COMBINING ZEROTH AND FIRST-ORDER ANALYSIS WITH LAGRANGE POLYNOMIALS TO REDUCE ARTEFACTS IN LIVE CONCATENATIVE GRANULATION

Dario Sanfilippo
Independent
Sicily, Italy
sanfilippo.dario@gmail.com

Julian D. Parker
Native Instruments
Berlin, Germany
julian.parker@cantab.net

ABSTRACT

This paper presents a technique addressing signal discontinuity and concatenation artefacts in real-time granular processing with rectangular windowing. By combining zero-crossing synchronicity, first-order derivative analysis, and Lagrange polynomials, we can generate streams of uncorrelated and non-overlapping sonic fragments with minimal low-order derivatives discontinuities. The resulting open-source algorithm, implemented in the Faust language, provides a versatile real-time software for dynamical looping, wavetable oscillation, and granulation with reduced artefacts due to rectangular windowing and no artefacts from overlap-add-to-one techniques commonly deployed in granular processing.

1. INTRODUCTION

1.1. Historical background

In 1946 and 1947, Dennis Gabor published two crucial articles that would lay the theories and foundations of granular processing [1, 2]. Gabor developed a conceptual framework for the analysis of audio signals linking quantum theory, Heisenberg's uncertainty principle, and Mach's analysis of sensation to formulate a key relationship between time and frequency domain in the sonic realm [3, 4, 5]. The formula from his publication from 1946,

$$\Delta t \Delta f \approx 1 \quad , \tag{1}$$

represents the interpretation of Heisenberg's uncertainty principle for acoustic signals. Δt and Δf are the uncertainties for the temporal and frequency locations of an oscillation, linked by multiplication and equality to a constant to express their inverse relationship, that is, that reducing the uncertainty in one quantity increases it in the other and vice versa.

Gabor describes the design for a mechanical machine capable of frequency compression and expansion in his publication from 1946. Years later, in the 1960s, Springer developed an analogue device, the Tempophon [6], based on the principles described by Gabor, which implemented a set of rotating reading tape heads for the modulation of, independently, tempo or pitch. A revolutionary aspect of the granular theories is that tempo and pitch are no longer linked and can be altered individually. The Tempophon worked with six reading heads arranged in a circle to rotate at different speeds and read portions of the tape, hence grains, at different pitches. The tape could change in speed too, and this allowed to change the tempo independently, provided that the speed of the reading heads relative to the movement of the tape remained constant.

1.2. Online and offline approaches to granular processing

Composer Horacio Vaggione is one of the primary proponents of a technique called *micromontage*, which he used in some of his works such as *Octuor* (1982), *Thema* (1985), and *Schall* (1995). The technique consists of extracting and rearranging sonic fragments from a sample, or in the generation of new fragments, that are specifically combined according to their characteristics at the microsound time scale to create complex higher-level sonic materials [7, 8].

In music and DSP, granular processing is a technique similar to micromontage for the automated generation and transformation of sounds through fragmentation and reorganisation. Signals are decomposed into sonic grains, each having independent features such as pitch, amplitude, and duration (typically below 0.05 seconds), and are then rearranged following different criteria to obtain a new stream with original global characteristics [9, 10].

In more modern times, granular processing criteria have evolved through different paths. One of the most relevant techniques today is concatenative sound synthesis (CSS), which is used for speech and music processing. The general working mechanism of CSS is based on large databases of signals that are divided into units, where specialised selection algorithms find the best-fitting fragments to be combined to achieve the desired target sound. Typically, the selection is performed on low-level and high-level descriptors extracted algorithmically or previously determined by the user. A thorough review of CSS techniques can be found in [11].

A discussion on other techniques that are related to granular processing and concatenative techniques can be found in [12]. These techniques include *granulation*, described as a granular synthesis technique similar to CSS based on the reorganisation of a recorded sound file [9, 10]. Alternatively, the technique called *brassage*, which consists of scrambling sounds [13]. Furthermore, the discussion in [12] focuses on adaptive concatenative sound synthesis (ACSS), also referred to as data-driven concatenative sound synthesis, which is often deployed for automated processes of sound concatenation for the generation of complex streams [14]. The technique presented in this paper can be considered a form of ACSS, albeit the information processing performed for the automatic selection is basic and low-level, that is, zero-crossing and first-order derivative analysis to reduce discontinuities at the junction points.

