Playability study of a bowed string physical model including finite-width thermal friction and hair dynamics

MSc Modelling for Science and Engineering

Quim Llimona Torras

Master Thesis / 2015

THESIS ADVISORS
Esteban Maestre, Susana Serna
Departament de Matemàtiques

Universitat Autònoma
de Barcelona

ABSTRACT

In this work, we present a playability study of a state-of-the-art bowed string computational model that includes finite-width thermal friction and hair dynamics, primarily targeted at sound synthesis. The concept of playability itself is controversial; it can be loosely defined as how large (and physically feasible) is the region of the possible input controls (namely bow force, bow-bridge distance, and bow velocity) to the instrument that result in good-sounding vibration regimes. There are numerous studies on the qualitative analysis of vibration regime distribution over different variants of the model, but here we aim at taking a more quantitative approach so that it is possible to compute a *functional playability* over model parameter ranges. We introduce an algorithm for automatic detection and classification of vibration regimes based on studying the periodicity of the simulated string velocity under the bow. Based on such automatic classification, we propose a scalar playability descriptor computed as the portion of the input control space that leads to each of the vibration regimes into consideration. Some preliminary results are presented by sweeping a range of thermal-related parameters, which display a clear region of easily-achievable Helmholtz regime. Other promising results related to patterns on the distribution of anomalous low frequency vibration regimes over input control space are also presented, opening ground for further investigations.

Contents

List of Figures

Chapter 1

INTRODUCTION

1.1 On physically-informed sound synthesis

Music has always depended on technology - even cavemen used their tools to build percussion instruments that could accompany their chants. It is no wonder, then, that as soon as computers appeared people started thinking about how to make music with them. At the beginning of the computer era, computational constraints enforced the used of fairly basic synthesis methods, such as wavetable lookup. Frequency Modulation synthesis was a great breakthrough because of the great range of timbres it gave for such an algorithmically simple method, thanks to its non-linearities.

Nowadays, the great compuational capabilities that even handheld devices such as smartphones offer allow real-time physically-based synthesis, that is, synthesis based on numerically solving mathematical models of how real instruments behave. As computing power grows and research on such methods advances, more and more refined methods become available, aiming always at what many argue impossible - to trick a musician into thinking a fully synthetic sound is actually the result of recording an acoustic instrument.

In the meanwhile, the community has developed numerous artistic applications of such methods, such as the synthesis of sounds physically impossible to produce at reasonable costs (i.e. without breaking the instrument), or the synthesis of an entire symphonic orchestra (where individual detail is no that notiecable, especially when all musicians play in a *tutti*) at a minimum expense.

Physically-informed sound synthesis has been very closely related to research on organology and musical acoustics, that is, on the physical foundations behind the sound producing mechanisms of musical instruments. Not only the synthesizers are based on equations that acousticians derived, but the numerical models are enabling more and more the validation of theoretical or experimental findings and actually contributing to science themselves. State-of-the-art physical models of instruments have such a detail that allow generating sounds with a physically unfeasible degree of control. This can lead to novel discoveries of the inner workings of musical instruments at a microscale, and these findings can be lately corroborated through extensive (and expensive) experiments involving purposely-built mechanical instrument players such as described in Schoonderwaldt et al. [2008], and viceversa.

These findings greatly contribute to our understanding of how musical instruments work, and can be later applied to a range of fields, such as instrument making and manufacturing (is it possible to *optimize* current instruments? To what extent can we predict how good an instrument is before playing it?) or pedagogy (how can a student choose a violin best suited to her needs?).

While explicitely not trying to modify or improve our physically-based bowed string synthesiser of choice, in this project we aim at exploring how it responds to different inputs and at discovering the possibilities that such a state-of-the-art artifact can offer. In plain words, we want to know what does it do when it is *played*.

1.2 Statement of purpose

The main goal of this project is to explore novel ways of determining the playability of a physical model of a bowed string instrument such as the violin. Although models (and implementations of those) that are computationally efficient enough to apply our analysis methods while being physically accurate are scarce, our objective is not to score the model itself but to compare different configurations of it.

In particular, our objectives are:

- To devise a method for efficiently and accurately determining under which vibration regime the model is.
- To build vibration regime maps that show how a specific configuration of the model reacts to different gestural inputs.
- To define exploratory playability descriptors based on such maps, that allow easier comparation of different configuration layouts.
- To present preliminary playability maps on a defined range over specific configuration parameters.

This report is structured as follows: after a brief context on sound, bowed strings, and the concept of playability, the method for estimating the vibration regime of a particular state of our model is presented, followed by some proposed computational indicators of playability. The results we obtained with these methods are presented afterwards, including well-known diagrams of unprecedented resolution that arise new, interesting questions. Finally, the main contributions of this thesis (both on playability and on generic understanding of the bowed string) are summarized.

Chapter 2

BACKGROUND ON BOWED STRING INSTRUMENTS

Given the nature of the field where this thesis is being presented, we deemed it appropiate to give a short introduction on what is sound, how do musical instruments produce it, and how can we represent it through mathematical formulations. It is very difficult to summarize an entire field in such as short space, but references are provided along the way for those interested.

This chapter is structured as follows: first, the basic principles of sound are presented, together with a rationale on how to treat it and a few notes on how musical instruments work from a conceptual point of view. Then, we pay more attention to string instruments, and in particular to free string vibration first and secondly to bowed strings, together with an outlook of how to model them efficiently as physical systems. Finally, we define the concept of playability, introducing some well-known diagrams and concepts that help explain how bowed string instruments such as the violin work.

2.1 Mathematical and computer representations of sound

Sound is a very complex phenomenon at the perceptual and cognitive levels, but based on a relatively simple physical event: the perturbation of air mollecules causing local deviations in air pressure, which moves a membrane located inside our ear and transmits this movement to the sensory system. Within the range of audible while non-damaging sound, perturbations on the fabric of air follow the wave equation; therefore, the perturbations expand in all directions, travelling and becoming locally fainter. Some may argue there is also an energy loss due to friction, but we can safely ignore this for now.

The way we perceive these perturbations is through oscillations; this is natural, since many of the air perturbations we are interested in perceiving (i.e. those resulting from something moving) will be of an oscillatory nature due to the elasticity or stiffness of objects. Therefore, when talking about a sonic event one is mainly interested in the spectral components of it: which oscillation frequencies are present, and at which intensity.

Ignoring now that we have two ears and therefore binaural hearing, a person hearing a sound can be modelled through a one-dimensional, time-dependent variable that indicates fluctuations around the mean air pressure at a particular point in space inside their ear. It is better, however, to think of this signal as the fluctuations a source (such as a speaker) adds to the environtment; then, these effects can be considered linear and therefore just additive. In the case of the listener wearing headphones, the two models can be considered equivalent; this is what we will be assuming, because although the propagation of sound through space by itself does not add much to the question (just linear attenuation), very complex phenomena arise when the sonic waves bounce off walls or in the ear pavement.

Moreover, the *room* where the acoustic event takes place further modifies the sound; its walls act as boundaries for the wave equation, and the wavefronts are reflected back. This creates both resonance modes, that emphatize certain frequential components of the sound, and *reverberation*, which plays a key role in music through softening note transitions. The reader is refeered to Kuttruff [2009] for a complete guide on room acoustics, where these effects are discussed in depth.

For further reading on musical acoustics, please refer to Benade [1990].

2.2 Musical instruments: the bowed string

Most musical instruments can be accurately described through the following framework:

- *Excitation source*: a device that makes use of a certain physical phenomena to generate a complex oscillation pattern at a certain rate using an input energy.
- *Radiator*: a device that is mechanically attached to the excitation source, and transmits the vibration to the air through a coupling device such as a plate.

While radiators can be represented safely and linearly through a convolution with a kernel, excitation sources require more advanced techniques, and their modelling using PDE's that reflect how the underlying phyisical system works is an area of active research. In the case of string instruments such as the violin, the string acts as an excitation source; its tension propagates perturbations as waves, and its quite rigid ends reflect the perturbations giving rise to periodic oscillations. The string transmits these perturbations to the bridge, which is connected to the body, our radiator. The body acts as an impedance adapter from the very narrow vibration of the string powerful enough to move such as big mass to the big plates of the violin that can move a bigger volume of air. Notice that while in plucked strings such as the guitar the string is given an initial excitation and then let resonate, in bowed strings there is a continued flow of energy, giving rise to self-sustained oscillations.

In this project, we will use a non-linear model for the excitation source, the string itself, coupled through one of its ends to a bridge modelled through a linear admittance filter measured experimentally and implemented as a battery of two-pole resonators. Further details about the model will be given throughout the chapter, and can be found in Maestre et al. [2014].

