COMPOSING HOLOCHORIC VISUAL MUSIC:
INTERDISCIPLINARY MATRICES

by

Michael Rhoades

An Exegesis

Submitted to the Faculty of Virginia Tech University

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

Schools of Visual Arts, School of Performing Arts, and School of Computer Science

Blacksburg, Virginia

November, 2020

THE VIRGINIA TECH GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL

Zachary Duer, MFA - Chair
    School of Visual Arts
Dr. Sang Won Lee
    School of Computer Science
Dr. Charles Nichols
    School of Performing Arts
Thomas Tucker, MFA
    School of Visual Arts
Dr. Nicholas Polys
    Advanced Research and Computing

Approved by:

Dr. Karen DePauw
    Vice President and Dean for Graduate Education
Dr. William Huckle
    Associate Dean of the Graduate School

Dedication

There are many individuals to whom this dissertation is dedicated. Certainly, it is dedicated to my beloved wife, partner, and best friend Janet. She has endured numerous trying effects in putting her life on hold to allow the pursuit of this work. The gratitude I feel toward her is immeasurable. It is also dedicated to my wonderful daughter Madison and to my wonderful son Jordan. They have each, in their own way, grown into capable, active, and dynamic human beings of whom I am so very proud. This dissertation is further dedicated to several friends and family members who have stuck by my side during my absence over the past several years. It is my steadfast wish that all of the above mentioned receive the extreme gratitude I extend in their direction. I could not have approached this work without their understanding and support.

Lastly, it is realized that we human beings experience our lives through the lenses of our particular perspectives. Therefore, it must be mutually understood here at the beginning that it is in accord with this realization that these perspectival views are collected and assembled. A clear perspective may lend objectivity to perspectives in general and, if it is possible, allude to that which is beyond them. Such is the intent of this work. However, though the notion of denoting a perceived truth is tantalizing, no matter how clearly symbolized by language and thought, it remains elusive in that we seem to exist within a constantly changing and evolving system. This ceaseless moment-to-moment rearrangement of the reality we seem to perceive renders any apparent realization in one moment impertinent in the next.

Therefore, the goal here is to achieve a generalized view from a limited position in space/time in order to contribute a voice to a very broad chorus of views… past, present, and future. It cannot be determined that such an endeavor consists of any inherent value, per se. We nonetheless persevere in attempting to organize the fathomless into some sort of a linguistically based (symbolic) structure. Such a structure is willingly relegated to the eventuality of failure in essence if for no other reason than the consistency of life occurring. (In this context the terms time, space, change, and life can be considered synonymous.) As such this work can be considered little more than the fleeting vision of the one inspired to it. We seem to possess a brief moment in that which we construe as space/time and so we must do something. This is an example, in an infinite sea of examples, of it being occupied and rearranged.

In summary, this work is ultimately dedicated to the ever-shifting sands of space/time to which each of us contributes but a momentary minute grain.

ACKNOWLEDGEMENTS

This research project would not exist in its current form without the steadfast and insightful input from Professor Zach Duer. His extremely perceptive and steadfast input and discussion regarding the composition of this document and the associated artifacts has been nothing short of excellent. He spent countless hours guiding me along the process. In addition to Professor Duer, Dr. Sang Won Lee, Dr. Charles Nichols, Dr. Nicholas Polys, and Professor Thomas Tucker have each contributed to the thought processes and to the creativity involved in the writing of this document. Through their extremely well-schooled minds they have provided inspiration and clarity to the compositions that are the primary constituents of the research. Their critiques were consistently on point, never negative, and always provided with the goal of excellence at heart. They have further instilled an ever-deepening self-critical perspective, yet one tempered by good humor thus enabling a reduction in the self-imposed pressure involved in this endeavor. The gratitude that I feel for their encouragement and faith cannot be over stated. Interacting with them through these past years has been a life-altering experience, one that I will carry throughout the remainder of my tenure here on planet Earth.

Dr. Simone Paterson, who was the original chair of my committee, should also be mentioned. She retired just before my first year in graduate school, however she was extremely instrumental, during the preceding year, in editing and critiquing my iPhD proposal. Dr. Paterson accompanied me to its presentation to the Committee for Graduate Studies and Policies, which was the governing body that ultimately approved my proposal. Her knowledge of practice-based research provided a solid and well-informed foundation in terminology and implementation of this pursuit, without which it would not exist in its present form.

Lastly, I would like to thank Tanner Upthegrove, an ICAT media engineer working in the Cube at Virginia Tech. Several discussions with Tanner during my graduate studies spurred new ideas and insights in the area of the spatialization of sound. I have learned much from Tanner’s extremely creative mind through the process, for which I am eternally grateful.

Thank you to each of you from the bottom of my heart.

TABLE OF CONTENTS

LIST OF FIGURES

ABSTRACT

With a lineage originating in the days of silent films, visual music, in its current incarnation, is a relatively recent phenomenon when compared to an historically broad field of creative expression. Today it is a time-based audio/visual territory explored and mined by a handful of visual and musical artists. However, an extensive examination of the literature indicates that few of these composers have delved into the associable areas of merging virtual holography and holophony toward visual music composition. It is posited here that such an approach is extremely rich with novel expressive potential and simultaneously with numerous novel challenges. The goal of this study is, through praxis, to instantiate and document an initial exploration into the implementation of holochory toward the creation of visual music compositions.

Obviously, engaging holochoric visual music as a means of artistic expression requires an interdisciplinary pipeline. Certainly, this is demonstrated in merging music and visual art into a cohesive form, which is the basis of visual music composition. However, in this study is revealed another form of interdisciplinarity. A major challenge resides with the development of the means to efficiently render the high-resolution stereoscopic images intrinsic to the animation of virtual holograms. Though rendering is a challenge consistent with creating digital animations in general, here the challenge is further exacerbated by the extensive use of multiple reflections and refractions to create complexity from relatively simple geometric objects. This reveals that, with the level of computational technology currently available, the implementation of high-performance computing is the optimal approach.

Unifying such diverse areas as music, visual art, and computer science toward a common artistic medium necessitates a methodological approach in which the interdependency between each facet is recognized and engaged. Ultimately, a quadrilateral reciprocative feedback loop, involving the composer’s sensibilities in addition to each of the other facets of the compositional process, must be realized in order to facilitate a cohesive methodology leading toward viability.

This dissertation provides documentation of methodologies and ideologies undertaken in an initial foray into creating holochoric visual music compositions. Interlaced matrices of contextualization are intended to disseminate the processes involved in deference to composers who will inevitably

follow in the wake of this research. Accomplishing such a goal is a quintessential aspect of practice-based research, through which new knowledge is gained during the act of creating. Rather than formulating theoretical perspectives, it is through the praxis of composing holochoric visual music that the constantly arising challenges are recognized, analyzed, and subsequently addressed and resolved in order to ensure progression in the compositional process. Though measuring the success of the resultant compositions is indeed a subjective endeavor, as is the case with all art, the means by which they are achieved is not. The development of such pipelines and processes, and their implementation in practice, are the basic building blocks of further exploration, discovery, and artistic expression. This is the impetus for this document and for my constantly evolving and progressing trajectory as a scholar, artist, composer, and computer scientist.

CHAPTER 1. INTRODUCTION

1.1 Introduction

This dissertation is intended to provide an interdisciplinary phenomenological and methodological framework and contextualization toward the creation of holophonic and (pseudo) holographic visual music compositions. Pairing holophons and pseudo holograms, diffused and projected in 3D/360 venues, creates an evolutionary progression of counterpoint in the field of visual music composition that approaches a multidimensional spatiotemporal experience. Intricately interwoven methodological ontologies involving generative audio algorithms, aleatoric visual procedures, and high-performance computing techniques pose a confluence of process that facilitates unique aesthetical and formal stratifications of structure in support of this approach. The interplay between the composer's sensibilities, audio and visual indeterminacy, and high-performance computing, necessitates a novel quadrilateral and reciprocative feedback loop actively engaging each constituent. Therefore, and due in part to lengthy rendering cycles, this work expands the traditional tone row methodology of musical composition, wherein thematic elements are constrained and then developed to the greatest feasible extent, to foster a novel counterpoint of intermedial foci.

The practice-based research (PBR) conducted to achieve these new approaches was derived in composing several works that interlace these elemental components. These artifacts consist of a set of seven, aesthetically related, holophonic and holographic visual music compositions. The title of this project is "Antitheses". The holographic visual aspects of these compositions are designed to be presented in an Extended Reality (XR) cinematic setting in which each member of the audience is wearing a head mounted display (HMD) or seated within a proprietary 360-degree stereoscopic visual projection system (SVPS). The holophonic musical aspect of the compositions is synchronously presented over a high-density loudspeaker array (HDLA). It is ultimately the amalgamation these artifacts that form the basis for artistic originality and to the contribution of new knowledge.

Though definitions for each tend to vary, here holograms are considered projections of three-dimensional virtual visual objects that appear to occupy a common three-dimensional space as the viewer. Holophons are considered their auditory counterparts. Holochory, in this sense employs stereopsis, or binocular perception, which is how the brain perceives three-dimensionality. With this paradigm, audio/visual expressions appear to occupy the space contiguous to the audience.

Visual music, as expressed in the pre-recorded moving image, began as an art form in the early 20th century (Kanellos). However, to date there appears little in the literature to suggest previous investigations consisting of the implementation of holophons and holograms toward the medium. Due to their common three-dimensional nature, it is here proposed that merging holophons and holograms is an obvious compositional approach. However, this territory appears relatively unexplored in terms of research and creative practice. This is perhaps due, in part, to the complex matrices of techniques required to achieve such a union… not the least of which involves high performance computing (HPC). Herein lines of interconnectivity between holophons, holograms, HPC, and numerous other constituents are discovered, exposed, developed, and documented. These lines dynamically extend the scope of each aspect, generally, and are intended to expand the expression of visual music specifically.

Here, in Chapter 1, fundamental information is disseminated before drilling down into the details in subsequent chapters.

1.2 Practice Based Research

Dr. Linda Candy, a leading expert in the area of practice-based research, states:

> Practice-based Research is an original investigation undertaken in order to gain new knowledge partly by means of practice and the outcomes of that practice. Claims of originality and contribution to knowledge may be demonstrated through creative outcomes which may include artefacts such as images, music, designs, models, digital media or other outcomes such as performances and exhibitions. Whilst the significance and context of the claims are described in words, a full understanding can only be obtained with direct reference to those outcomes (Candy).

Candy later states, “If a creative artefact is the basis of the contribution to knowledge, the research is considered practice-based” (Candy).

The research delineated in this exegesis is **practice based**. As a scholar of praxis, the foci of my research reside in the methodologies and artifacts that emerge from the creative process. The contribution of new knowledge is primarily based upon the (subjective) results, the artifacts, and this exegesis that provides an associated context. Novel questions naturally arise in the midst of any truly creative endeavor, which require novel solutions in order to move forward toward a sense of completion. These solutions bring to light not only new processes and techniques but, perhaps more importantly, they drive the creative process in previously unimagined directions. It is the culmination of this activity that results in new knowledge, both personal and societal, and the artifactual representation of it.

The act of composing can be viewed as one of problem solving (Laske). Certainly, when delving into unexplored regions in any discipline each step forward exposes novel challenges and requires novel solutions. It is the journey undertaken that necessitates the research. For this reason, composing holochoric visual music is a quintessential example of the need for a PBR approach since there are few, if any, previous examples to call upon.

It is important to keep in mind that theorizing based upon practice is very different than applying theory to practice (Bolt). In such an investigation it is through practice that a theory is derived, implemented, and then evaluated for efficacy. In actuality it is through countless cycles of praxis, theory, implementation, and evaluation that completion is achieved. It should be further kept in mind that the evaluation step in this approach is crucial although in many ways subjective. This is a major departure from quantitative or qualitative research where the results of testing are objectively evaluated, which may be described as evidence-based research. Conversely, PBR can be described as practice-based evidence (Sullivan).

The practice of creating a time-based visual music composition requires countless split-second iterations of this practice, theorizing, implementation, and evaluation cycle… many of which are barely perceptible. This renders the documentation of the creative process elusive at best. While in the act of composing, the composer makes countless such decisions toward the “rightness” of the composition. The results of decisions that are accepted are woven into the composition. Those that are not are discarded and others are attempted. Often, when a decision is fruitful, variations of it are instantiated, which contributes interest and cohesiveness to the composition. This decision-making process is non-linear in nature and occurs on multiple time scales during the act of composition. For example, a combination of events might seem “right” when creating a musical phrase in isolation and yet when later placed within a contrapuntal section, where the phrase becomes one of several layers of phrases occurring simultaneously, it might cause an unacceptable conflict and thus requires an alteration or exclusion.

As can be witnessed through this isolated example, the creative process is extremely complex and therefore challenging to document. This point is exponentially amplified when working within an interdisciplinary environment. However, all is not lost. For we are able to allude to the creative process by documenting many challenges experienced and in describing techniques developed to resolve them. These challenges become evident not only during the creative process but also, and perhaps more so, by experiencing the resulting artifact. In this case material not included in the final version of a composition is as important as that which is.

These and other contingencies are the basis for the decision to implement a PBR approach for this work.

1.3 Influences

That we stand on the shoulders of giants is a statement commonly attributed to Isaac Newton. This perspective could not be more appropriate than when considering the research herein described. We human beings are influenced by everything we have experienced and ultimately by everything that has ever been experienced. This research does not exist in a vacuum but instead is simply a minute contribution to the ever-evolving and -expanding human consciousness. There are many direct influences that contribute to it.

Musically, the inspiration for this research comes from several directions. Major influences come from the creative work, writing, and conversations with two very close friends, colleagues, mentors and mentees… Dr. Otto Laske and Dr. Sylvia Pengilly. Dr. Laske is a pioneer in the field of, among several other areas, computer-generated music. Dr. Pengilly is a dynamic force in the area of visual music and a pioneer in electroacoustic music and its relationship to movement. They have influenced and inspired this work more than any other individuals. However, there are many others who are directly connected to it.

Other musical aesthetic and theoretical influences are Natasha Barrett, Richard Boulanger, John Chowning, Francis Dhomont, John Ffitch, Jean Claude Risset, Curtis Roads, Pierre Schaeffer, Dennis Smalley, Morton Subotnick, Iannis Xenakis, and Steven Yi. Additional influences include musical groups that originated in the 1960's and 1970's that are now referred to as the Progressive Rock bands. These include Gentle Giant, Genesis, Yes, Frank Zappa, Jethro Tull, Pink Floyd, King Crimson, Mike Oldfield, Kansas… the list goes on. Their work incorporated rock, jazz, experimental, and classical influences with synthesizers, which at the time was unprecedented, into complex and intricately developed long form compositions. They were also true pioneers.

In the world of visual music, in addition to Dr. Pengilly, are Dennis Miller and Brian Evans. Each of whom has had a direct influence on this work. Also included is a wide variety of artists, all of whom worked with light in truly innovative and creative manners. The NYC Hoffman School, the Liquid Light Lab, Henri Matisse, Joan Mitchell, Claude Monet, Camille Pissarro, Milton Resnick, Gerhard Richter, Paul Friedlander, Char Davies, and Thomas Wilfred. This list only begins to scratch the surface.

In the area of computer science influences are derived from the teams at UCSD that developed the ROCKS supercomputing cluster and from William Robinson and his master's thesis as a student at Clemson University. More basically, inspiration was derived from the Linux, Unix, Windows, and Apple computing environments. Each of which contributed dynamically to this work. Certainly, knowledge of the Beowulf clustering paradigm from the mid 1990's is an influence.

Philosophically and spiritually there are Adyashanti, David Hawkins, J. Krishnamurti, and Eckhart Tolle. Their work and teaching in the area of the quiet mind and creative living are a constant source of inspiration.

There have been various influences in the area of physics and the sciences in general such as David Deutsch, Richard Dawkins, Albert Einstein, Brian Greene, Stephen Wolfram, and Evan Harris Walker. Their ability to convey deep physics-related topics to laymen is a cherished gift and inspiration.

Lastly, in the area of creative writing are Isaac Asimov, Sir Arthur Conan Doyle, Robert Jordan, Stephen King, Steven Pinker, Phillip Pullman, Kim Stanley Robinson, Gene Roddenberry, J.R.R. Tolkien, and Kurt Vonnegut to name a few.

There are perhaps more influences omitted here than included but this listing provides a sampling of those who directly influenced and inspired this work.

1.4 Motivations

From my perspective…

Consciousness in this multiverse is ever evolving. Here the words "life" and "change" are considered synonymous. We seem to exist within a continuum of events, which are at once minute and infinite in scope. Historically speaking, we human beings have lived in a world immutably consisting of three spatial dimensions and one linear temporal sequence for eons. We are genetically entrained toward experiencing our lives from this conditioned perspective, which was passed down through countless generations. However, recent empirical evidences suggest the discovery of non-linearity in the folds of this spatiotemporal environment. Observable evidence derived from experiments in quantum physics and astrophysics are revealing greatly expanded views that are creating chinks in the proverbial armor of our current experience of reality. It is therefore suggested here that humanity is on the verge of a dynamic shift in consciousness (Rhoades, 2013).

Heightened levels of awareness are often brought about by events that jolt us out of our habitual, day-to-day perceptual stances. My work is intended to catalyze such awareness in those who experience it. Non-linearity in spatiotemporal terms implies space that is not contiguous and time that is not sequential. The later might be described by the manner in which Kurt Vonnegut's protagonist in the Slaughter House Five, Billy Pilgrim, becomes unstuck in time" (Vonnegut). In this fictional novel, Pilgrim jumps back and forth through the events of his lifetime thus (perhaps) depicting time from a perspective of being stacked in a vertical manner such that any time might be accessed in any moment. Though a fictionalization, there may come a point in human history where this is not so farfetched. Who can say? The currently predominate human mindset finds it challenging to imagine, for instance, a particle that spins in more than one direction simultaneously although this is what physicist are discovering as they delve ever deeper into quantum mechanics. It is this rapidly evolving perspective that my work is intended to allude to if not to portray.

It is here suggested that we human beings are capable or much more than we commonly exercise in every aspect of our lives. As a simple example, we could perceive that human consciousness has greatly evolved in the past three hundred years. The innumerable advancements in our everyday experiences provide undeniable evidential support of this. However, the music most commonly composed and listened to today is still primarily based upon theoretical underpinnings that are three hundred if not three thousand years old. This belies a mind that has evolved in potential and yet has not embraced it. We are capable of much more and my work is intended to extend beyond these current limitations and to act as a departure from them. The success of this intention is in the eye, and ear, of the beholder… as it were.

Heightened states of awareness directly perpetuate evolutions of consciousness. States of consciousness, which are generally undulating in various directions, are expressed and experienced inwardly and outwardly, which are two sides of the same coin. So, more than simply expressing breakthroughs that move us forward technologically, I am interested in those that move us upward in all aspects of what is it to be human. For it is through self-growth that society grows and society cannot grow unless the individual grows.

Atop these intentions is a consistent reiterative attitude of expansion and refinement. It is the breathing of my essence to learn and to grow throughout my lifetime. It is this same breathing that is the impetus for self-expression. The drive for the beauty of new expression requires that expansion and refinement continue, which once again leads to new expression. An interruption of these cycles leaves death inevitable. Their continuance makes life livable. It is this evolution that makes creativity possible. Perhaps ultimately my goal is to express the inner beauty of this movement within.

An overarching theme in this work is that it be experienced as a fully immersive immaterial environment. Holochory lends toward this goal. It provides an environment that can be directly sensed with the eyes and ears, in a manner similar to our normal realities, and yet act as antitheses of them. Here an opening into a non-linear, yet corporeal, domain is provided… a domain that is rarely accessible from our normal states of consciousness.

In tandem with holochory, employing indeterminacy as a methodological approach lessens the conscious volition of the composer during the creation of the audio/visual material of which this project consists. It is suggested that this state of mind translates directly through the compositions to the audience thus reducing, if not eliminating, the conscious volition of the experiencer. It is posited here that in this loosening of our fixated levels of consciousness we achieve an opening toward their transcendence.

The hope is that when the experiencer temporarily steps into this nonlinear experience it catalyzes heightened states of awareness. We human beings commonly become fixated upon very limited perspectives of our selves and our realities. While these perceptions may be construed as stable and secure, truly the nature of our being remains virtually unknown to us. Therefore, it is in departing the reality we think we understand, if only for a few moments, that we may find the freedom and the energy to wake from this slumber of false security and assuredness to embrace our lives for the unfathomable mysteries they entail.

1.5 Visual Music

"*visual music should be visible music, music made visible or, to expand the term an equal and meaningful synthesis of the visible and audible, and is therefore ultimately its own art form* (Dahn in Lund & Lund in Watkins)". Visual music, in its current incarnation, is a relatively recent phenomenon with humble beginnings in silent films and other projections of the early 20th century (Kanellos). Then as now the visual subject matter is most commonly abstract or non-representational. It is generally situated within the Modernist perspective of art for art sake (Watkins). Visual music compositions usually mean nothing beyond that which is experienced. Any perceived meaning ascribed is that projected by the experiencer. Additionally, these compositions commonly consist of geometrically based abstract or non-representational moving objects and environments, which further lends to the apparent lack of narratively derived meaning. Many composers compare visual music to listening to instrumental music, which is experientially oriented, and consider this approach similar.

In his book, "Visual Music Masters", Adriano Abbado, traces the historical lineage of visual music to long before current views based upon recorded moving images. Abbado finds the essence of the art form beginning in ancient Greece, in the sixth century B.C.E. Pythagoras and his research with the numerological unity observed in nature found numerical similarities between the behaviors of light and sound. He termed this unity “the divine harmony”. Plato expanded upon his research in his dialogue “Timaeus”. Aristotle later mentions the relationship between color and harmony in his work. Throughout the subsequent centuries to the present, the perspective of this divine order appears in the work of various artists and scholars (Abbado). Relationships between sight and sound are essentially the basis of visual music today.

The artifacts associated with the research disseminated here, consist of a set of seven related, yet varied, visual music compositions. The title of this collection of compositions is “Antitheses” though each individual composition is titled as well. They are: His Dancers Three, Shards of Oil, Petal and Branch, Twisted Chambers, The Crimson Castle (An Escher Multiverse), Bending Glass Walls, and 03b_0208.

For Antitheses, an attempt was made to provide equal perceptual value between the audio and visual components. This is in many ways an unrealistic aim since we human beings perceive sound and image in very different manners. In my experience, from observing audiences of my own work in addition to that of others, when audio and visual work are presented together, the experiencer generally and primarily focuses upon and is consciously aware of the visual components and the audio tends toward background perception. On a less conscious level, the music has a powerful impact upon how the imagery is perceived and at points in the composition where the audio and visual are in sync, the moment is more heavily imprinted in the memory of the experiencer (Watkins). Nonetheless it was one aim of the research to compensate for this apparent disparity in perception.

Achieving such a goal is a challenge with any audio/visual work and is perhaps more so with holophonic and holographic work due to the uniqueness of the visual aesthetic. One technique, discovered during the practice of composing, to address this dilemma was to create sections in the composition where the visual component fades to translucency or is not present at all during brief sections while the music continues to evolve. Techniques such as these will be described in greater detail later in these pages. Their mention here is intended to briefly demonstrate the extent to which maintaining the dominance, or at least equality, of the music in the composition is important and to show that such a constraint is necessary in order to maintain, or at least attempt, the desired aesthetic.

In pursuing the development of interactions between the auditory and the visual for this project, several approaches were implemented. Primarily, a choreographic perspective was employed wherein the objects and environments are related in terms of movement and spatial location. Relationships ranging from synchronous to contrapuntal to completely disparate are determined as aesthetical choices between musical and visual objects and environments. Another approach involves parametric relationships such as timbre, amplitude, and frequency verses color, brightness, and hue. Ultimately, sound existing as light and light existing as sound is the desired effect. Creating these relationships is key to visual music expression.

1.6 Holophons and Holograms

Though much terminological ambiguity exists, true holograms and true holograms are directly related to the properties of waves. As demonstrated in Fig. 2 below, holograms are commonly created by recording interference patterns between laser beams (Kostuk). As shown in Fig. 3, holophons, also known as acoustical holograms, consist of the interactions of wavefronts, which are the points at which waves, propagating through the medium of air, converge (Roads, 2015). The term “holochory” is an umbrella term that includes holophony and holography and applies to any given physical field in which wave interaction is considered (Illényi).

Ultimately, holograms and holophons afford artists and composers mediums within which to create objects and environments, and subsequently artifacts, that closely align with the perceptual reality of human experience. This allows for deeply immersive experiences that place the viewer or listener *within* the artifact as opposed to experiencing it as being projected toward them from single or multiple positions. Further, it is reasonable to hypothesize that, similar to the manner in which a two-dimensional photograph alludes to the third spatial dimension of depth, a holographic or holophonic artifact might allude to a fourth spatial dimension.

Visual interference patterns are clearly demonstrated in the common double slit experiment in physics, Fig. 1. With it, a single beam of light, for instance from a laser beam, is projected through two slits in a barrier. Instead of two slits being projected on a surface beyond the barrier, as one would expect from a particle-like perspective of light, many slits and shadows are projected demonstrating a wave-like characteristic. Interference patterns, as waves of light pass through the two slits, create numerous bright and dark areas on the projection surface. As the troughs and peaks of the waves converge, they add, subtract, or cancel each other (Greene).

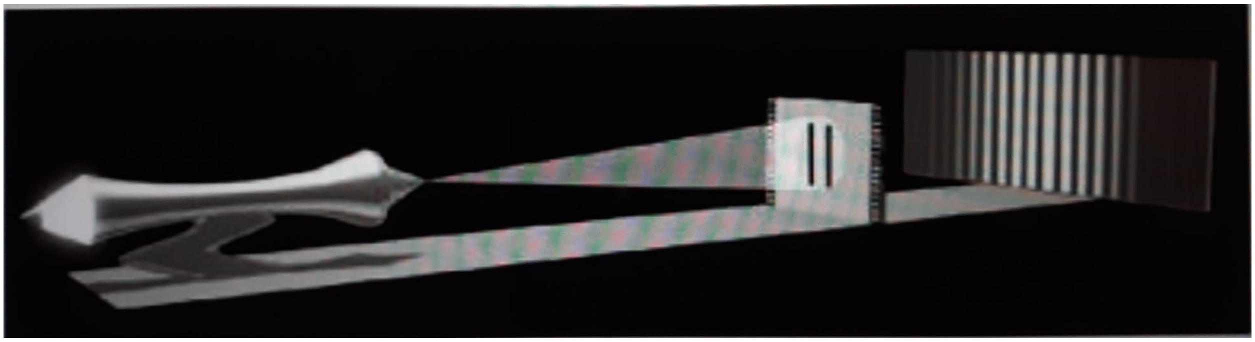

Figure 1 – Double-Slit Experiment Demonstrating Wave Interference (Greene)

"... by recording the interference pattern generated when the original recorded signal of a sound is combined with an inaudible reference signal, a genuinely realistic noise is produced. This is holophonic sound." (Howe). This is one perspective however, the Ambisonics and the wavefield synthesis paradigms are considered by many acousticians to be holophonic in nature due to the interaction of wavefronts inherent to their characteristics. For example, the Ambisonics paradigm is used to create spherical sound fields to which audio events are directed. Each loudspeaker receives a weighted sum of all the speakers in a venue, thereby increasing the size of the listening sweet spot and the quality of sound localization. Here, wave interference patterns are once again the primary determining factors. Perhaps the "how" of the creation of wave interaction is not as important as is the essence of that interaction in the association of the term "holophony" with a given auditory artifact.

"A hologram is a sculptural casting of the light waves." (Chaudhari, et al). Figure 2 depicts the basic process of recording a transmission hologram. To exhibit it, the process is simply reversed. Figure 3 illustrates acoustical wave interference patterns derived from 2nd order Ambisonics and a 6-channel sound system.

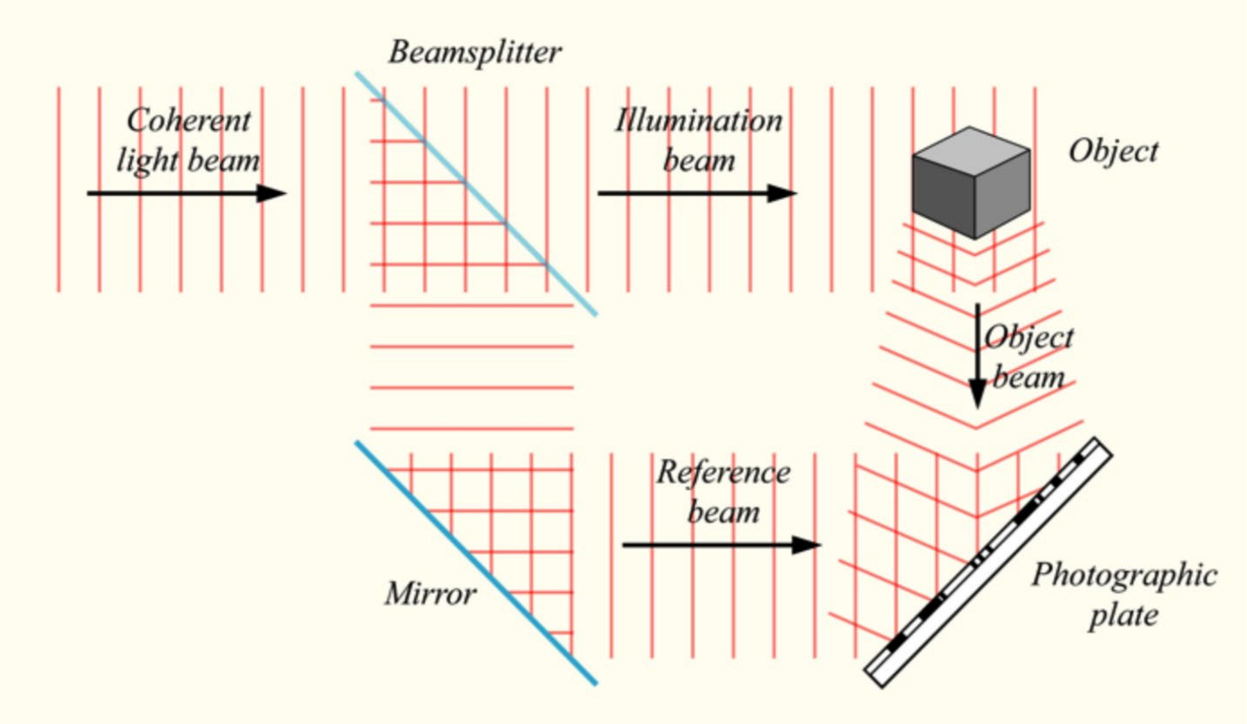

Figure 2 – How a Hologram Works (Chaudhari, et al)

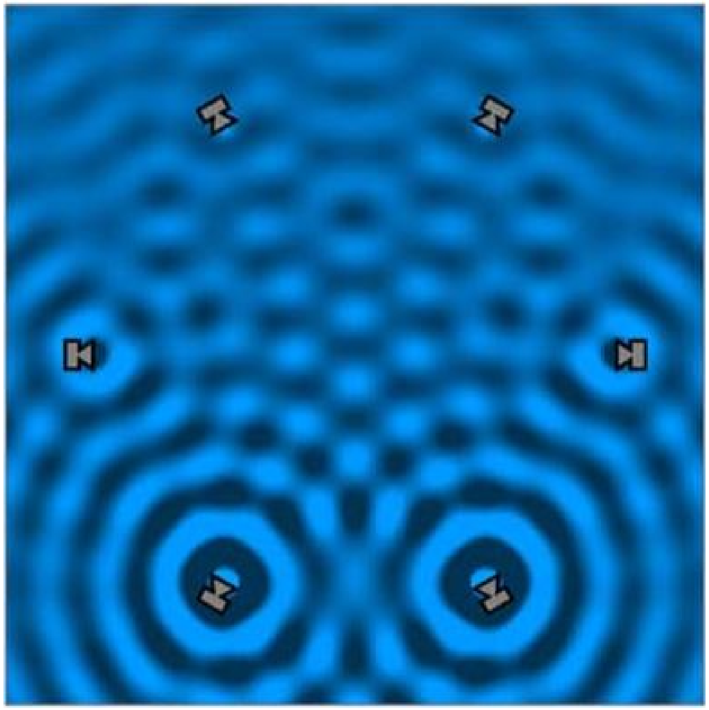

Figure 3 - Mathematically Derived Sound Field Interference – 2nd Order Ambisonics (Spors, et al)

As stated earlier, holophons are here considered auditory objects and environments that are perceived to occupy the same three-dimensional space as the observer. Similarly, holograms are their visual counterpart. Historically, in both the musical and visual arts, images and sound objects are generally projected from the front of the performance space toward the audience. This causes a separation, a dimensional gap, between the audience and the performance or presentation. Here the goal is to create an immersive environment that completely envelops the audience thus eliminating this gap. In doing so, the audience becomes an intrinsic aspect of the composition.

Spatialization of sound, via HDLAs, is utilized to optimally approach the goal of holophony. Instead of actually placing a loudspeaker in every desired position within a space, known as a point source approach, techniques are implemented to provide the perception of a given sound occupying the space between the loudspeakers. Techniques such as convolution reverb (CR), vector-based amplitude panning (VBAP), and high order Ambisonics (HOA) are here employed in harmony to advantage the composer toward this goal. Manipulation of amplitude relationships between large numbers of loudspeakers and further using them to create fields of perception is key to the perception of locality of a given sound object or environment. Providing a realistic relationship to distance perception is also a major factor.

Stereoscopic imaging is implemented to achieve a similar result with visual images. In order to control and determine the perceived depth of field necessary to provide the illusion of a virtual visual object or environment as occupying the same space as the observer, a left- and right-eye perspective image is rendered for every single frame of the video. Then, using an HMD or other less available proprietary 360-degree stereoscopic projection screens, the left- and the right-eye images are alternately turned off and on in the appropriate eye at a rate of ~60 times per second. This fools the brain of the observer into perceiving it is seeing both images simultaneously, as it would outside the virtual world, thus providing depth of field.

Composing visual music with holophons and holograms from a holistic perspective, as tightly unified perceptual objects, is the main thrust of this research. Though each has been extensively researched separately, this union is a relatively unexplored territory that presents a novel frontier.

1.7 High Performance Computing

In order to render the stereoscopic and high-resolution images required to adequately present the holograms in these seven compositions, high performance computing (HPC), which involves parallel computing, was required. HPC clustering can be traced back to the mid 1990’s and the first Beowulf clusters. Today the term Beowulf is in many ways a generic name for HPC clusters constructed from commercial level computers.

To facilitate the need to render numerous complex stereoscopic images for this project, a 48-node HPC cluster was built, configured, and employed by the author. Basically, the cluster consists of six compute nodes, each with eight virtual CPU cores. They were built based upon high-end gaming motherboards, CPUs, and RAM memory. Gaming components are commonly amongst the most heavily utilized, for extended lengths of time, computer systems available. For rendering, the cluster must be able to run at nearly 100% capacity for months or years on end in addition to being extremely fast. As can be imagined, they create extreme heat and so the components must be able to withstand it. These compute nodes are connected via Ethernet to a head node that does no rendering. It directs traffic on the cluster by scheduling rendering events as they progress through a render job and by collecting the resultant images after they are completed. More on the specifics of building, configuring, and utilizing these systems appear later in this document.

Even with this extraordinarily efficient computing methodology render times were lengthy. For instance, it was not unusual for an 80-second long stereoscopic video clip to require 2 – 6 weeks and sometimes more, rendering 24/7, to complete… as many as 20 of these renders might be required for a single composition. The waiting is in many ways the most difficult component of this work since it is very difficult to achieve a creative flow until the editing phase, which can only occur after the completion of the lengthy rendering process.

Rendering times are determined by the complexity of each scene. Scenes for this project regularly included multiple iterations of reflections and refractions to achieve complexity from simple objects and environments. Computing the position, angle of reflection, opacity, intensity, color, hue, and numerous other parameters of each and every pixel in such a complex scene, in each of 30 stereoscopic frames per second, is a highly demanding task for any contemporary CPU. The resolution for the stereoscopic images created for this project was 7000 x 1750 pixels. This means there are 12,250,000 pixels per eye or 24,500,000 total pixels that need to be calculated for each frame. Ultimately, it is safe to speculate that, these extremely complex numerical calculations are comparable with the calculative tasks required for other more scientifically oriented calculations.

It is because of these factors that the HPC cluster functioned as an intrinsic aspect of the creative process. Render times, even when scenes were optimized to the extreme limits, have a huge effect upon what can and cannot be accomplished. Extracting the most out of each and every CPU cycle requires creativity and ingenuity. As such, the cluster directly affects the outcomes of the composition. This aspect of the creative process is discussed in greater detail later. Here is it important to apprehend the critical role the HPC cluster plays. It is not just powerful tool but an intrinsic partner in a technological quadrilateral utilized in the execution of this project. It is asserted here that any artist working in the area of holophonic and holographic visual music composition must know how to integrate HPC clustering into their creative pipeline if they are to proceed.

1.8 Intricately Interlaced Compositional Matrices

Since a basic overview of the major components of the processes involved in conducting this research is in hand, let us proceed by gaining a basic understanding of the intricate interlacing of them into matrices of interactivity. Figure 4 provides a simplistic overview of the compositional flow. Each block in the diagram represents a world of depth not immediately apparent upon first inspection of the seemingly linear format. It is only with a thorough understanding of each that the truly non-linear complexity of the process in its entirety becomes evident. One goal of this document is to foster such awareness in anyone so inclined to apprehend it. The overarching trajectory is intended to expand from the general to the specific. Figure 4 functions as a procedural overview, in the form of a flowchart that acts as a guide from the general toward the specific. Additionally, it clearly delineates the manner in which the interdisciplinary areas of Music, Visual Art, and Computer Science interact to produce an amalgam of their essences. Therein lies the beauty and the power of this approach to visual music composition. Each block will be discussed in detail in the forthcoming chapters, however, by way of introduction, here a synopsis of the overall interconnected processes involved is provided.

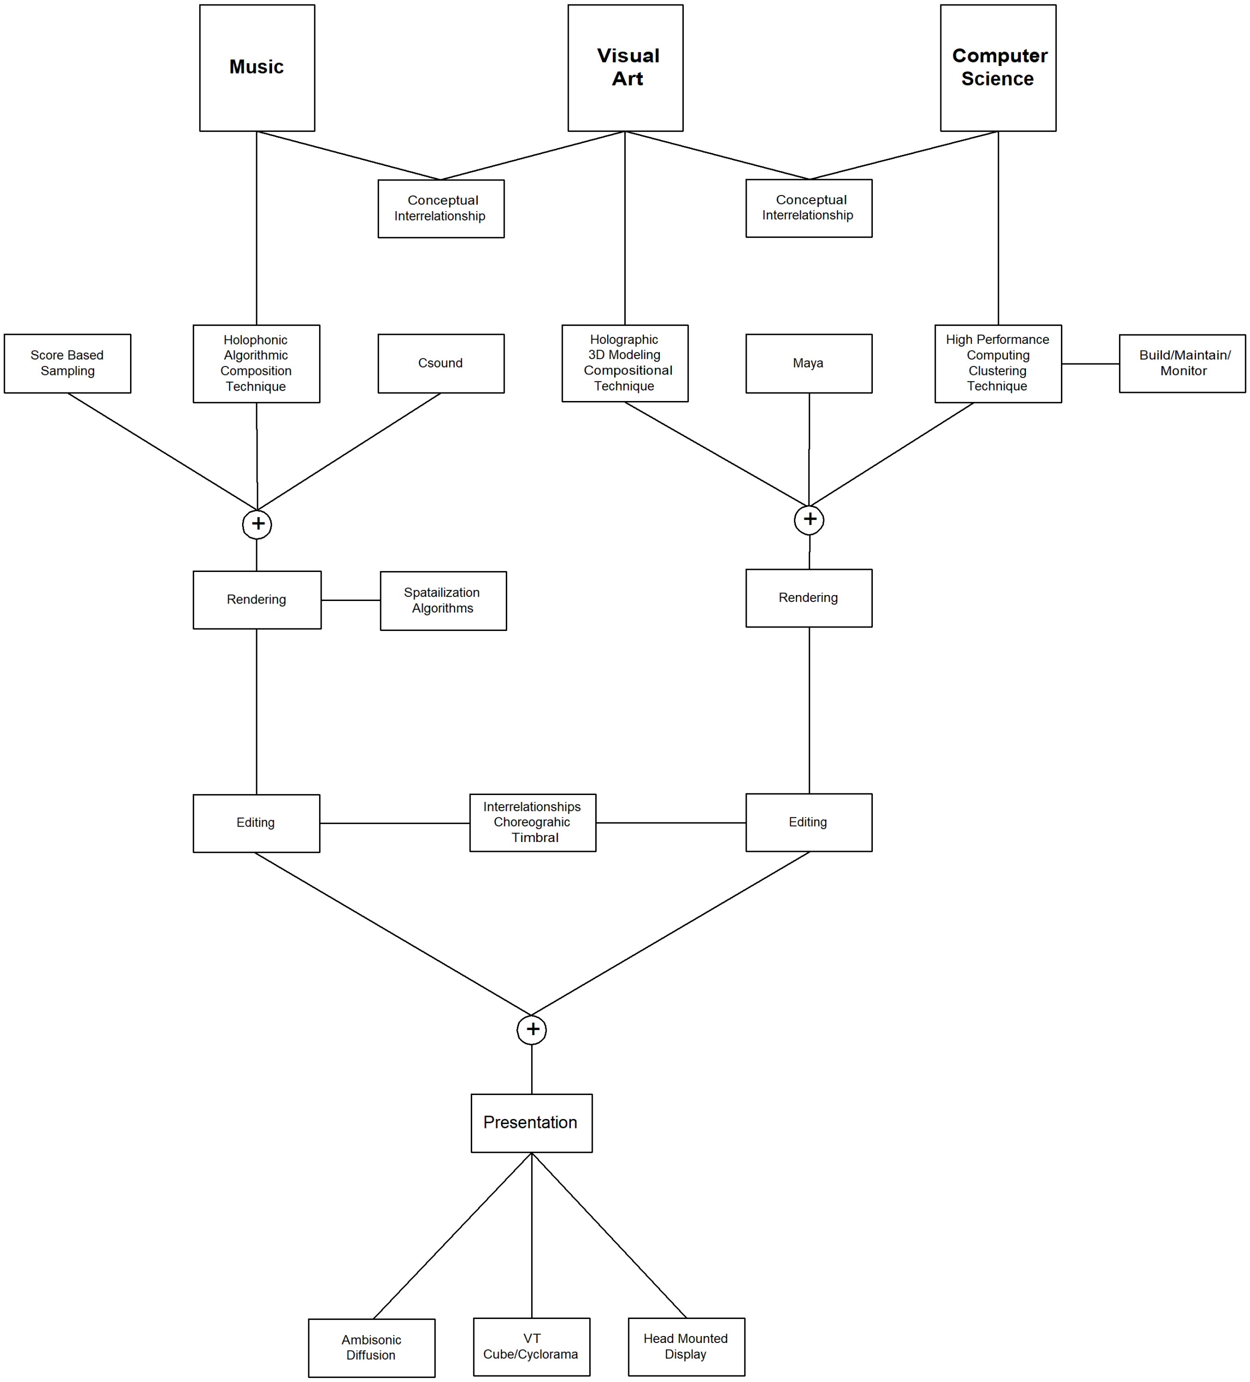

Figure 4 – Intricately Interlaced Compositional Matrices

1.9 Overview of Intricately Interlaced Compositional Matrices

From the top-down flowchart in Fig. 4, one can immediately ascertain that the goal is to unify three distinct disciplines toward the creation of a singular final artifact. Specifically, this is to say that, proceeding from perspectives in auditory art, visual art, and computer science, the intent is to create holochoric visual music compositions, ultimately, for public presentation. The remainder of the diagram is a stratified roadmap for traversing the extent of the endeavor. Yet, please keep in mind that, though this seems to imply a linear causality, this is a gross over simplification of what occurs in practice. Instead, for example, temporally speaking many of the steps denoted here could occur in any order. Nonetheless, the process is described here in a linear fashion.

The *Conceptual Interrelationships* between music, visual art, and computer science, usually begin as vague conceptions of a composition. At this point general frameworks with regard to the manner in which these three disciplines will interact throughout the composition are forming. Perhaps a schema designating various parametric intervallic relationships over time might be apparent at this point in the process. This might be formal or intuitive. Ideas about the visual content of the piece might occur, which then stipulate the manners in which initial rendering might proceed. This point in the process is deliberately left very open and pliable. As the composition develops it will take on a life of its own and determine much of the subsequent decision-making.

The *Holophonic Algorithmic Composition Technique* relies upon the use of *Csound* and the process of *Score Based Sampling*. Csound is a programming language intended for musical composition and Score Based Sampling is a technique whereby musical phrases and structures are derived through score manipulation. This occurs, primarily, through the use of stochastic techniques involving tendency masks. A program called Cmask, which is a stochastic numerical generator for Csound scores, implements the use of tendency masks to quasi-randomly generate values for each parameter of the score on an event-by-event basis. The inception of holophony is generated at this point through the use of convolution reverb and my own version of vector-based amplitude panning. It is here that the parameters for each are determined for each score event.

During the *Rendering* phase of the musical process, the Csound orchestra file and the algorithmically generated score file are referenced to render sound files using the Csound engine. It is in this step that the *Spatialization Algorithms* are implemented causing the holophonic aspects of the composition to begin to coalesce and to emerge into their final forms.

In tandem with Csound, *Maya*, a 3D modeling and rendering application, is employed to create the holographic (visual) scenes. To create a stochastic approach to the generation of visual material, numerous reflections and refractions, using glass, metal, and other reflective and refractive textures, are integral to each scene. Simple geometric objects are modeled, textured, lighted, and then placed into similarly reflective and refractive environments therefore attaining indeterminate visual complexity.

Pseudo *Holography* is achieved in this process through the use of stereoscopic techniques. A plugin for Maya called DomeMaster 3D includes a three-camera stereoscopic camera rig that is utilized. It consists of a left- and a right-eye camera each of which are parented to a center camera that is used for positioning the rig. Basically, the left- and right-eye perspectives are to be viewed alternately causing the brain perceive the two images as occurring simultaneously thus producing the perception of three-dimensionality.

The movements of the objects and environments in these scenes are achieved through the use of keyframing in Maya. The position of a given object at a given time is set with a keyframe and then its position at a later time is set. Maya interpolates from the first position to the second. This interpolation can be linear or techniques can be implemented that create exponential interpolations. Along with position, most any other parameter or characteristic of an object can be keyframed. Attributes such as rotation, color, lighting, scale, and etc. are easily set as keyframes thus creating extreme variability within the evolutions within a given scene.

High-Performance Computing Clustering Techniques are necessary to accomplish the high-resolution rendering required to generate the stereoscopic images for this project. Approximately 2400 frames, each consisting of two high-resolution images, are rendered per scene and often between 10 and 20 scenes are rendered for each composition.

The ROCKS clustering bundle was utilized to build a 48-node supercomputer capable of rendering the required image sequences within a reasonable timeframe. Being a Beowulf cluster, this super-computing paradigm utilizes multiple commercial grade computers, working as a single computation engine, to render the required images. The ROCKS bundle, maintained by the University of California in San Diego, is relatively straightforward to *build, maintain, and monitor*.

Rendering the Maya scenes is accomplished by the cluster. Each compute node of the cluster runs an installation of Maya and DomeMaster 3D. The scene is copied from the workstation, where it is conceived of and created, to the head node of the cluster, which acts as a scheduler for the rendering. It is responsible for assigning rendering tasks to each core of each compute node of the cluster in a comprehensive and organized manner. As each image is rendered it is written back to the storage hard drive on the head node and is then ready to be harvested by moving it to the workstation. Since the renders are extremely complex considering the numerous reflections and refractions, they often require weeks and sometimes months to complete.

Editing the sound files rendered by Csound is accomplished by importing the 32-channel sound files into Sonar®, which is a digital audio workstation. The various audio renders are cut, pasted, and copied, and rearranged in a multitude of manners as the composition takes form. Effects such as fades, reverb, filtering, delays, retrograding, transposing, and etc. are also applied during the editing phase. Where specific editing tasks are required, Sound Forge® is also a powerful tool that is often utilized.

Similar to editing the sound files, *Editing* the video renders is accomplished using Adobe After Effects® and Premiere®. In After Effects, the stereographic images are stacked, left-eye on top and right-eye on bottom, for each rendered image sequence. Effect renders of the stacked images can be created at this point to create variations of the lengthy Maya renders. The stacked stereoscopic images, along with low resolution monitoring versions, are imported into Premiere and then edited into a final composition.

Interrelationships are achieved when the rough edits of the music and video are complete. The music is imported into Premiere and then the video clips are synchronized with the music. Conversely, the rough edits of the video are imported into Sonar, and the music is then synchronized with the video. This activity reciprocates between the two applications until a finished version of each is accomplished.

Choreographic and Timbral relationships are established as a basis for tightly integrating music with video. As with ballet, for instance, the music is tightly interwoven with the movement of the dancers. The two elements at times move in unity, at others in harmony, at still others they move in counterpoint, and at others they are completely disparate. The key here is to relate music and video in terms of movement relationships.

Timbral relationships between the music and the video are also key factors in establishing tight interrelationships between the two. For instance, as the color of an object changes in the video so might the tonality of the associative musical phrase be altered. Though often subjective interpretations, the mood of a section of music might profoundly affect the visual qualities of the scene and vice versa.

It is an important consideration to apprehend that the relationships between the music and the videos is key to creating a visual music composition as opposed to a music visualizer. Here, the idea is to compose these relationships in a formal manner. Certainly, the spatial qualities of the work interact as well. Considering that the perceived location of a given sound event and a given visual event occurs within a 3D space shared by the audience, spatial relationships are an important interrelationship.

Presentation of the finished compositions is the final stage of this process. These compositions are inherently site-specific owing to the need for special equipment to accurately represent the intended form. *Head-mounted Displays* or proprietary 3D/360-degree stereoscopic projection systems like the *Cyclorama* in the *Cube* at Virginia Tech, are required for the intended video playback. High-Density Loudspeaker Arrays, consisting of 32 or more loudspeakers are required for the intended music playback implementing high-order *Ambisonics*.

The forms the presentations of this work can take include installations or virtual reality theaters. VR theaters, using head-mounted displays, are becoming ubiquitous at popular film festivals. As this trend continues the availability of venues for this work will increase. Many art museums recognize the need for presentation of these types of materials and are becoming open and equipped for them.

This brief introduction presents a general overview of the processes presented in detail throughout the remainder of this document. To apprehend a more complete apprehension of the intricately interwoven matrices involved in the composition of the artifacts involved and the practice-based research processes, it is necessary to delve into each of the following chapters.

1.10 Intermedial Counterpoint and the Tone Row Methodology

Throughout this document two concepts are consistently referred to that bear some clarification before proceeding. The concepts are extensively explored in relation to the manner in which they evolve the field of visual music later. Here we are concerned with a basic understanding of the terms themselves.

The first is that of "intermedial counterpoint." Counterpoint, in general, is the interplay of contrasting ideas. Musically, it can be described as two or more melodic phrases interacting within a common time sequence. As applied in this research and creative praxis, intermedial counterpoint refers to the interaction of multiple contrasting mediums toward a visual music expression. On the superficial level, this relates to the conjunction and juxtaposition of musical and visual art, high-performance supercomputing, and the composer's influence.

In the Antitheses project the interactions and interplay of various parameters of these four elemental constituents are exposed. Whereas in traditional musical counterpoint, the contrasting interplay occurs between melodic voices, such as the soprano, alto, tenor, and bass, in the intermedial counterpoint of the Antitheses, the counterpoint could be, as an example, between hues of light, computational rendering cycles, audio amplitudes, and the composer's sensibilities. These types of intermedial relationships reveal not only a unique stratification of process but they provide for multi-sensorial experiences. When projected and diffused within a 3D/360-degree venue, holochoric audio/visual compositions present an audience with a novel temporal expression occurring in multiple spatiotemporal dimensions. These ideas are explored throughout this document and, perhaps more importantly, they are expressed in the visual music compositions, which form the essence of this project.

The second concept is that of the "tone row" methodology. This terminology is utilized loosely here as an allusion to the efficient use of all materials the methodology produces. In traditional tone row musical composition, an ordered set of twelve notes is used exclusively for an entire composition. Numerous variations are traditionally employed requiring creativity and ingenuity to maintain interest and formal development throughout the composition. In the work associated with this project, the quantity of visual material available is constrained by extremely lengthy rendering time thus requiring an attitude of creativity similar to that involved with tone row composition. The traditional processes of variation include inversion, transposition, retrograde, and retrograde inversion. In the Antitheses project, this is expanded to include processes such as multiple layers of reiteration, effects such as reverberation, comb filter sweeps, harmonic manipulations, and time stretching or condensing on the musical side. Visually, effects such as beveling, hue and color alteration, various lensing and other distortion effects, and atmospheric additions such as cloud, smoke. The potential of the base material is exhaustively explored in this manner as a contribution to the composition. The methodologies employed to create the variations are detailed later in this document. As if by osmosis, though directly limited by the lengthy rendering cycles of the visual material, this approach indirectly transmits to the musical aspects of the compositions, thus providing a cohesiveness of compositional approach. It is considered desirable to work in this manner since it fosters an attitude of theme and variation, which is an age-old approach to musical composition. Such a methodology is based upon creating and including related yet varied raw material intended to be mixed and edited into a finished work.

1.11 Creativity

“All true artists, whether they know it or not, create from a place of no-mind, from inner stillness” Tolle).

“I think 99 times and find nothing. I stop, swim in silence, and the truth comes to me” (Einstein).

The practice-based research undertaken for this project hinges upon the ambiguous act of creativity. Therefore, it is prudent to present the author’s perspective on the subject by way of finalizing this introductory chapter. What follows is the delineation of such, which is intended as and realized as a singular perspective. It is not intended as an authoritative statement.

Creativity is intrinsic to the practice of composition. (Throughout this document the term “composition” is used in general terms to refer to the act of the organization of the both the auditory and the visual into a completed artifact.) More so, composition is a result of creativity. It requires fearlessness in facing the unknown and proceeding. Analogically, an act of creativity involves going out on a limb with little more than a glimmer of hope of some form of success. However, more than an action, creativity is a state of mind. This state of mind underlies, and is the impetus for, the work here described. Beginning with a blank page and a free and open mind we endeavor toward a certain goal knowing that it will never be achieved… for the act of creativity also involves and implies pliability. We may have a direction in mind, when we begin, that can act as a catalyst yet, if we are being creative, there must be a level of flexibility as the project develops. There is a feedback loop between the composer and the composition… each reacting to the direction established by the other and yet establishing new directions as well. This push and pull continues throughout the course of a project until there is nothing more to do. Indeed, knowing when this finalizing moment has occurred is a subtle yet powerful aspect of creativity, which of itself is timeless.

Each person who approaches a deliberate act of creativity does so from a unique perspective. The notion of creativity by its nature infers such individuality. There is no predetermined manner in which to proceed. This is an essential aspect of the utter beauty and joy involved. Yet with this unparalleled freedom comes a price… the price of uncertainty. Proceeding in a creative direction is not for the faint of heart. We human beings crave the illusion of security in our lives. Much energy is spent contriving and maintaining the illusion of control. The act of creativity is a confrontation with the reality that we are not in control of anything but also that we seem to nonetheless have an effect upon that with which we interact. This paradox explains, at least in part, why so many creative minds are often tortured minds. It can be difficult to reconcile the two perspectives.

Creativity, as previously mentioned, is a state of mind and yet beyond that it is a way of life. While maintaining a mindset that is conducive to creativity is healthy in the studio, it can be quite challenging in most other social settings. Further, sustaining creativity can be extremely challenging. When it temporarily stops, as is inevitable, we are left bereft of solace. There are often sine waves of psychological ups and downs that accompany deeply engaging a lifestyle entrenched in creativity. While in the act of creativity we are ultimately free and uniquely challenged. It is a form of euphoria with no equal. Yet moving in and out of the creative mindset, as life requires, can be tortuous. Nonetheless, it is conceivable that a mature creative mind can learn to bridge this chasm and to fluidly shift between the poles. As can be imagined, this kind of maturity is probably quite rare yet it is proposed here that it is attainable.

From the author’s perspective, creativity can only be derived from the quiet mind. By a “quiet mind” is intended a mind that is free of the domination of thought. Obviously, this is generally the antithesis of an academic perspective and yet it is postulated here that thought is not the epitome of human achievement. It is inarguably a tool worthy of an intense investment in development, and yet it nearly ubiquitously assumes control of one's actions, which can and does have disastrous effects. This can be witnessed by the current state of human affairs in general on planet Earth.

Creativity occurs when one is free of the domination of thought. Thought can be considered a collection of accumulated memories that are constantly rearranged to address the current situation. As such they are nothing new, per se, they are a result of the past… at best the reconfiguration of that which is known and which already exists. This could be considered re-creation… recreation. Creativity, by definition, instantiates the new. This is something of which thought is incapable. Considering this perspective, a lifestyle in which thought is not in control is one that is conducive to creativity. Upon experiencing this phenomenon, the creative mind becomes self-aware and is ready for the moment when a shift occurs. This shift is an inspired opening in the domination of thought… one in which creativity is catalyzed. The creative mind, a highly sensitive mind, is ready to seize such a moment and to fully engage with it (Krishnamurti).

1.12 Looking Ahead

As a preview, in Chapter 2 the technical aspects of the musical compositional process are delineated. Chapter 3, is an investigation of the technical aspects of the visual compositional process. The subject of Chapter 4 involves the process of building, configuring, implementing, and maintaining, the high-performance computing clusters required for rendering the visual images for this project. Chapter 5 describes the confluence of the afore mentioned technical processes into a unified compositional approach. There, questions of aesthetics, creativity, and challenges unique to the nexus of holophonic and holographic art are considered. In Chapter 6 modes of presentation are discussed. Chapter 7 is a summary and is where conclusions are drawn knowing that they will soon be dismissed as new data is acquired.

CHAPTER 2. MUSIC COMPOSITIONAL PROCESSES

2.1 Chapter Overview

As a fundamental element in the evolutionary progression of counterpoint in the field of visual music composition, this chapter endeavors to communicate the details of the musical compositional processes utilized to create the associated musical aspects. As will be demonstrated, stratifications occurring on multiple levels with regard to the methodologies involved and the mixing of the resultant rendered sounds play important roles in the process. Actively engaged in a feedback loop with the other major constituents, the composer's sensibilities, an aleatoric approach to the creation of visual material, and HPC, the auditory art of this project is an important catalyst and result of this intermedial interplay. Though rendering cycles involved with the musical production are not lengthy in comparison to the visual aspects, the auditory and the visual are intricately interconnected and so the musical material must be approached in a manner similar to the visual. As will be detailed in chapter 3, visual rendering times create a long delay, weeks or months, between artistic conception and realization. Because of this, a compositional process extended from the tone row compositional paradigm is necessitated in which each musical idea is subject to extensive variation and development.

It is challenging to separate the processes into categories since they are non-linear in nature, meaning that they may occur out of order and, at times, aspects occur simultaneously. Further, when embracing a creative mindset, the lines between commonly separated mediums blur into inconsequentiality. For example, a musical phrase and a visual segment might be seen as synonymous. Nonetheless, in this chapter the goal is to begin to extricate the technical underpinnings of such and describe them in isolation. Below, in Fig. 5, is an excerpt from Fig. 4 that forms a general basis of what will be covered in this chapter.

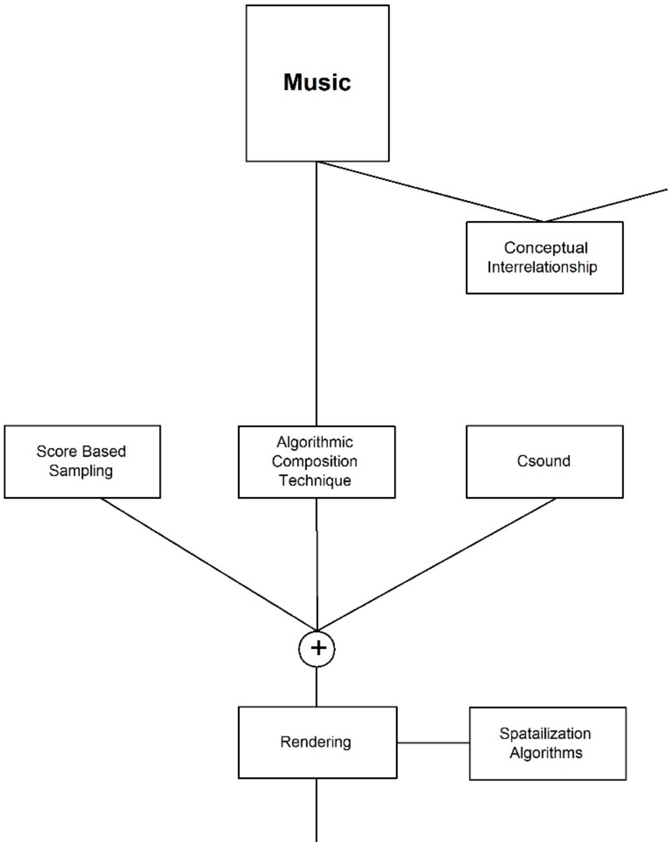

Figure 5 – The Generative Musical Compositional Process

An overriding and guiding principle in this generative approach to composition involves the relationship between the composer and the computer. As the composer's alter ego, the computer plays a role intrinsic to the generation, or creation, of musical material (Laske). As an extension of this role, the process could be considered that of AI. However, instead of the acronym referring to artificial intelligence, "augmented intelligence" is considered a more realistic view. Through the techniques described in-depth in this chapter the computer greatly extends the creative capability of the composer thus enabling the creation of musical compositions that could not otherwise exist. Further an attitude of "cinema for the ear" is implicit in this approach. This phrase serves the purpose of describing the results of composing in this manner by way of an analogy. Think of a stage play as compared to a movie. Though many innovative approaches to a stage production have emerged in the past few years, commonly a stage play is rather simplistic in terms of temporal and spatial locale as opposed to a movie. For instance, to change the location in which a scene is taking place on stage, the set must be changed in one of several ways. In a movie, such a change requires only shooting on location and then integrating the footage in the editing room.

Shifting from one world to another can occur in an instant. Employing the algorithmic techniques described below, music is created that could never be performed live on a stage. It is far too complex in a manner similar to the way that most movies would be impossible to perform on stage. The special effects commonly employed in any contemporary movie are further evidence of this parallel.