Non-real-time techniques for granular processing are typically based on fixed wavetables. In this case, sonic fragments can be read at different speeds using some interpolation scheme and at different positions depending on the array index pointer. For the live application of the algorithm presented here, we use circular buffers that can be updated with new data continuously [15]. Circular buffers are played back by indexing their content. If S is the size of the buffer in samples, $x[n]$ is an input, $w_{\text{idx}}[n]$ is the unit-increment writing pointer cycling from 0 to $S-1$, and $r_{\text{idx}}[n]$ is the reading pointer to access the table, then we can describe their behaviour as follows:

$$C_{\text{buff}}[n] = x[n - \text{mod}(w_{\text{idx}}[n] - \text{mod}(r_{\text{idx}}[n], S), S)] \quad . \tag{2}$$

We define

$$\text{mod}(n, d) = \text{fmod}(\text{fmod}(n, d) + d, d) \quad , \tag{3}$$

and

$$\text{fmod}(n, d) = n - \text{trunc}(n/d)d \quad , \tag{4}$$

which is the floating-point reminder. Fractional reading indexes are implemented via fifth-order Lagrange interpolation [16].

Specific configurations of digital granular processing can implement live pitch and time transpositions independently. The mechanical implementation based on Gabor's design produced a smooth transition between successive reading heads. A single buffer is adequate in the digital case to achieve streams without gaps as the reading head can move throughout the tape in virtually no time. However, the digital implementation requires a design that avoids discontinuities as they are likely to occur when transitioning between grains. The windowing of grains is a critical element of granular processing. Windowing means to ring-modulate two signals, so the window type and its spectral content play a significant role in the outcome. Windowing also results in a non-continuous stream as there are regions where the amplitude decreases drastically. Generally, the problem is resolved using two windowed reading heads that are out of phase by 180 degrees [17]. The overlapping grains then add up to a constant or quasi-constant amplitude envelope depending on the used windowing function.

The regularity or non-regularity of how pitch, rate, duration, and buffer position are selected for each grain defines important characteristics in granular processing. Randomness and stochasticity are often used to render asynchronous granular processing. The focus in the musical practice of the authors, in particular, is on having parameters such as pitch, rate, and grain position determined by recursivity and structural coupling between the sonic output and the parameters of the granulators themselves. Self-modulation through a nonlinear iterated configuration is the basis for chaotic behaviours [18] and the emergence of unpredictability, which are key for a wide range of electronic music works [19].

2. CONCATENATIVE GRANULAR PROCESSING

Our investigation on concatenative granular processing (CGP) focuses on avoiding the artefacts that may be introduced by the windowing of grains through non-rectangular windows and overlap-add-to-one (OATO) methods found in standard granular processing. Some of these artefacts are discussed in [20] and [21]. The pitch-synchronous technique (PSOLA) for granular synthesis is a solution to some of the artefacts in the overlap-add method, which allows to preserve the spectral characteristics of the input signal when, independently, varying either pitch or time. However, the PSOLA method is particularly well-suited for speech signals as it relies on the (quasi) periodicity of its input, which makes it less versatile [22]. Deploying trapezoidal windowing with short overlap portions may appear as an easy solution to the problem, although only rather large overlap lengths would effectively reduce the junction discontinuities. See 1, for example.

The authors' interest, thus, is to investigate algorithms that can operate with any signals for the generation of complex spectra by sequential combination of sonic fragments while preserving the characteristics of the individual grains. The results are deployed in the authors' music practice with self-modulating and adaptive audio feedback networks [23] to explore the potentials of the algorithm in the creative practice. The design presented here is not intended to replace well-established standard techniques for granular processing, although it offers an alternative mechanism through which novel sonorities and musical structures can be achieved.

Concatenative granular processing corresponds to rectangular-windowing granulation through non-overlapping grains. With this technique, the main issue concerns the signal discontinuities at the junction points due to the connection of uncorrelated sonic fragments. The spectrum of the splice artefact consists of the superposition of the spectrums of the discontinuities at all derivative levels up to N, including the zeroth discontinuity which is the basic continuity requirement. The N-th-order derivative discontinuity is integrated N times in the signal, and therefore has a spectral roll-off of $6N$dB per octave, which is also scaled by its magnitude. Hence, eliminating the zeroth-order discontinuity and minimising the first-order discontinuity has a significant attenuating effect on the overall artefact produced, for the higher-order discontinuities are heavily filtered.

