UNIVERSITY OF CALIFORNIA
RIVERSIDE

The Mathematical Life and Death Sequence of the Universe: Generating Audiovisual Art From the Mandelbrot Set

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Music

by

Charles A. Sykes

June 2023

Dissertation Committee:

Ian Dicke, Chairperson

Dana Kaufman

Bradly Butterworth

The Dissertation of Charles A. Sykes is approved:

Committee Chairperson

University of California, Riverside

Dedication

To my family, friends, mentors, and colleagues, without whom this work would not have been possible.

ABSTRACT OF THE DISSERTATION

The Mathematical Life and Death Cycle of the Universe

by

Charles A. Sykes

Doctor of Philosophy, Graduate Program in Music
University of California, Riverside, June 2023
Dr. Ian J. Dicke, Chairperson

The Mandelbrot set, a fractal well known in the world of mathematics, exhibits many interesting properties for music composition. In fact, when read as a visual score from top-left to bottom-right, this fractal not only creates form through its horizontal symmetry, but also has clear areas of tension and release ideal for creating exciting music. As the Mandelbrot set is visually beautiful, this project focuses on creating audiovisual works of art using data sets from the main set, zooms of the set, and both mild and extreme functional alterations. While basic sonifications of the set have been explored by other artists before, this project seeks to create a multi-movement work where each sonic and visual element is controlled by or directly relates to the Mandelbrot set. This raises a number of interesting questions, most importantly: How can one translate math into art in a way that there is a clear, meaningful connection between the art and the source? While answers may vary, the most effective answer for this project

was to focus on recreating my personal musical language in the form of a mathematical formula, through which data points from the Mandelbrot set are processed.

Table of Contents

Signature Page: -- iii

Dedication: --- iv

Abstract: --- v

List of Figures: -- viii

Introduction: --- 1

Chapter 1: Historical and Contemporary Contexts: -- 6

Chapter 2: The Mandelbrot Set: -- 14

Chapter 3: Musical Algorithms and Audio Synthesis: ------------------------------------- 17

Chapter 4: Audiovisual Art for Audiovisual Math: -- 29

Chapter 5: Analysis of Audiovisual Project Results: -- 34

Conclusion: --- 55

Bibliography: -- 58

List of Figures

fig. 1: The Mandelbrot set: -- 3

fig. 2: Sketch of glissandi from *Metastaseis*: --- 9

fig. 3: The logistic map with relation to the Mandelbrot set: --------------------------------- 9

fig. 4: A zoom of the original set: --- 18

fig. 5: A minor functional alteration of above zoom: -- 18

fig. 6: Division-altered ETA results: --- 19

fig. 7: Zoom of above featuring interlocking rings: --- 19

fig. 8: A Mandelbrot image with no values within range: ---------------------------------- 19

fig. 9: Non-fractal, heavily modified ETA: -- 20

fig. 10: Potentially fractal, heavily modified ETA: -- 20

fig. 11: Mandelbrot waveform: --- 24

fig. 12: Screenshot of lightshow from *Hours into Minutes*: -------------------------------- 42

fig. 13: Piano melody from *Transcendence*: --- 43

fig. 14: Score excerpt from Ornstein's *Piano Sonata No. 4*.: ------------------------------- 44

fig. 15: Spirits escaping the Mandelbrot set: -- 45

fig. 16: Image of the protagonist: -- 46

fig. 17: Image of the demon: -- 48

fig. 18: Image of the protagonist's family: --- 49

fig. 19: Image of animal guardians: --- 50

fig. 20: Black clover image: --- 51

fig. 21: Spectrograph from *Black Clover*: -- 54

fig. 22: Spectra of the Mandelbrot waveform: -------------------------------------- 54

Introduction

All structured music follows some form of procedure that dictates how musical elements are treated, from the chosen pitch collection, to functional or nonfunctional intervallic relationships, to the larger overarching form. Algorithmic music trends towards an extreme approach of adhering to these musical procedures, where a computer strictly follows the rules coded by the composer. As with any art, rules are made to be broken and challenged, and with the power of modern computing it is fairly simple for a composer to edit the MIDI output of their algorithms in order to suit their work's musical objectives. However, in regards to my own personal work, part of the beauty of algorithmic music is the relationship of sound to the original data set(s). Too many abstract edits or functional changes may create a disconnect between the original images and the resulting sonic materials.

This creates a problem: how can one create algorithmic compositions that conform to certain aesthetic sensibilities and still have the output represent the original data? Each composer has their own musical identity and must carefully define and understand their musical language before they can begin to build a facsimile algorithm. I chose the word language here deliberately, as the process of creating a musical algorithm is essentially a complex translation from one form of meaningful data (language) into a musical language. If both languages are well understood by the composer and represented accurately in the algorithm, the result will often be satisfactory with minimal edits needed. Compositional and musical elements that I find beautiful and incorporate in my musical language include:

- Layers of complexity, where my music may be both accessible at a surface level and rewarding for active listeners.
- A balanced juxtaposition of musical extremes, including consonance and dissonance, order and chaos, and familiarity and the unknown.
- An evolution of the materials so that the transitions between these extremes feels natural in spite of potentially dramatic shifts in tone and sound mass.

For this project, I used Max for Live and Max/MSP to handle the process of translating the musical language I feel best produces music that satisfies these ideals.

My background in music composition up until attending the Ph.D. program at UC Riverside had been almost exclusively acoustic, however I had wanted to work in an electronic medium for quite some time. While acoustic instruments will always have a place in my work, they are also saturated with limitations that I do not find particularly inspiring, and my compositional practice is typically bounded by self-imposed rules to restrain my writing. Once I was introduced to Max, I was able to relegate my computer to handle the algorithmic processes I'd previously generated myself. For this project, I created four Max patches to handle various tasks:

- M-Score Generator. Generate fractal scores derived from the Mandelbrot set (fig. 1).
- M-Score Reader: Translate and playback data extracted with M-Score Generator as MIDI.

- M-Audio: Take a black outline of these scores and convert them into a waveform. Load and playback these visual scores.
- Iterative Visualizer: Generate visuals based on similar functions used to create the music.

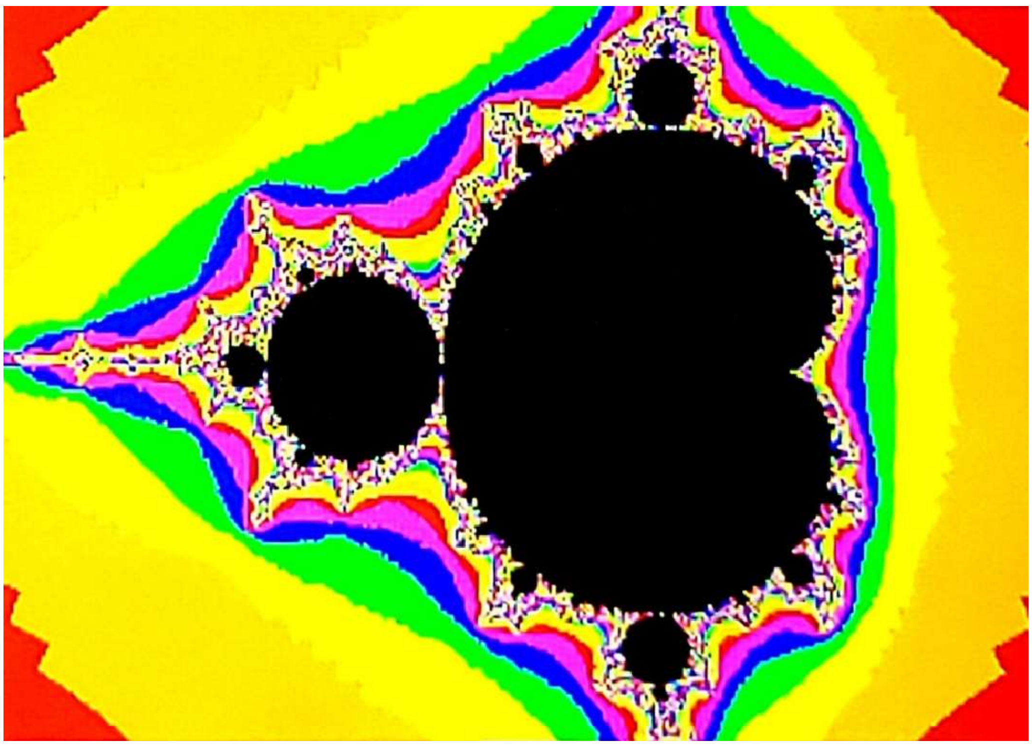

(fig. 1: The Mandelbrot set.)

These tools saved me a tremendous amount of time, but more importantly, made the fine-tuning process for my algorithm much simpler with quick rules editing and on-demand playback of the material. It also allowed me to handle a greater density of data than I'd ever be able to work with by hand. In fact, I would recommend more powerful programming languages when dealing with large data sets and high-quantities of mathematical iterations. I intend for these tools to be easy to learn for novice composers to experiment and discover their musical language with, while being deep enough to satisfy the wide-ranging needs of professional composers. For this project all of my data sets were derived from the Mandelbrot set.

The Mandelbrot set is a fractal image rendered by a fairly simple mathematical formula, where a complex coordinate is entered and the results are fed back into the function. If the result exceeds two within a set number of iterations, the result “escapes” the set, and is colored based on how quickly it escapes. Otherwise, the point is within the set and colored black. This set acts as an infinite data source in the context of my composition, and has multiple notable musical qualities that make it an excellent generator of algorithmic music. It is symmetrical about its horizontal axis, exhibits self-similarity at various zoom levels, and has multiple differentiable areas of intensity. In order to maintain the relationship between the mathematical and musical languages, it was important to use this source material in as many ways as possible. The Mandelbrot set controls all of the musical elements of my project, but more than that it is itself a musical element. Using my Max tools, I was able to create fractal waveforms using the Mandelbrot set and various zoom images. Finally, I used the Mandelbrot set to create three videos: Two being generated by the audio, and the other from AI images where the most visually interesting material acted as seed images.

As this work is generated from fractal imagery, it was important to include a visual component to this project to reinforce that connection. The video for the first movement is generated directly from the audio in a similar iterative process to the Mandelbrot set, creating a lightshow that starts with the same color palette from various mathematical functions of the right and left audio signals. When overlapped at lowered opacity, they blend and create even more complex and interesting colors and textures. I created the video for the second movement using AI images for a couple of reasons; first, the iterative process of generating a new image from the previous image and a text prompt is reminiscent of the function to generate the Mandelbrot set. Second, the level of control I had over the images and the quality of the images rendered was surprisingly high. In the final movement, I chose to have a simple spectrograph viewer, as I found it had the most visually interesting spectra.

This paper will discuss in detail the above elements, including the theory I created and utilized in order to produce them, the pitfalls encountered, and my solutions to them. Before diving in to my project, I will provide a brief historical context for my work, from the early mathematically-derived works of Iannis Xenakis. I will also provide examples of the works of various mathematicians and composers regarding the sonification of fractals and the translation process from a mathematical language into music.

