SCALABLE SPECTRAL REFLECTIONS IN CONIC SECTIONS

Declan Murphy, Georgios Marentakis, Tue Haste Andersen, Kristoffer Jensen

Music Informatics Laboratory,
Department of Datalogy, University of Copenhagen
{declan,babis,haste,krist}@diku.dk

ABSTRACT

The object of this project is to present a novel digital audio effect based on a real-time, windowed block-based FFT and inverse FFT. The effect is achieved by *mirroring* the spectrum, producing a sound effect ranging from a purer rendition of the original, through a rougher one, to a sound unrecognisable from the original.

A mirror taking the shape of a conic section is constructed between certain partials, and the modified spectrum is created by *reflecting* the original spectrum in this mirror.

The user can select the type and continuously vary the amount of curvature, typically 'roughening' the input sound quite gratifyingly. We demonstrate the system with live real-time audio via microphone.

1. INTRODUCTION

This work grew out of a composition whose entire score was constructed out of reflections of notes in conic sections. This concept piece, "Mirrors—Espejos" for flute, violin and live electronics, was composed by Maria Esteves in close collaboration with the first author. The mirrors were embedded in two dimensional "stave space", with horizontal length corresponding to note duration and vertical height corresponding to note pitch. Beginning with a single note chosen arbitrarily, this was repeatedly reflected on the stave, sometimes being discretised in pitch and time, and at other times being left as glissandi and curious note timings. Melodic sections were built up from iterations of rotations, translations and reflections of notes in differently curved mirrors, which were themselves rotating, translating and changing their curvature in time.

To continue this concept, some kind of analogous processing was sought for the live electronics. This paper presents such a technique, one which can be applied to process the timbre of the instruments in real-time. The piece ("Mirrors—Espejos") has had two recorded public performances [1], and further performances are scheduled to include the more developed spectral mirroring presented in this paper.

The order of the sections of this paper broadly correspond to the modular chronological processing of the system, and to its historical development. The next section, §2, describes the overall modes of operation of our system, and its first stage of spectral processing: constructing the mirrors. §3 then describes how the conic section reflections are calculated, and how this is applied to the spectrum. §4 follows with details of complications which arise in translating the reflected spectrum back into an audio signal, and how they are tackled. §5 outlines the physical and conceptual layout of our implementation, and supplies an evaluation of the resulting effect. The paper concludes with a brief summary (§6) and references.

2. PRE-PROCESSING

The modification is done on the spectrum, found by taking the real part of the logarithm of the FFT. The phase (the imaginary part of the log of the FFT) is not modified. This ensures that partials of subsequent blocks remain aligned, independent of the amount of modification.

There are two modes of functioning: the spectrum can be reflected in one global mirror extending over its entirety (§2.2), or in individual mirrors for sections between perceivable partials (§2.1). In addition, a fast mode of vertical projection is possible, which does not move the frequencies of the FFT bins [2] (§2.1.2).

When using individual mirrors, the spectrum is partitioned into a number of subsections, each having local maximum peaks as upper and lower bounds.

2.1. Sectional Mirrors

The prominent partials used to create the individual mirrors are found using the perceptual criterion of masking [3]. First of all the absolute maximum of the spectrum is determined. The weak spectral components masked by this peak are disregarded (*i.e.* they are skipped in the selection of peaks between which mirrors are constructed). The next maximum is then found and all remaining spectral components masked by it are similarly disregarded. This process is repeated until no further peaks are found.

2.1.1. Conic Section Reflection

The conic section is constructed (*cf* §3.1) so that it passes through two consecutive prominent partials. Therefore all non-masked partials are reflected onto themselves in the mirroring (effectively preventing their modification), ensuring the possibility of very subtle effects. See figure 5 for a schematic overview of a dramatic effect.

2.1.2. Vertical Projection

A simpler scaled projection – perpendicular to the frequency axis – is provided as an alternative to the conic sections in the sectional approach, since it gives more robust results (and also a much simpler implementation). In this mode, only the magnitude of each FFT bin's coefficient is changed, basically changing the signal to noise ratio but generally without adding dissonance to the sound. See figure 1 for an example.

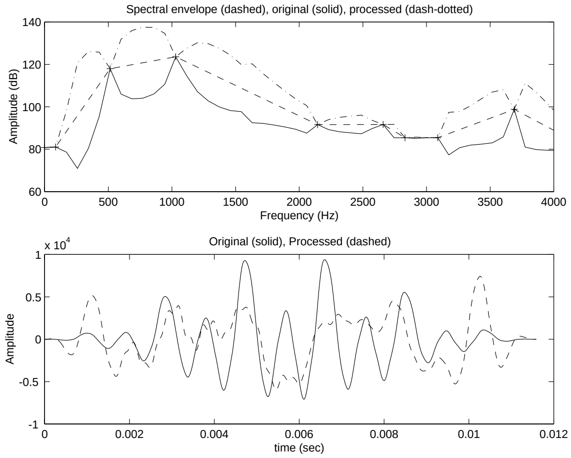

Figure 1: *Sectional vertical projection between prominent partials.*

2.2. Global Mirror

The user has the choice of constructing a single mirror specially chosen to suit the block in question.