The solution function u of our equations represents the transversal velocity of the string at time t and position x, with x being a scalar that indicates distance from the origin. Notice that strings can have transversal waves in two dimensions or *polarizations*, horizontal and vertical; they are usually simulated independently and interact at the boundaries, but we will for now assume that we only simulate the horizontal polarization, the one that follows the bowing (or plucking) direction. The solution will be defined within the domain $x \in [0, L]$, where L is the length of the string. The boundaries of such domain can for now be fixed to 0, that is, the

string is "glued" at the terminations and cannot move; this, as we will see, will cause waves to reflect back inverted. Notice that, as the model gets perfected, these boundaries will have additional conditions on the derivatives of u, representing the damping caused by the string stopper not being ideal (especially in the case of fingers) and introducing an energy loss, or the boundary not being fixed but slightly free and coupled with the movement of other strings on the instrument.

2.2.1 The wave equation as a string vibration model

A very basic model for the behaviour of a tense string is the 1D wave equation:

$$u_{xx} = c^2 u_{tt}$$

As shown by D'Alembert in the XVIII century, this equation has as a general solution that consists of two fronts propagating the initial excitation in opposite directions (travelling waves) and at speed v, that is:

$$u(x,t) = F(x - ct) + G(x + ct)$$

However, when a fixed (Dirichlet) boundary is added to both ends of the string, the waves reflect back with reversed sign. It can be shown that this originates *standing waves* on the string, or longitudinal oscillations that do not propagate in space. A very rough intuition about why this happens is that the two travelling waves "cancel out" after bouncing.

The propagation speed, together with the string length L, define the oscillation frequency, which has the perceptual equivalent of *pitch* (in plain terms, which note is being played); notice that the effective length is actually double, because for a full cycle to complete the wave needs to travel back and forth across the whole string:

$$f_0 = c/2L$$

When tuning a string, it is useful to select the desired frequency of oscillation given the string length, although we will refer to c in this report for simplicity:

$$u(x,t) = F\left(x - \frac{f_0}{2L}t\right) + G\left(x + \frac{f_0}{2L}t\right)$$

This general solution to the equation allows solving it without requiring expensive schemes such as the finite difference method; one simply needs to calculate the displacement during the time step (that is, $c\Delta t$), and make whatever memory arrangements are required. In the 1D case, this can be achieved with a computationally cheap delay line, which even allows interpolation in case the displacement is not an even number of samples. This method is known as Digital Waveguides.

Such an equation has a very fundamental issue: once excited, the oscillation on the string is stationary and does not decay. This is completely unnatural, since it would correspond to a note that never ends! However, it serves its purpose on showing how a string can *oscillate* given a once-only initial energy input.

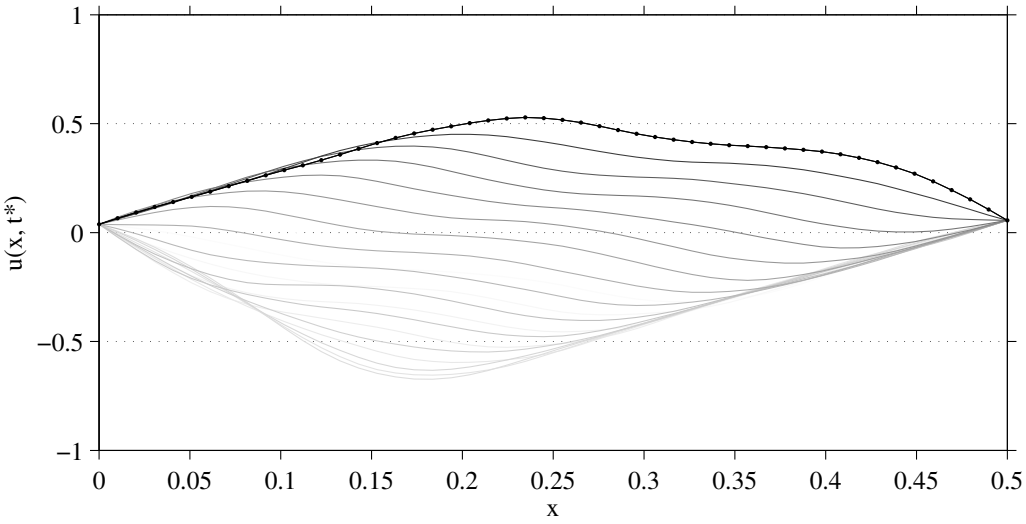

Figure 2.1: Sketch of the temporal evolution of the string profile.

The reason why real strings have a deacaying oscillation are losses mainly due to damping at the string ends, friction, and internal dispersion. They all contribute to uniform damping and to frequency-related damping, more related to higher order derivatives. However, they can all be modelled through linear and time-invariant digital filters (i.e. through a convolution with a specifically designed response). Thanks to the commutativity of these filters, it is possible to group (*lump*) them all in a single point, usually at the end of the string or at the reading point, and apply all the losses through a single convolution.

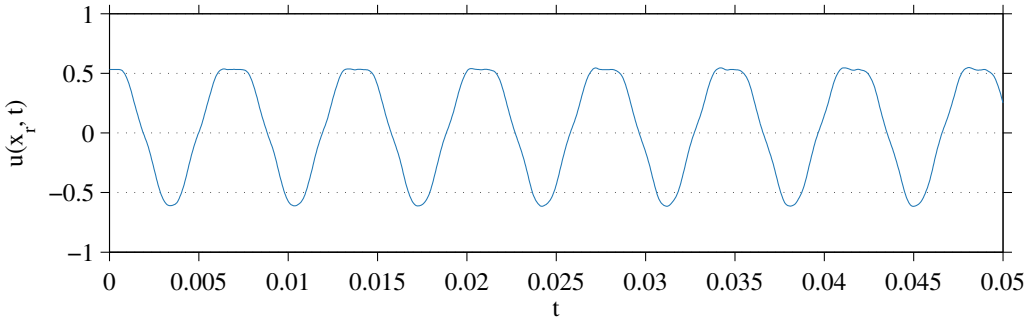

Figure 2.2: Temporal evolution of the string point selected for rendering for the basic equation.

For a more extensive but plain introduction to Digital Waveguide modelling, see Smith [1992].

The methods presented until now model the movement of the string itelf, but are not enough to produce musically pleasant sound. Real string instruments get all the color and tone from the resonance of their bodies, often made of wood. Moreover, some reverberation due to reflections on the walls of the hall or room usually needs to be added. Adding reverberation is usually accomplished by convolving the output signal with a measured impulse response from a concert or through specifically designed digital filters consisting mainly of delays and all-pass filters, but adding body resonance can be more complicated.

Although body resonance is often implemented using digital filters as well - as it physically consists on the superposition of selectively attenuated and delayed copies of the signal that bounce back and forth inside the body -, the tricky part is to know the desired response and therefore the coefficients to use. One of the most successful approaches - and the one we will be using - is based on measuring the admittance of the body at the bridge, which is where energy from the string vibration is transmitted to the body. This admittance, which is the impulse response of the bridge velocity as a response to a Dirichlet's delta excitation, is then fitted to a battery of parallel two-pole scalar resonating filters, each of them representing one of the modes of vibration of the body - hence its name, a *modal* approach Maestre et al. [2013]. However, this technique does not model how the instrument *radiates* sound to the outside world; it provides an output similar to that of a piezoelectric sensor mounted in the bridge, as it is usually done when amplifying bowed string instruments - especially in contemporary or jazz music traditions. In order to get an output closer to what a microphone would capture, further filtering is needed, although designing a filter by hand is usually enough to get pleasing

results. For a review of other body modeling techniques, see Karjalainen and Smith [1996].

Another source of damping mentioned before, the string ends, is modelled directly using hand-crafted low-pass filters. Although getting a realistic synthetic note is easy with this technique, it is very difficult to model the dynamics of this filter as the amount of damping varies greatly when what stops the string is a finger rather than the fingerboard end. Depending on the pressure and tilt of the finger, it can range from almost the same as the fingerboard end (very resonant) to a highly damped sound - and musicians use this fact a lot during performance (Kinoshita and Obata [2009]).

Shortly after Digital Waveguide methods appeared, a technique known as *commuted synthesis* became popular, which consisted in taking advantage of the fact that both the resonance of the instrument body and the room can be modelled as linear filters, and therefore could be actually convolved with the excitation signal *before* actually running the synthesis itself. Even more advanced methods that require non-linear interactions, such as bowed strings, can partially benefit from this approach by computing most of the linear elements together (Smith [1997]).