Chapter 1: Historical and Contemporary Contexts

Early Influences of Math on Music

Mathematics and music are inseparable, a relationship noted early in human history. Pythagoras is responsible for the discovery of the mathematical ratios that make up different musical intervals, for example a ratio of two-to-one being an octave and three-to-two being a fifth.[1] Euclid treated "musical acoustics as a branch of arithmetic and proposes a definition of consonance limited to intervals built on multiple or superparticular ratios."[2] Mizler, a polymath and associate of J.S Bach, "advocated what he felt was still an unattained goal: the establishment of a musical science based on mathematics and philosophy."[3] This goal to discover and work with the mathematical language of music developed rapidly with the aid of computers in the late 1950s onward. Max Mathews, for whom the program Max/MSP was named after, wrote on the potential of computers to synthesize sound and production of music: "A computer can be programmed to play "instrumental" music, to aid the composer, or to compose unaided. With the aid of suitable output equipment, the numbers which a modern digital computer generates can be directly converted to sound waves. The process is completely general, and any perceivable sound can be so produced."[4] My project is only possible through

[1] Andre Barbera, "Pythagoras," *Grove Music Online*, 2001.
[2] Andre Barbera, "Euclid," *Grove Music Online*, 2001.
[3] George J. Buelow, "Mizler von Kolof," *Grove Music Online*, 2001.
[4] Max V. Mathews, "The Digital Computer as a Musical Instrument." *Science*, Vol. 142 no. 3592 (Nov. 1, 1971), 553.

these advancements in our understanding of the mathematics of music and algorithms. One of the first notable composers of algorithmic music of this kind was Iannis Xenakis.

Iannis Xenakis

Iannis Xenakis was born in Romania in 1922.[5] From a young age he was exposed to all kinds of music, from classical piano to Roma bands to Greek folk music.[6] As a young adult, he studied architecture at Athens Polytechnic before it was shut down by the invading Italian army in 1940.[7] Escaping a death sentence in Greece for political activism, both on far-right and far-left sides of the political spectrum, he ended up stranded in France, working as an engineer and architect until 1959.[8] Still passionate about music, he took private lessons with Honegger and Milhaud, and through his relationship with Scherchen he was constantly in contact with experts in electro-acoustic music.[9] From 1957-1962 he worked in Pierre Schaffer's studio, and shortly after composed a suite of instrumental works on a computer at IBM Paris.[10] From 1967-1972 he ran the Center for Mathematical and Automated Music at the University of Indiana Bloomington.[11] He received his doctorate at the Sorbonne, where he taught from 1973-1989.[12]

---

[5] Peter Hoffmann, "Xenakis, Iannis," *Grove Music Online*, 2001.
[6] Ibid.
[7] Ibid.
[8] Ibid.
[9] Ibid.
[10] Ibid.
[11] Ibid.
[12] Ibid.

His works were heavily inspired by his study of architecture and math; both his works *Sacrifice* and *Metastaseis* feature the use of the Fibonacci series, a sequence of numbers that is generated by adding the previous number in the series to the current number, to control musical elements such as pitch, duration, and dynamics.[13] His macroscopic stochastic theory involved focusing on the overall quality of the sound mass: including texture and temperature, "like the pizzicato-glissando clouds in Pithoprakta, and their alteration in time."[14] His use of new computer technology and work regarding sound synthesis led him to the opposite extreme; microscopic stochastic theory, which focuses on the motion of an individual sound within the sound mass.[15] His theories that are most relevant to my project includes globally tempered sieves, or the use of a melodic line at various speeds superimposed on each-other, and cellular automata, "the development of sound in time by means of deterministic chaos."[16] Of Xenakis' work, I have chosen to examine *Metastaseis*, for these and other elements that are related to my project.

Metastaseis

Metastaseis was written by Xenakis in 1955.[17] It is approximately ten minutes in duration, and has two distinct sectional elements: long glissandi, and shorter notes accompanied by bursting pizzicati. When analyzing the graphic score, I noticed that the

[13] Ibid.
[14] Ibid.
[15] Ibid.
[16] Ibid.
[17] Iannis Xenakis, *Formalized Music: Thought and Mathematics in Composition*, (Stuyvesant, NY: Pendragon Press, 1992), 10.

glissandi sections resemble the bifurcation of the logistic map that corresponds to significant points of the Mandelbrot set (fig. 2 and fig. 3).[18][19]

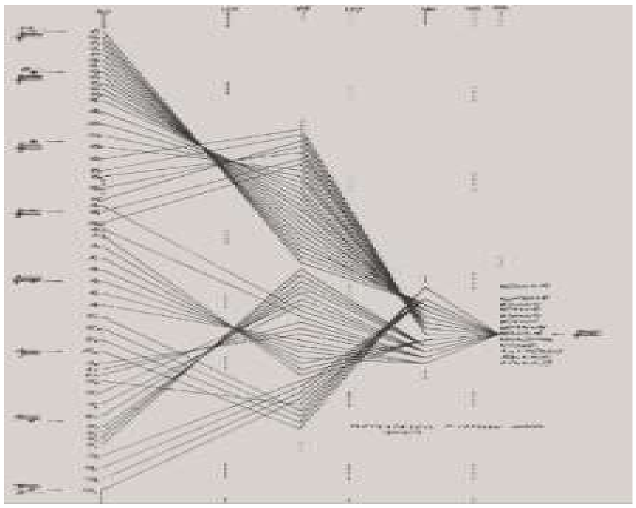
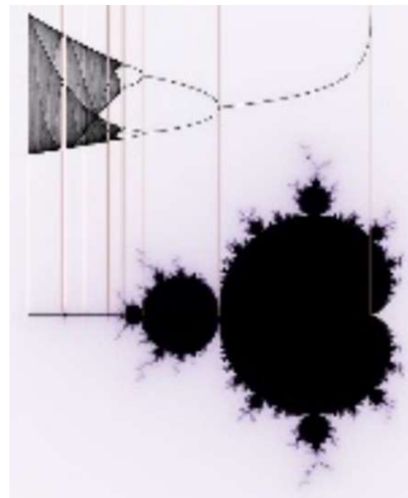

(fig. 2: A comparison of an early sketch of the ending glissandi section from Metastaseis. fig. 3: The logistic map, shown with relation to the Mandelbrot set.)

This is particularly fascinating, as this piece precedes the Mandelbrot set by multiple decades. It is possible these are related, as Xenakis' theory of cellular automata is centered on deterministic chaos, and the logistic map is "an archetypal example of how complex, chaotic behaviour can arise from very simple non-linear dynamical equations."[20] While this piece was written before his focus shifted to sieve structures and cellular automata in the eighties and nineties, I would argue that *Metastaeseis* lays the groundwork for his future compositional style.[21] Using the Fibonacci sequence to control

[18] Catherine Anne Hope and Michael Terren, "The Possibilities of a Line: Marking the Glissando in Western Art Music," *Proceedings of the International Conference on Technologies for Music Notation and Representations*, (ARU: 2016), 2.
[19] T.S. Sachin Venkatesh and Vishak Vikranth, "Investigating the Relation Between Chaos and the Three Body Problem," *arXiv: Chaotic Dynamics* (2020), 1.
[20] Ibid.
[21] Peter Hoffmann, "Xenakis, Iannis," Grove Music Online, 2001.

musical elements is a step towards this type of chaotic behavior, and the overlapping string glissandi at different rates is a precursor to melodic sieve structures.

The first section starts with long held glissandi in the string section, springing upward from a static line at different rates to create a web of sound. This is similar to the usage of time-stretching in my own project; where sections would overlap and play the same material simultaneously at different “tempi”. These glissandi culminate in a long held dissonant chord, breaking briefly before repeating with tremolo. The brass emulates this pattern, either playing two note phrases or creating short glissandi while the chord in the strings is held.

The middle sections maintain a strong feeling of serialism, both in the dissonance created through intervallic relationships, long leaps ignoring Western voice-leading rules, and the sudden and sharp textural shifts. Long-held notes are contrasted by short staccato and pizzicato bursts. This seemingly chaotic mass of sound is clearly representative of his macroscopic stochastic theory, where the greatest attention is paid to the sound mass itself as opposed to the motion of the individual instruments. This relates to my project by design, as the form of each of my pieces involves overlapping each section and adding multiple layers to play simultaneously.

The third section is similar to the second, with the inclusion of more brass instruments and at greater volume. The final section ends in a similar way as the piece began with long glissandi. However, instead of glissandi emerging from a static horizontal line, the glissandi move from a static vertical structure in both directions until uniting in a downward motion.

Math Inspired by Music

While music is constantly making use of new and exciting mathematics, mathematicians are also inspired by musical problems. Brady Haran interviews mathematicians on his YouTube channel *Numberphile*, where his interviews with Dr. Neil Slone occasionally include some sonification of the functions he chooses to share.[22] Slone's function titled "Fly Straight Damnit" was particularly fascinating to me, as it is seemingly random until a certain point, after which it becomes linear.[23] This was an exciting discovery for my project, as this sequence can be used to create two separate types of patterns, one being deterministic chaos, and the other reminiscent of a Shepard's Scale. Auditory illusions are fascinating and compliment the illusory nature common in fractal waveforms, so I included this sequence as the only non-Mandelbrot score option in my M-score reader, the device I created to read my Mandelbrot scores. Slone also discusses his website OEIS (On-Line Encyclopedia of Integer Sequences), where sequences can be played back audibly.[24] I used this resource to examine if a symmetrical scale I created (0, 3, 5, 6, 8, 11) had any interesting mathematical implications. The site confirmed two mathematical sequences associated with this scale, and I was able to play back each sequence via MIDI using a variety of simple instruments.

[22] "Amazing Graphs," YouTube, uploaded by Numberphile, Aug 8, 2019.
[23] Ibid.
[24] Ibid.

While Slone's interest in the sonification of mathematical sequences may or may not intended for music composition/production, others are attempting to write algorithms capable of generating aesthetically pleasing music. Most useful to my research was the article *L-Systems, Scores, and Evolutionary Techniques* by Bruno Lourenco et al., which discusses the use of a different type of iterative function for the purposes of creating music.[25] While L-Systems did not factor into my project, I was inspired by the concept of mutation rules.[26] The results of having the algorithm controlling every musical parameter, including which of these elements are controlled at any given time, was incredibly useful for expanding the types of lines that the device could produce. For example, allowing the device to enable and disable the melodic and rhythmic length functions meant that some areas of the piece would repeat and others would be a constant stream of new phrases. This creates contrasting areas of coherence and relative chaos that form an important part of my musical identity.

The Sonic Mandelbrot

While *The Mathematical Life and Death Cycle of the Universe* may be the first work of art music to use the Mandelbrot set to control every element, many others, including Craige Hales, have experimented with the Mandelbrot set in a sonic capacity and generated the resulting waveforms, which greatly assisted me when it came to creating my own.[27] Others, including Gustavo Diaz-Jerez, have sonified the set, often by

---

[25] Bruno F. Lourenco et al., "L-Systems and Evolutionary Techniques," *Proceedings of the 6th Sound and Music Computing Conference*, (SMC: 2009).
[26] Ibid, 115.
[27] "Mandelbrot Sound [seizure warning]," YouTube, uploaded by Craige Hales, Jan 12, 2019.

what appears to be taking the main image and tracing its borders, then mapping the change in amplitude from the center to frequency.[28] Each of the works that sonify the Mandelbrot set in this manner tended to present the set in a static way, however if one were to zoom in while their program continued to render audio, the results would likely have more variation and motion. Prior to this project I had a limited understanding of the Mandelbrot set, so I started with learning the basics of the formula in order to render my scores.

[28] "Fractal Music - Image Sonifications (II) - Mandelbrot Set," YouTube, uploaded by Gustavo Diaz-Jerez, Feb 20, 2009.

Chapter 2: The Mandelbrot Set

The Mandelbrot set is a surprisingly simple formula for the incredibly complex applications associated with it. I will not delve too deeply into the mechanics here; the focus of this project is the music that may be generated using the basic principles and extracted data sets. I will briefly define what a fractal is, provide a biography of Mandelbrot's life, and describe the basic formula and process for rendering the Mandelbrot set.