The global projection/reflection operates not only on the weak masked spectral components but also on the strong partials. The mirror ranges over the full spectrum, and is manually positioned. Alternatively, the magnitudes of the FFT coefficients can be projected onto a line created to have equivalent brightness and energy [4], or onto the spectral envelope (which can be created, for instance, by using Linear Predictive Coding).

In general, however, global mirrors tend to be rather crude: they lack articulate control, and their results lack subtlety.

3. MIRRORING

Having decided upon the sections to be processed, the selected conic section is constructed between each successive pair of local maxima (§2.1), and the reflection of each spectral amplitude-frequency point is computed.

Depending on the curvature of the mirror, and on the relative positions of its focus and the spectral point, the reflection may be either real or virtual (*i.e.* lying on the opposite side of the mirror).

3.1. Conic Sections

As known since antiquity [5], the intersection of a double right circular cone and a plane describes a point, line, double line, circle, ellipse, parabola or hyperbola (see figures 2 and 3). These *conic sections* occur ubiquitously in solutions to differential equations modelling natural phenomena, and they possess many interesting properties in their own right [6]. The general form of a conic section equation is

$$Ax^2 + 2Bxy + Cy^2 + 2Dx + 2Ey + F = 0,$$

although the particular normal forms are simpler: '$y^2 = 4px$' defines a parabola, for instance. Having been studied for at least 2 300 years, there is much literature available on the subject: [7] provides a colourful introduction.

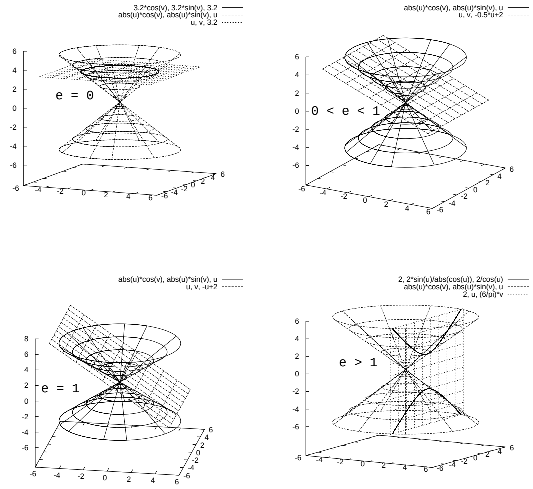

Figure 2: *How conic sections are formed by the intersection of a cone and a plane. Top row: circle, ellipse. Bottom row: parabola, hyperbola. The eccentricity, e, is defined by $e = \frac{\cos(\theta)}{\cos(\phi)}$, where θ is the angle between the axis of the cone and plane cutting it, and ϕ is the angle between the axis of the cone and a straight line on it.*

3.2. Image Computation

The reflection of a point is found by calculating the intersection of two carefully constructed rays (*i.e.* the image of a point, if it exists and is unique, may be calculated by tracing the intersection of two lines of reflection of that point, knowing that the angle of incidence equals the angle of reflection). As illustrated in figure 4, which rays are chosen depends upon whether the object point lies on the convex or concave side of the curve, whether it lies near the (major) axis or not, and on which side of the tangent to its turning point it lies.

On a historical note, understanding these complexities (avoiding division by zero situations in the derived equations) motivated the development of what is now the branch of mathematics known as complex analysis [8], in which functions having unbounded values at certain points can be dealt with rigorously.

Situations in which certain reflected points appear to be "at infinity", but which can be properly calculated by changing approach, correspond to the removable singularities of complex analysis. Other points exist, however, whose reflections really are boundlessly far away, which correspond to non-removable singularities, areas for which there really is no reflection. In the neighbourhood of such points, the images are not unique and appear blurred. This is reported as a warning by our implementation, and set to an approximate value.

