STANFORD UNIVERSITY, CCRMA

Research Report: Hybrid Clarinet Project

Romain Michon and John Granzow

December 4, 2013

In the past thirty years we have seen a wide range of studies on the physical modeling of musical instruments using the waveguide technique. We now have at our disposal sophisticated models encompassing a large array of musical instruments. While the waveguide technique is both efficient and allows for the creation of effective models, it remains dependent on the precise control of the parameters of the model as well as on the quality of the excitation that is used to drive them. Many works on the control of waveguide physical models and on the modeling of nonlinear excitations have been carried out. However, the link between these two parameters is generally understudied.
Indeed, for most musical instruments, the excitation is the element that has the greatest number of parameters to control and is thus the hardest element to model. The properties of the bore of a clarinet for example (but this also applies to most of the woodwind and brass instruments) can only be modified by tone holes, that are very discrete controllers (they can be opened or closed or half closed in some cases). On the other hand, the interactions between the mouthpiece and the player are extremely complex and difficult to simulate on a computer. This problem has been addressed in many ways in the past with different solutions for almost every case. Yamaha, for example, created a breath controller that worked with the VL1 synthesizer series[1]. For violins, Esteban Maestre developed a technique where the parameters of the physical model are controlled by gesture data acquired from real world performers[3]. The technique we present in this report uses a 3D printed mouthpiece combined with a system based on a piezo sensor and a loudspeaker to drive a simple physical model of a clarinet bore. The idea was to create a hybrid instrument based on a real mouthpiece and a virtual bore. We

[1] `http://www.patchmanmusic.com/yamahaVL1.html`.

discuss the results of various experiments we carried out and try to provide solutions to the problems encountered.

1 Description of the project

The use of "real world" excitation to drive virtual physical models has been experimented in the past [1] [5] [6] [8] [9] but not many publications exist on this topic. In the case of percussion and plucked string instruments, the process is very simple as any kind of audio impulse can be fed into a waveguide to excite it. For example, in the case of the Black Box project [5], piezo sensors are glued to acrylic plates that can be touched, stroke, etc. by users to drive a physical model of a metal plate. Almost no processing is carried out on the excitation signal that carries the "acoustical shape" of the plate. Similarly, the Kalickord [8] uses the audio signal created by piezo films to drive a string model.

While this technique can be theoretically applied to any physical model of musical instruments, things become complicated when the excitation signal partly determines the pitch of the sound generated by the system. This is the case for woodwind and brass instruments for example where there is a coupling between the mechanism that produces the excitation and the size of the bore of the instrument ([7] pp. 214-215). In other words, the length of the bore determines the frequency of vibration of the reed on a clarinet and of the lips on a trumpet. This is due to the negative pulse reflected by the end of the pipe of the instrument ([2] pp. 461-500) whose distance from the reed can be adjusted with the tone holes in the clarinet case and with the pistons on the trumpet. The goal of the project presented in this report is to leverage 3D printing technology to that of physical modeling to create a hybrid single reed instrument whose 3D printed mouthpiece is connected and coupled to a virtual computer modeled bore. The most challenging part of this work is to create a feedback system that reproduces the negative pulse created by the end of the bore and send it to the mouthpiece using a simple speaker.

The advantages of such a system (if it works :)) are multiple. It would make it possible to design hybrid instruments whose response was partly independent of its physical geometry. The player would have full control of the generated sound and experience some of the vibrotactile stimulation of a real instrument.

2 First model and experiments

2.1 3D printed mouthpiece and feedback system

Figure 2.1 depicts the computer model of a tenor saxophone mouthpiece fabricated for our experiment[2]. It is based on a real mouthpiece that was graciously lent to us by the Stanford

[2] Cf. `/models` folder.