Fractals

Before we discuss the Mandelbrot set, we must define what a fractal is. Fractals may be "loosely defined as having a characteristic form that remains constant over a magnitude of scales."[29] While some fractals are more self-similar than others, this is particularly true in the case of the Mandelbrot set, where zooming in and out of specific coordinates reveals the original Mandelbrot set image, regardless of depth. Fractals also have a fractional dimension between the topological and Euclidian dimensions, and therefor fractal dimensions serve "as a measure of the scale-independent irregularity, roughness, or variation of a system."[30] A fractional dimension indicates that these objects have an infinite perimeter: "We see then that we have a geometry which allows to see the infinite within the finite, the fractal geometry."[31]

---

[29] R W Glenny et al., "Applications of fractal analysis to physiology," *Journal of Applied Physiology* vol. 70, 6 (Bethesda, Md. 1991), 2351-67.
[30] Ibid.
[31] Marfa Antonia Castro and Marfa Jose Perez-Luque, "Fractal Geometry Describes the Beauty of Infinity in Nature," *Proceedings of Bridges Conference*, (Alhambra, Spain: 2003), 409.

Benoit Mandelbrot

While not the first to work with fractals, Mandelbrot first coined the term.[32] Born in Warsaw, Poland, in 1924, he would leave at the age of eleven to escape the Nazis.[33] He studied at École Polytechnique in Paris, and the California Institute of Technology in the United States.[34] He worked at IBM in New York for thirty-five years beginning in 1958, and taught at Yale from 1987 onwards, where he became the Sterling Professor of Mathematical Sciences in 1999.[35] His most notable works include *Fractals: Form, Chance and Dimension* (1977), and *The Fractal Geometry of Nature* (1982), where he argued that the behavior of some iterative processes in mathematics could be defined by a singular set, now known as the Mandelbrot set.[36] Mandelbrot's research achievements were celebrated throughout his career including fellowships at the American Academy of Arts and Sciences (1982) and of the National Academy of Sciences (1987), as well as being awarded the Wolf Foundation Prize for Physics (1993) for his work on fractals, and the Prize of the Science and Technology Foundation of Japan (2003).[37]

The Mandelbrot Set

The Mandelbrot set may be described as: "A fractal that when plotted on a computer screen roughly resembles a series of heart-shaped disks to which smaller disks are attached and that consists of a connected set of all points c in the complex plane for

---

[32] "'Fractal' mathematician Benoit Mandelbrot dies aged 85," *BBC News*, 17 October, 2010.
[33] Ibid.
[34] The Editors of Encyclopaedia Britannica, "Benoit Mandelbrot," *Encyclopedia Britannica*.
[35] Ibid.
[36] Ibid.
[37] Ibid.

which the recursive expression zn+1 = zn2 + c for n = 0, 1, 2, 3… with the starting value z0 = 0 remains bounded as n approaches infinity."[38] This function is iterative; the complex coordinate data (z + ic) is fed into the function, with the result feeding back into the formula until a desired number of iterations is reached. The number of iterations changes the "resolution" of the resulting image; the higher the number of iterations the more accurate the resulting image will be. If the number is stable after each iteration (<2), the complex coordinate becomes part of the set, and is generally colored black. Otherwise, it is typically colored based on how quickly it "escapes" the set. There are many alternate formulas that produce the same results, one of which I employed for my M-Score Generator.

[38] "Mandelbrot set," *Merriam-Webster Dictionary*.

Chapter 3: Musical Algorithms and Audio Synthesis

I used Max/MSP to create a suite of patches that could handle the intense processing load this project demanded. Aside from the ten-thousand processes per second limit, it was able to accomplish everything I needed. While this program would be suboptimal for generating high-quality Mandelbrot images, these limits ultimately led me to choose lower resolution imaging, which helped limit the amount of information required to generate music based upon the entire set. The suite of devices I created for this project include M-Score Generator, M-Score Reader, M-Audio, and an iterative visualizer.

M-Score Generator

The first patch I wrote for this project was a device to produce fractal scores. This is also where I ran into my first problem: Max doesn't understand the square root of negative one, producing the error message "-nan(ind)." I had to change my search query a couple of times before I found an answer: a parallel escape time algorithm (ETA) at the set level, which essentially allowed for me to split the complex coordinates.[39] This meant that I no longer had to handle imaginary numbers. After a few failed attempts to render, I was able to generate the Mandelbrot imagery and collected the coordinate data for each of the sixty-thousand points rendered in a data collection (coll). Next, I implemented a zoom function, allowing me to change the distance between each pixel and the x-y

[39] V. Drakopoulos et al, "An Overview of Parallel Visualisation Methods for Mandelbrot and Julia Sets," *Computers & Graphics*, Volume 27, Issue 4 (2003), 639.

coordinate range. I did not zoom in too deeply, as it requires specialized knowledge to accurately locate specific patterns and is more effective for experienced users to explore.

Once I had enough material from the Mandelbrot set, I manipulated the formula to see what interesting images would render. At first these changes were small, limited to singular arithmetic functions such as changing addition to subtraction. Surprisingly, this did not change the resulting image in most cases, however the colored areas outside of the set formed some interesting new patterns (fig. 4 and fig. 5):

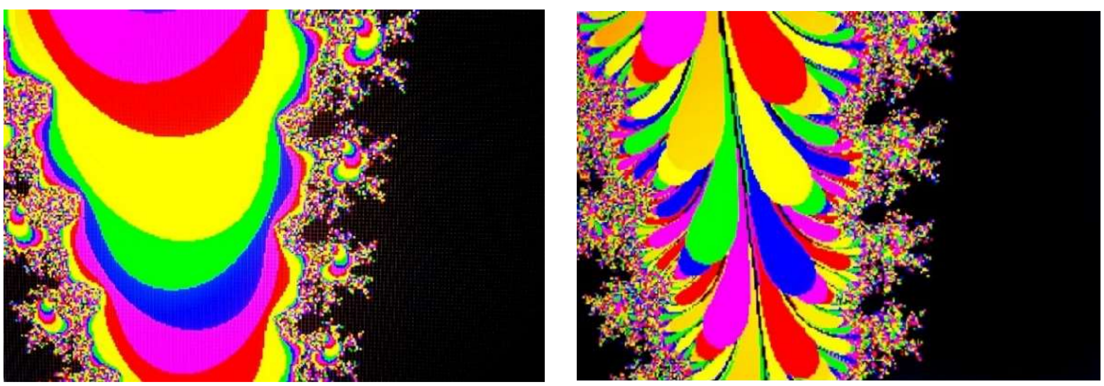

(fig. 4: A zoom of the original set. fig. 5: A minor functional alteration.)

The most interesting of these alterations was changing the addition of both sides of the ETA to division, resulting in a cross pattern. When zoomed into the background, they are revealed to be interlocking, disconnected rings that seem to descend infinitely (fig. 6 and fig. 7).

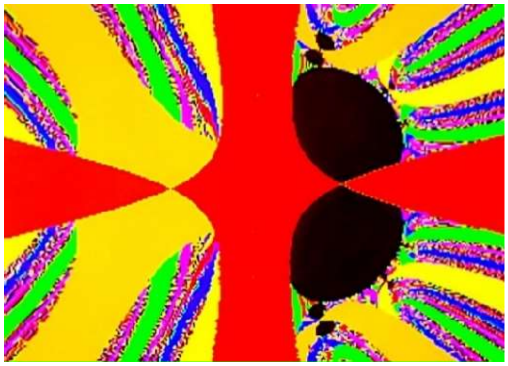 

(fig. 6: Division-altered ETA results. fig. 7: Zoom of division-altered ETA featuring interlocking rings.)

As I gained more experience, I experimented with more extreme alterations, including running the results through a second iterative stage, introducing a variable escape-time, and changing multiple elements of the base function at once. The results of these edits were fascinating, including Mandelbrot sets where none of the values were in range but the basic shape was maintained (fig. 8), to non-fractal imagery (fig. 9), and even recursive and potentially fractal imagery (fig. 10):

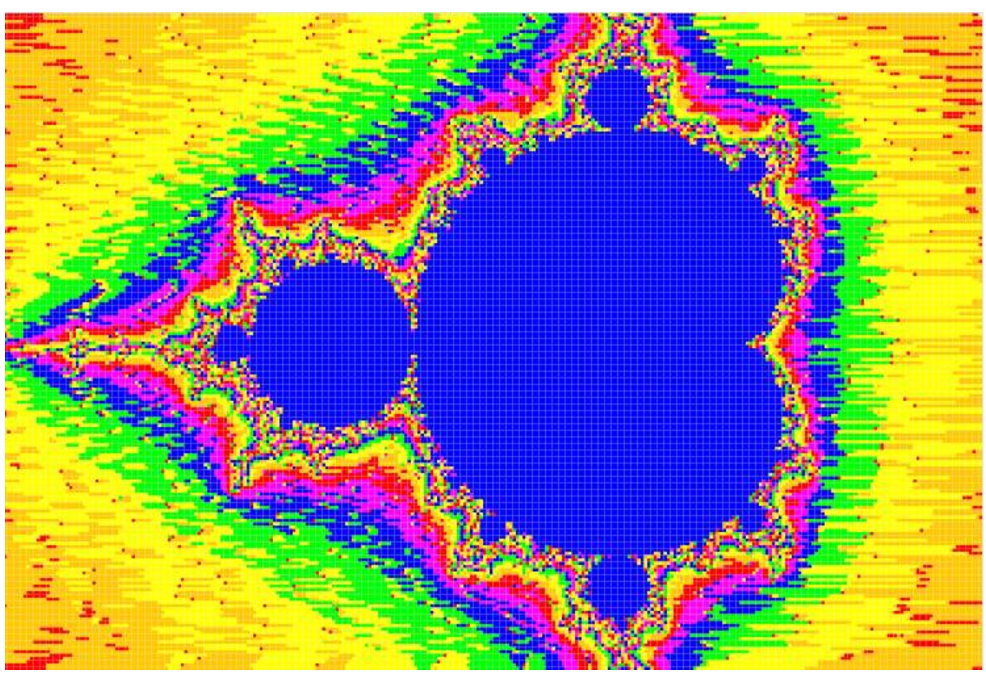

(fig. 8: A Mandelbrot image with no values within range. Generated with a heavily modified ETA.)

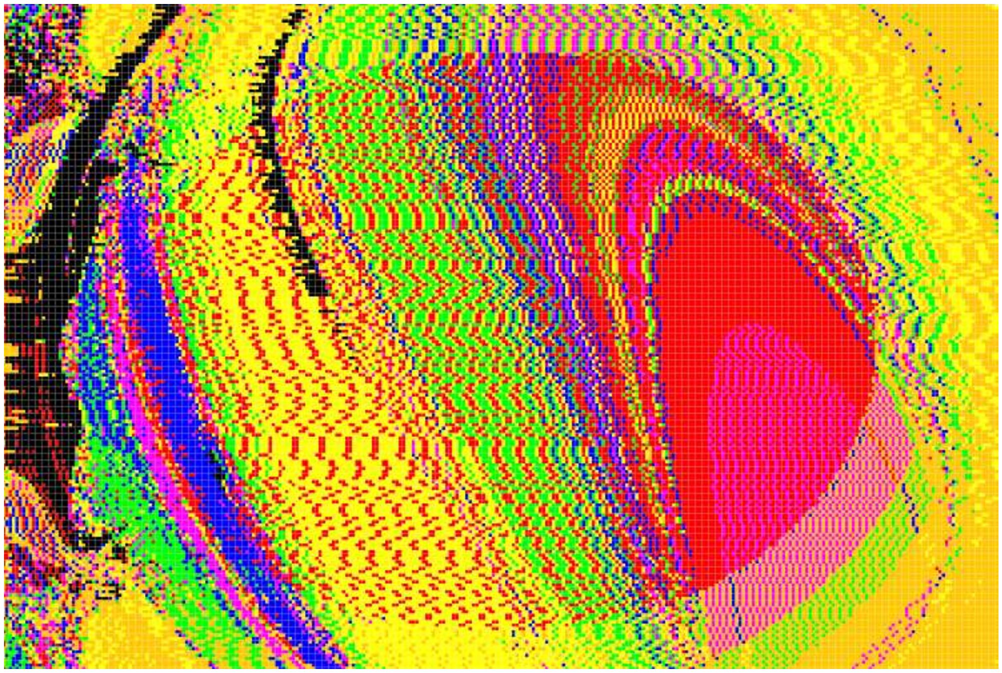

(fig. 9: Non-fractal imagery generated with a heavily modified ETA.)