2.2.2 Physical models of the bow-string interaction

While plucked string instruments, such as the guitar or the violin when playing *pizzicato*, can be implemented with what has been described until now - they only require an initial excitation -, bowed strings are basically an extension of them with a much more complex energy input mechanism, sometimes described as *periodic plucking*. The string is basically dragged and released at each vibration cycle by the bow, therefore getting a continuous input of energy and being able to sustain a note without decay undefinetely; that is why it is called a self-sustained oscillator. This kind of motion was first observed by Helmholtz [1862], and honoring his name is known as Helmholtz motion.

The first complete model of the bow-string interaction was provided by Raman [1918], and his analysis (performed by hand) already hinted some of the possible behaviours of the system. In his model, he assumed that the friction force is a known non-linear function of the string velocity. Some years later, Friedlander [1953] and Keller [1953] solved the model numerically by imposing a linear condition on the force-velocity relationship that depends on the past values, found through drawing characterstic lines and bouncing them at the boundaries avoiding

crossing the non-linear contact point. By finding the intersection between this line and the known non-linear function, both the friction force and the string velocity can be found at each time step.

McIntyre and Woodhouse [1979] later proposed an hysteresis rule to solve ambiguous cases of the intersection, together with a dispersive method for introducing corner rounding. In Woodhouse [2003], the friction coefficients were modulated according to heat, which was propagated with a finite difference scheme, instead of relying on the contact velocity alone for determining the friction force. This avoided having to solve the hysteresis explicitely on the $v - f$ curve. The rosin progressively warms up during the slipping phase, and it cools progressively during sticking; therefore, the hysteresis appears naturally. See Smith and Woodhouse [2000] for a discussion on how rosin affects the friction of the system.

Independently, other researchers had already developed models with finite-width contact point where the bow-string contact zone was represented as a line of nodes, with each node representing an equivalent hair, although most models exploring new friction formulations had ignored this fact (Piteroff and Woodhouse [1998]). This finite-width formulation demonstrated for the first time partial slips, where a portion of the bow is already slipping while another is still stuck.

For a complete guide on physically-informed musical instrument synthesis, the reader is referred to Smith [2010]. For further reading on the acoustics of bowed string instruments, please see Woodhouse and Galluzzo [2004] and Cremer [1984].

A relevant contribution posterior to such reviews is the simulation of the vertical string polarization (perpendicular to the bowing direction), as described in Mansour et al. [2013].

2.2.3 Model used in the experiments

The model we will be using in this project, described in Maestre et al. [2014], considers two string polarizations: horizontal, parallel to the bowing plane, and vertical, perpendicular to the bowing plane. The friction coefficient is modulated with heat along a finite-width contact zone that includes hair dynamics, and uses digital waveguides for corner propagation. At the bridge, a two-dimensional modal formulation of the admittance is applied to account for body effects as a frequency-dependent reflectance matrix. An overview of the model is provided in Figure 2.3.

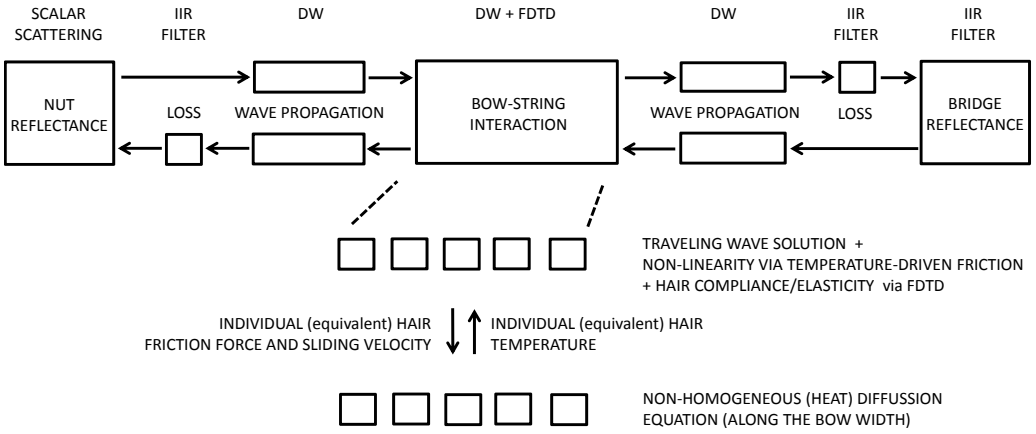

Figure 2.3: Overview of the bowed string model used in the simulations, taken from Maestre et al. [2014].

Horizontal and vertical transverse waves get coupled at the nut, at the bridge, and at the the finite-width bow-string contact. At the nut, both polarizations interact through a 2D reflectance matrix. At the bridge, both polarizations of all four strings get coupled through a lumped 2D admittance matrix that represents the instrument body, and contains complex, frequency-dependent elements.

This body representation is constructed with a modal approach, and is described in Maestre et al. [2013]. It consists of a battery of two-pole resonators in parallel parametrized by their central frequency, bandwith, and (complex) amplitude. Mode frequencies and bandwidths are first estimated in the frequency domain by solving a non-convex, constrained optimization problem. Then, mode amplitudes are estimated through semidefinite programming while enforcing passivity. This procedure is repeated for each combination of excitation and response polarizations, therefore providing the 4 elements needed for the admittance matrix.

Bow-string interaction is represented by a line of nodes, each modeling an equivalent hair-string contact. Here, bow and string interact in both polarizations: non-linear interaction happens on the horizontal plane, while linear interaction happens on the vertical plane. Polarization coupling happens through the normal force exerted by the string on the hair. In between nodes, we use the traveling wave solution provided by digital waveguides. Interaction in the horizontal plane takes place through a nonlinear friction characteristic that is solved at each node: the coefficient of friction is dynamically modulated by temperature changes. Then, hair compliance and elasticity are included in each node with finite differences,

leading to a non-linearity that can be solved graphically a la Friedlander. In the vertical plane, each node is represented as a mass-loaded spring-dashpot system.

Temperature at each of the nodes (including lateral nodes without bow contact to allow for diffusion) is increased via the conduction of heat due to the sliding friction, with the source term of the diffusion expression fed by sliding velocity and normal force. Temperature loss is caused either by convection during sliding, or by diffusion.

The model does not take torsional waves, hair-hair interaction, or finger-string interaction into account.

2.3 The concept of playability in bowed string instruments

Before understanding what playability means and how to measure it, it is important to give more context about bowed string models from a more functional perspective: beyond the equations, what kind of parameters do they have? How are they controlled during the simulation of a performance? Although not formally defined in the literature, we find the division in acoustic, parameter, and control spaces to be useful for writing about bowed strings and communicating how they work.

2.3.1 The acoustic, control, and parameter spaces

The complex nature of the bow and string mechanism are only enablers of what makes the bowed string a great instrument family. Non-linearities present in the system, especially in the string-hair interaction, produce very rich and diverse sounds. These sounds are not by themselves very complex, as they are just periodic oscillations, but they can vary greatly as a response to relatively small changes in the system state. We will call this range of different sounds, or regimes, that can be produced the *acoustic space* of the instrument. Since the goal of playing is to produce sound, we will also refer to it as the *target space*. This space can be described by diverse descriptors that represent attributes related to the spectrum of this oscillatory signal, although it is more convenient to refer in which of the possible regimes of oscillation the string is; this regime can be determined experimentally through those low-level descriptors; Raman [1918]

already provided a list of possible regimes in 1918, deduced entirely by hand with analytical formulations of a very simplifed model of string vibration.

On the other hand, this interaction is, as we mentioned, very sensitive to changes in the system. These changes can come from two sources: the instrument itself, such as in changes in the constants that drive the friction or elasticity components, or the player, through the way she controls (or plays) the instrument, such as the relative instantaneous velocity between bow and string, or the exerted force on it. We will call these sources of change the *parameter space*, for the ones due to the instrument, and the *control space*, for the ones due to the player. Notice that the parameter space is expected to be adjusted beforehand and mantained constant during the performance, while the control space is expected to vary and is where the art of playing bowed strings is.

2.3.2 Spatial mappings

In 1973, Schelleng defined a parametrized mapping between the control and acoustic spaces Schelleng [1973]: given a projection of the control space that only takes into account bow force and relative distance between the contact point of bow and string and the bridge, there are three regions separated by two lines that correspond to different sound characteristics: the expected periodic motion in the middle, raucous sound on one side, and surface sound on the other, as seen in Figure 2.4. These areas have been further divided in the literature: within the raucous region two different regimes can take place: aperiodic motion, and anomalous low frequency oscillations. Within the lower surface sounding region, two different regimes also exist: higher-frequency modes due to multiple stick-slip cycles per period, and regions where there is not enough energy to actually produce any oscillation (i.e. the bow is always slipping), usually labelled as having decaying energy.