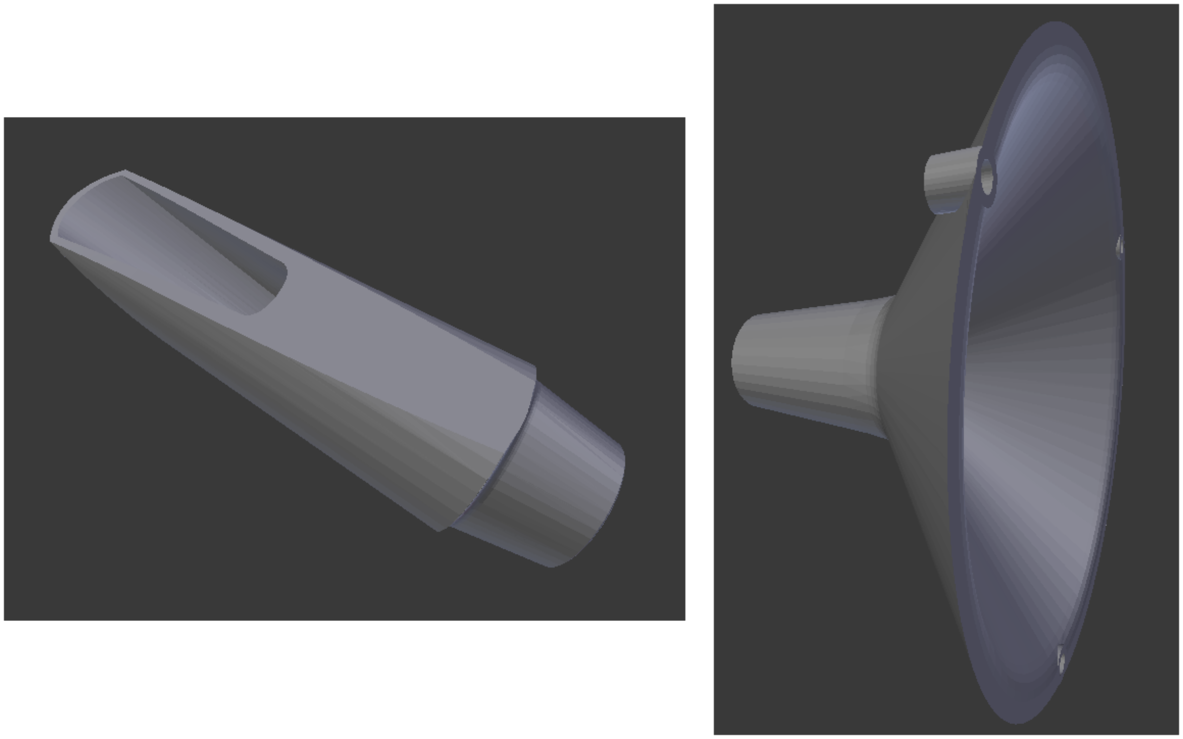

Figure 2.1: Computer model of our first mouthpiece feedback system. The `.stl` files of this model can be found in `/model`.

music department. We tried the printed mouthpiece it with a “real” saxophone and it appeared to be of better quality than the original. Although it would be interesting to know why, this is not the object of this study. The other object depicted in the figure connects to the mouthpiece and houses a speaker. Figure 2.2 presents the printed version of these elements with a bamboo reed mounted on the mouthpiece and the speaker attached to the other section. We used this system to conduct most of the experiments that are presented in this report. Originally, a microphone was placed inside the mouthpiece to pick up the audio excitation created by the reed. This was obviously a very bad solution as the microphone was too closed to the speaker, creating feedback. We first tried to solve this problem by replacing the microphone by a piezo film glued on the reed. However, the impedance between the reed and the piezo was too high and we were not able to get usable signals. The solution we adopted was to glue a very small piezo disc (whose much stiffer than a piezo film) on the reed. This enabled us to get high quality signals of the excitation without creating feedback with the speaker.

The position of the piezo on the reed was chosen so that it is as close as possible from the vibrating area without affecting the sound quality. As the piezo stays in a very humid environment, we had to find a way to protect it from becoming too wet. For that, we drilled a very small hole in the reed with an identical diameter than the one of the piezo. Finally, we covered the piezo with a protection paste (cf. figure 2.3).