To attempt to compose the music for the compositions created for this project without the aid of the computer would require a lifetime to accomplish. There is extreme power in the human mind working in conjunction with a computer. For instance, the computer is much more capable in the areas of memory and calculation than a human being unaided. So here the process involves handing off compositional tasks that require extensive memory and calculation to the computer while the composer maintains a creative overview of the results. The human composer is currently more adept at interpreting the computer output from a humanistic perspective than the computer is and so makes the decisions as to what to utilize and/or modify and what to discard. After these decisions the human composer organizes the material into a form that is palatable to other human beings.

With this approach to composition the use of the word “composer” to describe the human being involved in the process is perhaps a misnomer. The role of the person involved in this type of process becomes that of a programmer, a constrainer, and an editor. Fundamentally, the computer does the composing. Here we will continue to use the word “composer” to refer to the human involved in the process of creating the visual music compositions, but this is only for convenience. The common connotation associated with the word “composer” may not be accurate.

These are a few of the ideas behind the processes involved herein. As such the music composed by means of these processes currently remains on the fringes of possibility.

2.2 Csound

The music considered herein finds its technical roots in Csound, which is a programming language meant for musical composition. In order to apprehend the compositional processes involved in this project, a basic understanding of how Csound works is necessary. Here, in the first section of this chapter, we will endeavor to provide such an understanding.

Csound's inception can be traced back to Bell Labs in the early 1960's. There, Max Mathews, under the direction of his supervisor John Pierce, was working on the first notions of recording and reproducing sound digitally. This meant that a sound could be analyzed and then represented digitally as data and then that data could in turn be converted back into sound. The innovative idea was that text-based data was much easier to send over long distances compared to an analog signal, which was the standard telephony of the day. While working on this project Mathews conceived of a way that sound could be produced, without the initial analysis stage, using the code. His first composition using this method was the infamous "Bicycle Built for Two" that the AI HAL sang in the original 2001: A Space Odyssey movie. Through the subsequent years this Music11 language, as it was called, developed into Mathew's final version; the Music 4 language. In the mid 70's Barry Vercoe at MIT continued developing the Music 4 language into what is known as Csound today. Since then, the further development of Csound is continuously occurring and is driven by a group of dedicated programmers.

There exist several powerful frontends to the Csound programming language. Among them are Blue, developed by Steven Yi, Cabbage developed by Rory Walsh, and CsoundQT by Andres Cabrera. There are also API hooks into and out of Csound for various popular programming languages such as C#, C+, Python, Java, and etc. Though each extends the reach of the programming language through unique and dynamic features, it is not necessary to use them. Here the most basic implementation of Csound is utilized and yet with it an entirely unique compositional paradigm was instantiated. Every composer that uses Csound does so in a manner unique to his or her own sensibilities… thus revealing the true power of Csound. Of course, such power is accompanied by a steep learning/designing curve but it is well worth the effort to one dedicated to the process.

Csound can be used in real time, as midi input and output, live audio programming, used in conjunction with other programs such as MAX or Ableton Live, or to employ an out of time approach rendering sound files. It is the latter feature that is implemented for this research. To instantiate this most basic configuration one writes two plain text files. One is an orchestra and the other is a score. In the former, using Csound opcodes and syntax, one or more instruments are created using various synthesis techniques such as frequency modulation (FM), amplitude modulation (AM), granular synthesis, sample playback, and numerous others. By default, each event in the score determines when each instrument plays and its duration. Any other variable parameters, as stipulated by the instrument in the orchestra are also determined in the score on an event-by-event basis. After completing these two text files, the Csound rendering engine is instantiated to create a sound file referencing them. Below is a very simple orchestra/score example wherein a sample is played back:

```
;orchestra
;header
sr      = 48100         ;sample rate
kr      = 48100         ;control rate
ksmps   = 1             ;control rate/sample rate multiplier
nchnls  = 1             ;number of output channels

instr 1                 ;instrument number
idur    =       p3      ;duration
iamp    =       p4      ; amplitude
ifreq   =       p5      ; playback frequency
isfile  =       p6      ; file name

kenv1 linseg    0, idur * .001, iamp, idur * .799, iamp, idur * .2, 0     ;amplitude envelope
a1 diskin2 isfile,  ifreq,  0,      1                                     ;sample playback opcode
out     a1 * aenv1                                                        ;ouput
endin                                                                     ;end of instrument

;score
t 0 60                                                                    ;tempo
;p1     p2      p3      p4      p5      p6                                ;pfield number
i1      0       10      1       1       10                                ;instrument 1 instantiation
e                                                                         ;end of score
```

Figure 6 – Example Simple Csound Orchestra

This example produces a sound file 10 seconds in length beginning at time 0. It plays the sample labeled soundin.10 at a frequency of 1 and an amplitude of 1, which is the frequency and amplitude of the original sample. If the frequency value in (p5) were changed to 2, the sample would have been played back at twice the normal rate and so, when rendered, would produce the sound of the sample playing an octave higher than the original sample. There would be a similar response to the amplitude value in (p4). Incidentally, the individual columns in p1 – p6 of the score are called pfields.

Though we have glossed over many of the details in this example, the basic workings of the language are demonstrated. For more in-depth understandings of Csound there are numerous tutorials and other references that can be found with a simple search online. Such understandings are beyond the scope of this document.

Using the reiterative score-based sampling paradigm outlined in the next section, the first step involves an orchestra/score combination very much like the one above. The second step is extremely complex in comparison as is the process used to render it as a sound file. For it, the orchestra file consists of one instrument, which is approximately 10,000 lines of code in length, with 43 pfields, and renders an 8-channel sound file.

2.3 The Process - Score Based Sampling

In the mid 1970's Dr. Otto Laske was teaching, studying, and composing at the Institute of Sonology in the city of Utrecht in the Netherlands. One of his goals was in formulating a "theory of making music" (Laske, Tabor). One reason he considered this important was because he was forging an unprecedented path into the act of composing generative music… music composed through the use of artificial intelligence. The theories and perspectives he and the director of the institute, Gottfried Michael Koenig, developed have inspired and informed subsequent generations of composers.

Figure 7 – Laske (right), Barry Truax (center), and Curtis Roads (left) at Utrecht circa 1989 (Courtesy Otto Laske)

Dr. Laske and I began a friendship in 2002 that continues to this day. Over the years we have exchanged numerous ideas and theories. One of Laske's theories, which he had rudimentarily developed during his tenure at Utrecht, was that of "Score Based Sampling" (SBS). His notion was that through reiterative processes, musical sequences could be generated from a basic set of samples, using straightforward rule sets, which would at once achieve variability and yet maintain the type of similarity that composers typically employ to provide cohesiveness in a composition. From a formal perspective this was an incredibly enticing theory. Subsequently, the next twenty years was spent integrating it into my compositional processes and dynamically expanding and refining Dr. Laske's subtle yet extremely powerful compositional theory. Through the practice of composing in this manner, my process has evolved into the general structure represented in Fig. 8 below.

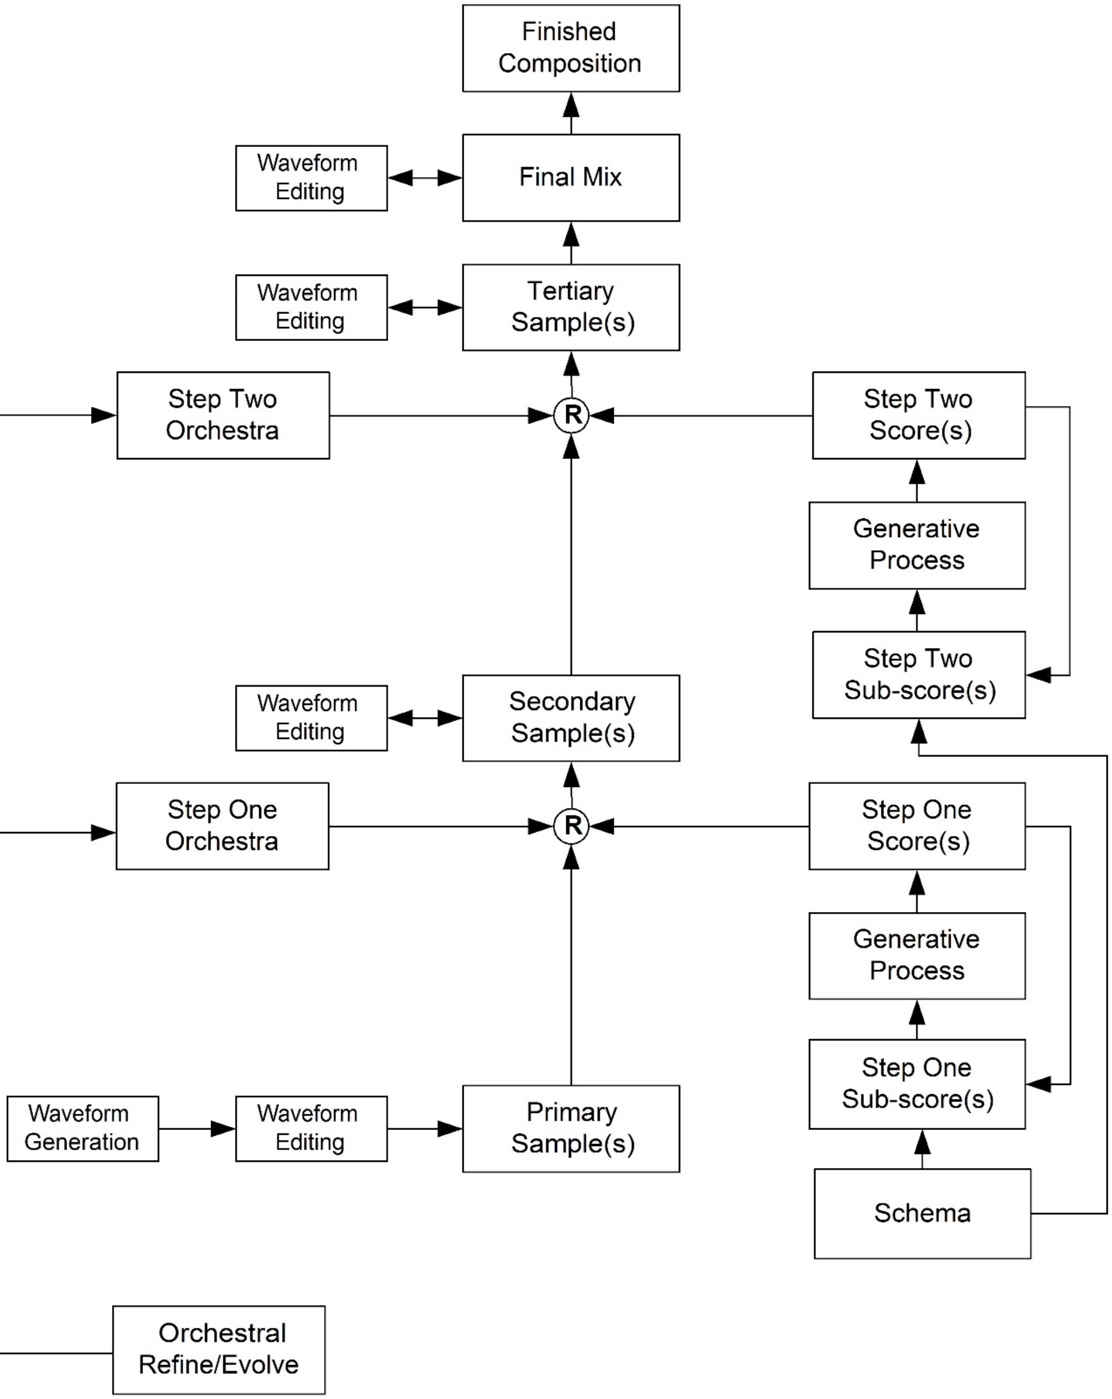

Figure 8 – Generative Algorithmic Compositional Paradigm

In order to apprehend the flow of the diagram in Fig. 8, one must first understand that it is "bottom up" in nature. Using a "top down" approach, a composer begins with a musical idea and then develops a composition from it. This could be described as "model-based" composition, which was the primary approach to composition until the mid-20th century, when “rule-based” composition became a driving force in the world of computer music (Laske, Tabor). To follow the rule-based compositional paradigm diagramed here, one must begin at the bottom, with an idea of the rule(s), or algorithms, to be employed, and work toward the top only then to fully realize the musical result. This approach follows nicely with the author’s notion of stepping back from the compositional process and allowing unexpected emergent quanta to occur and then to develop it. It is an approach based upon stochastic processes in an effort to engage in indeterminacy. This is intended as a deliberate attempt to remove the limited and limiting conscious mind of the composer from the process in lieu of an influence greater than the sum of the parts to emerge. It is the sincere intention to set up circumstances wherein the computer will produce emergent behavior and then the observant composer, who notices it, further develops it into a composition. This complex compositional process was required, and so developed and employed, to accomplish this task.

Though it is reasonable to presume that composing from a stratified structure such as this might lead to homogeneity of result, nothing could be further from the truth. This approach provides an extremely open concept wherein creative license is impossible to escape. In fact, it has been the experience that each new composition produces a unique syntax and therefore a similarly unique semantic. Though complexity is in itself not a worthy goal, nor does it necessarily evoke a desirable outcome, it is through a process such as this that a leading edge of musical composition can be approached. It is asserted here that we human beings are capable of much more by way of expression and comprehension than we currently realize or engage. Certainly, the human mind appears to continuously evolve and expand its capacity as can be witnessed by the increasingly rapid technological growth and implementation we are currently experiencing (Kurzweil). Should not our music also reflect and embrace this trend?

Considering that this process of musical composition is at the confluence of each facet of the research conducted in this study, it is advantageous to expand each category stipulated in the diagram above.

Orchestral Refine/Evolve

To begin each new composition, or more often each new set of compositions, it is the practice to expand and refine the Csound orchestra in Step 2 that was used for the previous set of compositions. This process has occurred over the past 20 years and so the changes have become increasingly subtle. They do occur nonetheless in order to both expand and refine the orchestra. For example, for the current research project and set of compositions, a convolution reverb has been integrated into the orchestra. Though this will be described in detail later in this chapter, for now we will state that the change was made to enhance the holophonic characteristics of the spatial aspects of the work. These types of changes are ongoing and are considered the beginning of the compositional process.

2.4 Waveform Generation

In order to collect the, usually four, primary samples (base samples) to be utilized for the composition, we begin with the generation or the recording of the basic waveforms. Often software applications such as UI Software's Metasynth®, Native Instrument's (NI) Absynth®, or NI's FM8®, or others, are used to generate rough versions of the base samples. Additionally, Csound programming language is often employed to create samples as are live recordings of naturally occurring sounds. Basically, we are looking for four sound samples, approximately six seconds in length, that seem in some manner interrelated and yet are diverse enough to be interesting. At this point we only care that the sounds evolve in a manner that feels "right". This highly subjective and intuitive determination will have a huge impact on the direction of the composition. Therefore, instinctive impulses are more important at this stage than thoughts about them. Often 10 – 20 samples are initially generated and then selectively eliminated over the course of several days, and several listening sessions, until the most desirable set of four remain. In the mid 1970's Laske determined that this base sample structure utilizing four samples was optimal and though variations to it have achieved excellent results, it usually provides for a very desirable underlying structure.

2.5 Waveform Editing

After or perhaps during the process of sample selection waveform editing is instantiated. To accomplish this Sound Forge Pro®, and/or Sonar® Professional, is utilized to conduct refinement of the samples. This refinement involves trimming any unwanted material or silences from the beginning or the end of the samples and refining the amplitude structure. This latter aspect involves making the attack portion of the amplitude envelope begin smoothly and then increase to its peak at a speed that is determined to be optimal for the composition. Another aspect of this step includes making sure that the relative amplitude of all four samples is generally equal. This establishes a control level that can be relied upon during the subsequent compositional process. Figure 9 is a screen shot of a set of four samples in Sound Forge®. As can be seen, the peak amplitudes are each at around -6db regardless of where they appear in the evolution of the sample over time. One caveat to this step is that at times it is desirable to create morphologies from one or more sample(s) to one or more other sample(s). There have been occasions when the base samples are morphed with each of the other samples to create 16 base samples total. This is accomplished by placing one sample over the other in the digital audio workstation (DAW), beginning at a predetermined time after the beginning of the first sample and then employing amplitude envelopes to fade between them. See Fig. 10. Although the first step in the Score Based Sampling paradigm is aimed specifically toward the quasi-random mixing of the samples, there are times when approaching it manually is preferable and is used in conjunction with it.

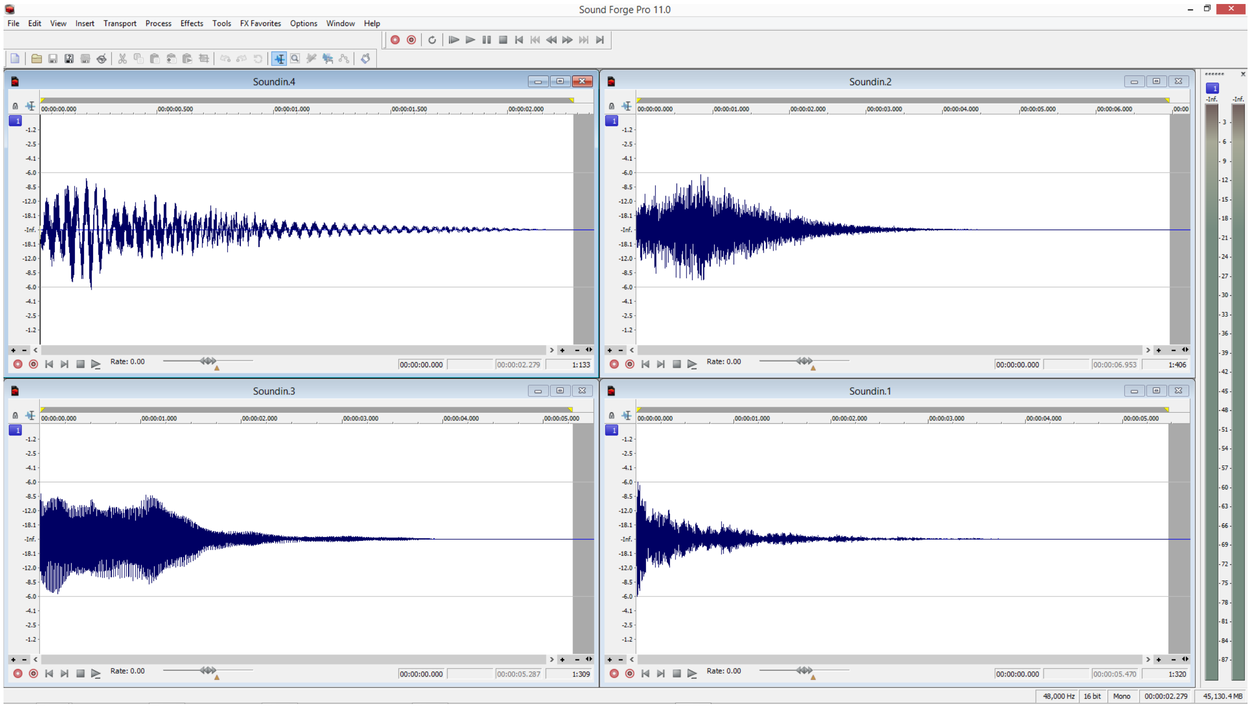

Figure 9 – Edited Base Samples

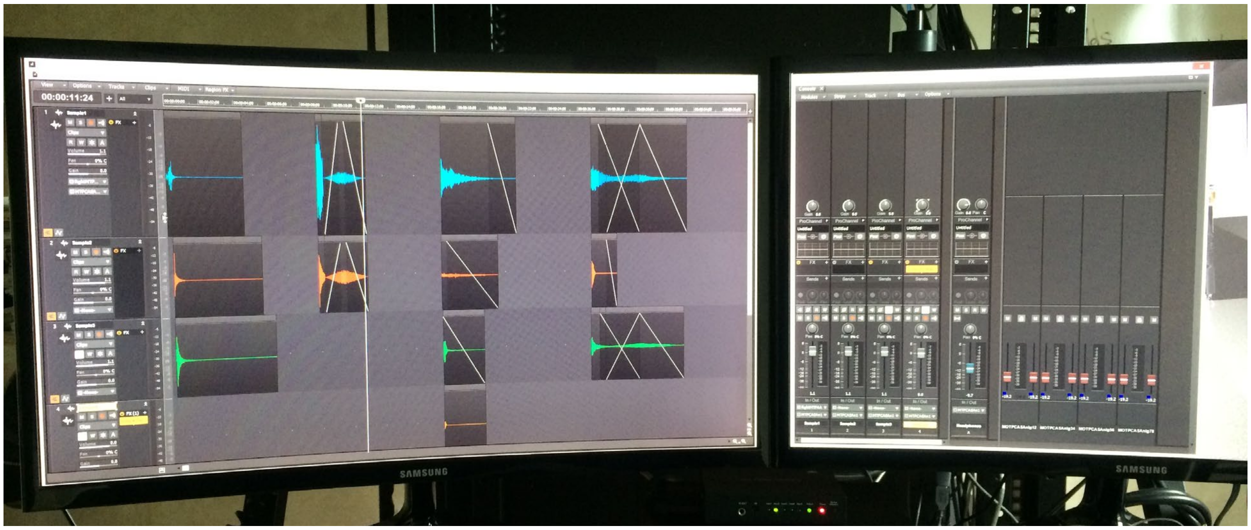

Figure 10 – Creating Four Morphologies of the Base Samples as Compound Samples

2.6 Primary Samples

With the base samples created, they need to be renamed from whatever they were named whilst working with them and put into the correct directory so Csound can find them without requiring a full path be stipulated in the orchestra file. Therefore, they will be renamed soundin.1, soundin.2, soundin.3, soundin.4, soundin.n… and put into the directory labeled SSDIR in the Csound directory. This will allow them to be referenced in the Csound orchestra or score file as 1, 2, 3, or 4. With this step completed we are ready for the next, which involves the score generation process.

2.7 Schema

The schema step in the process can be extremely helpful but it can also be quite constraining and so it may or not be utilized as desired. If not, then an intuitive real-time approach is employed. Basically, the schema is a methodology for stipulating, monitoring, numbering, and the overall organization of the reiterative process of sample mixing. It can be quite helpful later in the process when there are many samples and, for example, one finds it necessary to derive the origin of a specific sample. Figure 11 below is a screenshot from an example of a typical schema. Creating it on a spreadsheet is an extremely helpful organizational technique.

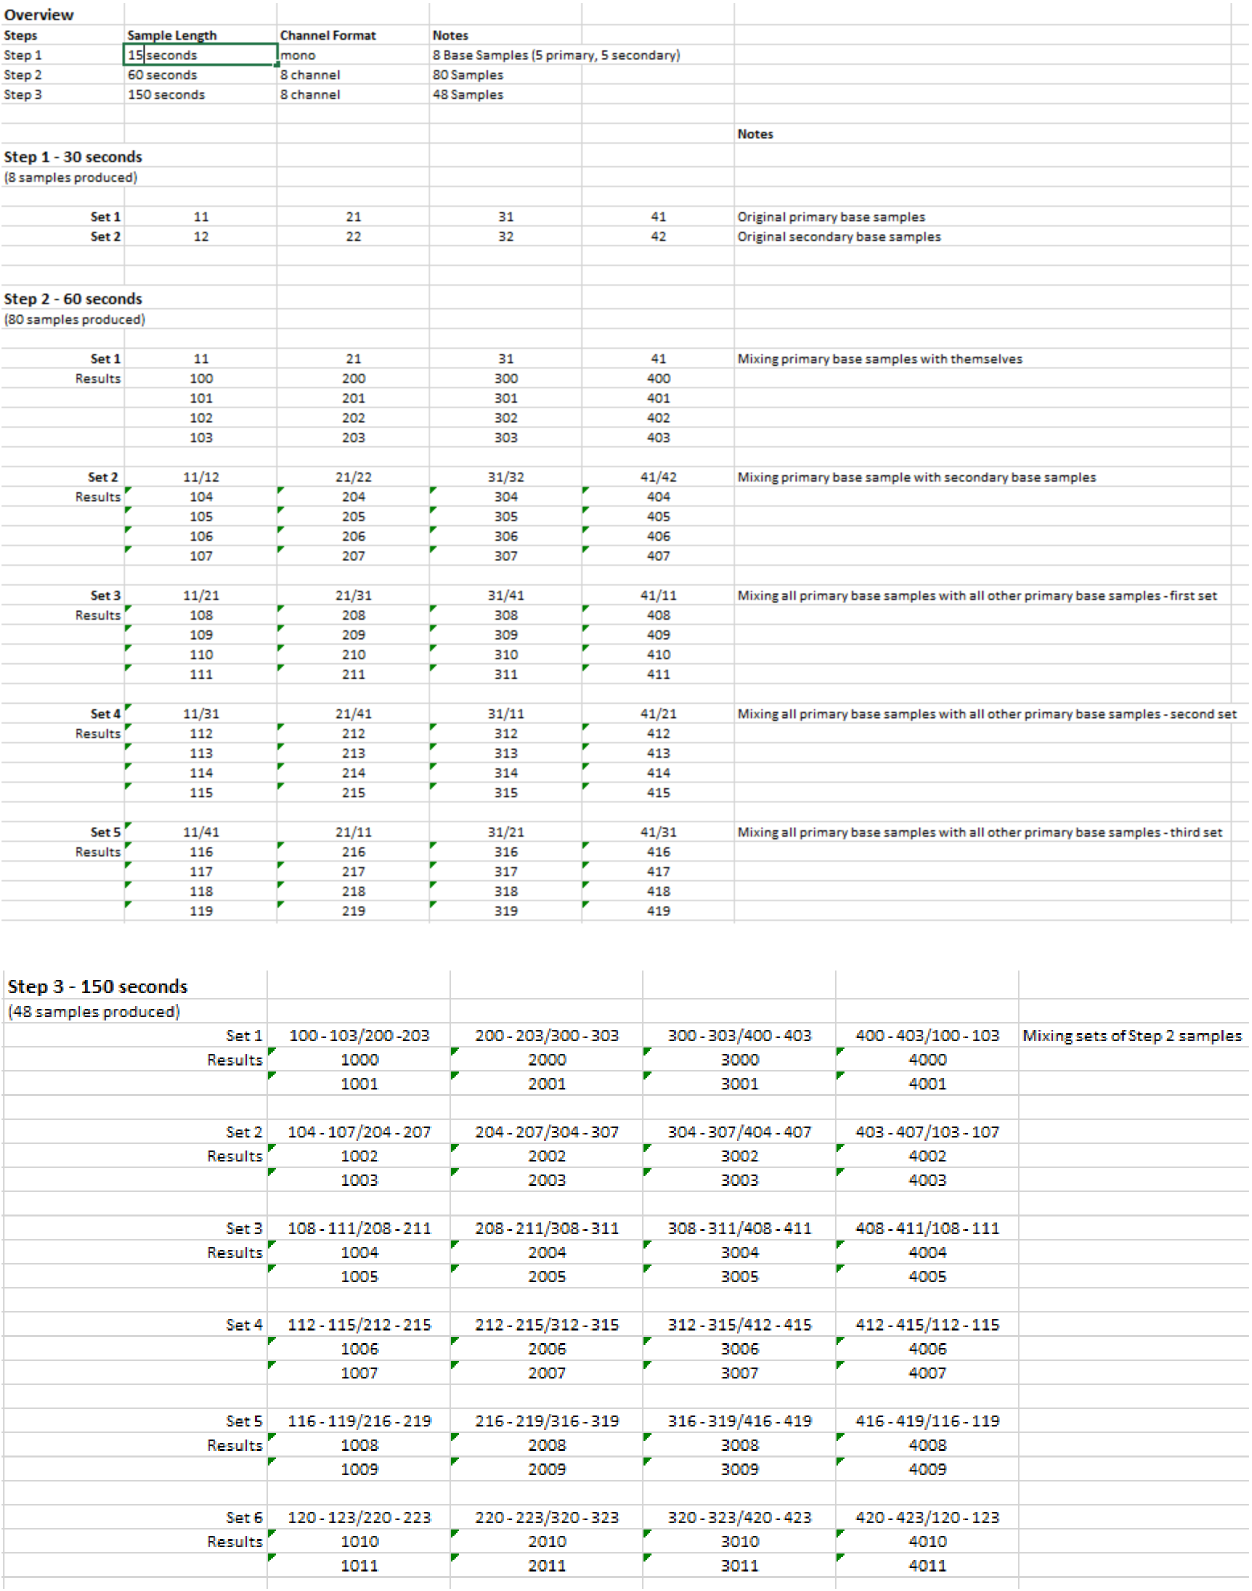

Overview

Steps	Sample Length	Channel Format	Notes
Step 1	15 seconds	mono	8 Base Samples (5 primary, 5 secondary)
Step 2	60 seconds	8 channel	80 Samples
Step 3	150 seconds	8 channel	48 Samples

Step 1 - 30 seconds
(8 samples produced)

					Notes
Set 1	11	21	31	41	Original primary base samples
Set 2	12	22	32	42	Original secondary base samples

Step 2 - 60 seconds
(80 samples produced)

Set 1	11	21	31	41	Mixing primary base samples with themselves
Results	100	200	300	400	
	101	201	301	401	
	102	202	302	402	
	103	203	303	403	
Set 2	11/12	21/22	31/32	41/42	Mixing primary base sample with secondary base samples
Results	104	204	304	404	
	105	205	305	405	
	106	206	306	406	
	107	207	307	407	
Set 3	11/21	21/31	31/41	41/11	Mixing all primary base samples with all other primary base samples - first set
Results	108	208	308	408	
	109	209	309	409	
	110	210	310	410	
	111	211	311	411	
Set 4	11/31	21/41	31/11	41/21	Mixing all primary base samples with all other primary base samples - second set
Results	112	212	312	412	
	113	213	313	413	
	114	214	314	414	
	115	215	315	415	
Set 5	11/41	21/11	31/21	41/31	Mixing all primary base samples with all other primary base samples - third set
Results	116	216	316	416	
	117	217	317	417	
	118	218	318	418	
	119	219	319	419	

Step 3 - 150 seconds
(48 samples produced)

Set 1	100 - 103/200 -203	200 - 203/300 - 303	300 - 303/400 - 403	400 - 403/100 - 103	Mixing sets of Step 2 samples
Results	1000	2000	3000	4000	
	1001	2001	3001	4001	
Set 2	104 - 107/204 - 207	204 - 207/304 - 307	304 - 307/404 - 407	403 - 407/103 - 107	
Results	1002	2002	3002	4002	
	1003	2003	3003	4003	
Set 3	108 - 111/208 - 211	208 - 211/308 - 311	308 - 311/408 - 411	408 - 411/108 - 111	
Results	1004	2004	3004	4004	
	1005	2005	3005	4005	
Set 4	112 - 115/212 - 215	212 - 215/312 - 315	312 - 315/412 - 415	412 - 415/112 - 115	
	1006	2006	3006	4006	
	1007	2007	3007	4007	
Set 5	116 - 119/216 - 219	216 - 219/316 - 319	316 - 319/416 - 419	416 - 419/116 - 119	
Results	1008	2008	3008	4008	
	1009	2009	3009	4009	
Set 6	120 - 123/220 - 223	220 - 223/320 - 323	320 - 323/420 - 423	420 - 423/120 - 123	
Results	1010	2010	3010	4010	
	1011	2011	3011	4011	

Figure 11 – Typical Schema

Though the schema above is the most commonly implemented in this work, there are numerous manners in which a composition is originally conceived of… each requiring a customized approach. Figure 12 is an example of one such deviation. The composition that ultimately became, “Shards of Oils”, began with this formal flowchart. In it, various parameters were stipulated in general and subjective terms.

One reason it is important to consider this particular schema is that it documents a direct correlation between the musical and the visual aspects of the work. Though such correlations are present in each composition, they are much too multifaceted to attempt representation other than through the work itself. (Thus, the importance of the artifact as the evidence and the result of scholarly research and praxis.) Such inclinations operate on multiple levels simultaneously and are constantly developing and evolving. They are worked with in, what is usually described as, an intuitive manner. Intuition from this perspective is perhaps a higher form of intelligence that allows for complexity of action far beyond that of which the intellect could accommodate. Nonetheless, this overly simplified view was created as a general reminder of the manner in which the composition was originally imagined and as a communication tool.

Constraints such as these are often enjoyable to work with since they require a particular approach. We are never so creative as when we are faced with such obstacles. It is the drive to complete the composition coupled with the unwavering intent that it must be “right” that forces creativity into action.

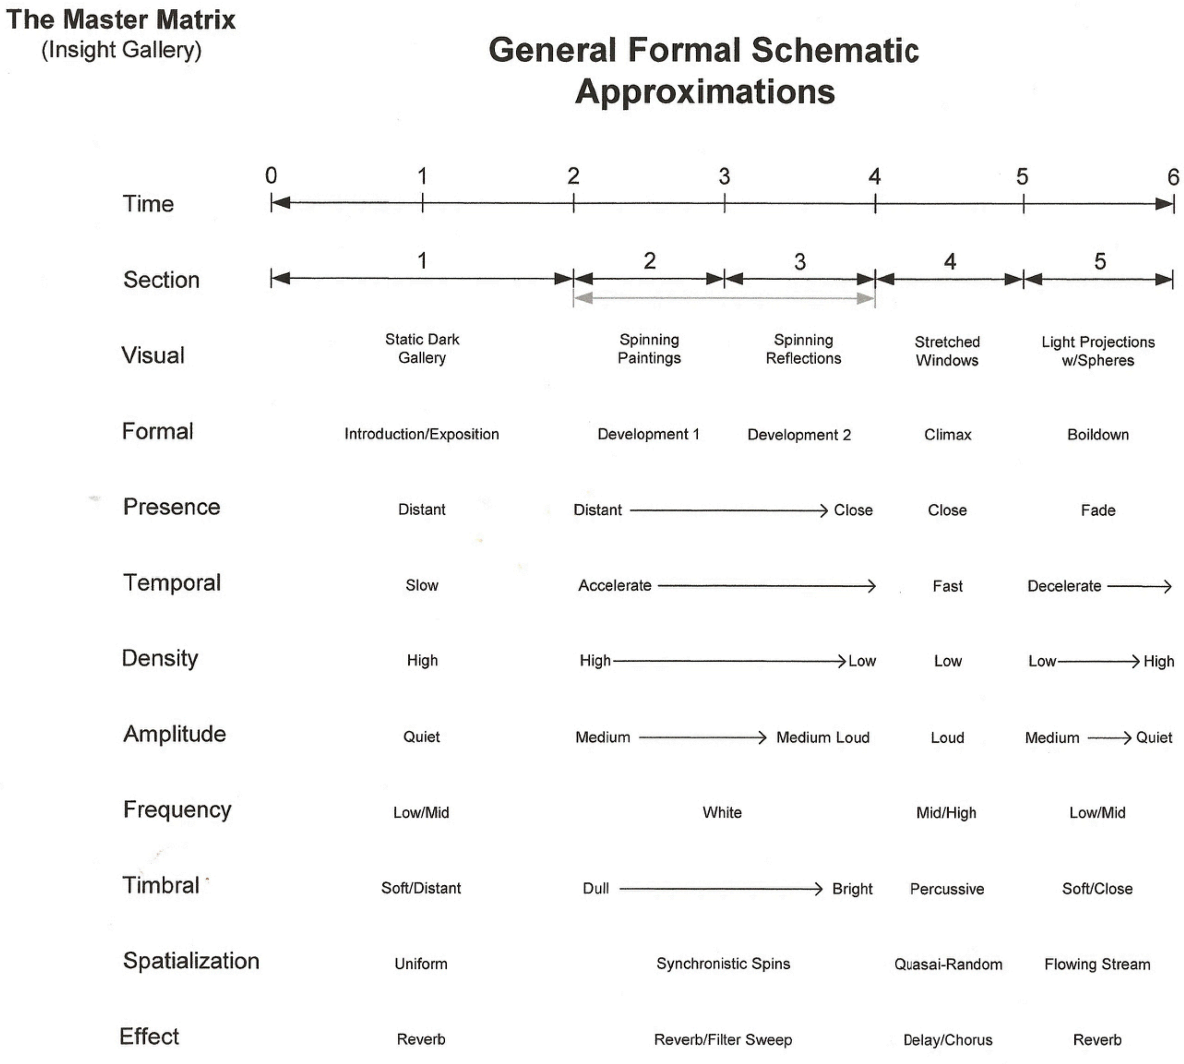

The Master Matrix
(Insight Gallery)

General Formal Schematic Approximations

Time: 0 1 2 3 4 5 6

Section	1	2	3	4	5
Visual	Static Dark Gallery	Spinning Paintings	Spinning Reflections	Stretched Windows	Light Projections w/Spheres
Formal	Introduction/Exposition	Development 1	Development 2	Climax	Boildown
Presence	Distant	Distant → Close		Close	Fade
Temporal	Slow	Accelerate →		Fast	Decelerate →
Density	High	High → Low		Low	Low → High
Amplitude	Quiet	Medium → Medium Loud		Loud	Medium → Quiet
Frequency	Low/Mid	White		Mid/High	Low/Mid
Timbral	Soft/Distant	Dull → Bright		Percussive	Soft/Close
Spatialization	Uniform	Synchronistic Spins		Quasai-Random	Flowing Stream
Effect	Reverb	Reverb/Filter Sweep		Delay/Chorus	Reverb

Figure 12 – An Alternative Schema

2.8 Step One Sub-score

The creation of the sub-scores is the essence of the score-based sampling process. We utilize the subscore to generate the Csound score, which, in addition to the orchestra file, is referenced when rendering the audio files. Cmask is a stochastic event generator for Csound scores (Rhoades, 2009). It generates score sequences based upon tendency masks, which are constrained by values stipulated in the sub-score. The tendency masks constrain the numerical score values, generated by Cmask, based upon weighted probabilistic distributions. For example, Fig. 13 below illustrates a uniformly distributed tendency mask. For it, the range of the values to be generated is between

0, stipulated by the bottom line, and 1, stipulated by the top line. Since both lines are horizontally parallel, every value within the range has an equal probability of being generated when Cmask is instantiated. Conversely, Fig. 14 represents a triangular distribution. Using it as a constraint within the same range, the value of .5 has the highest probability of being generated and values of 0 and 1 have the lowest. There are several distribution patterns that are readily available in Cmask such as Gaussian, Cauchy, Exponential and Reverse Exponential, Weibull, and etc. In addition to these pre-existing constraints, Cmask can be used to manually create other, more esoteric distributions. These are explained in the Cmask manual, which at the time of this writing, can be found here: www2.ak.tu-berlin.de/~abartetzki/CMaskMan/Cmask-Manual.htm.

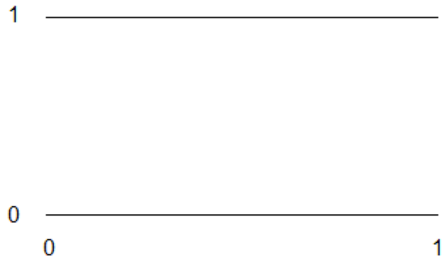

Figure 13 – Uniform Distribution

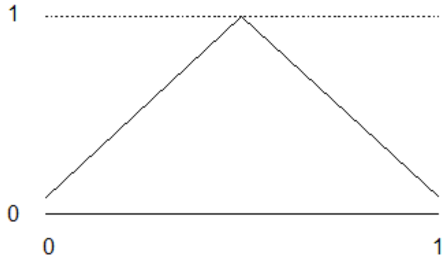

Figure 14 – Triangular Distribution

Figure 15 depicts an example of a Cmask sub-score file that will generate a Step 1 Csound score file. The Csound score file generated is then referenced by the Csound rendering engine, in conjunction with a Csound orchestra file, to generate a sound file. It is beyond the scope of this document to detail the low-level underpinnings of Cmask, however an overview of a few aspects of the syntax involved in its implementation follows.

To begin, it should be noted that semicolons are comment designators in Cmask syntax. Moving on, line 3 stipulates the tempo of the score file, i.e. tempo = 60 beats per second, beginning at time 0. Though this, perhaps archaic, perspective of musical notation has long been abandoned by the author, it is here simply considered a point of reference. Line 8 stipulates that events generated, each of which is intended for the Csound orchestra instrument designated in the following line, will begin at time 0. Events are to be generated for fifteen seconds of score time. Line 10 stipulates a value of 1 for the first pfield, p1, in the Csound score. This means that the score file generated is intended for instrument 1 as it exists in the Csound orchestra. Pfield 2 is stipulated by the next set of statements. It constrains start time values with a triangular distribution. Beginning at time 0 the range of values is between 2 and 4. The prec 2 is syntax that stipulates that the values generated will have a precision of two decimal points. The subsequent statements are similar to this one with the exception of the last one. It stipulates a quasi-random, uniformly distributed, selection of the values 1, 2, 3, and 4 for P6, in the score. This stipulates the sound file number used for each event. The sound files are named soundin.1, soundin.2, soundin.3, and soundin.4 and they are the samples that were created at the beginning of this process and stored in the SSDIR directory.

Incidentally, every time the Cmask program is instantiated it renders a different set of values for a new score file. So, after listening, if the values rendered do not yield desirable results when Csound renders the sound file, the Cmask process can be run again and then Csound can render the new sound file based upon the same constraints. If the results are consistently undesirable or if a change in their output is desired for any other reason, it is a simple matter to change the sub-score and run Cmask again. It is this process that yields the desired number of Csound scores and subsequently the desired number of sound files to be generated for the Step One portion of the overall process as stipulated by the schema whether created formally, before the rendering process, or in real-time during it.

Figure 16 below is a screenshot of a command line execution of Cmask. As is clear, it generates values for each of 6 pfields for 5 events each with an approximate cumulative 15 second duration.

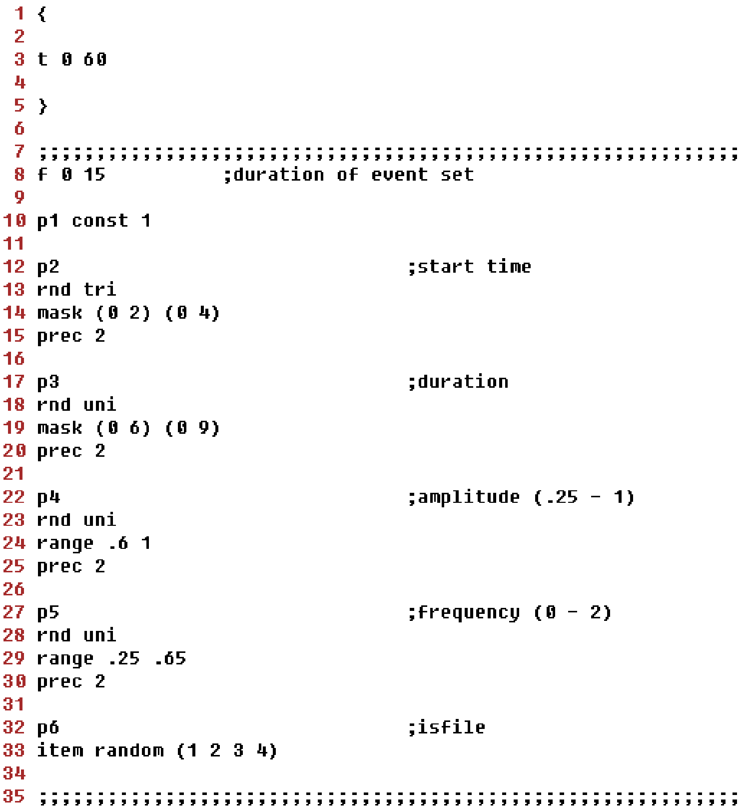

```
{

t 0 60

}

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
f 0 15          ;duration of event set

p1 const 1

p2                              ;start time
rnd tri
mask (0 2) (0 4)
prec 2

p3                              ;duration
rnd uni
mask (0 6) (0 9)
prec 2

p4                              ;amplitude (.25 - 1)
rnd uni
range .6 1
prec 2

p5                              ;frequency (0 - 2)
rnd uni
range .25 .65
prec 2

p6                              ;isfile
item random (1 2 3 4)

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
```

Figure 15 – Cmask Sub-score Step One

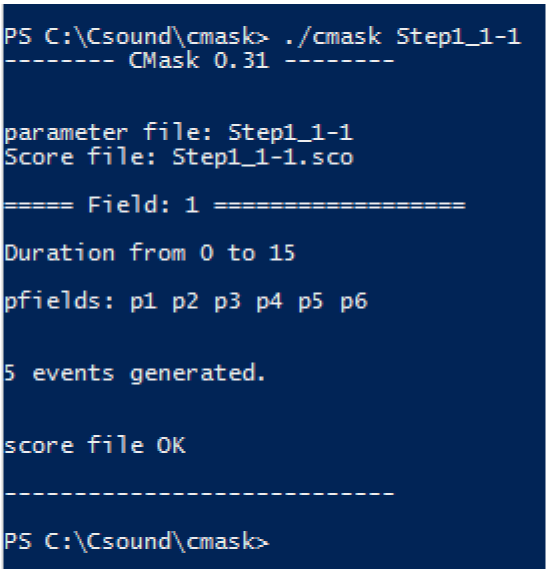

```
PS C:\Csound\cmask> ./cmask Step1_1-1
-------- CMask 0.31 --------


parameter file: Step1_1-1
Score file: Step1_1-1.sco

===== Field: 1 ==================

Duration from 0 to 15

pfields: p1 p2 p3 p4 p5 p6


5 events generated.


score file OK

----------------------------

PS C:\Csound\cmask>
```

Figure 16 – Executing Cmask from Command Line

2.9 Step One Score(s)

The final aspect of Step One to consider, before rendering the sound files, is the Csound score file that is generated by the Cmask sub-score. Figure 17 below is an example of a score file that might be generated by the Cmask process. Referencing section 2.2 above, it should be self-explanatory, however, to reiterate:

- pfield 1 = instrument number
- pfield 2 = start time
- pfield 3 = duration
- pfield 4 = amplitude
- pfield 5 = frequency
- pfield 6 = sample number

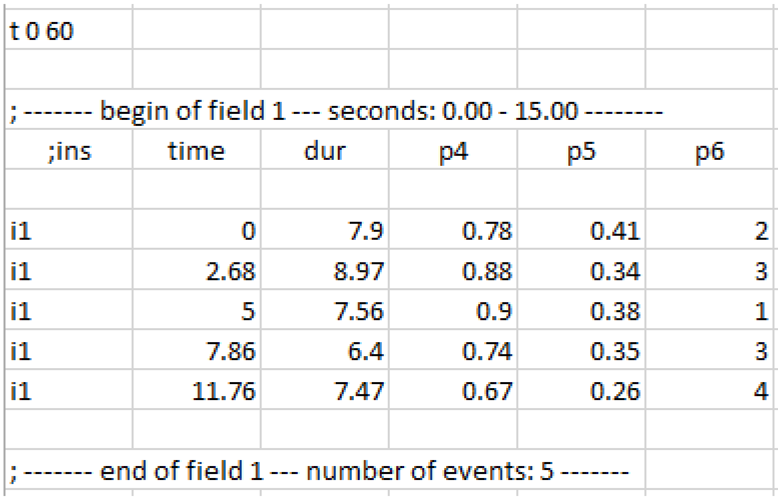

t 0 60					
; ------- begin of field 1 --- seconds: 0.00 - 15.00 --------					
;ins	time	dur	p4	p5	p6
i1	0	7.9	0.78	0.41	2
i1	2.68	8.97	0.88	0.34	3
i1	5	7.56	0.9	0.38	1
i1	7.86	6.4	0.74	0.35	3
i1	11.76	7.47	0.67	0.26	4
; ------- end of field 1 --- number of events: 5 -------					

Figure 17 – Example of Step One Csound Score Generated by Cmask Sub-score

2.10 Step One Orchestra

Figure 6, in section 2.2 above, illustrates and example of the Step One orchestra file that could be rendered by Csound in conjunction with the score file that was generated by Cmask. Since it was briefly explained there, it is not discussed here further. Suffice it to say, however, that Step One is extremely simplistic compared to Step Two though it provides an appropriate starting point toward a thorough apprehension of the process.

2.11 Secondary Samples

Now that a Step One Csound score has been generated utilizing Cmask and the Step One orchestra is complete, it is time to render a sound file. Though, as previously stated, a Csound front end is not used, a GUI is. This makes it quick and easy to change the referenced file names and to seamlessly move from one file to the next while rendering. Nonetheless it should be noted that the rendering could be accomplished with a Csound frontend or from command line. It is a matter of personal preference.

The GUI, shown below in Fig. 18, is called WinSound. The file paths are of great importance at this point in the process. Clicking the yellow "Render" button instantiates the Csound render engine. Referring to the file name for the orchestra and the score, it renders the resulting sound file to the Output File path. On the right is the output window that shows Csound in action. Any errors are printed here, which is extremely helpful when troubleshooting code. For instance, in this example there were no sound files in the SSDIR directory that were referenced in the score, which the output is complaining about. Another good use of this window is in the indented lines, near the end, that designate the "overall amps" and "overall samples out of range". One should be aware of the values here in order to avoid hearing damage especially if headphones are being used. Any values above 0 in the "overall samples out of range" line will cause aliasing in the sound file produced and can damage one's ears. Watching that line when every render is finished will act as a safeguard against hearing loss brought about by excessive amplitude.

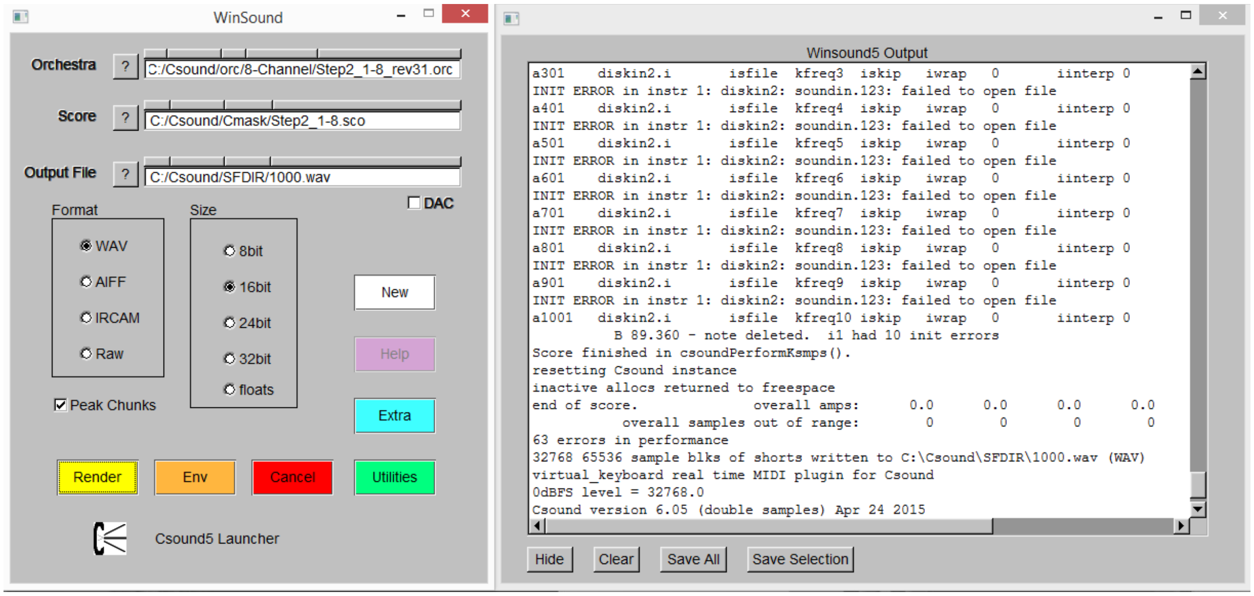

Orchestra: C:/Csound/orc/8-Channel/Step2_1-8_rev31.orc

Score: C:/Csound/Cmask/Step2_1-8.sco

Output File: C:/Csound/SFDIR/1000.wav

☐ DAC

Format: ● WAV ○ AIFF ○ IRCAM ○ Raw

Size: ○ 8bit ● 16bit ○ 24bit ○ 32bit ○ floats

☑ Peak Chunks

```
a301    diskin2.i       isfile  kfreq3  iskip   iwrap   0       iinterp 0
INIT ERROR in instr 1: diskin2: soundin.123: failed to open file
a401    diskin2.i       isfile  kfreq4  iskip   iwrap   0       iinterp 0
INIT ERROR in instr 1: diskin2: soundin.123: failed to open file
a501    diskin2.i       isfile  kfreq5  iskip   iwrap   0       iinterp 0
INIT ERROR in instr 1: diskin2: soundin.123: failed to open file
a601    diskin2.i       isfile  kfreq6  iskip   iwrap   0       iinterp 0
INIT ERROR in instr 1: diskin2: soundin.123: failed to open file
a701    diskin2.i       isfile  kfreq7  iskip   iwrap   0       iinterp 0
INIT ERROR in instr 1: diskin2: soundin.123: failed to open file
a801    diskin2.i       isfile  kfreq8  iskip   iwrap   0       iinterp 0
INIT ERROR in instr 1: diskin2: soundin.123: failed to open file
a901    diskin2.i       isfile  kfreq9  iskip   iwrap   0       iinterp 0
INIT ERROR in instr 1: diskin2: soundin.123: failed to open file
a1001   diskin2.i       isfile  kfreq10 iskip   iwrap   0       iinterp 0
          B 89.360 - note deleted.  i1 had 10 init errors
Score finished in csoundPerformKsmps().
resetting Csound instance
inactive allocs returned to freespace
end of score.              overall amps:      0.0       0.0       0.0       0.0
           overall samples out of range:        0         0         0         0
63 errors in performance
32768 65536 sample blks of shorts written to C:\Csound\SFDIR\1000.wav (WAV)
virtual_keyboard real time MIDI plugin for Csound
0dBFS level = 32768.0
Csound version 6.05 (double samples) Apr 24 2015
```

Figure 18 – Csound Rendering Using the WinSound GUI

If the Schema from Fig. 11 above was being used, the Step One (base) samples would have been the basis of the 80 Step Two samples created. These samples are morphologies of various combinations of the Step One samples that occur over a 15-second-long interval. At this point the samples are all still monophonic, single channel, sound files.

2.12 Waveform Editing

Before proceeding further, it is often prudent to import all 80 sound files into Sound Forge®, which is a professional audio editing program, for an inspection. It is important to locate any peaks that rise above 0dB. Additionally, as accomplished with the original base sample editing, this is the time to establish a general amplitude for each sample of around -6dB in order to provide an equal basis for comparison for between the samples during the next step in the process. After these and any other necessary adjustments are made, it is time to proceed to Step Two.

2.13 Step Two Orchestra

The Step Two orchestra is the result of ~20 years of expansion and refinement. It is a complex instrument consisting of nearly 10,000 lines of code. Going through it in its entirety would be impractical for this document. Instead it will be divided into sub-sections and summarized accordingly. It is referenced when generating the tertiary samples that are ultimately edited into the final composition. The process is again one of reiteration. Here the samples resulting from Step One are mixed in a wide variety of ways with others from the same group and used to generate several ~90-second-long, eight channel samples.

Figure 19 is a screenshot from the beginning of the Step Two orchestra file. As is directly apparent there are 41 pfields. These pfield are basically variables the values of which are determined by the score on an event by event basis. In general, pfields p4 – p20 are each associated with audio qualities of the sound events and pfields 21 – 41 are associated with the spatial characteristics of each event.

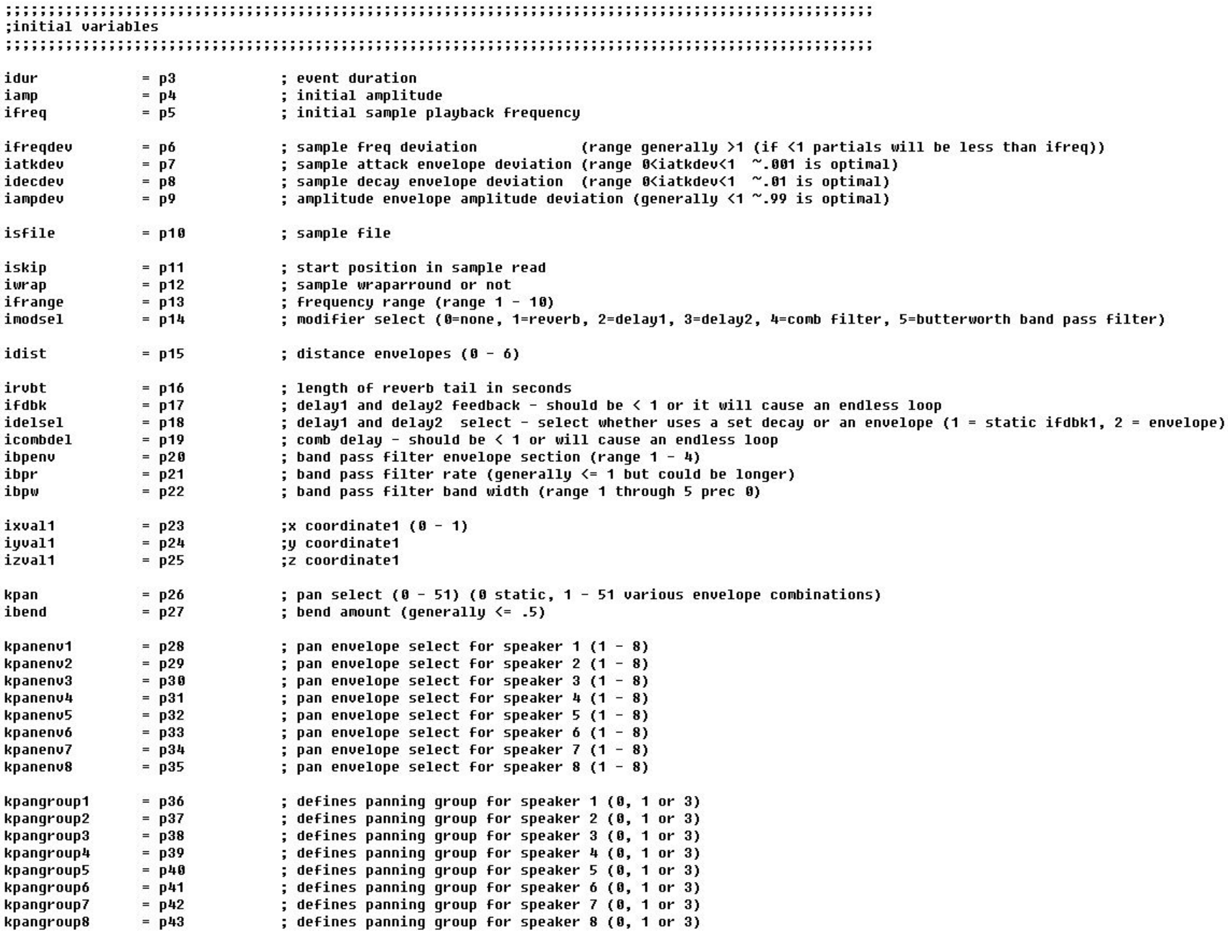

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;initial variables
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

idur            = p3            ; event duration
iamp            = p4            ; initial amplitude
ifreq           = p5            ; initial sample playback frequency

ifreqdev        = p6            ; sample freq deviation            (range generally >1 (if <1 partials will be less than ifreq))
iatkdev         = p7            ; sample attack envelope deviation (range 0<iatkdev<1  ~.001 is optimal)
idecdev         = p8            ; sample decay envelope deviation  (range 0<iatkdev<1  ~.01 is optimal)
iampdev         = p9            ; amplitude envelope amplitude deviation (generally <1 ~.99 is optimal)

isfile          = p10           ; sample file

iskip           = p11           ; start position in sample read
iwrap           = p12           ; sample wraparround or not
ifrange         = p13           ; frequency range (range 1 - 10)
imodsel         = p14           ; modifier select (0=none, 1=reverb, 2=delay1, 3=delay2, 4=comb filter, 5=butterworth band pass filter)

idist           = p15           ; distance envelopes (0 - 6)

irvbt           = p16           ; length of reverb tail in seconds
ifdbk           = p17           ; delay1 and delay2 feedback - should be < 1 or it will cause an endless loop
idelsel         = p18           ; delay1 and delay2  select - select whether uses a set decay or an envelope (1 = static ifdbk1, 2 = envelope)
icombdel        = p19           ; comb delay - should be < 1 or will cause an endless loop
ibpenv          = p20           ; band pass filter envelope section (range 1 - 4)
ibpr            = p21           ; band pass filter rate (generally <= 1 but could be longer)
ibpw            = p22           ; band pass filter band width (range 1 through 5 prec 0)

ixval1          = p23           ;x coordinate1 (0 - 1)
iyval1          = p24           ;y coordinate1
izval1          = p25           ;z coordinate1

kpan            = p26           ; pan select (0 - 51) (0 static, 1 - 51 various envelope combinations)
ibend           = p27           ; bend amount (generally <= .5)

kpanenv1        = p28           ; pan envelope select for speaker 1 (1 - 8)
kpanenv2        = p29           ; pan envelope select for speaker 2 (1 - 8)
kpanenv3        = p30           ; pan envelope select for speaker 3 (1 - 8)
kpanenv4        = p31           ; pan envelope select for speaker 4 (1 - 8)
kpanenv5        = p32           ; pan envelope select for speaker 5 (1 - 8)
kpanenv6        = p33           ; pan envelope select for speaker 6 (1 - 8)
kpanenv7        = p34           ; pan envelope select for speaker 7 (1 - 8)
kpanenv8        = p35           ; pan envelope select for speaker 8 (1 - 8)

kpangroup1      = p36           ; defines panning group for speaker 1 (0, 1 or 3)
kpangroup2      = p37           ; defines panning group for speaker 2 (0, 1 or 3)
kpangroup3      = p38           ; defines panning group for speaker 3 (0, 1 or 3)
kpangroup4      = p39           ; defines panning group for speaker 4 (0, 1 or 3)
kpangroup5      = p40           ; defines panning group for speaker 5 (0, 1 or 3)
kpangroup6      = p41           ; defines panning group for speaker 6 (0, 1 or 3)
kpangroup7      = p42           ; defines panning group for speaker 7 (0, 1 or 3)
kpangroup8      = p43           ; defines panning group for speaker 8 (0, 1 or 3)
```

Figure 19 – Step Two Orchestra Pfields

Csound parses through the orchestra for each event sequentially, in a linear manner, from top to bottom. Each variable with an "i" at the beginning of its name is instantiated during the initial parsing of the orchestra allowing each i-rate variable to be established upon the initial rendering pass only. The k-rate variables are parsed according to its related stipulation in the orchestra file header which is directly proportionate to the sample rate. The a-rate variables are parsed at the sample rate.

The use of code that branches is an effective manner with which to utilize variables in the score to instantiate various options afforded in each branch of code in the orchestra. In Fig. 20 the top to bottom progression of parsing first defines line segment envelopes kenv400 – kenv405 and subsequently variables kenv43 – kenv40. Then begins the process of defining kdampenv, which will afterward be multiplied times the audio output a1000 to define a new audio output… a1001. Then begins a series of if/else statements that could be considered to work their way back up the code. That is not what actually happens because the Csound render engine only parses from top to bottom but to human perception this is an effective manner with which to apprehend the code. In verbal pseudo code we would say the flow follows this path beginning at the bottom: Does kdist, pfield 15 in the score, equal 0? If so then kdampenv equals kenv400. If not (else) then it equals kenv40. Then, kenv40 asks, does kdist equal 1? If so then kdampenv equals 401, else it equals kenv42… so on until the final statement that if kdist does not equal 1, 2, 3, or 4 then kdampenv equals kenv405. This approach to branching is extremely effective and takes advantage of the fact that Csound parses from top to bottom. It is used in various sections of the Step Two orchestra.

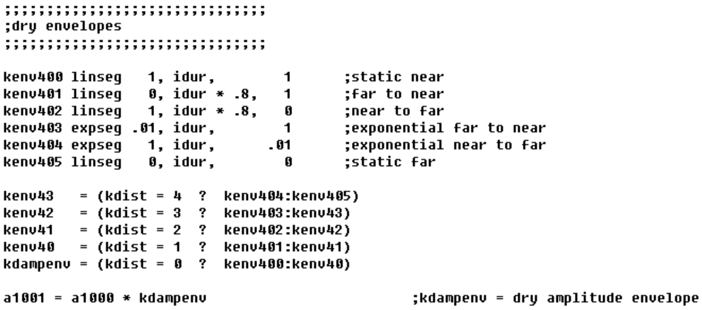

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;dry envelopes
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

kenv400 linseg   1, idur,        1      ;static near
kenv401 linseg   0, idur * .8,   1      ;far to near
kenv402 linseg   1, idur * .8,   0      ;near to far
kenv403 expseg .01, idur,        1      ;exponential far to near
kenv404 expseg   1, idur,      .01      ;exponential near to far
kenv405 linseg   0, idur,        0      ;static far

kenv43   = (kdist = 4  ?  kenv404:kenv405)
kenv42   = (kdist = 3  ?  kenv403:kenv43)
kenv41   = (kdist = 2  ?  kenv402:kenv42)
kenv40   = (kdist = 1  ?  kenv401:kenv41)
kdampenv = (kdist = 0  ?  kenv400:kenv40)

a1001 = a1000 * kdampenv                         ;kdampenv = dry amplitude envelope
```

Figure 20 – Csound Code Excerpt Demonstrating Branching

For each event the value in pfield 13, between 1 and 10, determines the number of samples that will be played back for any given event. This layering of samples provides a wide variety of timbral characteristics ranging from an enhanced overtone series to a chorusing effect and to other effects that must be heard to comprehend. The values of p6 – p9 determine the characteristics of each layer's amplitude envelopes and frequency envelopes. Figure 21 provides an excerpt of an example of the code related to this. Notice the series of “if” statements, beginning with “out1:”, which is another manner in which to branch the code.