4. POST-PROCESSING

4.1. Bin Realignment

The nature of curved surface reflections (*viz.* how the reflected light (or spectral points in our case) becomes more dense or more dispersed) is of course what makes them interesting, but it is also a

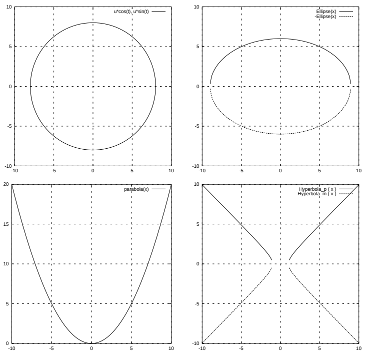

Figure 3: *Top to bottom, left to right: Circle, Ellipse, Parabola, Hyperbola.*

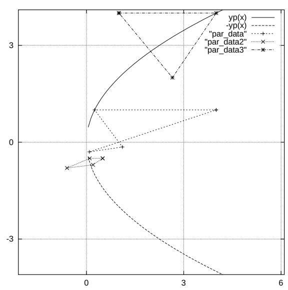

Figure 4: *Ray tracing at different points about a parabola.*

cause of some complications for us. Our reflected values no longer lie at evenly spaced FFT bin frequencies, preventing immediate use of the FFT^{-1}.

The simple solution to this problem is to perform individual inverse Discrete Fourier Transforms (DFT^{-1}) on the new frequencies, but this method is rather costly and it also has perceptual draw-backs, caused by the difficulty of calculating the phases on the non-bin frequencies. Current research in phase modeling [9] may however prove useful in further development in this area.

Therefore, we use an interpolation scheme to cope with the case of sparse reflections (based on polynomial splines and insertion of zero values) and a method to cope with the case of dense reflections (simply taking the maximum contending value, since this would tend to dominate). See figure 5 for a schematic overview of this interpolation post-processing.

Having recreated amplitudes at the original FFT bin frequencies, the inverse FFT is performed, using the original phases to ensure that phasing effects do not occur.

4.2. Scaling and Normalising

Scaling is performed, as per user GUI input, by linear interpolation between the original and the reflected points in Hz and dB. A scaling value of zero gives the original spectrum, and a scaling value of one renders the mirrored spectrum. Of course, nothing prevents the user of scaling more than one, or less than zero.

Amplitude normalisation is not performed, because of the difficulty of implementing normalisation over blocks without prior knowledge of the level of the forthcoming signal. Since the most interesting and desired effects are usually rather subtle, this is not

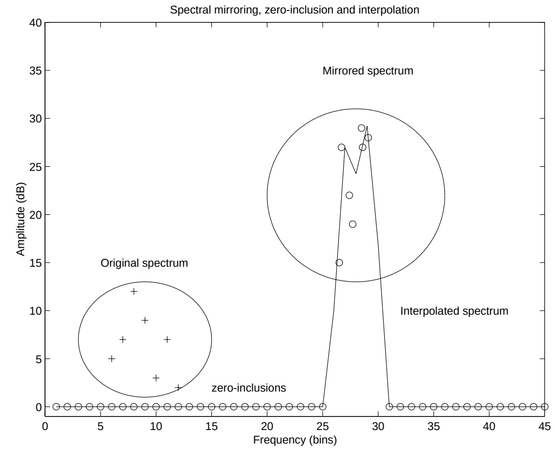

Figure 5: *Schematic overview of insertions and interpolation of bin-realignment.*

believed to be an important issue. On the other hand, if radical sound modifications are desired, then loudness drift generally goes unnoticed.

5. IMPLEMENTATION

5.1. Platform

An initial prototype system was developed in Common Lisp, dealing with mirroring at the note level. The system described here, is based on a translation of this system into C++.

The system is implemented using C++ and Java on Linux. However, the libraries and languages are chosen to ensure portability, and that the system can run in real time on many platforms. For this reason the only native library used is ALSA, the Advanced Linux Sound Architecture, chosen to ensure minimum system latency [10]. For other platforms this can be replaced by the cross platform audio API, PortAudio [11]. Apart from the user interface the mirroring is implemented in C++ and as such forms a separate module which is loaded in a general sound engine. The sound engine performs analysis/synthesis using FFT based overlap and add, interfacing with the sound hardware independently of the loaded modules.

The graphical user interface is implemented in Java using the Swing library and interfaces with the C++ module using the Java Native Interface (JNI).

5.2. Semantic

In short, the system accepts a signal from a microphone or digital source, and pipes it through the different modules (FFT, mirroring, interpolation, FFT^{-1}) sequentially. The result is forwarded for sound output. The user can change the values of the different mirroring and signal processing parameters with the user interface, and the changes are instantly propagated to the relevant modules, thus affecting the output in real time. These parameters include the block size, the conic section form and size, and the scaling coefficient.

5.3. User Interface

The user is presented with GUI controls and feedback as follows:

- for the parameters of the analysis,
  - slider adjusting resolution,
  - bar display showing latency,
- for determining the conic section,
  - toggle between global or sectional mirrors,
  - toggle between vertical projection and reflection,
  - toggle between flat and curved mirrors,
  - slider adjusting eccentricity,
  - slider adjusting the perpendicular distance of its center from the bounding peaks,
  - slider adjusting the lateral placement of its center,
  - labelled icon displaying type,
- slider adjusting the scaling coefficient.

