

Lyapunov Space of Coupled FM Oscillators

Claude Heiland-Allen
claude@mathr.co.uk

Abstract

I consider two coupled oscillators, each modulating the other's frequency. This system is governed by four parameters: the base frequency and modulation index for each oscillator. For some parameter values the system becomes unstable. I use the Lyapunov exponent to measure the instability. I generate images of the parameter space, implementing the number crunching on graphics hardware using OpenGL. I link the mouse position over the displayed image to realtime audio output, creating an audio-visual browser for the 4D parameter space.

Keywords

chaos, DSP, GPU

1 Introduction

In Soft Rock EP [ClaudiusMaximus, 2005] and Soft Rock DVD [ClaudiusMaximus, 2006] I explored the transitions between order and chaos in coupled FM oscillators, implemented in Pure-data using GridFlow for visualization of the output waveforms over time. More recently I developed the idea to map the parameter space on a perceptually relevant level and in performance choose parameters on the basis of desired sound character.

Seeing a bifurcation diagram produced by an analogue Moog synthesizer [Slater, 1998] and images of Lyapunov fractals [Dewdney, 1991], I decided to apply the latter technique to coupled FM oscillators in the digital realm.



Figure 1: Example output.

I had originally experimented with one Pure-data batch mode instance per CPU core each sending analysis data to a realtime Pure-data instance. The analysis used various methods (including FFT for spectral statistics and the sigmund external for pitch tracking) to classify points into pitched (ordered, stable) or unpitched (chaotic, unstable) with measures of distortion or noisiness. Sadly this approach proved impractical as it achieved only tens of pixels per second, even with a fast multi-core CPU, and porting these signal analysis algorithms to massively-parallel programmable graphics hardware seemed to be too difficult.

2 Formulation

2.1 Coupled FM Oscillators

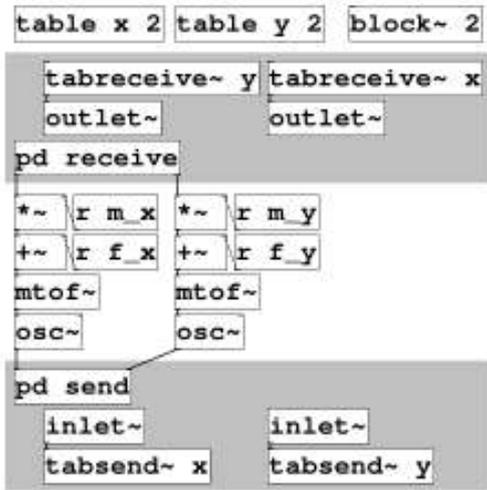


Figure 2: Coupled FM oscillators in Pure-data.

Consider the two coupled oscillators in Figure 2. I formulate this into a mutual recurrence relation:

$$\begin{aligned} x_{n+1} &= \%_0(x_n + I(f_x + m_x \cos(2\pi y_{n-d}))) \\ y_{n+1} &= \%_0(y_n + I(f_y + m_y \cos(2\pi x_{n-d}))) \end{aligned} \quad (1)$$

where

$$\%(t) = t - [t], \quad I(t) = \frac{440}{\text{SR}} 2^{\frac{t-69}{12}}$$

Here x_n, y_n is the phase of each oscillator at time step n , d is a delay measured in samples, f_x, f_y is the base frequency of each oscillator as a MIDI note number, and m_x, m_y is the modulation index of each oscillator as a MIDI note number.

I'll write the four-dimensional parameter space vector

$$a = (f_x, f_y, m_x, m_y)$$

and the $(2d+2)$ -dimensional phase space vector

$$z = (x_n, y_n, x_{n-1}, y_{n-1}, \dots, x_{n-d}, y_{n-d})$$

with sample rate $\text{SR} = 48000$. I'll fix $d = 1$ for reasons explained in Section 5.1.

2.2 Lyapunov Exponents

A good introduction is found in Chapter 4.3 *Lyapunov Exponent* [Elert, 2007]. The definition is covered in Chapter 13.7 *Liapounov exponents and entropies* [Falconer, 2003] which also relates it to measures of fractal dimension.

The Lyapunov exponent λ measures divergence in phase space:

$$|z_1(t) - z_0(t)| \approx e^{\lambda t} |z_1(0) - z_0(0)|$$

$$\lambda = \lim_{\substack{t \rightarrow \infty \\ z_1(0) \rightarrow z_0(0)}} \frac{1}{t} \log \frac{|z_1(t) - z_0(t)|}{|z_1(0) - z_0(0)|} \quad (2)$$

An attracting orbit has $\lambda < 0$ and a divergent (chaotic) orbit has $\lambda > 0$.