The Schelleng diagram alone is not enough to fully represent how the control and acoustic spaces map. For instance, it is known that, when the bowing speed is changed, the slope of the line that separes surface sound from the Helmholtz region also moves. Moreover, it represents a *static* portrait of the mapping; what makes bowed strings really difficult to master are their dynamics. The Guettler diagram is another proposal that focuses on, assuming there is Helmholtz motion, how long does it take to get there starting with no velocity; it sweeps the space based on the acceleration of the bow and the bow pressure.

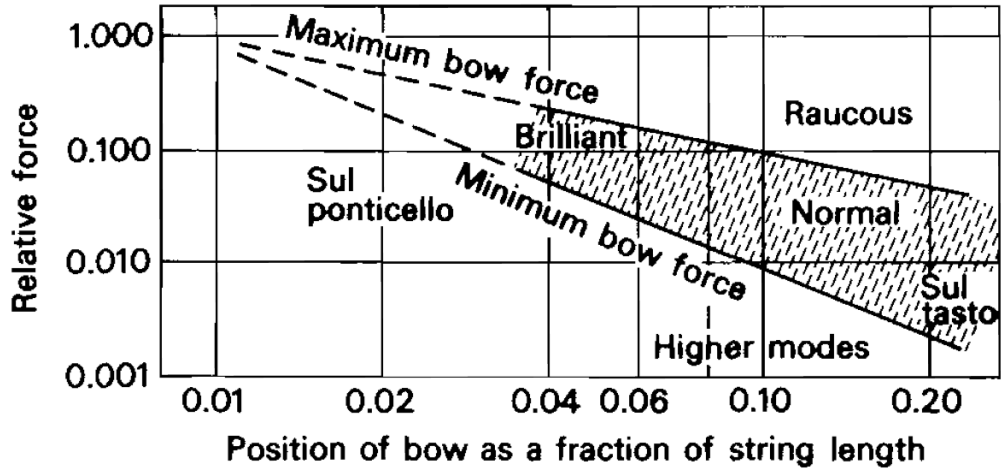

Figure 2.4: A replica of the original Schelleng diagram, published in Schelleng [1973].

These acoustic and control space mappings are valid for a given point of the parameter space, or a given *configuration* of the model. Different configurations will be indeed fully characterized by their acoustic-control mapping, expressed through tools such as the mentioned diagrams - even more than by the coefficients themselves that one has to input to the system.

2.3.3 What is playability?

According to Woodhouse and Galluzzo [2004], what makes a violin (or any bowed string instrument, for that matter) good is a combination of its *tone* and its *playability*. The instrument's tone refers to its sonic signature, to how listeners may rate its sound, and is usually associated with the spectral shape of the sound, closely related to the resonance modes of the body and therefore to its impulse response. Playability, by contrast, refers to how *easy* it is to obtain good tone from the instrument, where good tone means Helmholtz motion, regardless of its quality. These two dimensions are not necessarily correlated; it is well known that Stradiavari violins, considered amongst the best in the world, have actually bad playability; unless played with great care and good tuning, they will not deliver what one would expect from them. This leads to a very important - and well-known by violin professors - corollary: the dream violin of a consolidated concertist might be a very bad choice for a beginner, since she would find it too difficult to get good tone

from it (especially considering how expensive such a violin would probably be).

It follows from intuition that an instrument with a very wide region of the control space where there is Helmholtz motion (that is, with a wide Helholtz band on a hypothetical multi-dimensional Schelleng diagram) will be easier to play, since a broader range of gestures will end up providing "good tone". Therefore, one may argue that the Schelleng diagram can be used as a (partial) playability indicator, assuming there is not much variability when changing other control variables such as terminal velocity. Similarly, instruments that show very short transient times in the Guettler diagram, that is, that very quickly achieve Helmholtz motion, are also very playable.

Some may argue that playability is more than easily reaching Helmholtz motion; a fundamental part of bowed string performance consists precisely on playing with the duration and character of transients, art known as *articulation*, and having an instrument that systematically switches to Helmholtz motion as soon as the bow starts to move would not be the best choice. It would then be necessary to:

1) Define acoustic and control space mapping diagrams that reflect the different articulations that a musician uses.
2) Study what do musicians prefer, i.e., what is their *sweet spot* amongst the possible shapes of such diagrams.
3) Find a reliable distance metric that can estimate how close a given mapping is to what a musician expects.

This is beyond the scope of our work, but it raises an interesting point when interpreting the playability results that will arise after the experiments.

Notice that we are talking about playability from an acoustician's point of view. Another issue is the playability of a virtual instrument from a usability point of view: are the controllers that the player uses to control the model easy to use and intuitive? Are they expressive (i.e. do they allow exploring a wide range of sonic possiblities)? For a control-oriented discussion on playability and a more detailed distinction between what can people mean by playability, please refer to Young and Serafin [2003].

Chapter 3

REGIME ESTIMATION AND PLAYABILITY SCORING

In this chapter, we present the tools we chose for estimating the vibration regime of the string given a numerical simulation, together with a brief overview of the literature on this specific topic, and our approach to quantifying playability to be able to better compare different versions of the model.

3.1 String regime estimation

The core of this project is the computation of a Schelleng-like diagram by numerically playing our model using different combinations of bow-bridge distance and bow force controls given a model state or snapshot, i.e., a complete parameter-space configuration, and constant controls such as pitch or bow velocity. The diagram sweeps, as proposed by Schelleng, the bow-bridge distance relative to the effective string length and the bow pressure, both in logarithmic scale. For each evaluated point on the lattice, we compute the regime of the string.

There are different methods mentioned in the literature for computing the regime of a simulated bowed string. The first article that mentions the numerical estimation of vibrational regimes in simulated bowed strings is Woodhouse [1993], in which the author suggests driving a nonlinear model with a variety of inputs that sweep a space in order to get a map, a la Mandelbrot. He suggests simulating a note with a

duration of 100 expected period lengths. Then, the string velocity corresponding to the last five periods is used for good tone estimation. He computes the correlation between what are expected consecutive periods and, therefore, a high value will indicate that there is indeed periodic motion. This oscillation will have a rate either equal to the expected or to one of its multiples. If that is the case, the number of slipping regions within each period is estimated (a slipping region corresponds to a negative velocity, opposed to the bow movement), and if it is greater than one the segment is discarded, attributing to a higher regime of multiple stick-slips. The author already acknowledges that there are difficulties when partial multiple slips appear, and that this algorithm is very rough and calls for future revision. Nevertheless, the authors of Serafin et al. [1999] used the very same method when comparing different formulations of the model and different trade-offs (the most notable one being whether the addition of torsional waves plays a big role in playability, which apparently does not).

We decided to build a custom detector, based on analyzing the differential string velocity at the contact point after a fixed transient time and looking at the estimated fundamental frequency (or oscillation rate), aperiodicity, and energy. The fundamental frequency and aperiodicity are estimated using the well-known YIN algorithm, described in de Cheveigné and Kawahara [2002].

A sketch of the raw signal used for detection is outlined in Figure 3.1; notice how under the decaying regime the magnitude of the signal is very low, and under Helmtholz and multiple stick-slip regimes there is oscillatory motion at different rates.

The detector determines the mean square energy of the selected portion of signal, as well as the mean pitch and aperiodicity as computed using the YIN algorithm. If the energy lies below a given threshold, the note is discarded as having low energy, which corresponds to the *decaying* regime. If the signal is not decaying and the aperiodicity is above a given threshold, the signal is classified as *aperiodic*, which corresponds to part of the raucuous regime in the classic Schelleng classification. The other raucous regime occurs when the signal is not aperiodic and has a pitch significantly lower than the expected one (half an octave seems to be a good threhsold), which corresponds to *anomalous low frequencies (ALF)*. On the contrary, if the pitch is significantly higher the note is classified as having *multiple stick-slip*, which corresponds to the surface sounding region in the classic Schelleng classfication. Finally, if the pitch is close enough to its expected value the signal is classified as displaying full *Helmholtz* motion, which corresponds to the playable region in the classical diagram.