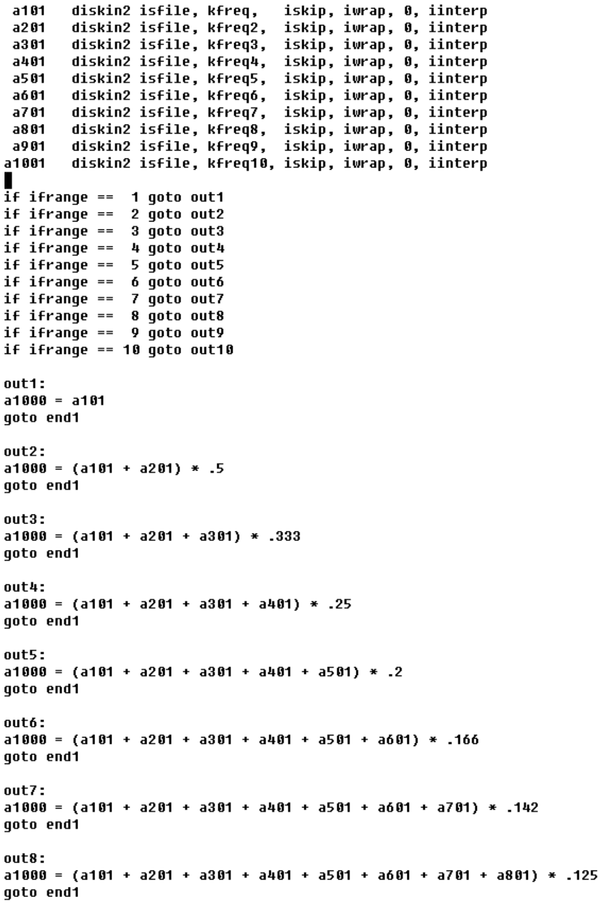

```
 a101   diskin2 isfile, kfreq,   iskip, iwrap, 0, iinterp
 a201   diskin2 isfile, kfreq2,  iskip, iwrap, 0, iinterp
 a301   diskin2 isfile, kfreq3,  iskip, iwrap, 0, iinterp
 a401   diskin2 isfile, kfreq4,  iskip, iwrap, 0, iinterp
 a501   diskin2 isfile, kfreq5,  iskip, iwrap, 0, iinterp
 a601   diskin2 isfile, kfreq6,  iskip, iwrap, 0, iinterp
 a701   diskin2 isfile, kfreq7,  iskip, iwrap, 0, iinterp
 a801   diskin2 isfile, kfreq8,  iskip, iwrap, 0, iinterp
 a901   diskin2 isfile, kfreq9,  iskip, iwrap, 0, iinterp
a1001   diskin2 isfile, kfreq10, iskip, iwrap, 0, iinterp

if ifrange ==  1 goto out1
if ifrange ==  2 goto out2
if ifrange ==  3 goto out3
if ifrange ==  4 goto out4
if ifrange ==  5 goto out5
if ifrange ==  6 goto out6
if ifrange ==  7 goto out7
if ifrange ==  8 goto out8
if ifrange ==  9 goto out9
if ifrange == 10 goto out10

out1:
a1000 = a101
goto end1

out2:
a1000 = (a101 + a201) * .5
goto end1

out3:
a1000 = (a101 + a201 + a301) * .333
goto end1

out4:
a1000 = (a101 + a201 + a301 + a401) * .25
goto end1

out5:
a1000 = (a101 + a201 + a301 + a401 + a501) * .2
goto end1

out6:
a1000 = (a101 + a201 + a301 + a401 + a501 + a601) * .166
goto end1

out7:
a1000 = (a101 + a201 + a301 + a401 + a501 + a601 + a701) * .142
goto end1

out8:
a1000 = (a101 + a201 + a301 + a401 + a501 + a601 + a701 + a801) * .125
goto end1
```

Figure 21 – Code Excerpt of Sample Playback Section

Pfields 10 – 12 determine aspects of the sample playback such as the sample file to be used, the time within the sample from which playback is instantiated, and whether or not the sample playback loops back to the beginning of the sample if the duration of the event is longer than the sample considering its playback frequency. For example, a playback frequency of 1 plays the sample at its original pitch and for its original duration. A frequency of .5 plays the sample an octave lower and for twice the original duration. A frequency of 2 plays the sample back at an octave higher than the original sample with 50% of the duration.

Pfield 14 determines which, if any, effects are used in conjunction with the samples. They include reverb, delay, chorusing, comb filtering, and band pass filtering. Pfields 16 – 20 determine specific parametric characteristics of each of the effects. Figure 22 demonstrates the use of if/then statements in series to define the choice of effect to be used for the sound event including the choice of no effect.

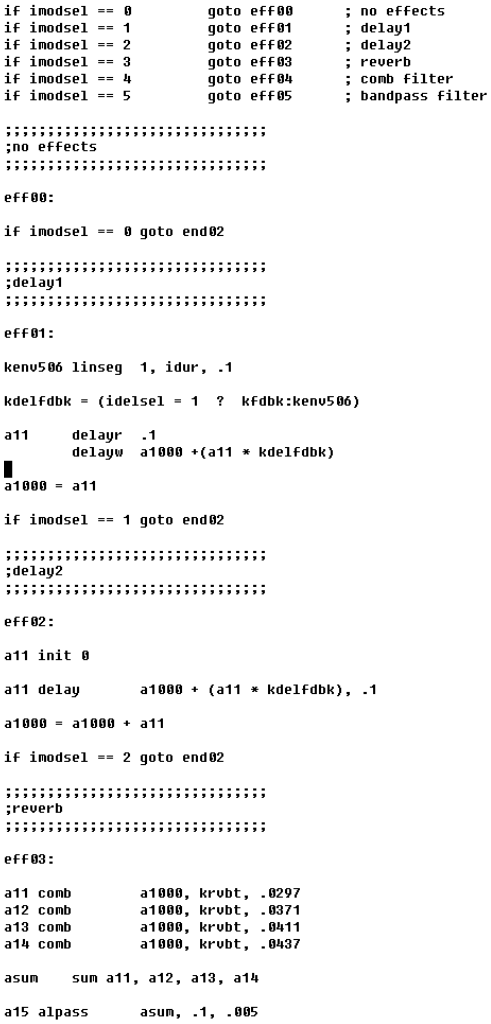

```
if imodsel == 0         goto eff00      ; no effects
if imodsel == 1         goto eff01      ; delay1
if imodsel == 2         goto eff02      ; delay2
if imodsel == 3         goto eff03      ; reverb
if imodsel == 4         goto eff04      ; comb filter
if imodsel == 5         goto eff05      ; bandpass filter

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;no effects
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

eff00:

if imodsel == 0 goto end02

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;delay1
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

eff01:

kenv506 linseg  1, idur, .1

kdelfdbk = (idelsel = 1  ?  kfdbk:kenv506)

a11     delayr  .1
        delayw  a1000 +(a11 * kdelfdbk)

a1000 = a11

if imodsel == 1 goto end02

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;delay2
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

eff02:

a11 init 0

a11 delay       a1000 + (a11 * kdelfdbk), .1

a1000 = a1000 + a11

if imodsel == 2 goto end02

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;reverb
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

eff03:

a11 comb        a1000, krvbt, .0297
a12 comb        a1000, krvbt, .0371
a13 comb        a1000, krvbt, .0411
a14 comb        a1000, krvbt, .0437

asum    sum a11, a12, a13, a14

a15 alpass      asum, .1, .005
```

Figure 22 – Code Excerpt of Effects Section

Pfield 15 in the score determines the characteristics of distance algorithms. Please note the use of the Csound opcode “pconvolve”, which implements convolution reverb. Again, extensive use of if/else statements branches the code to access the various options based upon values in the score.

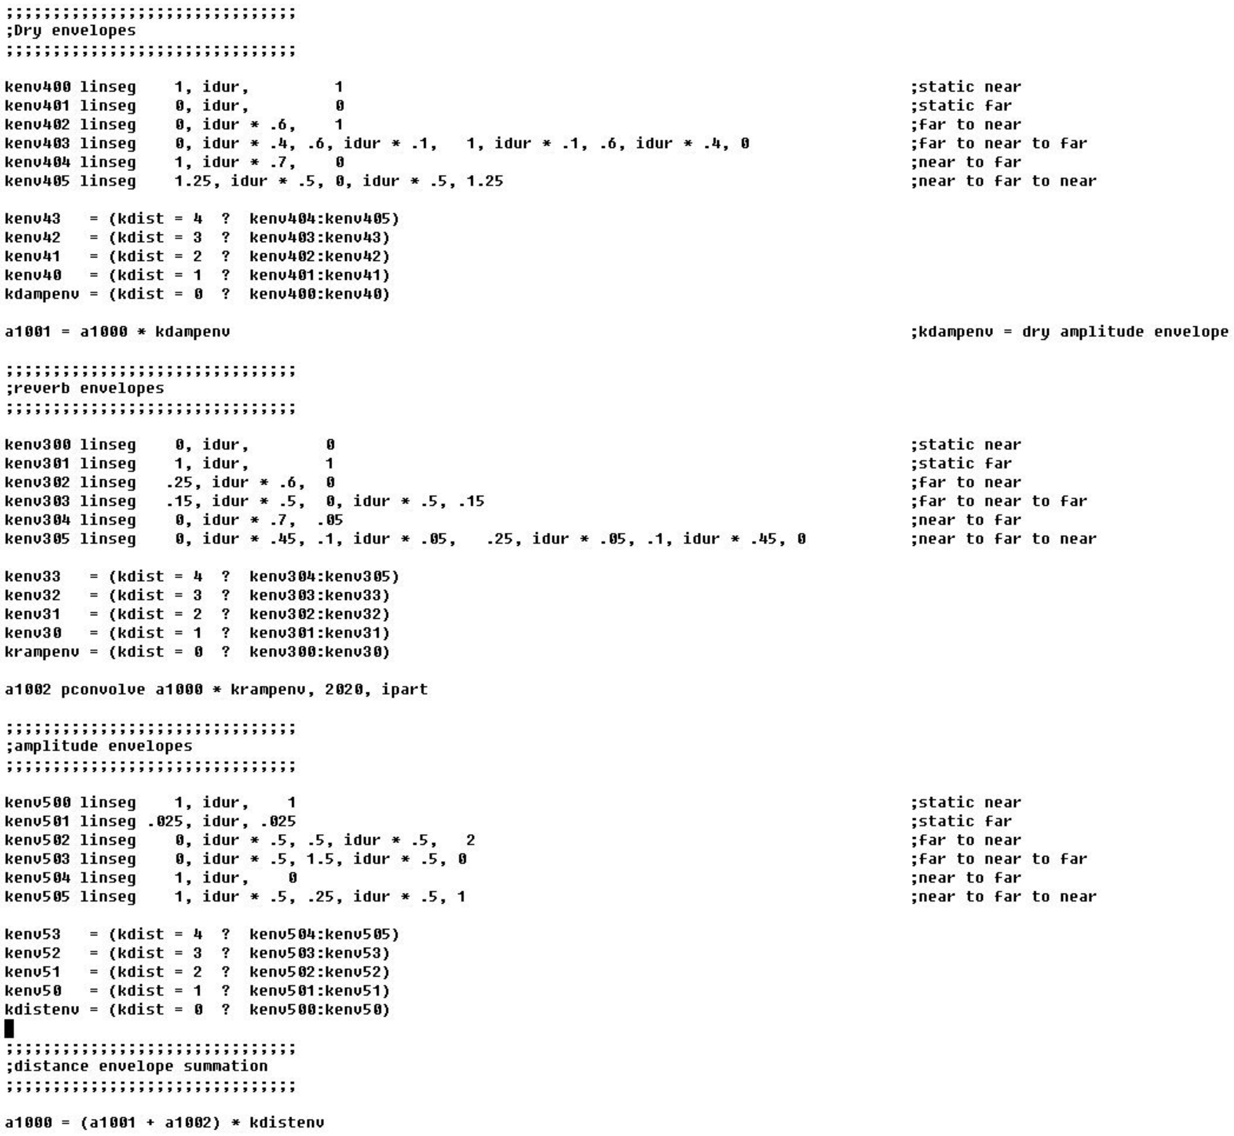

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;Dry envelopes
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

kenv400 linseg    1, idur,         1                                                            ;static near
kenv401 linseg    0, idur,         0                                                            ;static far
kenv402 linseg    0, idur * .6,    1                                                            ;far to near
kenv403 linseg    0, idur * .4, .6, idur * .1,   1, idur * .1, .6, idur * .4, 0                 ;far to near to far
kenv404 linseg    1, idur * .7,    0                                                            ;near to far
kenv405 linseg    1.25, idur * .5, 0, idur * .5, 1.25                                           ;near to far to near

kenv43   = (kdist = 4  ?  kenv404:kenv405)
kenv42   = (kdist = 3  ?  kenv403:kenv43)
kenv41   = (kdist = 2  ?  kenv402:kenv42)
kenv40   = (kdist = 1  ?  kenv401:kenv41)
kdampenv = (kdist = 0  ?  kenv400:kenv40)

a1001 = a1000 * kdampenv                                                                        ;kdampenv = dry amplitude envelope

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;reverb envelopes
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

kenv300 linseg    0, idur,        0                                                             ;static near
kenv301 linseg    1, idur,        1                                                             ;static far
kenv302 linseg   .25, idur * .6,  0                                                             ;far to near
kenv303 linseg   .15, idur * .5,  0, idur * .5, .15                                             ;far to near to far
kenv304 linseg    0, idur * .7,  .05                                                            ;near to far
kenv305 linseg    0, idur * .45, .1, idur * .05,   .25, idur * .05, .1, idur * .45, 0           ;near to far to near

kenv33   = (kdist = 4  ?  kenv304:kenv305)
kenv32   = (kdist = 3  ?  kenv303:kenv33)
kenv31   = (kdist = 2  ?  kenv302:kenv32)
kenv30   = (kdist = 1  ?  kenv301:kenv31)
krampenv = (kdist = 0  ?  kenv300:kenv30)

a1002 pconvolve a1000 * krampenv, 2020, ipart

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;amplitude envelopes
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

kenv500 linseg    1, idur,    1                                                                 ;static near
kenv501 linseg .025, idur, .025                                                                 ;static far
kenv502 linseg    0, idur * .5, .5, idur * .5,   2                                              ;far to near
kenv503 linseg    0, idur * .5, 1.5, idur * .5, 0                                               ;far to near to far
kenv504 linseg    1, idur,    0                                                                 ;near to far
kenv505 linseg    1, idur * .5, .25, idur * .5, 1                                               ;near to far to near

kenv53   = (kdist = 4  ?  kenv504:kenv505)
kenv52   = (kdist = 3  ?  kenv503:kenv53)
kenv51   = (kdist = 2  ?  kenv502:kenv52)
kenv50   = (kdist = 1  ?  kenv501:kenv51)
kdistenv = (kdist = 0  ?  kenv500:kenv50)
█
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;distance envelope summation
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

a1000 = (a1001 + a1002) * kdistenv
```

Figure 23 – Distance Algorithms

Pfields 21 – 41 determine the vector-based amplitude panning characteristics of each sample playback on an event-by-event basis. Instead of using the vector-based amplitude panning (VBAP) opcodes that are a part of Csound, I created my own version of VBAP based upon the Cartesian coordinate structure shown in Fig. 24 below. As opposed to traditional VBAP that commonly uses a triangulation method to project the perception of the location of a given sound event, this version implements an octagonal approach.

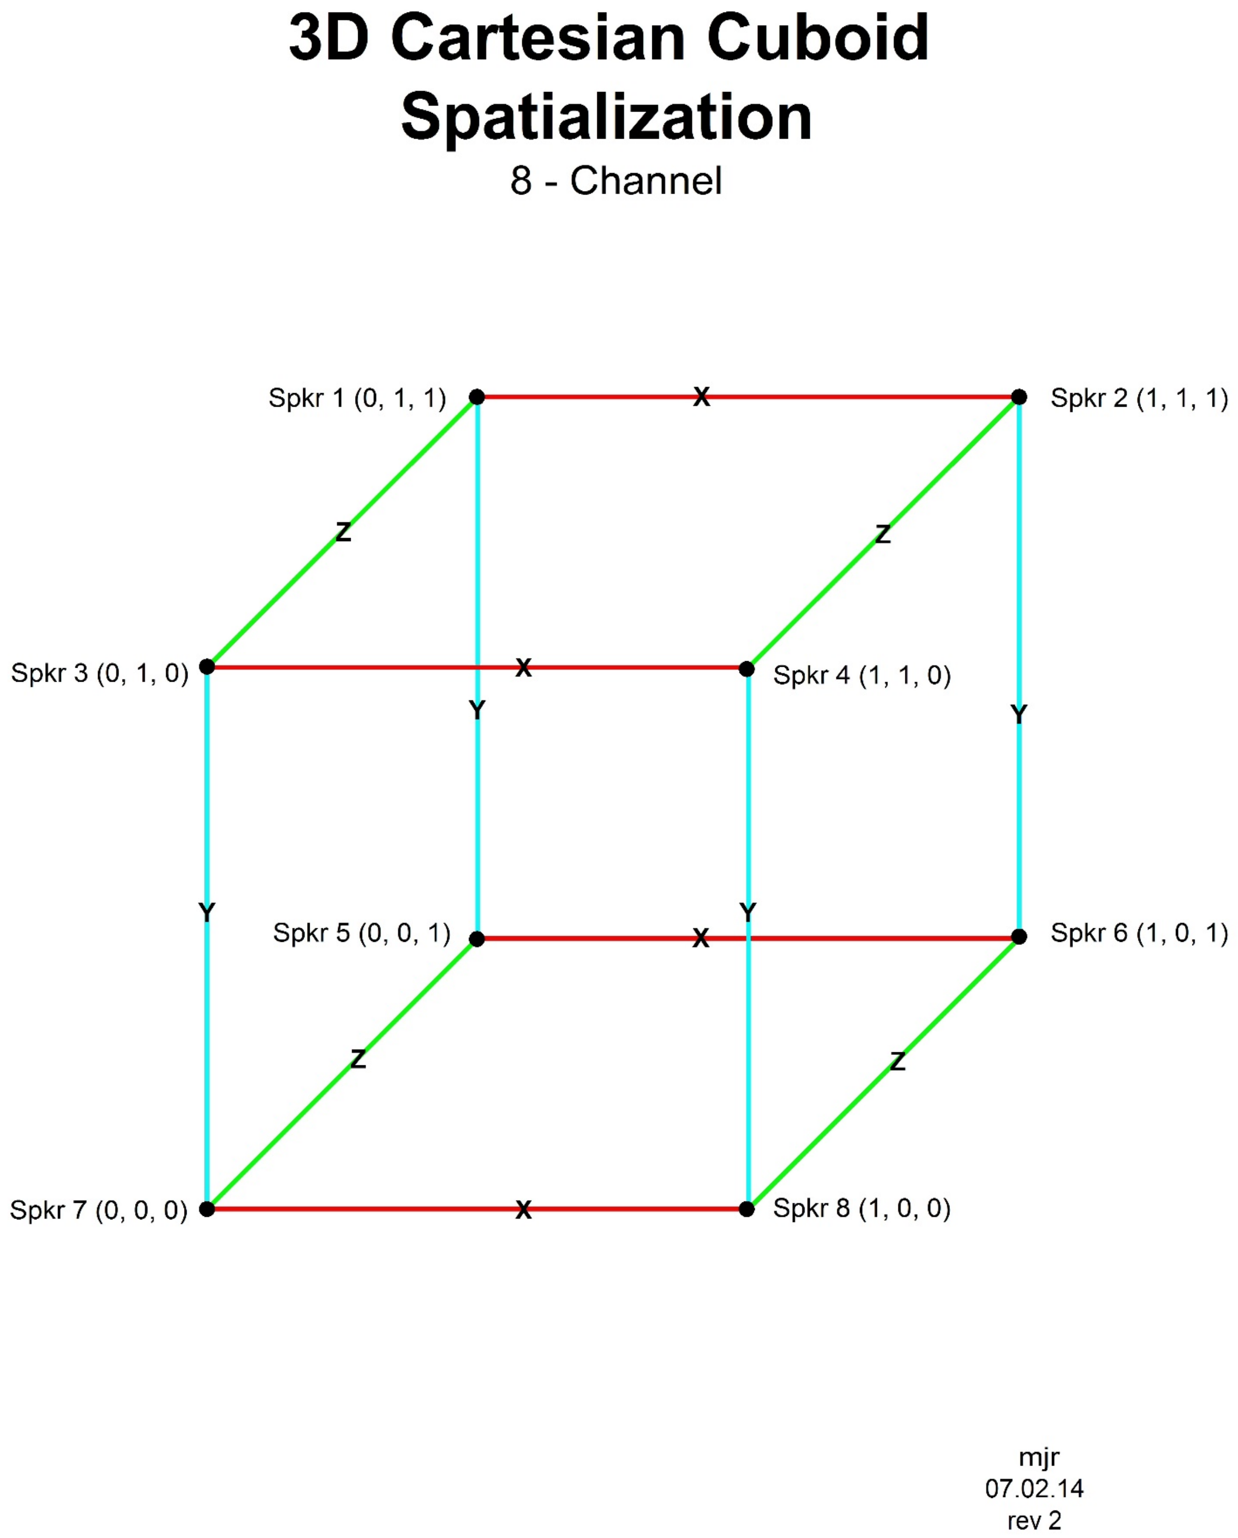

Figure 24 – Vector Based Amplitude Panning Structure

Figure 25 is a code excerpt from the beginning of the spatialization algorithm section. The first line of code could be read as follows; if ixval1 equals 0 and iyval1 and izval1 equals 1 the go to out001. In the section below, Out001: states that audio signal a1000 is sent to audio output a10000 only. Though this is a static audio output, the same general code structure is employed for the dynamic audio outputs throughout the panning algorithm section.

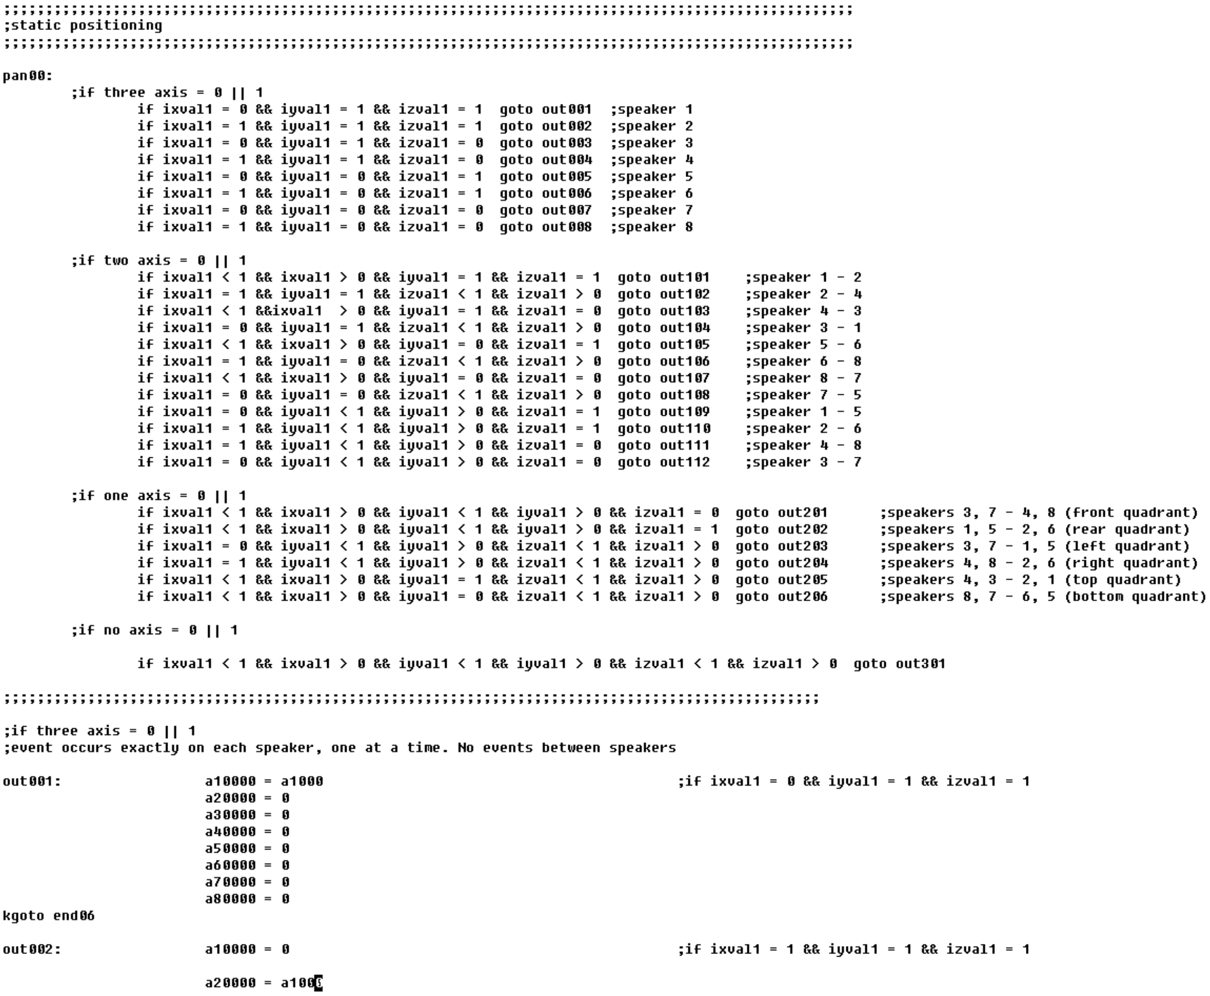

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;static positioning
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

pan00:
        ;if three axis = 0 || 1
                if ixval1 = 0 && iyval1 = 1 && izval1 = 1  goto out001  ;speaker 1
                if ixval1 = 1 && iyval1 = 1 && izval1 = 1  goto out002  ;speaker 2
                if ixval1 = 0 && iyval1 = 1 && izval1 = 0  goto out003  ;speaker 3
                if ixval1 = 1 && iyval1 = 1 && izval1 = 0  goto out004  ;speaker 4
                if ixval1 = 0 && iyval1 = 0 && izval1 = 1  goto out005  ;speaker 5
                if ixval1 = 1 && iyval1 = 0 && izval1 = 1  goto out006  ;speaker 6
                if ixval1 = 0 && iyval1 = 0 && izval1 = 0  goto out007  ;speaker 7
                if ixval1 = 1 && iyval1 = 0 && izval1 = 0  goto out008  ;speaker 8

        ;if two axis = 0 || 1
                if ixval1 < 1 && ixval1 > 0 && iyval1 = 1 && izval1 = 1  goto out101    ;speaker 1 - 2
                if ixval1 = 1 && iyval1 = 1 && izval1 < 1 && izval1 > 0  goto out102    ;speaker 2 - 4
                if ixval1 < 1 &&ixval1  > 0 && iyval1 = 1 && izval1 = 0  goto out103    ;speaker 4 - 3
                if ixval1 = 0 && iyval1 = 1 && izval1 < 1 && izval1 > 0  goto out104    ;speaker 3 - 1
                if ixval1 < 1 && ixval1 > 0 && iyval1 = 0 && izval1 = 1  goto out105    ;speaker 5 - 6
                if ixval1 = 1 && iyval1 = 0 && izval1 < 1 && izval1 > 0  goto out106    ;speaker 6 - 8
                if ixval1 < 1 && ixval1 > 0 && iyval1 = 0 && izval1 = 0  goto out107    ;speaker 8 - 7
                if ixval1 = 0 && iyval1 = 0 && izval1 < 1 && izval1 > 0  goto out108    ;speaker 7 - 5
                if ixval1 = 0 && iyval1 < 1 && iyval1 > 0 && izval1 = 1  goto out109    ;speaker 1 - 5
                if ixval1 = 1 && iyval1 < 1 && iyval1 > 0 && izval1 = 1  goto out110    ;speaker 2 - 6
                if ixval1 = 1 && iyval1 < 1 && iyval1 > 0 && izval1 = 0  goto out111    ;speaker 4 - 8
                if ixval1 = 0 && iyval1 < 1 && iyval1 > 0 && izval1 = 0  goto out112    ;speaker 3 - 7

        ;if one axis = 0 || 1
                if ixval1 < 1 && ixval1 > 0 && iyval1 < 1 && iyval1 > 0 && izval1 = 0  goto out201      ;speakers 3, 7 - 4, 8 (front quadrant)
                if ixval1 < 1 && ixval1 > 0 && iyval1 < 1 && iyval1 > 0 && izval1 = 1  goto out202      ;speakers 1, 5 - 2, 6 (rear quadrant)
                if ixval1 = 0 && iyval1 < 1 && iyval1 > 0 && izval1 < 1 && izval1 > 0  goto out203      ;speakers 3, 7 - 1, 5 (left quadrant)
                if ixval1 = 1 && iyval1 < 1 && iyval1 > 0 && izval1 < 1 && izval1 > 0  goto out204      ;speakers 4, 8 - 2, 6 (right quadrant)
                if ixval1 < 1 && ixval1 > 0 && iyval1 = 1 && izval1 < 1 && izval1 > 0  goto out205      ;speakers 4, 3 - 2, 1 (top quadrant)
                if ixval1 < 1 && ixval1 > 0 && iyval1 = 0 && izval1 < 1 && izval1 > 0  goto out206      ;speakers 8, 7 - 6, 5 (bottom quadrant)

        ;if no axis = 0 || 1

                if ixval1 < 1 && ixval1 > 0 && iyval1 < 1 && iyval1 > 0 && izval1 < 1 && izval1 > 0  goto out301

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

;if three axis = 0 || 1
;event occurs exactly on each speaker, one at a time. No events between speakers

out001:                 a10000 = a1000                                          ;if ixval1 = 0 && iyval1 = 1 && izval1 = 1
                        a20000 = 0
                        a30000 = 0
                        a40000 = 0
                        a50000 = 0
                        a60000 = 0
                        a70000 = 0
                        a80000 = 0
kgoto end06

out002:                 a10000 = 0                                              ;if ixval1 = 1 && iyval1 = 1 && izval1 = 1

                        a20000 = a1000
```

Figure 25 – Excerpt of Code from Panning Algorithms

To continue describing the diffusion aspects determined in this section of the process, the values in pfields 21 – 41 determine the perceived position of a given event and its perceived movement characteristics. Here the use of the word "perceived" is important because to actually determine the static or dynamic location of a sound it is necessary to physically locate or move the loudspeaker accordingly, which is generally very limited if at all feasible. This approach is called point source diffusion. Instead we provide the illusory perception of the static or dynamic location of a sound using techniques such as VBAP in addition to others described later to approach the perception of holophony.

With amplitude panning a variance in amplitude between two of more speakers provides the perception of a sound moving across the space between them. For instance, as a simple example, in Fig. 26 below, consider a listener seated in the center between four speakers. From the listener's position, known as the sweet spot, as the amplitude of speaker 1 decreases over the duration of an event and the amplitude increases inversely in speaker 3 the listener will perceive the sound event as moving from speaker 1 to speaker 3. Further, when the amplitude is multiplied by .5 in each of the speakers simultaneously, the sound event will be perceived to exist in the center between them… in the space between. Note that to achieve the modulation of the speaker amplitudes from 0, no amplitude, to 1, full amplitude, the amplitude of each is multiplied by a Csound envelope. Variables k1 and k2, also demonstrated in Fig. 20 below, stipulate line segment envelopes using the Csound linseg opcode.

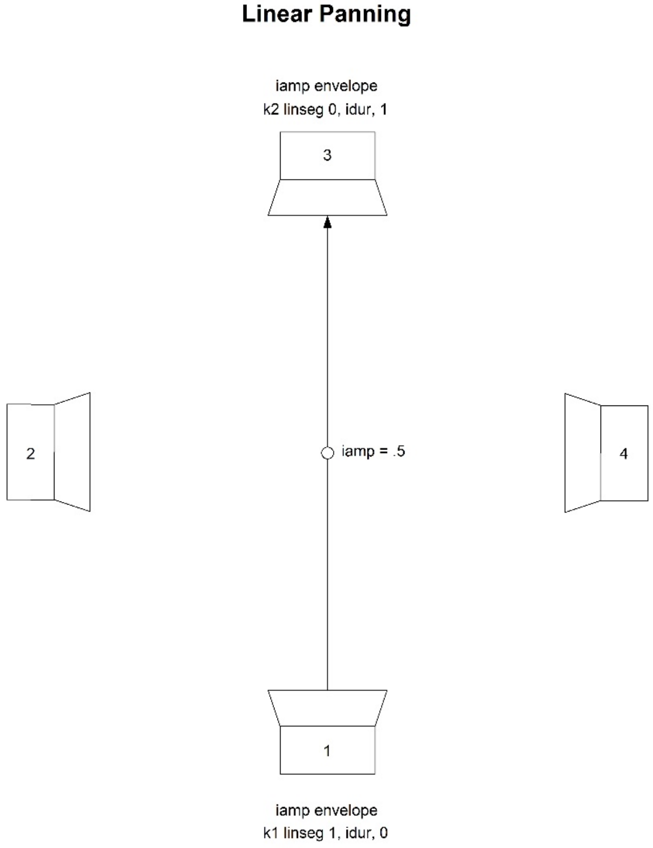

Figure 26 – Linear Amplitude Panning

Along with the perception of linear motion between speakers the notion was conceived of and then implemented to bend the perception of movement in an exponential manner. In this simple example the implementation of a third speaker is required. Though within a small listening venue and using only four speakers, the perceivable difference between a linear and exponential panning algorithm might be subtle, in a larger venue using a HDLA and augmenting the VBAP approach with Ambisonics, as is demonstrated later, it can be quite effective. Figure 27 illustrates the basic principle behind an exponential vector-based amplitude panning approach. In it as the amplitude in speakers 1 and 3 are modulated by amplitude envelopes, as with the linear panning example, during the duration of the event speaker 4 is also being modulated, acting as an attractor. This causes the perception of the location of the sound to circumvent the listener seated in the sweet spot.

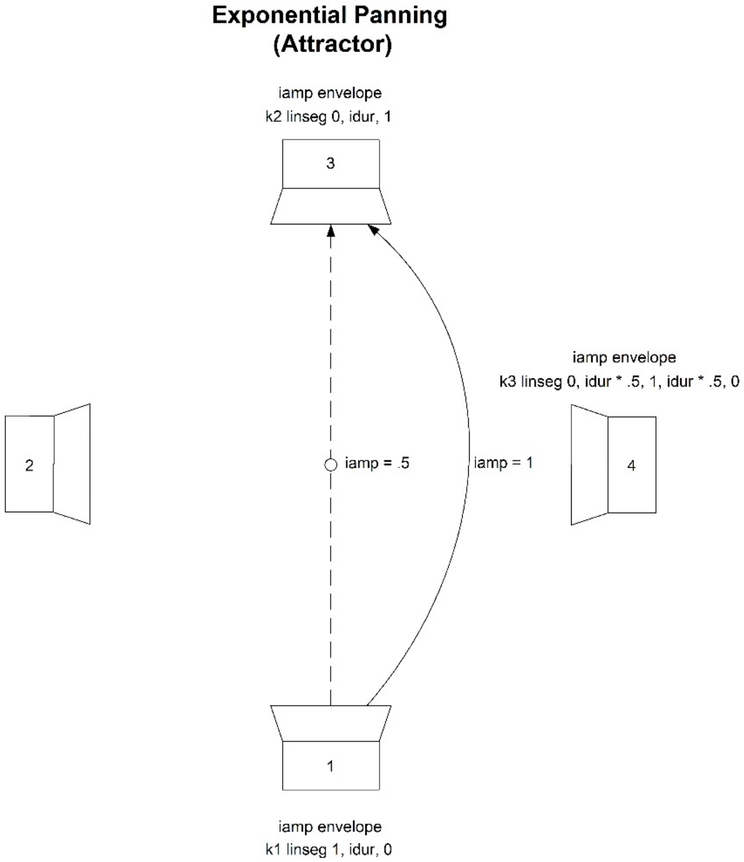

Figure 27 – Exponential Amplitude Panning Using an Attractor

Though simplistic, these examples provide a basic understanding of the panning algorithms implemented in the Step Two Csound orchestra. In it there are 50 potential patterns of perceivable movement, each of which can be applied to varying combinations of the 8 speakers. This provides a very wide variety of perceived patterns of movement within a composition, again, on an event by event basis. This approach is a significant aspect in the equation for creating the perception of holophony that will be discussed later in this document. When using eight speakers, or 32 as we shall see later, the possibilities for perceived patterns of movement exponentially increase. It creates the basis and potential for very rich spatial listening environments and experiences.

The versatility of this orchestra in tandem with the use of innumerable combinations of base samples provides an incalculable potential for musical results. Combined with the upcoming generative score techniques it this composer's dream come true.

2.14 Step Two Sub-score(s)

The Step Two subscore, though similar in principle to the Step One subscore, is the extremely complex powerhouse behind the formal and expressive aspects of this compositional process. Through numerous sequences of iterations of the subscores and subsequent listening, and through the compositionally aware direction of the tendency masks in each by the composer, a naturally occurring musical form emerges. Through this process, new musical material that is related and yet varied from the previously generated material, is constantly and consistently being generated. These variations are the very essence of the joy and expression of musical composition through the centuries and yet here being intensely augmented by the computer.

Fig. 28 below prvides an example of a Step Two subscore. Since there are 41 pfields in the Step Two orchestra the subscore, as opposed to six in the Step One orchestra, it is much more complex to implement. The comments in the subscore provide sufficient description of each pfield range so it will not be further elucidated here. The idea is to create multiple 8-channel sound files, ~90 seconds in length, using the 80 samples generated in Step One. This is a further mixing and morphing of them that often affords dynamic emergent results.

```
{
t 0 60
}
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
f 0 90                                  ;duration of event set
p1 const 1

p2                                      ;start time
```

```
rnd uni
mask (0 1.06) (0 1.9)
prec 2

p3                                      ;duration
rnd uni
mask (0 2) (0 4)
prec 2

p4                                      ;amplitude (.25 - 1)
rnd uni
range .5 2

p5                                      ;frequency (0 - 2)
rnd uni
range .48 .78
prec 2

p6                                      ;ifreqdev (generally close to 1)
rnd uni
range .95 .99

p7                                      ;iatkdev
item random (.001 .002 .003 .004)

p8                                      ;idecdev
item random (.01 .011 .012 .013 .014)

p9                                      ;iampdev
rnd uni
range .95 1
prec 2

p10                                     ;isfile
rnd uni
item rand (122 123 124 125 126)

p11  const 0                            ;iskip
;range 0 3
;prec 2

p12         ;const 1                    ;iwrap
item random (0 1)

p13                                     ;ifrange
rnd uni
range 1 10
prec 0

p14                                     ;imodsel
item random (0 1 1 2 2 3 3 3 4 4 4 5 5 5 5)

p15                                     ;idist
item random (0 0 0 0 1 1 2 3 4 5)

p16                                     ; length of reverb tail in seconds
rnd uni
range 0 2
prec 2

p17                                     ; delay1 and delay2 feedback - should be < 1 or it will cause an endless loop
rnd uni
range 0 .99
prec 2

p18                                     ; delay1 and delay2  select - select whether uses a set decay or an envelope (1 = static ifdbk1, 2 =
envelope)
item rand (1, 2)

p19                                     ; comb delay - should be < 1 or will cause an endless loop
```

```
rnd uni
range 0 .99
prec 2

p20                                ;band pass filter envelope section (range 1 - 4)
rnd uni
range 1 4
prec 2

p21                                ;x coordinate1 (0 - 1)
rnd uni
range 0 1
prec 2

p22                                ;y coordinate1 (0 - 1)
rnd uni
range 0 1
prec 2

p23                                ;z coordinate1 (0 - 1)
rnd uni
range 0 1
prec 2

p24                                ;pan select (0 - 51) (0 static, 1 - 51 various envelope combinations)
rnd uni
range 0 51
prec 0

p25                                ;bend amount (generally <= .5)
rnd exp
range 0 1
prec 2

p26                                ;pan envelope select for speaker 1 (1 - 8)
rnd uni
range 1 8
prec 0

p27                                ;pan envelope select for speaker 2 (1 - 8)
rnd uni
range 1 8
prec 0

p28                                ;pan envelope select for speaker 3 (1 - 8)
rnd uni
range 1 8
prec 0

p29                                ;pan envelope select for speaker 4 (1 - 8)
rnd uni
range 1 8
prec 0

p30                                ;pan envelope select for speaker 5 (1 - 8)
rnd uni
range 1 8
prec 0

p31                                ;pan envelope select for speaker 6 (1 - 8)
rnd uni
range 1 8
prec 0

p32                                ;pan envelope select for speaker 7 (1 - 8)
rnd uni
range 1 8
prec 0

p33                                ;pan envelope select for speaker 8 (1 - 8)
```

```
rnd uni
range 1 8
prec 0

p34                                   ;defines panning group for speaker 1 (0, 1 or 2)
item rand (0, 1, 2)

p35                                   ;defines panning group for speaker 2 (0, 1 or 2)
item rand (0, 1, 2)

p36                                   ;defines panning group for speaker 3 (0, 1 or 2)
item rand (0, 1, 2)

p37                                   ;defines panning group for speaker 4 (0, 1 or 2)
item rand (0, 1, 2)

p38                                   ;defines panning group for speaker 5 (0, 1 or 2)
item rand (0, 1, 2)

p39                                   ;defines panning group for speaker 6 (0, 1 or 2)
item rand (0, 1, 2)

p40                                   ;defines panning group for speaker 7 (0, 1 or 2)
item rand (0, 1, 2)

p41                                   ;defines panning group for speaker 8 (0, 1 or 2)
item rand (0, 1, 2)
```

Figure 28 – Step Two Subscore Example

2.15 Generative Process and Step Two Score(s)

Similar to Step One, this process is one of executing Cmask using the subscore(s) to produce Csound scores. At this point in the compositional process this does not occur in the straightforward linear manner in which it is being described here. Instead it occurs in conjunction with rendering the resultant sound files using the Csound rendering engine. As the sound files are created and listened to more sound files are being rendered. Using a powerful workstation computer Csound can render the sound files, generated by this complex orchestra, in close to real time. However, in keeping with the flow of this document, in Fig. 29 and Fig. 30 are screenshots of the command line Cmask execution and screenshot of an excerpt of a resultant Csound score imported into a spreadsheet. Incidentally, spreadsheets provide a powerful way to edit Csound scores if it is desirable to do so. They provide not only the usual editing of rows of data but also the columns. The latter allows for formulae to be applied to specific pfields without affecting the others. This provides quite an effective manner in which to create further variations of a score that is working well. Transposition is an example of one such variation.

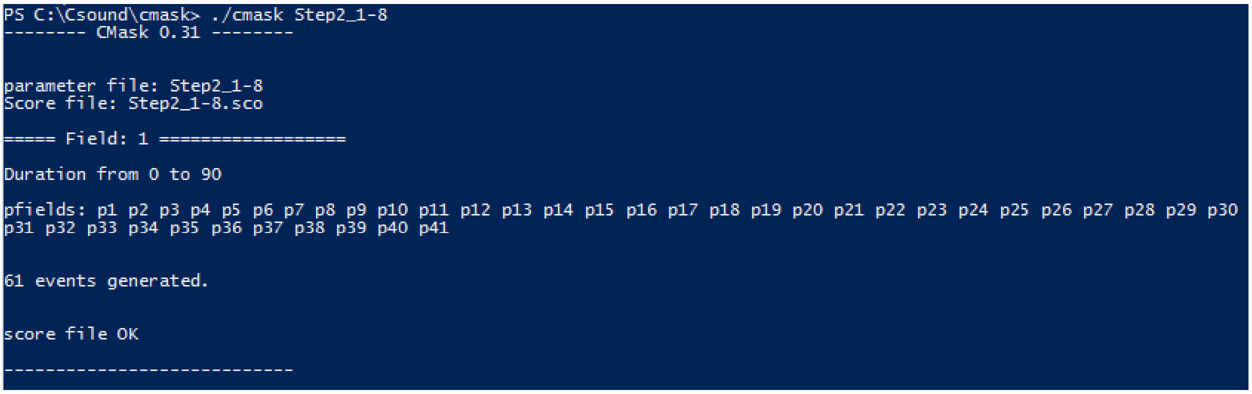

```
PS C:\Csound\cmask> ./cmask Step2_1-8
-------- CMask 0.31 --------


parameter file: Step2_1-8
Score file: Step2_1-8.sco

===== Field: 1 =================

Duration from 0 to 90

pfields: p1 p2 p3 p4 p5 p6 p7 p8 p9 p10 p11 p12 p13 p14 p15 p16 p17 p18 p19 p20 p21 p22 p23 p24 p25 p26 p27 p28 p29 p30
p31 p32 p33 p34 p35 p36 p37 p38 p39 p40 p41


61 events generated.


score file OK

----------------------------
```

Figure 29– Step Two Cmask Execution

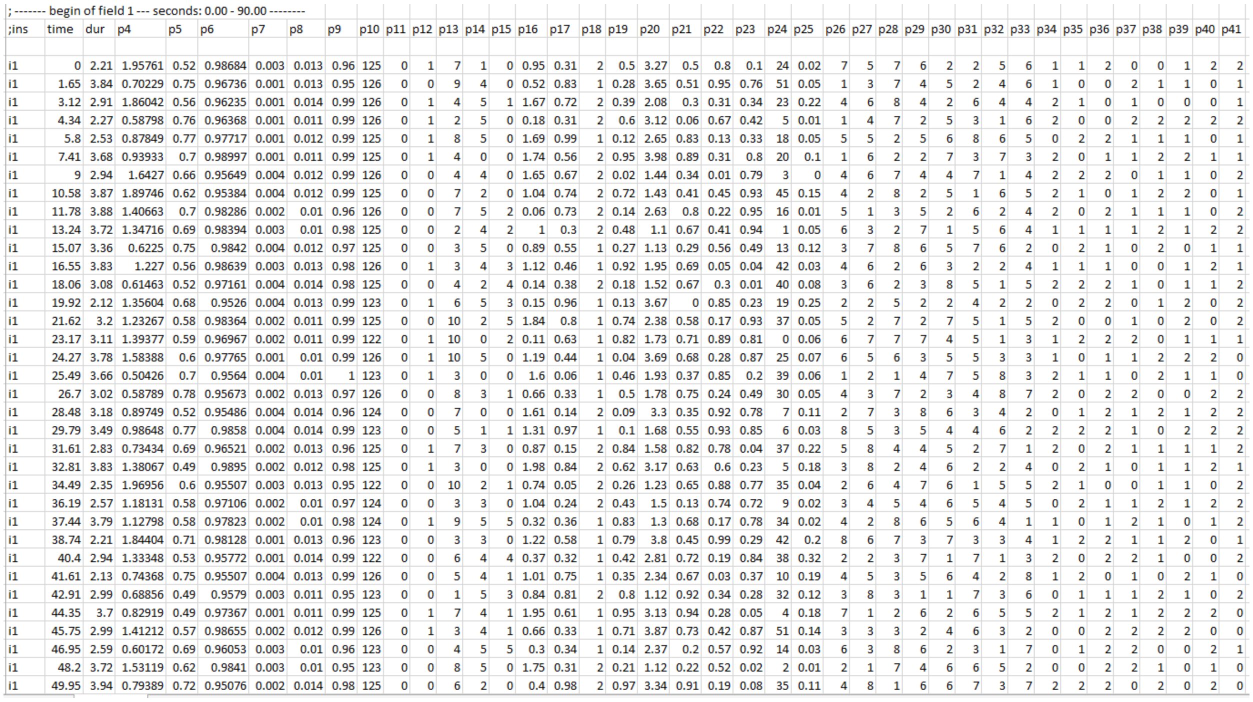

; ------- begin of field 1 --- seconds: 0.00 - 90.00 --------

;ins	time	dur	p4	p5	p6	p7	p8	p9	p10	p11	p12	p13	p14	p15	p16	p17	p18	p19	p20	p21	p22	p23	p24	p25	p26	p27	p28	p29	p30	p31	p32	p33	p34	p35	p36	p37	p38	p39	p40	p41
i1	0	2.21	1.95761	0.52	0.98684	0.003	0.013	0.96	125	0	1	7	1	0	0.95	0.31	2	0.5	3.27	0.5	0.8	0.1	24	0.02	7	5	7	6	2	2	5	6	1	1	2	0	0	1	2	2
i1	1.65	3.84	0.70229	0.75	0.96736	0.001	0.013	0.95	126	0	0	9	4	0	0.52	0.83	1	0.28	3.65	0.51	0.95	0.76	51	0.05	1	3	7	4	5	2	4	6	1	0	0	2	1	1	0	1
i1	3.12	2.91	1.86042	0.56	0.96235	0.001	0.014	0.99	126	0	1	4	5	1	1.67	0.72	2	0.39	2.08	0.3	0.31	0.34	23	0.22	4	6	8	4	2	6	4	4	2	1	0	1	0	0	0	1
i1	4.34	2.27	0.58798	0.76	0.96368	0.001	0.011	0.99	126	0	1	2	5	0	0.18	0.31	2	0.6	3.12	0.06	0.67	0.42	5	0.01	1	4	7	2	5	3	1	6	2	0	0	2	2	2	2	2
i1	5.8	2.53	0.87849	0.77	0.97717	0.001	0.012	0.99	125	0	1	8	5	0	1.69	0.99	1	0.12	2.65	0.83	0.13	0.33	18	0.05	5	5	2	5	6	8	6	5	0	2	2	1	1	1	0	1
i1	7.41	3.68	0.93933	0.7	0.98997	0.001	0.011	0.99	125	0	1	4	0	0	1.74	0.56	2	0.95	3.98	0.89	0.31	0.8	20	0.1	1	6	2	2	7	3	7	3	2	0	1	1	2	2	1	1
i1	9	2.94	1.6427	0.66	0.95649	0.004	0.012	0.99	126	0	0	4	4	0	1.65	0.67	2	0.02	1.44	0.34	0.01	0.79	3	0	4	6	7	4	4	7	1	4	2	2	2	0	1	1	0	2
i1	10.58	3.87	1.89746	0.62	0.95384	0.004	0.012	0.99	125	0	0	7	2	0	1.04	0.74	2	0.72	1.43	0.41	0.45	0.93	45	0.15	4	2	8	2	5	1	6	5	2	1	0	1	2	2	0	1
i1	11.78	3.88	1.40663	0.7	0.98286	0.002	0.01	0.96	126	0	0	7	5	2	0.06	0.73	2	0.14	2.63	0.8	0.22	0.95	16	0.01	5	1	3	5	2	6	2	4	2	0	2	1	1	1	0	2
i1	13.24	3.72	1.34716	0.69	0.98394	0.003	0.01	0.98	125	0	0	2	4	2	1	0.3	2	0.48	1.1	0.67	0.41	0.94	1	0.05	6	3	2	7	1	5	6	4	1	1	1	1	2	1	2	2
i1	15.07	3.36	0.6225	0.75	0.9842	0.004	0.012	0.97	125	0	0	3	5	0	0.89	0.55	1	0.27	1.13	0.29	0.56	0.49	13	0.12	3	7	8	6	5	7	6	2	0	2	1	0	2	0	1	1
i1	16.55	3.83	1.227	0.56	0.98639	0.003	0.013	0.98	126	0	1	3	4	3	1.12	0.46	1	0.92	1.95	0.69	0.05	0.04	42	0.03	4	6	2	6	3	2	2	4	1	1	1	0	0	1	2	1
i1	18.06	3.08	0.61463	0.52	0.97161	0.004	0.014	0.98	125	0	0	4	2	4	0.14	0.38	2	0.18	1.52	0.67	0.3	0.01	40	0.08	3	6	2	3	8	5	1	5	2	2	2	1	0	1	1	2
i1	19.92	2.12	1.35604	0.68	0.9526	0.004	0.013	0.99	123	0	1	6	5	3	0.15	0.96	1	0.13	3.67	0	0.85	0.23	19	0.25	2	2	5	2	2	4	2	2	0	2	2	0	1	2	0	2
i1	21.62	3.2	1.23267	0.58	0.98364	0.002	0.011	0.99	125	0	0	10	2	5	1.84	0.8	1	0.74	2.38	0.58	0.17	0.93	37	0.05	5	2	7	2	7	5	1	5	2	0	0	1	0	2	0	2
i1	23.17	3.11	1.39377	0.59	0.96967	0.002	0.011	0.99	122	0	1	10	0	2	0.11	0.63	1	0.82	1.73	0.71	0.89	0.81	0	0.06	6	7	7	7	4	5	1	3	1	2	2	2	0	1	1	1
i1	24.27	3.78	1.58388	0.6	0.97765	0.001	0.01	0.99	126	0	1	10	5	0	1.19	0.44	1	0.04	3.69	0.68	0.28	0.87	25	0.07	6	5	6	3	5	5	3	3	1	0	1	1	2	2	2	0
i1	25.49	3.66	0.50426	0.7	0.9564	0.004	0.01	1	123	0	1	3	0	0	1.6	0.06	1	0.46	1.93	0.37	0.85	0.2	39	0.06	1	2	1	4	7	5	8	3	2	1	1	0	2	1	1	0
i1	26.7	3.02	0.58789	0.78	0.95673	0.002	0.013	0.97	126	0	0	8	3	1	0.66	0.33	1	0.5	1.78	0.75	0.24	0.49	30	0.05	4	3	7	2	3	4	8	7	2	0	2	2	0	0	2	2
i1	28.48	3.18	0.89749	0.52	0.95486	0.004	0.014	0.96	124	0	0	7	0	0	1.61	0.14	2	0.09	3.3	0.35	0.92	0.78	7	0.11	2	7	3	8	6	3	4	2	0	1	2	1	2	1	2	2
i1	29.79	3.49	0.98648	0.77	0.9858	0.004	0.014	0.99	123	0	0	5	1	1	1.31	0.97	1	0.1	1.68	0.55	0.93	0.85	6	0.03	8	5	3	5	4	4	6	2	2	2	2	1	0	2	2	2
i1	31.61	2.83	0.73434	0.69	0.96521	0.002	0.013	0.96	125	0	1	7	3	0	0.87	0.15	2	0.84	1.58	0.82	0.78	0.04	37	0.22	5	8	4	4	5	2	7	1	2	0	2	1	1	1	1	2
i1	32.81	3.83	1.38067	0.49	0.9895	0.002	0.012	0.98	125	0	1	3	0	0	1.98	0.84	2	0.62	3.17	0.63	0.6	0.23	5	0.18	3	8	2	4	6	2	2	4	0	2	1	0	1	1	2	1
i1	34.49	2.35	1.96956	0.6	0.95507	0.003	0.013	0.95	122	0	0	10	2	1	0.74	0.05	2	0.26	1.23	0.65	0.88	0.77	35	0.04	2	6	4	7	6	1	5	5	2	1	0	0	1	1	0	2
i1	36.19	2.57	1.18131	0.58	0.97106	0.002	0.01	0.97	124	0	0	3	3	0	1.04	0.24	2	0.43	1.5	0.13	0.74	0.72	9	0.02	3	4	5	4	6	5	4	5	0	2	1	1	2	1	2	2
i1	37.44	3.79	1.12798	0.58	0.97823	0.002	0.01	0.98	124	0	1	9	5	5	0.32	0.36	1	0.83	1.3	0.68	0.17	0.78	34	0.02	4	2	8	6	5	6	4	1	1	0	1	2	1	0	1	2
i1	38.74	2.21	1.84404	0.71	0.98128	0.001	0.013	0.96	123	0	0	3	3	0	1.22	0.58	1	0.79	3.8	0.45	0.99	0.29	42	0.2	8	6	7	3	7	3	3	4	1	2	2	1	1	2	0	1
i1	40.4	2.94	1.33348	0.53	0.95772	0.001	0.014	0.99	122	0	0	6	4	4	0.37	0.32	1	0.42	2.81	0.72	0.19	0.84	38	0.32	2	2	3	7	1	7	1	3	2	0	2	2	1	0	0	2
i1	41.61	2.13	0.74368	0.75	0.95507	0.004	0.013	0.99	126	0	0	5	4	1	1.01	0.75	1	0.35	2.34	0.67	0.03	0.37	10	0.19	4	5	3	5	6	4	2	8	1	2	0	1	0	2	1	0
i1	42.91	2.99	0.68856	0.49	0.9579	0.003	0.011	0.95	123	0	0	1	5	3	0.84	0.81	2	0.8	1.12	0.92	0.34	0.28	32	0.12	3	8	3	1	1	7	3	6	0	1	1	1	2	1	0	2
i1	44.35	3.7	0.82919	0.49	0.97367	0.001	0.011	0.99	125	0	1	7	4	1	1.95	0.61	1	0.95	3.13	0.94	0.28	0.05	4	0.18	7	1	2	6	2	6	5	5	2	1	2	1	2	2	2	0
i1	45.75	2.99	1.41212	0.57	0.98655	0.002	0.012	0.99	126	0	1	3	4	1	0.66	0.33	1	0.71	3.87	0.73	0.42	0.87	51	0.14	3	3	3	2	4	6	3	2	0	0	2	2	2	2	0	0
i1	46.95	2.59	0.60172	0.69	0.96053	0.003	0.01	0.96	123	0	0	4	5	5	0.3	0.34	1	0.14	2.37	0.2	0.57	0.92	14	0.03	6	3	8	6	2	3	1	7	0	1	2	2	0	0	2	1
i1	48.2	3.72	1.53119	0.62	0.9841	0.003	0.01	0.95	123	0	0	8	5	0	1.75	0.31	2	0.21	1.12	0.22	0.52	0.02	2	0.01	2	1	7	4	6	6	5	2	0	0	2	2	1	0	1	0
i1	49.95	3.94	0.79389	0.72	0.95076	0.002	0.014	0.98	125	0	0	6	2	0	0.4	0.98	2	0.97	3.34	0.91	0.19	0.08	35	0.11	4	8	1	6	6	7	3	7	2	2	2	0	2	0	2	0

Figure 30 – Excerpt from Step Two Generated Score

Figure 31 and Fig. 32 below are screenshots from my workstation while producing the Step Two samples. Missing from these images, the Sonar DAW is minimized. It is, however, ready to be opened between each render so the previously rendered 8-channel sounds files may be imported and listened to. If inspiration motivates during this action, the sound files may be edited before moving on to the creation of the next set of samples. Otherwise, the sample is imported and then the application is again minimized until the next sample is completed and the cycle repeats.