2.1. Zeroth-order continuity

The first mechanism to guarantee zeroth-order continuity is zero-crossing synchronicity. The fundamental concept of zero-crossing synchronous granular processing is that grains start and end at a zero-crossing (ZC). For a fixed audio source on an array, we can scan the table and store all the ZC positions in a second array of as many elements as the ZC occurrences to be recalled later. Since signals and ZC occurrences can be irregular, the duration of each grain is variable and dependent on the signal itself. One essential condition to generate a sequence of grains of duration D without zeroth-order discontinuities is that each successive grain should be triggered after the time D has passed, at the first ZC occurrence. Hence, the output of the granulator must be inspected continuously to detect a ZC, and the information must be sent back to the section that triggers each grain. It is the minimum requirement for a stream without major amplitude discontinuities at the grains junction, although further actions must be taken to reduce artefacts.

Particularly, grains ending and starting at a ZC position result in a two-sample zero-order hold, as the amplitudes at ZC positions are likely to be close. When the first-order derivative at the end and start of consecutive grains is the same, then the successive grain start can be positioned one sample ahead to improve smoothness. However, when transitioning between grains with different first-order derivatives at the junction, the amplitudes at the ZC can change significantly. Thus, the position correction mechanism shifts the starting position of the next grain by a number of samples given by the ratio between the first-order derivatives at the junction. This improves zero-order continuity and preserves first-order polarity if combined with the mechanism discussed next.

2.2. First-order continuity

As discussed earlier, continuity in the first-order derivative is also essential to prevent artefacts caused by the concatenation of sonic fragments. A basic criterion to reduce first-order discontinuity is to select grains with matching first-order polarity at the junction. This, combined with the position correction mechanism discussed above, will guarantee that the first derivative sign is consistent while transitioning to the next grain. To effectively reduce the junction discontinuities, we deploy interpolation combined with zeroth and first-order analysis. Several interpolation schemes are available; we have chosen Lagrange polynomials for both efficiency and accuracy. Lagrange interpolation is based on an FIR structure and is thus stable for any time-variant case, which is a reason why it was preferred to Thiran interpolation. Furthermore, low-order Lagrange schemes performed adequately to reduce first-order discontinuities, while Sinc interpolation would have require a larger set of points. A crossfade at the junction combined with zeroth and first-order analysis was considered, although the design would have presented two issues: a discontinuity in the first derivative introduced by the crossfade, which would have been larger for smaller crossfade lengths, and a necessary lookahead delay equal to the size of the crossfade. Hence, we opted for the interpolation scheme.

Specifically, fifth-order Lagrange interpolation is used to reconstruct the starting segment of the next grain using three points from each side. We can see comparisons between different Lagrange orders below. Essentially, fifth-order Lagrange appeared to be a good compromise to demonstrate these results, but the interpolation order is a parameter that the user can change at will depending on their specific needs. For a reconstruction of M samples and an order N, the coefficients of the Lagrange interpolator for arbitrary spacing of the points can be written in closed form as follows:

$$h_\delta(m) = \prod_{\substack{k=0 \\ k\neq m}}^{N} \frac{\delta - \xi(k)}{\xi(m) - \xi(k)}, \quad m = 0, 1, 2, ..., N \quad . \tag{5}$$

For this application, N is assumed to be odd for symmetry and we have that:

$$\delta = 0, 1, 2, ..., M - 1 \quad , \tag{6}$$

$$\xi(m) = \begin{cases} m - (N+1)/2, & \text{if } m < (N+1)/2 \\ m - (N+1)/2 + M, & \text{otherwise} \end{cases} \quad , \tag{7}$$

where δ is a set of M interpolated positions, and ξ is a set of $N + 1$ elements indicating the spacing of the side points on the x-axis. Later, we will use (5) and (7) to define the output of the granulator.