The wrapping of phase into $[0, 1)$ requires a modified norm:

$$|z|_{\%_0} = \sqrt{\sum_i \min(\%(z_i), 1 - \%(z_i))^2}$$

2.3 Viewing Planes

An image is 2D, which requires choosing a subset of the 4D parameter space to visualize. I chose two particular planes:

$$\begin{aligned} A_+(a_0, r_0) &= a_0 + r_0 \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \\ A_-(a_0, r_0) &= a_0 + r_0 \begin{pmatrix} 1 & 0 \\ -1 & 0 \\ 0 & 1 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \end{aligned} \quad (3)$$

where (u, v) is the coordinates of the pixel, a_0 is the centre of the view, and r_0 is the radius of the view.

3 Results

3.1 Examples

Figure 3(a) shows the initial view on starting the interactive browser. Low frequencies to the left are stable even at high modulation index away from the central axis. High frequencies to the right become chaotic at progressively lower modulation index. (b) shows the A_- plane at the same location. (c) shows bands alternating between stability and chaos. The bands become distorted and collapse as the modulation index and frequency increase. (d) shows its A_-

plane, bands become rings. When the frequency is greatly increased, the shapes become more intricate. (e) exhibits spirals of stability, with similar spirals in the A_- plane in (f).

When $f_x = f_y$ and $m_x = m_y$ the A_+ plane has mirror symmetry about its horizontal axis, and the A_- plane has two-fold rotational symmetry about its centre. Breaking the symmetry and setting $f_x \neq f_y$ or $m_x \neq m_y$ leads to diverse forms, shown in Figure 4. In particular Figure 4(c) has shapes that resemble those of Lyapunov space images of the logistic map.

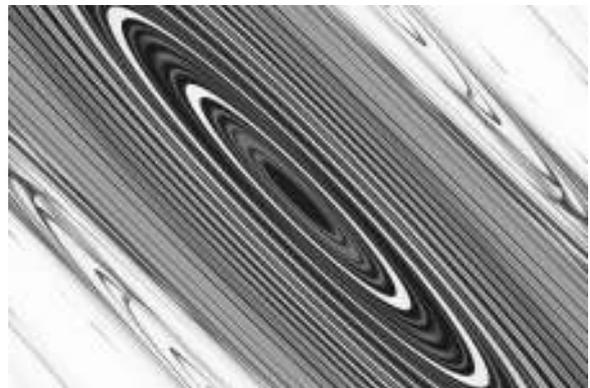
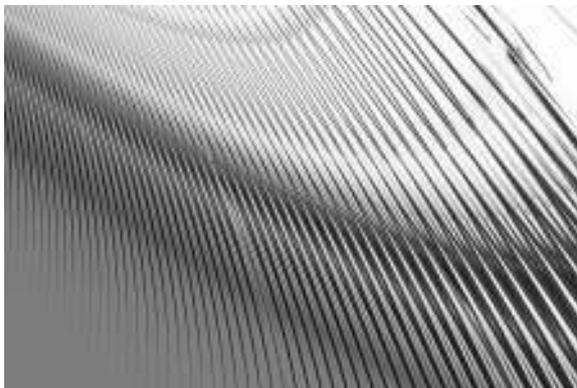
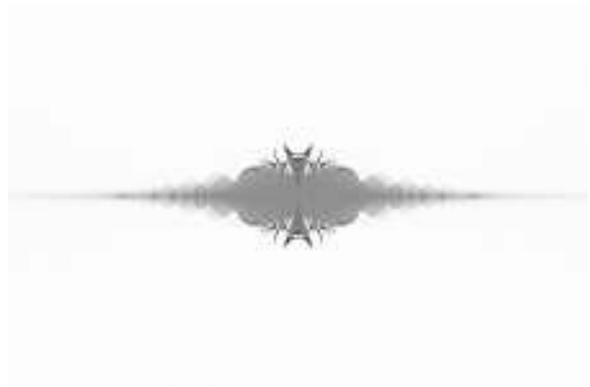
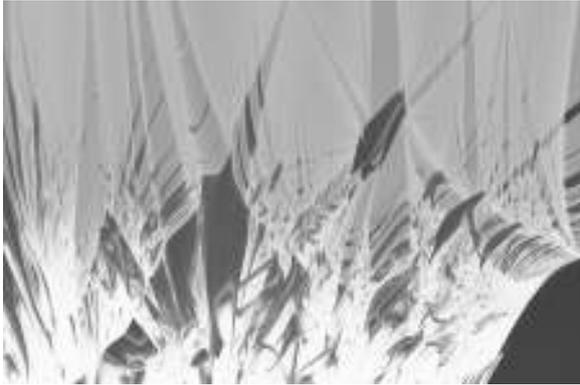