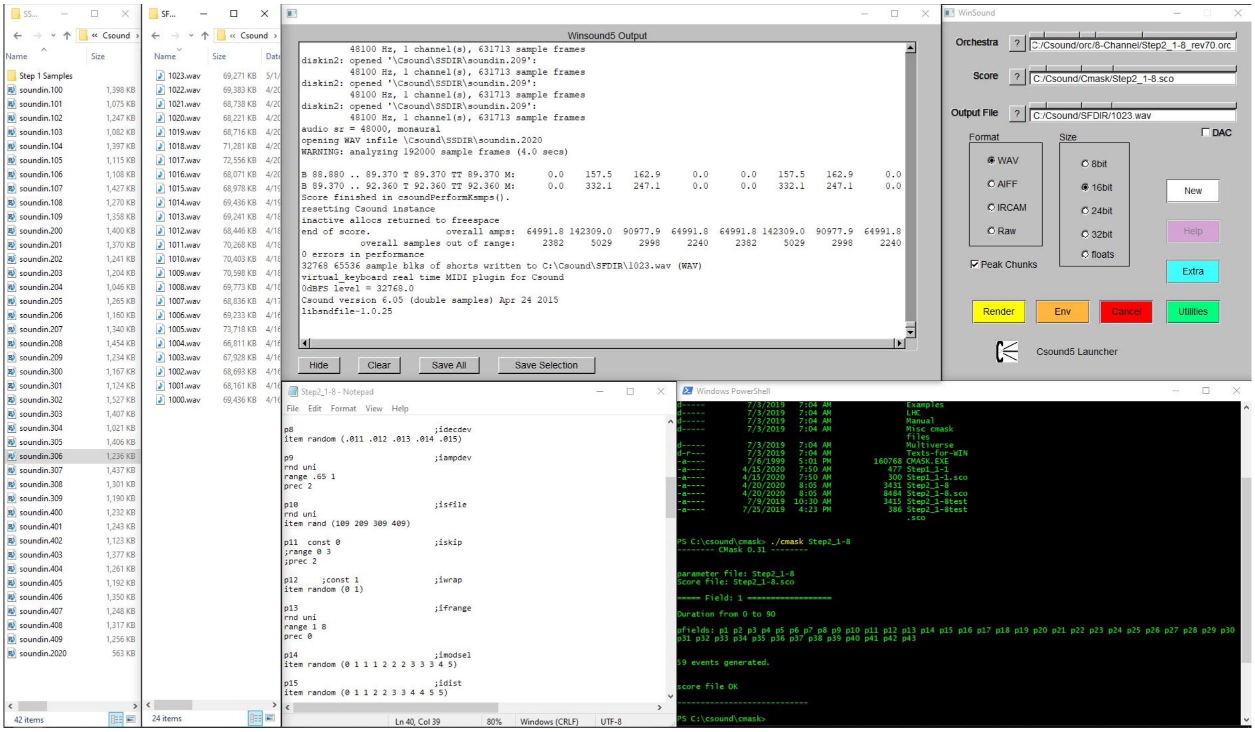

Name	Size
Step 1 Samples	
soundin.100	1,398 KB
soundin.101	1,075 KB
soundin.102	1,247 KB
soundin.103	1,082 KB
soundin.104	1,397 KB
soundin.105	1,115 KB
soundin.106	1,108 KB
soundin.107	1,427 KB
soundin.108	1,270 KB
soundin.109	1,358 KB
soundin.200	1,400 KB
soundin.201	1,370 KB
soundin.202	1,241 KB
soundin.203	1,204 KB
soundin.204	1,046 KB
soundin.205	1,265 KB
soundin.206	1,160 KB
soundin.207	1,340 KB
soundin.208	1,454 KB
soundin.209	1,234 KB
soundin.300	1,167 KB
soundin.301	1,124 KB
soundin.302	1,527 KB
soundin.303	1,407 KB
soundin.304	1,021 KB
soundin.305	1,406 KB
soundin.306	1,236 KB
soundin.307	1,437 KB
soundin.308	1,301 KB
soundin.309	1,190 KB
soundin.400	1,232 KB
soundin.401	1,243 KB
soundin.402	1,123 KB
soundin.403	1,377 KB
soundin.404	1,261 KB
soundin.405	1,192 KB
soundin.406	1,350 KB
soundin.407	1,248 KB
soundin.408	1,317 KB
soundin.409	1,256 KB
soundin.2020	563 KB

42 items

Name	Size	Date
1023.wav	69,271 KB	5/1/
1022.wav	69,383 KB	4/20
1021.wav	68,738 KB	4/20
1020.wav	68,221 KB	4/20
1019.wav	68,716 KB	4/20
1018.wav	71,281 KB	4/20
1017.wav	72,556 KB	4/20
1016.wav	68,071 KB	4/20
1015.wav	68,978 KB	4/19
1014.wav	69,436 KB	4/19
1013.wav	69,241 KB	4/18
1012.wav	68,446 KB	4/18
1011.wav	70,268 KB	4/18
1010.wav	70,403 KB	4/18
1009.wav	70,598 KB	4/18
1008.wav	69,773 KB	4/18
1007.wav	68,836 KB	4/17
1006.wav	69,233 KB	4/16
1005.wav	73,718 KB	4/16
1004.wav	66,811 KB	4/16
1003.wav	67,928 KB	4/16
1002.wav	68,693 KB	4/16
1001.wav	68,161 KB	4/16
1000.wav	69,436 KB	4/16

24 items

Winsound5 Output

```
         48100 Hz, 1 channel(s), 631713 sample frames
diskin2: opened '\Csound\SSDIR\soundin.209':
         48100 Hz, 1 channel(s), 631713 sample frames
diskin2: opened '\Csound\SSDIR\soundin.209':
         48100 Hz, 1 channel(s), 631713 sample frames
diskin2: opened '\Csound\SSDIR\soundin.209':
         48100 Hz, 1 channel(s), 631713 sample frames
audio sr = 48000, monaural
opening WAV infile \Csound\SSDIR\soundin.2020
WARNING: analyzing 192000 sample frames (4.0 secs)

B 88.880 .. 89.370 T 89.370 TT 89.370 M:      0.0    157.5    162.9      0.0      0.0    157.5    162.9      0.0
B 89.370 .. 92.360 T 92.360 TT 92.360 M:      0.0    332.1    247.1      0.0      0.0    332.1    247.1      0.0
Score finished in csoundPerformKsmps().
resetting Csound instance
inactive allocs returned to freespace
end of score.              overall amps:  64991.8 142309.0  90977.9  64991.8  64991.8 142309.0  90977.9  64991.8
           overall samples out of range:     2382     5029     2998     2240     2382     5029     2998     2240
0 errors in performance
32768 65536 sample blks of shorts written to C:\Csound\SFDIR\1023.wav (WAV)
virtual_keyboard real time MIDI plugin for Csound
0dBFS level = 32768.0
Csound version 6.05 (double samples) Apr 24 2015
libsndfile-1.0.25
```

Orchestra ? C:/Csound/orc/8-Channel/Step2_1-8_rev70.orc
Score ? C:/Csound/Cmask/Step2_1-8.sco
Output File ? C:/Csound/SFDIR/1023.wav
☐ DAC

Format: ● WAV ○ AIFF ○ IRCAM ○ Raw
Size: ○ 8bit ● 16bit ○ 24bit ○ 32bit ○ floats
☑ Peak Chunks

```
p8                              ;idecdev
item random (.011 .012 .013 .014 .015)

p9                              ;iampdev
rnd uni
range .65 1
prec 2

p10                             ;isfile
rnd uni
item rand (109 209 309 409)

p11  const 0                    ;iskip
;range 0 3
;prec 2

p12     ;const 1                ;iwrap
item random (0 1)

p13                             ;ifrange
rnd uni
range 1 8
prec 0

p14                             ;imodsel
item random (0 1 1 1 2 2 2 3 3 3 4 5)

p15                             ;idist
item random (0 1 1 2 2 3 3 4 4 5 5)
```

Ln 40, Col 39 | 80% | Windows (CRLF) | UTF-8

```
d-----        7/3/2019   7:04 AM                Examples
d-----        7/3/2019   7:04 AM                LHC
d-----        7/3/2019   7:04 AM                Manual
d-----        7/3/2019   7:04 AM                Misc cmask
                                                files
d-----        7/3/2019   7:04 AM                Multiverse
d-r---        7/3/2019   7:04 AM                Texts-for-WIN
-a----        7/6/1999   5:01 PM         160768 CMASK.EXE
-a----       4/15/2020   7:50 AM            477 Step1_1-1
-a----       4/15/2020   7:50 AM            300 Step1_1-1.sco
-a----       4/20/2020   8:05 AM           3431 Step2_1-8
-a----       4/20/2020   8:05 AM           8484 Step2_1-8.sco
-a----        7/9/2019  10:30 AM           3415 Step2_1-8test
-a----       7/25/2019   4:23 PM            386 Step2_1-8test
                                                .sco


PS C:\csound\cmask> ./cmask Step2_1-8
-------- CMask 0.31 --------


parameter file: Step2_1-8
Score file: Step2_1-8.sco

===== Field: 1 =================

Duration from 0 to 90

pfields: p1 p2 p3 p4 p5 p6 p7 p8 p9 p10 p11 p12 p13 p14 p15 p16 p17 p18 p19 p20 p21 p22 p23 p24 p25 p26 p27 p28 p29 p30
p31 p32 p33 p34 p35 p36 p37 p38 p39 p40 p41 p42 p43


59 events generated.


score file OK

----------------------------

PS C:\csound\cmask>
```

Figure 31 – Step Two - Left Workstation Computer Monitor

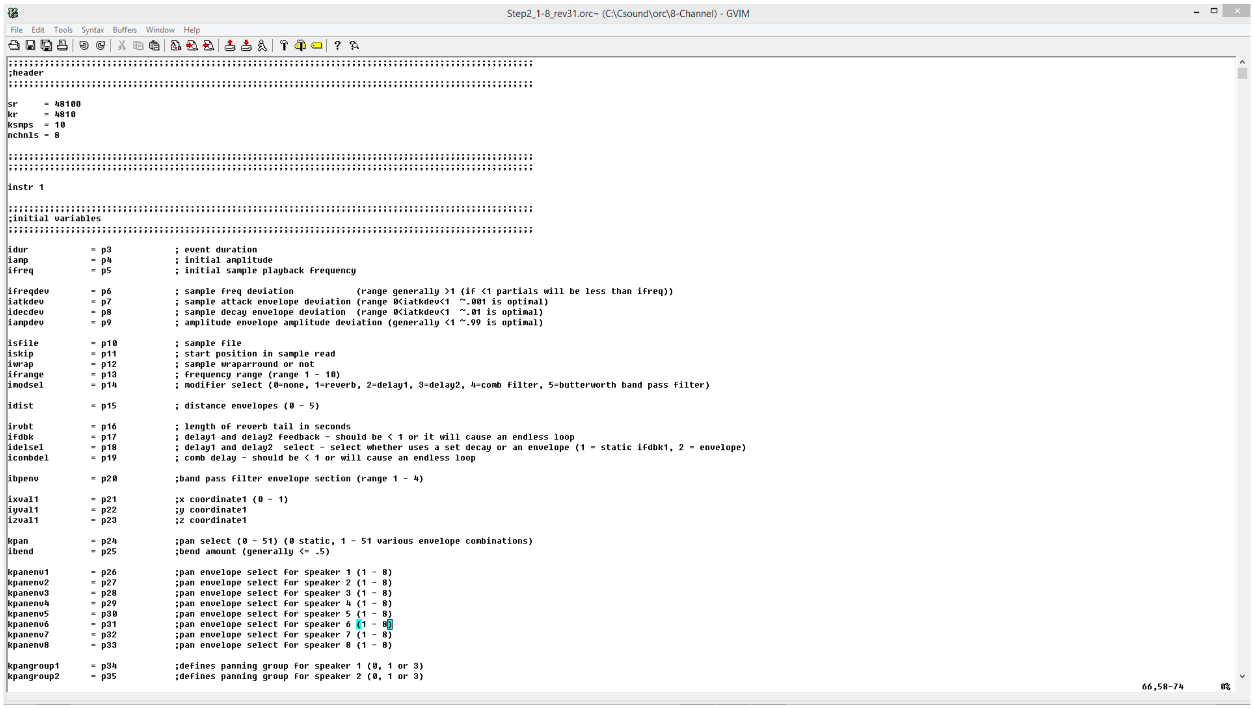

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;header
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

sr     = 48100
kr     = 4810
ksmps  = 10
nchnls = 8

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

instr 1

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;initial variables
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

idur            = p3            ; event duration
iamp            = p4            ; initial amplitude
ifreq           = p5            ; initial sample playback frequency

ifreqdev        = p6            ; sample freq deviation            (range generally >1 (if <1 partials will be less than ifreq))
iatkdev         = p7            ; sample attack envelope deviation (range 0<iatkdev<1  ~.001 is optimal)
idecdev         = p8            ; sample decay envelope deviation  (range 0<iatkdev<1  ~.01 is optimal)
iampdev         = p9            ; amplitude envelope amplitude deviation (generally <1 ~.99 is optimal)

isfile          = p10           ; sample file
iskip           = p11           ; start position in sample read
iwrap           = p12           ; sample wraparround or not
ifrange         = p13           ; frequency range (range 1 - 10)
imodsel         = p14           ; modifier select (0=none, 1=reverb, 2=delay1, 3=delay2, 4=comb filter, 5=butterworth band pass filter)

idist           = p15           ; distance envelopes (0 - 5)

irvbt           = p16           ; length of reverb tail in seconds
ifdbk           = p17           ; delay1 and delay2 feedback - should be < 1 or it will cause an endless loop
idelsel         = p18           ; delay1 and delay2  select - select whether uses a set decay or an envelope (1 = static ifdbk1, 2 = envelope)
icombdel        = p19           ; comb delay - should be < 1 or will cause an endless loop

ibpenv          = p20           ;band pass filter envelope section (range 1 - 4)

ixval1          = p21           ;x coordinate1 (0 - 1)
iyval1          = p22           ;y coordinate1
izval1          = p23           ;z coordinate1

kpan            = p24           ;pan select (0 - 51) (0 static, 1 - 51 various envelope combinations)
ibend           = p25           ;bend amount (generally <= .5)

kpanenv1        = p26           ;pan envelope select for speaker 1 (1 - 8)
kpanenv2        = p27           ;pan envelope select for speaker 2 (1 - 8)
kpanenv3        = p28           ;pan envelope select for speaker 3 (1 - 8)
kpanenv4        = p29           ;pan envelope select for speaker 4 (1 - 8)
kpanenv5        = p30           ;pan envelope select for speaker 5 (1 - 8)
kpanenv6        = p31           ;pan envelope select for speaker 6 (1 - 8)
kpanenv7        = p32           ;pan envelope select for speaker 7 (1 - 8)
kpanenv8        = p33           ;pan envelope select for speaker 8 (1 - 8)

kpangroup1      = p34           ;defines panning group for speaker 1 (0, 1 or 3)
kpangroup2      = p35           ;defines panning group for speaker 2 (0, 1 or 3)
                                                                                                    66,58-74        0%
```

Figure 32 – Step Two - Right Workstation Computer Monitor

2.16 Tertiary Waveform Editing

For me, the most enjoyable and rewarding aspect of this compositional process is when it all comes together in the editing sessions associated with this final step in the process. It is here that the coding, the generative processes, and rendering finally culminates as music. The numerous 8-channel .wav files generated by the previous reiterative processes are imported into Sonar®, which is a Digital Audio Workstation (DAW) and then edited into a composition.

An aspect of the subscore manipulation of the sound file output now becomes important to mention. While generating the tertiary sound files each subscore was utilized to generate a score four times. The only change to the subscore was the tendency mask for the frequency range of the sample playback. As stated earlier, each time a sub-score is instantiated to generate a score, the results are unique. Similar to the manner in which soprano, alto, tenor, and bass (SATB) choral scores are generally separated by frequency ranges, the four renders are aligned for use as a group when editing. In other words, four sets of eight sound files are stacked in the same time range in the DAW and each set of eight is then edited separately. It was found that this provides a richness to the spatialization and to the polyphonic timbre. Further, it adds an extensive contrapuntal dimension to the mix. Below in, Fig. 33 and Fig. 34, are examples of the workstation monitors depicting this 32-channel diffusion.

Another process often applied in this step is to copy one set of the eight rendered sound files and paste it three times to fill the other 24 tracks. Then, using the transpose function in SONAR, they are transposed into three distinct frequency ranges roughly emulating the SATB paradigm. This allows for a lush unison section. Further, each set of eight sound files may also be offset by a few milliseconds from the others. This slight temporal separation creates a “fullness” to the sound of the final mix. The temporal offset is also used in other areas of the mix to the extent that occasionally each of the 32 channels is offset from the others in this manner in the final mix.

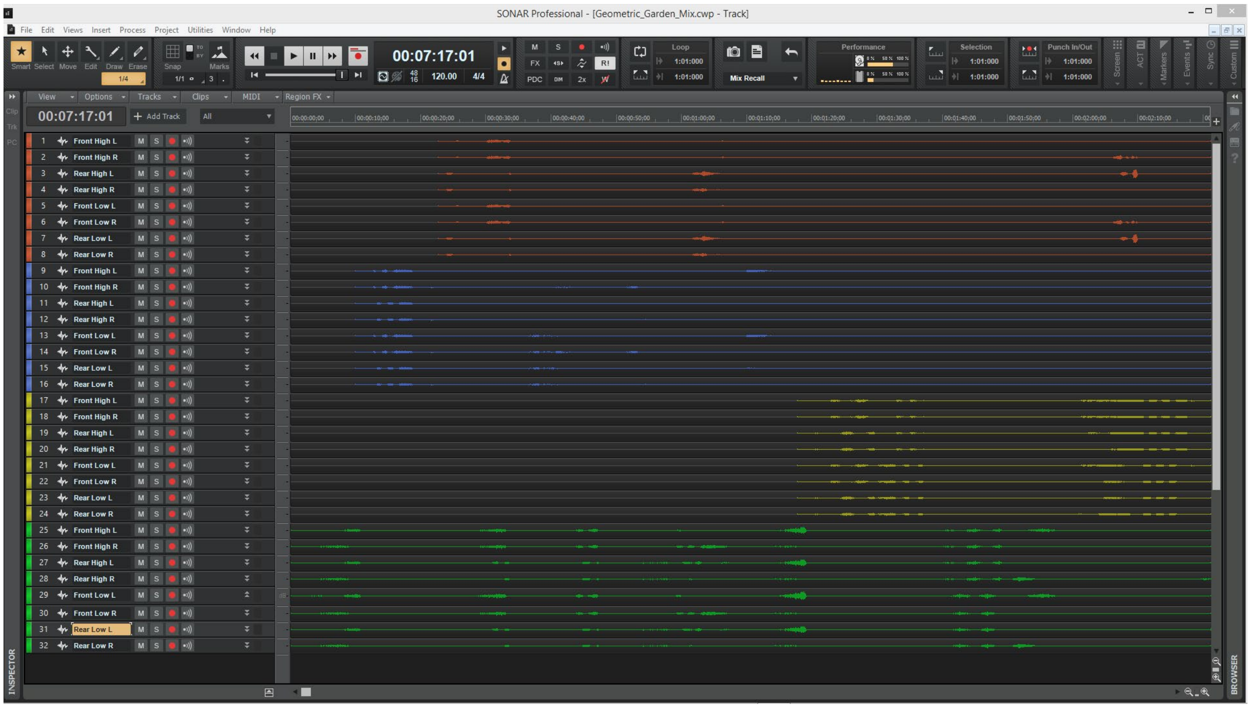

Figure 33 – Tertiary - Left Workstation Computer Monitor

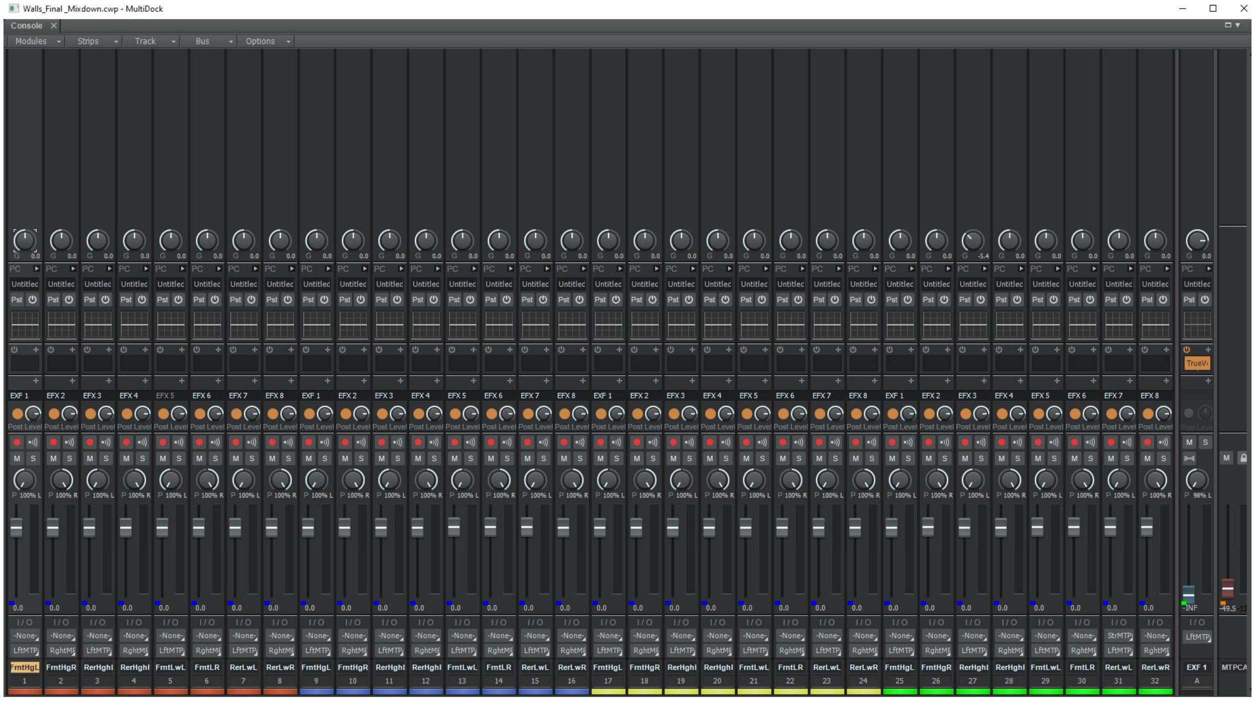

Figure 34 – Tertiary - Right Workstation Computer Monitor

The raw renders are imported into the DAW and then ruthlessly edited. Anything that sounds familiar, cliché, or just bad is immediately removed and discarded. Any sound segment that does not "work" with the composition as a whole is discarded. These decisions are determined intuitively. Then what is left is cut, moved, stretched, duplicated, transposed, retrograded… edited in any manner deemed necessary by the creative mind at work. Ultimately the segments, sometimes thousands of them, are all fit together until the finished composition emerges. It is extremely difficult to describe this process as it is intuitively driven and countless decisions are made in rapid succession. It is as if the final composition is being sculpted from the computer-generated sound files. All the unwanted material is chipped away until the "Michael", to cite Michelangelo, inside them, the essence of the composition, is revealed. Fortunately, thanks to digital editing, chips can also be added during the process as well. This step is where the artist meets the computer and interjects the humanistic aesthetic.

"Phrasing is the difference between related sound combinations that simply exist in adjacent spaces and those that connect and interact in a meaningful way" (Rhoades, 2009). During the editing phase of this compositional paradigm, it is important to maintain a focus upon the meaningful interconnection of the numerous audio clips. It is a common pitfall to connect audio sequences in a very choppy manner. It sounds as if the composer is saying, "here is one sequence and now here is another, and here is still another". Instead, I am interested in presenting a composition consisting a flow of progressions that seem to form a single unified event. These sequences should be fluid in the transition from one to the next so that it is difficult to discern where one begins and another ends. This a major challenge. Through an understanding and development of phrasing a collection of sound events becomes music.

If the visual aspect of the composition is available, when the rough draft of the musical composition is ready, a low-resolution version of the visual composition is imported into the DAW and bears influence over the final form of the final mix. The opposite also occurs when a stereo mix of the rough draft of the completed musical composition is imported into Adobe Premier® and exerts influence over the current draft of the visual aspect of the composition. As will be elaborated upon later in this document, in this manner there is a feedback loop created between the music, the visual, and the composer. Indirectly the high-performance computing cluster is also involved in this interactive and interdisciplinary process. This multifaceted process culminates in a final draft of the musical composition and eventually to the completed visual music composition.

2.17 Final Mix

When a final version of the musical composition is completed, each of the 32 tracks is exported as mono, 32 bit, 48k sample rate, .wav files and then imported into Sound Forge®. There each is combed for any glitches, or overs, and anything else that might be out of place. Also, the amplitude of each file is raised to an optimum value… each relative to the others. An optimum amplitude value is as high as possible to keep the main aspects of the composition near the -6db range. Often there is a small amount of compression employed to accomplish this. The compression algorithm, for which I created a preset in Sound Forge®, reduces the peaks by 2db… nothing more. This is a soft compression that has little effect on the sound files and yet brings them to within an acceptable perceived amplitude range. Certainly, the more compression employed proportionately reduces the dynamic range of a composition. Though this is often desirable in other musical genres, here a wide dynamic range is considered an important aspect of the expressiveness of a composition and so should be preserved. Another consideration is the perceived amplitude of the overall piece should match or fit well with any other compositions within the current project. This may be a throwback to the days when it was desirable to mix for CDs where all tracks might be listened to in a sequence. Nonetheless, if any two or more of the compositions in a project are presented together, this is a desirable practice.

2.18 Finished Composition

At this point in the process the composition is considered complete. As can be seen, this process, though complex, yields related yet varied material that is the essence of a viable compositional process. That combined with the spatial aspects afford the opportunity for a rich and unique flora of sonic material.

2.19 Holophony

Now that a basic comprehension of the compositional techniques utilized to create the musical side of the visual music compositions associated with this research has been established, we can proceed to the specifics of one of the pivotal attributes involved… holophony. To be very clear, holophony is a result of the interference patterns caused by wave propagation. This was described in the introductory chapter of this document. From this perspective, the term refers to a mechanical process. In this study, the focus is upon the perception of a sound event existing within the space adjacent the listener as a result of auditory stereopsis. As discussed later in this chapter, the employment of Ambisonics in this process renders the term holography appropriate here. Nonetheless, the primary factor I am concerned with is perception. Therefore, the term holography is used from a perceptual perspective and not a mechanical one.

To proceed, much work accomplished to date in the area of 3D, or periphonic, spatialization of sound, including my own, could be considered holophonic in nature. Backing up another step, nearly all sound could be considered holophonic since the source sound and/or its reflections, in the form of reverberations and delays, are perceived within the listening space. Is holophony only a technically based term that belies the theoretical existence of a sound within a space as with, for example using VBAP? (See Fig. 24 and Fig. 25 above.) Or can holophony exist as a perceptually based phenomenon?

This study is ultimately concerned with the phenomenological aspects of the projection of sound events, including their associated indirect spatial qualities such as reverb and delay, as clearly existing at times within the listening space. A methodological perspective as to how to approach such perceptions are described in detail here. Since the nature of sound is both temporally-based and spatially-based, these sound events need not be present within the listening environment for the entire duration of a composition, however they must be perceived to be so at times. Further, these occurrences should happen in relation to or in correlation with visual events.

It is a relatively straightforward exercise to conceptualize holophony. That is not the goal here. Instead we are looking for the perceptual phenomenon. Though commonly the 3D spatialization of sound is theoretically holophonic, it is the author’s general experience that, perceptually, the sound events tend to reside primarily on the perimeter of the loudspeaker arrangement. (An exception to this may reside in the area of Wave Field Synthesis (WFS).) By this is intended that the primary amplitude foci of the sound events, the source sounds, most often tend to reside on the perimeter of the loudspeaker arrangement, directly from or between loud speakers. This is especially noticeable in a large venue where there are more cubic feet of air to move in the center as opposed to between the loudspeakers. Proceeding from this premise, how can we bring the perception of the location(s) of sound events within the listening space? This is a key question with regard to this research and within a movement toward creating a deeply immersive listening experience.

The techniques employed here to approach a more perceptually vivid holophonic experience involve interweaving several existing spatialization techniques and blending them toward this goal. The combination of the author's version of Vector Based Amplitude Panning (aVBAP), High Order Ambisonics (HOA), Convolution Reverb (CR), and High-Density Loudspeaker Arrays (HDLA) are the primary methodologies involved in this endeavor. The aVBAP and CR aspects are incorporated at the sound-generation level with the Score Based Sampling (SBS) process and Csound. HOA, specifically targeting a given site-specific HDLA configuration, is incorporated into the process during diffusion of the composition using MAX. The amalgam of these approaches culminates in the presentation of the composition. These techniques are outlined below.

As mentioned in the previous "Tertiary Sample(s)" section, the score generation process culminates as 32-channel sound files in a SATB-like configuration. This extended aVBAP approach is the first step toward highly granulized spatial content. By this is intended that by generating 32-channel sound files we increase the number of sound fields employed, which decreases the distances between them. This allows for greater flexibility with regard to the perceived location and movement of the sound events. Figure 35 is a diagram that demonstrates how the output of each of the 32 channels is assigned to specific set of Cartesian coordinates within a cuboidal space.

It can be seen that the original 8-channel system, expressed in red lines and consisting of channels 1 - 8, remains in place and three additional similar systems are rotated approximately 22.5, 45, and 67.5 degrees off axis, channels 9 – 16, 17 – 24, and 25 -32 sequentially, keeping their coordinates within the confines of the Cartesian cuboidal space. This allows the coordinates of each of the quasi-SATB sound files to be evenly spread throughout the space and also provides a mechanism to maintain the height dimensional relationships, which are the most difficult to diffuse perceptibly (Hollerweger). This issue is further addressed and resolved with the MAX diffusion patch described later in this section.

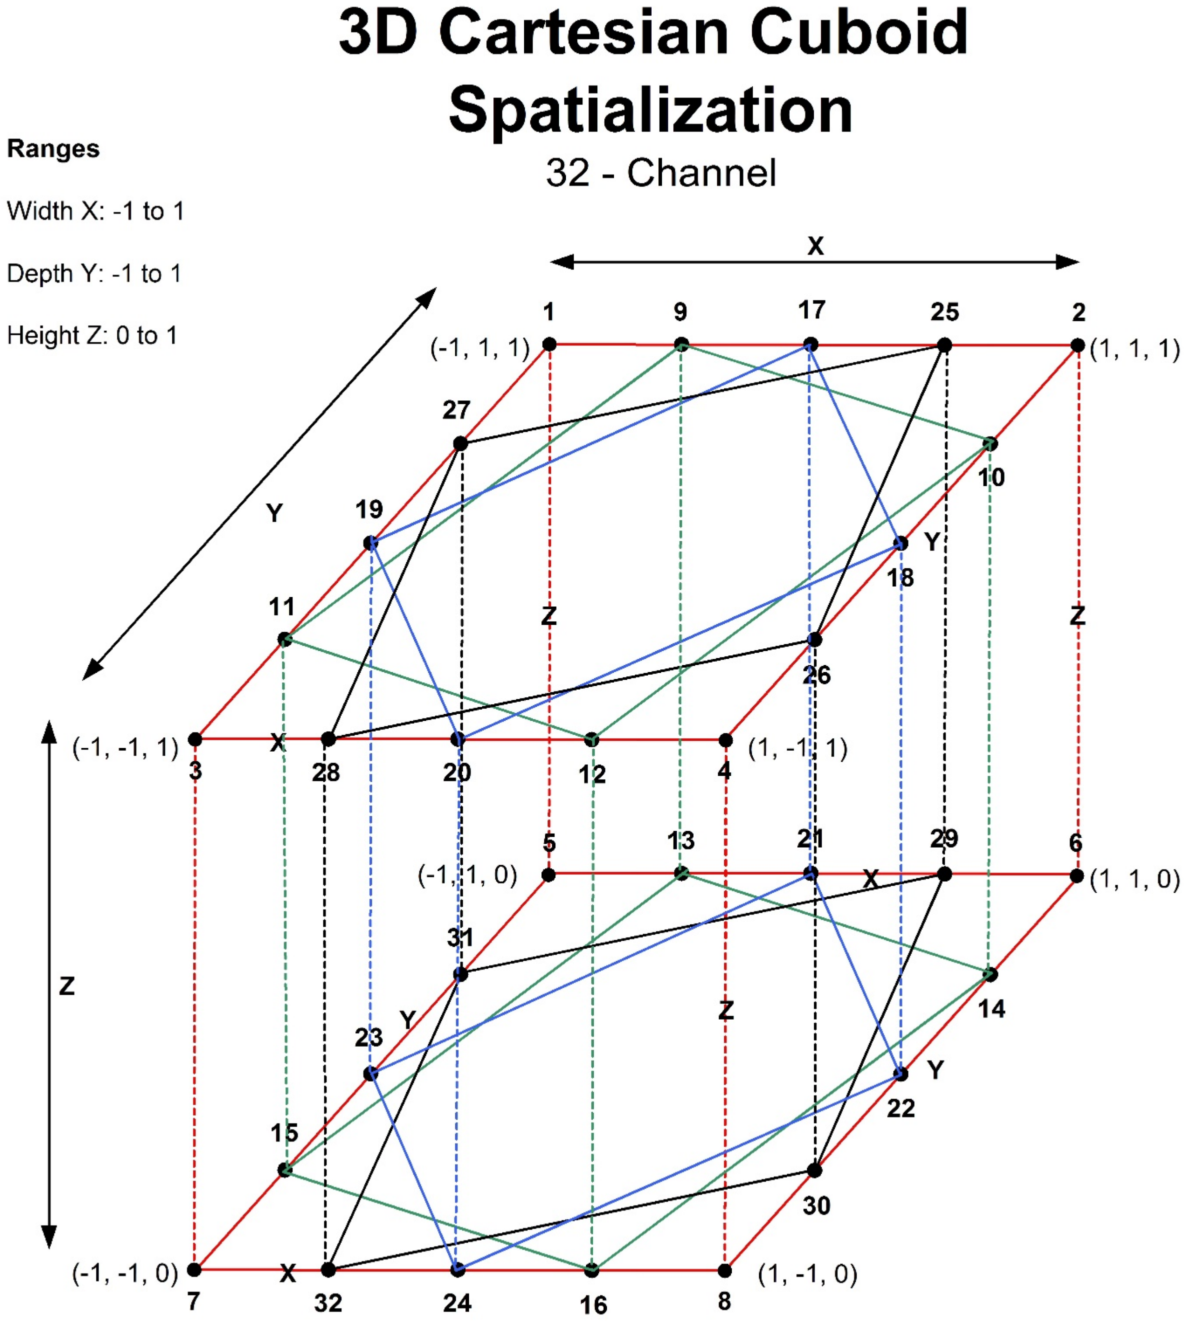

Figure 35 – Quasi-SATB 32-Channel Cartesian Coordinates

Referring to Fig. 35 we can see that if 100% of the current amplitude of a sound event is routed to channel 1, it will theoretically be perceived to exist within the top left rear side of the sound environment. This is a straightforward perception. When we modulate the amplitude between various channels, we perceive movement of the sound source between the speakers. When a sound event is modulated between channels that exist at diagonals to each other the perception of the sound, theoretically, exists within the listening space. This seems to generally be true although one must be listening for the effect, since it is often quite subtle, to recognize. Therefor aVBAP provides a basis for holography yet it is asserted here that alone it is not, experientially, 100% convincing.

To further enhance the perception of holophony, distance algorithms previously employed in my Csound orchestra code were modified for this project. The previously utilized opcode called nreverb, which is a filter-based reverberation approach, was replaced by a CR opcode called pconvolve. Nreverb is a reverberator consisting of 6 parallel comb-lowpass filters being fed into a series of 5 allpass filters (Csound Manual). The Pconvolve opcode utilizes convolution based on a uniformly partitioned overlap-save algorithm (Csound Manual).

Convolution reverb consists of mapping the reverberation signature of an impulse response (IR) recording onto another sample. This approach ascribes realistic reverberation characteristics to a sound event as opposed to more conventional forms of reverberation, where filters are employed to imitate generic qualities of an imaginary reverberant space. Personal experience has demonstrated that CR augments the perception of holophonic distance dimension in a much more convincing manner than standard filter-based reverb paradigms, leaving the latter the obvious choice for this work.

“Convolution is a mathematical procedure whereby one function is modified by another” (Heintz). Convolution reverb consists of the mapping of the reverberation signature of an impulse response sample onto another sample. This allows a recording of an impulse, a very short sound source with a fast attack portion of an amplitude waveform such as a hand clap or a gun shot, to stipulate the reverberation characteristics of a sound source not recorded in the same sound space as the impulse. This approach provides much more realistic reverberation characteristics to a sound event thus, for the purposes of this study, providing a much more realistic holophonic distance representation.

The venue in which the impulse response (IR) sample was recorded is key to the CR characteristics. After experimenting with several different IR files, it was found that the ones recorded in various cathedrals around the world most favorably augmented the distance aspect of the holophonic effect. Several examples of these cathedrals include, St. Mark's, St. Luke's, several Notre Dame, and York Minster, which were each recorded at 24 bit, 48k sampling rates. After numerous tests, the latter was arrived upon as the most convincing to utilize although the others can be quickly stipulated in the orchestra file instead if desired. One reason different IR samples might be desirable is when used in relationship to a specific visual VR space. The correlation between the two is interesting and further discussed later in this document.

The distance algorithm code implemented here consists of three sets of modulations each applied to the same sound input. The over simplified premise is that if a sound event occurs close to the listener the source of the sound is primarily perceived and with very little reverberation. Conversely, when a sound source is far away from the listener, the reverberation is prominent in the listener's perception and the sound source is minimal. Though simplistic, this approach seems to provide a strong augmentation to the desired holophonic distance characteristics. Below is the Csound code for these algorithms. It simply provides a way to modulate between wet, meaning reverb only, and dry, meaning no reverb, outputs over the duration of the event as related to the Cartesian coordinates of the event stipulated by the Csound score. It also implements overall amplitude modulation control in parallel with the wet/dry modulation. Obviously, the farther away a sound source is relative to the listener, the quieter it is. This is related to the overall amplitude of the sound in general. The kdist parameter, pfield 15, is stipulated in the score and provides the event by event choice with regard to which algorithm is implemented. Since it is germane to the holophonic aspect of this project, the entire distance algorithm portion of the Csound orchestra is in Fig. 36 below.

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;distance algorithms
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;dry envelopes
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
kenv400 linseg     1, idur,    1                                  ;static near
kenv401 linseg     0, idur * .7,   1                              ;far to near
kenv402 linseg     1, idur * .7,   0                              ;near to far
kenv403 expseg .01, idur,     1                                   ;exponential far to near
kenv404 expseg   1, idur,  .01                                    ;exponential near to far
kenv405 linseg     0, idur * .45, 1, idur * .55, 0                ;far to near to far
kenv406 linseg     1, idur * .45, 0, idur * .4, 1                 ;near to far to near
kenv407 linseg     0, idur,     0                                 ;static far

kenv45      = (kdist = 6  ?  kenv406:kenv407)
```

```
kenv44      = (kdist = 5  ?  kenv405:kenv45)
kenv43      = (kdist = 4  ?  kenv404:kenv44)
kenv42      = (kdist = 3  ?  kenv403:kenv43)
kenv41      = (kdist = 2  ?  kenv402:kenv42)
kenv40      = (kdist = 1  ?  kenv401:kenv41)
kdampenv = (kdist = 0  ?  kenv400:kenv40)

a1001 =  a1000 * kdampenv                               ;kdampenv = dry amplitude envelope
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;reverb envelopes
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
kenv300 linseg     0, idur,        0                     ;static near
kenv301 linseg     1, idur * .7,   0                     ;linear far to near
kenv302 linseg     0, idur * .7,   1                     ;linear near to far
kenv303 expseg   1, idur,    .01                         ;exponential far to near
kenv304 expseg .01, idur,       1                        ;exponential near to far
kenv305 linseg    1, idur * .45, 0, idur * .55, 1        ;far to near to far
kenv306 linseg    0, idur * .45, 1, idur * .4, 0         ;near to far to near
kenv307 linseg    1, idur,        1                      ;static far

kenv35      = (kdist = 6  ?  kenv306:kenv307)
kenv34      = (kdist = 5  ?  kenv305:kenv35)
kenv33      = (kdist = 4  ?  kenv304:kenv34)
kenv32      = (kdist = 3  ?  kenv303:kenv33)
kenv31      = (kdist = 2  ?  kenv302:kenv32)
kenv30      = (kdist = 1  ?  kenv301:kenv31)
krampenv = (kdist = 0  ?  kenv300:kenv30)

a1002 pconvolve a1000 * krampenv, 2020, ipart
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;amplitude envelopes
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
kenv500 linseg    1, idur,        1                      ;static near
kenv501 linseg    0, idur * .75,  .7, idur * .25, 1      ;linear far to near
kenv502 linseg    1, idur * .1,   .3,  idur * .9, 0      ;linear near to far
kenv503 expseg .01, idur,   1                            ;exponential far to near
kenv504 expseg   1, idur,  .01                           ;exponential near to far
kenv505 linseg    0, idur * .45, 1, idur * .55, 0        ;far to near to far
kenv506 linseg    1, idur * .45, .15, idur * .4, 1       ;near to far to near
kenv507 linseg  .25, idur, .25                           ;static far

kenv55      = (kdist = 6  ?  kenv506:kenv507)
kenv54      = (kdist = 5  ?  kenv505:kenv55)
kenv53      = (kdist = 4  ?  kenv504:kenv54)
kenv52      = (kdist = 3  ?  kenv503:kenv53)
kenv51      = (kdist = 2  ?  kenv502:kenv52)
kenv50      = (kdist = 1  ?  kenv501:kenv51)
kdistenv = (kdist = 0  ?  kenv500:kenv50)
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;distance envelope summation
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

a1000 = (a1001 + a1002) * kdistenv
```

Figure 36 – Excerpt of the Distance Algorithm Portion of the Csound Orchestra File

The HDLA is a key component in projecting holophony. Though there is no standard, generally an HDLA can be thought of as comprised of 40 or more loudspeakers (Garavaglia). Certainly, an audience located within an HDLA is guaranteed a highly immersive experience even if the musical material is stereophonic… limited to two channels. As the number of channels utilized increases so does the immersive quality of the experience. Increasing the number of loudspeakers in a venue increases the spatial granularity. Theoretically, this means that the perceived location of a sound event may be produced with increasing precision.

The number of loudspeakers in an HDLA is not the only factor in its efficacy as a holophonic instrument. Also, of high importance is the placement of the loudspeakers. One possible configuration is planar in nature whereby each speaker is placed at the same height above the floor surrounding the listeners. Referencing Fig. 35, such a loudspeaker arrangement may provide excellent diffusion on the x and y axes, the width and depth of the space, and yet contains little or no reference to the z-axis, the height. Generally, however, an HDLA implements a polar or a cuboidal arrangement of the loudspeakers thus providing perception of the z-axis in addition to the x- and y-axes. After much experimentation with the audio system in the Cube at Virginia Tech, which consists of over 144 loudspeakers (depending upon the choice of configuration), it was determined that the cuboidal arrangement best provides a perceivable holophonic experience with regard to the approach employed here.

The height dimension is usually the most challenging to diffuse in a manner that is clearly perceived by the audience (Hollerweger). In a cuboidal loudspeaker arrangement, the top corners of the space are utilized thus providing the greatest distance from the listener's ears. In a polar arrangement the corners of the space are sloped toward the center of the top of the space thus eliminating access to this extra space in the corners. As will be seen in the next section, coupled with aspects of the MAX patch designed for holophonic diffusion, this corner space becomes advantageous toward holophony.

Pictured below in Fig. 37 are two-dimensional representations of the polar and the cuboid loudspeaker configurations. The black squares represent the loudspeakers. As can be seen, the cuboid arrangement more fully utilizes the extent of the space.

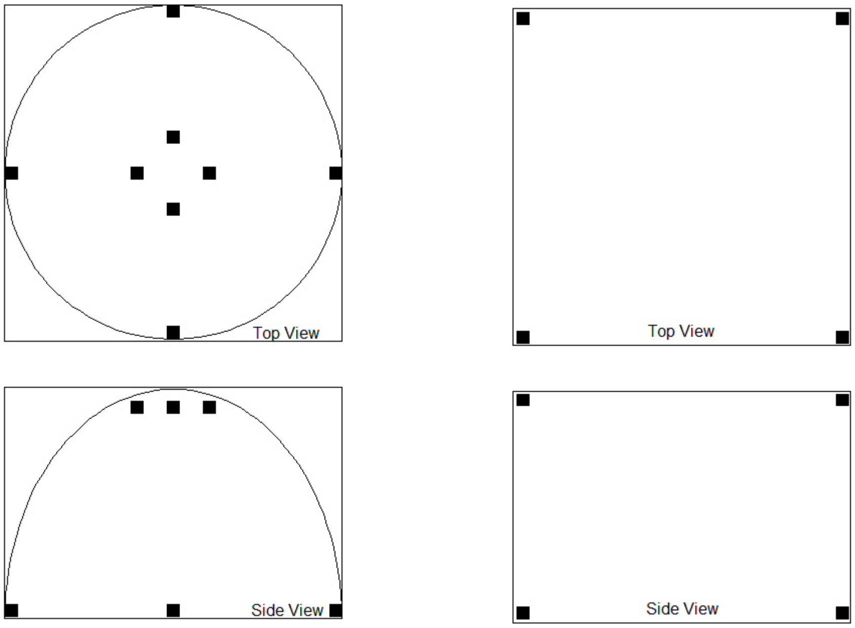

Figure 37 – Left: 8-Channel Polar Venue, Right: 8-Channel Cuboid Venue

The MAX Patch utilized for concert or installation diffusion of the compositions created for this project is the final major influence toward a perceivable holophony to be discussed here. In addition to diffusing the audio, it acts a master controller for stopping, starting, pausing, and resuming the video in sync with the audio utilizing open sound control (OSC) or a ping sensing sub-patch. The former works with the Cyclorama and the latter with HMDs. This functionality will be discussed later in this document. It should be clearly noted here that Tanner Upthegrove, who is a Media Engineer in the Cube at Virginia Tech, created the basis of this patch. It was expanded and modified by the author for the purposes of this research. Using it, various spatial parameters can be modified quickly and easily using sliders so that experimentation can occur swiftly and easily. Further, this will allow the system to exactly adjusted for any venue including the Cube.

We will first proceed with an overview of the patch and then provide a more detailed examination of the more pertinent parameters. Figure 38 is a screenshot of the main window.

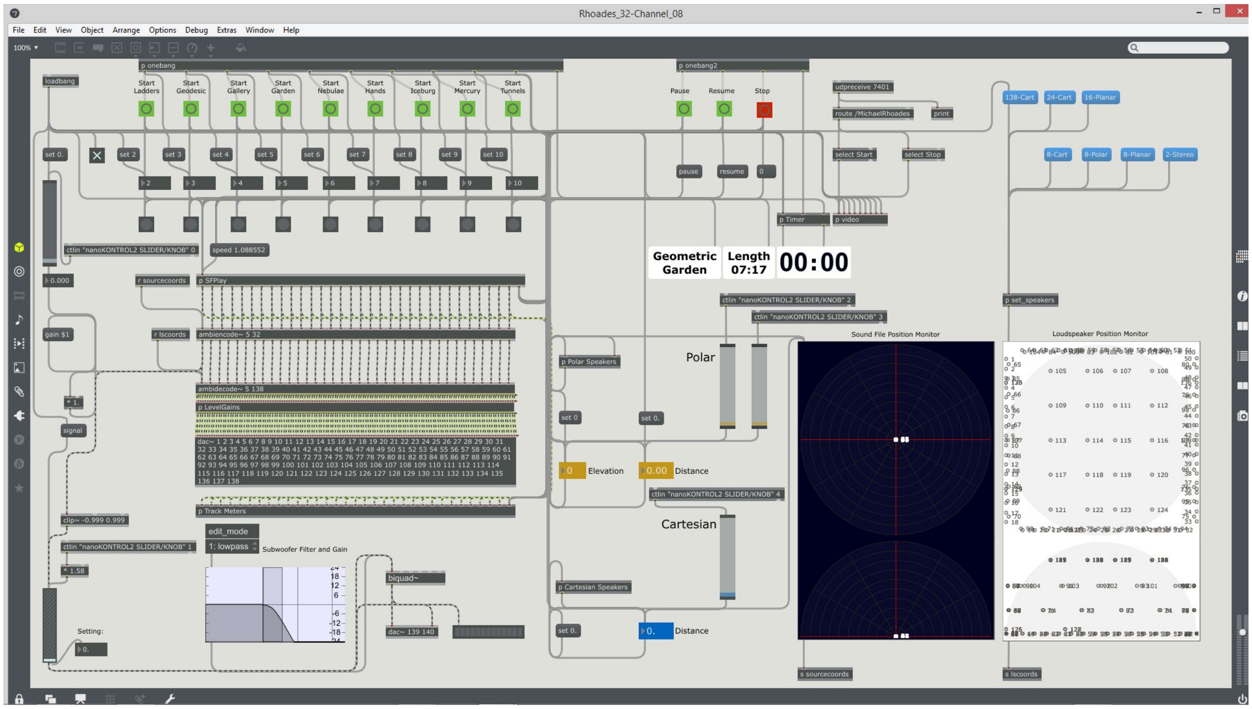

Figure 38 – Screenshot of Diffusion MAX Patch – Main Window

The sliders on the left side of the patch are amplitude controllers. The uppermost slider controls the master overall amplitude output of the patch and the lowermost controls the amplitude of the sub woofers. Though they can be operated using the workstation mouse, it is advantageous to control them using the KORG® nanaKONTROL2 (nK), which is a midi controller. It sends values between 0 and 127, for each of the various sliders, knobs, and buttons to the patch through a USB cable connected to the workstation, which can be scaled as needed for the parameter it controls. Figure 39 is a photograph of the controller and Fig. 40 provides a layout of each of the actuators and their associated midi channel numbers. For example, looking at the master volume, the "ctlin" object routes the output from the slider using midi channel 0 on the nK into the MAX slider thus controlling the amplitude. This allows a single slider to control the overall amplitude of the sound system.

Figure 39 – The KORG® nanoKONTROL2 Photograph

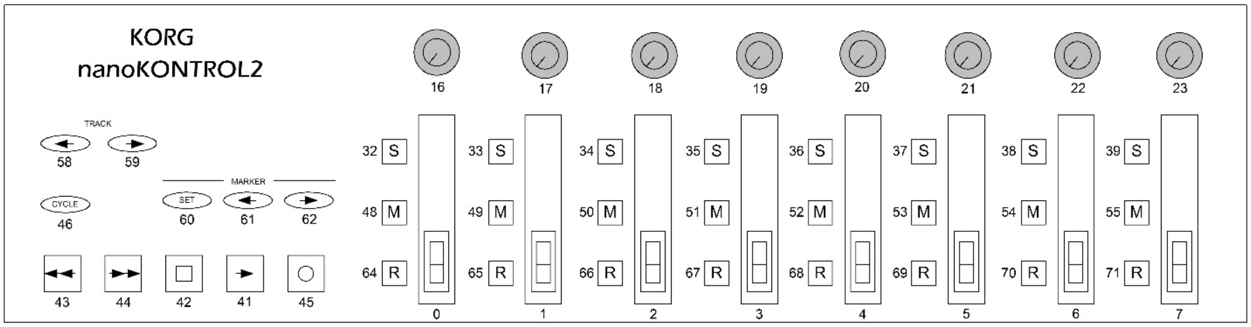

Figure 40 – The KORG® nanoKONTROL2 Midi Channel Assignments

The green buttons at the top of the MAX patch allow for the selection of which composition is instantiated. Both the audio and the video begin simultaneously and, if necessary, a delay in the sub-patches can be set for perfect synchronization. Each of the MAX buttons is assigned a button on the nK. To the right of those are the Stop, Pause, and Resume buttons. The center left area contains the ICST (B-format) Ambisonics objects. The "SFPlay object sends audio from each of 32 sound files to the Ambisonic encoder and decoder and then out to each of the 138 audio outputs on the host computer. From there the audio signals are routed to the loudspeakers.

The Ambisonic encoding and decoding objects and those to the right of them is where the final aspect of holophonic magic occurs. It is these areas upon which we will focus. High Order Ambisonics (HOA) creates one or more spherical harmonics sound fields to which audio output is sent. In general, each speaker receives a weighted sum of all Ambisonic channels, which increases the size of the listening sweet spot and the quality of sound localization (Hollerweger). The theoretical aspects of Ambisonics are a complete study and are beyond the scope of this practice-based research. Knowing how to implement them in practice toward holophony is the focus here.

After numerous listening tests it was determined that 5th order Ambisonics was the most effective choice for the final component of the projection of a holophonic listening experience. Though HOA are mathematically based, which is theoretically non-subjective, subjective listening is still the most effective manner with which to decide how they are utilized in a practical sense. A complete B-Format 5th order Ambisonics configuration utilizes a maximum of 36 channels. Here we are using 32, which works well.

One of challenges with the Cube, at least to the author's ears, has to do with the presence of the audio. Though the audio representation is pleasant, it has always seemed distant. Increasing the amplitude seems to rectify this issue to an extent, however, it also tends to reduce spatial clarity… the location of a sound event tends to become indeterminate at high amplitudes. Using the ICST MAX objects, it is possible to bring the location of the Ambisonic sound fields toward the center of the listening space. In this manner, we can increase the presence without increasing the amplitude. By doing so it was found that the spatial location of the sound events became more distinct and the holophonic quality of the sound events became more perceivable.

We can see in the images below that by adjusting the sliders in either for the polar Ambisonic fields or for the, preferred, cuboidal Cartesian fields, we can bring them closer to the center of the listening space. Pictured in Fig. 41 and Fig. 42 the cuboid Cartesian sound fields are brought closer to the center of the space on all three axes. In these figures, the top of the dark rectangle is the top view of the listening space and the bottom is a side view. The numbered small white circles represent the placement of the sound fields.

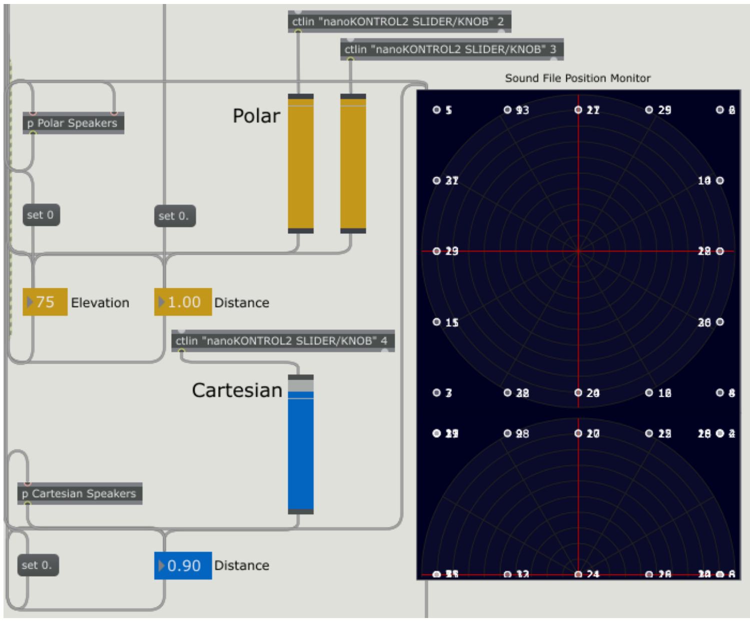

Figure 41 – The Cartesian Cuboid Sound Fields in the Perimeter of the Listening Space

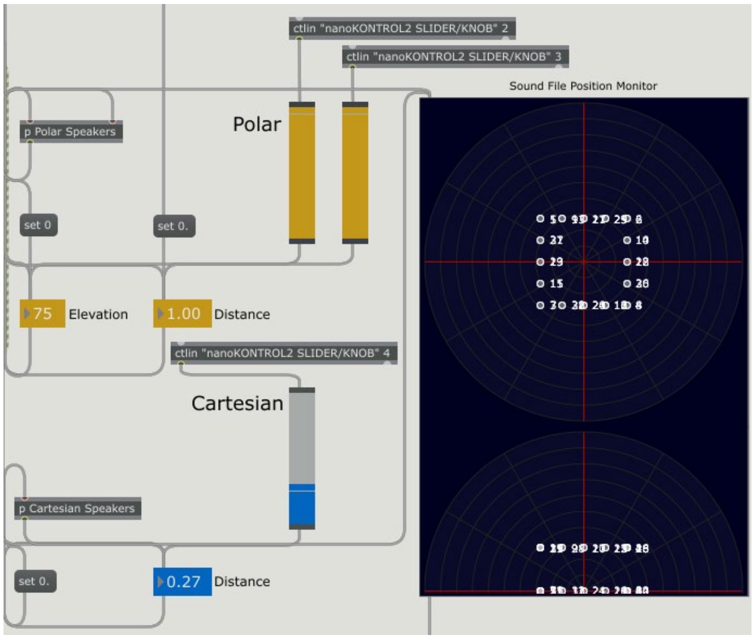

Figure 42 – The Cartesian Cuboid Sound Fields in the Center of the Listening Space

The distance value of 27, seen in Fig. 42, was found to be undesirable because the perceived spatial location of sound events became smeared. A value of around 50 seemed to be optimal in each studio space in which it was tested. Nevertheless, this particular aspect of the patch enables the presenter to optimize the listening experience based upon his or her perception of the response qualities of any given listening space.

The speaker selection objects, to the right of the “Sound File Position Monitor”, determine the actual relative position of the loudspeakers within a given space. Selecting the blue message boxes above the monitor reconfigures the loudspeaker positioning depending upon the space being utilized. In Fig. 43, on the left the configuration for the Cube at Virginia Tech is pictured and on the right is the configuration for the author’s studio using a Cartesian configuration. Also included in the patch are the speaker configurations for the 24-channel Perform Studio and the 16-channel DISIS studio, both located at Virginia tech. It is relatively straightforward to add additional listening venues as needed. In Fig. 44 is pictured the sub-patch that stipulates the relative coordinates of the speakers for each venue.

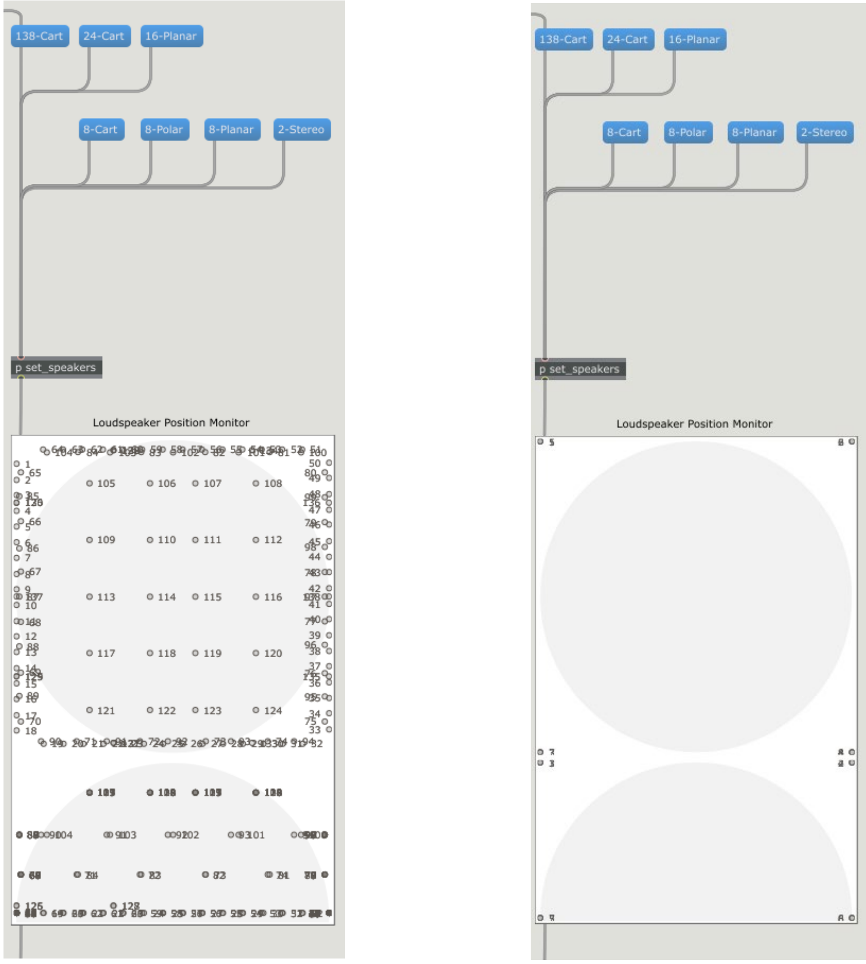

Figure 43 – Loudspeaker Positioning for the Cube (left) and the Author’s Studio (right).