2.3. Storing and recalling zero-crossing indexes

The ZC positions of the input signal are sampled from the signal $w_{idx}[n]$ when specific conditions are satisfied, and they are stored in two separate buffers to distinguish between positive and negative first-order polarity. The ZC positions, once recalled, must be exact; hence, non-interpolated buffers can be used to reduce CPU load. Since ZC positions are likely irregular, when a ZC is detected, the position will be sampled and held and written in the buffer continuously until a new ZC occurs. This way, the buffers containing ZC positions can be recalled using any reading index, and it is guaranteed that the output will be a ZC position, which can then be used to determine the starting position of a grain. In particular, the buffer position of the grain will be the last ZC occurrence before the desired grain position or the specified position itself if that corresponds to a ZC.

The condition for storing a position is that the input signal is at a ZC, while the polarity of the first derivative determines the corresponding buffer. The indicator functions for these cases can be implemented via logic operators checking if the product between current and past samples is negative, for the ZC, and whether the first difference, based on backwards Euler's method, is positive or negative.

2.4. Grains starting position processing

A new grain is triggered when two conditions are satisfied: at least D samples of the current grain must have been read, and the output of the granulator, $G_y[n]$, must be at a ZC. To read a grain, the determined starting position is sampled and held to offset a line that reads the samples in the buffer. The algorithm implements a trigger indicator that is used to trigger grains and synchronise parameters variations, including the grain positions. The trigger indicator consists of a unit-increment counter and a ZC indicator inspecting the output of the system. If the counter is above a value representing the grain size and the output is at a ZC, a new grain is triggered and the counter is reset.

The grain positions are recalled from the buffer containing ZC indexes with first derivative polarity corresponding to that of the current output grain. Once a grain is triggered, a correction mechanism adjusts the selected ZC position according to the ratio between the first derivatives of current and next grains at the junction: $G_{\text{ydiff}}[n]/G_{\text{xdiff}}[n]$. This preserves zeroth-order continuity and first-order polarity. We can call $G_{\text{p}}[n]$ the calculated starting position of each grain, and $G_{\text{l}}[n]$ the line function to index the buffer and read the samples.

$G_{\text{r}}[n] = G_{\text{p}}[n] + G_{\text{l}}[n]$ is then the granulator's reading head. Based on (2), the non-interpolated version of the system is given by:

$$G_{\text{y}}[n] = x[n - \text{mod}(w_{\text{idx}}[n] - \text{mod}(G_{\text{r}}[n], S), S)] \quad . \tag{8}$$

2.5. Junction interpolation

The Lagrange polynomials are deployed as a switching section between the interpolated and non-interpolated outputs. When $1_{\text{T}}[n] = 1$, the system outputs the interpolation for M samples, while it moves back to $G_{\text{y}}[n]$ afterwards. We can apply (5) and (7) and define the interpolation section, $L_{\text{y}}[n]$, as follows:

$$L_{\text{y}}[n] = \sum_{m=0}^{N} h_\delta(m) G_{\text{y}}[n + \xi(m) - L_{\text{count}}[n]] \tag{9}$$

$$L_{\text{count}}[n] = \begin{cases} 0, & \text{if } 1_{\text{T}}[n] = 1 \\ 1 + L_{\text{count}}[n-1], & \text{otherwise} \end{cases} \quad . \tag{10}$$

Finally, the switching section, $G_{\text{yL}}[n]$, results in an M-sample reconstruction of the starting segment at the junction for a smooth

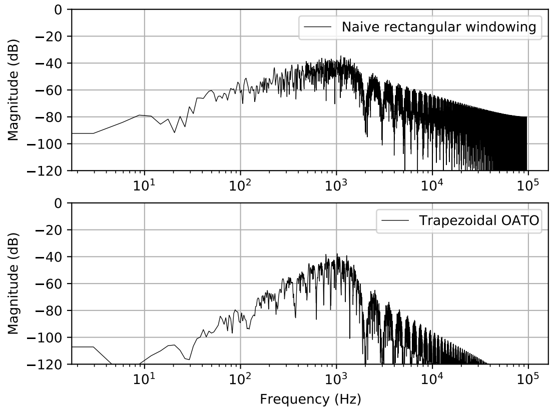

Figure 1: *Reconstruction of a 999 Hz sinusoid input with uniformely distributed random positions with 192-sample grains at 192 kHz sample rate. Comparison between spectra of naive non-overlapping rectangular windowing (top) and 32-sample overlap with trapezoidal windowing (bottom).*

transition between grains. The output of the system can then be written as follows:

$$G_{\mathrm{yL}}[n] = \begin{cases} L_{\mathrm{y}}[n], & \text{if } L_{\mathrm{count}}[n] < M \\ G_{\mathrm{y}}[n], & \text{otherwise} \end{cases} . \quad (11)$$

