a sensitive influence on the size and shape of the basin of attraction of the Helmholtz motion, as seen in a subspace of the player's control space which is arguably related to subjective judgements of "playability." The various types of "terrain" found in this control subspace have been discussed.

The large number of system parameters have been grouped into a classification scheme which will form the basis of later studies using the same computer methods described in detail here. The most interesting studies will concern the influence of parameters which govern the properties of the string and the instrument body. The former are the concern of string manufacturers, and the latter the principal concern of violin makers. If the widespread notion that one violin is "easier to play" than another is really rooted, in part at least, in the ideas explored here, then there must be some significant influence from the parameters determined by the construction and setup of the violin body: in particular, the frequencies, damping factors, and normalised amplitudes at the string terminations of the various vibration modes of the body. If some such link could be established, this would form a very significant advance for the study of violin acoustics.

ACKNOWLEDGMENTS

The simulations for this work were carried out on the CM2 Connection Machine operated by the Pittsburgh Supercomputer Center, Pittsburgh, PA. The authors gratefully acknowledge the grant from the PSC which made this work possible. RTS further acknowledges support from the National Science Foundation for work on bowed string instruments.

APPENDIX: SIMULATION DETAILS

All the simulations shown in this paper have some parameters in common. The bowed point divides the string in the ratio 15:124, and there are 139 time samples per nominal period of transverse string motion. Frequencies are all quoted based on an assumed sampling rate of 61.16 kHz, which makes this nominal fundamental frequency 440 Hz. The time decay rate of all bow-force transients is $\tau_f = 400$ time samples.

The transverse reflection functions have been described in the text: details of Cremer's model are given in section 4 of Ref. 1, the "stiff" reflection function is described and plotted in the Appendix of the same reference. The torsional reflection functions are identical Gaussian functions, of the form $\exp[(-t/\tau)^2]$ where $\tau = 1.4$ in units of time samples, normalised such that the sum of the discrete points is -0.95. The torsional wave speed is 2.4 times the transverse wave speed, and the ratio of characteristic admittances is 0.617. The time delays for the Gaussian functions on the two sections of the string are adjusted to match the assumed speed and bow position by ensuring the vanishing of the first moment (calculated in discrete form) at the desired time delay.

The slipping portion of the friction curve is modelled as a rectangular hyperbola, which is uniquely specified by three quantities: the limiting coefficient of sliding friction, on the point of sticking, is 0.8 (when the string surface speed is equal to the bow speed); the asymptotic coefficient of friction at high relative sliding speeds is 0.3; and the value of string speed when the coefficient of friction passes through the mean of these values, denoted v_{mid}, is given a variety of values which are specified in individual figure captions. Units are assumed such that the bow speed $v_b = 1$ and the characteristic admittance for transverse motion $Y_0 = 2$. In these units, the Schelleng maximum bow force[9,26] is approximately 12, which lies within the limits of Figs. 10 and 11.

[1] J. Woodhouse, "On the playability of violins, Part 1 Reflection functions," Acustica **78**, 125-136 (1993).
[2] J. Woodhouse, "On the playability of violins, Part 2 Minimum bow force and transients," Acustica **78**, 137-153 (1993).
[3] H. Helmholtz, *On the Sensations of Tone* (Dover, New York, 1954) (English translation of the German edition of 1877).
[4] L. Cremer, *The Physics of the Violin* (MIT Press, Cambridge, MA, 1985), see Chap. 3.

[5]M. E. McIntyre, R. T. Schumacher, and J. Woodhouse, "On the oscillations of musical instruments," J. Acoust. Soc. Am. **74** 1325-1345 (1983).
[6]The Connection Machine is made by the Thinking Machines Corporation. These computations were done on the CM2 owned by the Pittsburgh Supercomputer Center.
[7]M. E. McIntyre and J. Woodhouse, "On the fundamentals of bowed-string dynamics," Acustica **43**, 93-108 (1979).
[8]M. E. McIntyre, R. T. Schumacher, and J. Woodhouse, "Aperiodicity in bowed-string motion," Acustica **49**, 13-32 (1981); see also Acustica **50**, 294-295 (1982).
[9]J. Woodhouse, "Idealised models of a bowed string," Acustica **79**, 233-250 (1993).
[10]J. H. Smith, "Stick-slip vibration and its constitutive laws," Doctoral dissertation, University of Cambridge, 1990.
[11]J. Woodhouse, "Physical modelling of bowed strings," Comput. Music J. **16** (4), 43-56 (1992).
[12]F. G. Friedlander, "On the oscillations of a bowed string," Proc. Cambridge Philos. Soc. **49**, 516-530 (1953).
[13]J. B. Keller, "Bowing of violin strings," Comm. Pure Appl. Math. **6**, 483-495 (1953).
[14]J. Woodhouse, "On the stability of bowed string motion," Acustica **80**, 58-72 (1994).
[15]A fast, light bow on a cello C string will sometimes elicit nothing but a slight rushing sound, presumably from the rosin irregularities on the bow.
[16]C. V. Raman, "On the mechanical theory of vibrations of bowed strings," Ind. Assoc. Cult. Sci. Bull. **15**, 1-158 (1918).
[17]F. S. Gillan and S. J. Elliott, "Measurement of the torsional modes of vibration of strings on instruments of the violin family," J. Sound Vib. **130**, 347-351 (1989).
[18]R. T. Schumacher, "Self-sustained oscillations of the bowed string," Acustica **43**, 109-120 (1979).
[19]See Ref. 4, section 5.4.
[20]Values measured by one of the present authors (RTS).
[21]Lord Rayleigh, *The Theory of Sound* (Dover, New York 1945), article 189.
[22]O. Krigar-Menzel and A. Raps, "Ueber Saitenschwingungen," Ann. Phys. Chem. **44**, 623-641 (1891).
[23]A. H. Benade, *Fundamentals of Musical Acoustics* (Oxford University Press, London, 1976; and Dover, New York, 1990).
[24]X. Boutillon, "Analytical investigation of the flattening effect," J. Acoust. Soc. Am. **90**, 754-763 (1991).
[25]See Ref. 4, section 5.2.
[26]J. C. Schelleng, "The bowed string and the player," J. Acoust. Soc. Am. **53**, 26-41 (1973).
[27]J. C. Schelleng, "The violin as a circuit," J. Acoust. Soc. Am. **35**, 326-338 (1963).
[28]See Ref. 4, section 4.6.