Cube Loudspeaker Coordinates (138 Channels)

```
1 -0.997986 0.849206 0.040984, 2 -0.997986 0.745714 0.040984, 3 -0.997986 0.648571 0.040984, 4 -0.997986
0.547937 0.040984, 5 -0.997986 0.447619 0.040984, 6 -0.997986 0.347302 0.040984, 7 -0.997986 0.246984
0.040984, 8 -0.997986 0.143492 0.040984, 9 -0.997986 0.046032 0.040984, 10 -0.997986 -0.054286 0.040984,
11 -0.997986 -0.157778 0.040984, 12 -0.997986 -0.254921 0.040984, 13 -0.997986 -0.355238 0.040984, 14
-0.997986 -0.459048 0.040984, 15 -0.997986 -0.55619 0.040984, 16 -0.997986 -0.656508 0.040984, 17
-0.997986 -0.76 0.040984, 18 -0.997986 -0.857143 0.040984, 19 -0.824809 -0.94254 0.040984, 20 -0.697543
-0.94254 0.040984, 21 -0.570278 -0.94254 0.040984, 22 -0.44261 -0.94254 0.040984, 23 -0.315344 -0.94254
0.040984, 24 -0.184052 -0.94254 0.040984, 25 -0.060814 -0.94254 0.040984, 26 0.066452 -0.94254 0.040984,
27 0.190093 -0.94254 0.040984, 28 0.321385 -0.94254 0.040984, 29 0.448651 -0.94254 0.040984, 30 0.579944
-0.94254 0.040984, 31 0.703182 -0.94254 0.040984, 32 0.83085 -0.94254 0.040984, 33 1 -0.850476 0.040984,
34 1 -0.746984 0.040984, 35 1 -0.649841 0.040984, 36 1 -0.549524 0.040984, 37 1 -0.446032 0.040984, 38 1
-0.348889 0.040984, 39 1 -0.248254 0.040984, 40 1 -0.144762 0.040984, 41 1 -0.047619 0.040984, 42 1
0.052698 0.040984, 43 1 0.15619 0.040984, 44 1 0.253651 0.040984, 45 1 0.353968 0.040984, 46 1 0.460635
0.040984, 47 1 0.554603 0.040984, 48 1 0.654921 0.040984, 49 1 0.75873 0.040984, 50 1 0.855873 0.040984,
51 0.827225 0.940317 0.040984, 52 0.69996 0.940317 0.040984, 53 0.568667 0.940317 0.040984, 54 0.445026
0.940317 0.040984, 55 0.317761 0.940317 0.040984, 56 0.186468 0.940317 0.040984, 57 0.06323 0.940317
0.040984, 58 -0.064035 0.940317 0.040984, 59 -0.195731 0.940317 0.040984, 60 -0.318969 0.940317
0.040984, 61 -0.446234 0.940317 0.040984, 62 -0.577527 0.940317 0.040984, 63 -0.700765 0.940317
0.040984, 64 -0.828433 0.940317 0.040984, 65 -0.9706 0.794286 0.286885, 66 -0.9706 0.481905 0.286885, 67
-0.9706 0.159683 0.286885, 68 -0.9706 -0.162222 0.286885, 69 -0.9706 -0.487619 0.286885, 70 -0.9706 -0.8
0.286885, 71 -0.613774 -0.926349 0.286885, 72 -0.209424 -0.926349 0.286885, 73 0.211438 -0.926349
0.286885, 74 0.607733 -0.926349 0.286885, 75 0.972211 -0.8 0.286885, 76 0.972211 -0.490794 0.286885, 77
0.972211 -0.159048 0.286885, 78 0.972211 0.156508 0.286885, 79 0.972211 0.475556 0.286885, 80 0.972211
0.794286 0.286885, 81 0.619815 0.92381 0.286885, 82 0.20741 0.92381 0.286885, 83 -0.201369 0.92381
0.286885, 84 -0.605719 0.92381 0.286885, 85 -0.980266 0.64127 0.540984, 86 -0.980266 0.307619 0.540984,
87 -0.980266 -0.000635 0.540984, 88 -0.980266 -0.31873 0.540984, 89 -0.980266 -0.636508 0.540984, 90
-0.841321 -0.926349 0.540984, 91 -0.421667 -0.926349 0.540984, 92 -0.030608 -0.926349 0.540984, 93
0.372533 -0.926349 0.540984, 94 0.775675 -0.926349 0.540984, 95 0.972211 -0.639683 0.540984, 96 0.972211
-0.309206 0.540984, 97 0.972211 -0.000635 0.540984, 98 0.972211 0.317143 0.540984, 99 0.972211 0.634921
0.540984, 100 0.820781 0.921905 0.540984, 101 0.421667 0.921905 0.540984, 102 0.002416 0.921905
0.540984, 103 -0.400725 0.921905 0.540984, 104 -0.804269 0.921905 0.540984, 105 -0.531615 0.724127
0.811475, 106 -0.144986 0.724127 0.811475, 107 0.144986 0.724127 0.811475, 108 0.531615 0.724127
0.811475, 109 -0.531615 0.361905 0.811475, 110 -0.144986 0.361905 0.811475, 111 0.144986 0.361905
0.811475, 112 0.531615 0.361905 0.811475, 113 -0.531615 -0.00127 0.811475, 114 -0.144986 -0.00127
0.811475, 115 0.144986 -0.00127 0.811475, 116 0.531615 -0.00127 0.811475, 117 -0.531615 -0.361905
0.811475, 118 -0.144986 -0.361905 0.811475, 119 0.144986 -0.361905 0.811475, 120 0.531615 -0.361905
0.811475, 121 -0.531615 -0.726349 0.811475, 122 -0.144986 -0.726349 0.811475, 123 0.144986 -0.726349
0.811475, 124 0.531615 -0.726349 0.811475, 125 -0.998792 -0.511111 0.087432, 126 -0.998792 0.6 0.087432,
127 -0.378574 -0.936508 0.087432, 128 -0.378574 0.936508 0.087432, 129 -0.998792 -0.511111 -0.122951, 130
-0.998792 0.6 -0.122951, 131 -0.378574 -0.936508 -0.122951, 132 -0.378574 0.936508 -0.122951, 133
0.507451 -0.936508 -0.122951, 134 0.507451 0.936508 -0.122951, 135 0.998792 -0.511111 -0.122951, 136
0.998792 0.6 -0.122951, 137 -0.998792 0 -0.122951, 138 0.998792 0 -0.122951
```

24-Cartesian

```
1 -1. -0.5902 0.5238, 2 -1. -0.0164 0.5238, 3 -1. 0.694
0.5238, 4 -0.6494 1. 0.5238, 5 0.0115 1. 0.5238, 6
0.7586 1. 0.5238, 7 1. 0.3115 0.5238, 8 1. 0.0164
0.5238, 9 1. -0.4481 0.5238, 10 0.5345 -1. 0.5238, 11
0.0115 -1. 0.5238, 12 -0.5345 -1. 0.5238, 13 -1.
-0.6393 1., 14 -1. -0.0164 1., 15 -1. 0.6557 1., 16
-0.7471 1. 1., 17 0.0172 1. 1., 18 0.7586 1. 1., 19 1.
0.6339 1., 20 1. 0. 1., 21 1. -0.6776 1., 22 0.6207 -1.
1., 23 -0.0345 -1. 1., 24 -0.6322 -1. 1.
```

16-Dual Planar

```
1 -1 1 1, 2 0 1 1, 3 1 1 1, 4 1 0 1, 5 1 -1 1, 6 0 -1 1, 7 -1
-1 1, 8 -1 0 1, 9 -1 1 0.5, 10 0 1 0.5, 11 1 1 0.5, 12 1 0
0.5, 13 1 -1 0.5, 14 0 -1 0.5, 15 -1 -1 0.5, 16 -1 0 0.5
```

8-Polar

```
1 -45. 30. 1., 2 45. 30. 1., 3 -135. 30. 1., 4 135. 30. 1., 5
-45 0. 1., 6 45. 0. 1., 7 -135. 0. 1., 8 135. 0. 1.
```

8-Cartesian

```
1 -1 1 1, 2 1 1 1, 3 -1 -1 1, 4 1 -1 1, 5 -1 1
0, 6 1 1 0, 7 -1 -1 0, 8 1 -1 0
```

8-Planar

```
1 315. 0. 1., 2 0. 0. 1., 3 45. 0. 1., 4 90. 0. 1., 5
135 0. 1., 6 180. 0. 1., 7 225. 0. 1., 8 270. 0. 1.
```

2-Stereo

```
1 -1 0 0.5, 2 1 0 0.5
```

Figure 44 – Max Sub-Patch for Stipulation of Relative Speaker Positions for Each Venue

The final aspect of the MAX patch that is important to the holophonic characteristics of the audio events concerns controlling of the Z-axis, the height dimension, amplitudes. As previously mentioned, the Z-axis is perhaps the most challenging to present in a manner that is perceivable to the listener. In 3-dimensional listening environments, the z-axis is represented on at least two levels. In the Cube, there are five levels of loudspeakers. Balancing the amplitudes of each level is essential for height location perception. To address this, the first step was to create separate amplitude control for each level. Figure 45 is a screenshot from the sub-patch labeled "p LevelGains". It allows the amplitude for each of the five levels of loudspeakers to be controlled by the top left five knobs on the nK.

Though this subpatch allows for the adjustment of the amplitude on each height level of loudspeakers, deep listening is required to make use of it. The process was to first drop the amplitude of each level to -∞. Also, the sub-woofers should be set to -∞ as well. (The lower frequencies tend to reduce the perception of spatial differentiation. It is for this reason that the sub-woofers should be used sparingly in general when providing a richly spatial environment.) Then, after starting the playback of a composition, the amplitude of the ceiling level of loudspeakers should be increased to full volume. Next, one begins bringing up the Tier 3 loudspeakers until both they and the Ceiling level loudspeakers can be heard distinctly and at relatively equal amplitudes. At this point one should be able to distinguish between sonic material on each. The process is repeated for the other levels until the Floor speakers are set. Often several attempts are required before the optimal level settings are achieved. Ideally, by the time this step is completed one's ears should be acclimated to the point of being able to distinguish between the various sound events occurring on each level.

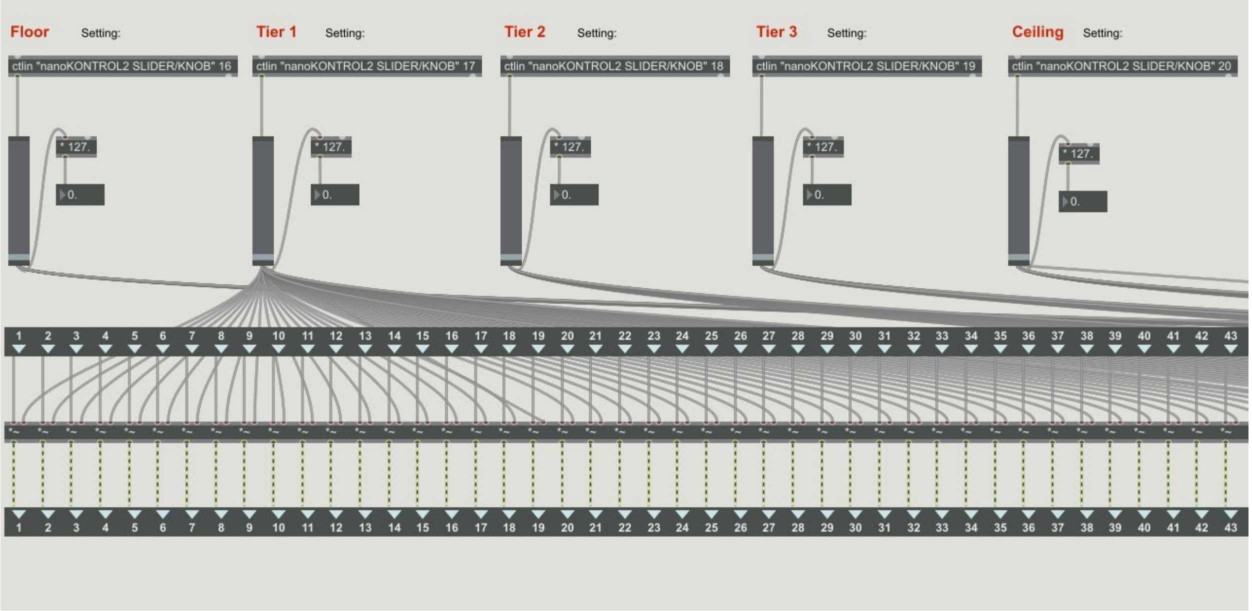

Figure 45 – Max Sub-Patch for Individual Amplitude Control of Each Speaker Level

By implementing care with regard to each aspect of this MAX patch when presenting compositions in concert or installation venues, the underlying spatialization qualities of the compositional process are augmented in a manner that aptly extends the perception of holophony. Holophony is perhaps the most viable, currently known, approach to audio immersion. It is by intelligently and perceptually combining several very effective tools that are available to the composer that the highest level of immersion is possible. As we shall see in subsequent chapters, holophony utilized in tandem with holography is a deeply engaging artistic endeavor… one that, from the author's perspective, is currently unrivaled in providing a deeply immersive audience experience.

2.20 Chapter Conclusions

As can be realized in this chapter, the auditory aspect of this project can be considered a major stratification in terms of the aesthetical and formal development of the visual music compositions in general. As such it forms one major stratum that contributes to the composition as a whole. However, more deeply considered, one can apprehend the numerous substrata involved in this generative approach to musical creation in isolation. For example, consider the SATB-like approach to layering the rendered Csound material in the DAW, wherein 32 channels of sound files are interweaved. Also consider the generative compositional process that consists of numerous levels of iteration. The 3D/360-degree spatial qualities of the holophons lend themselves toward the perception of auditory multidimensionality through the extensive foci toward exacting techniques that produce a clear perception of the location of the multi-layered auditory events.

Convolution reverberation, an extended version of VBAP, and HOA, engage a confluence and hybridization of process that empowers a facet of the contrapuntal interplay between each of the constituents involved in the intermedial interaction. As such, it provides a powerful contribution toward a multidimensional spatiotemporal experience in that this musical stratum is perceived to occur in multiple independent layers. Add to these contributions the generative algorithmic approach involving stochastic processes, that attempt to remove the composer’s conscious volition from the process, and the evolutionary progression of creative counterpoint becomes apparent.

CHAPTER 3. VISUAL COMPOSITIONAL PROCESS

3.1 Chapter Overview

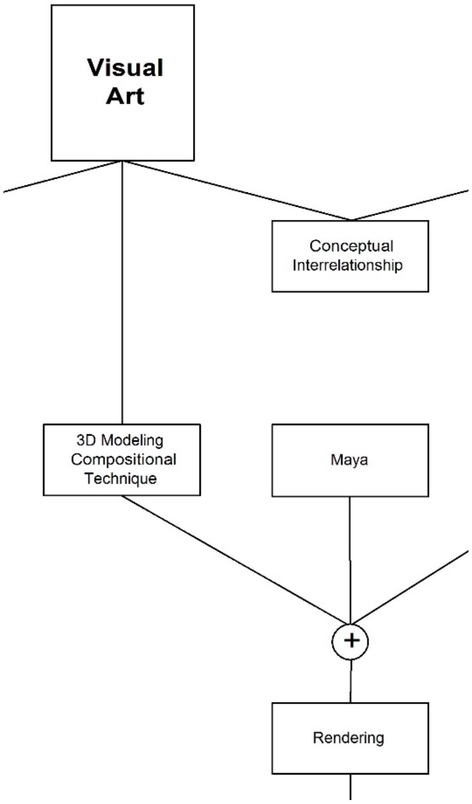

Figure 46 – The Aleatoric Visual Compositional Process

The methodologies involved in the contribution of the visual aspects of the evolutionary progression in the field of visual music this project produces offer major insights into it. Certainly, deeply investigating the creation of 3D/360-degree stereoscopic (pseudo) holograms as a major factor in a contrapuntal, with regard to the interplay of the various objects, environments, and methodology, and interdisciplinary approach fosters a multidimensional spatiotemporal paradigm that provides for a fully immersive visual music experience. Toward this, numerous strata of reflective and refractive visual environments and objects are catalyzed to create numerous other indeterminate and varied stratifications of the material produced. This aleatoric approach releases much of the composer's conscious volition to indeterminacy, which acts in a relationship of augmented intelligence between the composer and the computer.

As mentioned in chapter 2, lengthy rendering cycles require an approach to final editing that resembles the traditional tone row approach, mentioned in the Introduction of this document, to musical composition. Extensive development and variation of the expensive, in terms of CPU cycles, image sequences are required in order to facilitate the rendering of material toward a composition. It is these and other facets involved in the creation of the visual materials that will be considered in this chapter.

The term “aleatoric”, with regard to the compositional process, was a term likely first used in lectures in the 1950’s by acoustician Werner Meyer-Eppler. According to Meyer-Eppler, a process is said to be aleatoric if its course is determined in general but depends on chance in detail (Meyer-Eppler). Though he was referring to musical composition, the term may be realistically applied to any creative endeavor in general and specifically to the visual compositional approach described herein. The latter is intended to create circumstances that are conducive to emergent quanta. “*In general, emergence may be described as the effect of the outputs being greater than the sum of the inputs. Something happens… something beyond what one would expect the combination of algorithms we set into motion should be able to produce*” (Rhoades, 2009). In more colloquial terms emergence belies unanticipated events that surprise the composer.

In many statistically based studies the foci are upon the statistically relevant norms and the outliers are for the most part ignored. This is a logic-based approach. For the logician, that which does not make logical sense is often discarded. With a dialectic-based approach, we recognize the outliers as significant contributors, in this case, toward the creation of the composition. This is not to state that the norms are without merit. For it is they that provide a foundation from which the contrast of the outliers is recognizable. Not every outlier is considered relevant since not all are considered aesthetically desirable. Nevertheless, it is the outliers… the unexpected results of the process, that we are primarily concerned with here. Interestingly, this process of discovering, developing, and extending the outliers in turn creates new norms of them. As such this procedural approach properly describes and exploits an evolutionary process.

To set up the conditions conducive to emergence in the creation of the visual material for this project, we implement an indeterminate process of scene creation that involves countlessly reiterative reflections and refractions. It is through these manipulations of light that the unexpected can and often does occur. Practically speaking, it is a nearly impossible task to anticipate the results of the numerous parameters utilized in any given frame of video wherein countless photons are passed through multiple iterations of variably angled refractive and reflective objects… all of which are in motion. These constantly varying parameters consist of position, color, hue, brightness, reflectivity, opacity, and several others. Through the properties of these reflections and refractions, quasi-randomness occurs and then the resultant images are edited into a (subjectively) meaningful cosmology of photons that interact within the confines of the virtual space and objects therein.

Below, in Fig. 47 and Fig. 48, are two examples of emergent behavior derived from experimenting with two separate MAYA scenes. Ultimately, they functioned as a starting point for the visual aspects of the first two compositions associated with this project, “Twisted Chambers” and “his Dancers Three”, respectively.

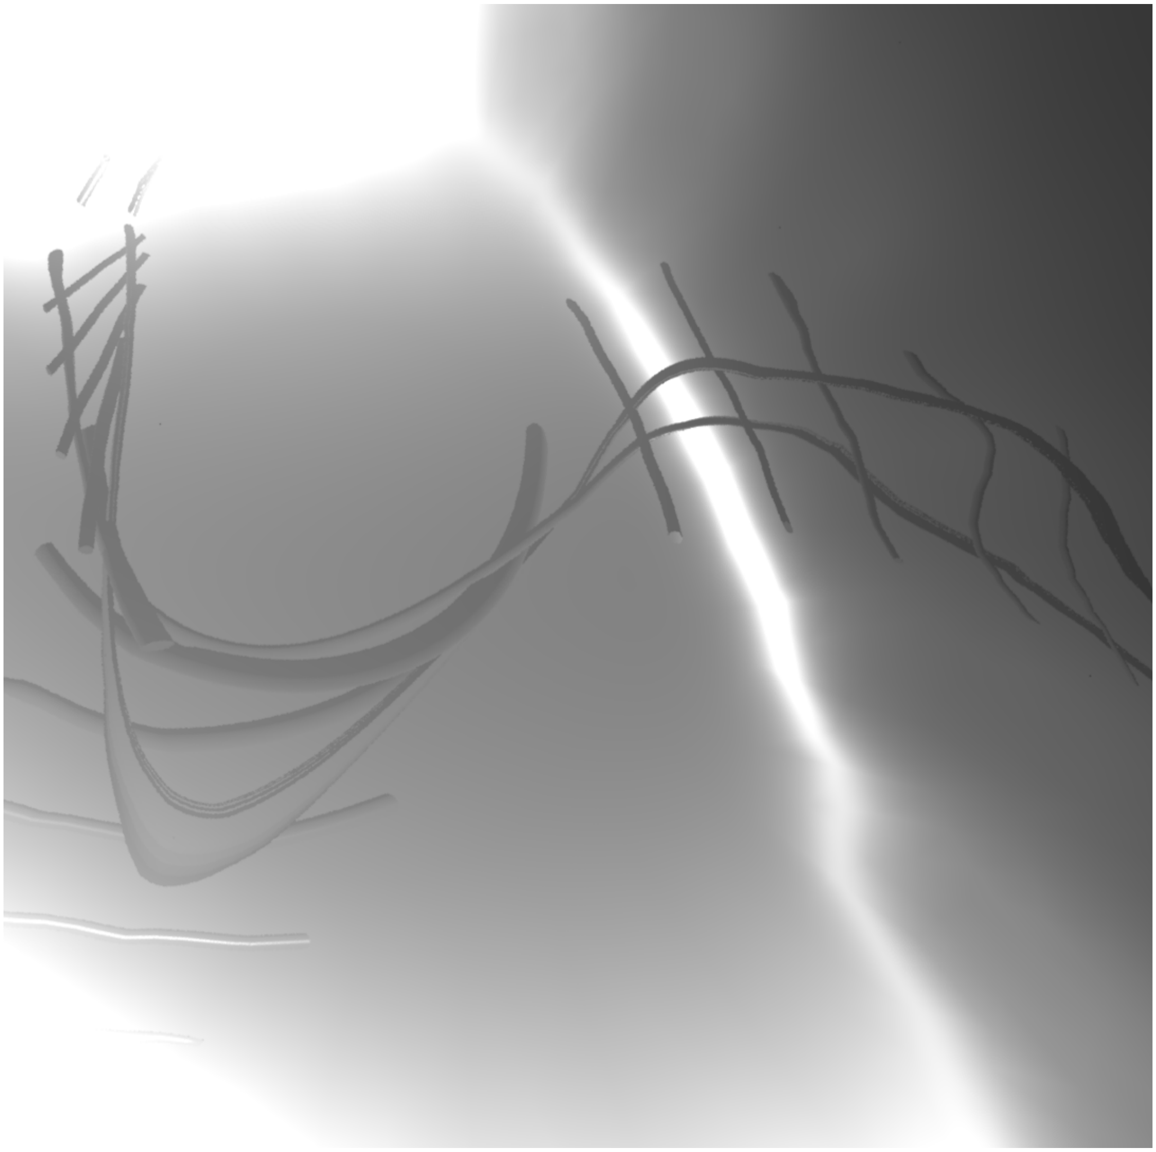

Figure 47 – Initial MAYA Sketch Derived from Emergent Behavior

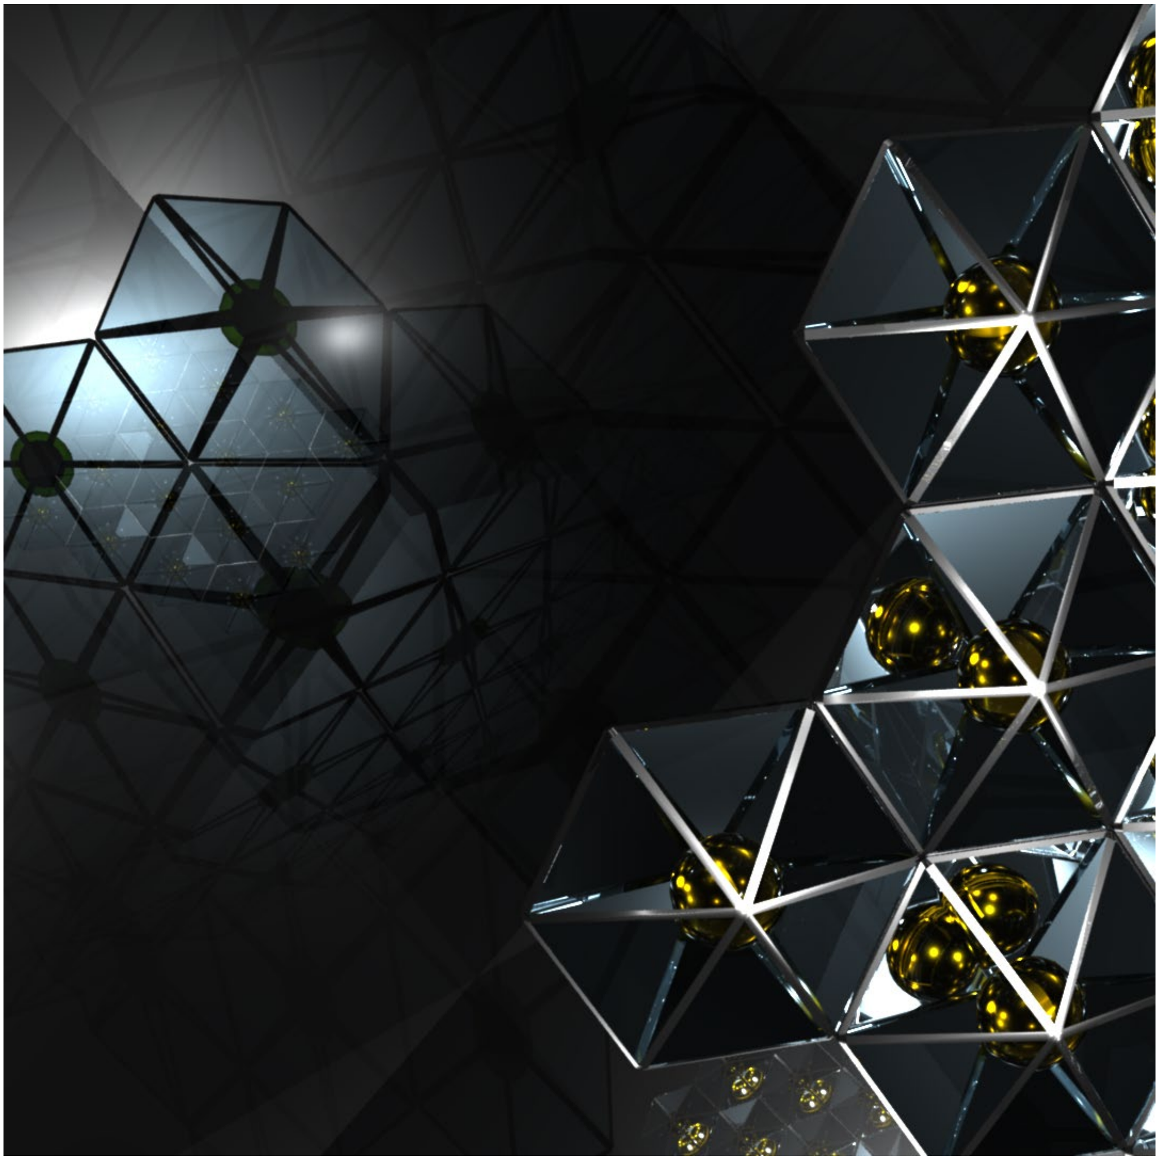

Figure 48 – Initial MAYA Sketch Derived from Emergent Behavior

This chapter will focus upon a detailed look at the technical processes involved in this endeavor in addition to providing an historical contextualization.

3.2 Historical Backdrop

The historical roots of holography extend back to 1948 and to Dennis Gabor. Gabor, a noteworthy twentieth century electrical engineer and physicist, was attempting to improve the efficiency of the electron microscope when he discovered the basic principles of holography. His first experiments were conducted using mercury vapor lamps (Abbasi, et al). Later, in the mid 1960's,

E. N. Leith and J. Upanieks successfully applied Gabor's principle of wavefront reconstruction, in practice, in the form of laser transmission holograms. During the following years many experimental, and expensive, approaches to the creation of holograms were employed. Among those were reflection holograms (as depicted in Fig. 2 in the introductory section of this document), white light transmission holograms, integral holograms, multiplex holograms, and numerous others (Benyon). Many holographic artists began finding ingenious methods to create them. One such notable artist was Walter Spiering, the founder of the Dutch Holographic Laboratory. His company was funded nearly entirely on the sale of holograms, intended for homes and offices, throughout Europe. Spiering's humble beginnings can be traced to his graduate school studies and a shed behind his student housing. There he used a broken stone tabletop and a concrete fence post to fashion the base for his vibration-sensitive equipment. A chemical engineer, Spiering's holographic techniques utilized chemical processes in photographic film plates to create his holograms (Spiering). In the 1980's holograms became increasingly desirable artifacts. They were small, often in the range of 8" x 10", typical photograph sizes for the time period, and so easily transported.

Today, with the advent of digital technology, (pseudo) holograms are much easier to create than they were in the days of using lasers, chemical processes, and photographic plates. Stereoscopic techniques employed utilizing computers and various projection devices have enabled artists, game developers, moviemakers, scientists, medical professionals, and others, to implement a wide range of 3D projection techniques. For the purposes of this research, the foci are upon the artistic use of stereoscopic imagery to create virtual 3D/360 holographic art.

Though at the time of this writing there are relatively few artists using the virtual holographic medium, the potential is increasing and so is the implementation. As head-mounted displays become less expensive, less cumbersome, and more effective, it is reasonable to anticipate increasing interest in the development of holographic art. Work in this area also sets the precedence for holographic movie making as large-scale projection systems, such as those currently being built and installed by the Elumenati®, become viable and perhaps eventually ubiquitous. Certainly, it is one goal of this research to contribute a perspective, on both a creative and a technical level, for the advent of this foreseeable future in the area of holography.

3.3 Maya Version

After the initial creative concepts for a scene and its variations are conceived of, which will be discussed in Chapter 5, the technical process begins with, and is centered around, Autodesk® Maya 2015 – Student Version. Though as of this writing Autodesk has offered new versions of Maya every year up to and including Maya 2020, Maya 2015 is being used for this project for several reasons. Firstly, it is a very stable version of Maya. When working on renders that will require weeks to complete, it is highly undesirable for crashes or glitches to occur during the process. Maya 2015 is not only utilized for scene creation but also on each node of the rendering cluster. The latter will be discussed in detail in Chapter 4. Secondly, Maya 2015 was the last stable version to accommodate the Mental Ray (MR) renderer. The textures and materials afforded by Mental Ray are considered, as a personal preference, superior in many ways to more recently prominent renderers such as Renderman® and Arnold®. Though the latter are arguably far easier to use than MR they seem to tend toward a homogeneity of aesthetic whereas MR is more conducive to stylization. Again, this is a personal choice, perhaps based upon familiarity. It is not a judgment nor is it authoritative.

Perhaps the most important reason to continue use of Maya 2015 for this project is the DomeMaster3D® plugin. It is available for Maya 2015 – 2018 but is not supported for Maya 2019 and beyond. Considering the available renderers, this makes Maya 2015 the obvious choice. The DomeMaster3D plugin, produced by Andrew Hazelden and Roberto Ziche, is key to the stereoscopic 3D rendering for this project. It also allows for both 2D and 3D full dome rendering providing another outlet for this work in the planetarium arena. As of this writing no plugin with similar features is available.

It is perhaps generally considered unfashionable to continue working with older versions of operating systems and programs when newer and supposedly more highly optimized versions exist and yet, if a version works for a given project, if it is available, and if it is comfortable to work with, it is argued here that there is no reason to upgrade unless forced to do so. One such event would be the necessity to upgrade to a newer operating system because new hardware does not support the previous versions. However, less computer maintenance and more creating is the central focus here. Certainly, as we shall see later, upgrading to a newer version of Maya is a major undertaking considering that it involves doing so not only with the workstation but also with the rendering cluster, which is a much more complex process. This represents less time for rendering material for this, or any other, project.

3.4 The Process

Figure 49 below is a flowchart that represents the workflow for creating the visual aspects of the visual music compositions that are the foci of this research. As with the Score-Based Sampling flowchart in Chapter 2, we will drill down into the blocks in the diagram and provide detailed descriptions of the activities involved in each. As a reminder, we are making linear descriptions of non-linear processes in order to provide a symbological verbal framework of understanding. In practice, though this pipeline is generally not sequential, the creation of a composition occurs on many levels simultaneously and in innumerable directions. The reader should consider the following a simplification of the processes involved.

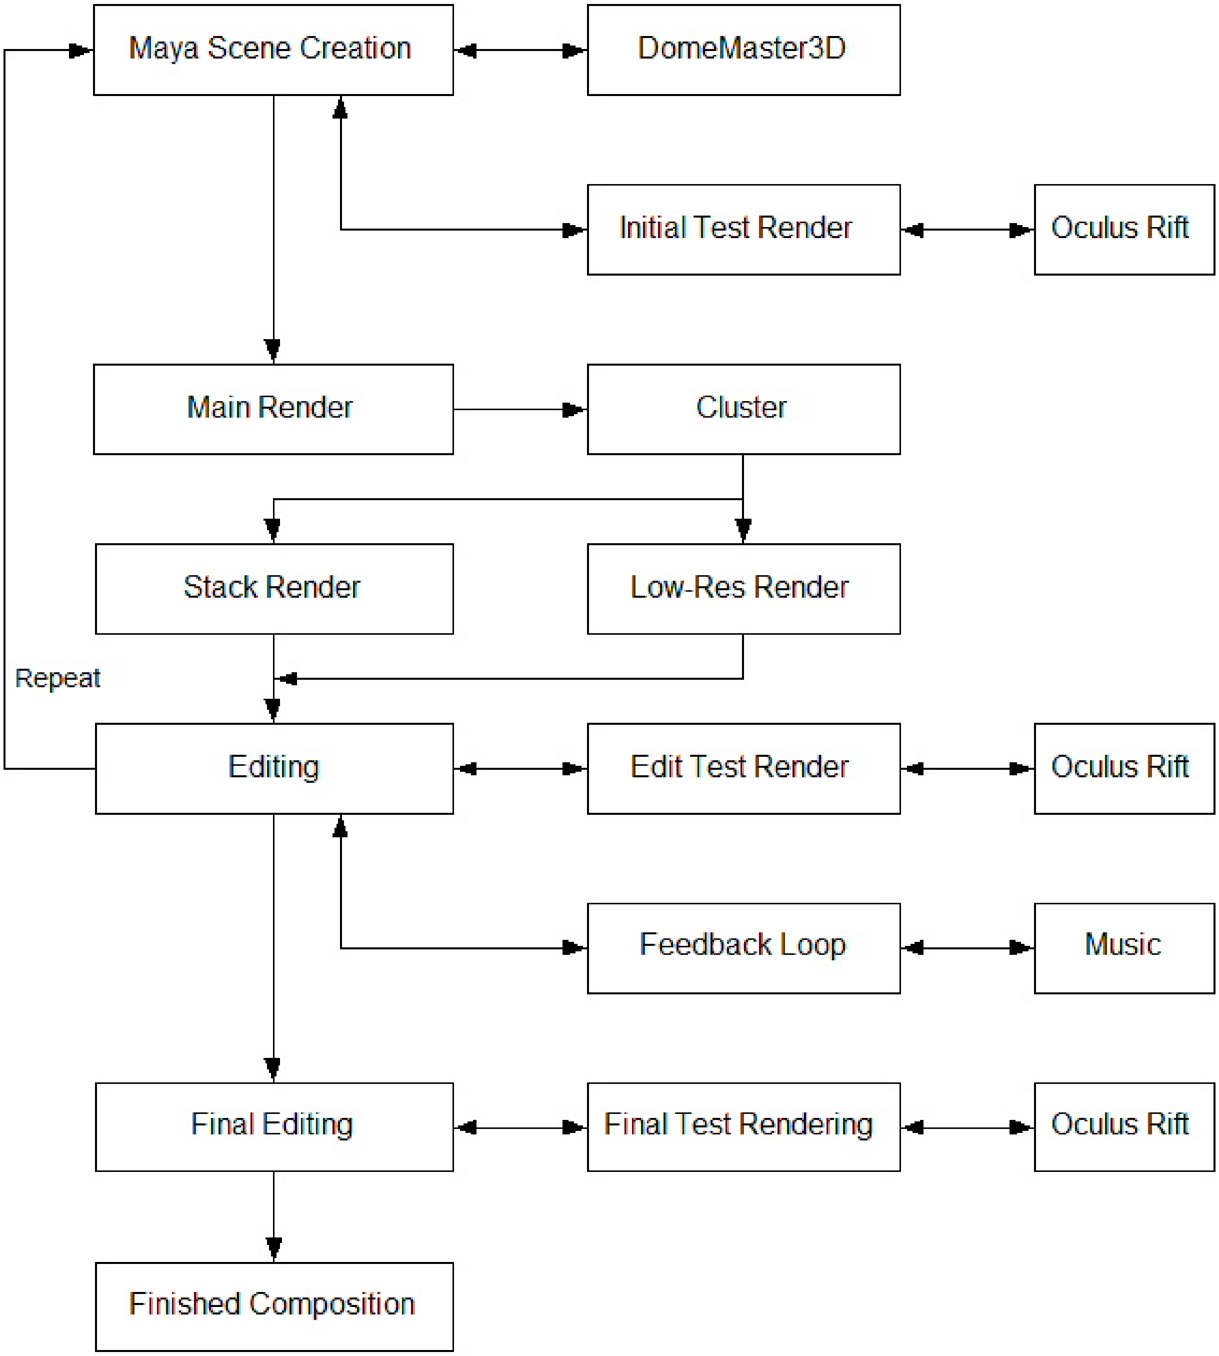

Figure 49 – Overview of the Visual Compositional Process

3.5 Maya Scene Creation and DomeMaster3D

As previously mentioned, aesthetical perspectives on scene creation are discussed in Chapter 5. Here we discuss the mechanics of working with Maya and DomeMaster 3D thus enabling a creative flow. Pictured in Fig. 50, we begin with the default Maya layout as seen through the perspective view.

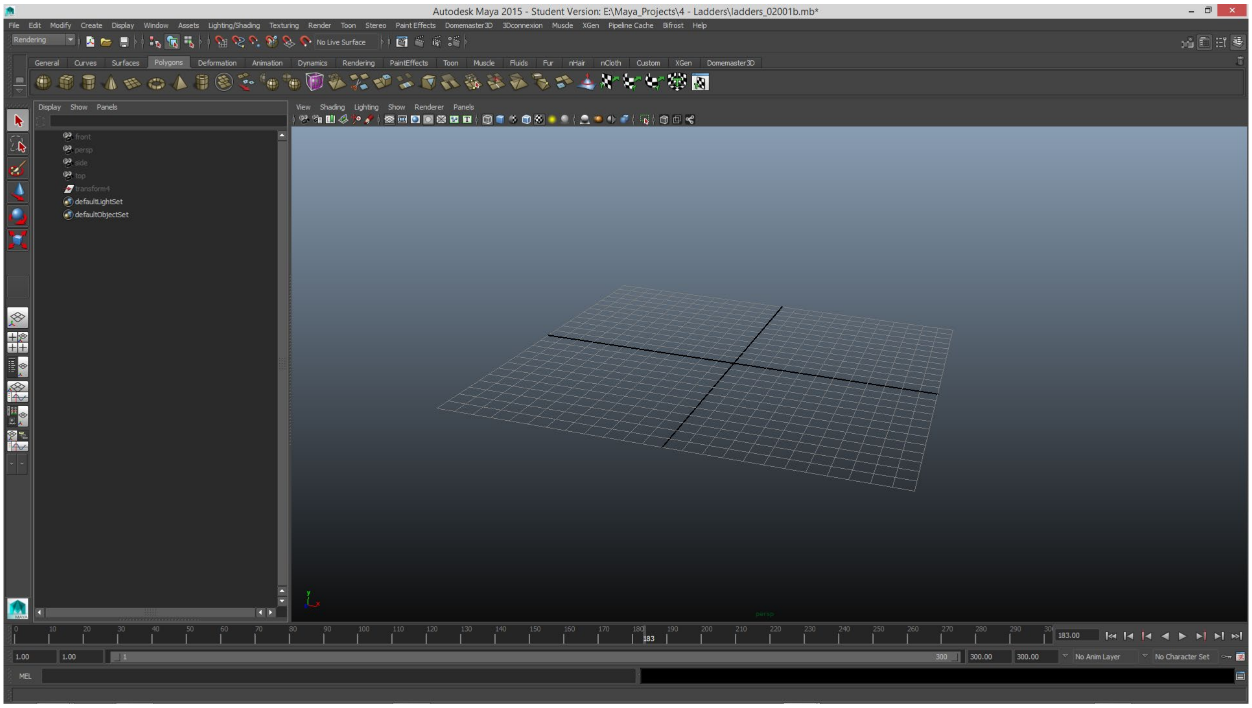

Figure 50 – Basic Maya Layout

There is a necessity to steadfastly maintain the basic layout and parametric relationships in a consistent manner for every scene. Numerous hours of testing and rendering were initially invested to establish the exact parameters related to this layout. Beginning with the default size for the grid, shown in Fig. 50, in the center of the perspective view window, the proportional dimensional relationships between all objects, environments, and cameras must be maintained throughout the entire project. These relationships are critical when working with stereoscopic rendering. If one aspect is out of sync with the others, the subsequent renders will not be viewable from a stereoscopic perspective. A render that looks fine in 2D may be completely unusable in stereoscopic 3D/360. It is through maintaining this consistency from scene to scene that the tedious and time-consuming work of re-testing is avoided.

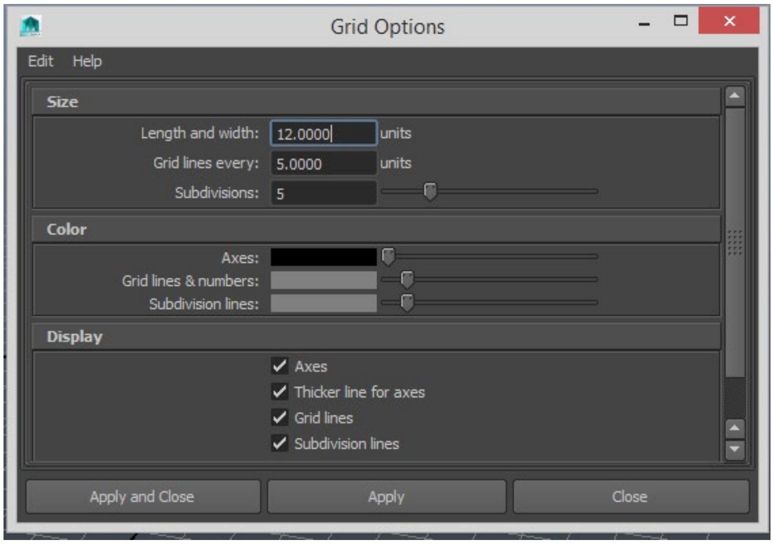

Length and width: 12.0000 units
Grid lines every: 5.0000 units
Subdivisions: 5
Axes:
Grid lines & numbers:
Subdivision lines:
☑ Axes
☑ Thicker line for axes
☑ Grid lines
☑ Subdivision lines

Figure 51 – Default Grid Settings

A set of keyboard shortcuts, Hot Keys, is configured in the Preferences menu that support the author's working process. In addition to the usual, w – Move, e – Rotate, and r – Scale, of primary importance are the following custom Hot Keys:

- g - grouping
- u - ungrouping
- d - duplicate
- p - parenting
- z - center pivot
- b - combine
- s – separate

The 3DConnexion® SpacePilot™ Pro, Fig. 52, is an invaluable and integral tool in this process. Of main focus is the ergonomic left-hand perch that is extremely comfortable over long periods of scene creation. The center knob features a "six degrees of freedom" paradigm that allows for moving the currently used camera in a manner not afforded any other method. There is a learning curve when first using the device but it is well worth the investment to become accustomed to it. It allows for the camera to be moved in the following directions with the single knob:

- Up/Down
- Spin Right/Left
- Push Forward/Reverse
- Raise/Lower
- Tilt Right/Left
- Tilt Forward/Reverse

This ability allows for smooth, intuitive, and effortless workflow not possible with more standard manipulation tools.

Figure 52 – 3D SpacePilot™ Pro Mouse

The DomeMaster 3D plugin is a major key to the stereoscopic rendering in this project. Though it offers several options, the LatLong Stereo camera rig is essential. It is a three-camera setup consisting of left- and right-eye cameras and a center camera. The latter is not functional; instead it serves as a parent to the left- and right-eye cameras. Moving it moves the entire three-camera rig. The viewable left- and right-eye cameras capture images at 360 degrees on the horizontal plane and 180 degrees on the vertical. Figure 53 pictures the camera rig in its default configuration both with regard to its orientation on the grid and its relative scale.

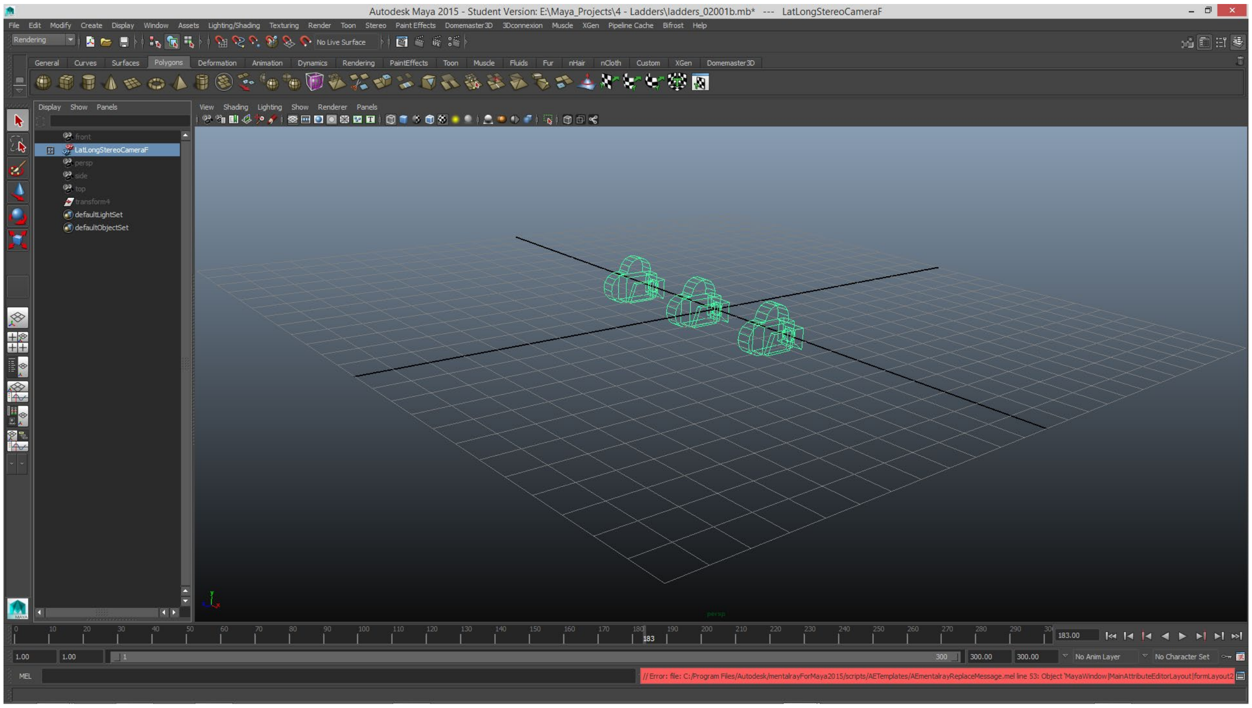

```
// Error: file: C:/Program Files/Autodesk/mentalrayForMaya2015/scripts/AETemplates/AEmentalrayReplaceMessage.mel line 53: Object 'MayaWindow|MainAttributeEditorLayout|formLayout2
```

Figure 53 – DomeMaster 3D LatLong Camera Default Configuration

Figure 54 demonstrates the default LatLong camera rig settings. Of particular interest are the "Camera Separation", also known as the interaxial separation, and the "Zero Parallax Distance" values. The former determines the distance between the left eye and right eye cameras. This correlates to the distance between the eyes of a human as is relative to the size of objects in the scene and the dimensions of the overall space, which in this case generally occupy confines of the grid. The Zero Parallax Distance determines the distance at which the convergence of the left- and right-eye images occurs (Zicher3d-org). Any stereoscopic image consists of two images… the left- and right-eye. These images, when overlapped, are slightly offset depending upon the distance of the object(s) and environments(s) from the eyes.

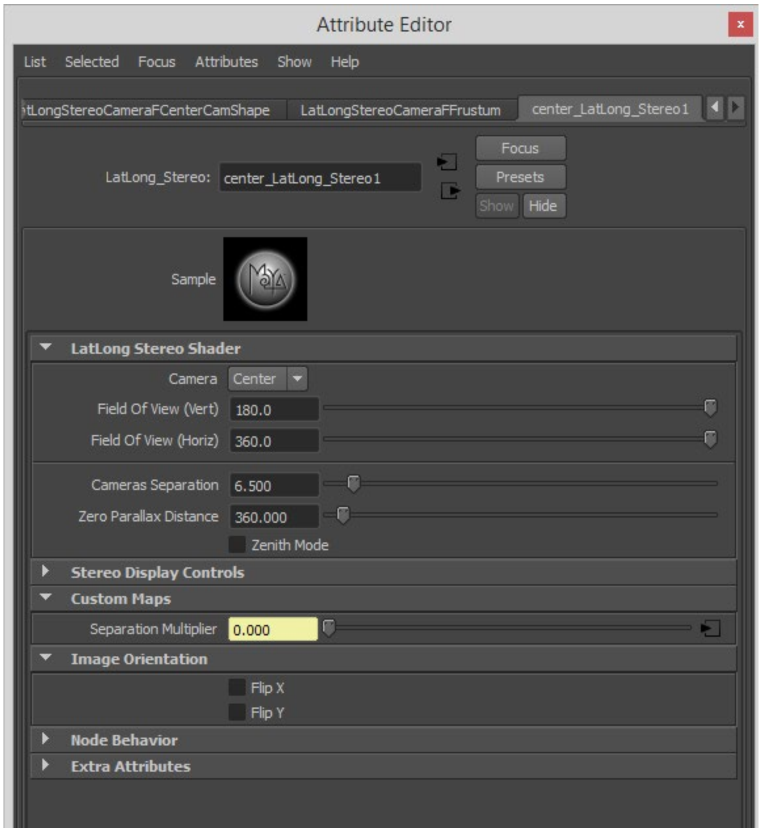

LatLong_Stereo: center_LatLong_Stereo1

Camera: Center

Field Of View (Vert): 180.0

Field Of View (Horiz): 360.0

Cameras Separation: 6.500

Zero Parallax Distance: 360.000

☐ Zenith Mode

Separation Multiplier: 0.000

☐ Flip X

☐ Flip Y

Figure 54 – DomeMaster 3D LatLong Camera Default Attributes

After numerous hours of testing, the standard values arrived upon for this research were .65 for the "Cameras Separation" and 36.5 for the "Zero Parallax Distance". These values make the cameras appear as pictured in Fig. 55 below in Maya's perspective view as compared to the default appearance. They provide the relationships between object size, grid size, and camera thus establishing constants that allow for reliably feasible stereoscopic rendering.

Using these standard parameters and tools it is possible to remove questions of these technicalities from the process and allow creativity to take charge. Though it engages numerous other technical decisions, since the basic configuration can be relied upon the creative process can proceed with fewer concerns as to the viability of a given scene. Using these simple standards, we have not only established a sound foundation from which to operate but we have also created a consistency that will pervade the entire process. Making the technical basics as transparent as possible for each scene allows the least impedance to the overall creative process. As competence grows within these parameters a transparency of process ensues. Figure 56 demonstrates a completed Maya scene in wireframe view, which is ready for the next step.

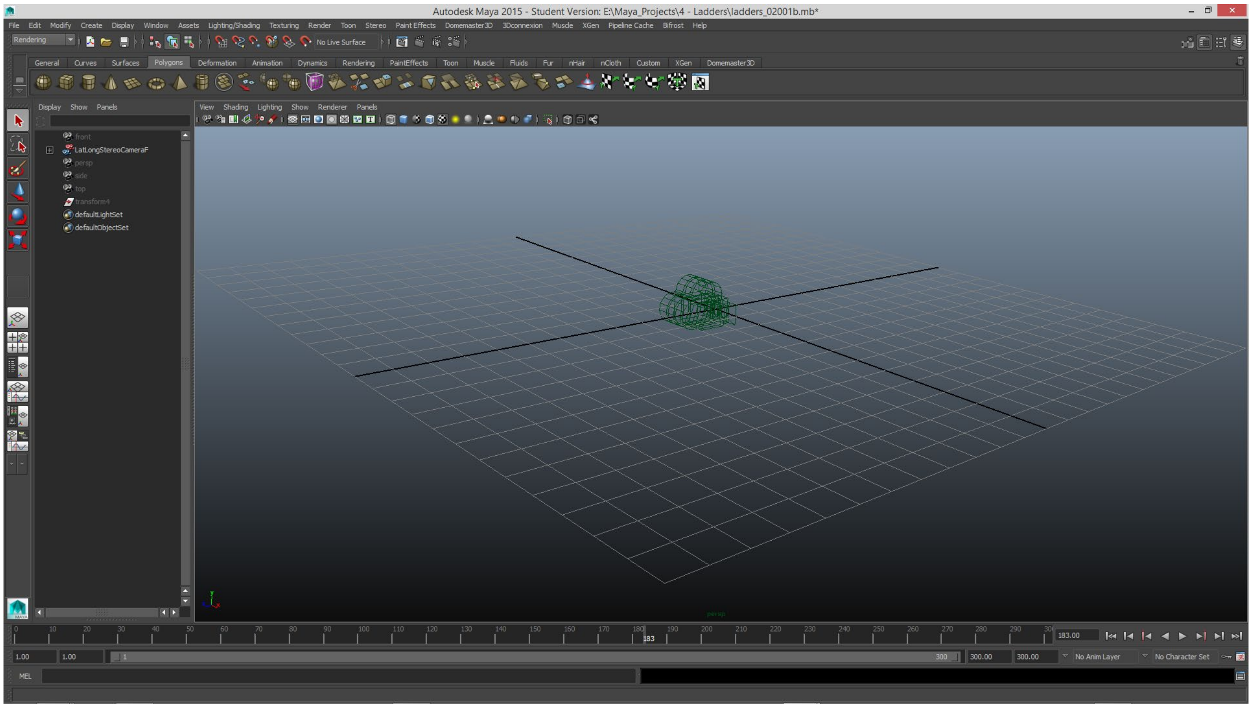

Figure 55 – DomeMaster 3D LatLong Camera Applicable Configuration

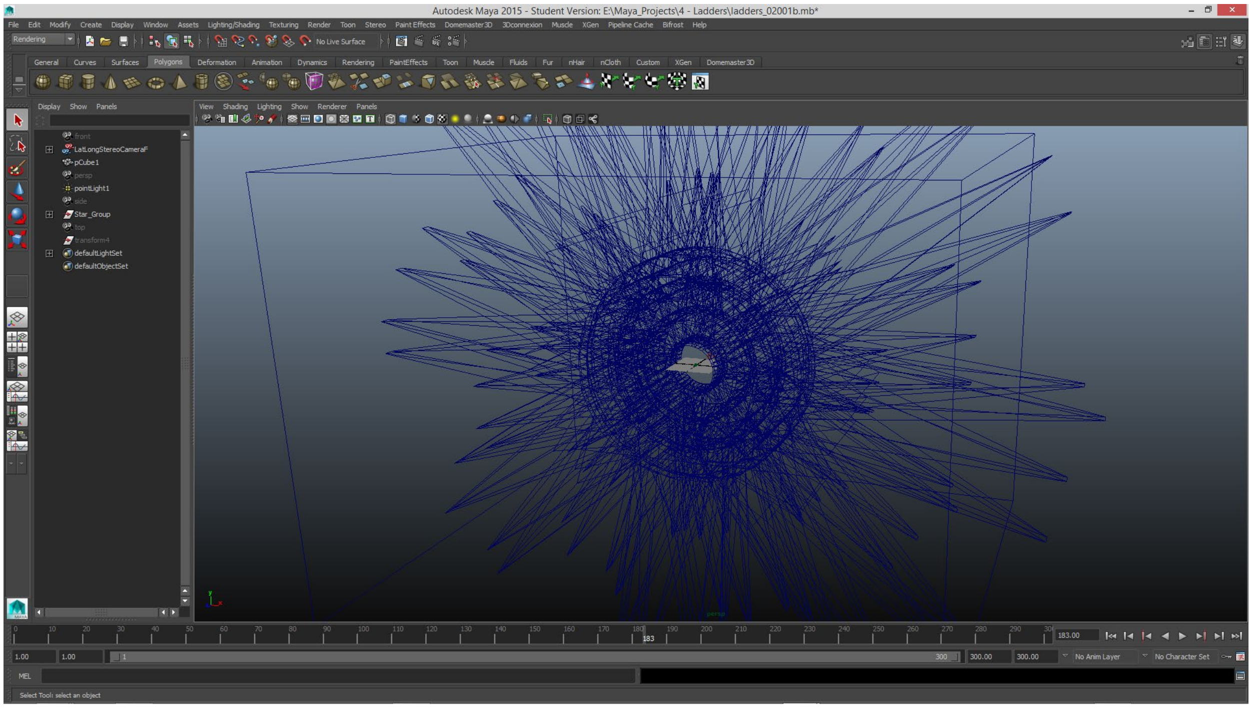

Figure 56 – An Example of a Completed Maya Scene – Wireframe View

One ancillary approach to 3D modeling was obtained through the use of ZBrush, which is a 3d modeling/sculpting program. ZBrush can be utilized to create .obj objects that can be imported into MAYA. Using ZBrush, one begins with a virtual lump of clay, of various sizes and shapes, and then sculpts it into the desired shape using a multitude of sculpting tools called brushes. This process is predisposed to very organic shapes. Generally, one adds “clay” as needed to fill in the gaps as one extends the geometry of the object. However, nearly exclusively for this project, the “move” brush was utilized, which simply stretches and compresses the clay. The result of this approach was very thin clay. When the number of polygons was increased using this very thin clay model, holes occurred in the geometry that, though usually unwanted, created an effect that was quite compelling and inspiring. From the image in Fig. 57, one can see the resultant meshing effect. This approach to 3D modeling proved quite effective and was implemented numerous times throughout this project. It greatly extended the potential of MAYA.

One desirable aspect of the renders using this technique is that these holes in the geometry become completely transparent when rending with the alpha channel enabled. The Image Based Lighting (IBL) images that are occasionally used to light the scenes provided very interesting backgrounds to the renders. Where no IBL was employed, a layering of various renders was possible and could be viewed as layers behind the foreground layers and objects. This provides for several creative openings to be explored further in the future.

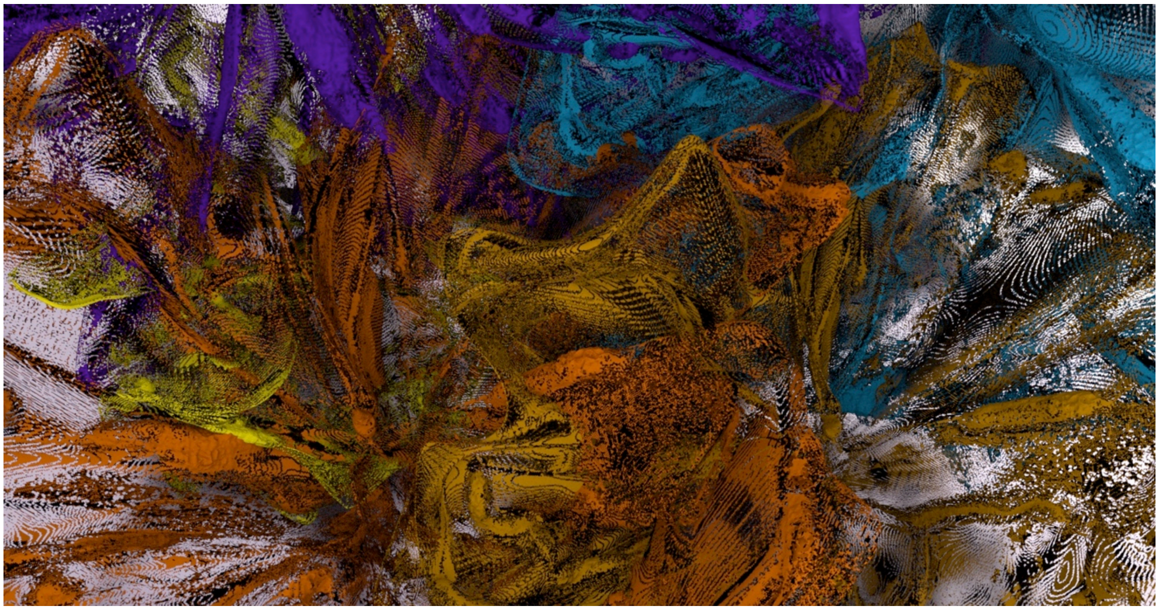

Figure 57 – Meshing Caused by the Extensive Use of the ZBrush Move Tool

3.6 Texturing and Lighting

In Maya, texturing and lighting are critical in creating beautiful scenes. As previously mentioned, Mental Ray is the author's preferred renderer. Along with it comes great flexibility in creating stylized scenes. Mental Ray consists of a rendering engine and a powerful set of default textures and lights. One specific texturing set is particularly useful and is focused upon nearly exclusively. It is called mia_material_x.

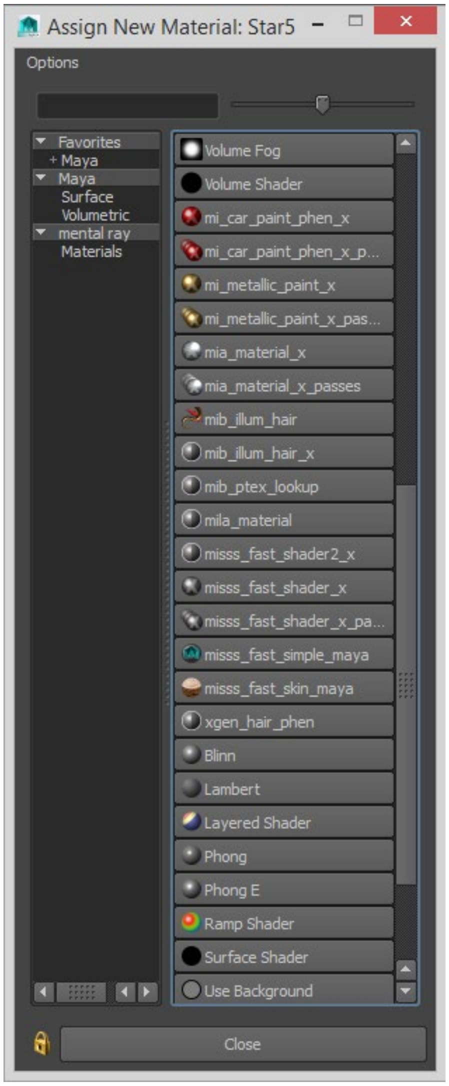

Figure 58 – Assign New Material Menu - mia_material_x

With the mia_material_x group comes textures such as chrome, copper, several types of glass, plastic, and rubber. Choosing the defaults is quite viable however the real power comes with adjusting the parameters using the Hypershade window or in the Attribute Editor, the latter of which is pictured in Fig. 59. Beautiful mirror, chrome, and glass textures are relatively simple to achieve and apply to geometry. The combinations of colors used for the diffuse, reflected, and the refracted properties of each associated virtual photon are important attributes as are other properties such as transparency and glossiness. If it is desirable one can delve even deeper into other properties to achieve more extensive control. Further, in the "Render Settings" menu there are controls for indirect lighting options that can make a huge difference in render quality. These include Global Illumination, Final Gathering, and Ambient Occlusion. Each of these can powerfully affect the outcome of a render yet, for this project, they were avoided due to the increase in render times required to achieve them.

As an oil painter, my work is primarily an exploration of color and space. This perspective is directly pertinent to my work with Maya where the relationships between the color structures are extremely important. The size and shape of objects, whether abstract or non-representational, and the colors that combine to form them, and in some circumstances illuminate them, create a vibrational counterpoint. Stephen Quiller's (Quiller) color theories certainly demonstrate this and his scientific use of exact opposite colors and their exact compliments are pertinent when considering any image… including those produced in Maya. For various aspects of this project color pallets, derived from photographs of my oil paintings, were used to create alternatives to those that are a part of MAYA. Using oil paints to create color swatches that are accurately photographed demonstrated a greatly extended color range and sensitivity compared to using the colors included in Maya alone.