2.6. Pitch and time transpositions

Pitch and time transpositions can be implemented easily within this design. Pitch factors can be varied by merely multiplying $G_{\mathrm{I}}[n]$. However, negative pitch factors require some changes. Particularly, since negative pitch factors play grains in reverse, the positive and negative first-order polarity indicators should be inverted for negative pitch values. Furthermore, considering that the grains selected for negative pitches crossed zero in the opposite direction, one sample should be substracted to correctly align the junction, which is then corrected by the first-order ratio. Note that in the case of negative pitches, the first-order ratio is negative, moving the position towards the correct direction with regard to the reverse playback.

Similarly, time transposition can be implemented by modulating the desired grain position through a line function. A line function can be implemented as an integrated increment, where the incremental value corresponds to the transposition factor.

Positions can also be determined recursively using the output of the system. Amplitude values, which will correspond to ZC occurrences in the system's output, can be magnified arbitrarily as they will be wrapped around and fed back to determine the position of the next grain. The iterative and nonlinear characteristics of the process are the basic requirements for the emergence of chaotic and dynamical behaviours [24], which is often a sought-after quality in contemporary electronic music [25]. Furthermore, the signal in the feedback path can be low-passed, and the cut-off of the filter can be used as a parameter to smoothly transition between synchronous and asynchronous behaviours in the grain positions.

Pitch factors can be determined recursively according to the magnitude of the first-order ratio at the junction. At the beginning of the next grain, the first-order derivative is forced to match the first-order derivative at the ending grain, resulting in an improved continuity and smoother transition while self-modulation generates chaotic behaviours.

3. RESULTS

Discontinuities detection techniques for concatenation speech signals can be found in [26]. For the evaluation of the results in the musical domain, we provide spectral analysis of the algorithm and compare it with naive non-overlapping rectangular windowing and OATO with Hanning, triangular, and trapezoidal windowing. To quantify the suppression of the grain junction artefact achieved by applying our smoothing method, we evaluate the technique with a sinusoidal input. This allows the distortion introduced by the algorithm to be clearly visible in the frequency domain. Moreover, a set of audio examples are provided to compare audio qualities between zeroth and first-order analysis with and without fifth-order Lagrange interpolation. The audio files are stereo, where the right channel corresponds to the algorithm with derivative analysis and interpolation, and the left one corresponds to the algorithm with analysis only, so that panning can easily show the differences. Audio examples will be based on sinusoidal inputs and complex musical signals, particularly from the piece Lunacy by the band Swans. The examples can be found at the following URL:
`https://soundcloud.com/dario-sanfilippo/sets/dafx21`; they demonstrate the basic features of the granulator, such as pitch and time transposition and reverse effects, as well as asynchronous behaviours based on recursive mechanisms, click-free looping, and wavetable oscillation as zero-factor time stretching. A creative practice example where the granulator is used extensively can be found in one of the authors' pieces, *Constructing Realities* (2019-2020), that can be listened to at the following URL: `https://soundcloud.com/dario-sanfilippo/constructing-realities-2019-2020`.

In Figure 1, we can see the spectrum generated with 1000 grains per second without pitch transposition and positions randomly selected from the entire buffer, using a 999 Hz sinusoid as an input signal. The spectra of the naive non-overlapping rectangular windowing and trapezoidal OATO with 32-sample overlap are compared. In Figure 2, we have the spectra of the OATO technique with Hanning and triangular windowing. In all of these cases, we can see that the input signal is heavily masked by broad noise, showing that OATO mostly attenuates high-frequency components.