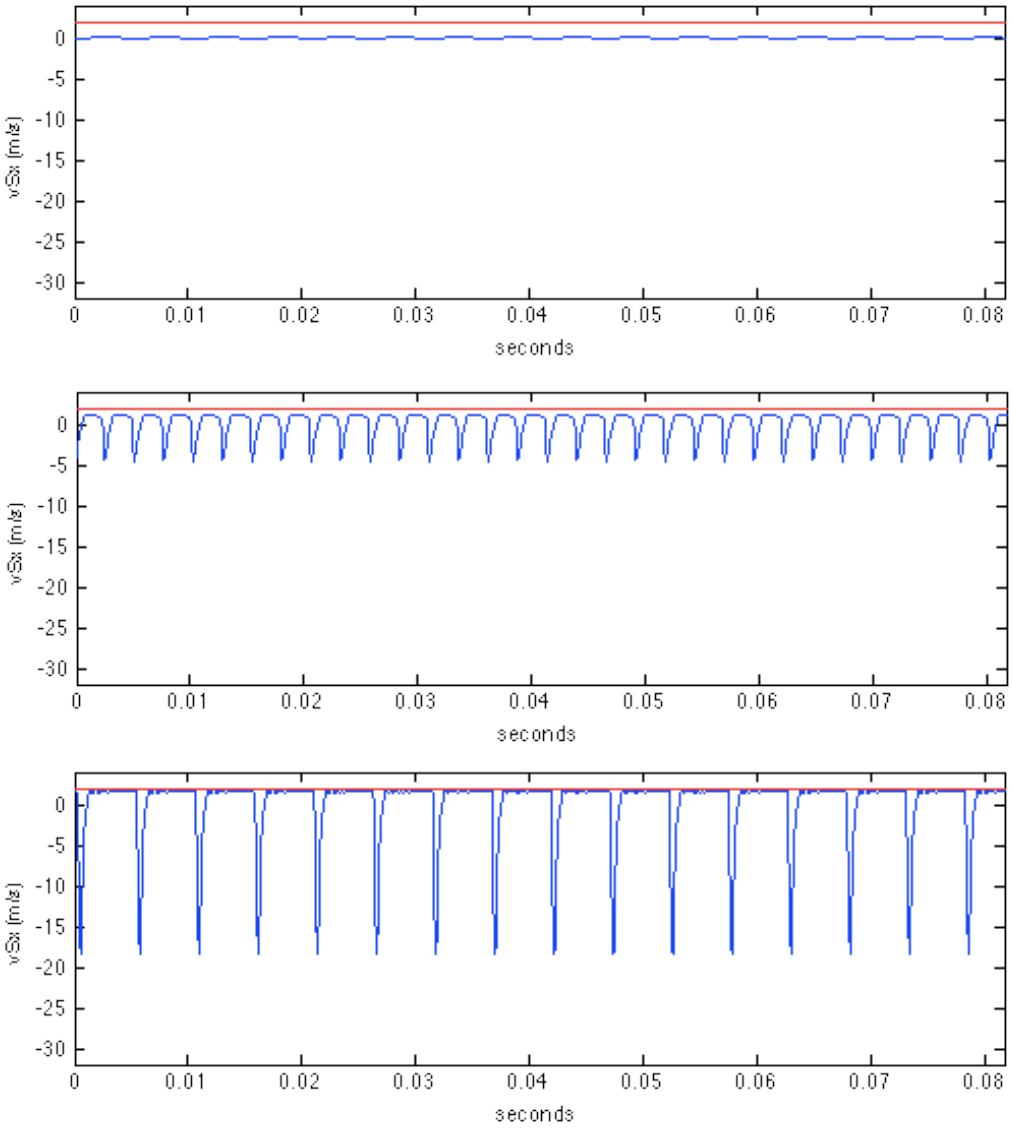

Figure 3.1: Temporal evolution of the differential string velocity under different vibration regimes. The top plot corresponds to decaying sound, the middle one to regular Helmholtz motion, and the bottom one to multiple stick-slip.

Figure 3.2 outlines the values that these features used during the classification take along the Schelleng diagram space. Notice how regions with low energy get indeed classified as decaying, regions with high aperiodicity as aperiodic, and so on. The multiple stick-slip and ALF regions match very closely twice the fundamental frequency and half of it respectively, which are their expected values on the most prominent regions.

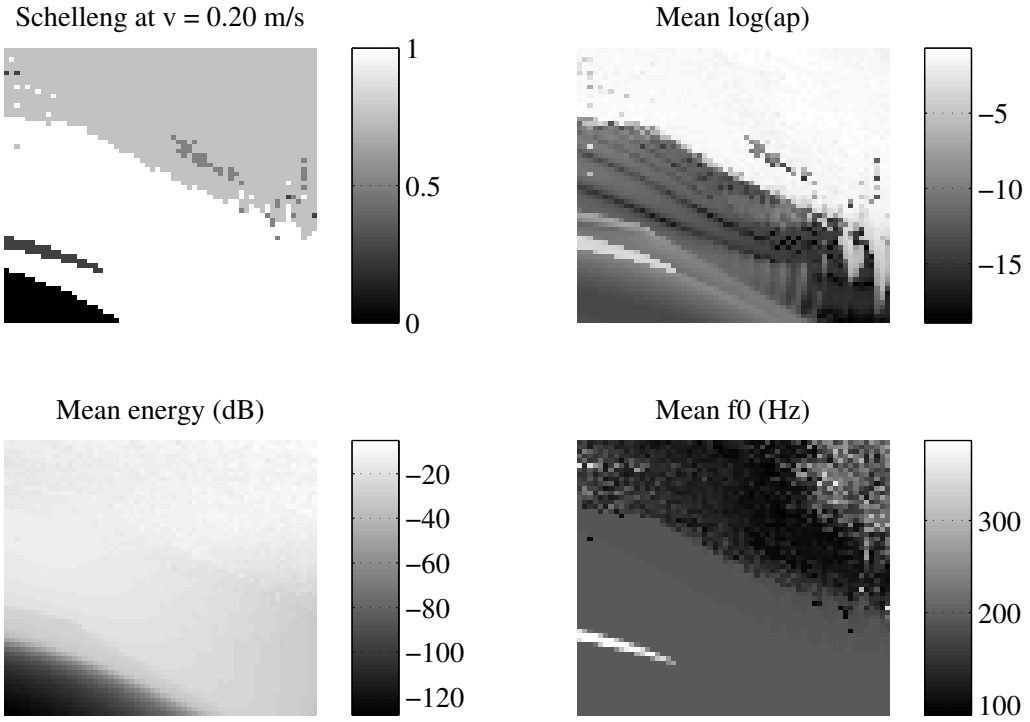

Figure 3.2: Features used for determining the vibration regime. The top plot on the left side shows the detected regime, with white being Schelleng, black being decaying, light gray aperiodic, gray ALF, and dark grey multiple stick-slip. The top right plot shows the detected aperiodicity of the signal in logarithmic scale for better contrast, the bottom left plot shows the mean energy, and the bottom right plot shows the detected fundamental frequency. The model was played at a bowing velocity of 0.2 m/s on its G string, with an expected fundamental frequency of 196 Hz.

Notice that the decision threshold for determining the octave (which discriminates Helmholtz, ALF, and multiple stick-slip regimes) needs a certain margin because of the *flattening* effect, which lowers the pitch when the bow force is high. This effect can be observed in the detected fundamental frequency subplot of Figure 3.2, and is thoroughly explained in McIntyre and Woodhouse [1979].

Because of the dynamic nature of the model, it is crucial to take care during the excitation or transient phase. At first, we probed the system by starting with no force or velocity, and applying a constant acceleration to both of them until the desired value. However, depending on the target values, the system got stuck in aperiodic motion at points where it was obvious that regular Helmholtz motion could be established. We found another pattern to work better, based on starting with the target force already on place, and applying dynamics only on the bow velocity.

Changing these profiles corresponds to playing with different gestures or *articulations*, that musicians traditionally use to give different characters to notes. The discussed excitation patterns correspond quite closely to the two main groups of bowing styles that musicians alternate during performance: the *legato* stroke, with a soft attack and a gradual force buildup, and the *martele* stroke, with a strong attack granting an explosive onset that adds "punch". See Galamian [1999] for a treaty on different bow strokes for a more complete description.

3.2 Playability measures

Most of the literature on playability has focused on rendering diagrams such as the ones proposed for different settings or variants of the models, and then qualitatively discussing which one is better. While this approach provides with great insight on the models and allows for extensive comments and discussions, we believe that a quantified approach is sometimes needed, especially when the number of trials to evaluate grows larger. Our goal is not only to run playability evaluations on a large number of versions of the model, but to construct *functional* maps that can hint how the playability changes according to model-related variables (as one smoothly navigates the parameter space). Therefore, we have devised some preliminary *scalar* indicators of bowed string playability.

As detailed previously, if one understands the playability of an instrument as a measure of how big are the regions of the control space that map to musically valid regions of the acoustic space, it becomes obvious that the Schelleng diagram can be of great help for a first approximation. It is only required to define a metric that, given the Schelleng diagram for a particular position in the parameter space, hints how well it fits into what is expected.

As an example, see Figure 3.3, which displays the computed Schelleng diagram

under identical conditions except for the "warm" coefficient of the parameter space, which controls how quickly the bow hair warms up due to friction. It seems, for instance, that the warmest version has less regions with decaying sound, and the Helmholtz region is therefore slightly wider.