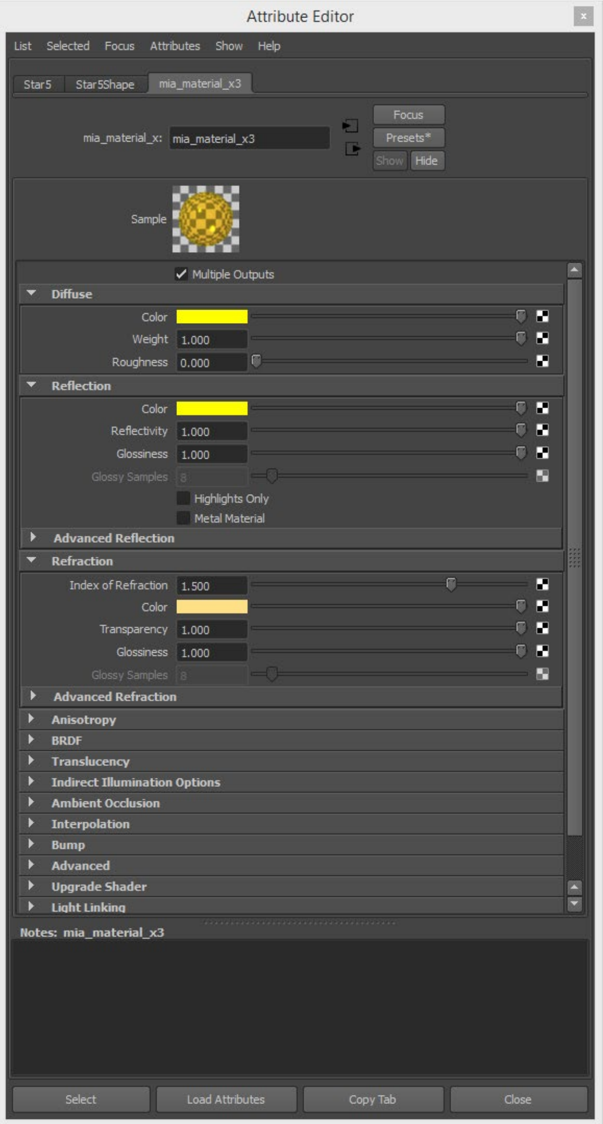

mia_material_x: mia_material_x3

☑ Multiple Outputs

Diffuse
Weight 1.000
Roughness 0.000

Reflection
Reflectivity 1.000
Glossiness 1.000
Glossy Samples 8
☐ Highlights Only
☐ Metal Material

Refraction
Index of Refraction 1.500
Transparency 1.000
Glossiness 1.000
Glossy Samples 8

Notes: mia_material_x3

Figure 59 – Material Attribute Editor

The textures alone are not enough to create viable scenes. Proper lighting is what brings the them to life. Lighting in Maya resembles direct lighting techniques used in photography and filmmaking (Derakhshani). Lighting can make or break a scene. Countless hours have been spent on this project working with the lighting, arranging and rearranging it until the scene is perfectly lit. Especially when working with reflections, lighting can be challenging. This is further compounded when working with 3d/360 scenes, which can be viewed from many perspectives. Though others are often implemented, the preferred light for the reflective and refractive environment is the point light. A point light emits light in all directions emanating from its center thus making it a good choice for a holographic scene. Intensity is an important parameter of a point light as is color and shadowing… each of which can be keyframed. An image file may also be used as a reference color for a light, which can create very interesting lighting effects.

A lesser-known attribute of lighting in Maya, that is extremely useful when working with reflections, is “Light Linking” in the “Relationship Editor”. With it one can turn off and on the effects of any light on any of the various objects in a scene. This allows, for example, an abstract tree object to be illuminated by a light but the reflective background can be omitted from the influence of the light. From this simple example it is obvious that lighting is important and complex. It could be, and has been, a complete study on its own (Lanier). Figure 60 demonstrates a rendering of a well-lit scene. The object and the background use the mia_material_x glass and chrome textures. It is a simple cylinder and a sphere that achieve complexity through reflection and refraction.

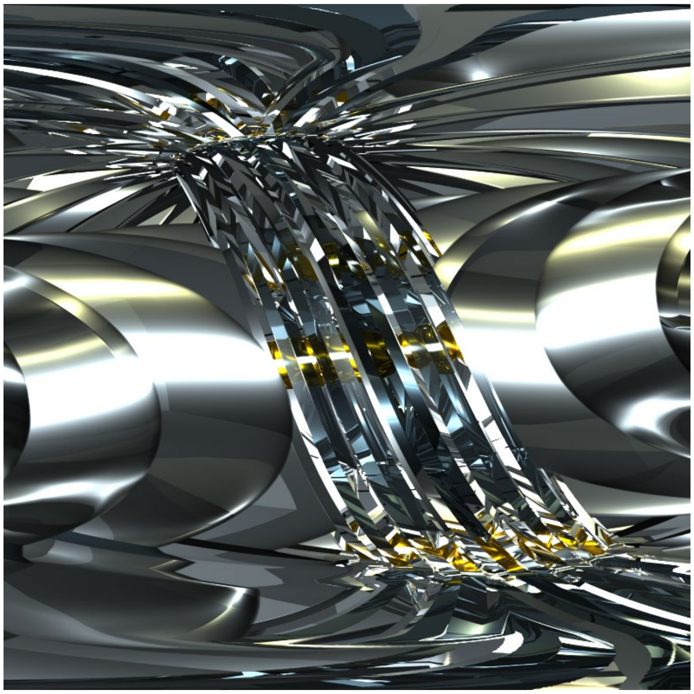

Figure 60 – Example Render using mia_material_x Glass and Chrome Materials

3.7 Basic Maya Movement – Keyframing

The concept of keyframing is common in most forms of animation and it is a basic tenet of movement in Maya. The simple premise is that the animator determines the position of an object during a specific frame time in a series of frames that make up the scene. Another position is then determined in another frame time and then Maya interpolates the movement between the two frames. The author generally works with 2400 frame scenes and each frame requires two images to be rendered… a left-eye perspective image taken from the left camera and a right-eye perspective image taken by the right camera. In a 2400 frame scene there are usually numerous keyframes. In addition, the animator can change numerous parameters for each keyframe including or excluding them as desired. These parameters range from rotation, position, and scale on the x-, y-, and z-axis material attributes, lighting attributes, and etc. Virtually any parameter in Maya can be keyframed making for richly complex scenes. Figure 61 demonstrates a keyframed scene. The red lines in the timeline at the bottom of the image are the keyframes. For simplicity here, it is a 34-frame scene with 34 keyframes.

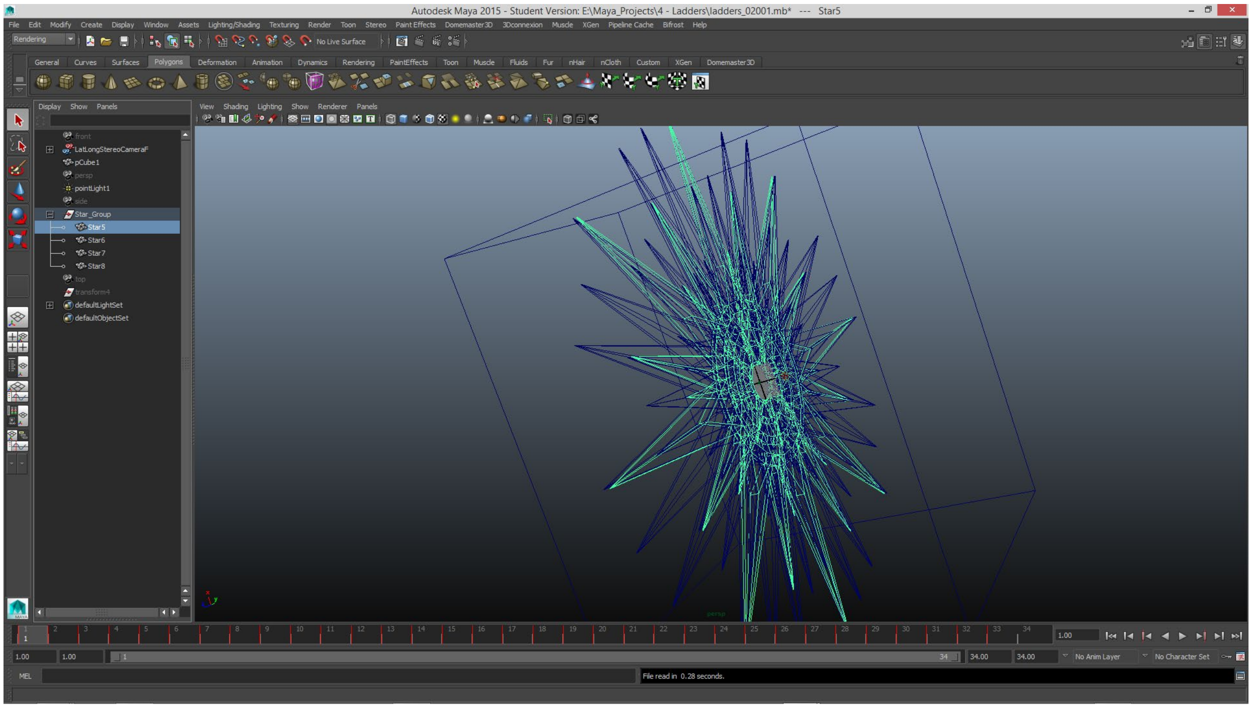

File read in 0.28 seconds.

Figure 61 – Example of Keyframing in Maya

3.8 Initial Test Render and the Oculus Rift

After a scene is textured, lighted, keyframed, and saved, it is time to conduct test renders. Usually, while setting up a scene, many lower resolution 2d single-frame renders have been run. These are accomplished using one of the LatLong cameras. The author prefers to do this with the left camera for consistency. The final resolution of the scenes is 7000 x 1750 pixels, more on this later, however to save time for the test renders 1752 x 438 renders are usually run. This maintains the 4:1 ratio and provides a reasonable version of what is happening in the camera. The renders are fisheye rather than flat, however with a little practice it is not difficult to interpret what is happening with the image prior to seeing it displayed in a projection format that shows it flattened.

Though a baseline with regard to camera/object/grid relationships is well established and test renders have been run, there remains a wide margin for error when creating these scenes. Usually these errors are in the area of objects being placed too close to the cameras causing a cross-eyed effect, or an object being too close to one eye causing an eye-patch effect. Yet through countless test renders myriad other potential issues have arisen that only became apparent after rendering stereoscopic test images and seeing them in 3D. Head-mounted displays are perfect monitors for such test renders. Firstly, they are one of the target mediums for this project and secondly, they offer a very close equivalent to, for example, what will be projected onto the Cyclorama in the Cube.

To render stereoscopic test frames in Maya we can render the left eye camera view using the renderer, save it and then render the right eye camera view. Or we can use the Maya batch renderer and chose only single frame to render. Either way we end up with two images that we stack using After Effects and then view them using the HMD. The techniques utilized to achieve this test will be disseminated in detail in section 3.10 below. Ultimately, single frames will consist of a configuration similar to the image in Fig. 62. If the image is viewable in the HMD without issue and if it looks as intended or, better yet, surprises in a positive manner, it is time to render the Maya scene using the high-performance computing cluster. If not, it is back to the drawing board. The only manner in which one can fail in this work is to give up too soon. Otherwise, there are no mistakes, only work that remains finished.

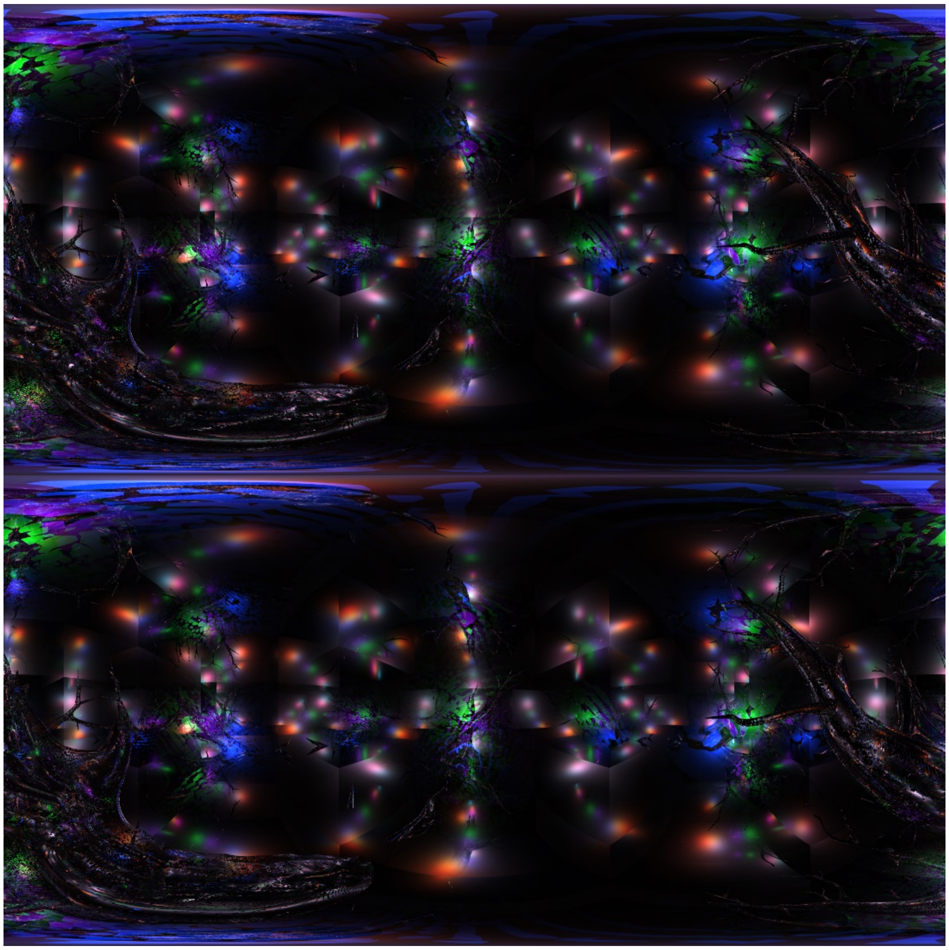

Figure 62 – Example of Stacked Test Image for Viewing on the Oculus Rift

3.9 Main Render and the Cluster

After several test renders were deemed successful for a given scene, both high- and low- resolution, the scene is nearly ready to send to the cluster for the final high-res rendering. There are a few things that need to be double-checked before doing so. Firstly, since both high- and low-resolution tests were conducted we need to make sure that the render settings are correct. They will be included with the scene when sent to the cluster. The resolution of the renders should be set to 7000 x 1750 pixels and the renderable camera should be set to the LatLong camera for the scene. Also, since the HPC cluster is Linux-based, one very important aspect of the render has to do with the path to the DomeMaster 3D separation map situated in Maya. The path can be found in the Window/Rendering Editors/Hypershade/mental ray/textures menu, as depicted in Fig. 64 below. There are usually two boxes… each with names that begin with “separation”. When selected there is a path window for each in the Attribute Editor. Make sure that any absolute windows paths are removed. All that should remain is “sourceimages\latlong _separation_map.png” with no beginning back slash. In that manner, if the cluster is set up correctly, the various Maya instances on the cluster nodes will be able to find the files, that DomeMaster 3D needs to operate, in the directory in which the scene is located. As will be discussed in detail later, DomeMaster 3D is installed on each render node of the cluster.

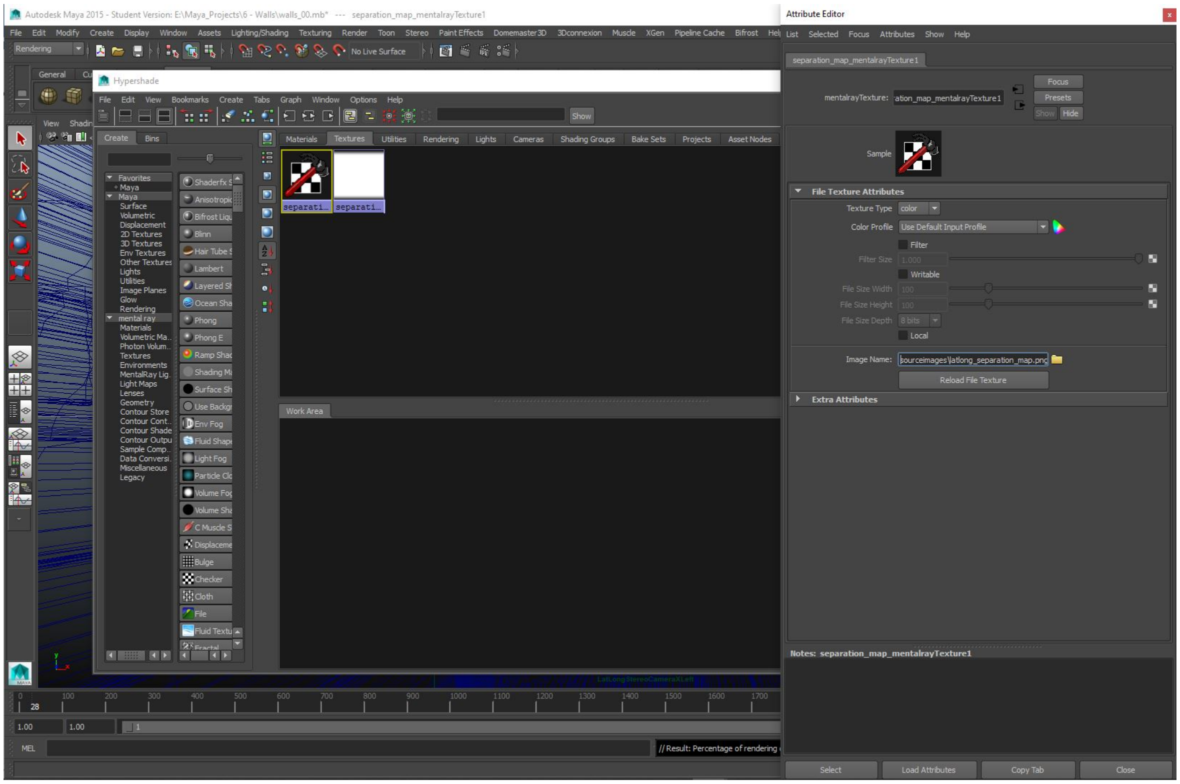

Autodesk Maya 2015 - Student Version: E:\Maya_Projects\6 - Walls\walls_00.mb* --- separation_map_mentalrayTexture1

mentalrayTexture: ration_map_mentalrayTexture1

Texture Type: color

Color Profile: Use Default Input Profile

Filter Size: 1.000

File Size Width: 100

File Size Height: 100

File Size Depth: 8 bits

Image Name: sourceimages\latlong_separation_map.png

Figure 63 – Ensuring the Correct Path for the DomeMaster 3D Separation Image

The Maya scene is then ready to be copied from the workstation to the hard drive on the head node of the cluster. To do that we use SFTP, secure file transfer protocol, on port 21. This can be accomplished using command line or using an SFTP client like FileZilla®, Fig. 64. The file is transferred into the appropriate directory, the “scenes directory”, on the cluster head node. The file structure will be discussed in Chapter 4. Via command line using either SSH (secure shell) or directly on the head node, the cluster is directed to execute the render. When the render is completed, many days later depending upon the complexity of the images, the left and right camera images are harvested, which means that they are transferred using SFTP back to the workstation.

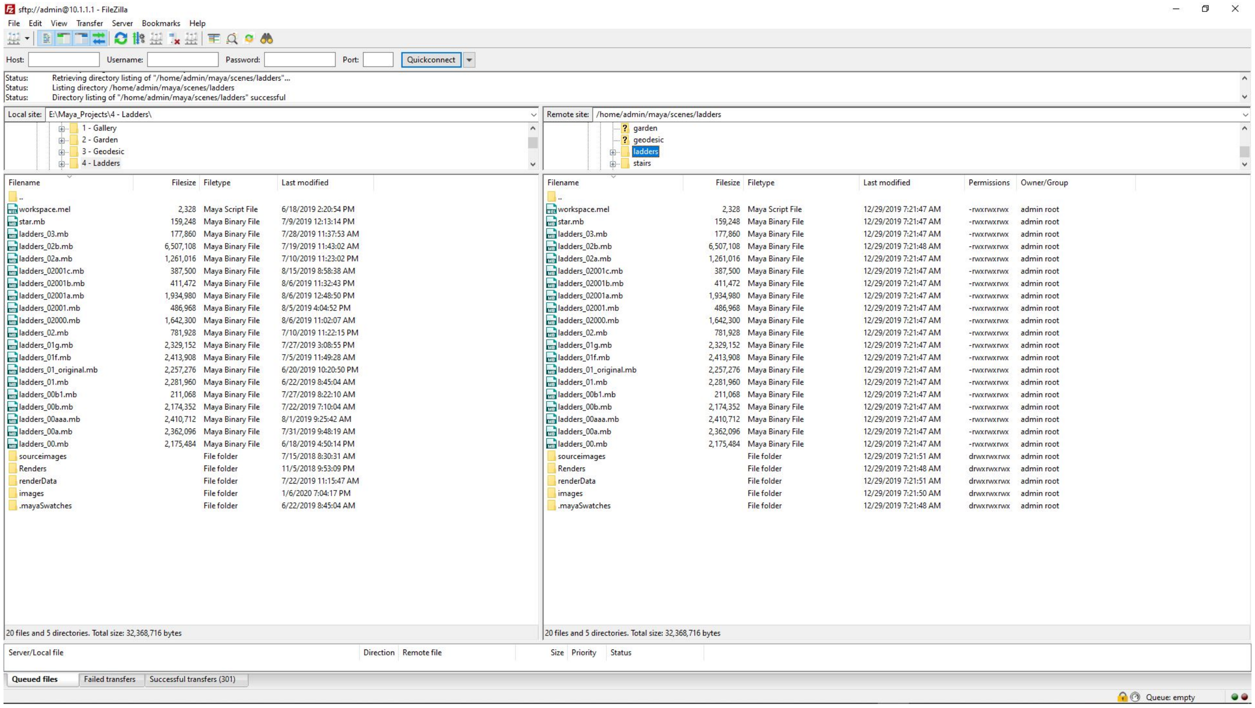

```
Status:    Retrieving directory listing of "/home/admin/maya/scenes/ladders"...
Status:    Listing directory /home/admin/maya/scenes/ladders
Status:    Directory listing of "/home/admin/maya/scenes/ladders" successful
```

Local site: E:\Maya_Projects\4 - Ladders\

1 - Gallery
2 - Garden
3 - Geodesic
4 - Ladders

Filename	Filesize	Filetype	Last modified
..			
workspace.mel	2,328	Maya Script File	6/18/2019 2:20:54 PM
star.mb	159,248	Maya Binary File	7/9/2019 12:13:14 PM
ladders_03.mb	177,860	Maya Binary File	7/28/2019 11:37:53 AM
ladders_02b.mb	6,507,108	Maya Binary File	7/19/2019 11:43:02 AM
ladders_02a.mb	1,261,016	Maya Binary File	7/10/2019 11:23:02 PM
ladders_02001c.mb	387,500	Maya Binary File	8/15/2019 8:58:38 AM
ladders_02001b.mb	411,472	Maya Binary File	8/6/2019 11:32:43 PM
ladders_02001a.mb	1,934,980	Maya Binary File	8/6/2019 12:48:50 PM
ladders_02001.mb	486,968	Maya Binary File	8/5/2019 4:04:52 PM
ladders_02000.mb	1,642,300	Maya Binary File	8/6/2019 11:02:07 AM
ladders_02.mb	781,928	Maya Binary File	7/10/2019 11:22:15 PM
ladders_01g.mb	2,329,152	Maya Binary File	7/27/2019 3:08:55 PM
ladders_01f.mb	2,413,908	Maya Binary File	7/5/2019 11:49:28 AM
ladders_01_original.mb	2,257,276	Maya Binary File	6/20/2019 10:20:50 PM
ladders_01.mb	2,281,960	Maya Binary File	6/22/2019 8:45:04 AM
ladders_00b1.mb	211,068	Maya Binary File	7/27/2019 8:22:10 AM
ladders_00b.mb	2,174,352	Maya Binary File	7/22/2019 7:10:04 AM
ladders_00aaa.mb	2,410,712	Maya Binary File	8/1/2019 9:25:42 AM
ladders_00a.mb	2,362,096	Maya Binary File	7/31/2019 9:48:19 AM
ladders_00.mb	2,175,484	Maya Binary File	6/18/2019 4:50:14 PM
sourceimages		File folder	7/15/2018 8:30:31 AM
Renders		File folder	11/5/2018 9:53:09 PM
renderData		File folder	7/22/2019 11:15:47 AM
images		File folder	1/6/2020 7:04:17 PM
.mayaSwatches		File folder	6/22/2019 8:45:04 AM

20 files and 5 directories. Total size: 32,368,716 bytes

Remote site: /home/admin/maya/scenes/ladders

garden
geodesic
ladders
stairs

Filename	Filesize	Filetype	Last modified	Permissions	Owner/Group
..					
workspace.mel	2,328	Maya Script File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
star.mb	159,248	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_03.mb	177,860	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_02b.mb	6,507,108	Maya Binary File	12/29/2019 7:21:48 AM	-rwxrwxrwx	admin root
ladders_02a.mb	1,261,016	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_02001c.mb	387,500	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_02001b.mb	411,472	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_02001a.mb	1,934,980	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_02001.mb	486,968	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_02000.mb	1,642,300	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_02.mb	781,928	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_01g.mb	2,329,152	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_01f.mb	2,413,908	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_01_original.mb	2,257,276	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_01.mb	2,281,960	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_00b1.mb	211,068	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_00b.mb	2,174,352	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_00aaa.mb	2,410,712	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_00a.mb	2,362,096	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
ladders_00.mb	2,175,484	Maya Binary File	12/29/2019 7:21:47 AM	-rwxrwxrwx	admin root
sourceimages		File folder	12/29/2019 7:21:51 AM	drwxrwxrwx	admin root
Renders		File folder	12/29/2019 7:21:48 AM	drwxrwxrwx	admin root
renderData		File folder	12/29/2019 7:21:51 AM	drwxrwxrwx	admin root
images		File folder	12/29/2019 7:21:50 AM	drwxrwxrwx	admin root
.mayaSwatches		File folder	12/29/2019 7:21:48 AM	drwxrwxrwx	admin root

20 files and 5 directories. Total size: 32,368,716 bytes

Figure 64 – Example Screenshot from FileZilla SFTP Transfer GUI

3.10 Stacking and Low-Res Test Render

The next step in the process is to stack the rendered left- and right-eye images, render them as a png image sequence, import the sequence into Premier, and render a movie viewable on the HMD for final testing.

To begin the test render, images must be stacked using After Effects (AE). In AE import the left-eye images as a sequence being sure to name them so it is clear they are the left-eye images. It is critically important that the left- and the right-eye images do not get switched. If this happens, they will look horribly wrong when viewing and will require back tracking through the process to find the place the error occurred to fix it. The author ALWAYS begins with the left-eye images. This kind of consistency tends to reduce the margin of error. Next the image sequences are dragged into the "composition" window and a new composition is created. The composition settings should be set to 7000 x 7000 pixels. Open the attributes of each image sequence in the composition window and in the "Transform" window change the scale of each from 100%, 100% to 100%, 200%. This will stretch the renders into the proper aspect ratio. They were rendered "squashed" in order to reduce render time. Then, in the left image transform window, change the "position" settings from

3500, 3500 to 3500, 1750. In the right image transform window, change the position settings to 3500, 5250. This will properly position and stack the images into a single 7000 x 7000-pixel frame. See Fig. 65 and Fig. 66 below. Next select the left- and right-eye sequences and add them to the render queue. Then render the stacked images as a png image sequence into the desired directory using the desired name. See Fig. 67 below.

At this point in the process it may be desirable to render the image sequence as a movie sequence to verify the render's viability. To accomplish this, the sequence is imported into Premier where it is exported as an h.265 file, which can then be viewed using an HMD or the Cyclorama. The process involved in this step is discussed in the next section.

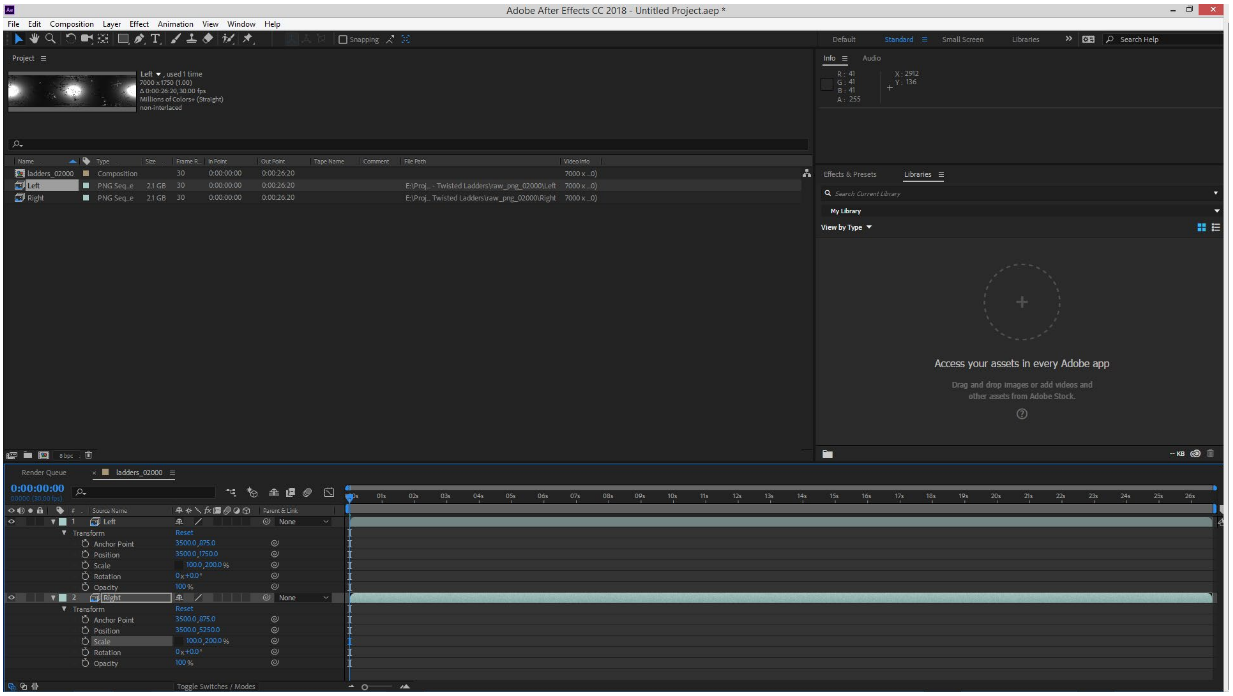

Left, used 1 time
7000 x 1750 (1.00)
Δ 0:00:26:20, 30.00 fps
Millions of Colors+ (Straight)
non-interlaced

Name	Type	Size	Frame R...	In Point	Out Point	Tape Name	Comment	File Path	Video Info
ladders_02000	Composition		30	0:00:00:00	0:00:26:20				7000 x ...0)
Left	PNG Seq...e	2.1 GB	30	0:00:00:00	0:00:26:20			E:\Proj... - Twisted Ladders\raw_png_02000\Left	7000 x ...0)
Right	PNG Seq...e	2.1 GB	30	0:00:00:00	0:00:26:20			E:\Proj... Twisted Ladders\raw_png_02000\Right	7000 x ...0)

0:00:00:00
00000 (30.00 fps)

#	Source Name	Parent & Link
1	Left	None
	Transform	Reset
	Anchor Point	3500.0,875.0
	Position	3500.0,1750.0
	Scale	100.0,200.0 %
	Rotation	0x+0.0°
	Opacity	100 %
2	Right	None
	Transform	Reset
	Anchor Point	3500.0,875.0
	Position	3500.0,5250.0
	Scale	100.0,200.0 %
	Rotation	0x+0.0°
	Opacity	100 %

Figure 65 – Stacking Left- and Right-Eye Images in After Effects

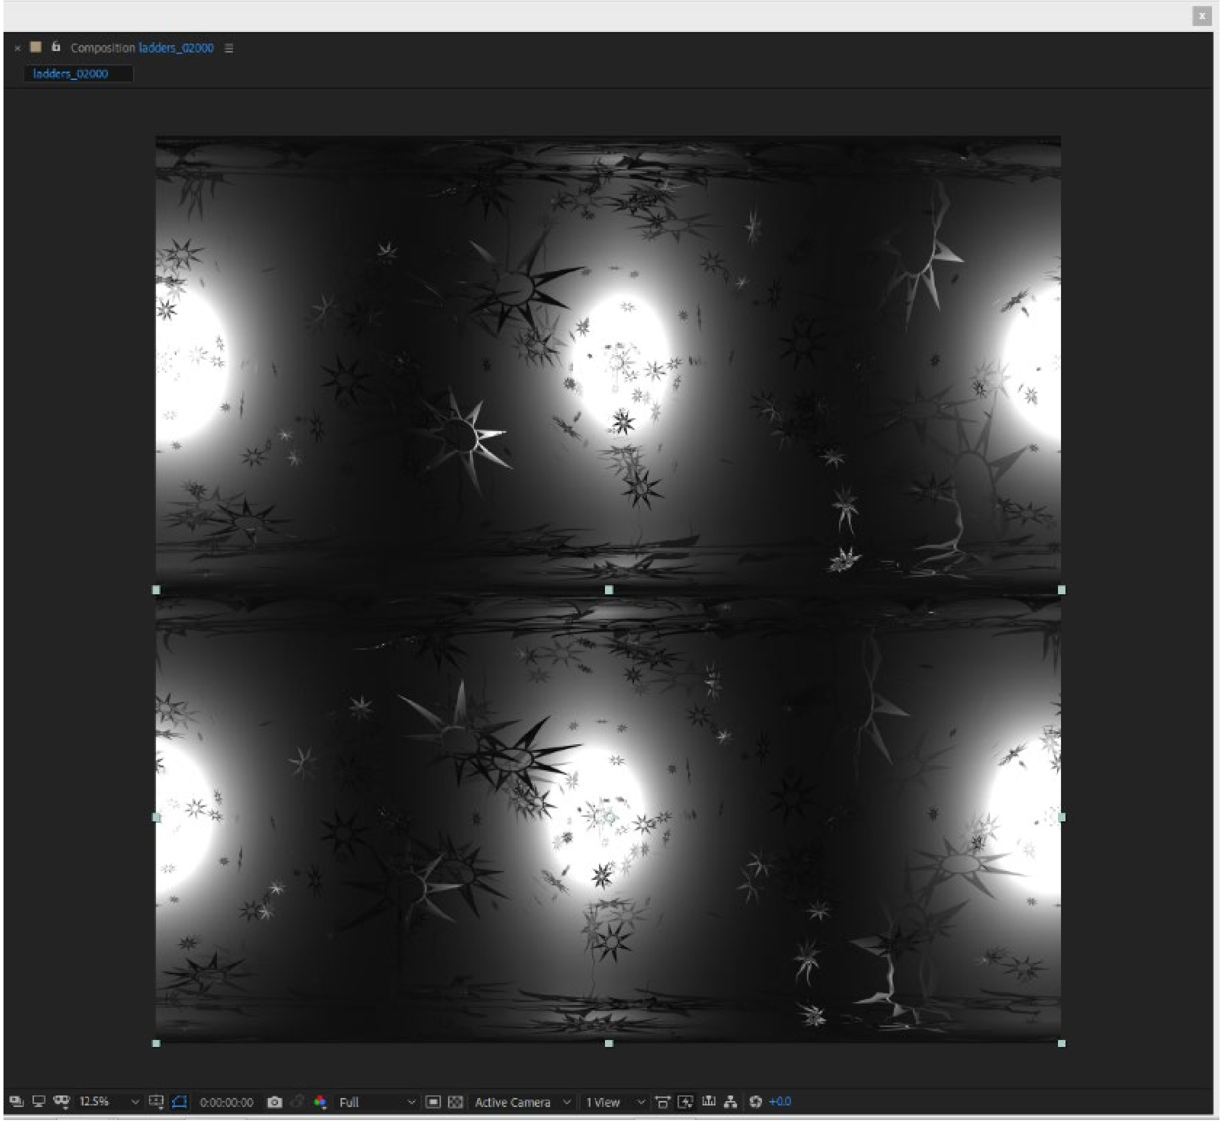

Figure 66 – Correct Image Stacking in After Effects

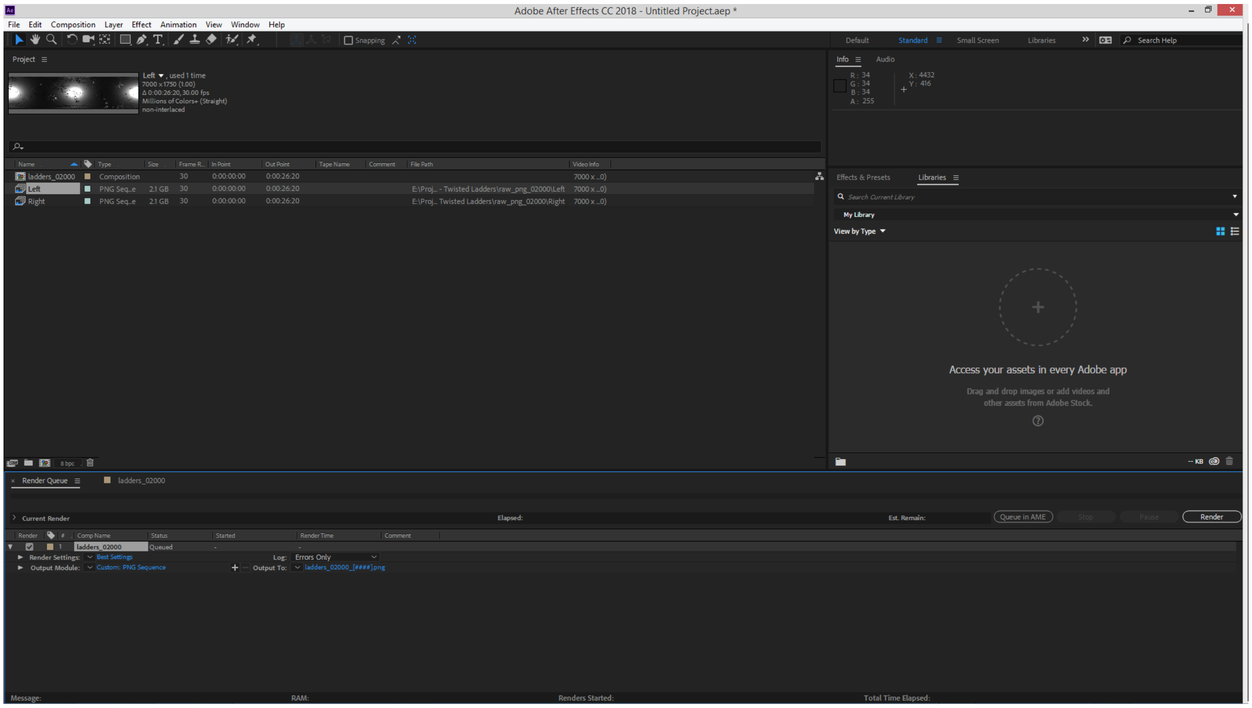

Left, used 1 time
7000 x 1750 (1.00)
Δ 0:00:26:20, 30.00 fps
Millions of Colors+ (Straight)
non-interlaced

Name	Type	Size	Frame R...	In Point	Out Point	Tape Name	Comment	File Path	Video Info
ladders_02000	Composition		30	0:00:00:00	0:00:26:20				7000 x ...0)
Left	PNG Seq...e	2.1 GB	30	0:00:00:00	0:00:26:20			E:\Proj... - Twisted Ladders\raw_png_02000\Left	7000 x ...0)
Right	PNG Seq...e	2.1 GB	30	0:00:00:00	0:00:26:20			E:\Proj... Twisted Ladders\raw_png_02000\Right	7000 x ...0)

Render	#	Comp Name	Status	Started	Render Time	Comment
☑	1	ladders_02000	Queued	-	-	

Render Settings: Best Settings
Log: Errors Only
Output Module: Custom: PNG Sequence
Output To: ladders_02000_[####].png

Figure 67 – Naming and Rendering the Stacked Images in After Effects

3.11 Editing

After the Maya rendering is complete and the rendered images have been stacked and saved as a png image sequence it is time to import it into Premier for editing. This step is where numerous image sequences, rendered over the course of several months, are edited together into a cohesive movie that can be viewed on a HMD, the Cyclorama, or on any other similar system capable of 3D/360 stereoscopic projection. In many ways this step is the most enjoyable. The culmination of many weeks or months of creating the Maya scenes and then waiting for the rendering to complete is nearly finished and it is time to begin the process of assembling the parts into the final creative expression. While this step in the process is in progress it may be necessary to re-render a given section or perhaps to render a variation of one of the sequences. Nonetheless, this is where creativity is fully engaged and all the hard work comes to fruition.

The first step is to create a low-res left-eye version of each sequence. The reason for this is because, even with a professional quality video adapter in the workstation, current computers are not capable of playing back a 7000k image sequence in real time without stuttering. Also, editing would be difficult when viewing stacked images as opposed to viewing a single image sequence.

Therefore, it is standard practice to import the left-eye image sequence from each render separately into Premier and export it as a 1752 x 438 h.264 mp4 movie. This movie will be imported into the Premier file for the project and put on a parallel timeline with the hi-res stacked image sequence. As such it is acts as a proxy sequence. The high-resolution stacked sequences are turned off so they are not viewable during the editing process, however any editing done to the low-res sequence is also done to the hi-res sequence and, in the end, it is the edited hi-res sequences will be exported to the final movie.

In Fig. 68 can be seen the left-eye image sequence in the timeline near the bottom of the image. The export settings can be seen in the window above the timeline. Figure 69 is a screenshot of a typical editing session in Premier. Notice the list of sequences on the left side of the screen. Zoom in to see that there are two timelines for every sequence. One, the low-resolution movie, is being viewed during the editing process and the other, labeled “Video1 Stacked”, for example, is turned off as indicated by the blue circle with the line through it. Again, any editing done to the low-res movie is also done to the high-res sequence. When the editing is completed, the low-res lines will be turned off and the high-res turned on before exporting the final movie.

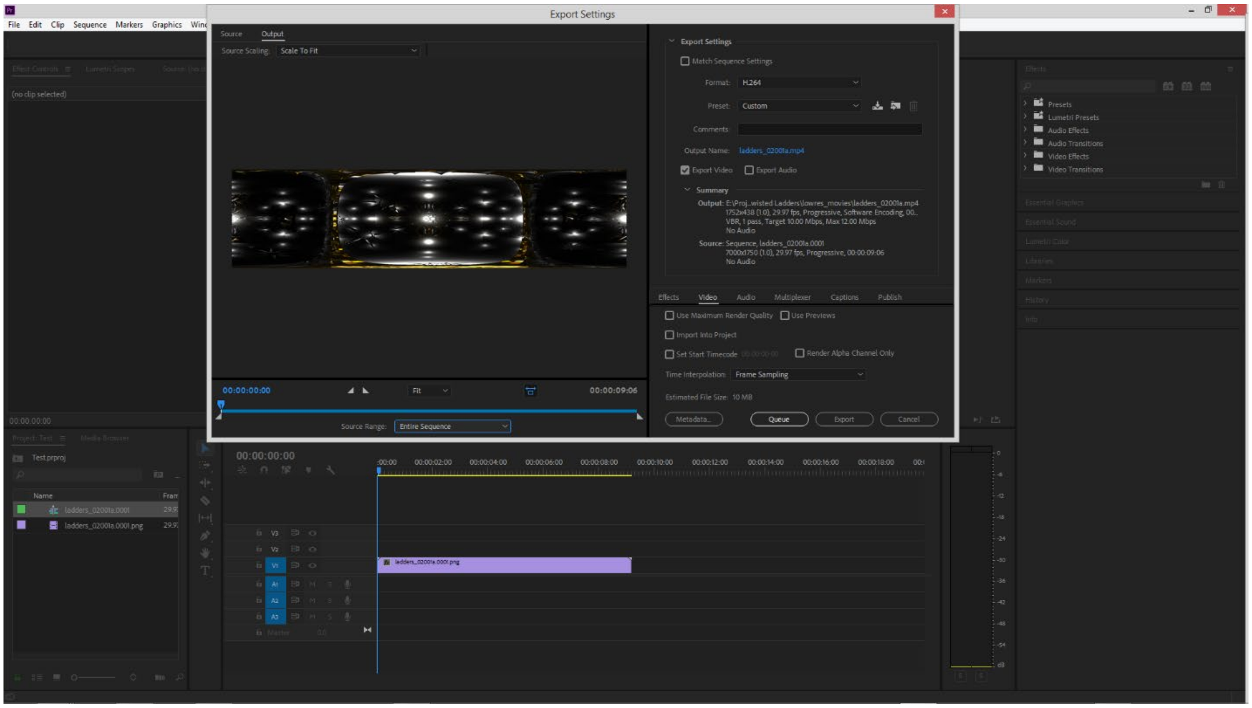

Source Scaling: Scale To Fit

Source Range: Entire Sequence

Export Settings

☐ Match Sequence Settings

Format: H.264

Preset: Custom

Comments:

Output Name: ladders_0200la.mp4

☑ Export Video ☐ Export Audio

Summary

Output: E:\Proj...wisted Ladders\lowres_movies\ladders_0200la.mp4
1752x438 (1.0), 29.97 fps, Progressive, Software Encoding, 00...
VBR, 1 pass, Target 10.00 Mbps, Max 12.00 Mbps
No Audio

Source: Sequence, ladders_0200la.0001
7000x1750 (1.0), 29.97 fps, Progressive, 00:00:09:06
No Audio

☐ Use Maximum Render Quality ☐ Use Previews

☐ Import Into Project

☐ Set Start Timecode ☐ Render Alpha Channel Only

Time Interpolation: Frame Sampling

Estimated File Size: 10 MB

Figure 68 – Exporting an Image Sequence as a Low-Resolution Movie

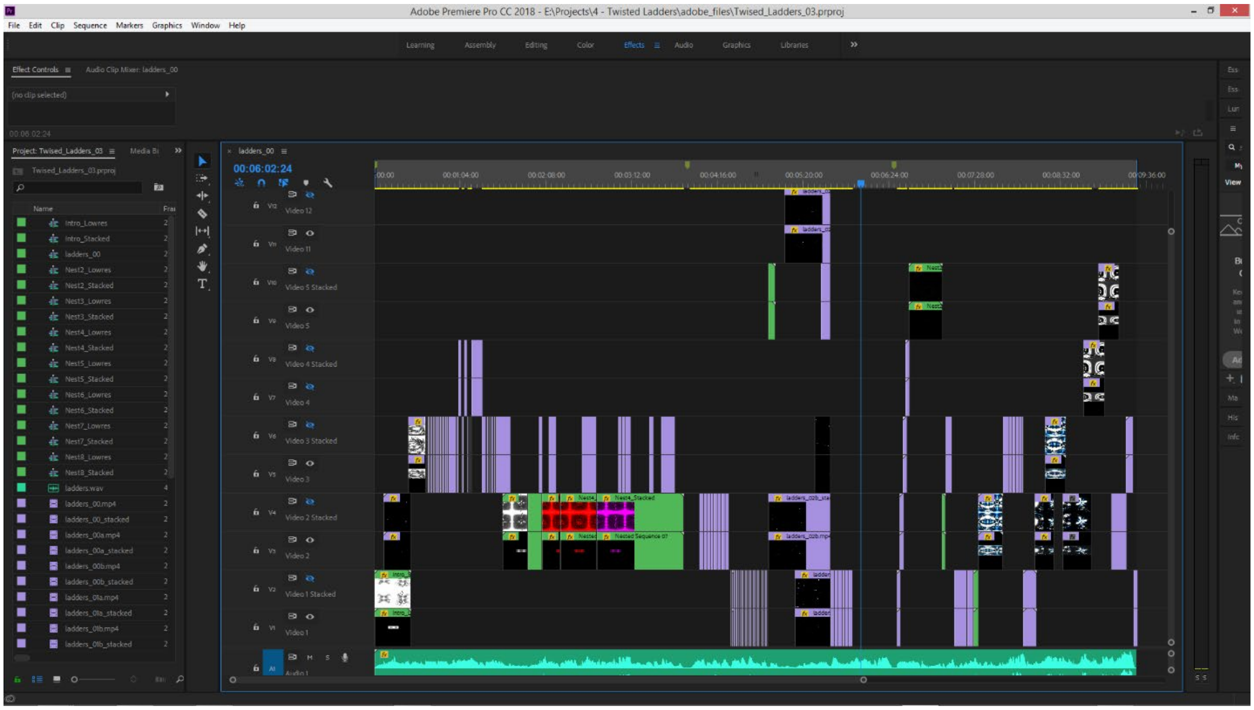

Figure 69 – Standard Premier Layout for Movie Editing

As can be seen in Fig. 69, editing involves cutting the image sequences into pieces that are fit together in a manner that is considered aesthetically desirable. The bottom most line on the timeline, colored green, is the music clip, which is a stereo mix of the 32-channel musical component of the composition. The various sections of video are aligned to the music in varying choreographic relationships, which will be discussed further in Chapter 5. Transitions between sequences are accomplished in several ways the most prominent being with fade ins and outs. Above the timeline in Fig. 69 is the editing window. This is where effects can be added to the various sections on the timeline and each effect can be keyframed. The sequence highlighted in white on this section of the timeline is the one being edited. As can be seen in the editing window, there are keyframes represented as the white triangles, which correlate to various points on the timeline. In this example, each of the keyframes contains a specific value for the opacity of the clip at that time. As in Maya, Premiere interpolates between the keyframes making smooth transitions from one value to another. Using this method fade ins and outs are accomplished. Additionally, overlapping various sequences on the timeline can be accomplished by reducing the opacity of the sequences above the bottom sequence. For instance, the sequence highlighted in white in Fig. 70 is set to a maximum opacity of 40% while the bottom sequence is set to 100%. This is just one example of the creative control that is availed the composer at this point in the process.

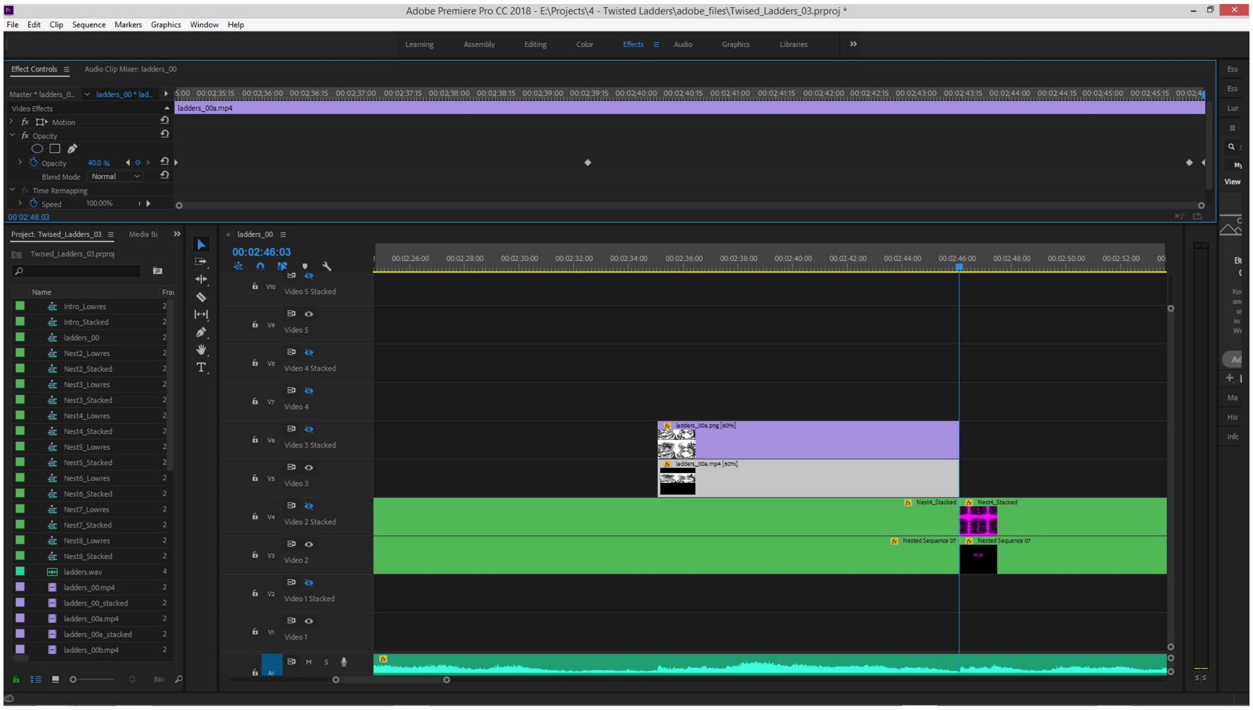

Opacity 40.0 %
Blend Mode Normal
Speed 100.00%

Figure 70 – Fade Ins and Outs in Premier by Keyframing Opacity

It should be mentioned here that the music and the video interact during this process in a feedback loop. This means that the video might come first and then the music is aligned to it. Or, perhaps the music is present during the first edit of the video and then the video is used to edit the music. This back and forth process creates a feedback loop between the music and the video, each informing and influencing the other. The process will be described further in Chapter 5.

3.12 Edit Test Render and the Oculus Rift

After sequencing the various sections of the video into a cohesive and aesthetically desirable whole, it is exported as a video for viewing in 3D/360. Up until this point, except when testing on the HMD, it has been viewed only as a 2D movie. This is the point in the process when it can finally be experienced as originally intended… as a binocular holographic 3D/360 composition. To do this, it must first be exported as an 4k, h.265 movie. Currently available head-mounted displays and the Cyclorama are each capable of 4096 x 4096, playback resolutions. However, since the images are rendered at a resolution of 7000 x 7000 originally, any resolution up to that will work with minimal degradation of image quality. So far, in test renders, the author has not found a computer capable of playing back an h.265 movie, which is the most desirable codec, at a resolution higher than 4k.

To view the first final edit draft, it must first be exported. To do this we must make sure that the low-res movies in the timeline are deactivated and the high-res sequences are activated. The sequence settings also need to be verified. They should be set at 7000 x 7000. Then the "Export Media" window is instantiated, see Fig. 71. On the bottom left side, we select "Entire Sequence" as the Source Range. Then for the Format, we select AfterCodecs, which is a 3rd party plugin for Premier and After Effects that currently allows for resolutions at and above 4k for the h.265 format. The Adobe renderer is not currently capable of this resolution and AfterCodecs is a highly optimized and efficient rendering engine. AfterCodecs renders the h.265 videos much faster than Premiere's rendering engine and it offers greater flexibility in features… see Fig. 72. For instance, it is standard practice for this project to render at a bit rate of 100mbps and with the full color range. Next, we select that path and file name for the render, shown in blue. The resolution is changed from 7000 x 7000 to 4096 x 4096. Finally, the check box for "Use Maximum Render Quality" is checked. The "Render" button is clicked and then we wait for the render to complete. Depending upon the length of the movie it can require an average of two to six hours to render.

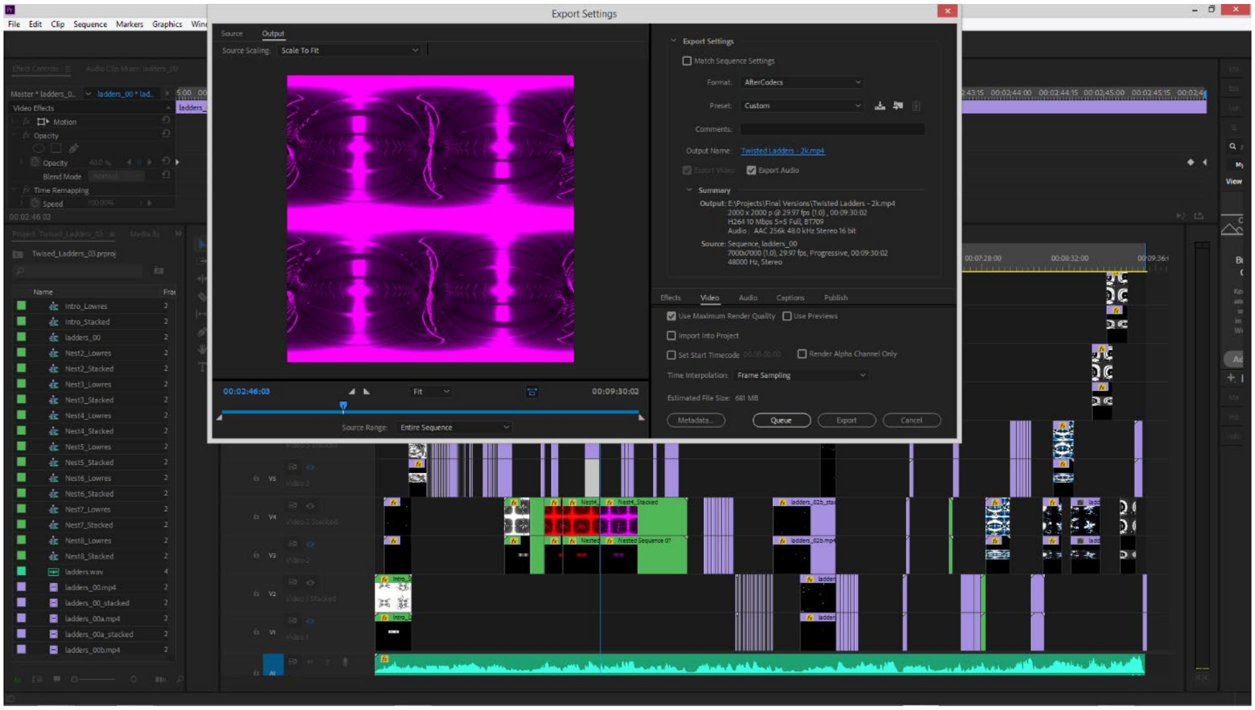

Export Settings

Export Settings

☐ Match Sequence Settings

Format: AfterCodecs

Preset: Custom

Comments:

Output Name: Twisted Ladders - 2k.mp4

☑ Export Video ☑ Export Audio

Summary

Output: E:\Projects\Final Versions\Twisted Ladders - 2k.mp4
2000 x 2000 p @ 29.97 fps (1.0) , 00:09:30:02
H264 10 Mbps S=5 Full, BT709
Audio : AAC 256k 48.0 kHz Stereo 16 bit

Source: Sequence, ladders_00
7000x7000 (1.0), 29.97 fps, Progressive, 00:09:30:02
48000 Hz, Stereo

☑ Use Maximum Render Quality ☐ Use Previews

☐ Import Into Project

☐ Set Start Timecode 00:00:00:00 ☐ Render Alpha Channel Only

Time Interpolation: Frame Sampling

Estimated File Size: 681 MB

Source Scaling: Scale To Fit

Source Range: Entire Sequence

Figure 71 – Export Settings in Premiere

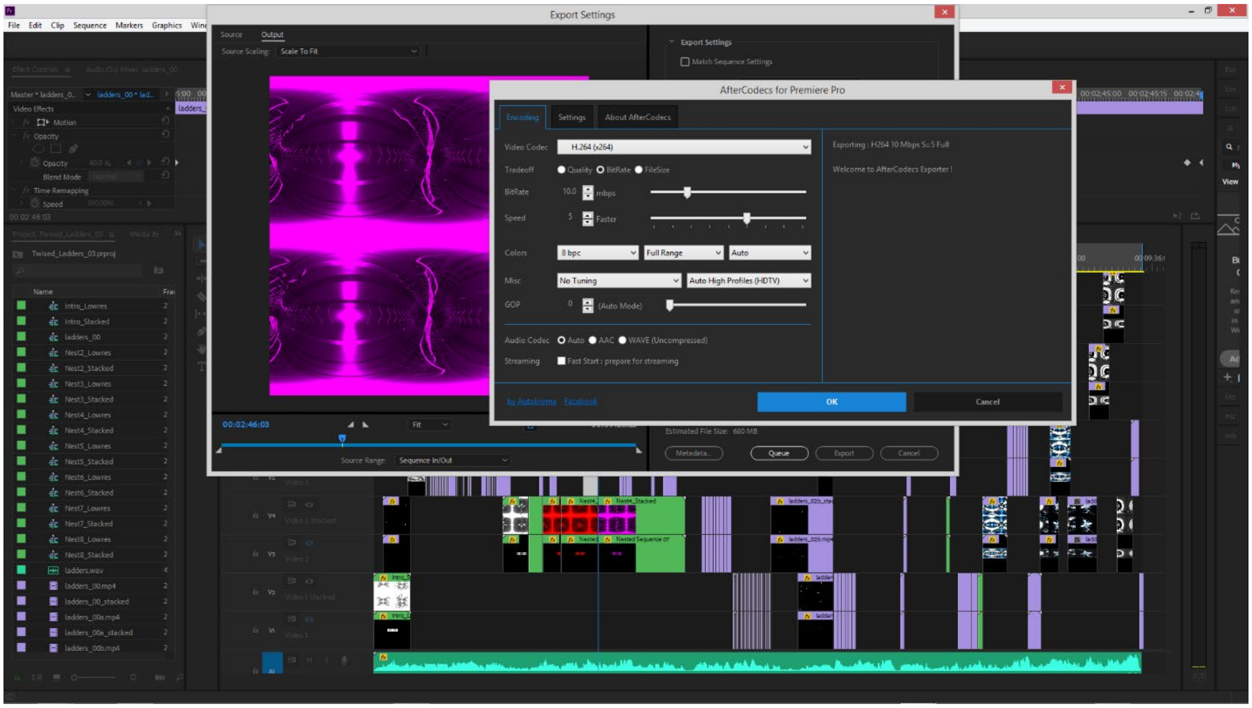

AfterCodecs for Premiere Pro

Video Codec: H.264 (x264)

Tradeoff: ● Quality ○ BitRate ● FileSize

BitRate: 10.0 mbps

Speed: 5 Faster

Colors: 8 bpc | Full Range | Auto

Misc: No Tuning | Auto High Profiles (HDTV)

GOP: 0 (Auto Mode)

Audio Codec: ○ Auto ● AAC ● WAVE (Uncompressed)

Streaming: ☐ Fast Start : prepare for streaming

Exporting : H264 10 Mbps S=5 Full

Welcome to AfterCodecs Exporter !

Estimated File Size: 680 MB

Source Range: Sequence In/Out

Figure 72 – AfterCodecs Settings in Premiere

To view the exported video using the Oculus Rift, for instance, the Oculus program must be instantiated. The program consists of configuration procedures and driver instantiation. Most conveniently, it also acts as a marketing page for promoting the games Oculus is selling. After plugging in the Oculus' HDMI and the USB3 cables to the host workstation, we bring up the Oculus application, Fig. 73.