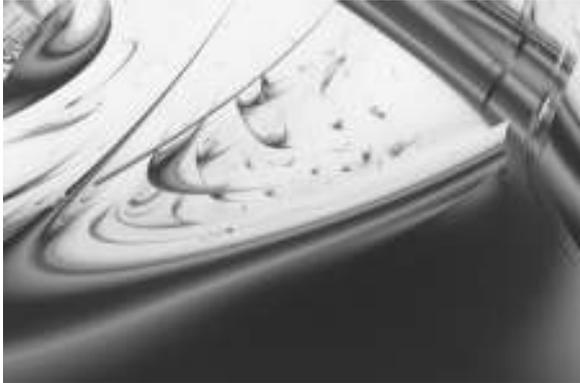
Figure 3: Example images. Darker shades are stable, lighter shades chaotic.



(a) $A_-((117.0, 148.4, 20.4, 2.7), 1.8)$



(b) $A_-((141.46, 146.22, 22.76, 0.27), 0.14)$



(c) $A_+((103.65, 108.41, 33.42, 10.93), 0.14)$



(d) $A_-((89.8, 137.5, -17.5, -7.1), 3.7)$

Figure 4: More examples.

3.2 Interactive Explorer

I implemented an interactive audio-visual explorer for the parameter space of coupled FM oscillators. Clicking with the mouse zooms the view about the clicked point. The left button (or scroll up) zooms in, the right button (or scroll down) zooms out, the middle button centers the view on the target point. Pressing the TAB key toggles between the A_+ and A_- planes in Equation 3, and F11 toggles full screen operation.

While the GPU simulates and analyses one oscillator pair per pixel, the CPU simulates one oscillator pair with a determined from the pixel under the mouse pointer. The image acts as a map, a reference frame for choosing parameters to audition by moving the mouse.

4 Implementation

I used OpenGL [Segal, 2013] and OpenGL Shading Language [Kessenich, 2013], with GLUT [Kilgard, 1996] for windowing and input event handling, and JACK [Davis, 2013] for audio output.

4.1 Introduction to Modern OpenGL

Modern OpenGL has a programmable shader pipeline. Vertex attributes are read from vertex buffers and processed by vertex shaders. The outputs of the vertex shader (called varyings) are further manipulated by an optional geometry shader stage. Geometry shaders can output a different vertex count to their input count, whereas vertex shaders are one-in one-out. The result of the geometry shader can be captured into another vertex buffer using transform feedback. Following the geometry shader the primitives (points or triangles) are rasterized, and varyings interpolated across each primitive. Finally a fragment shader takes these values and computes the colour at that pixel. The output of a fragment shader can be captured by attaching a texture to a framebuffer.

4.2 Computation Overview

To render an image I first fill a texture with (u, v) coordinates using a framebuffer object and a fragment shader. I copy this texture to a vertex buffer object, interleaving an initial phase space vector $z = (0, 0, 0, 0)$ and Lyapunov exponent statistics vector $l = (0, 0, 0, 0)$.

Using a vertex shader, I calculate a from (u, v) using Equation 3, and then compute a rough estimate of the Lyapunov exponent using Equation 2 by perturbing $z_1(0) = z_0(0) + \delta$

with δ small and performing $t = 256$ iterations of Equation 1. I discard the first few repetitions and those resulting in $-\infty$, and accumulate the rough λ estimates in l .

Between each repetition I compact the working set using a geometry shader, plotting points whose mean Lyapunov exponent estimate changed very little during the previous step. I keep the others to refine further, directing the computational effort on the points that need it most: those slow to converge.

To ensure user interface responsiveness, the computation is amortized over several frames. I divide the target frame period by the measured time for one repetition to compute how many repetitions to perform that frame. The repetitions-per-frame increases as the working set becomes smaller.

4.3 Noise Increases Stability

At the end of each repetition I keep z_1 instead of z_0 . This effectively adds a small amount of noise, counter-intuitively increasing stability. Noise allows more of the phase space to be explored, and makes it more likely for the perturbed orbit to reach an attracting part of the phase space.