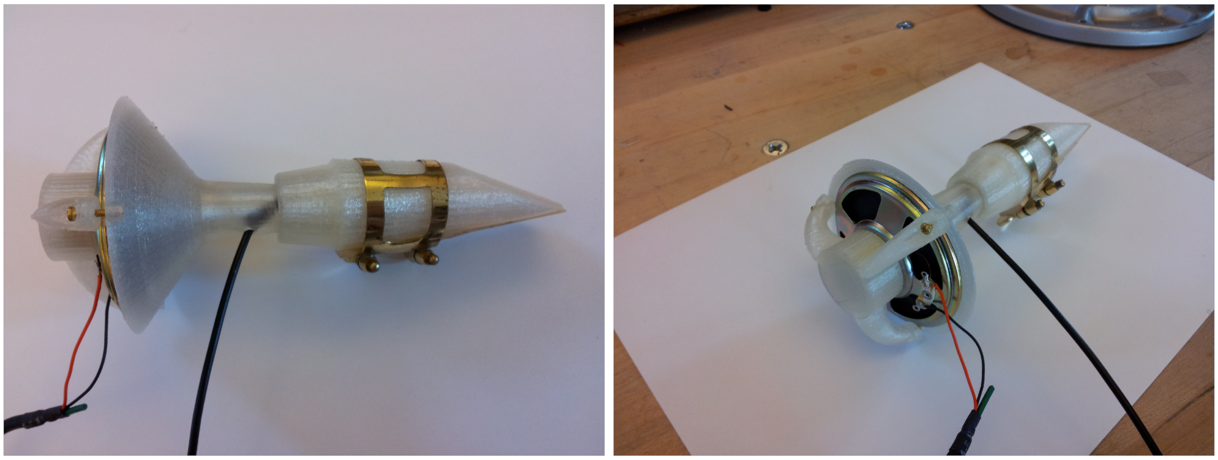

Figure 2.2: 3D printed mouthpiece feedback system.

In practice, after several minutes of playing, the entire reed becomes wet which introduces artifacts in the signal created by the piezo. We temporarily solved this problem by "grounding" the mouth of the player, achieved by putting a wire linked to the ground in the mouth. However, a better solution will have to be found in the future.

2.2 The physical model

The physical model we used for our experiments is based on the clarinet model of the Faust-STK [4]. Thus, it is implemented in Faust[3]. Figure 2.4 presents the top level block diagram of this algorithm. The input signal (x), for instance the signal generated by the piezo on the reed is fed into a waveguide. The reflection filter is a simple one-zero filter which also reverses the sign of the signal. The output of this filter is retrieved and sent to the speaker placed in front of the mouthpiece (cf. figure 2.2). Before that, it is scaled and it goes through a delay line that allows to control the phase of the signal (however, this proved to be totally useless as changing the phase had absolutely no effect on the behavior of the reed).

2.3 First experiment

In order to achieve the best latency with the audio interface we used[4], we had to use a very high sampling rate of 192KHz with buffers of 128 samples and 2 periods per buffer. Unfortunately, we were never able to use the PD-externals compiled from the Faust code at such a high sampling rate. The trick we used to to compensate for this was to compile our Faust code as jack applications that we controlled from PD[5] sending OSC messages. We chose PureData

[3] `http://faust.grame.fr`.
[4] We used a Roland UA-101 for all our experiments.
[5] *PureData*: `http://puredata.info/`.

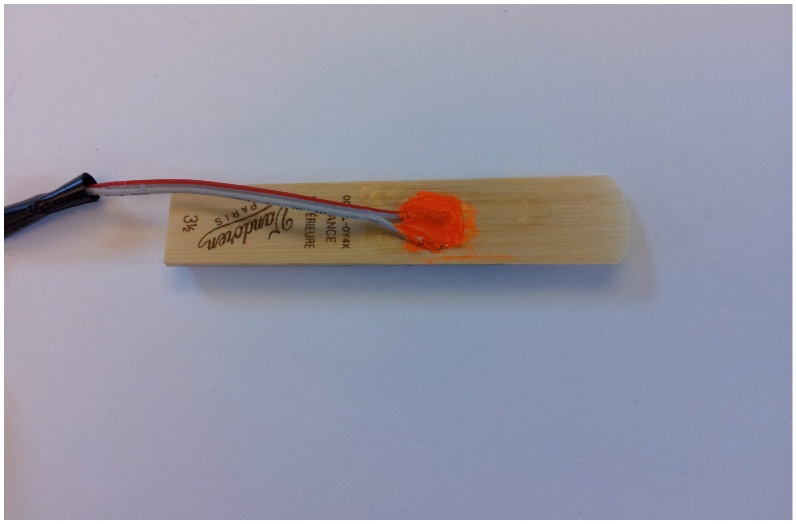

Figure 2.3: Tenor saxophone reed with a piezo disc glued on it. The orange material is a paste that protects the piezo from humidity.