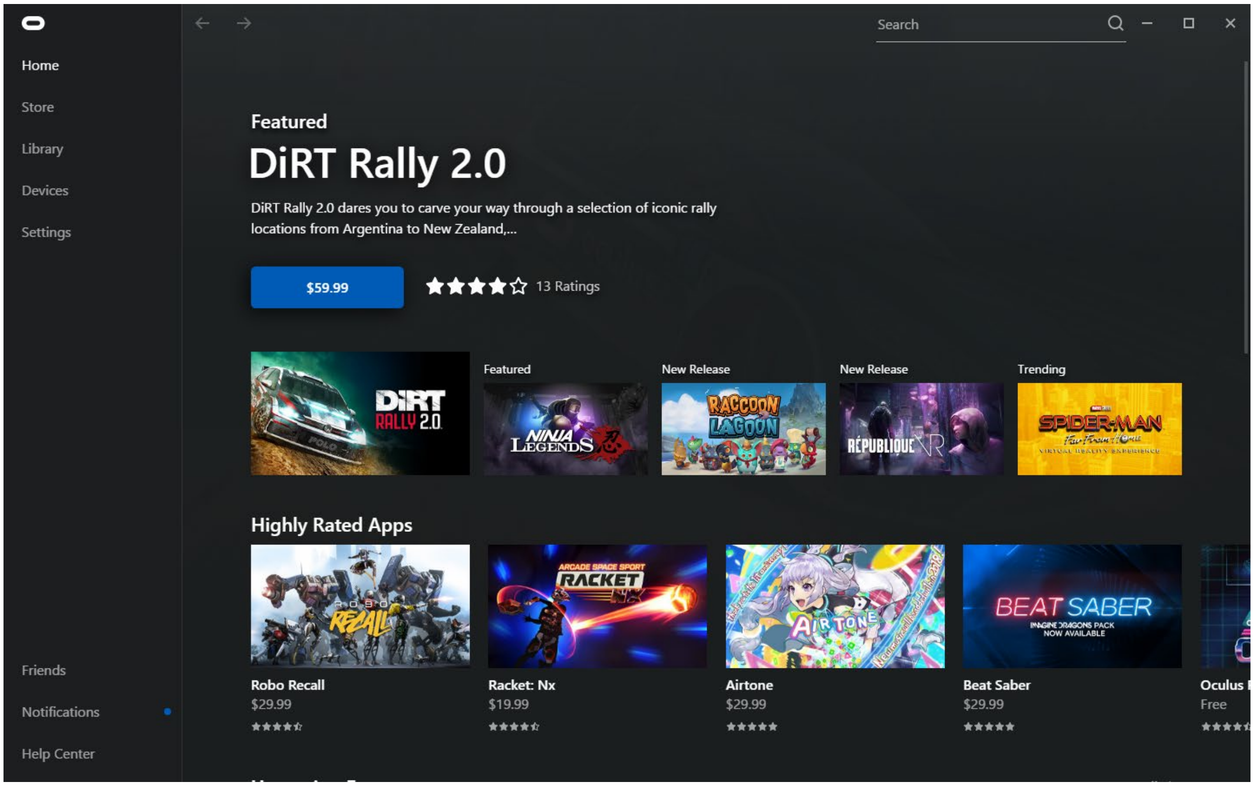

DiRT Rally 2.0 dares you to carve your way through a selection of iconic rally locations from Argentina to New Zealand,...

Figure 73 – Oculus Application

Finally, the GoPro VR Player application is instantiated. Fig. 74, Fig.75, and Fig. 76 demonstrate the pertinent configuration settings. After selecting the movie and then putting on the Oculus and moving it up and down until settling upon a position that allows for clear vision, we hit the keyboard spacebar and view the holographic video on the HMD and in 2D on the workstation monitor as seen in Fig. 77.

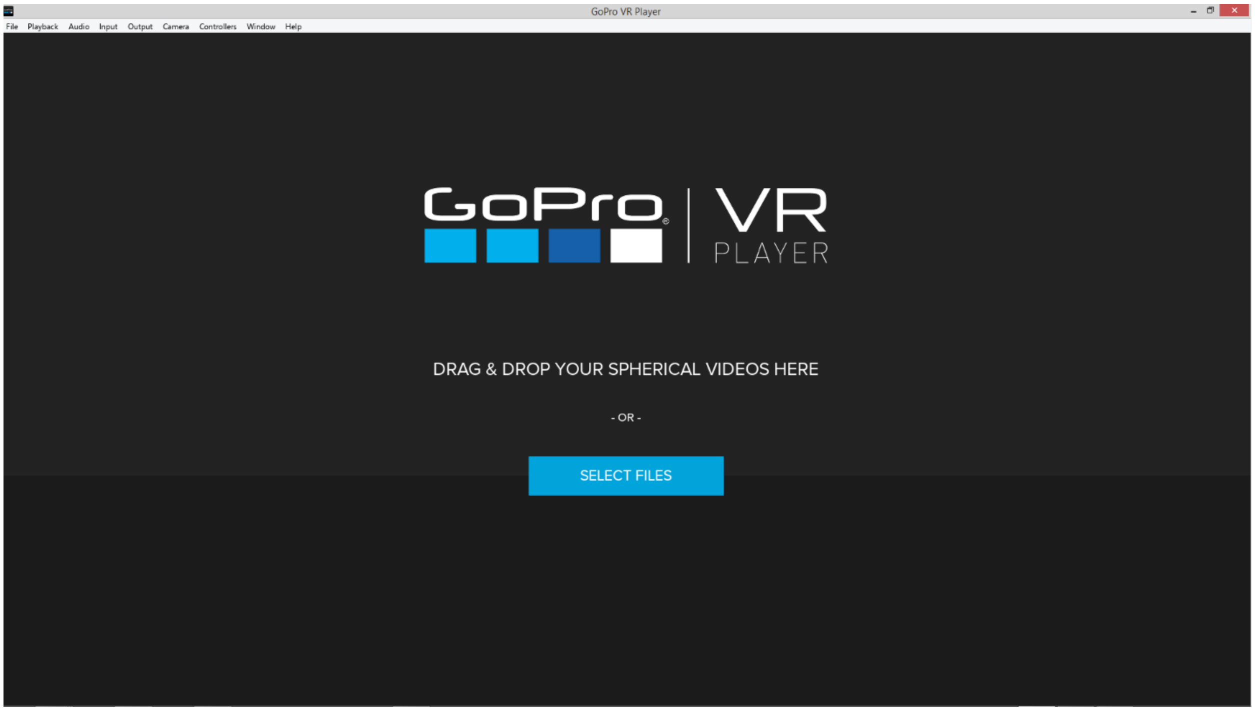

Figure 74 – GoPro VR Player Application

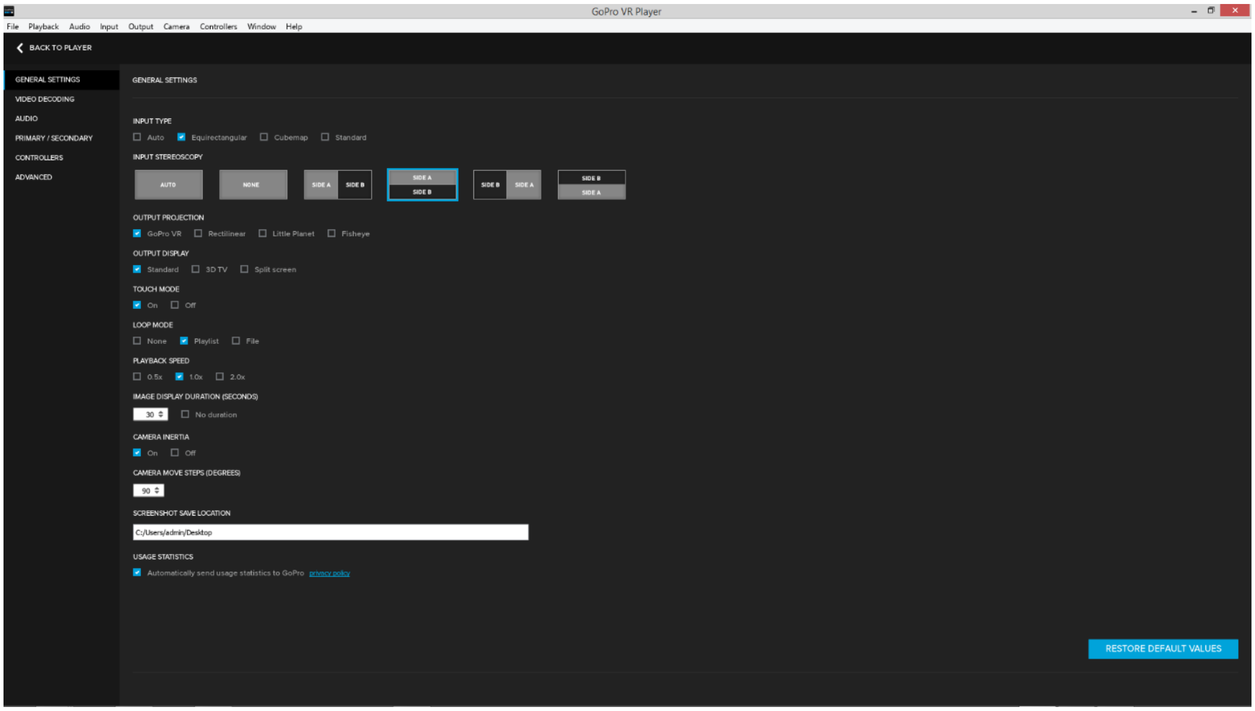

GENERAL SETTINGS

INPUT TYPE
☐ Auto ☑ Equirectangular ☐ Cubemap ☐ Standard

INPUT STEREOSCOPY
AUTO | NONE | SIDE A SIDE B | [✓ SIDE A / SIDE B] | SIDE B SIDE A | SIDE B / SIDE A

OUTPUT PROJECTION
☑ GoPro VR ☐ Rectilinear ☐ Little Planet ☐ Fisheye

OUTPUT DISPLAY
☑ Standard ☐ 3D TV ☐ Split screen

TOUCH MODE
☑ On ☐ Off

LOOP MODE
☐ None ☑ Playlist ☐ File

PLAYBACK SPEED
☐ 0.5x ☑ 1.0x ☐ 2.0x

IMAGE DISPLAY DURATION (SECONDS)
30 ☐ No duration

CAMERA INERTIA
☑ On ☐ Off

CAMERA MOVE STEPS (DEGREES)
90

SCREENSHOT SAVE LOCATION
C:/Users/admin/Desktop

USAGE STATISTICS
☑ Automatically send usage statistics to GoPro privacy policy

Figure 75 – GoPro VR Player General Settings Window

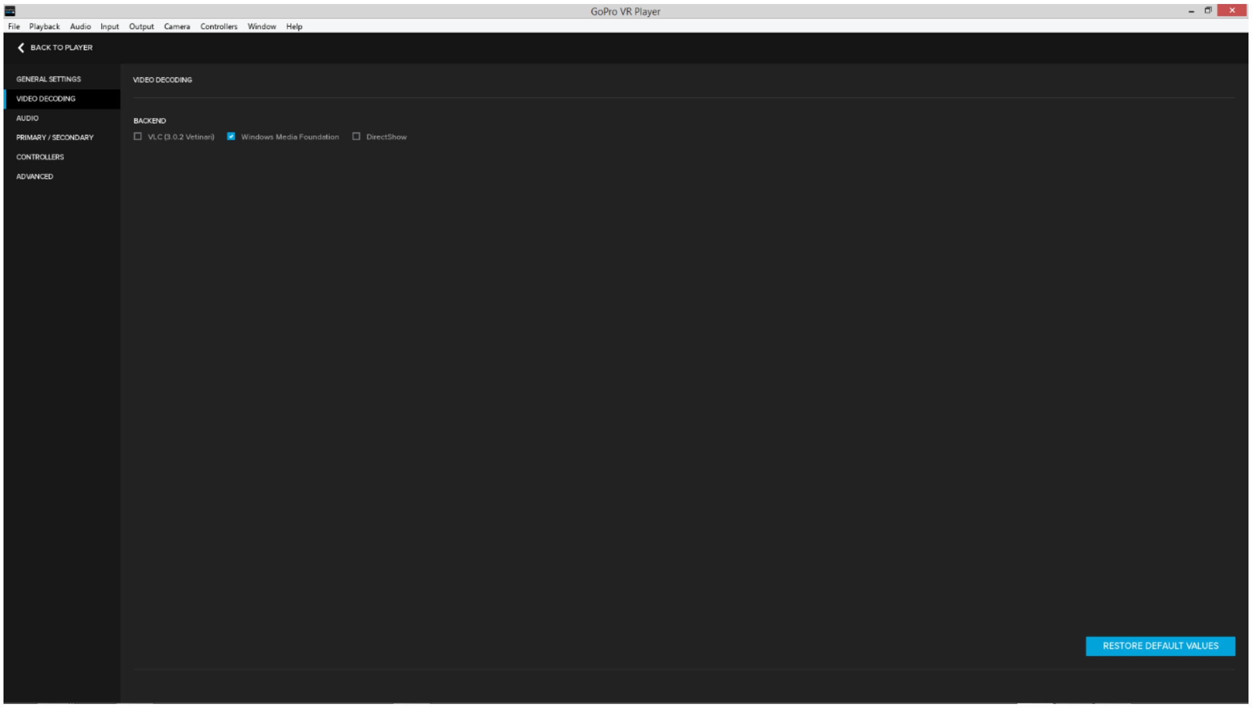

VIDEO DECODING

BACKEND

☐ VLC (3.0.2 Vetinari) ☑ Windows Media Foundation ☐ DirectShow

Figure 76 – GoPro VR Player Video Decoding Window

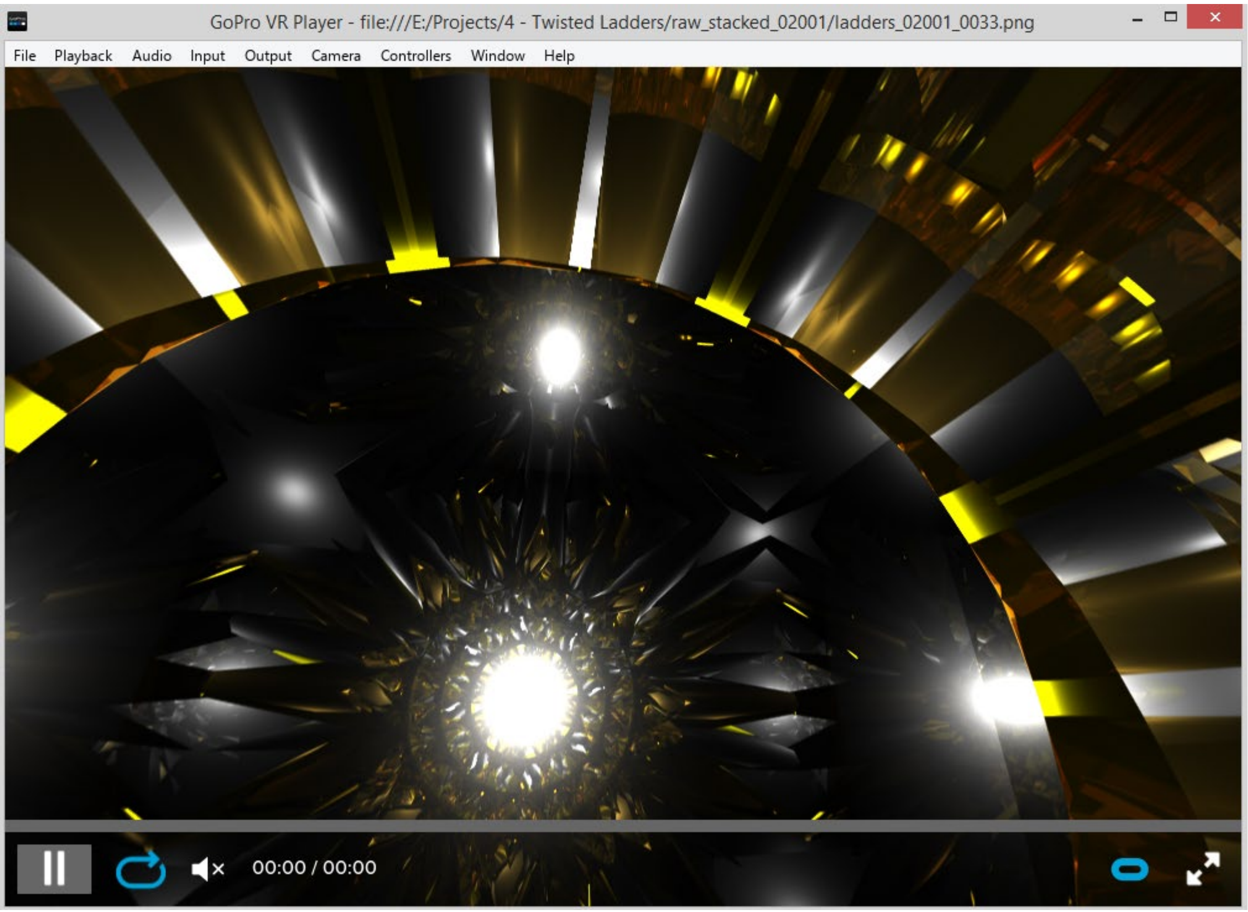

Figure 77– GoPro VR Player Viewing Video

3.13 Final Editing

After viewing various drafts and refining them until satisfied, the composition is completed with a final render in the desired resolution and format. Here there may be a need to re-render a sequence or an inspiration to make variations of some of the existing sequences may occur. The need to make some completely new scenes may arise, which will be rendered to added to the mix. One manner in which to create variations of the original sequences is to bring them back into After Effects, add in effects as desired, and then export as a new png sequence. There are no rules. Often following whatever minute impulse occurs can lead to exciting new area of exploration.

3.14 Creating an Edit Render in After Effects

I should state that many results of editing video or still images in AE have become cliché and, of course, it is not these types of effects that interest me or are pertinent to my work. On the other hand, as will be demonstrated, some of the most basic editing techniques using AE, when coupled with images that are already quite interesting on their own, can be a method for creating thematic variations that powerfully compliment and extend the original renders.

At first thought it may seem like a simple task to bring the material into AE and create variations. The intuitive approach would be to simply import the low-res movies and to import the stacked sequences and use the former as a reference to affect the latter and then export the results of each. But there is an issue with that. We want the render to occur on each left and right eye image separately and equally. If an effect is implemented on the stacked images, then it will affect the images as one image thus creating different effect on each eye. This is not acceptable. Therefore, it is necessary to import each left- and right-eye rendered sequence separately into AE along with the low-res movie. Then we apply the desired effect to the latter and then apply the same effect to each of the former images. We next export the low-res movie as another low-res movie, with a different name. Finally, the left- and the right-eye sequences, now with the effects applied, must be stacked and then exported as a stacked sequence. Though protruded, this process makes it possible to integrate AE effects into the 3D/360 stereoscopic visual composition editing workflow. It is protruded because rendering a stacked version of an eighty-second-long left eye/right eye sequence usually requires several hours to accomplish. However, the often-stunning results reveal that it is well worth the investment. Here is the work flow described systematically:

In AE:

1. Import low-res mp4 version of the sequence
2. Import left-eye high-res png sequence
3. Import right-eye high-res png sequence
4. Mute the left- and right-eye sequences
5. Apply the desired effect to the low-res mp4 version
6. Export the now effected low-res mp4 version as another (renamed) low-res mp4 version
7. Unmute the left- and right-eye high-res png sequences and apply the same effect as applied to the low-res mp4 version
8. Stack the left- and right-eye images in a 7000x7000 format

9. Export the stacked images as a png image sequence
10. Import the effected mp4 version and the effected stacked png sequence into Premier, one stacked directly above the other, and proceed with editing the composition.

3.15 Various Rendering Levels

When referring to “rendering” throughout this document, the generic term extends well beyond the lengthy rendering required to create the initial stereoscopic image sequences created in MAYA. Without being privy to the inner workings of the visual compositional side of the processes involved in this intricately interlaced compositional approach, there is little way to apprehend the extent of the term as applied here. Throughout this process, rendering is instantiated on five different levels. In order to clarify this, I will present a brief description of each here.

- Maya Rendering: Using the HPC cluster in my studio, MAYA scenes are rendered as png sequences. Each frame consists of a left eye and a right eye image. Each is 7000 x 1750 pixels in dimension.
- After Effects Rendering: Each left- and right-eye image sequence is stacked within a 7000 x 7000 image. The left-eye image is on top of the right-eye image. The left-eye image is centered at 1750 pixels from the top of the new image. The right-eye is centered 5250 pixels from the top of the new image. These stacked images are exported as png sequences.
- Premier rendering: The left-eye png sequence is imported into Premier and then rendered as an H.264 mp4 low-resolution reference movie.
- After Effects Rendering: Any sequence that has effects applied through the use of After Effects processing to create variations of the initial MAYA renders, is rendered both as a low-resolution mp4 reference movie and as a stacked png high-resolution sequence
- Premier Rendering: The final edited version of the composition is exported (rendered) as a high-res png sequence.

As can be realized, rendering throughout this process is a multi-faceted task.

3.16 Completed Composition

After the final editing is complete, after there seems nothing more to do, the composition is complete. Often at this point it is prudent to render a 2D version of the visual music composition. For one, the composition should stand on its own as a 2D video. If it does not then one is likely relying too heavily upon the gimmickry of 3D. Also, many venues are incapable of projecting a 3D video therefore it is prudent to have a 2D version on hand to extend the presentational reach. Finally, often when viewing the project in 2D, issues become obvious that may have been missed while viewing in 3D. In 3D since one cannot view the entire screen at the same time, there are often issues that are overlooked unless one can complete several full viewings. At this stage, technically speaking, the process is complete.

3.17 Holography

Now that we have established a baseline for the technical processes involved in composing 3D stereoscopic movies, how do we set our sights toward holography? All 3D visual work could be considered virtual holography. Even more so, all 3D/360 visual works can surely be considered as such. However, for this project the intention is to proceed beyond this simplistic perspective and create a novel perspective on holography. Here we are interested in placing the audience within the holochoric environment. The idea is to create an environment in which the audience is completely immersed. Further, it is considered important that the objects within this completely immersive environment share the space with the viewer. It is a goal to place the objects close enough to the viewer that they are inclined to reach out to touch them even though they are realized as virtual. At times an object should move through the observer and, perhaps through a synesthetic sense, they should actually “feel” it moving through or in very close proximity to them, perhaps with a sense of an accompanying air movement. Some audience members have reported experiencing these types of feelings, which are conditioned responses to visual or auditory stimuli. There are other goals for the compositions as well, which will be discussed later, but many of them are more deeply enabled, by this holochoric approach.

3.18 Chapter Conclusions

Within the framework of this chapter, numerous stratifications of contrapuntal process lend toward the composition of the visual aspects involved with this project. Evolutionary progressions of pseudo holographic imagery are a key factor in the creation of a multidimensional spatiotemporal experience. Placing an audience within a visual world that is highly stratified, both on a micro and macro level, and one that could not exist within the day-to-day visual experience, facilitates an environment unique to this endeavor. Add to this the aleatoric and indeterminant procedures, the actively engaged interplay that functions within the quadrilateral creative feedback loop that produces it, and the greatly expanded tone row paradigm, one can envisage an evolutionary contrapuntal leap in the field of visual music.

CHAPTER 4. CLUSTERING PROCESS

4.1 Chapter Overview

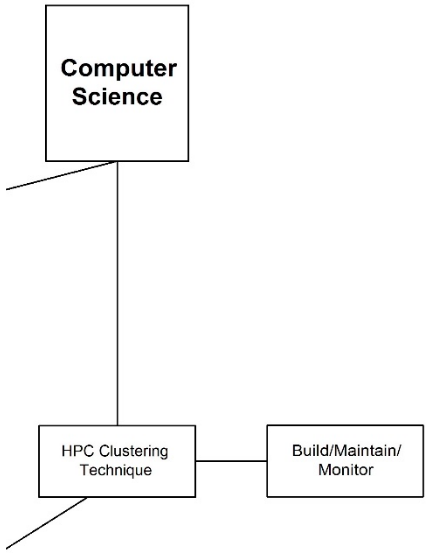

Figure 78 – HPC Integration into the Creative Process

Producing the high-resolution stereoscopic (pseudo) holographic 3D/360-degree image sequences intended to be projected in parallel with the diffusion of the holophonic participants in this project would be extremely challenging without integrating an HPC paradigm into the quadrilateral reciprocative feedback loop. Without HPC, rendering visual work that enables multidimensional spatiotemporal experiences would take years instead of weeks or months. The HPC rendering methodologies developed and employed here are intricately interwoven into the audio, the visual, and the composer's contribution to this project.

Because the rendering process is nonetheless lengthy, with regard to the times required to complete rendering cycles, the HPC cluster often acts as a constraint to the extent to which the visual material may be developed. This factor necessitates the implementation of the extended tone row approach discussed in the Introduction. Fortunately, constraints are often the basis for the composer's active engagement with creativity (Stravinsky). Therefore, intricately interweaving the rendering methodologies developed here into the overall compositional process is not only possible but necessary and contributes dynamically toward the evolutionary progression of counterpoint in the field of visual music.

The Maya rendering aspect of creating this extended reality (XR) visual work is critical to its success. It is reasonable to suggest that the level of computational power accessible to the artist and his or her ability to implement it to advantage belies the extent to which creativity is empowered. Though it is often considered outside the realm of the artist's world to pursue knowledge and experience in the area of computer science, it is herein suggested here that it is not only desirable but also necessary. Artists often off load such tasks to others; however, rendering can be considered a part of the creative process. Many fortuitous creative discoveries are made during each step of the creative process and if the artist is not present, they are missed opportunities. Today most every artist working within a digital medium must be cognizant of the rendering process. In order to render the high resolution and complex visual images for this project, high-performance computing was a necessity. The project would not have been possible to realize without it.

The alternative to HPC is the rendering farm. Two of the most prominent rendering farms for Maya projects are Backburner and Tractor. They have each been implemented for my work in the past. Though similar in principle to HPC, there are significant differences that affect efficacy. It was discovered that when using HPC there were significant performance gains, which can be mainly attributed to very low overhead imposed by the CentOS operating system that has been reduced to only the absolutely necessary components and to extremely efficient scheduling of rendering events using the Sun Grid Engine. Though the documentation of bench testing will be left to those involved in quantitative/qualitative empirical research, it was found by general comparison in initial testing that there was a ~300% increase in rendering speeds using the HPC approach as opposed to a typical rendering farm approach. This was a significant performance gain. Certainly, reducing rendering times is of utmost importance to the overall process for this project.

High-performance computing is a clustering process, which is why an HPC system is often referred to as a cluster. The type of cluster utilized here is specifically termed a "Beowulf Cluster", which is a cluster based upon commercial grade computers. Basically, it consists of several compute nodes that are controlled by a single computer called a head node. This paradigm creates a unified super computer from the culmination and interconnectivity of the individual computers involved. Basically, the head node sends computational commands to each of the compute nodes, which respond by carrying out the task and then report back to the head node when ready for another. Indicative of a cluster is the network configuration that places each component on the same private local area network (LAN) and a shared file system. The file system on the head node is linked and shared with each of the compute nodes such that each of the latter can write directly to the specified hard drive on the head node. This is an extremely efficient manner with which to accumulate the results of each computational task from each compute node into a single directory. The head node is additionally placed on the public LAN that is accessible by the appropriate workstation(s) and has access to the Internet. This "dual homing" approach allows for a flexible system and yet isolates the cluster so there is no network traffic congestion that would reduce the efficiency of the cluster communications and file transfers. It also provides a layer of network security. See section 4.2, Fig. 82, for a wiring diagram of the Ethernet configuration.

In searching for an optimal HPC approach, many packages were considered. The pure Beowulf cluster was the obvious and well-known choice. However, from the documentation studied, implementing it seemed to involve a very manual and tedious process. It was desirable to find a package that was simpler to install and implement. Finally, after a protruded research effort, it was decided that the ROCKS clustering bundle was a viable choice. ROCKS is developed and maintained by the computer science department at the University of California in Dan Diego. The ROCKS package is relatively easy to install and highly customizable. The latter aspect allows one to install only the applications it offers that are pertinent to the computing tasks for which the cluster is intended. Later in this chapter the installation process is delineated in a general sense.

For this project a cluster consisting of six compute nodes was implemented. Each of the compute nodes is based upon an Intel Core-i7 processor, which consists of four dual cores. This provides a total of forty-eight processors to carry out computational tasks. ASUS Sabertooth Z87 gaming motherboards and 16GB of Crucial® DDR3-1866 gaming RAM, were utilized both for their speed and their durability. For this project, the cluster was utilized 24/7/365 at ~87% load average for over three years so durability was a critical factor in the choice of components. To the credit of these components, there was not a single hardware failure over the course of the three years of continuous use. Of course, it is realized that these components will soon be archaic given the steady advances in computational hardware, however, at the time of this work they are the most effective and viable choice especially when considering the financial constraints of a grad student working independently.

Cooling the hardware was an important factor. A fan was placed directly above the cluster rack. It evacuates the heat directly from the room into a vented attic space and pulls relatively cool air into it. This cooling arrangement is necessary not only for longevity of the components but also for the speed at which they operate. Nonetheless, a more effective cooling system would be advantageous and, as it becomes financially feasible, will be implemented. There are safeguards in place on the motherboards such that when CPU temperatures rise above a designated threshold CPU usage is reduced by deactivating processor cores. This diminishes the operating temperature of the CPU; however, it also impedes the efficiency of the rendering process. The temperature ranges of these safeguards may be altered or altogether eliminated utilizing the Basic Input/Output System (BIOS) configuration of the motherboard. However, this is not a wise choice given the computational demands placed upon the compute nodes and the necessity of their participation in the creative process. A CPU failure would seriously hamper the progress of a rendering process that requires months to complete without such complications.

Below, in Fig. 79 and Fig. 80, are photos of the cluster. Note the attic opening in the upper right-hand corner, the UPS (uninterruptable power supplies) units on the lower left, and, atop the computer chassis on the bottom of the rack, the external backup drives consisting of nearly 30TB of backup storage space. Storage, backup, and redundancy are necessities. The head node is at the top of the rack followed by six compute nodes. The workstation is at the bottom of the rack. Note also, the wireway holes behind the rack. They allow the passage of wiring between the datacenter room and the studio.

Figure 79 – HPC Clustering Rack, Workstation, Cooling, UPS Units, and Wireway

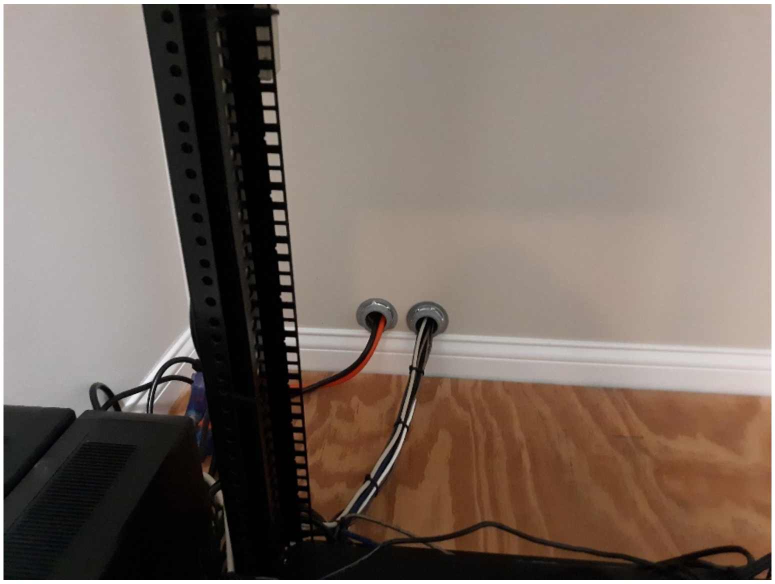

Figure 80 – Wireway Between Datacenter and Studio

4.2 Building/Configuring the Cluster

This section will provide a general overview of the process involved in building and configuring the head node, installing the compute nodes, and configuring the cluster to render Maya scenes. This process is not intended as a recipe but instead as an example of how it was approached here. Anyone interested in a similar approach will likely discover numerous variations due to the equipment used, the version of ROCKS employed, and the intended use of the cluster.

The head node is where the majority of configuration takes place. It can be thought of as a traffic director (scheduler) and as a collector of the Maya scenes and the rendered image sequence results. The hard drive should be sizeable, >= 1TB, and solid state if possible. It is where the rendered results are stored until moved (harvested) to the workstation. The machine itself, with regard to hardware, need not be as substantial as the compute nodes because it is doing relatively little work. It is the compute nodes that are the workhorses. The head node utilized for this project consists of an older dual Xeon processor on a SuperMicro® server motherboard, with 8GB of DDR2 RAM. Again, there is little demand for memory on the head node. It was housed within a rackmount 4U chassis to provide plenty of cooling space. The head node does not create much heat, however since the environment it is in can be quite hot it is prudent to provide plenty of cooling. Though the head node is not the workhorse, if it were to fail, the entire cluster would be lost and the process of re-building it would be lengthy. Therefore, it is important that it be robust if not fast.

Next the compute nodes are assembled. Below, the basic parts are listed that were utilized for each of them. Though it is not a necessity, it was considered important to utilize matching components so that the renders will all be operating on the same timeline. This makes it easier to calculate render times and to anticipate that all renders will complete at approximately the same time. This also simplifies watching load levels and monitoring performance graphs, while a render is being executed, since each compute node is operating at the same level.

- Motherboard: AUSU Sabertooth Z87 LGA 1150 Intel Z87 SATA 6
- Processor: Intel Core-i7 4770 Haswell Quad Core (8 virtual) 3.4 GHz LGA 11
- Noctua NH-L12S CPU Cooler w/Quiet 120mm PWM Fan
- RAM Memory: 16 GB Crucial Ballistix Tactical DDR3-PC3-12800
- Hard Drive: Crucial MX300 2.5” 1TB SATA Internal Solid-State Drive (SSD)
- Power Supply: Antec EarthWatts Green EA380D 380 watts Continuous Power Intel Haswell Compatible
- Chassis: Rosewill RSV-L4000 4U Rackmount Case / Chassis 8 Internal Bays

As previously noted, these parts were specifically chosen for speed and durability. For instance, the black capacitors on the motherboard have a tighter tolerance than those typical of lesser motherboards and are rated at a much higher temperature. Both features lend toward speed and durability. The RAM memory has special cooling fins intended to dissipate heat faster and more efficiently. The larger than necessary 4U rackmount chassis allow for a high volume of air flow through them and include four 5” ball bearing fans also increasing the cooling power. Figure 81 below provides a view to the inside of the chassis.

After the hardware is configured it is prudent to check the BIOS of each computer to ensure it is up-to-date. It was determined that the most recent version of the BIOS was not installed on the motherboards for this project. This caused an issue with the PXE boot, which is required when adding a compute node to the head node’s database. PXE boot simply refers to a server obtaining its operating system from the network instead of from its hard drive. More information on how this is pertinent later in this section. The BIOS is the basic operating system for the motherboard and one has to boot the computer into the BIOS configuration on startup to change the settings and to update the BIOS. Each manufacturer has a unique manner to accomplish this. For the ASUS Sabertooth motherboard one simply holds down the F2 key during the initial boot sequence to enter the BIOS. After the BIOS has been updated, any other tweaking to the configuration can be accomplished before proceeding. For instance, while in the BIOS configuration, it is a good idea to set the boot order to include attempting a PXE boot first before booting from the hard drive.

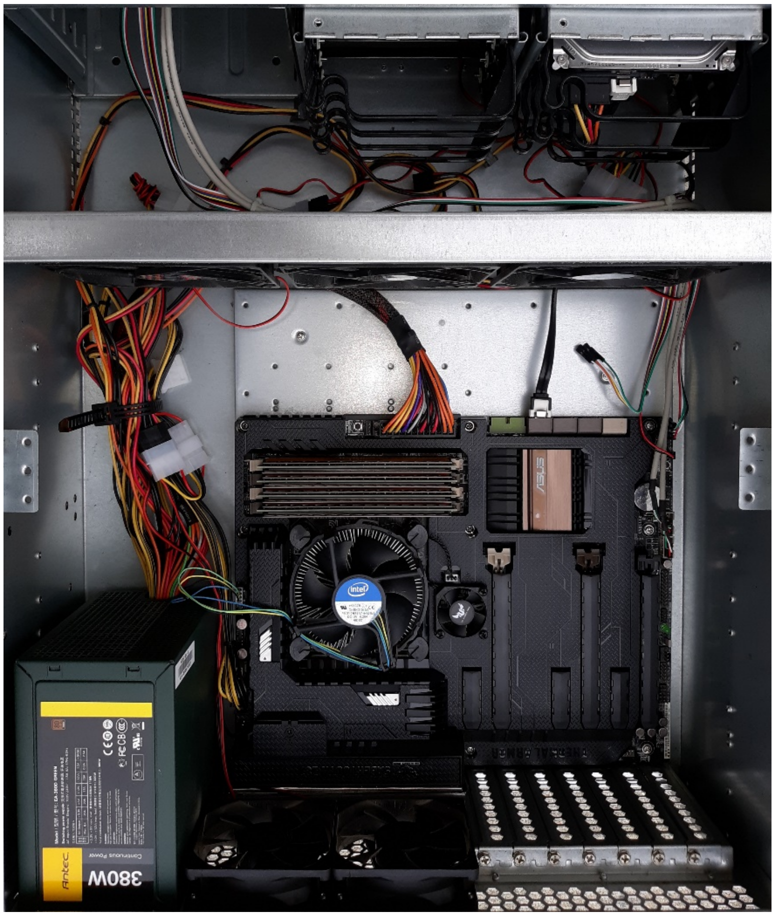

Figure 81 – View of the Inside of a Compute Node Chassis

After the hardware for the head and compute nodes is setup and ready to go, the next step is to determine and configure the wiring for the network topology. In Fig. 82 below, it can be observed that there are two networks present… the blue lines represent the public LAN that connects all computers, phones, printers, and other devices on the network to the internet and, depending upon sharing determinations, to each other. The red lines represent the private cluster LAN. It should be noted that the head node is connected to both thus utilizing two network interface cards. As previously stated, this allows the cluster to operate on a clean network that is free of any network traffic that might slow communication and file transfers to and from the head node and the compute nodes. As demonstrated in the Fig. 82, the public network in this case implements a class C TCP IP range of 192.168.1.1 – 192.168.1.255 and a netmask of 255.255.255.0. Similarly, the cluster private LAN is set up to utilize the class C TCP IP range of 10.10.10.1 – 10.10.10.255 with a netmask of 25.255.255.0. When configuring the network topology in ROCKS, these ranges need to be stipulated. In each case a class C network allows for up to 256 IP addresses. The reason to stipulate the private cluster address range is so that the head node can hand out addresses using DHCP as new compute nodes are added to the system. Conversely, the head node requires access to the public LAN so that it can phone home to obtain updates to the rolls, the programs, utilized on setup. Further, the studio workstation, in this case, will need to upload Maya scenes to the head node for rendering and to download the rendered images for processing and integration into the pertinent composition. Both of the latter processes occur using secure file transfer protocol (SFTP). As previously stated, this can be accomplished using command line or one of the popular file transfer applications such as FileZilla. Most communication with the cluster is conducted via ssh2.

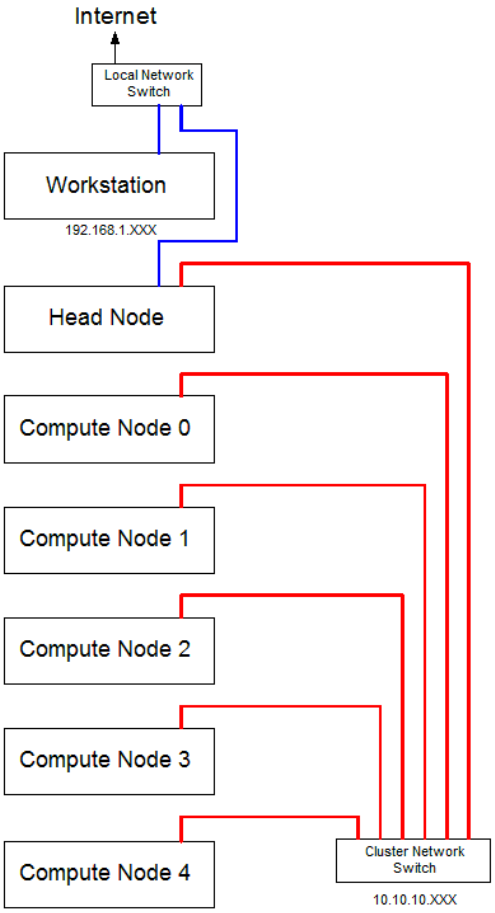

Figure 82 – Network Topology for the HPC System and Integration to the Local LAN

When the hardware is setup and ready to go and the network topology is set, the next step is to install the software on the head node. There is little reason to provide a step-by-step procedure here as the process is delineated very clearly in the ROCKS documentation, which is available in document format and on their web site. Here is the URL for the web site: http://central-7-0-x86-64.rocksclusters.org/roll-documentation/base/7.0/install-frontend.html. Figure 83 provides a photograph of the boot image for ROCKS 6.1.1, each version of which is the name of a snake. For instance, this one is called "Sand Boa". There is a newer version of ROCKS, 7.x, which is implemented on CentOS 7.x, a Linux variant, as opposed to CentOS 6.3. Since 6.3 is working with the hardware being used for this project there is little reason to update and thus take time away from rendering. At such a time as the head node is updated or replaced then version 7.x would be installed and configured.

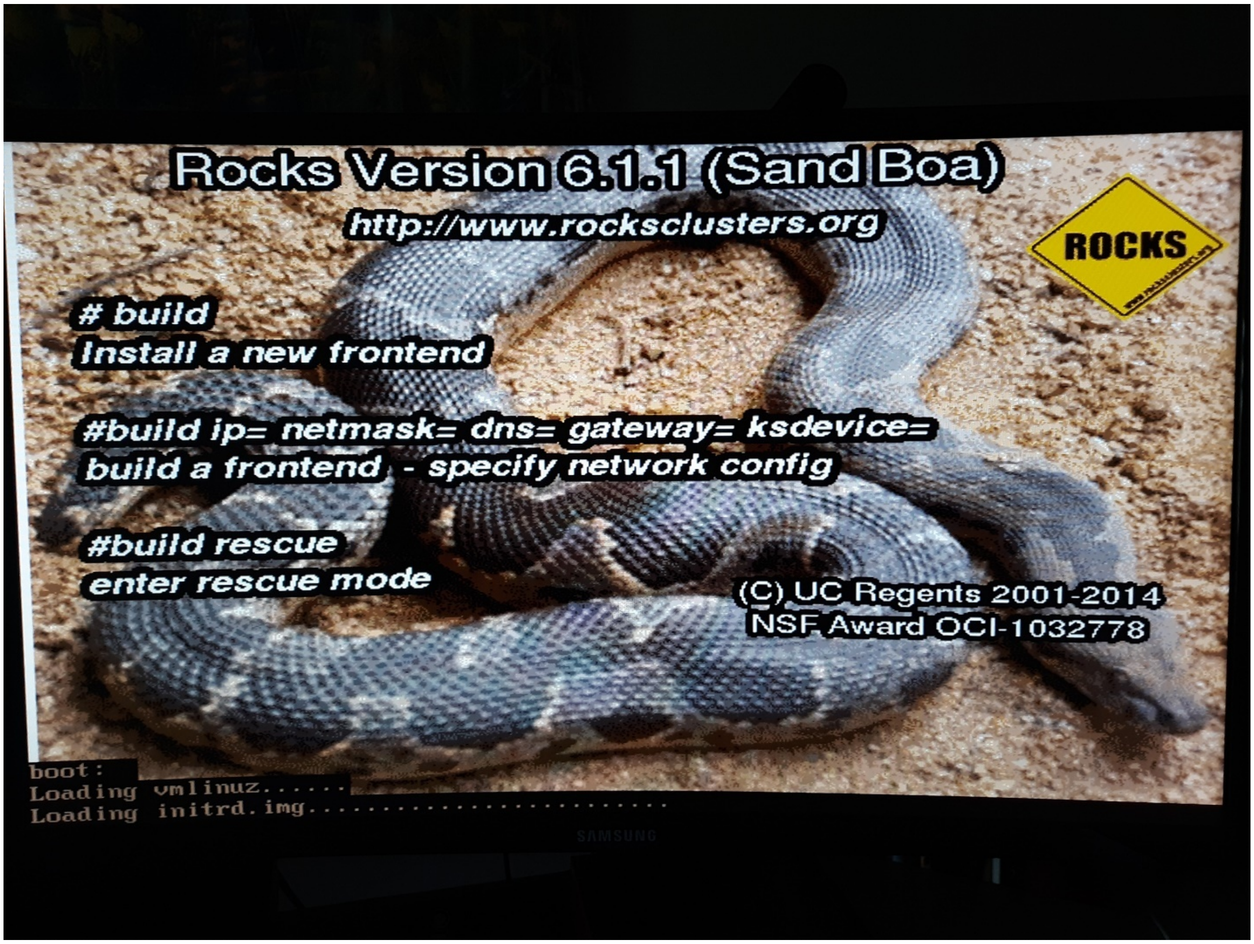

Rocks Version 6.1.1 (Sand Boa)

http://www.rocksclusters.org

build
Install a new frontend

#build ip= netmask= dns= gateway= ksdevice=
build a frontend - specify network config

#build rescue
enter rescue mode

(C) UC Regents 2001-2014
NSF Award OCI-1032778

```
boot:
Loading vmlinuz......
Loading initrd.img........................
```

Figure 83 – ROCKS 6.1.1 Initial Installer Screen

After the head node is set up and configured the compute nodes can be added to the cluster. To do this the program “insert-ethers” is execute from the command line of the head node as root. Then the compute node computer is booted. Insert-ethers instantiates DHCP and provides a PXE boot image that the compute node sees and installs on its hard drive. After this process completes, the compute node reboots and since it has already been configured, bypasses the PXE boot and boots from its hard drive image. A DNS name was assigned to the compute node in the process of its configuration. The first one is designated compute-0-0. The process is repeated for each compute node that is added to the cluster. For the cluster utilized for this project, the compute node names are compute-0-0, compute-0-1 and etc. through compute-0-5.

At this point the compute nodes have been included in the head node database and the cluster is nearly ready to render the high-resolution images for the holographic aspect of the visual music compositions.

4.3 Building and Rebuilding Compute Nodes and Maya Installation

After the cluster head node was set up and configured, the compute nodes, which ultimately conduct the computational work of rendering the images as directed by the head node, must be setup and configured. Maya 2015, the same version used to create the scenes to be rendered, must be installed on each compute node along with the DomeMaster 3D plugin. To facilitate this, a directory was created in the /home/admin/apps directory on the head node that is shared to all the compute nodes. In it, the Maya 2015 installation executable, a licensing bug fix for Maya 2015 called Poodle, and the DomeMaster 3D plugin were each placed. To build or re-build the compute nodes, the following steps are required:

1. On head node, as root, instantiate insert-ethers and select compute.
2. Start up new compute node and, if it is not already, set up PXE boot in bios.
3. Reboot compute node and see that it is getting dhcp from head node and then it will begin installing the ROCKS compute node image.
4. After completing the installation of the image, the compute node will reboot and then come up to a command line that will provide its compute node name (ea. compute-0-0). This will also be seen in the head node insert-ethers window.
5. From the head node, ssh to the newly imaged compute node.
6. Install Maya 2015 from /home/admin/apps/
7. Maya will ask for the license number and the product key.
8. Install Poodle bug fix: `cd /home/admin/apps/      ./installerR9.sh`
9. Run Maya and let it activate the license:

   ```
   cd /usr/autodesk/maya/bin
   ./maya
   ```

   **After activation, running Maya will fail with the error: "Xlib: Extension "XInputExtension" missing on display "localhost: 10.0". Segmentation fault."*
   This happens because the compute node is running headless… disregard.
   Maya is ready to go.

10. Install Domemaster3D:

    ```
    cd /home/admin/apps/
    cp –R Domemaster3D   /opt
    cp Domemaster3D.mod   /usr/autodesk/modules/maya/2015/
    ```

11. Test each new compute node by instantiating a test render. It should show up when running qstat on the job from the head node command line.

This process is repeated for each compute node and then the cluster is ready to render the images from a Maya scene.

4.4 Scripting and Scheduling

Scripting and scheduling are utilized to execute each render job. The Sun Grid Engine (SGE) is the master scheduler. Each rending job SGE is assigned is provided a numerical designation and, according the command line instructions for the job, it is placed into a queue, where it awaits execution. If there is no job ahead of it in the queue, execution begins immediately. There are numerous command line options that allow the user to work with the SGE scheduler. Among them three are utilized frequently for this project. The “qstat” command displays the status of the queue. The “qdel” command used with the queue number deletes a render task from the queue. In order to monitor the queue, which is most often accomplished using the secure shell (SSH) the watch application is utilized. The command used is “watch -d qstat”. This allows the queues to monitored directly and by default updates every two seconds. The “qsub” command is utilized to submit a job to the queue. It will be demonstrated later in this section.

To add a rendering job to SGE a script is utilized in conjunction with a command line command that stipulates the parameters of the script. The origin of the script utilized was found in a thesis written by William Robinson toward the fulfillment of the requirements of an MFA at Clemson University (Robinson). The title of his thesis is “Single Command Rendering with the Sun Grid Engine”. The script Robinson created was modified to integrate it into the processes required for this project. The modified script can be seen in Fig. 84 below. For more information regarding the specifics of how the script works, please refer to Robinson’s thesis.

```
########
```

```
#render.sge script
#! /bin/bash
#$ -o /home/admin/maya/logs
#$ -j y
#$ -S /bin/bash
#$ -l arch="linux-x64"

frame=$SGE_TASK_ID
case "1$" in
        -[0-9]*) let frame=$frame$1; shift
esac

nice -n 15 /usr/autodesk/maya2015-x64/bin/Render -s $frame -e $frame "$@"
########
```

Figure 84 – Sun Grid Engine Script (Robinson)

An example of the command line command that executes the script and creates a job for the SGE queue is shown below in Fig. 85.

```
########
#SGE Maya Render
qsub -t 1-2400 /home/admin/maya/scripts/render.sge -proj /home/admin/maya/renders
/home/admin/maya/scenes/test1.mb
########
```

Figure 85 – Command for Executing and Stipulating the SGE Script

These two key components, modified as necessary for each rendering job provide access to the cluster.

4.5 Monitoring the Render and the Cluster Hardware

It is important to monitor the progress of a render for several reasons and on two levels. First, the status of the hardware during the render is important. If at any time in the process a given compute node is not operating at or near full capacity something has gone wrong and needs to be corrected. This can be the result of a power outage that lasted longer than the batteries on the UPS units or an overheated CPU or motherboard or other potential reasons. Secondly, the status of the files during a render needs to be monitored. What frames are being rendered and how long they require to render are important factors. For instance, if a render is taking an inordinate length of time to complete the render job needs to be stopped and the scene re-evaluated. Some renders might take several months to complete, which would generally not be practical. Additionally, a render might fail due to some file management error, or perhaps a shortage of disk space on the head-node. For these and numerous other reasons monitoring is important. Though the specifics are difficult to see from such a small image, Fig. 86 is a screenshot taken from the workstation that is monitoring the status of the SGE queue. There are three command line windows, each of which are using SSH to access different areas of the head node. The left is monitoring the status of the SGE queue. The center is monitoring the directory where the left camera images for each frame are either in the process of being rendered or are complete. The window on the right side is similar to the center window except it is monitoring the right camera images for each frame. Figure 87 is a screenshot of the Ganglia monitoring system that shows hardware status of each compute node as well as the head node. An inexpensive video monitor is located next to the wall in the studio that is constantly displaying the Ganglia graphs. At a glance, one can discern the status of the compute nodes and the head node.

It should also be noted that one should monitor the size of the log directory. For each frame rendered a log file is generated and, though small, these files can accumulate over several renders and consume more disc space than expected. If these files were to consume all the disc space a render failure would occur.

Watch -d qstat, watch -d ls -la

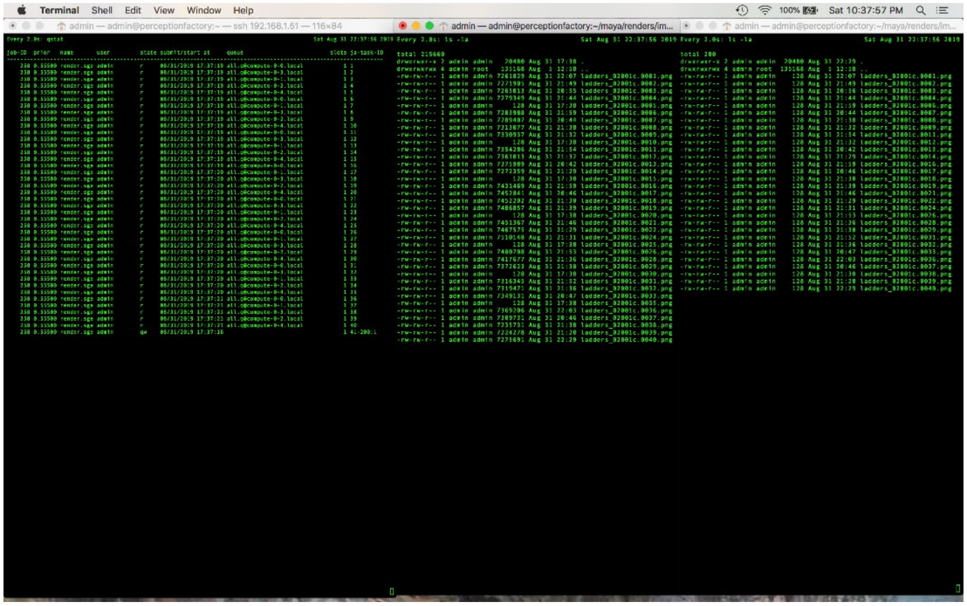

Figure 86 – SGE Rendering Monitor on Head Node Via SSH

CPUs Total: 44
Hosts up: 6
Hosts down: 0

Current Load Avg (15, 5, 1m):
193%, 196%, 195%
Avg Utilization (last hour):
191%

Cluster load_one last hour sorted by name

Metric: load_one
Show Hosts Scaled: Same
Size: small
Columns: 6 (0 = metric + reports)
Show only nodes matching
Max graphs to show: all
Sorted: by name

(Nodes colored by 1-minute load) | Legend

Ganglia Web Frontend version 3.5.12 Check for Updates.
Ganglia Web Backend (gmetad) version 3.6.0 Check for Updates.
Downloading and parsing ganglia's XML tree took 0.1826s.
Images created with RRDtool version 1.4.8.
Powered by Dwoo 1.1.1.

Figure 87 – Ganglia Cluster Hardware Monitor on Head Node

4.6 Importance of RAM Memory on the Compute Nodes

At one point, during the rendering of the images for this project, Maya was unexpectedly not rendering two or three images out of the eight that it should render. The reason for this was not readily apparent since Maya 2015 had always been the most stable version I had worked with including the more recent ones. After much consideration and some rather extensive trouble shooting it was discovered that each compute node required more RAM memory to alleviate the issue.

To explain this phenomenon, during a render, each compute node is running eight instances of Maya… one for each core of the processor. This means that for each instance of Maya the entire scene is loaded into RAM memory. If the scene was less than around .9 GB in size, this was not an issue since the compute nodes originally housed 8 GB of RAM and the stripped-down CentOS operating system installed on each compute node, by design, requires minimal memory. Yet, when the Maya scene file size was near or over 1 GB in size, Maya was unable to render all eight of the images… it could only render the images for which it had enough RAM.

After the source of the issue was discovered, another 8 GB of RAM memory was purchased and installed in each compute node. This allowed the file size limit to approach 2 GB. Since 32 GB of RAM was prohibitively expensive, 16 GB would have to suffice. Since this discovery, the project has pushed toward this constraint of file size per scene and so at times has had a direct influence on the content of the scene. As a work around, some scenes have been divided in such a manner as to allow two renders that could be edited together in the editing phase of the project. Nevertheless, this is an example of how a rendering paradigm is monumentally important and is a factor that an artist working in this medium needs to be aware of.

4.7 Technically Imposed Constraints

Due to rendering times, file size constraints, and other limiting factors that constrain the ability to render extremely complex scenes, the cluster is an integral part of the creative process. Specific scenes that were produced and then either abandoned or modified by reduction, have been directly affected by the limitations of the cluster. The scenes for this project, as mentioned previously, often contain multiple refractions and reflections thus greatly increasing the computational horsepower required to render them in comparison to scenes without such challenges. If one can imagine the monumental task of calculating the position of every virtual photon in a moving scene that has innumerable reflections and refractions, an understanding of the challenge is apprehended. On the other hand, it is quite amazing that the clustering paradigm empowers one to create a visual complexity that would have been previously unimaginable and impossible just a few short years ago. As Igor Stravinsky, the well-known composer from the early 20th century stated, albeit paraphrased here, it is constraints that engender creativity (Stravinsky). Thusly, though perhaps not directly apparent, the HPC cluster directly plays an important role in the creativity involved in this project.

4.8 Chapter Conclusions

Evidenced by the information covered in this chapter, even though high-performance computing primarily involves a technical perspective, it is nonetheless extremely stratified within the multiple dimensions that integrate it into this project. Between database management, hardware implementation and maintenance, network facilitation, and scripting and software configuration, this technical paradigm is intricately interwoven into the compositional procedural stratifications. As such it becomes an active member of the quadrilateral reciprocative feedback loop in which it poses constraints that engender a deeper sense of creativity from the composer, which ultimately integrally tethers the other constituents of the loop it. It is required and merged into the process that produces the evolutionary progression of counterpoint in the field of visual music composition. Though perhaps indirectly, it is an essential factor in the production of the multidimensional spatiotemporal experience through its connection to rendering the stereoscopic (pseudo) holographic image sequences.

The limitations posed by the currently available computational systems are a direct influence in the compositional decisions involved in employing an expansion of the traditional tone row approach to musical composition. Due to the fact that rendering cycle times can require between two weeks and four months, as implemented here, the quantity of stereoscopic visual material available for a given composition is limited. Therefore, such an approach is not only warranted but required. Herein is demonstrated the manner in which the contrapuntal stratifications of the HPC rendering is integrated within the highly stratified compositional approaches with the other involved participants.

CHAPTER 5. THE NEXUS – INTERDISCIPLINARITY

5.1 Chapter Overview

The crux of pairing holophons and holograms, to be diffused and projected in 3D/360-degree venues, toward an evolutionary progression of counterpoint in the field of visual music, is in the nexus of the stratifications involved in each constituent. This pairing, and the methodologies involved in their creation, provides the potential for a multidimensional spatiotemporal experience by placing an audience within a fully immersive environment consisting of the confluence of each of the numerous strata created. These environments are outside the normal perspectival views afforded daily perception and so tend to move the audience outside their fixated, limited, and perhaps even (soon-to-be) outmoded views of reality… if for a few moments.

As will be described in this chapter, the quadrilateral reciprocative feedback loop, between the composer, the auditory and visual indeterminate processes, and the HPC clustering, actively engages each in a loop of constantly exchanging and evolving influence. It is the confluence of the countless interwoven methodological ontologies, including the composer's sensibilities, that facilitates unique stratifications of aesthetical and formal elemental structures thus fostering an intricate and novel counterpoint of intermedial foci.

After achieving an understanding of the technical undergirding of the processes involved in this research, we can now venture into the murky waters of creative impetus. Please keep in mind that such an endeavor is by nature fraught with uncertainty and communicative complexity. The constant flow of moment-by-moment problem solving and decision-making occurs on multiple levels simultaneously thus rendering it challenging to capture and then represent with the syntactical symbology of the written word toward a semantic apprehension. We will nonetheless persevere in an attempt to allude to the intersections and interior cross-influences between the various elements of the overarching processes. As previously stated, they are non-linear in nature, many of them occurring on a nearly subconscious level. To further add to the challenge is the praxis and subsequent aesthetic of integrating indeterminacy-based systems. The nature of such systems lends toward reducing, if not entirely eliminating, conscious volition from the process. Though it may run counter to contemporary academic perspectives, here the practice is keenly focused upon removing "knowledge" from the process considering it an impedance to creativity.

This may be considered an unrealistic goal and yet it is the major driving factor behind the author's work. Given these conundrums, the intent of this chapter is to broach a level of understanding in these areas, which will of necessity proceed through both reductionist and relativist perspectives.

5.2 Formal Considerations

In the early twentieth century, Thomas Wilfred created numerous visual compositions, which he called "Lumia" (Orgeman, et al). They were realized using mechanisms he designed and built to manipulate light in beautiful undulating non-objective forms. Today, he is considered a pioneer of contemporary visual music. Yet, one interesting fact is that Wilfred allowed no sound to accompany his work, considering it a distraction from it. Wilfred and other artists of the time, including painter and synesthete Wassily Kandinsky, considered the visual image musical (Abbado). Certainly, especially in non-objective compositions, light and sound are extremely challenging to merge into a cohesive artifact. As previously mentioned, we tend to be consciously aware of the visual elements of our environments more so than the auditory. On the other hand, the auditory aspects powerfully inform us often on a subliminal level. It has been my personal experience that when presented audio-visual material I naturally tend to give the visual precedence over the auditory. In contemporary cinema, the musical score is most commonly viewed as an accompaniment to visual elements. This implementation is quite successful in conveying and supporting the narrative.

For this project, the intent is to provide a perceptually balanced relationship between the audio and the video on conscious and subliminal levels. This is not to intimate that each element is equally prominent during every moment of a composition. Often, one or the other shifts in and out of prominence as the composition progresses. An example of attempting this type of relational balance might be that at a given moment, the visual may fade to gray in order to allow the music to move to the perceptual foreground. At times, perhaps the music is very quiet or completely absent allowing the visual to occupy the foreground. This perceptual interplay of prominence is viewed as a formalistic compositional element.

In general, the formal characteristics of these visual music compositions are anchored by musical perspectives. Though the possibilities are endless, the ancient sonata form, and myriad variations of it, are favored as examples of the general approach. Simply put, the sonata form is in three sections consisting of the exposition of the theme(s), the development of the theme(s), and a recapitulation of the theme(s). Since the themes in this work can be perhaps less obvious than those based upon melodic content, as described below, the sonata form is intended as a pertinent example, in the general sense, of describing audio-visual relationships that recur in varied manners and combinations throughout the duration of a composition. This approach can be thought of broadly as “theme and variation”.

Themes, in the sense referred to herein, may be viewed as parametric motifs. From a musical perspective, spatial orchestration is considered a formal element in which a composer might consider sound object perceived location and movement within a composition (Lyon). Here spatialization is considered a formally available parameter from audio and visual perspectives alike. Along with effects such as delay, filtering, frequency (including rhythmic structures), amplitude, density, object choices and morphologies, and others, spatialization is considered a parameter worthy of thematic focus. This formal approach may not be directly recognizable to the untrained audience. However, it is reasonable to anticipate that it does provide an organizational structure that will propel the piece through the temporal dimension in a manner intuitively recognizable to the listener.