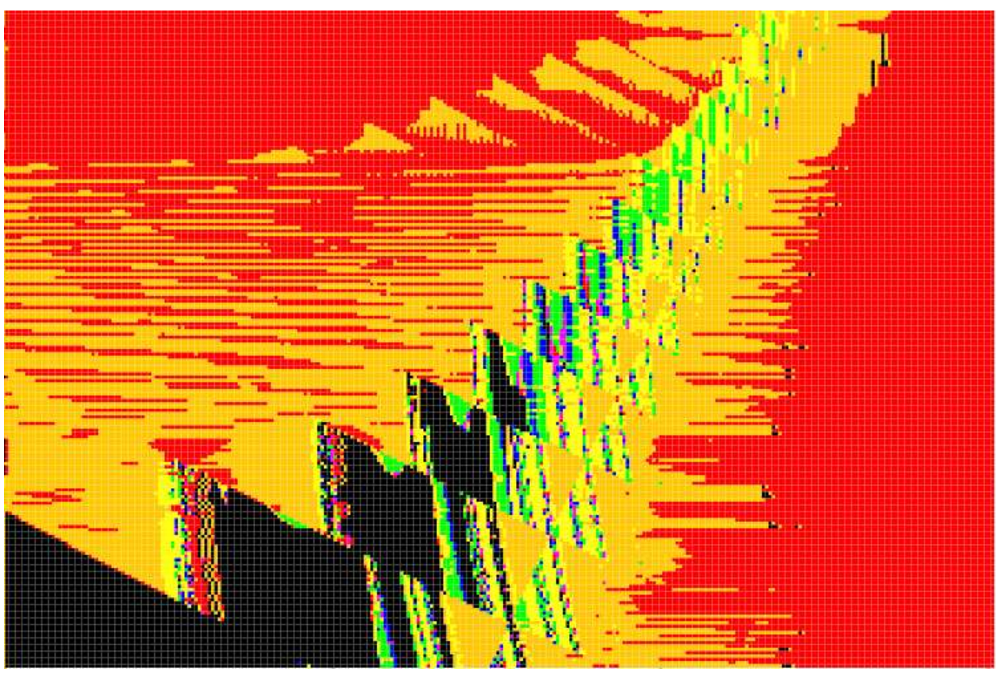

(fig. 10: Recursive, if not fractal imagery generated with a heavily modified ETA.)

M-Score Reader

As we have examined in previous chapters, the translation of data into music is not a new concept. Neither is the idea of translation on a smaller scale; taking qualities of a data set and using them to control individual elements of music. For example, quantization is a common effect in digital music that takes sound outside of a scale or rhythm and sets it within those musical bounds. However, the great appeal of experimental algorithmic music is that no theory is one-size-fits-all.

For example, working on my master's thesis I used pi as a data set to modulate pitch in what I consider an early step towards algorithmic music. The process was exceptionally simple; I took each digit of pi in order and took the modulus about seven to generate a scale degree. This provided me an additional chance at the tonic and second scale degrees, but otherwise this was almost a direct translation of each digit into pitch. I allowed myself to choose all other musical elements, from the octave to rhythm to dynamics. Surprisingly, I found this procedure highly effective, especially for how simple the processing turned out to be. This would not be the case for my other algorithmic works, with this project diverging the most from simple processes to a more complex translation that takes every musical element into account.

To begin with, this process produces two different types of data: the escape value, and the coordinate data. To generate pitch, the escape value is sent through a similar modulus about seven to start with, both corresponding to pitch and the color of the pixel that is being played. However, there is also a gate that sends the output to one of two translator destinations, depending upon whether the number equals the previous value. If the number is repeated, the value will continue through to a process that chooses the next pitch based on the previous scale degree and a function of the coordinate data. This keeps the phrase coherent by limiting the pitch options to theoretically sound choices, as well as integrating a separate but related data set into the equation. This creates variety that would be more challenging to produce with a single data set, as using the same number to select the intervallic change would produce monotonous results without significant abstract processing.

If the colors are different, the phrase will reset by bypassing this process and generating a new initial seed value. When the value is within a solid color field of the Mandelbrot set, the algorithm generates a distinct phrase, and once it leaves that field, a new phrase begins. This value then passes through a voice leading check, which tests the current value against the previous to determine if the pitch has gone beyond a major second interval. If it has, it will repeat that value plus or minus one to resolve that leap in the correct direction. This ignores the incoming pitch, but is ultimately necessary to the creation of more coherent musical phrases. Similarly, octave leaps are also considered in the same manner, adding or subtracting twelve semitones to keep pitches within a reasonable range. Finally, the value passes through optional processes such as a sequencer which saves a set number of pitches and loops them as opposed to allowing new pitches through, creating repeated motivic material derived from the color field.

Rhythm is controlled by a function of all data sets through a weighted routing system that multiplies or divides the user's set pulse by multiples of two. This limits the frequency of extreme rhythmic changes while allowing moderate variations to occur more frequently. Similar to pitch, there is an optional rhythmic sequencer that collects a number of notes and saves their values, repeating a set number of times. Velocity and pitch durations are also controlled, where both values are lowered on the perimeter of the score, and increased as the player approaches the center. This was decided with the base Mandelbrot set in mind, as I wanted values within the set to be the most significant. While the regions outside the set are more visually interesting, the area's musical stability was prioritized, as the algorithm would generally play longer continuous phrases within the set compared to the individual colored areas.

M-Audio

Creating fractal waveforms from these images was one of the more complicated elements of my project. My initial plan was to cut the Mandelbrot set in half and flip one of those halves horizontally, however this would have created vertical overlapping of the wave that is not possible with audible waveforms. While my research uncovered methods of developing fractal waveforms unrelated to existing fractals, or that were already functions suited to the creation of waveforms, I discovered video resources by Craige Hales that not only sonified the Mandelbrot waveform at varying frequencies, but also provided visuals that clearly demonstrated how fractal waveforms could be generated.[40][41][42] I started by creating a device that would trace the black-bordered area of

[40] Gordon Monro, "Fractal Interpolation Waveforms," *Computer Music Journal* 19, no. 1 (1995): 88–98.
[41] Manfred R. Schroeder, "Auditory Paradox Based on Fractal Waveform" *The Journal of the Acoustic Society of America* 79, no. 1 (January, 1986): 186–189.
[42] "Mandelbrot Sound [seizure warning]," YouTube, uploaded by Craige Hales, Jan 12, 2019.

any image I fed into it. This was fairly simple, essentially checking for black points that bordered colored areas and excluding everything else. Next came a considerably challenging task: finding the distance from each point from the center in order. I knew I would have to translate to polar coordinates, but wasn’t sure which function to use to calculate the distance between each point and the center, and how to keep the data in order. After relearning the trigonometric functions, it seemed like the correct option was arctan. This function created shapes that were visually promising, however they were in the wrong order. More troubleshooting and research lead me to atan2, which takes the same process as arctan and adds a second argument that handles the ordering without having to sort it manually. After confirming the waveform was accurate in a visualizer, I loaded the resulting data into a wavetable. I converted the main set, a zoom, and a mutation of the set into waveforms in this way.

The main Mandelbrot waveform (fig. 11) is a bright saw with a lot of buzz and noticeable beating properties.

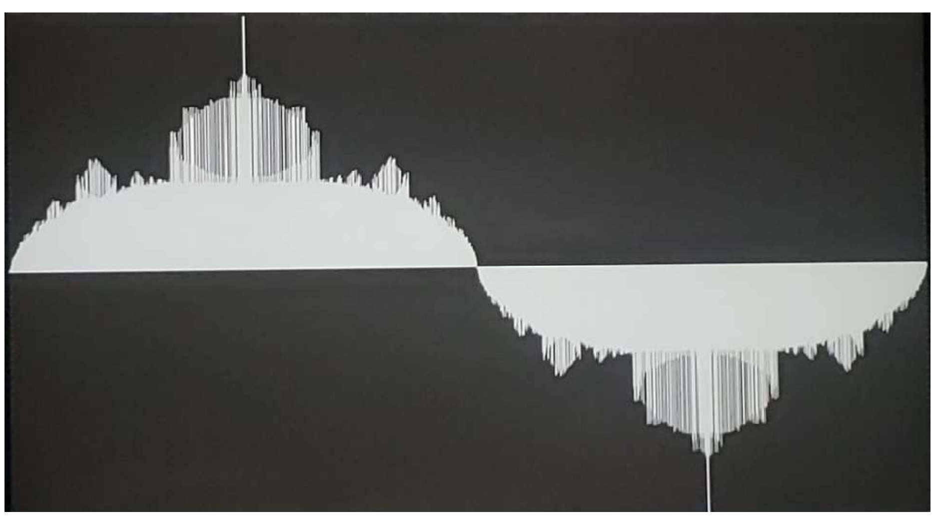

(fig. 11: Image of the Mandelbrot waveform.)

I had a difficult time getting this waveform to fill a lead instrument role, as it seems to take on different timbral qualities based on register. However, it makes a rich pad that is successful in various roles. For example, in *Cave Paintings*, I use it in the opening section to create a bright and triumphant explosion of energy. Later on, it has more of a whining insect sound, buzzing in and out at high frequencies. I employed this waveform in all audio generated for the project. The zoom that I used was significantly less effective: similar to the Mandelbrot waveform but dull and comparatively lifeless. I rarely employ this waveform other than to tame higher harmonics, and it will not be in the released version of my software. Finally, the mutated score I used to create the third waveform was just as interesting as the Mandelbrot waveform, only with a darker, heavier color to it. I use this in the first movement of my project, both by itself and blended with the Mandelbrot waveform.

The device itself is fairly simple. It has one oscillator with the option to add a second, allowing the user to combine the fractal waveforms. I included an amplitude envelope and basic filters, including high-pass, band-pass, and notch filters, which can be combined with each other as well as the oscillator's output. Users may switch between mono and stereo signals, with the stereo option alternating between left and right channels. This, along with the detune option, is incredibly useful for creating binaural beats. The maximum number of polyphonic voices can also be controlled, which is primarily useful in keeping the audio from creating clusters at high tempos.

Iterative Visualizer

I wanted to present three different visual components with my project, one of which was a visualizer that would translate the audio back into a visual format. I take the Fast Fourier Transform (FFT) of some arithmetic function of both channels, then periodically send out the maximum number of resulting numbers each second. I run the result through a basic iterative process, and then take the modulus to route each number to paint the current pixel a primary color, secondary color, or black. The algorithm determines whether the video renders horizontally or vertically, which creates variety within the imagery. When this value is alternated rapidly, a diagonal effect is created. The results are a vivid lightshow that meaningfully responds to the musical input.