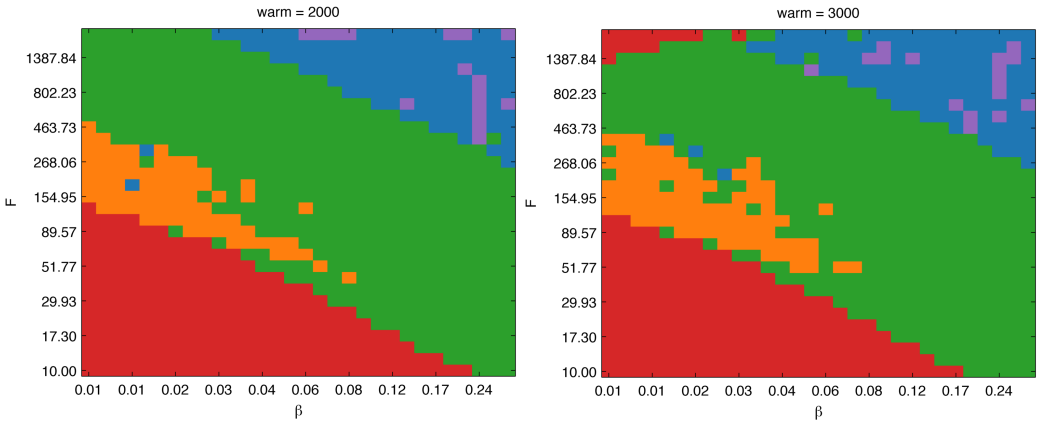

Figure 3.3: Computed Schelleng diagram for the model with two different "warm" coefficients. Different colors indicate different regimes; green means Helmholtz motion, blue means aperiodic, purple means ALF, orange means multiple stick-slip, and red means decaying.

If the region of the bow force and bow-bridge distance plane covered by the computed diagram is big enough, a valid first approximation to estimating how easy it is to lie inside the Helmholtz regime region would be to just count the portion of space that the Helmholtz regime (and each of the other regimes, just for completeness) takes. This will be our proposed playability metric for a given snapshot of the model, the results of which will be presented in the next chapter. This reflects what Serafin et al. [1999] proposes as playability; namely, "Playability *can be loosely defined as the* volume *of the multidimensional parameter space in which* good tone *is produced.*". Notice that their parameter space is what we call the *control* space, in order to distinguish it from the instrinsic instrument parameters.

Indeed, by estimating this playability indicator on a large number of adjacent points on the parameter space of the model (e.g. by sweeping a chosen friction-related coefficient), we can construct a *functional playability* map of our model, although for now limited to one and two-dimensional projections of the parameter space. Such functional maps have two main applications that makes them highly interesting:

- Automatic model calibration using non-convex optimization techniques on

the parameter space.
- Deeper qualitative understanding of the effect of the model parameters on the control-acoustic spatial mapping, and possible empirical law extraction.

Notice that when automatically calibrating the model, the objective function should not only take into account playability as has been defined here (a maximizer on the area of the Helmholtz region of the Schelleng diagram), but also the achieved tone quality - same as happens with acoustic instruments. Moreover, more realistic playability measures should be developed that also take into account the desired and musically useful aperiodic region, navigated especially in note transitions.

Unfortunately, these applications - especially the first one - lie outside the scope of this project, and are therefore left for future investigations.

The problem with our approach lies, of course, on determining what *big enough* means when deciding the portion of the control space to cover during playability estimation. In our case, we decided to select a range that makes sense both from a physical point of view, i.e., within the limits of what a musician could use (and adequately control) during performance - although we pushed bow force a little further, to see well the effect of aperiodic and ALF regions that are scarce on moderate and high bow velocities.

Chapter 4

COMPUTED SCHELLENG DIAGRAMS AND PLAYABILITY MAPS

In this chapter, we present the results obtained with our simulations using a complete model of the bowed string that includes thermal friction, hair dynamics with finite hair width, and a modal formulation of the bridge admittance. First, some individual results are presented of Schelleng diagram estimations, followed by functional playability maps generated by sweeping regions of the parameter space of the model and driving them with different gestures on the control space.

4.1 High-resolution regime maps

The computational efficiency of the model, especially thanks to its native implementation in C and its posterior parallelization in OpenMP, allowed us to compute Schelleng-like regime maps of great resolution, only matched by previous experiments run in expensive parallel computers. Actually, the bottleneck ended up being the regime analysis rather than the simulation itself. Figure 4.1 displays a rendering of the Schelleng space using a 50 by 50 lattice of control permutations.

Although this high resolution rendering was just a by-product, because our goal was to simulate lots of such lattices with less resolution, it provided us with deep

insights about the qualitative behaviour of the system. Two very important features that need further study arise at first sight: the shape of the playable region, and the distribution of the anomalous low frequency regions.

It has been long believed that there is a playable region in the Schelleng diagram separated by straight lines in a log-log representation. However, our detailed simulation shows a pronounced curvature on both limits of the region. This curvature varies depending on the constants of the model, and is a subject worhty of further study.

Another feature that requires more consideration are the anomalous low frequency regime (ALF) regions that arise. Traditionally, some sporadic points of the raucous region have been considered as displaying ALF, while others as aperiodic without any apparent order. This fact was probably attributed to the excitation of particular features of the system. However, our detailed analysis bowing at relatively low speed reveals several dense regions where low-frequency regimes are established. These regions have stripe-like shapes, and present a very clear pattern - in a direction roughly normal to the edge of the playable region, stripes are alternated with pure aperiodic regions. If one looks at the detected fundamental frequency of each low-frequency zone, it decreases as one goes further from the playable region, that is, as the force increases and the bow-bridge distance decreases.

In the regime classification scheme proposed, we purposely avoided mentioning something that most violin players are well aware of: if the bow force is too high, especially at low velocity, it gets to a point where the sound is no longer aperiodic or raucous for that matter, but just a series of clicks corresponding to the bow dragging the string along with it several millimeters, and then releasing it. This can easily go as slow as two or three clicks per second. Although it is easy to come upon this behavior when playing, there are not many references to it in the literature. However, our diagram proposes an explanation: they are actually anomalous low frequencies themselves, only at such a low frequency that they lie way beyond the hearing threshold of 20 Hz, and are therefore perceived as series of clicks.

Another simulation that was possible was the so-called 3D Schelleng diagram, which adds a third dimension that sweeps different terminal velocities. In Schoonderwaldt et al. [2008], the authors mechanically bow a string at different velocities, and present the results in separate Schelleng diagrams; we decided to take this concept further and present a single 3D rendering of the resulting volumes, where the manifolds that contain regions of the same regime are displayed. In the mentioned article the authors observed that, as the velocity increases, the maximum

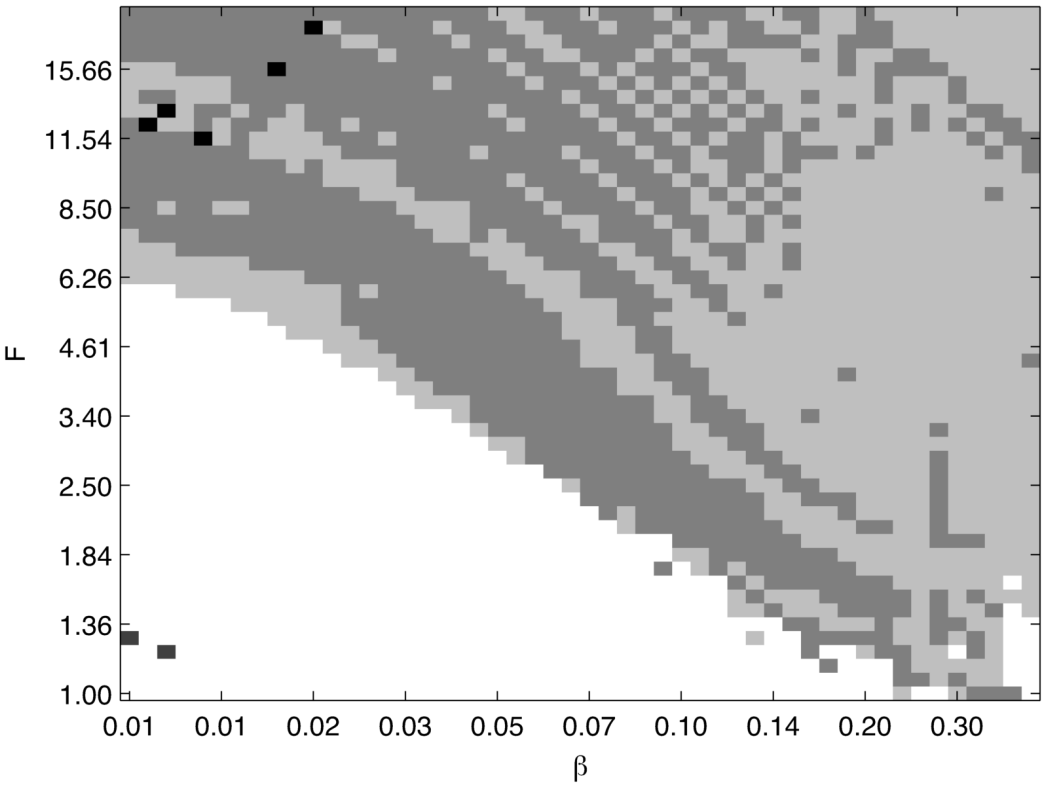

Figure 4.1: High-resolution rendering of the Schelleng diagram, focused on the raucous region. The pitch detector has been configured to detect very low pitches. White indicated Helmholtz motion, light grey aperiodic motion, dark grey anomalous low frequencies, and black other artifacts such as multiple stick-slip.

force before breaking Helmholtz motion is increased, but the slope of the minimum required force to achieve it is also modified; we can report similar results using our virtual synthesizer. Figure 4.2 provides a representation of such diagram, where the spatial distribution of the different modes is graphed; the change in slope can be noticed in the tf0 subplot (corresponding to Helmholtz motion), while the change in maximum force is more obvious in the aperiodic subplot. Notice also how regions with multiple stick-slip regime increase greatly at higher terminal velocities.