In Figure 3, we can see the spectrum generated with 1000 grains per second without pitch transposition and positions randomly selected from the entire buffer, using a 999 Hz sinusoid as an input signal. The spectra of the non-interpolated analysis-only design and the algorithm deploying fifth-order junction interpolation can be compared. The frequency of the sinusoid and the grain rate are non-even ratio to maximise the artefacts and the consequent noise introduced by the process, as the algorithm must actively operate to achieve zeroth and first-order continuity. We can see that the input signal is clearly visible and that the noise resulting from the concatenation is heavily filtered from about 5 kHz onwards, showing an overall improvement in the resulting signal-

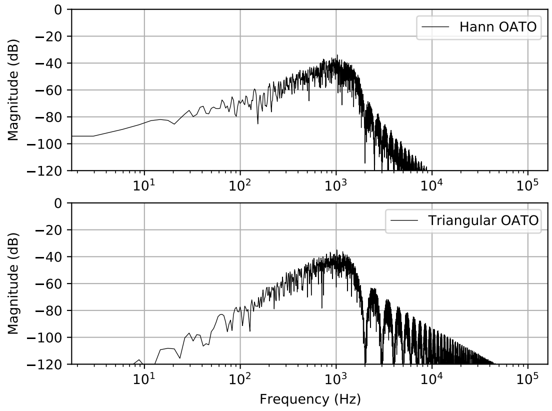

Figure 2: *Reconstruction of a 999 Hz sinusoid input with uniformely distributed random positions with 192-sample grains at 192 kHz sample rate. Comparison between spectra of the OATO technique with Hanning windowing (top) and OATO with triangular windowing (bottom).*

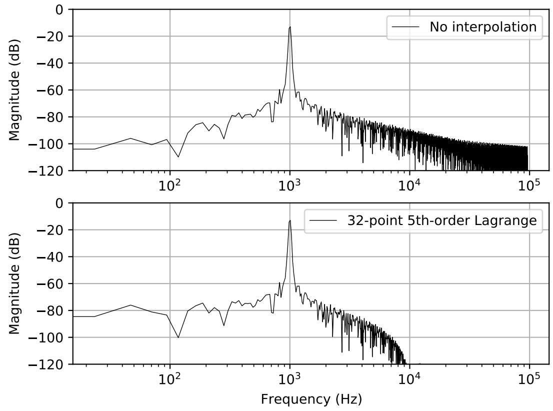

Figure 3: *Reconstruction of a 999 Hz sinusoid input with uniformely distributed random positions with 192-sample grains at 192 kHz sample rate. Comparison between spectra of the algorithm without interpolation (top) and 32-point interpolation with fifth-order Lagrange polynomials (bottom).*

to-noise ratio. The mean square error between input sinusoid and granular output is 0.01100429. Finally, Figure 4 shows a comparison between linear interpolation and third-order Lagrange interpolation.

A disadvantage of this design for granular processing is that the actual length of each grain depends on the output of the granulator, as the current grain must be at a ZC before being able to trigger the next one. Especially with pitch transpositions that decrease the signal frequency, the actual length of each grain may differ substantially from the desired one. This is mostly negligible for high grain rates or for high-frequency input signals. On the other hand, a less homogeneous grain size distributes the noise artefacts over a broader spectral region, which is perceptually more desirable.

In Figure 5, we have samples generated with 1000 grains per second with randomly selected pitch and position for each grain, using a 999 Hz sinusoid as an input signal. It is possible to see that, while transitioning between grains with different derivatives, the correction mechanism is appropriately adjusting the starting position of the successive grains to preserve zeroth-order continuity and first-order polarity. Furthermore, we can see that the interpolation mechanism at the junction generates a smooth transition between grains, preserving first-order continuity.

Figure 6 shows the granulation of a complex signal with reversed playback and random positions for each grain at a rate of 1000 grains per second.

In Figure 7, we have an example of the granulator operating as a wavetable oscillator where a short fragment from Swans' *Lunacy* is repeated. We can see a smooth transition between the fragments as well as an exact repetition of the waveform.

The Faust program is available on Github: `https://github.com/dariosanfilippo/concatenative_granulation`.

4. CONCLUSION

We have discussed a technique for generating granular streams based on zeroth and first-order analysis combined with Lagrange polynomials to reduce concatenation artefacts. This technique allows for sequences of non-overlapping sonic fragments without significant signal discontinuities while avoiding some of the artefacts common in standard granular processing techniques due to overlapping. Furthermore, this design allows for multi-modality as the system, depending on the parameters settings, can operate as a wavetable oscillator with complex waveforms, a click-free looper, and a granulator having standard features such as pitch and time transposition, reverse effects, as well as asynchronous behaviours.