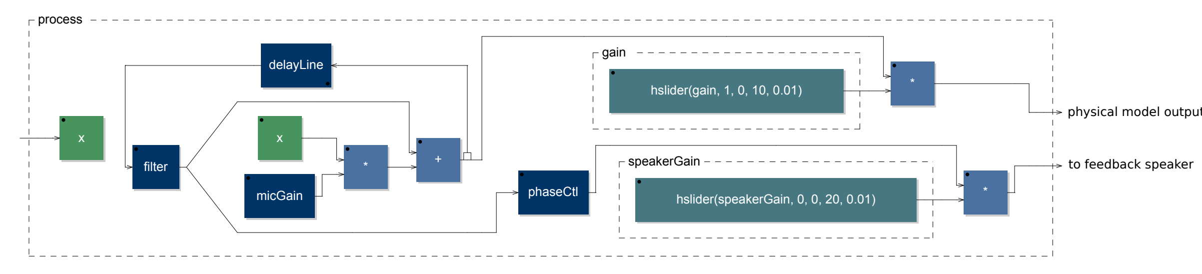

Figure 2.4: Block diagram of the Faust code `hybrid.dsp` (available in `/realHybrid`).

to carry out this task because it is easy to quickly create controllers in this environment.
For our first experiment, we tried the system with different lengths for the virtual bore, using a very high amplitude mouthpiece feedback signal. While we were not able to change the frequency of vibration of the reed, the saxophone player found it very hard in some cases to hold the pitch and felt really disturbed by the behavior of the reed when the length of the bore was set such that it approached the natural frequency of vibration of the reed.
After these primarily experiments, we figured out that the latency introduced by our audio interface, even though it was very small (around 2ms with the configuration described at the beginning of this section), was one of the problems in the system. Indeed, a latency of 2ms in an environment where speed of sound is 340m/s corresponds to a delay of 68cm which is huge in the case of musical instruments.
Also, we thought it would be interesting to do more basic experiments like sending a square wave signal into the mouth piece. The implementation and the results of these experiments are presented in the two following chapters.

3 Square wave experiments

The waveform of an audio signal created by a clarinet is similar to a square wave. For this reason, we thought it would be interesting to try to send a square wave audio signal at different amplitudes and frequencies in the mouthpiece while blowing in it.
In our first experiment[6], we tried to send a high amplitude signal where the frequency of the square wave was the same as the natural frequency of vibration of the reed (665Hz) and slowly increased it to 1000Hz. We recorded the signal from the piezo on the reed and plotted the spectrogram (cf. figure 3.1).
We can see that after 3 seconds, the frequency of vibrations of the reed starts to be modulated, creating a vibrato effect that evolves into some kind of frequency modulation behavior after 12 seconds. Thus, somehow the square wave acts as a modulating signal on the reed which can be compared to a carrier in this case.
Another experiment[7] where a square wave signal with a constant frequency of 670Hz and an increasing gain is sent to the mouthpiece. The spectrogram of the signal recorded on the reed is plotted in figure 3.2 along with the waveform of the square wave. We can see that increasing the gain of the square wave broadens the range of the modulation, exactly like a frequency modulation synthesizer when the index of modulation is changed.
These results prove that the frequency of vibration of the reed can be modulated by the signal coming from a speaker if the amplitude of the wave it creates is big enough. However, it is important to note that this experiment was particularly unpleasant and unnatural for the saxophone player.

[6] The files related to this experiment can be found in `/testSquare/matlab`.
[7] *Ibid.*