The Future of this Project

While these devices are in a workable state for my own personal use, I intend to distribute them to a broader audience, and there are a number of issues that must first be resolved. M-Score Reader has a number of memory issues; even after the reset button is pressed there are some previous values that play an undesirable influence. While these artifacts typically have a minor impact and can have musically interesting results, most users will likely expect the output to remain the same if the settings are unchanged. To compromise, I intend to make this a feature as opposed to a bug by having a button to toggle whether some leftover data remains or not.

Even with the lower resolution imagery there is still far too much data, to the point that I rarely represented the entire set. In fact, reading the fractal scores pixel by pixel, it would take approximately five hours at the current fastest tempo in order to accomplish this goal. I created a coordinate scaler for my score reader in order to reduce the amount of data to more manageable sizes, however if the scalar value is too high, the parameters have a tendency to change too quickly in sections where I intend stability. One possible solution could include reading each column, either taking an average of all points or using different clusters of points to control different musical elements.

Rhythmic functions in M-Score Reader are not bound to a time signature and do not indicate a default note value, both features that will be crucial options for the commercial model. While polyrhythms are fairly simple to produce between two devices by offsetting the pulse of one device to the desired ratio, an option to create polyrhythms with a single device would be more data efficient and ensure that these rhythms are in time with each other. These would be additional preset sequences within the weighted rhythm selection that inject into the rhythmic sequencer when chosen by the algorithm. I would give the user the option to choose which note lengths and polyrhythmic sequences the selector could choose at any given time. Allowing the mutate function to modulate all of these parameters would greatly expand upon the variety and complexity of possible rhythms, as well as improve differentiation between formal sections within the musical ouput.

While the mutate function for M-Score Reader is extremely useful, some features that should be able to be algorithmically controlled independently are tied to this function, such as key changes. Key changes themselves can be jarring, as the function that controls them is fairly simple, checking for leaps of a fifth and modulating to the dominant key. There are also elements that are controlled via mutation that I would remove, such as the algorithmic velocity, as this tends to create sudden spikes in volume that can cause clipping. While I tried to keep the rules as simple as possible to maintain the relationship between the fractal score and the generated music, more complex sets of guidelines could provide further assistance.

M-Audio currently only has three fractal waveforms, and they sound fairly similar. I will research further into adding new fractal waveforms that would have more differentiated timbral qualities. I am particularly interested in adding Julia set-derived waveforms, as these sets are related to the Mandelbrot set. However, I have also started developing my own fractal imager, and would be interested in creating waveforms derived from these images. While I am significantly less experienced with audio compared with MIDI, I am constantly discovering new options for shaping sound.

Both M-Score Generator and my Iterative Visualizer currently do not have a GUI, as I wrote these programs in Max/MSP as opposed to Max for Live. For both of these devices, I plan on utilizing jitter objects in order to display graphics in the future, as live.cellblock was not designed to be used as a visual display. This will also allow me to reduce the size of the screens for these devices to fit into the standard Max for Live dimensions, as well as create small score samples to display in M-Score Reader so the user can differentiate which score they are using.

Chapter 4: Audiovisual Art for Audiovisual Math

History

Before discussing the techniques used to create my videos, I will briefly discuss the history of audiovisual art. One of the most notable examples of early film set to music is Disney's *Fantasia* (1940), which sets themed cartoons to classical music. However, music had been set to silent videos decades earlier to create "song-plug" films.[43] The Beatles experimented with promotional videos for their songs meant for TV that included many of the standard features of a music video, however these "did not attract any wider attention or achieve tangible commercial effects."[44] The video for Queen's *Bohemian Rhapsody* may be considered one of the first music videos, both for the use of visual effects, and for having a significant impact on record sales.[45] Music videos became commonplace in the eighties as an advertising method for popular music, and the TV channel MTV hosted these videos twenty-four hours a day beginning in 1981.[46] These videos became far more advanced as technology became more powerful; most notably in the 2000s with digital film and CGI.[47] The most relevant audiovisual work that influenced my project was Ikue Mori's 2007 masterpiece *Bhima Swarga*. This piece uses a glitchy, psychedelic animation style to tell the story of Bhima from the Indian epic the *Mahabharata*. This video heavily inspired my work Transcendence for its use of visual

---

[43] John Richardson, "Video [music video]," *Grove Music Online*, 2020.
[44] Ibid.
[45] Ibid.
[46] Ibid.
[47] Ibid.

elements to tell a story without explicitly stating the narrative, as well as for its style of animation.

Project Videos

There are three video components to my project, one for each movement. I wanted the visual elements of my project to be as closely related to the sonic elements as possible. For example, the form for the second movement's video mimics the musical form, and both the first and third videos are generated directly from the audio signals. This was crucial for creating a natural synchronicity between what is seen and heard, and saved a lot of time and energy compared with animating and syncing the material manually.

The video for the first movement, *Hours into Minutes*, is a lightshow generated in Max/MSP. As each pixel is rendered individually, I sized my screen so that Max could render the full display quickly enough to create an animated experience while still being scalable to the aspect ratio I wanted. The left and right signals of the audio are run through various mathematical functions and translated using an FFT, or fast Fourier transform. This algorithmic function allowed me to break the waveform apart into its frequency components, which are then translated into coordinates and run through an iterative process that determines the color of each large pixel.[48] The choice of colors and iterative function tie the elements of my project to the subject matter in as many ways as possible.

---

[48] G. D. Bergland, "A Guided Tour of the Fast Fourier Transform," *IEEE Spectrum*, 6, no. 7, (1969), 1, 2.

I captured multiple videos using different mathematical functions, keeping the more interesting variations for the final video. Each video sequence is then overlayed on top of one another using various blending types to lighten or darken specific takes. I used various color correction effects to differentiate colors and provide variety, and flipped some of the videos horizontally to create both variation and a level of symmetry.

The video for the second movement, *Transcendence*, was constructed using a far more intensive process. I fed images of the Mandelbrot set, as well as various zooms and functional alterations, into the *Deep Dream Generator*, an AI image generator derived from Google's *Deep Dream* project. In order to be able to monetize this project in the future, I used their paid services so that I would own any images generated from my source images. I purchased bulk packs of energy, which are tokens that are spent to generate images, in order to generate as many images as I could as quickly as possible. That said, Deep Dream is open source, and for future projects I plan to learn how to install and train AI software to have greater freedom and control, as well as minimize cost.

I used *Deep Dream Generator's* text-to-image function, which allowed me to have a tremendous amount of control over the generated images. While it occasionally struggled with text prompts, simply editing my text instructions to be more specific typically resulted in images that were highly effective. I selected themes that I felt best fit the mood of the formal section they accompany. For example, the A section of this movement is light and carefree, and has a choral quality that was reminiscent of western religious music. My text prompts included the phrases “clergy in a cathedral” and “angels flying,” which generated highly effective images that illuminate the music.

Once I had generated and organized approximately 700 images, I arranged them in a logical order to achieve an effective flow. One of the beautiful things about iterative generation is that visual changes from one frame to the next are often fairly minimal when the text prompt remains the same, which leads to naturally animating sequences without the need for extensive reordering. Between each frame I added a crossfade to create a flickering effect, which I attempted to synchronize with the pulse of the audio. To accomplish this, I set each frame’s duration to a fixed time based on estimates of when significant events would occur. In this case, I would argue that working in a time-based music system made these estimates significantly easier than if I had used a tempo-relative system, as fewer unit conversions were required.

I matched the visual form to the musical form. Once each section concluded, I would repeat that section time-stretched and at a lower opacity. This created a tremendous amount of variation for a video with minimal frames, where sectional chunks could repeat verbatim and still be fresh based on new elements that were overlayed in the background. As each sectional and even sub-sectional chunk often played back at different speeds, I was able to create a project full of motion in spite of the elements typically running at a much lower than normal framerates.

The final video is a simple spectrograph of the piece, and while the process for generating the visuals is not particularly involved, I will briefly discuss the considerations that were made. For example, the name of this movement is *Black Clover*, so I colored the spectrograph yellow, green, and a dark blue-green. While there is audio above the 9.5kHz range, I limited the project to the most audible range. While this process was exceptionally simple, I believe that the results are far more interesting than the process.

Chapter 5: Analysis of Audiovisual Project Results

The title of the complete project is *The Mathematical Life and Death Cycle of the Universe*. In its entirety this work is ninety-six minutes and thirty seconds long. However, I subgroup the movements so that the first two are the "A-side" and the final movement is the "B-side," and in traditional concert setting only one side should be played in a single sitting. Movements may also be played individually, for a minimum playtime of twenty-four minutes. Each work was initially a rondo form (ABACABA), however edits to the third movement made it more of a arch form (ABCBA). While these pieces are generated exclusively from the M-Score Reader, each has a unique sound and thematic place within the project as a whole. Each section includes fractal waveforms, and most include various sonic elements being tuned to create binaural beats, or "a perceptual auditory illusion occurring when presenting two neighboring frequencies to each ear separately."[49] While I may have my own interpretation of *Transcendence's* narrative, the work's program notes only provide some context of the theme and invites listeners to consider that a story is being told. Before I discuss this project, I would like to provide an analysis of the first work I generated as a study, since I learned a lot from this initial process.

Cave Paintings

The first attempted work using my Max devices, *Cave Paintings*, is likely one of the purist translations from the entire Mandelbrot set data into sound, as I used an early

[49] Hector D Orozco Perez et al., "Binaural Beats through the Auditory Pathway: From Brainstem to Connectivity Patterns," *eNeuro* 7, no. 2, (Mar. 2020).

version of my software with a limited number of musical translators compared to the current version. As I had initially intended to create a large-scale single-movement work, this is the longest piece I generated at one hour and six minutes long. In fact, due to the sheer amount of data from my low-resolution render, this piece was originally twice as long. To reduce play time, I cut what had rendered in half and used it as a separate track. This piece is now part of a separate project and is not included in *The Mathematical Life and Death Cycle of the Universe*. However, I will provide an in-depth analysis and submit this piece with my project, as it is the only track that features the entire set render and is an early step towards the pieces that currently make up my project.

Cave Paintings is the second movement in my project *Prehistory*, although I would generally recommend that it be performed as a standalone in traditional concert settings. It is meant to represent a number of critical moments in early human history, starting with the discovery of fire and ultimately leading to perhaps the earliest form of figurative art in cave paintings.[50] This is the only algorithmic piece where I made significant edits to the music itself, taking my favorite elements and repeating them in moments where it felt they would fit.

The piece begins with a rhythmic bubbling, the Mandelbrot waveform buzzing white-hot as an airy synth plays a slow melodic line. My initial intention was to follow a similar form to my work *Evalyn, Lost in the Fractal Fog*; each element from the piece is an evolution of what came before it. This iterative form not only fits well within the

[50] Jo Merchant, "A Journey to the Oldest Cave Paintings in the World," *Smithsonian Magazine*, (January 2016).

context of my project in general, but also with the theme of the piece, as the majority of technological developments in human history could not have been discovered without prior advancements to build upon. The set waveform here acts to represent fire, or perhaps a lightning strike that causes fire. A pan-flute instrument triggers so rapidly that it becomes pitched percussion. At about the forty-five second mark the set waveform repeats in a higher register, creating a triumphant feeling as the rest of the sonic elements build and release simultaneously. This may represent early human discovery of fire, and immediately recognizing its practical applications: warmth, protection, and light. Nature audio is present in this piece, starting with bird sounds. A piano is introduced, creating a heavy, lumbering melody amongst the chaotic pan-flute percussion. I intended for this piece to act as a piano concerto, where the piano would take the leading role to represent early human subjects, which becomes apparent as the piano part increases in tempo, frequency, and complexity.