Controls not applicable (*e.g.* lateral placement for parabolic mirrors) are grayed-out.

The latency of the system is controlled by the block size and by the amount of overlap. The block size is adjusted directly by the resolution slider; the overlap is set automatically by the system. Better frequency resolution corresponds to higher latency.

The type of curvature is determined by the eccentricity, e:

$$\begin{array}{rcl} e = 0 & \Rightarrow & \text{circle} \\ 0 \leq e < 1 & \Rightarrow & \text{ellipse} \\ e = 1 & \Rightarrow & \text{parabola} \\ e > 1 & \Rightarrow & \text{hyperbola.} \end{array}$$

(*Cf* figure 2.)

5.4. Sound Evaluation

Used without restraint, the effect of reflecting the spectrum in conic sections can be such that it distorts the original signal beyond recognition. Inclusion of a scaling parameter however, or using the scaled vertical projection mode, can render interesting subtle effects. These range from "purifying" the sound (*i.e.* removing roughness and background noises), to roughening it (basically performing the opposite, see figure 6). The more extreme effects do seem to have musical interest, and this is being explored further in artistic collaborations.

6. CONCLUSION

This paper has presented a novel sound effect method. The spectra of overlapped blocks are reflected on different types of conic section shaped mirrors. Continuity between blocks is guaranteed by retaining the original phases. The issue of realigning sparse and dense spectra is tackled.

The modification of the sound ranges from subtle changes in the signal to noise ratio, to creation of sounds far from the original. Decreasing the weaker parts of the spectrum creates a purer sound, whereas increasing them makes for a rougher sound. Spectral points may be moved far from their original frequency/amplitude positions giving pronounced effects, the extent of which can be scaled. Retaining the phase of the original spectral point frequencies prevents unwanted effects such as phasing.

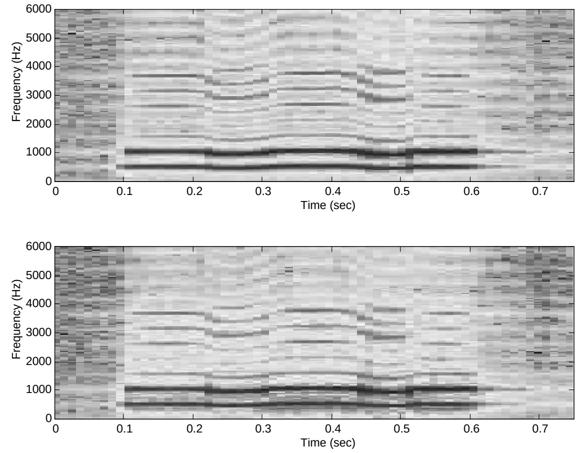

Figure 6: *Example of spectrograms of a soprano sample: original on top and roughened below.*

This technique, with its broad range of effects, is already in use in artistic collaboration, and sounds promising.

7. REFERENCES

[1] Maria Esteves, "Mirrors—Espejos", Performed in The Hague, The Netherlands, 25 May 2001 and 07 Jun 2001.

[2] J. B. Allen, "Short term spectral analysis, synthesis and modification by discrete Fourier transform". IEEE Trans. on Acoustics, Speech and Signal Processing, Vol. ASSP-25, No. 3, June 1977.

[3] E. Zwicker and H. Fastl, "Psychoacoustics: Facts and Models", Springer-Verlag, 2nd ed, 1999.

[4] Jensen, K., "Timbre Models of Musical Sounds", PhD. Dissertation, DIKU Report 99/7, 1999.

[5] Apollonius of Perga, "Conics, Books I–III", Green Lion Press, 1998.

[6] George Salmon, "A Treatise on Conic Sections", 6th ed. Chelsea, 1960.

[7] Xah Lee, "Conic Sections", `<http://xahlee.org/SpecialPlaneCurves_dir/ConicSections_dir/conicSections.html>`, 28 Jun 2001.

[8] John B. Conway, "Functions of One Complex Variable", Springer-Verlag, 1987.

[9] T. H. Andersen and K. Jensen, "Phase modeling of instrument sounds based on psycho acoustic experiments", Proc. Workshop on Current Research Directions in Computer Music, Barcelona, Spain, 2001.

[10] Karl MacMillan, Michael Droettboom and Ichiro Fujinaga, "Audio Latency Measurements of Desktop Operating Systems", Proc. ICMC-2001, Havana, Cuba.

[11] Phil Burk, "PortAudio – Portable Audio Library", `<http://www.portaudio.com>`, 30 May 2001.