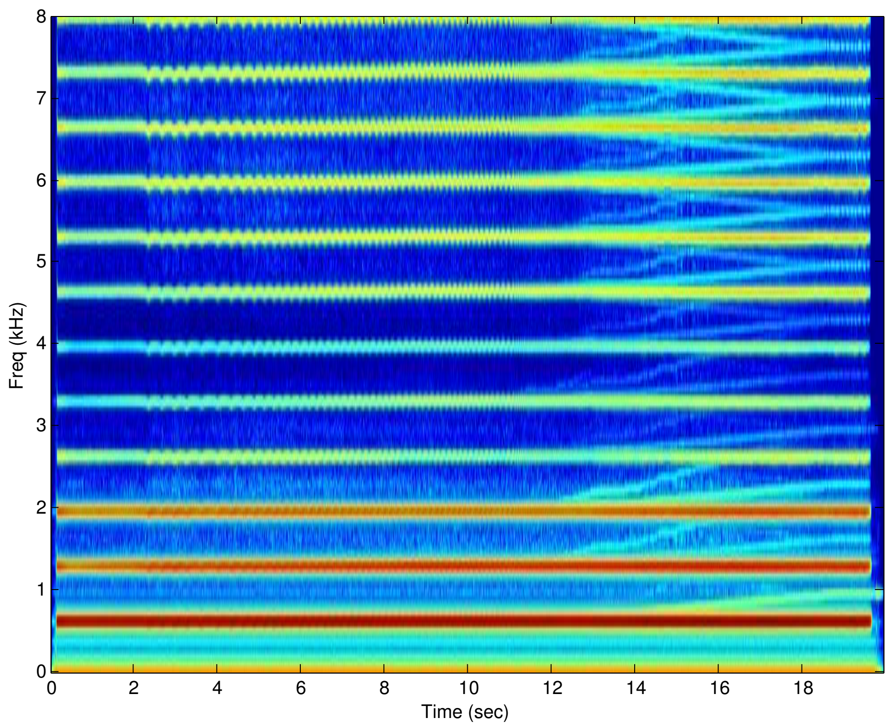

Figure 3.1: Spectrogram of `staticAmp-incFreq-reed.wav` (available in `/testSquare/matlab`). The frequency of a square wave starting at 665Hz is increased to reach 1000Hz and is sent to the mouthpiece.