4.4 Dither Increases Quality

To reduce grid sampling artifacts, I perturb (u, v) within the bounds of its corresponding pixel before calculating the a parameter vector for each repetition.

5 Conclusions

5.1 OpenGL Issues

The current implementation is hardcoded with delay $d = 1$ and would be very awkward to generalize. OpenGL architecture limits each vertex attribute to four components with the maximum number of attributes typically limited to sixteen. This totals 64 floats per vertex, 6 of which are needed for the pixel coordinates and Lyapunov exponent statistics accumulation. Therefore using OpenGL imposes a limit $d < 28$. For comparison my original experiments in *Soft Rock EP* used Pure-data's default block size of 64, with $d = 32$. Moreover, increasing d increases video memory consumption. With the maximum $d = 27$, browsing at 1920×1080 resolution would require over 1GB.

My future work on this project will look into using OpenCL, which provides a heterogeneous CPU and GPU computation framework. I hope

it will avoid the inherent awkwardness of abusing OpenGL shaders to perform calculations.

5.2 Audio Issues

While the implementation works as intended, with $d = 1$ the sound is nowhere near as rich and varied as with $d = 32$. With small d there is much more very high frequency content in interesting-looking regions. I've not found any regions of the parameter space with both interesting appearance and palatable audio frequencies at $d = 1$, while with high d there are parameters that generate sounds that fluctuate intermittently between smooth tones and noise. I haven't been able to visualize with high d to know whether their neighbourhoods look as interesting as they sound.

Unfortunately, heavy use of the GPU in the interactive browser can block the operating system for too long and cause audible glitches (JACK xrums). This situation may change as free drivers continue to improve, allowing use of the browser in a live situation.

5.3 Pretty Pictures

Despite these shortcomings, I think the images look good. I plan to render a selection at high resolution and print postcards and posters. For huge images I divide the image plane into tiles and compute each tile in succession, finally combining the pieces into one large picture.

There is also scope for video work, moving and rotating the viewing plane through the 4D parameter space, with different shapes forming and collapsing over time. My rough benchmarks take 5-10 seconds per frame at 1920×1080 , so for now I'll wait until faster cheaper graphics cards become available.

6 Obtaining the Implementation

I wrote the implementation on GNU/Linux Debian Wheezy using proprietary drivers running on a quad-core AMD64 processor with NVIDIA GTX 550Ti graphics card.

The source code for the implementation is available at: <https://gitorious.org/maximus/lyapunov-fm>

7 Acknowledgements

Thanks to Rob Canning, Adnan Hadzi and Joanne Seale for their helpful feedback.

References

- ClaudiusMaximus. 2005. *Soft Rock EP*.
[http://archive.org/details/
ClaudiusMaximus-SoftRockEP](http://archive.org/details/ClaudiusMaximus-SoftRockEP).
- ClaudiusMaximus. 2006. *Soft Rock DVD*.
[http://archive.org/details/
ClaudiusMaximus-SoftRockDVD](http://archive.org/details/ClaudiusMaximus-SoftRockDVD).
- Paul Davis. 2013. *The JACK Audio Connection Kit*. <http://jackaudio.org>.
- A. K. Dewdney. 1991. Mathematical Recreations: Leaping into Lyapunov Space. *Scientific American*, 265:178–180.
- Glenn Elert. 2007. *The Chaos Hypertextbook*.
<http://hypertextbook.com/chaos/>.
- Kenneth Falconer. 2003. *Fractal Geometry: Mathematical Foundations and Applications, Second Edition*. Wiley.
- John Kessenich. 2013. *The OpenGL Shading Language*.
[http://www.opengl.org/registry/doc/
GLSLangSpec.4.30.8.pdf](http://www.opengl.org/registry/doc/GLSLangSpec.4.30.8.pdf).
- Mark J. Kilgard. 1996. *The OpenGL Utility Toolkit (GLUT) Programming Interface*.
[http://www.opengl.org/documentation/
specs/glut/glut-3.spec.pdf](http://www.opengl.org/documentation/specs/glut/glut-3.spec.pdf).
- Mark Segal. 2013. *The OpenGL Graphics System: A Specification*.
[http://www.opengl.org/registry/doc/
glspec43.core.20130214.pdf](http://www.opengl.org/registry/doc/glspec43.core.20130214.pdf).
- Dan Slater. 1998. Chaotic Sound Synthesis. *Computer Music Journal*, 22(2):12–19.