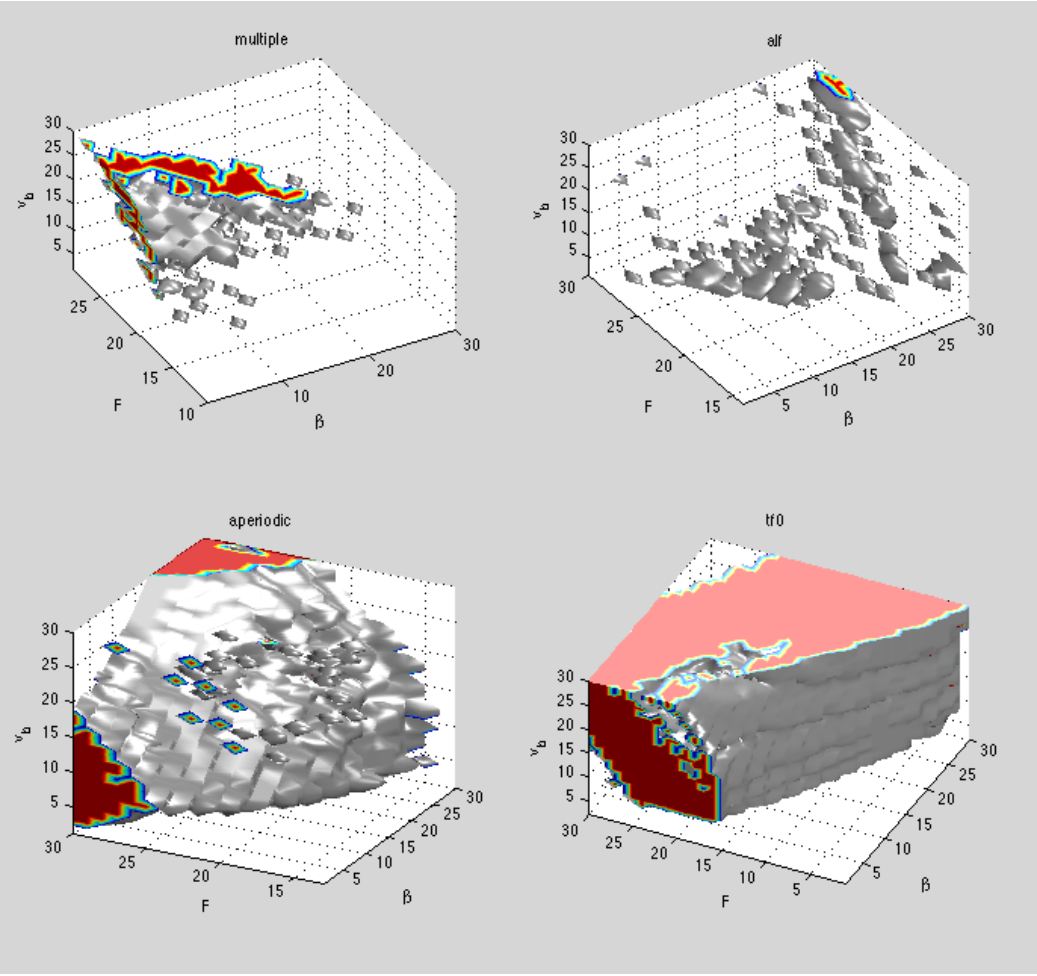

Figure 4.2: 3D Schelleng diagram displaying how the regime distribution on the force vs bow-bridge distance plane varies according to velocity. Axes do not indicate the actual units, but just the simulation index.

4.2 Playability

As a very rough measure, we will sweep a large portion of the feasible parameter space, and we will count the percentage of this space that lies in each of the possible regimes. Figure 4.3 shows these counts and how they change when the "warm" coefficient of the parameter space is modified. We can see that the portion of full Helmholtz motion and aperiodic regions grows with the coefficient, while the portion of decaying sound shrinks. It is difficult to give an explanation to this fact, but since large portions of space with decaying sound have it because of numerical artifacts, it could be the case that the system is more stable when it heats up faster. It could also have a more classical explanation whereby at low forces, which is when we would expect decaying sound, the contact does not heat up enough to trigger oscillatory motion; if it heats faster, it will be easier to reach that threshold.

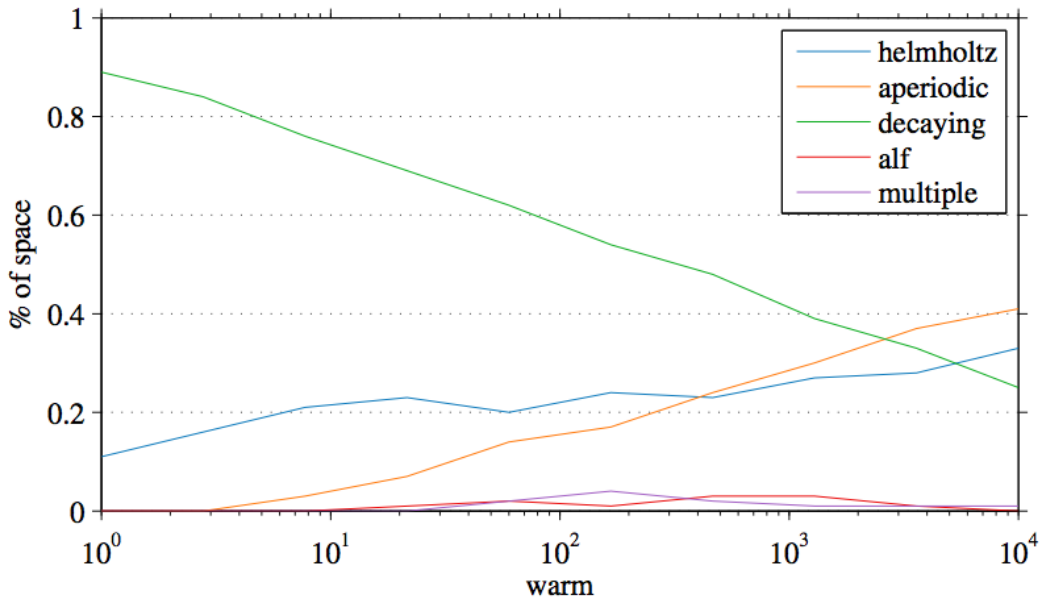

Figure 4.3: Functional playability of the model, depending on how quickly the bow hair heats up during friction.

This analysis can be expanded to a 2D parameter sweep; in Figure 4.4, we can see how the regime distribution changes depending on both the "warm" (already presented) and "cool" (which controls how quickly the system cools down) coefficients. Notice how the decaying regime shows that the model barely works when the coefficients are opposite. In the top left corner, where the system heats quickly and cools slowly, the rosin does not cool down enough during sticking and heat builds up; in the real world, the string would melt! This is of course unrealistic. In

the bottom right corner, where the system heats slowly and cools quickly, the rosin does not warm up enough during slipping to initiate periodic motion. As we go to the center from this bottom right corner, regions of the Schelleng diagram with high force (which heat up more) already bridge the critical point, and aperiodic regions appear. The Helmtholz region is similar, because these bridged points correspond to high forces unlikely to result in Helmholtz motion. It is also interesting to see how multiple stick-slips do not happen at all when cooling is slow. Finally, the most playable heat configuration seems to be near the top right corner, where the Helmholtz region is broader.