Future developments for this algorithm will guarantee higher-order continuity by deploying more advanced grain selection adaptive algorithms for transitions among most compatible grains, i.e., grains with best-matching derivatives at the junction.

The live nature of the system implies that discontinuities may result from the buffer being updated with new data from the input signal. While the software allows freezing the buffer preventing new data from being written, the authors plan on implementing guarding mechanisms to guarantee that the granulator's reading head does not overlap with the writing one, hence making sure that grains will be reproduced in their entirety correctly.

Lastly, adaptive algorithms will be developed to determine the best-fitting interpolation length and maximise the effectiveness of Lagrange polynomials based on the characteristics of each grain.

5. ACKNOWLEDGMENTS

We want to thank Stéphane Letz for his help with Faust, and Julius Smith and Diemo Schwarz for their precious insights.

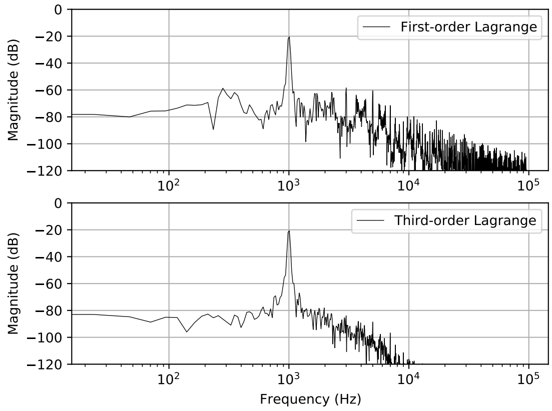

Figure 4: *Reconstruction of a 999 Hz sinusoid input with uniformely distributed random positions with 192-sample grains at 192 kHz sample rate. Comparison between first-order (top) and third-order 32-point interpolation with Lagrange polynomials (bottom).*

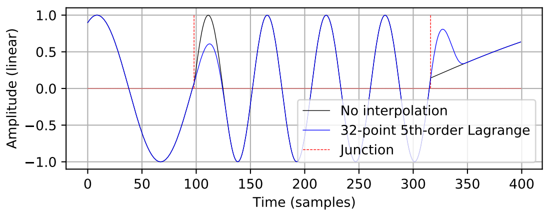

Figure 5: *Granulation of a 999 Hz sinusoid input with uniformely distributed random positions and pitches with 192-sample grains at 192 kHz sample rate. Comparison between non-interpolated and interpolated transitions.*

6. REFERENCES

[1] Dennis Gabor, "Theory of communication. part 1: The analysis of information," *Journal of the Institution of Electrical Engineers-Part III: Radio and Communication Engineering*, vol. 93, no. 26, pp. 429–441, 1946.

[2] Dennis Gabor, "Acoustical quanta and the theory of hearing," *Nature*, vol. 159.4044, pp. 591–594, 1947.

[3] Agostino Di Scipio, "I quanta acustici di Gabor nelle tecnologie del suono e della musica," *Musica/Tecnologia*, vol. 10, pp. 17–42, 2016.

[4] Werner Heisenberg, "Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik," in *Original Scientific Papers Wissenschaftliche Originalarbeiten*, pp. 478–504. Springer, 1985.

[5] Ernst Mach, *The analysis of sensations, and the relation of the physical to the psychical*, Open Court Publishing Company, 1914.

[6] Pierre Dutillieux et al., "Time-segment processing," *Digital Audio Effects. Udo Zölzer. 2a ed. John Wiley & Sons, Ltd. Cap*, vol. 6, pp. 185–217, 2011.

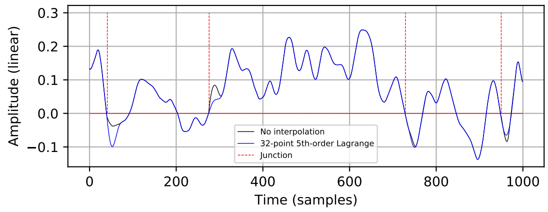

Figure 6: *Randomly selected 192-sample grains played in reverse from Swans' Lunacy.*

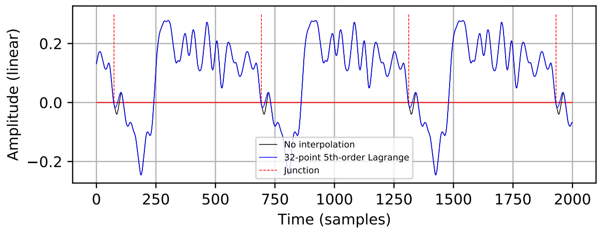

Figure 7: *Performing a zero-factor time transposition on Swans' Lunacy for complex wavetable oscillation.*

[7] Curtis Roads, *Microsound*, MIT press, 2004.