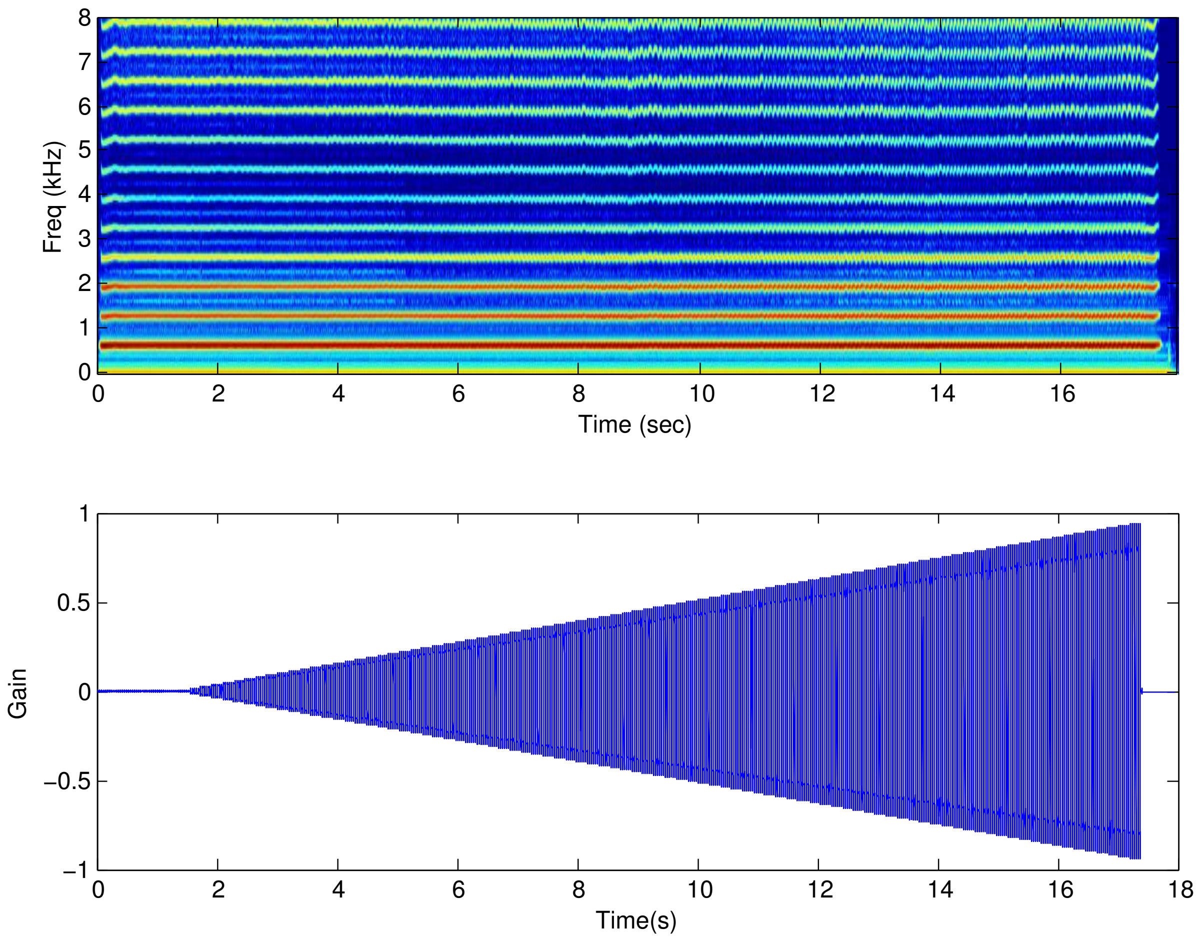

Figure 3.2: Spectrogram of `ampChange-reed.wav` (available in `/testSquare/matlab`). The amplitude of a 670Hz square wave is slowly increased and sent to the mouthpiece. The lowest graph shows the waveform of the square wave sent into the mouthpiece. It is aligned with the spectrogram.

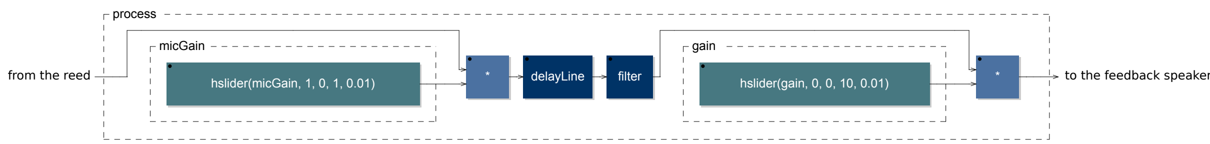

Figure 4.1: Block diagram of the Faust object `zeroLatency.dsp` (available in `/zeroLatency`).

4 Limited "zero-latency" system

One way to achieve 0ms latency in our system was to create a model[8] where the latency of the audio interface is used as the delay line of the waveguide. A diagram of this model can be seen in figure 4.1. We basically take the signal from the reed and use the same filter as the one used in the model depicted in figure 2.4 and send it back to the speaker. Some delay can be added to increase the size of the simulated bore whose minimum length is defined by the latency of the audio interface (65cm/2 = 32.5cm).

We tried to use this model by slowly increasing the length of the virtual bore starting at 32cm (the minimum length we can achieve with this configuration) and ending at 132cm. The spectrogram of the signal recorded from the reed can be seen in figure 4.2.

We can see that as the length of the virtual bore is increased, the frequency of the reed is shifted down by a semitone and comes back to its original state every time the minimum size of the bore is doubled (every time 32cm is added). This proves that despite the fact that our system is not able to control the frequency of vibration of the reed, it can shift it a little bit.

5 Current study and experiments

5.1 Diminishing latency further

While the trick we presented in the previous chapter to get rid of latency issues works, it is not optimal because it imposes a minimal length to the virtual bore. Moreover, it increases the unpredictability of the model as the addition of the feedback signal with the reed signal is not carried out on the computer but in the real world.

The only solution to this problem is to reduce the latency of the system (adc -> computing -> dac) we're using. One of the possibility we are currently studying is to use a Digilent Atlys FPGA development board[9] with the WaveCore technology developed by Math Verstraelen. Theoritically, we should be able to achieve a latency inferior to 500 micro seconds using this system. The main advantage of this solution is its (theoretical) compatibility with Faust and

[8] The files related to this experiment can be found in `/zeroLatency`.

[9] `http://www.digilentinc.com/Products/Detail.cfm?NavPath=2,400,836&Prod=ATLYS&CFID=3252896&CFTOKEN=81292424`.