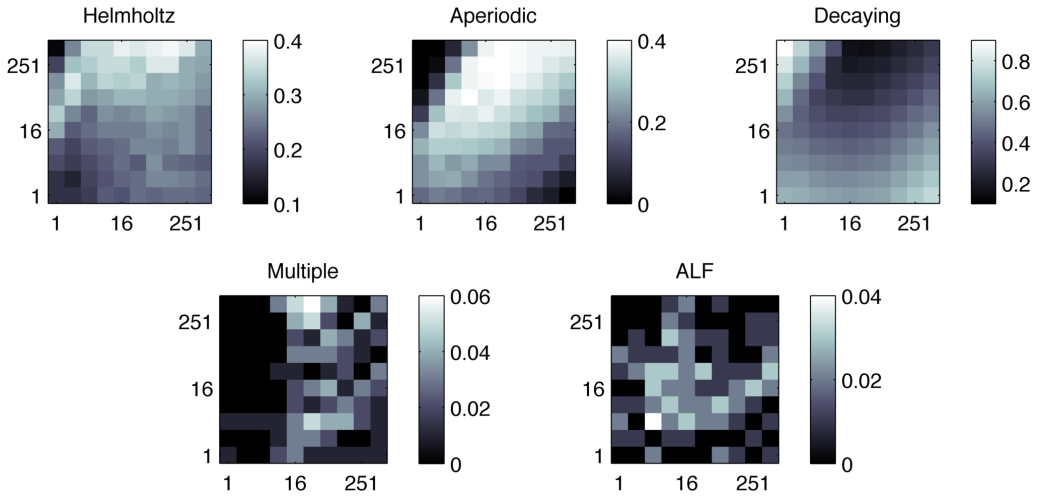

Figure 4.4: Two-dimensional functional playability map, displaying the regime distribution on the Schelleng diagram for differentand *cool*, horizontal axis (how quickly the hair cools), and *warm*, vertical axis (how quickly the bow hair heats up with friction) model parameters. The units have been reduced two orders of magnitude for better display. The color intensity represents the percentage of the selected Schelleng diagram that each of the regimes covers.

Chapter 5

OUTLOOK

Although this work only scratches the surface of the experimental analysis that can be performed in our bowed string model, we have successfully achieved our objectives and demonstrated that it is indeed possible to extract meaningful information from computer simulations similar to the ones described.

We have successfully computed Schelleng diagrams for given configurations of the model, which has lead to the discovery of Anomalous Low-Frequency stripes in the raucuous region with an unexpected organization.

The efficiency of the diagram extraction method has also allowed us to compute it in three dimensions, incorporating terminal bow velocity as another of the variables rather than performing a handful of simulations at different speeds, although we found visualization issues and only extracted from it a portion of the information the diagram for sure conveys.

We have also experimented with how the diagram changes when specific parameters of the model are changed, taking the "warm" coefficient as a case study, and from this we have constructed a *functional playability* indicator that shows how the overall playability changes when the parameter space is navigated.

These functional playability maps hint towards automatic model calibration, by exploring the parameter space while optimizing the playability metric on it. This could give rise to important future work, since model calibration is still one of the major issues in using physically-based synthesis for musical purposes; most of the parameters are too far from the physical reality to be determined experimentally,

especially when they relate to very intrinsic properties of the materials that make up the violin, which tend to be compound and very complex.

Besides this and the obvious future research lines that the ALF strips and the 3D representation of the Schelleng diagram offer, there are several aspects of this project that call for refinement:

- The incorporation of control dynamics to the playability metric through Guettler-like diagrams.
- The evaluation of playability beyond the Helmholtz regime, i.e., by studying how players use aperiodic transients in a musical context.
- A depper sweeping of the parameter space, by taking into account all possible parameters. Notice that this is no small task, since the parameters are most likely dependent (so sweeping them separately will be helpful but not complete), and the number of parameters can be increased by reconsidering some design choices during the design of the model that admit other options.

While it may seem there are lots of future work to do and that may make the project look incomplete, I think that's quite the contrary; realizing that there are so many *really* open questions has given me the energy to pursue them even further, and I sincerely hope to do my best at giving an answer to them in the future.

I would like to express my deepest gratitude to Esteban Maestre, for teaching me all I know (especially what I was not aware that I did not know) about what happens when you play the violin and how to tell a computer to do it. It is being really rewarding to see my intuition as a violinist to make sense with what science says and feel like it all "clicks". I would also like to thank all the professors of the Master's, for making me realize mathematics and modelling are more than equations and for providing me with intuition and insight about models that may look complex on paper but actually represent everyday phenomena. Finally, I would also like to thank Marina, my wife, for supporting and taking care of me during these period of intensive work. This thesis was started in Barcelona and finished in Montréal, and moving across continents has not been an easy task.

Barcelona / Montréal, 2015

Bibliography

A. Benade. *Fundamentals of Musical Acoustics*. Dover Publications, second edition, 1990.

L. Cremer. *The physics of the violin*. MIT Press, 1984.

A. de Cheveigné and H. Kawahara. YIN, a fundamental frequency estimator for speech and music. *Journal of the Acoustical Society of America*, 111(4): 1917–1930, 2002.

F. G. Friedlander. On the oscillations of the bowed string. *Proceedings of the Cambridge Philosophical Society*, 49:516–530, 1953.

I. Galamian. *Principles of Violin Playing and Teaching, 3rd edition*. Shar Products Co., 1999.

H. von Helmholtz. *Lehre von den Tonempfindungen*. Braunschweig, 1862.

M. Karjalainen and Julius O. Smith. Body modeling techniques for string instrument synthesis. In *Proceedings of the 1996 International Computer Music Conference*, Hong Kong, 1996.

J. B. Keller. Bowing of violin strings. *Comm. Pure and Applied Maths*, 6:283–495, 1953.

H. Kinoshita and S. Obata. Left hand finger force in violin playing: Tempo, loudness, and finger differences. *Journal of the Acoustical Society of America*, 126(1):388–395, 2009.

H. Kuttruff. *Room acoustics*. Taylor & Francis, New York, USA, 5th edition, 2009.

E. Maestre, G. Scavone, and J. O. Smith. Digital modeling of bridge driving-point admittances from measurements on violin-family instruments. In *Proceedings of the Stockholm Music Acoustics Conference*, 2013.

E Maestre, C. Spa, and J. O. Smith. A bowed string physical model including finite-width thermal friction and hair dynamics. In *Proceedings of the 2014 International Computer Music Conference*, 2014.

H. Mansour, J. Woodhouse, and G. Scavone. Time-domain simulation of the bowed cello string: dual-polarization effect. In *International Congress on Acoustics*, 2013.

M. E. McIntyre and J. Woodhouse. On the fundamentals of bowed string dynamics. *Acustica*, 43(2):93–108, 1979.

J. Piteroff and J. Woodhouse. Mechanics of the contact area between a violin bow and a string, Part II: simulating the bowed string. *Acustica united with Acta acustica*, 84:744–757, 1998.

C. V. Raman. On the mechanical theory of the vibrations of bowed strings. *Bulletin of the Indian Association for the Cultivation of Science*, 15:1–158, 1918.

J. C. Schelleng. The bowed string and the player. *Journal of the Acoustical Society of America*, 53:26–41, 1973.

E. Schoonderwaldt, K. Guettler, and A. Askenfelt. An empirical investigation of bow-force limits in the schelleng diagram. *Acta Acustica united with Acustica*, 94:604–622, 2008.

S. Serafin, J. O. Smith, and J. Woodhouse. An investigation on the impact of torsion waves and friction characteristics on the playability of virtual bowed strings. In *Proceedings of the 1999 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, New Platz, New York, USA, 1999.

J. H. Smith and J. Woodhouse. The tribology of rosin. *J. Mech. Phys. Solids*, 48: 1633–1681, 2000.

J. O. Smith. Physical modeling using digital waveguides. *Computer Music Journal*, 16(4):74–91, 1992.

J.O. Smith. Nonlinear commuted synthesis of bowed strings. In *Proceedings of the 1997 International Computer Music Conference*, 1997.

Julius O. Smith. *Physical Audio Signal Processing: for Virtual Musical Instruments and Digital Audio Effects*. W3K, 2010.

J. Woodhouse. On the playability of violins. part ii: Minimum bow force and transients. *Acustica*, 78:137–153, 1993.

J. Woodhouse. Bowed string simulation using a thermal friction model. *Acta Acustica united with Acustica*, 89:355–368, 2003.

J. Woodhouse and P.M. Galluzzo. The bowed string as we know it today. *Acustica - Acta Acustica*, 90(4):579–589, 2004.

D. Young and S. Serafin. Playability evaluation of a virtual bowed string instrument. In *Proceedings of the 2006 International Conference on New Interfaces for musical Expression*, pages 104–108, Montréal, 2003.