The pan-flute percussion begins to subside around the eleven-minute mark, allowing the piano to shine through. The percussion fades in and out throughout the piece, providing moments of both relative calm and heightened tension. An eerie synth chord accompanies the piano here, and can be heard throughout the piece, vaguely sounding like a human voice singing "Oooo." An elk call can also be heard around this moment, and shortly after the lumbering piano returns. This may imply that the human tribe is hunting, and perhaps have discovered the usefulness of fire in cooking their food. The bird sounds return, time stretched to resemble a harsh wind blowing. This may foreshadow the final movement in the Prehistory titled *Weather the Storm*. A new melody is introduced in the piano a few minutes later, with alternating moments of sweetness and violence. This represents the modern human spirit just as much as that of our ancestors, creating a connection between the audience and the subjects of the piece. I repeat this moment throughout the remainder of the piece, as I found it to be both musically and thematically effective.

The Mandelbrot waveform noticeably picks up around the twenty-minute mark, along with a strange arpeggiated synth in the background. The piano melody repeats here, pitch shifted from its original presentation. The synth becomes much clearer and more prominent and provides consonance in its individual line, but also dissonance as it phases in and out of key with the rest of the material. This has a strong futuristic feel to it that is a stark contrast with the primitive theme of the piece, perhaps acting as the reflection of the viewer as they watch their ancestor's work unseen from a time machine. There is a passage of reversed percussion here as well, which adds a layer of noise from the cymbals. Noise is inherently fractal, both infinitely complex and self-similar, and as such I wanted this and future pieces on this subject to include a degree of noise. This adds an additional timbre and texture within the overall noise object, metallic and fine as opposed to the thicker airy sound of the slowed bird audio.

New piano elements occur around twenty-eight minutes in, and foreshadow the final triumphant piano theme at the end of the piece, which bursts from a repetition of this theme. The percussion starts to build in intensity, rhythmic wooden crashes enter and exit before the drums as a whole finally exploding in tempo and volume at around thirty-four minutes. Perhaps this could be a conflict with a large predator, who they were able to ward off with force and fire. The final piano melody enters shortly after this; however, it is presented at a much lower volume to act as a background element. This has dramatically less impact than it does towards the end of the piece, where the rest of the instruments provide room for this moment to shine.

The video for this project is currently a clay-colored spectrograph to tie into the ancient human theme. As interesting as I find the spectra, I plan on creating either another AI-generated animation or a live-action film for this work. As this piece's form is significantly more difficult to follow than the others, I would match the video to the significant sonic events described above.

Hours into Minutes

The first movement of *The Mathematical Life and Death Cycle of the Universe* is titled *Hours into Minutes*, and represents the neutral element of time in the life and death cycle. The name comes from the twenty-four-minute length, and is meant to imply the compression of a full day from hours into minutes. Making an electronic piece last a set amount of time would normally be simple, however I wanted each overlapping section to be given the same amount of time. Each section for each movement in this piece is six minutes long, and using a rondo form, I had seven six-minute sections. I wanted to maintain a high density throughout this movement, so I overlapped the sections as an even split down the middle. While I carefully removed most obstructive noise elements from the other two movements, I wanted to emphasize the weight of a typical day by keeping each section densely overlapped until significant noise makes it harder to hear the more musical, pleasurable moments of the piece. While they are still present, this noise makes it disorienting to try to follow until the end of the day when there is a rest period and we have a chance to relax and revisit the previous musical sections with clarity.

The A sections notable sonic elements include a thick panning bass, a repeated vocal sample, a drum synth, and a spatial synth. The synth layers start with a tonal harmony but quickly become dissonant as more long-sustained pitches are added. This creates a wall of sound where individual pitches may fade in and out and create different harmonic profiles, but the overall texture remains thick and static. The vocal elements are the only live audio in this piece, which separates it from the other two movements that are full of nature audio. The voice repeats the same phrase over and over, however the attacks overlap and cut each other off, creating choppy variations. The bass has the lone melodic function of this section, repeating variations of pitch classes 0-5-0-5-9-11-9-5. The drums provide a slower and more complex rhythm underneath, as well as a much-needed percussive texture. While there are some conventional rhythmic timings, they vary frequently and the tempo is quite slow, which creates a floating atmospheric environment. Fractal waveforms are audible throughout, acting as a bright saw-like noise pad. This is most prominent just before the B section starts to overlap around three minutes and fifty seconds in.

The B section features fractal waveforms more prominently than the prior section, and notably features sonic elements such as a bell, a sustained bass, new drum sounds, and a gliding synth. The fractal waveforms in this movement appear and disappear throughout, with their distinct buzzsaw sounds giving a scratchier and harsher texture in contrast to the smoothness of the bells. The bells are run through a major arpeggiator in a minor key, creating strange tonal and chromatic relationships. Rapid motion helps drive the piece forward and contrasts with the static A section. The drums start in a similar slow and measured manner as the A section. However, there are more intense moments in this section where either the same element is repeated rapidly, or there's a cascade of multiple different percussive instruments. The rhythms are well-defined and have noticeable meters, including 2/4 and 3/4 time, and are more in sync with the other instruments. The glide synth has a similar timbre to the bells with more sustain, and usually contains short repeated cells. The second half of this movement has the A section overlapped twice, once for the natural rondo repetition and the other for the halftime repeat, increasing the overall density by fifty percent.

The C section begins around the ten-minute mark. This section introduces a piano and a guitar, and contains various synth sounds, including a new drum set. This section has the highest sonic density of the piece with the presentation of the halftime B section played in reverse. The sound is so thick some elements of each section are lost in the noise, however the light melodic material in the piano and guitar strumming peak through. The piano melodies are somewhat reminiscent of the gliding synth from the B section in their pitch and rhythmic relationships, creating cohesion throughout the movement without the loss of interest of verbatim repetition. The subdued guitar strums two major chords a whole step apart, filling out the harmony in a more subtle way than the previous section.

The combination of these elements with each other and time stretched variants creates an overlapping form (rondo and ABC), producing a ramping in the sonic density of this movement as it progresses until receding back into a single stable sub-section at the work's conclusion. As the noise profile changes and grows with each new section, instruments that are not lost to the wall of sound pop in and out of what is audible. Some instruments, particularly drum and bass sounds, seem to cut through more easily, while slower, mid to high-mid frequencies are more likely to fade in and out of the sound mass. The bells from the second movement are pitched high enough and have enough motion that they are present when time stretched, giving a better sense of the arpeggiated harmonies and dissonances. The vocals are not omnipresent, however when they do cut through, they have different timbral qualities depending on factors such as direction, time stretching, and other instruments present. The majority of my edits throughout this piece, and this project as a whole, include these types of edits and variations, which help to reinforce the form and sense of location within the piece for active listeners.

Getting the visuals for this video sized to the aspect ratio I wanted was a problem I initially encountered. The first video had noticeable horizontal seams that were a result of overlaying poorly fitted videos on top of each other. The initial video was chromatically a lot cleaner, which, while excellent for relating back to the colors I chose for the Mandelbrot set, was not as visually interesting. The video for this movement (fig. 12) overlays screen recordings of my visualizer with multiple different mathematical functional processes, which produced a lot of varying textures that helped create coloristic variety when properly overlayed.

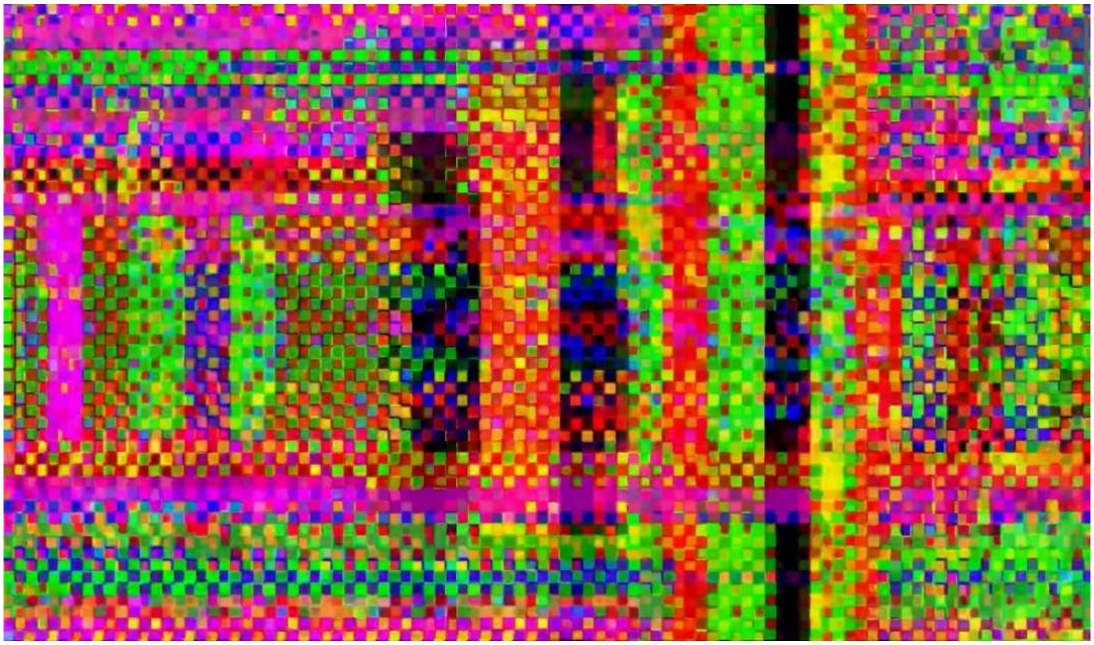

(fig. 12: Screenshot of the lightshow for Hours into Minutes.)

Transcendence

The second movement titled *Transcendence* runs thirty minutes and thirty seconds long, and is the most musically clear and thematically rich, meant to represent life within the life and death cycle. Thematic elements include the many forms of spirituality, karma, and love conquering all. None of the real-world religious associations are meant to be taken literally, and are only meant to provide a clear sense of the spiritual setting of the piece. In order to make this piece the most traditionally musical movement of the work, I carefully went through each track and removed any elements that contributed only noise. Some fractal audio was removed, however I made sure that at least one track of these elements was present in each section. In order to make each section clearly stand out from each other, I made sure that they were dramatically different sonically while still being compatible in the overarching theme for the movement. The nature sound samples in this movement mimic the overlapping form of the music, ensuring the same level of coverage and saturation.

The A section starts with a piano, with each pitch duration being extended by midi effects to create harmony with a single melodic line. This line begins with ascending arpeggiation, producing an uplifting sensation to set the overall tone of the piece. Fractal audio again acts as a pad and pans rhythmically back and forth, creating an insect-like buzz that is both hypnotic and fits into the nature audio. Various bird and frog sounds are initially presented with limited effects, but are time-stretched and reversed as the movement progresses. The frog samples are binaurally tuned with the intention of creating a beat, but as this audio has a natural pulse the results are difficult to detect. While there is no choral synth present, the composite audio creates a sense of a choir that fits into registers around the piano part, both alluding to high angelic singing and a lower tenor choir reminiscent of western religious music. There is a brief moment of tension, created by a dark Phrygian rumbling in the piano before bursting into the main melodic theme. In my mind, this theme (fig. 13) is reminiscent of the melody in the fourth movement of Leo Ornstein's *Piano Sonata No. 4* (fig. 14), albeit significantly simplified.

(fig. 13: Main piano melody in the A section of *Transcendence*.)

(fig. 14: The opening of the fourth movement to *Ornstein's Piano Sonata No. 4*.)

The intervallic and rhythmic motion is similar, particularly when comparing the cascading eight-note phrases. As one of my favorite composers, finding elements of Ornstein's work in my interpretation of the Mandelbrot set is both satisfying and fascinating; what algorithmic rules would be required to emulate specific composers? Could one design a universal algorithm capable of representing each known composer's style without the aid of AI? Towards the end of the first section the static bass gets to shine, as most other instruments except the percussion fade away. This is the only truly static moment in the piece, and has a digeridoo-like quality that provides a brief meditative pause with each repeat of this section.