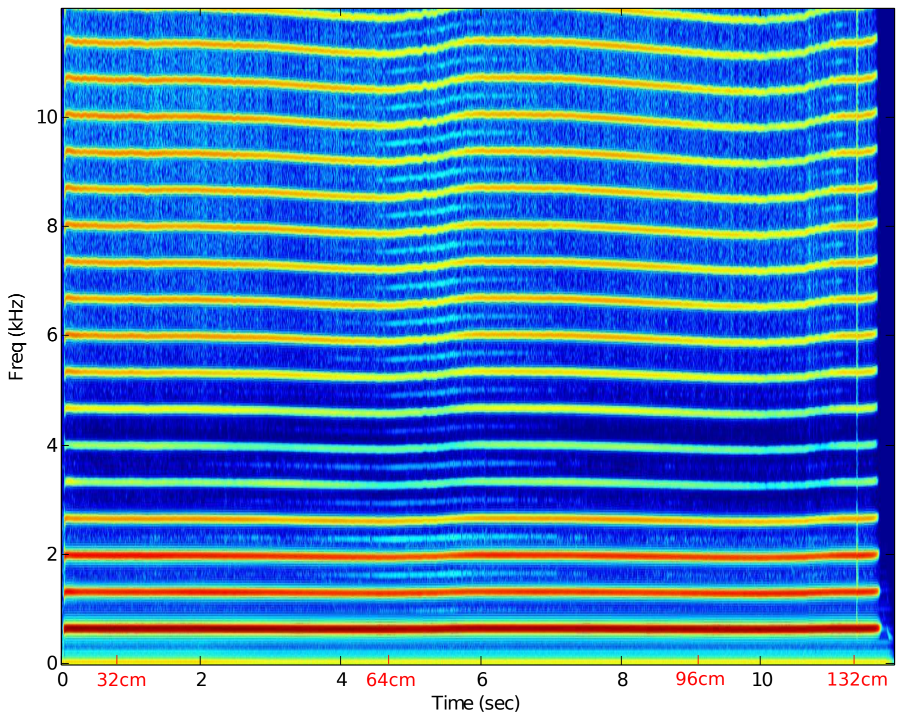

Figure 4.2: Spectrogram of `zeroLatency.wav` (available in `zeroLatency/audio/matlab`). The length of the virtual bore is slowly increased from to 32cm to 132cm.

the fact that it would make it possible to build standalone instruments that would not require the use of a laptop.

5.2 Improving the mouthpiece feedback system

We think that one of the main reason our system is not working is because the “real” reflected wave created by the mouthpiece and the small pipe that links it to the speaker chamber arrives earlier and is stronger than the virtual reflection wave generated by the speaker. Therefore, we designed a new system (cf. figure 5.1) where we will try to drive the wave created by the reed in a pipe. An acoustical damping material will be placed at the end of the pipe to cancel this wave in order to prevent it from being reflected. The speaker will be placed at the middle of the pipe and thus will be able to send the virtual reflection before the “remains” of the damped wave reach the reed. We hope this solution will work or provide new clues for the next step. In the worst case scenario, it can be used as a vuvuzela for football games.

Creating hybrid musical instruments where the excitation is generated in the “real world” and where the instrument itself is computer modeled opens new possibilities of interactions and playability. The generated sounds are usually more realistic than with traditional physical models as the performer has more input on the instrument. While this technique works very well with plucked strings and percussion instruments, lots of work remains to be done to apply this technique to a broader array of physical models.

References

[1] Berdhal, Edgar J. *Applications of Feedback Control to Musical Instrument Design.* Stanford University: PhD. Thesis, 2009.

[2] Fletcher, Neville H., and Rossing Thomas D. *The Physics of Musical Instruments.* Second Edition. Springer: New York, 1998.

[3] Maestre Gomez, Esteban. *Modeling Instrumental Gestures: An Analysis/Synthesis Framework for Violin Bowing.* Universitat Pompeu Fabra: PhD. Thesis, 2009.

[4] Michon, Romain. “Faust-STK: a Set of Linear and Nonlinear Physical Models for the Faust Programming Language.” In *Proceedings of the 14th Int. Conference on Digital Audio Effects (DAFx-11), September 2011.* IRCAM, Paris (France), 2011

[5] Michon, Romain, and Borins Myles, and Meisenholder David. “The Black Box.” In *Proceedings of the International Conference on Computer Interfaces for Musical Expression (NIME), May 2013.* KAIST, Daejon (South Korea), 2013.

[6] Puckette, Miller. “Playing a virtual drum from a real one.” *The Journal of the Acoustical Society of America,* no. 130(4) (2011): 2432-2432.