[8] Curtis Roads, "The art of articulation: The electroacoustic music of horacio vaggione," *Contemporary Music Review*, vol. 24, no. 4-5, pp. 295–309, 2005.

[9] Curtis Roads, "Introduction to granular synthesis," *Computer Music Journal*, vol. 12, no. 2, pp. 11–13, 1988.

[10] Barry Truax, "Real-time granular synthesis with a digital signal processor," *Computer Music Journal*, vol. 12, no. 2, pp. 14–26, 1988.

[11] Diemo Schwarz, "Concatenative sound synthesis: The early years," *Journal of New Music Research*, vol. 35, no. 1, pp. 3–22, 2006.

[12] Bob L Sturm, "Adaptive concatenative sound synthesis and its application to micromontage composition," *Computer Music Journal*, vol. 30, no. 4, pp. 46–66, 2006.

[13] Trevor Wishart and Simon Emmerson, *On sonic art*, vol. 12, Psychology Press, 1996.

[14] Diemo Schwarz, *Data-driven concatenative sound synthesis*, Ph.D. thesis, University of Paris 6, 2004.

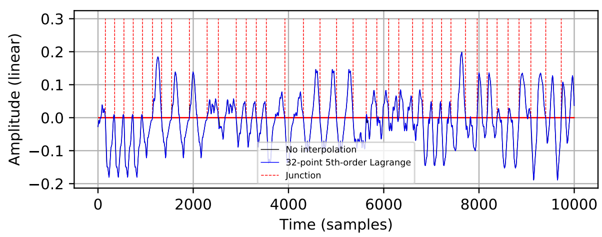

Figure 8: *Performing a 0.1-factor time transposition on Swans' Lunacy.*

[15] Davide Rocchesso, *Introduction to sound processing*, Mondo estremo, 2003.

[16] Francis Begnaud Hildebrand, *Introduction to numerical analysis*, Courier Corporation, 1987.

[17] Miller Puckette, *The theory and technique of electronic music*, World Scientific Publishing Company, 2007.

[18] James Gleick, *Chaos: Making a new science*, Open Road Media, 2011.

[19] Dario Sanfilippo, "Complex adaptation in audio feedback networks for the synthesis of music and sounds," *Computer Music Journal (forthcoming)*, 2021.

[20] Sergio Cavaliere and Aldo Piccialli, "Granular synthesis of musical signals," *Musical Signal Processing*, pp. 155–186, 1997.

[21] Damián Keller and Chris Rolfe, "The corner effect," in *Proceedings of the XII Colloquium on Musical Informatics*, 1998.

[22] Norbert Schnell, Geoffroy Peeters, Serge Lemouton, Phillippe Manoury, and Xavier Rodet, "Synthesizing a choir in real-time using pitch synchronous overlap add (psola).," in *Proceedings of the International Computer Music Conference*, 2000.

[23] Dario Sanfilippo, "Time-variant infrastructures and dynamical adaptivity for higher degrees of complexity in autonomous music feedback systems: the Order from Noise (2017) project," *Musica/Tecnologia*, vol. 12, no. 1, pp. 119–129, 2018.

[24] Stephen H Kellert, *Borrowed knowledge: Chaos theory and the challenge of learning across disciplines*, University of Chicago Press, 2009.

[25] Dario Sanfilippo and Andrea Valle, "Feedback systems: An analytical framework," *Computer Music Journal*, vol. 37, no. 2, pp. 12–27, 2013.

[26] Yannis Pantazis, Yannis Stylianou, and Esther Klabbers, "Discontinuity detection in concatenated speech synthesis based on nonlinear speech analysis," in *Ninth European Conference on Speech Communication and Technology*, 2005.