Images created for the first section include spirits ascending from the Mandelbrot set, the protagonists, various clergy and angel scenes, the main demon theme, and wedding scene. The spirits leaving the Mandelbrot set (fig. 15) is one of my favorite individual scenes in the video, as it shows a clear progression from the base image into the prompt.

(fig. 15: Image from Transcendence of spirits escaping the Mandelbrot set.)

The clergy scenes also do an excellent job of this, where the Mandelbrot variant is visibly transformed into the stained glass of the cathedral. One error encountered here was a floating clergyman, however I liked it so much I changed the prompt to produce variants until I could animate him spinning in the air. The protagonist (fig. 16) is vaguely alluded to in the demon’s theme, only appearing later when overlapping material is repeated at a higher framerate.

(fig. 16: Image of the protagonist from Transcendence.)

I see this character as a gifted human, able to communicate with the many spiritual entities around them. As such, she chooses to live in a cathedral to perform various spiritual services, in this case to officiate the wedding ceremonies. The visuals for the wedding theme are a link to a spiritual past; many of the characters have neanderthal-like facial features and appear to be in ceremonial attire, indicating the importance of this event. Multiple characters can be seen kissing or leaning in for a kiss, most of which are same-sex couples. This is meant to serve as a testament to the timeless sanctity of true love, regardless of external social factors. While none of these scenes are exclusive to any movement due to the overlapping form of the music carrying over into the visuals, the demon theme is also a major element of the second section.

The second section starts just before the five-minute mark, with a fast, steady pulse from the percussion and a bouncy synth mallet functioning as a minimalistic, repetitive melodic line. The panning fractal audio is present here as well, adding contrasting texture to the more spherical, bubbly sound of the melody. The nature audio is heavily slowed down at this point, where the birds sound like prolonged whistling through a tunnel. A separate synth is present that provides long, awkward leaps, which occasional answer the mallet sounds but are most noticeable during the change in mood towards the end of the section. While still chaotic and wild, with many of the former elements still present, this section adds an intense level of darkness in the form of a low layered drone. This is not dissonant necessarily, although it demands attention as both the only static element and the lowest pitch. This brings the listener's attention from high to low and creates a general sense of minor modality.

The B section video begins with the demon theme, entering a black hole portal and teleporting away from the previous scene. The new scene unfolds with colorful groups of various sizes dancing to the chaotic bouncy beat, overlooked by a group that radiate unamused aristocratic energy. Hourglasses also fade in and out here, possibly indicating a transition from a passive state to an active state. The demon appears from a black hole, the dark drone and rapid visuals producing an intense flickering effect indicating that an active state has arrived, and the plot is moving forward. The demon (fig. 17) desires the protagonist but their love is not reciprocated; the protagonist loves someone else.

(fig. 17: Image of the demon from Transcendence.)

The Demon wants the help of the aristocrats, relatives on one side of the protagonist's family, to force her into an arranged marriage with it in exchange for power. They accept, and set out with their followers to bring her to them. In the repeat of the A section, the protagonist and clergy are warned by angels of the demon's intentions, as the demon and protagonists themes overlap much more prominently. The protagonist's theme repeats towards the end at dramatically reduced opacity, indicating the clergy have helped her escape from the area. The hourglass from the B section can be seen here, and shortly after the protagonist flees the aristocrats and their followers arrive.

The C section begins just before the twelve-minute mark, with slowly flowing flute synths breathing the spiritual world into existence, percussion fading in shortly after to keep a steady beat. The flute switches between long held notes, arpeggiations, and oscillations between two distant pitches. A brass synth plays perfect intervals, both in and out of key with the flute to fill out consonant harmonies and create strange dissonances. This perfect intervallic relationship is reminiscent of hunting horns that would signal with fifth intervals, helping to reinforce the strong connection to nature in this section. Extremely slowed birds whistle deeply and wolves howl in slow motion throughout most of the movement, adding a similar texture and even timbre to the flute parts. Towards the end of the section, the wolves howl at normal speed, both signaling a transition back to the A section and that the fight to protect the protagonist has begun.

The video for the C section starts with the protagonist's family (fig. 18) greeting her in the spirit world, their bodies covered in flowers as they emerge through a floral Mandelbrot set.

(fig. 18: Image of the protagonist's family from Transcendence.)

They offer asylum and aid, bringing the protagonist back to their village that exists on the backs of a flock of living birds. The spirits of their tribe perform a ritual, summoning all types of animal guardians to protect the protagonist. An hourglass fades in and out starting around halfway through the section, indicating that their time in the spirit world is nearing an end. The protagonist is returned to her ancestral home deep in the forests a short way from the cathedral.

The animal guardians (fig. 19) persist into the return of the A section visuals, repelling the invaders and restoring peace to the congregation.

(fig. 19: Image of animal guardians from Transcendence.)

During the return of the B section, the aristocrats can be seen vibrating in the same way as the demon casting his spell, indicating that they are the new target as punishment for their failures. While the demon appears as a silhouette in the cathedral during the final scene, his powers alone are no match for the congregation as a whole, and he does not engage them. He watches as the protagonist marries her lover before skulking back to hell alone.

Black Clover

The final movement is titled *Black Clover*, and represents misfortune and death. It's the first piece I wrote for this project, and it is approximately forty-two minutes long. I chose the title Black Clover over other well-known options to represent the theme for two reasons; the first being that the majority of musical elements in this work are repeated four times, once for each leaf. The second, while coincidental, solidified my choice; while I was modifying the formula for the Mandelbrot to generate alternate scores, I came across this clover-like pattern (fig. 20):

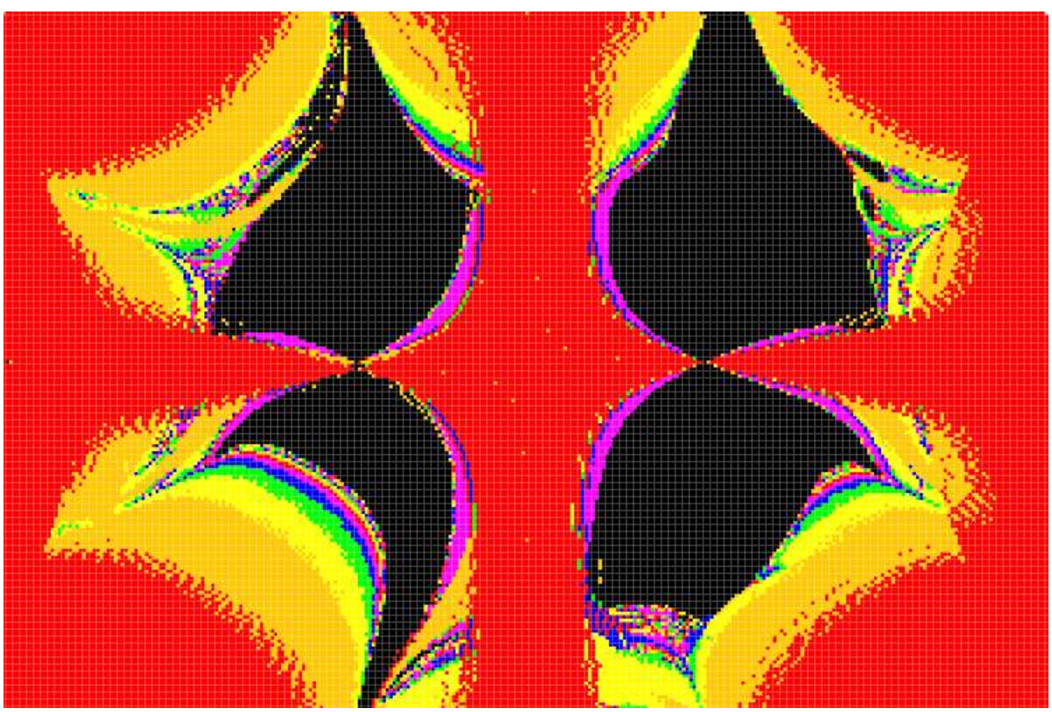

(fig. 20: Black clover image generated from a heavily-modified ETA)

While this image is not used as a score to generate the music, it served as retroactive affirmation for the title. Unlike the other rondo form movements, this piece is cast in an arch form. I wanted to include three separate types of videos for this project, and as this piece had the most visually appealing spectra, I chose the video for this piece to be a scrolling spectrograph.

The A section starts off with the tail end of the removed section fading in, which may indicate that in order for death to exist, life had to exist before it. This deleted section is still present in the piece, however only the time-stretched version is present. Notable elements include nature audio, fractal audio, drums, a pad, and a lead synth. The nature audio for this movement is its own separate track, so what is present here is time-stretched birds and frogs that would have been normal speed in the original A section. The Mandelbrot audio is difficult to hear, however it is faintly audible in the background. The lead synth is mellow, creating a pleasant contrast to the otherwise tense dissonance of this section. This dissonance is a result of harmonic pads from both the removed initial section and the A section overlapping. The melody repeats and has variable length, while the rhythmic elements seem to shift, starting or stopping earlier than anticipated. This shift is an error in my M-Score Reader; however, this error provides an interesting variety for the otherwise fairly static repetitions. This section feels like a combination of many of the qualities of the music that follows; it is not as static and mellow as the B section, but not as intense, mechanical, and choppy as C section.

The B section begins with the introduction of the piano just before the nine-minute mark. The melody in the piano is relatively short lived, giving the elements that would otherwise be supporting the melody to receive the limelight, such as the nature audio and the synth sound wall. The drums seem to fall away here as well, creating a loss of rhythmic sense. While I changed each instrument for each section of this project, the piano is the only noticeable new instrument here. This may be an error, as some instruments chosen may simply not have been powerful enough to cut through the wall of sound. However, while the overall colors used feel relatively unchanged, this section differs from the A section in that there is significantly less activity and texture overall. This creates an oddly tense peacefulness in this section; while it is fairly dissonant and quite dark, the static quality feels less menacing than the A section, and the lack of rhythm allows the user to get lost within the soundscape.

The C section begins around the eighteen-minute mark, and is the most dense and active section of the piece. The exclusive use of four repetitions paired with a new sax synth sound creates some fascinating glitch aesthetics. This is in stark contrast to the prior movement, helping differentiate sections where there are fewer noticeable instrumental changes compared with the other two movements. While I had this section configured to mutate parameters, the outcome resulted in predominantly short rhythmic lengths, creating explosive and glitchy loops.

The spectra for this movement are visually stunning, with these symmetrical geometrical patterns appearing most prominently during the third section. Each pattern within the spectra (fig. 21) repeats approximately four times; however, it is clear that some sections are cut off early or have an additional section that is similarly cut short.

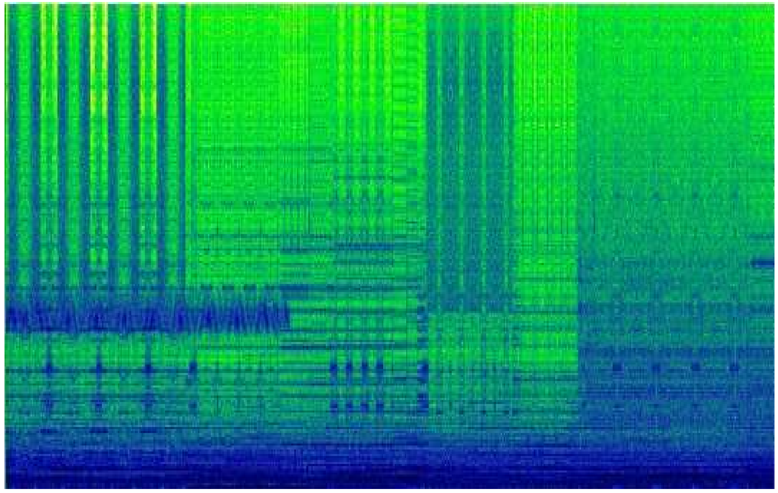

(fig. 21: Spectrograph from Black Clover at 26:33.)