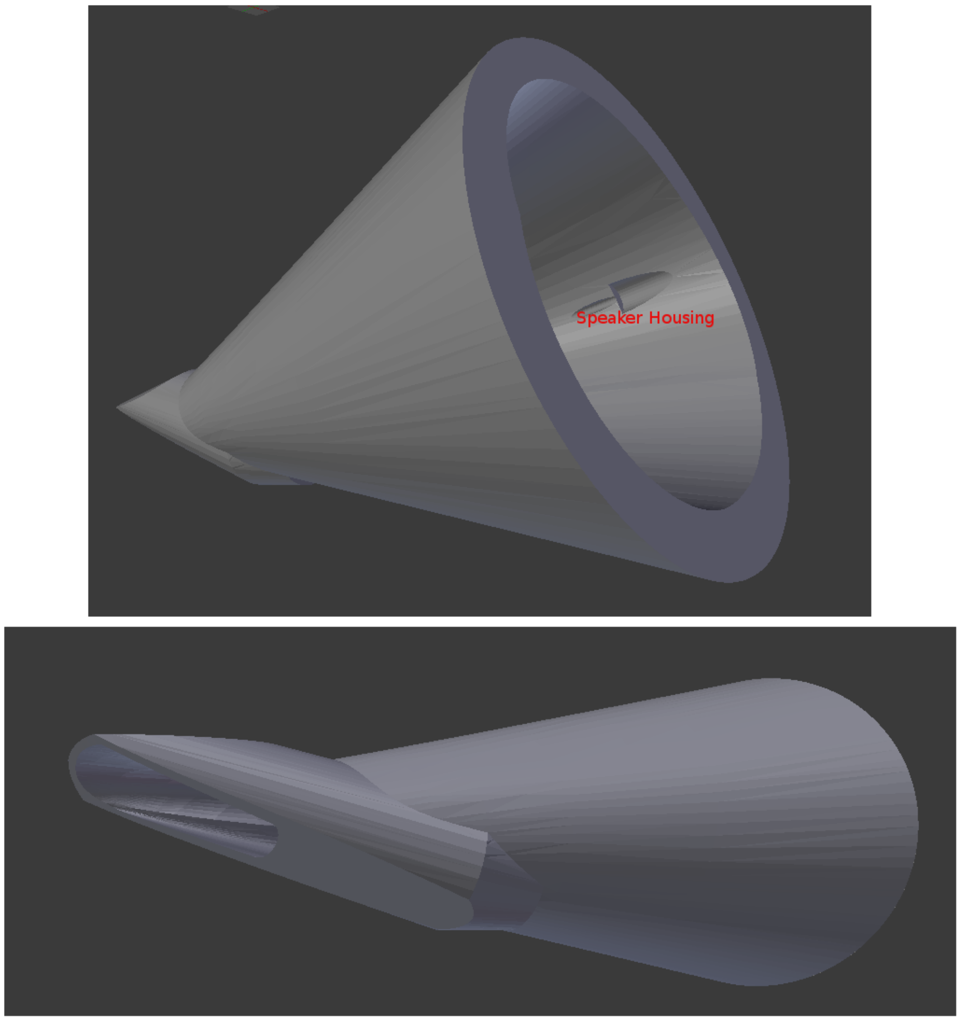

Figure 5.1: Our future mouthpiece feedback system. The wave created by the reed is carried until the end of the cone where it is attenuated by a damping material. The speaker is placed at the middle of the cone and is able to send the virtual reflection wave before the real one.

[7] Rossing, Thomas D. *The Science of Sound*. Addison-Wesley Publishing Company: Reading (Massachusetts), 1983.

[8] Schlessinger, Daniel, and Smith Julius O. III. “The kalichord: A physically modeled electro-acoustic plucked string instrument.” In *Proceedings of the 9th International Conference on New Interfaces for Musical Expression, June 2009*. Carnegie Mellon University, 2009.

[9] Serafin, Stefania, and Bottcher Niels, and Alberto Cruz Mario, and Fredslund Anders, and Gelineck Steven, and Lau Lorensen Christen, and Talaganov Goce, and Pozza Thomsen Peter. “Novel Interfaces for Controlling Sound Effects and Physical Models” In *Proceedings of 1st Nordic Music Technology Conference (NoMuTe), October 2006*. Trondheim (Norway), 2006.