This is likely a visual representation of the uncleared-data bug in M-Score Reader. I find this only serves to make these sounds and visuals more interesting; the work is highly repetitive, so variations in these repetitions are important to keep the material from becoming too static. I believe the broad shape of these patterns are formed by a combination of the Mandelbrot set waveform (fig 22.), which has similar sectional patterns throughout the visible spectra, and ping-pong delay, which would make these patterns symmetrical.

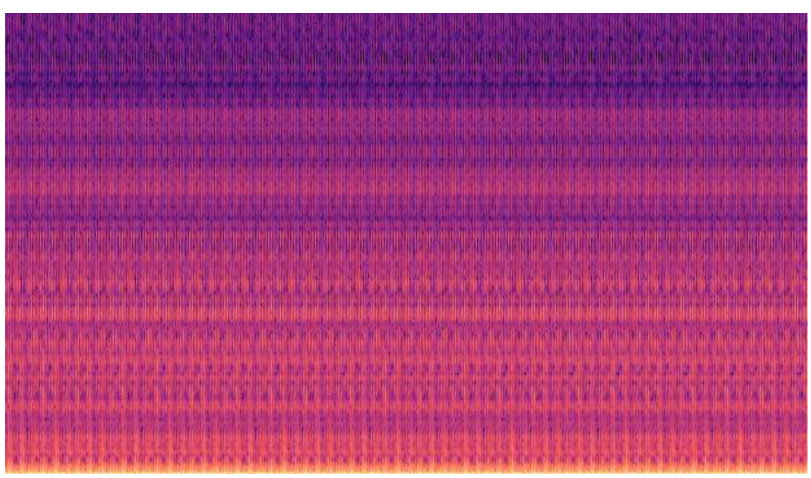

(fig. 22: Spectra of the Mandelbrot set waveform.)

Conclusion

While the M-Score Reader cannot comprehend musical languages it is not programmed to understand, the potential for this type of device is as infinite as the Mandelbrot set itself. With a strong understanding of any specific musical language, I believe it is possible to not only create original works in the style, but also approximately recreate every major work within that genre an infinite number of times. The strong law of large numbers states that, given a non-zero probability with an infinite number of chances to occur, the number of occurrences of that event will almost certainly be equal to its probability.[51] While the Mandelbrot set is chaotic and the data is independent of each other, it is not "identically distributed random" as this law requires.[52] However, the addition of a random variable to the Mandelbrot set data, where each sonic element has the same chance of occurring, would theoretically allow for my device to produce every piece of music ever written. The relationship to the Mandelbrot set would be maintained through the random seed and device settings used to generate it. One practical application for this type of randomization could be encryption, where random numbers are sent to be decoded by the user using the random seed as the key.

Regardless of one's interest in algorithmic music, the most important outcome of this process was to carefully and meticulously define my own compositional ruleset. Through the study of algorithmic composition, I now have the ability to prolifically

---

[51] H. Li et al., "Music as Mathematics of Senses. Advances in Pure Mathematics," Abstract, *APM* 8, no. 12, (2018).

[52] Walter L. Smith, "Renewal Theory and Its Ramifications," *Journal of the Royal Statistical Society*, Series B (Methodological) 20, no. 2 (1958): 245.

create large-scale works in my own. To recreate these works by hand would be nearly impossible, as the data sets are far too expansive.

I believe my method of compartmentalizing these algorithmic processes based on the previous pitch results ensured that my musical identity would be heard. The visual representation of the data sets I used allowed me to make associations based on patterns that would have been significantly more difficult to see in the numerical data, which made routing information to the correct function quite simple. It is unlikely that I would be able to feed my project back through a reverse algorithm to print back the scores that generated it; much of my project uses the mutate function that would make imaging any one individual score impossible. That said, I believe the sonic representation conveys the information of these fractal scores in an aesthetically pleasing way. For example, regardless of score, long, stable areas are represented by repeated phrases, and the color of each coordinate is represented in the motivic material. My primary goal is to generate music that I find beautiful in one way or the other; representation of the source material, while satisfying, will always come second.

Working with AI was incredibly satisfying; being able to give a prompt and receive artistic material within minutes is extraordinarily useful for creating audiovisual art. Not all images rendered perfectly, as many displayed extra extremities, while some figures overlap and fuse together; however, as a surrealist, these types of inaccuracies are a feature and not a bug. The next step in working with AI will be to download and train these tools to work in a way that elevates these surrealistic qualities as opposed to suppressing them.

While understanding all of the intricacies of the Mandelbrot set is beyond the scope of this dissertation, I learned a lot about fractal generation methods and enjoy employing them in my recent projects. It was fascinating to learn how much you could alter the escape time algorithm and still produce imagery that was stable and resembled the original set. I am most curious about the sets which included little or no areas that were stable, but still clearly resembled the Mandelbrot set. Altering the formula in more dramatic ways produced some fascinating imagery as well, and I am particularly impressed with the coincidental finding of the *Black Clover* image.

I hope this work provides a roadmap and further inspiration for other composers to delve deeper into the world of fractals as data sets to drive their algorithmic compositions. The power of both self-similarity and differing areas of stability and instability produces quite natural sounding musical phrases given the right boundaries. At the present time rulesets should be altered to suit the musical language of each individual composer, however I believe we are not far from a singular device that would be able to produce music of any style imaginable.

Bibliography

"Amazing Graphs." YouTube, uploaded by Numberphile, Aug 8, 2019, https://www.youtube.com/watch?v=pAMgUB51XZA&t=2

Barbera, Andre. "Pythagoras." *Grove Music Online*, 2001, https://www.oxfordmusiconline.com/grovemusic/display/10.1093/gmo/9781561592630.001.0001/omo-9781561592630-e-0000022603?rskey=njNRzD&result=1. Accessed 15 April, 2023.

Barbera, Andre. "Euclid." *Grove Music Online*, 2001, https://www.oxfordmusiconline.com/grovemusic/display/10.1093/gmo/9781561592630.001.0001/omo-9781561592630-e-0000009065#0000009065. Accessed 15 April, 2023.

Bergland, G. D. "A Guided Tour of the Fast Fourier Transform." *IEEE Spectrum*, 6, no. 7, (1969), 41-52.

Buelow, George J. "Mizler von Kolof." *Grove Music Online*, 2001, https://www.oxfordmusiconline.com/grovemusic/display/10.1093/gmo/9781561592630.001.0001/omo-9781561592630-e-0000018812?rskey=njpVOQ&result=1. Accessed 15 April, 2023.

Castro, Marfa Antonia and Marfa Jose Perez-Luque. "Fractal Geometry Describes the Beauty of Infinity in Nature." *Proceedings of Bridges Conference*, (Alhambra, Spain: 2003), 407-414. https://archive.bridgesmathart.org/2003/bridges2003-407.pdf

Drakopoulos, V. et al. "An Overview of Parallel Visualisation Methods for Mandelbrot and Julia Sets." *Computers & Graphics* 27, no. 4 (2003), 635-646. https://www.sciencedirect.com/science/article/pii/S0097849303001067

"'Fractal' mathematician Benoit Mandelbrot dies aged 85." *BBC News*, 17 October, 2010, https://www.bbc.com/news/world-europe-11560101.

"Fractal Music - Image Sonifications (II) - Mandelbrot Set." YouTube, uploaded by Gustavo Diaz-Jerez, Feb 20, 2009. https://www.youtube.com/watch?v=3Br57CsDAFw

Glenny, R W et al. "Applications of fractal analysis to physiology." *Journal of Applied Physiology* 70, no. 6 (Bethesda, Md. 1991), 2351-67.

Hoffmann, Peter. "Xenakis, Iannis." *Grove Music Online,* 2001, https://www.oxfordmusiconline.com/grovemusic/view/10.1093/gmo/9781561592630.001.0001/omo-9781561592630-e-0000030654. Accessed 17 April, 2023.

Hope, Catherine Anne and Michael Terren. "The Possibilities of a Line: Marking the Glissando in Western Art Music." *Proceedings of the International Conference on Technologies for Music Notation and Representations*, (ARU: 2016). https://www.semanticscholar.org/paper/THE-POSSIBILITIES-OF-A-LINE-%3A-

MARKING-THE-GLISSANDO-Hope-Terren/6c66b629477f9db48f670d988ae766479c4c8e44

Li, H. et al. “Music as Mathematics of Senses. Advances in Pure Mathematics.” Abstract. *APM* 8, no. 12, (2018) 845-862. https://file.scirp.org/Html/2-5301555_89311.htm

Lourenco, Bruno F. et al. “L-Systems and Evolutionary Techniques.” *Proceedings of the 6th Sound and Music Computing Conference*, (SMC: 2009). https://www.researchgate.net/publication/228936738_L-Systems_Scores_and_Evolutionary_Techniques

“Mandelbrot set.” *Merriam-Webster Dictionary*, https://www.merriam-webster.com/dictionary/Mandelbrot%20set. Accessed 18 May. 2023.

“Mandelbrot Sound [seizure warning].” YouTube, uploaded by Craige Hales, Jan 12, 2019. https://www.youtube.com/watch?v=2wJY4BZhQzc&list=LL&index=423

Manfred R. Schroeder, “Auditory paradox based on fractal waveform” *The Journal of the Acoustic Society of America* 79, no. 1 (January, 1986): 186–189. https://pubs.aip.org/asa/jasa/article/79/1/186/638324/Auditory-paradox-based-on-fractal-waveformLetters

Mathews, Max V. “The Digital Computer as a Musical Instrument.” *Science* 142, no. 3592 (Nov. 1, 1971), 553-557. https://www.jstor.org/stable/1712380?origin=JSTOR-pdf.

Merchant, Jo. “A Journey to the Oldest Cave Paintings in the World.” *Smithsonian Magazine*, (January 2016). https://www.smithsonianmag.com/history/journey-oldest-cave-paintings-world-180957685/

Monro, Gordon. “Fractal Interpolation Waveforms.” *Computer Music Journal* 19, no. 1 (1995): 88–98. https://www.jstor.org/stable/3681302?seq=9

Orozco Perez, Hector D et al. “Binaural Beats through the Auditory Pathway: From Brainstem to Connectivity Patterns.” *eNeuro* 7, no. 2, (Mar. 2020). https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7082494/

Richardson, John. “Video [music video].” *Grove Music Online*, 2020.

Smith, Walter L. “Renewal Theory and Its Ramifications.” *Journal of the Royal Statistical Society*. Series B (Methodological) 20, no. 2 (1958): 243–302. http://www.jstor.org/stable/2983891.

The Editors of Encyclopaedia Britannica. "Benoit Mandelbrot." *Encyclopedia Britannica*, https://www.britannica.com/biography/Benoit-Mandelbrot. Accessed 18 May 2023.

Venkatesh, T.S. Sachin and Vishak Vikranth. “Investigating the Relation Between Chaos and the Three Body Problem.” *arXiv: Chaotic Dynamics* (2020). https://www.researchgate.net/publication/343986739_Investigating_the_relation_between_chaos_and_the_three_body_problem

Xenakis, Iannis. *Formalized Music: Thought and Mathematics in Composition*. (Stuyvesant, NY: Pendragon Press, 1992).