Within the compositional structure the afore mentioned parameters are employed in a contrapuntal manner. The term “parametric counterpoint” describes the interplay of said parameters (Laske). For instance, the density of a given set of contiguous events, the manner in which they overlap, may be set against the frequency ranges of the events. If this isolated example is generalized to the full range of available parameters, one can begin to realize the innumerable combinations that can be related in specific contrapuntal conceptual frameworks. Viewed from the perspective of formal thematic material, one is able to realize nearly infinite possible combinations.

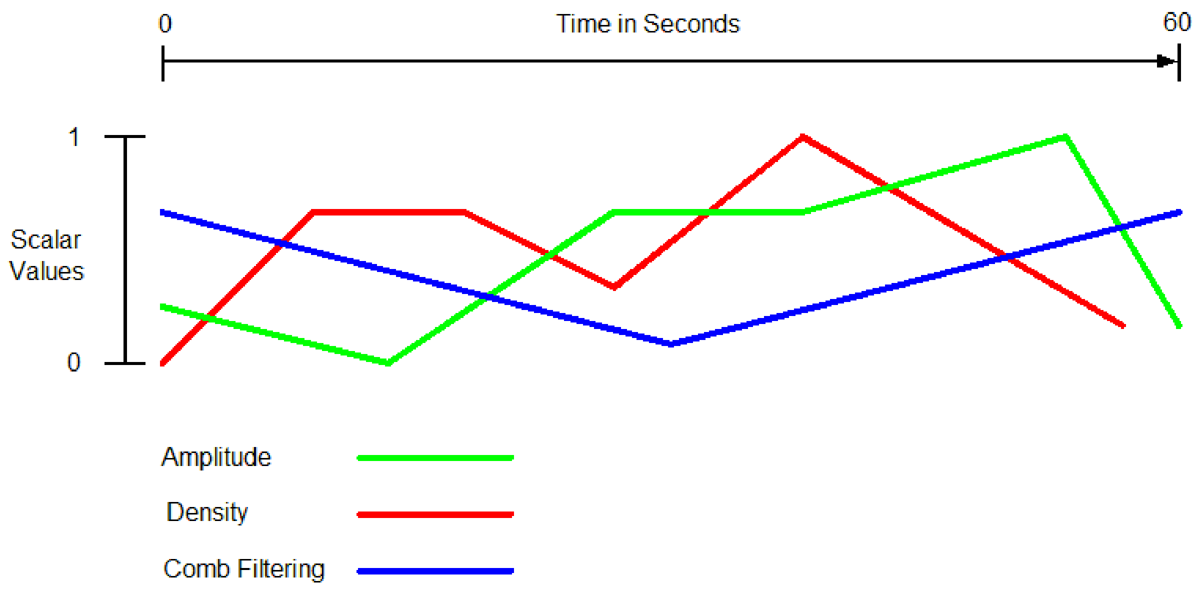

Figure 88 – Graphic Representation of a Simple Parametric Counterpoint Example

From the perspective of parametric formalization and counterpoint, one can further generalize this approach to include the visual objects and events in relation to one another and in relation to the musical evolutions. Though these two basic elemental constructs seldom act in unity, which would be nothing more than a visualizer program good for a screensaver, they do interact in manners related to their origins. Recurring visual thematic materials and structures occur in relation to musical counterparts and do so with structurally based intent. The two elemental dynamics interact in tandem to create an overall structural schema that often could be diagrammatically construed. As will be described later in this chapter, the two inform and influence each other in an interesting interplay that could be described as a feedback loop.

It is considered the responsibility of the composer to lead the listener through each composition in a manner that offers numerous signposts and landmarks along the course of a composition. Since both the music and visual work range from abstract to non-representational, it is intended that each new composition establish novel syntactical representations through which novel semantics are achieved. These compositions therefore require an investment on the part of the listener. Their depths cannot be easily fathomed in a single experience. It is through repeated listening and viewing that they can be apprehended, which renders their true characters revealed.

5.3 Material Derivation – Complexity and Emergence

We begin a description of the manner in which the raw material is generated for these holochoric visual music compositions with a discussion on the subject of complexity and then move on to the subject of emergence. Though both belie a theoretical perspective, which could be considered antithetical to practice-based research, it is through these theories that the basis of praxis is derived.

Complexity is most commonly thought of as derived from complex systems. For instance, if one is to attempt the design of a computer, the first step would be to decide upon the utility of such a device, its intended use, and design from there. From the onset of the project one would proceed with the constraint of designing toward a specifically complex system. This is a top down process in which the final goal is known from the beginning and worked toward. If one were, however, to consider the manner in which the complexity of our current world was created, such a constraint need not be present. In other words, the evolution of the world as we know it today, which occurred over billions of years, need not have an end goal. Unlike designing a system with a specific intent, the natural work is simply evolving (Wolfram). This bottom up approach simply occurs with no apparent specific end goal.

In his book, A New Kind of Science, Stephen Wolfram recounts an in-depth research methodology, conducted over a period of twenty years that employed very simple reiterative rule-based algorithms to create extremely complex behavior in the evolution of cellular automata. These rules consisted of the black or white coloration of the individual cells in a row of cells based upon the cells adjacent to them in the preceding row. Below, in Fig. 89, is an example of what he determined to be the 18th rule.

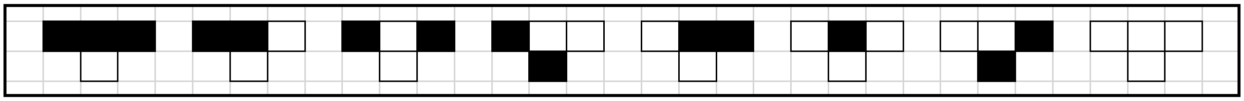

Figure 89 – Wolfram’s Rule #18

Here we can see that if the cell directly above and directly adjacent to the left and to the right are black then the new cell is white. A similar rule is stipulated for each of the eight possible contingent combinations of the upper cells. Each time a new row is created, its state is dependent upon the state of the previous row. Though Wolfram used Mathematic® to conduct his research, I created a similar rule set using Processor to generate a similar system, see Fig. 90. Below is an image that portrays a small portion of the results. For my algorithm Rule 18 was employed as a basis but a tendency mask was inserted into the rule set to determine the results. For example, if the three adjacent cells above the new cell were black then the tendency of the new cell was for it to be white. The tendency mask was based upon a triangular distribution in the same ilk as those mentioned in Chapter 2.

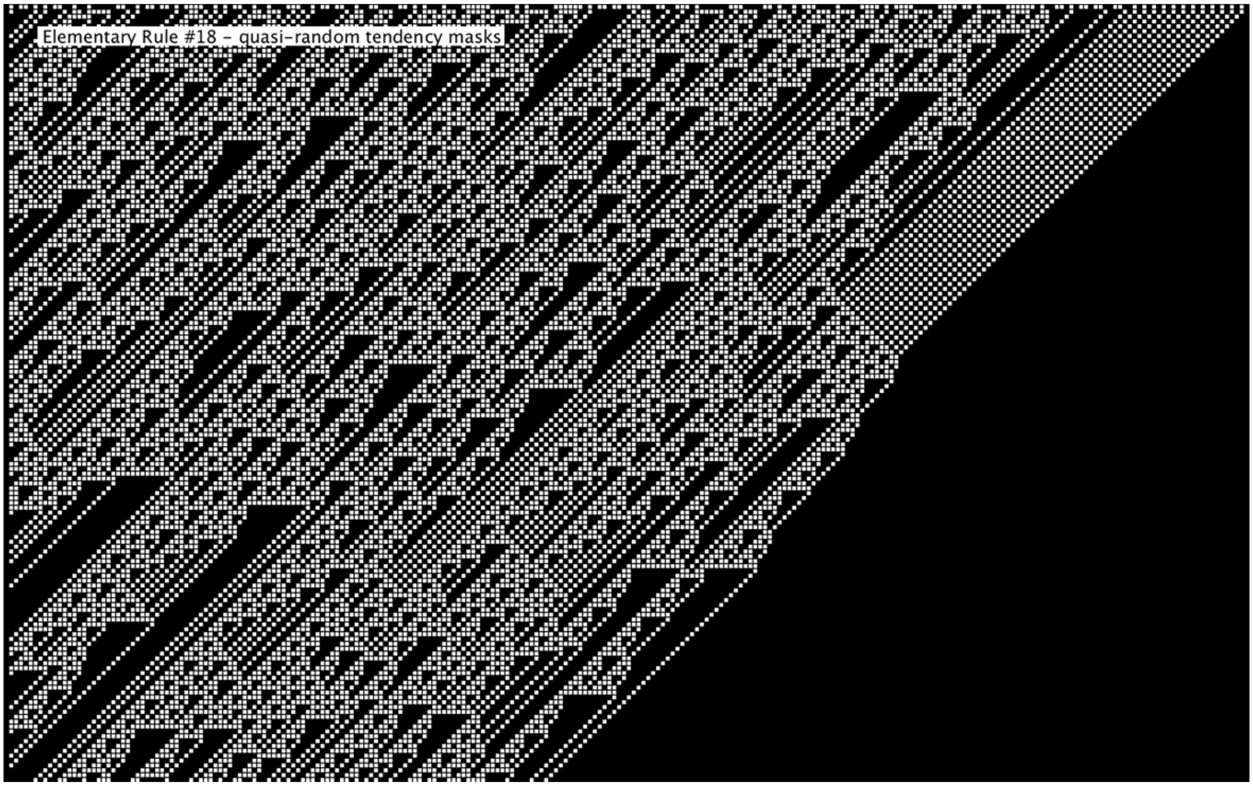

Figure 90 – Excerpt from an Evolving Cellular Automaton Based Upon Wolfram's Rule # 18

As Wolfram pointed out, this extremely complex development of the patterns of the cellular automatons was derived from a very simple rule set that, through innumerable reiterations, developed toward complexity. For the research described herein, similarly simple rule sets, a modicum of indeterminacy, and reiteration generated the complexity and richness of the compositions.

As described in Chapter 2, musically the simple rule sets were formed using tendency masks, a massively branching Csound orchestra, and a reiterative process. As described in Chapter 3, the visual side of this research was generated based upon premises of reflection, refraction, and reiteration. The latter was achieved with high-performance computing, which was a powerful enabler but also a constraint. The constraint was determined by the speed of rendering based upon the complexity of the scene. There were many scenes throughout the process that were discarded due to the impractical rendering time requirements. Though it was not unusual for a 120-second scene to require perhaps up to eight weeks to render, some of the scenes that were ultimately discarded would have required a year or more to render. Though highly desirable they were obviously not practical. As mentioned previously, there were also several other very desirable parameters in the scenes that were not included for similar reasons. Among them were ambient occlusion, final gathering and etc. In this manner the HPC cluster was quite influential in the final outcome of the work.

It should be pointed out that it is not the steadfast perspective of the author that complexity equals quality. Throughout the history of art, music, and computer science, there have been very simple end results that were considered by many people to be beautiful. Conversely, there have also been many artistic expressions that were extremely complex that were for many people quite unpleasant. Here the drive toward complexity is not for complexity sake but instead for the richness that it affords. Though simplicity can be beautiful it is extremely difficult for very simple expressions sustain a lengthy composition. Simplistic expressions generally become uninteresting after a few minutes and eventually, through repetition, they can become tortuous. An example could be popular music. Though much of it is extremely desirable, depending upon the listener's preferences, most popular music is no more than three or four minutes in length. Similarly, though complexity is more likely to sustain more lengthy composition, it nevertheless generally requires recognizable characteristics along its expanse in order to sustain attention.

Another aspect of complexity that is pertinent to this research is the phenomena of emergence. When one employs a compositional methodology that derives from the generation of complexity, there is much material generated that is undesirable from an aesthetical perspective. Considering this, the author is a miner for emergent phenomena (Rhoades, 2013). In all the myriad material generated in this process there occurs an enigmatic segment or phrase that shines in the otherwise inert set of results. Most of the material generated is ruthlessly discarded in lieu of the emergent gems. It is these gems that are first recognized and then brought out to play… developed.

Developing an eye and ear to recognize emergent behavior is an art in itself. In fact, in this artistic practice, it may be *the* art. An extreme sensitivity, and awareness, is required to recognize the potential and pertinence of a segment, section, or phrase that is located within numerous others of less value. The art of recognizing emergent behavior and then the energy to develop it comes from a deep sense of quietness within. This awareness is not possible when the mind is actively participating in the usual, energy-sucking, internal dialog and its accompanying impedances. Though such a state of mind cannot be developed or controlled, the author is vigilantly cultivating a state of mind where quiet creativity can occur. Then it is a matter of working with it. Thinking is the antithesis of creating (Krishnamurti).

It is through the deliberate implementation of the generation of complexity, and the awareness to recognize emergent behavior when it occurs, that this work has been accomplished. Though no claims are being stated, it is nonetheless apparent that state of mind in conjunction with process is the prime determinant of this creative expression.

5.4 Choreographic Ontologies

Reduced to its essence, choreography could be thought of as the planned, improvised, or designed relationship between one or more mediums of movement. Of course, the most common use of the term is in reference to dance. There it refers to designed sequences of movements of one or more dancer often in relation to music. For this study, the term is used in regard to the designed relationship between virtual holographic objects, including the environments in which they exist, and holophonic objects and the environments in which they exist. Though these relationships are separated for the purpose of discussion, here they are intrinsically melded to the point of union. It is the visual music composition as a whole that is the culmination of the creative investment and it is intended that the audience experience it as such.

The elemental aspects of this research are ultimately intended to be integrated into a whole by way of establishing choreographic relationships between them. Obviously, this approach can be likened to dance and it is possible to draw direct parallels between the two mediums.

However, one major difference that should be kept in mind is that in this holochoric expression the music is perceived as moving within the space along with the visual objects. We could describe the music of more traditional forms of dance as occurring in three dimensions as opposed to the four-dimensional movement of holochoric visual music compositions. Traditional music is primarily perceived on the horizontal plane, reverberation notwithstanding, and on a temporal dimension. In a traditional ballet, for instance, the dancers are perceived in four dimensions. These are left/right, forward/backward, to a lesser degree, up/down, and then temporally. The movement of the music is perceived temporally and perhaps on the horizontal plane due to the various instruments playing at various times across the stage or orchestra pit. In this regard holochoric visual music affords the composer the novel opportunity to relate visual and audio objects on all four commonly perceived dimensions. This provides an unprecedented aesthetical territory rich with unexplored possibilities. Nevertheless, this work is clearly influenced and informed by more traditional choreographic techniques and simultaneously it is a departure from them.

The chorography of Vaslav Nijinsky, Alwin, Nikolais, Merce Cunningham, and Pina Bausch have had a major impact on this work. More contemporary choreographer/artists such as Bill T. Jones, Shen Wei, Scotty Hardwig, and others have had a dynamic influence on the author’s aesthetic with regard to the relationship of audio and visual movement. From conversations with Hardwig, who is a dancer, choreographer, and professor at Virginia Tech, it was realized that each choreographer maintains a focus upon four primary elements in shaping their work. Certainly, they come into play in the author’s work.

- Time
- Energy/Quality
- Shape
- Space

In watching any ballet or other form of dance, we can see that, considering these four elements, there are times that the dancer(s) move in synchronization to the music. At other times they move in what could be described as a contrapuntal, harmonic, or otherwise complementary manner with the music. During still other moments, the dancers are moving in a completely disparate relation to the music. Though perhaps simplistic, these are very basic choreographic tenets. Combining each of these perspectives one can apprehend a universe of intricate relationships for dance. Generalize this to holochoric expressions wherein these relationships exist not only with the visual and auditory objects but also in the visual and auditory environments within which the objects exist and move, we see the potential of a nearly infinite multiverse of novel creative expression. As previously mentioned, the work should be viewed, both from the perspective of the composer and from that of the audience… holistically. The dance and music are ultimately a singular expression taking myriad forms.

Relating audio and visual elements in the context of this project is complicated since there are numerous levels of elements in each domain occurring during nearly every moment of each composition. Coupling a given audio object with a given visual object, for instance, often leaves several others, occurring in the same time domain, uncoupled. Additionally, though the relationships between coupled objects may be in unity, as is, for instance, common in several early compositions by Oscar Fischinger, often the relationships are less obvious and also evolving throughout a phrase or sequence. Even if multiple objects are related in some manner an audience member might not be able to ascertain a direct correlation.

Relationships that are constantly changing between object combinations that are constantly changing can at times be difficult to notice. One formal device that is often incorporated is that of establishing moments during which audio objects and visual objects directly and obviously align. This allows the audience to achieve junctures of formal clarity thus providing signposts that can be referred to literally, or by way of allusion, as the composition unfolds. It is these moments of establishing recognizable elements during a composition that allows a composition to be non-objective and yet, at least somewhat, palpable.

Within the choreographic perspective, apophenia is the composer's friend. Apophenia is the tendency we human beings have to perceive relationships between occurrences even when they may not exist. It is the natural action of the human brain to make sense of and to organize situations into apparent comprehension. In this manner, a composition provides a basic grid work within which each audience member fills in the spaces according to his or her personal propensities. Acknowledging this frees the composer to explore areas that might otherwise be avoided due to preconceived notions of possible disassociations. This infers an audience that is anxious to listen and watch.

Certainly, audience members may never consciously apprehend innumerable choreographic, and other relationships that are established between elemental constituents throughout a composition. Nonetheless, perhaps on a less than conscious level, it is these relationships that allow for sense making or at least for a feeling of rightness, direction, and completion within the composition. It is here considered the composer's responsibility to be aware of this active side of listening and watching and to provide the necessary formal occurrences during the course of a composition in order to create tangibility in an otherwise potentially intangible experience.

5.5 Novel Aesthetics

The incorporation of holography and holophony into visual music compositions is accompanied by unprecedented aesthetic choices. In addition to the novel potentials with regard to movement in the dimensions of time and space, as mentioned in section 5.4, there are other questions that have arisen during the course of this research. A few of these choices will be broached in this section.

What could be considered the "cross-eyed effect" is when visual objects are placed so close to the cameras within a Maya scene that the viewer's eyes cross when viewing the scene. This is never an issue in 2D video since objects simply become larger on the screen perhaps providing the perception of moving closer to the viewer. Working in stereoscopic video this is not the case. Instead, due to depth perception afforded by stereoscopy, the object is perceived as moving toward and eventually past the eyes and behind the viewer. The time during which the object is so close that the eyes cross is an uncomfortable, or at least unusual, feeling. Speed of movement is a factor in this effect. It appears that an object moving quickly in and out of the cross-eyed zone is much more acceptable compared to one that lingers there during which time the eyes attempt, unsuccessfully, to focus upon it. Since here the quality of holography is dependent upon proximity to the viewer, this is a critical aesthetic choice. It needs to be monitored carefully in order to take advantage of the effect of closeness to the viewer without overwhelming him or her… unless this is the intention.

Similar to the cross-eyed effect is the "monocular effect". Here a visual object is so close to one of the viewer's eyes that it gives the viewer the perception of wearing a patch over one eye. This is again a time dependent factor. When it occurs for a short amount of time it is acceptable but when it lingers it once again potentially produces discomfort in the viewer.

Similar to both previous examples is a phenomenon that could be termed the “claustrophobia effect”. This occurs when either the space or several large objects, visual and/or auditory, are located within close proximity to the listener for a relatively long period of time. The effect can produce a claustrophobic effect.

Another phenomenon that is a potential aesthetic pitfall could be termed the “syncing effect”. This is perhaps more subtle than the others yet it can nonetheless be disorienting. This is when the movement of the audio is contrary to the video in a manner that is in direct contrast to the real world. For instance, if an object is seen to be moving toward the viewer and the music is heard to be moving away, this disparity can be disorienting. It is a form of spatial decoupling.

When a large open environment is spinning while objects are moving in and through it the viewer may perceive the “dizzying effect”. Here he or she loses track of their equilibrium becoming uncertain as to which direction they are facing and which direction is up or down.

The “strobe effect” occurs under several conditions. Perhaps many objects are perceived as moving rapidly through the viewer’s head or perhaps the lighting is such that the scene is alternating between bright and dark. These effects can also be quite disorienting to the viewer.

There are other aesthetic considerations when creating audio/visual material for immersive environments, which are referred to as “frameless environments”. For instance, it has been noted that pacing of material in a full dome, as in a planetarium, is normally required to proceed much slower than in a 2D environment (Yu and Wyatt 2016). This is pertinent as well in a holographic/holophonic virtual environment where it is not possible to view the entire screen at the same time. The audience can be overwhelmed by the realization that they cannot see everything that is occurring on the screen.

This is a brief and general list of some aesthetical challenges of the work involved in this research. Perhaps a theoretical study is called for in order to document these and other effects in detail. However, here the concern is with praxis. This relegates these effects to being aesthetic questions. For instance, is it aesthetically useful to employ the dizzying effect to produce a disorientation that is then resolved in the next section? Considering this simple question, it can be observed that each of these effects can be seen as formal elements. Used judiciously, they could produce an interesting visual music experience.

Certainly, with each of these and other phenomena that are products of 3D/360 holographic/holophonic compositions user experience plays a part. During the course of conducting this research, it was discovered that as I became accustomed to viewing material using a head mounted display, my mind was much more accepting of that which would have disoriented me in the beginning. It could be that holochoric visual music compositions specifically and holochoric art in general may be the birthright of a new generation of viewers… those that are accustomed to virtual world gaming. Time may reveal the validity of such suppositions.

5.6 Reciprocation – The Feedback Loop

Derived from very different processes, the music from stochastic algorithms and the visual from parametric key framing, the two main constituents of this work must interlace in a transparent manner. In order to achieve this, a feedback system is implemented during each available intersection. This system relies upon the composer recognizing the potential for any given aspect of the video, the audio, or the rendering to influence the others and vice versa. This is actually a four-way feedback loop in considering that the audio and video content is created from indeterminate processes thus relegating the composer as a direct aspect of the feedback loop. Since the composer cannot anticipate the results of the generative processes, he becomes a functional part of them. Figure 91 depicts a graphic representation of the relationships referred to here. Let’s proceed with a straightforward example.

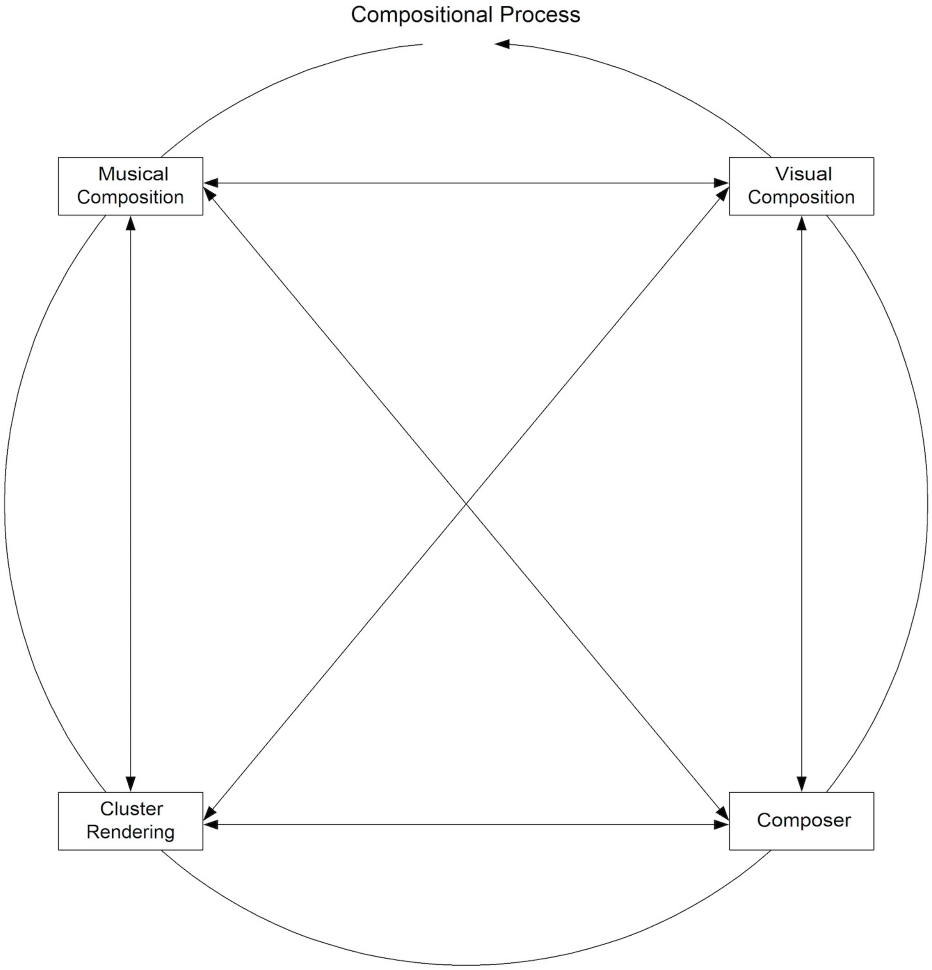

Figure 91 – Quadrilateral Reciprocative Feedback Loop

Let's first state that the music for a given composition was completed first. It was composed with perspectives similar to a sonata form and consists of three sections. These formal elements then become the structural framework within which the video, rendered next, will be placed. This is the first instance of feedback. During the Maya scene creation of the video segments this is kept in mind. So now we have three varied but related sections of videos rendered. It turns out that each section takes place in an environment specific to it. The environment in the first section is large as if the viewer is outside. The second section is a very small space in which the view feels almost claustrophobic. The third section's environment is a huge enclosed area perhaps similar to a cathedral. The first two environments posed no issue when rendering, however the third, though beautiful, was not possible to render because the scene was 20gb in size and there is only 16gb of available RAM memory in each of the compute nodes of the cluster. It had to be modified to fit the hardware constraints of the cluster. Thus, the cluster is now involved in the creative process and takes an active role in the feedback loop between the composer, the video, and, as we will see, the music.

The video segments are next imported into Premier. A stereo mixdown of the music is imported and we being editing the video segments to match the music… another aspect of the feedback loop. Very soon it becomes apparent that the reverb quality of the 3rd section of the music does not fit the reverb normally associated with the qualities of the visual environment intended for the 3rd section… if this is desirable to the composer… no problem. However, if the composer intended to have the reverberant qualities of the music in the 3rd section be somewhat realistic or directly related to the visual, then the reverb utilized for the final mix of the music for the 3rd section must change. In this situation we can observe that the composer, the visual components, the audio components, and the HPC cluster are each being influenced by the other. Perhaps later in the editing of the 2nd section of the same composition, the composer determines that there is additional visual material that needs to be included that extends the original length and content of the music for that section. Then the video, including the additional material, is exported as a low-res mp4 file and then imported into Sonar. There the music can be edited, and more material generated if necessary, to match the change in the video.

This type of feedback loop between the four major constituents in the compositional process is essential to serve the compositional process but also, and more importantly, to serve the composition itself. Though it may sound implausible, the author at times seems to feel as though the composition is a living thing and that he participates in the process simply to serve it and its needs. This is not such a different perception from that of Michelangelo's regarding his sculpting. When asked how he could carve from marble such a beautiful statue, like his Michael, he replied that he did not carve the statue but instead simply removed the marble that blocked Michael from being seen. Truly many creative artists have spoken or written of similar experiences and perhaps it is not so farfetched as might initially be perceived. Ultimately the proof is in the result. From this perspective, the composition itself is also a participant in the feedback loop of its creation… directing the others to its own means.

5.7 Extensive Methodological Hybridization

In retrospect, through the analysis and documentation of the processes involved in the creation of this project, a trend can be observed that is related to the interdisciplinary nature of the work. This tendency is best described as a methodological hybridization, which is an extremely powerful amalgam of processes toward a singular creative expression.

Most of the key processes utilized throughout this document are based upon theoretical constructs that are typically implemented as (primarily) individual approaches. However, it was discovered that this hybridization of numerous processes dynamically contributed toward a much more intricate and diverse tapestry of content. Such diversity lends toward the creation of a uniquely interesting artifact.

By "hybridization" is intended the blending of existing techniques and applications on multiple strata. For instance, consider the use of VBAP and Ambisonics in tandem. Many composers utilize either separately but seldom are they used together. This hybrid approach to spatialization extends the reach of either component used separately, arguably, in an exponential as opposed to an additive manner as might be initially expected. Adding convolution reverb into the spatialization paradigm extends this hybridization further. Certainly, using both Csound and MAX as major components of the same musical composition further exemplifies this tendency.

Consider the use of AE as an integral aspect of the rendering process as opposed to using it as a utility acting as an ancillary component of the post-rendering process. Certainly, it is a necessary component of post-rendering considering its role of stacking the left- and right-eye images, rendered by Maya, to form image sequences to be edited in Premiere. However, when used to add effects to the original image sequences rendered by Maya, AE takes on the role of a primary content co-creator. Since the source of much of the effected material cannot be recognized, AE is contributing basic project material previously only achievable through the use of Maya. Using Maya and AE in tandem in this manner is indeed a hybridization that, again, extends the reach of both systems in a manner not available to either independently.

Reverting to the basics, the act of joining music and visual toward a singular compositional medium is itself a hybridization. This is commonly done in television and movies, yet the act of each aspect serving and being integral to the other, solely as a creative expression instead of serving a narrative, is not as common. This basic tenet of the expression of visual music portrays a hybridization that is fundamental to the work achieved herein.

There are numerous other examples that can be cited from the previous chapters in this writing that demonstrate this approach. However, the point is that this type of hybridization of processes, on numerous compositional strata, exponentially expands the potential and diversity of material afforded the composer. Each process is in itself an area worthy of specialization. As this compositional approach continues to evolve it continues to expand by reaching into new modes of expression, each a compositional world unto itself that directly contributes toward the creative outcome of each project. Each new hybridization of process dynamically contributes toward the final artifact.

As opposed to more singular approaches, it is these hybridizations that bring a greatly expanded potential and interest to the compositional process. When competence is achieved in the various constituents of hybridization and when that competence is integrated into the overall compositional process, fluidity is required in shifting between what can ultimately be seen as widely varied mediums of expression. Though widely varied, these mediums nonetheless contribute toward an intricately unified result. By constantly shifting the mindset between them and doing so with a focus on the overall compositional direction, one can explore great unforeseen expanses along the journey toward the final artifact. Often it cannot be anticipated where such a journey will lead, which is the excitement and intrigue experienced throughout the process. It is this approach that contributes toward this composer being interested and fully present in every moment of this highly hybridized, highly stratified, and highly satisfying methodological approach.

5.8 Perspectives on the Subsequent Compositions

Throughout this project I have endeavored to forge an initial foray into a novel world of visual music expression. Advantaging a venue that consists of 3D/360-degree stereoscopic projection and a high-density loudspeaker array, holochoric objects, events, and environments fully immerse the audience within an all-encompassing experience. The audience is thrust into non-linear environments that require different manners of perceiving. It is intended that this shift in perception be a catalyst for a state of altered consciousness not induced by drugs or other similar substances, but instead doing so in a safe and risk-free manner. Ultimately, it is intended to raise the level of consciousness in those experiencing the work.

The seven compositions that comprise the Antitheses project are related in several ways. For instance, each are based upon relatively simple geometric audio and visual objects that achieve an indeterminate complexity through reiteration, reflection, and refraction. Visually, these objects commonly consisted of spheres, cubes, toroids, ladders, cylinders, cones, and pipes. Such simple objects rarely occur in isolation in our daily perceptions and so there are relatively few connotative associations that could be construed from them. Thus, a modicum of openness is attained that would not be present if more complex objects, derived from realism, were utilized. Such basic shapes are however simplifications of everything we see and experience in the natural world and so create indirect associations that make the abstractions approachable even within their unnatural complexity. Finally, these shapes are the basis for much of the visual music of the past, situating my work squarely within the idiom.

Commonly used textures included several forms of glass and also metals such as chrome and copper. These were enabling factors in emphasizing reflection and refraction as catalysts of complexity. Though as few one object might be present in the original Maya scene, these characteristics caused innumerable versions of objects and environments to be present in the renders. Objects reflecting off, or refracted through, other objects and bouncing back and forth throughout each scene achieve a complexity that could not be achieved using textures that do not consist of these characteristics. Further, the variations caused by inaccurate algorithms and subsequent degradations due to numerous reflective and refractive iterations provided a visual expression impossible to achieve in a predetermined manner. In each area and moment of each scene numerous variations of the thematic material are presented thus providing a depth and breadth of interest not commonly found.

Saturated colors were often favored over more subdued colors, the former of which contrast in interesting manners with the chrome and glass textures. The lighting in each scene was often white, and obvious, as reflected bright areas in many scenes thus providing another level of complexity. These aspects resulted in an amplification of a non-linear expression that forms a departure from the visual experience of our linear daily lives. However, since we human beings seem to possess a predisposition to creating order out of such complexity, apophenia plays an important role in the apprehension of the material.

In line with the visual material, the music consists of simple synthesized base samples that were occasionally mixed with samples of concrete objects such as a grill grate or water dripping upon a metal surface. Within each composition delays, reverberations, filter sweeps, comb filtering, retrograding, and transposing, combine to create a basis for sound events that achieve the perception of multidimensionality and, to some listeners, lead to a feeling of otherworldliness. To further enhance these effects, reverberation was customized, and applied to the final 32-channel mixes, intended to compliment the overall environment in which each composition exists. This provides the feeling that each work resides in it own unique space and therefore alleviates the monotony that might otherwise occur when experiencing the entire set of compositions sequentially. As with the visual work, reiteration was a standard extension of the approach toward perceptual complexity. These elements, combined with non-linear spatialization techniques, formed a combinatorial audio expression that would be highly unlikely to experience in our daily world. Layers upon layers of sound events are interwoven in relationships that are constantly changing and evolving. Inherent to these elemental constituents are rhythmic structures that are quite subtle and that occur on each stratum. Certainly, aligning them with visual rhythms aims to bring them to the forefront. Out of time structural components form in-time perceptions requiring deep listening if a holistic participation in the overall composition is desired. The culmination of the multiple dimensions apparent in the compositions could be listened to over a lifetime and new perspectives would be achieved each time.

The lengthy render times of the visual material had many unexpected and positive influences over the compositions. For instance, since the aggregate visual material available for each composition was limited by temporal practicality, an efficiency in its implementation was necessitated. Perhaps comparable to the tone row approach that accompanied much serial music from the mid-20th century, each render required exhaustive exploration and formal development for its contribution to the composition. As is apparent in each, the basic material was limited but nonetheless creative variations abound. Though initially perceived as a limitation, this catalyzed the use of After Effects in the creation of additional variations of the original rendered scenes. Here distortions in the materiality, in the forms of textures and shapes, were achieved in both objects and environments. Additional materials, not present in the original renders, were also added such as atmospheric textures like smoke. Additionally, color pallets were able to be altered. These and other variations of the originally rendered material added depth and interest in phrasing and overall thematic development.

Themes and variations are the heart and soul of the art of composition. With them, the audience is moved though the material in a manner that is relatable and yet is constantly changing. Executing such changes are a mainstay of the artist’s creative expression. For instance, fading back and forth from a version of one scene that is time aligned with another version of the same scene causes a very interesting transitional element. These types of transitions, if applied repetitively through a set of phrasings, can be used to emphasize the rhythmic structure of a section. Though a singular example, one can gain a sense of a deep and broad range of potential in the techniques involving material development and formal structuralization. Further, this aspect of the visual material required a similar approach to the production and treatment of the music in order to maintain a relatable balance. Subsequently, the resultant tone and development defines the compositional elements as members of the same family.

Rendering times also influenced the process by providing me an extended creative gestational period in each composition. Much time was available to consider each subsequent render and its relationship to the overall composition. Further, by the time all the renders were completed for a given composition, the final editing was obvious and occurred at a rapid pace… the time had allowed for the questions and solutions to any issues to be envisioned and resolved making the final editing and compositing step extremely brief. This also considers the unconscious mind as a partner in the final outcome since in many ways it seemed at work even when the compositions were not being considered consciously. In the end, each piece flowed as through a funnel, wherein the speed increased with each advancement toward the final outcome.

Another factor played an important role in the final outcome of the project. The myriad procedural strata, consisting of numerous techniques and processes, all completed by a sole composer, allowed for the confluence of a unified intent. Here a highly focused approach was possible. Implementing multiple layers of techniques in such areas as Csound, MAX, MAYA, Domemaster 3D, After Effects, Premiere, Sound Forge, Sonar, and others without the hinderance of second guessing or compromise, allowed the outcome to be greater than the sum of the individual constituents. Expertise in each area was allowed to fully blossom free encumbrance.

Holochory, was of course a major influence in the ultimate realization of the work. Though there are exceptions, visual music compositions are generally presented on a stage in front of the audience… they are projected from the stage to the audience. Placing the audience *within* a holochoric environment, placing them within the composition instead of seeing it from the outside, is a game changer for future visual music composers and their audiences. The possibilities change drastically within the 3D/360 stereoscopic, holographic and holophonic, paradigm. Experiencing the work from an internal perspective is fully immersive. Auditory and visual objects move from left to right, near to far, and high to low and as the audience perspective changes added dimensions are revealed and experienced. Documented in this exegesis, are methodologies, strategies, and hypotheses that form a springboard for future composers, empowering them into exploring their own voice in this novel area of expression.

The aesthetical potential of this paradigm opens worlds of previously unimagined cinematic expression and exploration. In a holochoric environment, that which is commonplace in a 2D composition is often no longer viable. For instance, in a holochoric environment, when a wall covers one of the viewers eyes but not the other a monocular effect takes place. Or if the environment spins and the viewer’s perspective remains constant, the viewer loses track of what is moving. Certainly, in a frameless environment, there is no indication to the audience as to which direction they are facing. There is no apparent front to face. In a 2D expression, there is the floor, the ceiling, and perhaps the walls on either side of the screen that indicate where the audience is located in the experience. In the frameless world this mode of reference is maintained by the composer and, as an aesthetical choice, it can be manipulated in previously unimagined ways.

The relationships between audio and visual objects and environments in a holochoric expression stretch the very notion of choreography. Coupling sight and sound, no matter where it exists in spatiotemporal relationship, was a common approach achieved in this project. For instance, as a visual object rises toward the top of the environment, perhaps the location of a audio object begins to rise along with it but then abruptly falls and, simultaneously, the room begins to spin forward from its static position. Afterward, the location of the next dominant audio event couples with different visual objects or objects in a very different manner. From this perspective, movement becomes an important, and at times, primary formal element. This is not achievable in a 2D environment except in, at best, two-dimensions, which is quite simplistic in comparison.

Though programmatical approaches were experimented with, the relationships between the audio and visual were dealt with in a manual methodology by the composer. It is possible to program the relationships between the movements of the two main kinetic mediums. However, at this time, a computer cannot be creative beyond its programming. I found that by dealing with these kinetic relationships manually, my creativity is engaged and flourishes. Actually, it is my perspective that transcending my own programming is the essence of creativity. This is something that contemporary computers are incapable of.

As can be seen throughout this work, novel vistas of creative expression are exposed, not only as potential but as necessity of the medium. From my perspective, the deeper I have delved into this work, and the more deeply I will do so in the future, a new artform is emerging. It is in fact based upon the visual music of the past and yet it goes beyond it into areas that I have barely scratched the surface of. Certainly, as the technology becomes more available and more advanced in terms of computational power and speed, this approach will become more widely accepted and more deeply explored. Being able to interact with the high-resolution holochoric images in real time, will be afforded by this development. Real time rendering of the visual, which we can currently accomplish with the auditory, will drastically change the possibilities herein described. Many desirable current attributes afforded the extended rendering times will be lost as rendering speeds increase. Yet, new attributes will become obvious, explored and employed as well. This work provides an initial step into this new frontier and it is exciting to anticipate that which will come next.

5.9 Chapter Conclusions

In this chapter we have endeavored to describe a few of the thousands of influences that contribute to the fruition of a composition. This is a monumental task and one doomed to failure when viewed as a comprehensive formula… thus justifying the practice-based quality of this research. For even if it were possible for someone to stand alongside the author during every moment of the creation of a composition, it would be impossible to ascertain the source, if even the existence, of every decision made during its course. Yet we are able to gain a sense as to the tendencies and motivations behind the work. We can see aesthetic biases along with influences based upon technological limitations, competence, and numerous other factors, all of which come into play. It is these influences and the feedback between the various activities that create a composition. It is said that no human face is perfectly symmetrical. The symmetrically related imperfections are the characteristics that make each of us who we are. Certainly, in the artistic and even in the technical arenas, it is embracing and developing these apparent limitations that set us free to create authentically rich expressions of who we are.

Overall, in this chapter we have endeavored to delineate the interdisciplinarity of the work involved in this research. For it should be readily recognizable that each and every aspect of the creative process is a complex and powerful matrix that, when intricately interlaced transparently with the others of equal complexity and power, provide a nexus of creative expression unique to the artist that created it and to the results of the endeavor.

The endeavor here was focused upon numerous stratifications of intricately interwoven methodological ontologies created by praxis in the areas of generative audio algorithms, aleatoric visual production procedures, high-performance computing implementation, and the application of the composer’s sensibilities as each was extended toward the culmination of processes in the conclusion of a completed composition. It is such aspects as kinetic and timbral ontological relationships, the search for and awareness of emergent quanta and complexity, the extensive hybridization of process, and novel aesthetical considerations that come to bear on the final outcome of a unique and multidimensional spaciotemporal experience. For it is in the creation of non-linearity in creative expression that transcends our linear Newtonian perspectives and opens us to those that go beyond them. Developing these disparate, yet enjoined, constituents to the greatest feasible extent that fosters the experiential counterpoint of unique intermedial expression.

CHAPTER 6. PRESENTATION

6.1 Chapter Overview

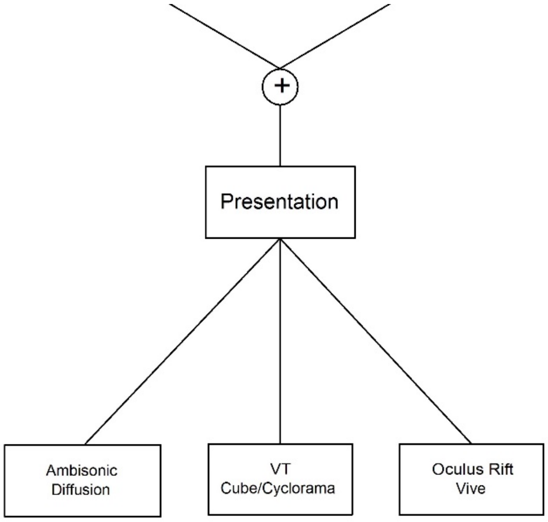

Figure 92 – Presentation of the Final Set of Compositions

To reiterate, a holochoric approach to visual music composition, intended for diffusion and projection in 3D/360 venues, creates an evolutionary progression of counterpoint in the field of visual music composition that approaches a multidimensional spatiotemporal experience. The intricately interwoven methodological ontologies, created, discovered, and developed through this research and creative praxis, pose a confluence of process that facilitates unique aesthetical and formal stratifications of structure in support of this approach. Therefore, the question remains as to the means of their public presentation. In this chapter several options will be examined. Each is accompanied by both positive and negative attributes. Weighing the benefits against the detriments, several options will be addressed.

Ultimately, since this type of art is currently on the far side of the cutting edge, options are limited and quality is a constant concern. It is reasonable to surmise that this will not remain the case for long. The human race is currently experiencing exponentially increasing technological growth and there is no end in sight (Kurzweil). Forward thinking individuals, organizations, and corporations, are making tremendous headway in the area of holography and holophony. It is reasonable to envision a day when going to a theater will involve being completely immersed in the movie rather than seeing it projected onto a screen in front of the viewer as it is now. We are at the inception of a medium that will soon be ubiquitous. However, for now while we are laying the groundwork for this future, we must make use of the technology currently at hand.

6.2 Site-Specific Cube/Cyclorama

The audio system and the visual system called the "Cyclorama", in the Cube, located in the Moss Art Center at Virginia Tech, are the basis of the site-specific nature of this research project. Its features are extremely unique in offering the ability to project holochoric material. Though the medium for this work has since branched into several areas, the Cube was the initial inspiration for this research project and is considered the ideal venue for it.

The High-Density Loudspeaker Array (HDLA) audio system in the Cube is one of few of its kind currently in existence. It consists of ~144 loudspeakers and several sub-woofers arranged in a three-tier system including additional floor and ceiling levels. Figure 93 provides a 2D representation of the basic layout of the loudspeaker arrangement. As such it provides a high-definition spatialization environment.

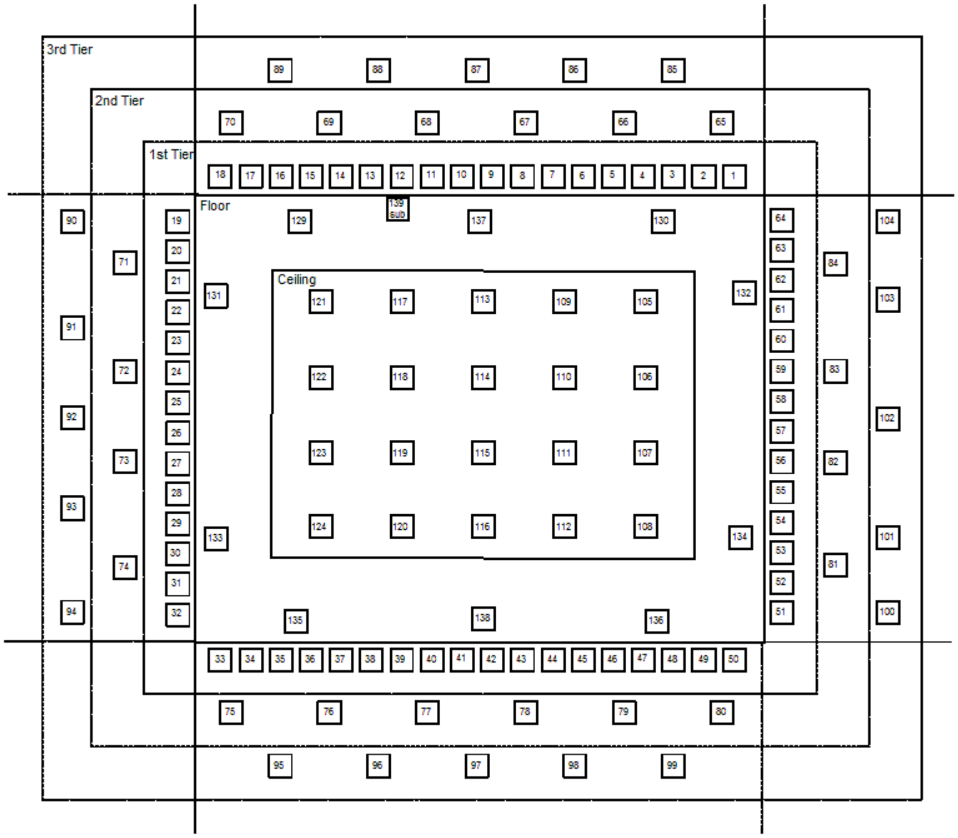

Figure 93 – Cube Speaker Layout – Three Tiers Plus Floor and Ceiling

The Cyclorama (aka Cyc) is a cylindrical projection screen with two projectors facing each of two halves of the screen. The images on each set of projectors are stitched into a single image that is projected onto the screen. The screen is thirty-eight feet in diameter and sixteen feet tall. It is separated in the center of each side to allow audience access. The material of which the screen consists is supposedly 100% permeable to sound although there is a slight high frequency attenuation. A company named the Elumenati designed the Cyc hardware and software, the latter of which they are constantly improving. Figure 94 illustrates the relation of the Cyc, the circle in the center of the Cube, and the sound system.

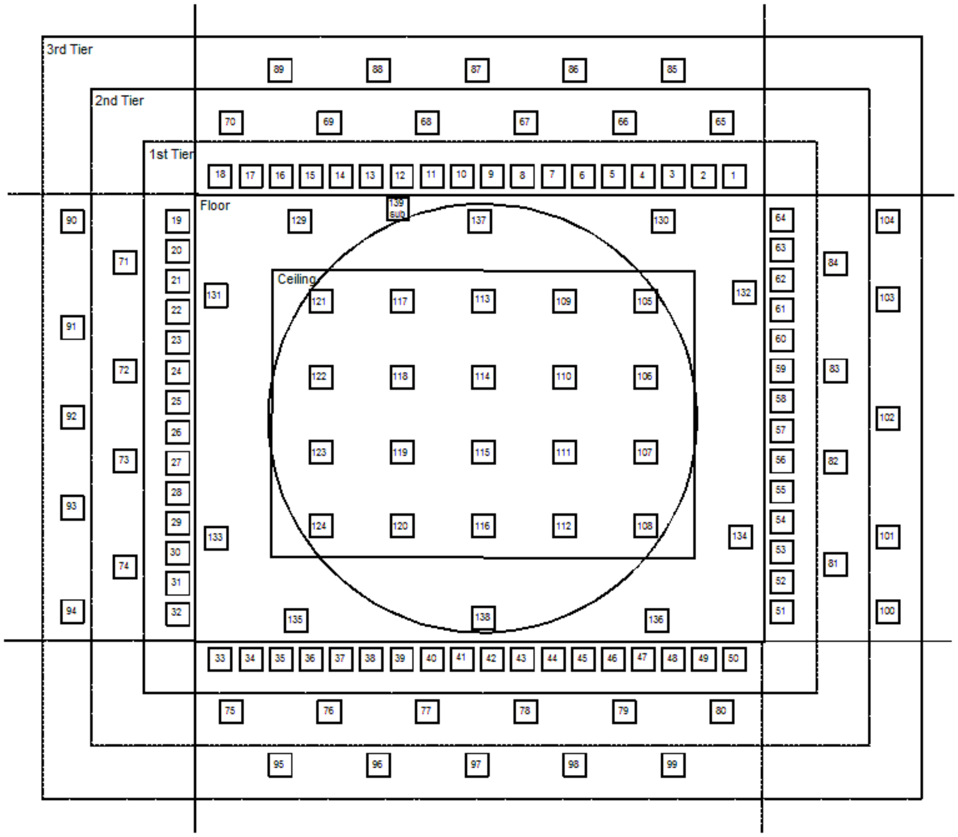

Figure 94 – Cube Cyc/Speaker Layout

There are two computer workstations involved in using the audio/visual system in the Cube. A windows workstation runs Worldviewer, which is the program developed by the Elumenati to project material onto the Cyc. It also consists of programs for operating the projectors and for calibrating the Cyc. This workstation also has an instance of MAX installed. There is a MAX patch that can be run as the OSC protocol that listens for OSC commands from the audio workstation, described next, which then pipes them to Worldviewer. In this manner synchronization between the sound system and the video system is achieved.

The sound system computer workstation implements Dante, an audio networking system, which allows the interconnection of all audio inputs and outputs. This workstation consists of an instance of MAX that allows the MAX patch detailed in Chapter 2 to be instantiated. When the "Play" command is instantiated on the MAX patch an OSC "Play Movie" command is sent via the network and the listener on the video workstation hears it and responds by playing the chapter and scene referenced in the command. It is sometimes necessary, due to latencies inherent to the systems, to adjust command delays in order to start the two systems in sync.

The audience must don a pair of actively polarized glasses to view stereoscopic videos on the Cyc. The projectors send a Radio Frequency (RF) signal to the glasses causing the right eye to be blacked out when the left eye image is being projected and vice versa. This alternation occurs at sixty times per second, which is rapid enough to convince the brain it perceives the images simultaneously. This is manner in which the stereoscopic images are seen as three-dimensional. The processes for creating the video mentioned in Chapter 3 here reach fruition and the perception of holography is achieved.

There is a similar phenomenon that occurs with the sound system. The MAX patch, using the Ambisonics tools mentioned in Chapter 2, projects the thirty-two channels of sound into thirty-two Ambisonic spherical sound fields, which are tuned to the Cube's architecture. This is accomplished by bringing the sound fields toward the center of the room for specifically increased presence and by carefully balancing each of the five levels of speakers. Here this approach, in addition to the compositional approach delineated in Chapter 2 reaches fruition and the perception of holophony is achieved.

The Cube, with the Cyc installed, can comfortably seat twenty audience members on swivel chairs. Considering that the set of compositions is rather lengthy and that the audience members could easily become disoriented and lose equilibrium during the presentation of the material, seating the audience is the option most conducive to both safety and comfort. Though an installation of the Antitheses project consists of all seven compositions, audience members are allowed to enter and exit the Cube at will. Some prefer to experience the content through multiple viewings, while others leave quickly after experiencing a feeling of disembodiment that might occur.

The Cube is a phenomenal space in which to experience the results of this research. However, not everyone can visit Blacksburg, Virginia to experience it. Therefore, other modes of presentation are available.

6.3 Head Mounted Displays

The various currently available versions of head mounted displays (HMD) offered by Oculus, VIVE, and Valve Index afford an aspect of portability to the projects created through this research and creative praxis. Using these head mounted displays and applications such as GoPro VR Player, the viewer can experience an audio-restricted version of the compositions in the comfort of their own home.

Steps to viewing the compositions using an HMD:

1. Copy the mp4 version of the composition into the directory of their choice,
2. Instantiate the Oculus or Steam programs
3. Configure and instantiate GoPro VR Player
4. Select the file to play and then watch/listen

This is perhaps the most straightforward manner in which to experience the visual music compositions. There is however a tradeoff in that the music is heard in a stereo format instead of as a 32-channel expression. This dynamically impacts the immersive and unified quality of the experience.

Figure 95 – Photograph of the Oculus Rift HMD

6.4 Virtual Theater

The virtual theater is a relatively new paradigm being used in major film festivals and conferences around the world such as the Tribeca Film Festival, Sundance Film Festival, and SIGGRAPH to name a few. For them, numerous HMDs are connected to a common computer and then the compositions are played synchronously on each. The audience has a more standard cinema experience in that they view the movie together.

This type of presentation of the material is well-suited, for instance, in the Cube where the sound system can be used and then the visual material is viewed in the HMD. Using the two in conjunction is a blending of the best of both worlds. This is also an excellent manner in which to create an installation in a museum or other public spaces.

In order to test it, the author obtained 10 – Oculus Go (wireless) head mounted displays and the Showtime VR (STVR) application. The latter registers HMDs that share a common local area network with the server. Each video that is loaded onto each of the Go units is available for video playback via STVR. In this manner each viewer sees the same video at the same time.

In order to sync the video with the audio for projection on a HDLA, I created an additional patcher object in the main MAX patch that listens for a sequence of tones from an HMD that has its audio output plugged into a microphone input in the audio interface. When a video is instantiated in STVR a tone associated with it begins. For instance, if there is a single tone that is sounded at the beginning, MAX hears it and begins the audio playback for the associated musical aspect in sync with the video. The syncing is accomplished using playback delays.

This system works very well. It provides portability of the compositions and maintains the cinematic, social, aspect of viewing art or film. It promises to be a viable venue for the presentation of the material created for this project.

6.5 Virtual Theater - Live Streaming

There is also a need for the ability to stream the holochoric videos in order for them to be experienced remotely. Using a similar approach to the virtual theater, the server may be located on a virtual private network (VPN) or accessible via port forwarding. In either case the HMDs are registered with the STVR server remotely, which then instantiates the videos that are located on a streaming server. Currently, one viable service for this is Microsoft's AZURE. Below is a proposed diagram of this configuration, which has not, as of this writing, been tested.

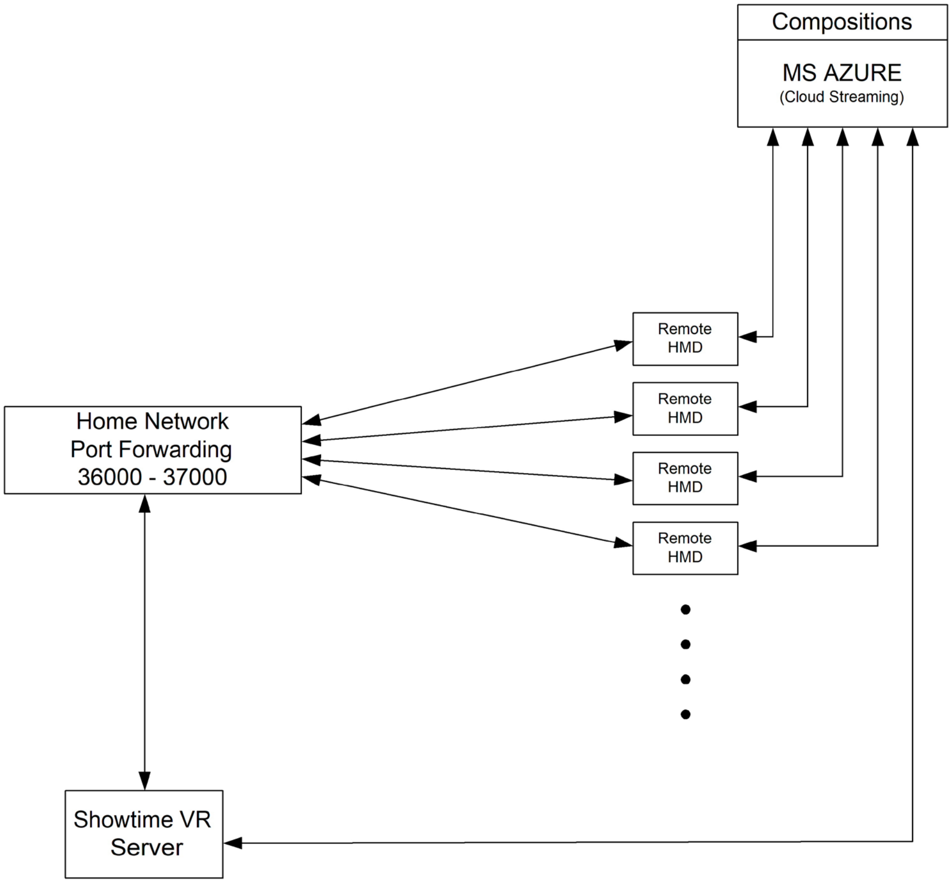

Figure 96 – Flowchart of the Data Streaming Paradigm

6.6 Planetariums

Planetariums are increasing in popularity for their ability to provide highly immersive visual experiences and their relatively wide availability. However, this paradigm limits much of the power of the compositions since there are very few planetariums that offer stereoscopic viewing and most tend to have, at best, a 5.1 sound system. Nonetheless, they do afford an opportunity for the essences of the compositions to be presented publicly and on a huge dome screen that surrounds the audience.

Though I have rendered a few full dome versions of some of the video sections of this research, further testing is required to ascertain its viability as a presentation option. One caveat is that the to produce a full dome version of the compositions each section of video would need to be re-rendered specifically for it. The rendering requirements are not as extensive as those for high-resolution stereoscopic images and so require less time and CPU power.

6.7 2D

The final method of public presentation, and certainly the most accessible, is that of 2D video. Most pertinent conferences are set up to be able to present 2D video on a large screen and usually provide at least 8 channels of audio. I have continued to render 2D versions of each composition for this purpose. It is a fairly straightforward process that includes importing the stacked 7000x7000 video into AE and then masking off the largest 16:9 ratio area available and then exporting it as a 1920 x 1080, which is currently the standard 2D video format.

Interestingly, since I began creating holochoric compositions, the 2D videos rendered from them belie a significantly influence on my stereoscopic work due to the simplified format. The 2D versions usually stand on their own very well and they offer an increase in perceived depth as compared with 2D compositions realized before working with holochory.

The 2D versions of these projects are available to a much broader audience though the limitations are obvious. In many ways it is similar to looking at a photo of a sculpture as opposed to the sculpture itself.

6.8 Pros and Cons

The Cube/Cyc combination offers several positive attributes. The polarized glasses are much lighter than the HMD devices, the audio environment is second to none, and each viewer can see the others in the venue and where they themselves are located in relation to the space. The latter is highly desirable. Some of the downsides include a 40% reduction in color and brightness due to the polarized glasses and the properties of the projectors. Also, the venue is limited to viewers who are within local proximity of the Cube.

The HMD devices provide a very vivid, color rich and bright, video display. One can experience the compositions within the comfort of their home and can control which compositions they watch and when. The downsides are that the HMDs are heavy to wear for extended periods of times and they are prone to video glaring effects. The audio is limited to restricted to stereo, unless viewed in a VR theater environment that implements a HDLA system. Lastly, the video is limited to resolution of the particular device.

The virtual theater provides a communal cinematic experience and allows for the potential of a HDLA sound system when used in conjunction with a venue such as the Cube. There is a hygiene concern with regard to multiple viewers using the same device over multiple showings. This can be counteracted by frequent cleanings and by using disposable face cover head masks. The latter, however, increases the awkwardness of wearing the HMD.

The main disadvantage of the streamed virtual theater over the local one is that the musical aspects of the visual music compositions can only be heard in stereo via the headphone on the listener's HMD. Otherwise it is quite viable and provides a potential means with which to reach a more broadly distributed audience.

There are advantages to the planetarium dissemination of the material. Planetariums are widely available and planetarium managers are always looking for quality material to present. They provide an immersive environment. On the downside, planetariums seldom have HDLA sound systems, thus limiting the spatial audio experience. Similarly, they are seldom offer stereoscopic projection. Also, in order to adapt the compositions for a planetarium a full dome rendering of the composition must be accomplished.

Lastly, the 2D video arena provides a wide array of possibility. Most concert, conference, and festival venues offer some level of multi-channel sound system implementation and usually can provide a large 2D projection system. There are numerous quality opportunities to submit work and they often include opportunities for paper presentation. The obvious downside is the lack of 3D/360 stereoscopic projection systems and often though the sound systems are multi-channel they are seldom set up in a cuboidal or spherical manner enabling holophony. This limits the immersive quality of the compositions to a reasonable imitation of the original work.

As can be seen there are options available for dissemination of the work produced by this research and subsequent compositions that will follow. As technology continues to increase in quality and capability, more and more of these types of venues will become available. It is clear that this is a future direction of visual art. This is evidenced in the gaming industry that is seeing a huge increase in the use of HMDs for gaming. Gaming producers have become increasingly interested in creating not only a holographic visual environment for their games but also a similar approach to the audio aspects. Recent developments in Microsoft's "Project Acoustics" demonstrates the focus upon such sound systems and, in general, increasing the flexibility and extent of control over the virtual word experience (Stack). As this area of development continues to grow and expand, new forms of artistic expression will emerge. This work is intended as a contribution toward this imminent future.

6.9 Chapter Conclusions

As portrayed in this chapter, venues capable of diffusion and projection of 3D/360-degree of holophons and (pseudo) holograms are essential toward their exposure as an evolutionary and unique progression of counterpoint in the field of visual music. Though options are currently limited in this novel artistic medium, the right venue, whether VR theaters engaging head-mounted displays or proprietary frameless projection systems that include a HDLA, provides the opportunity to create a multidimensional spatiotemporal experience. Through such an environment, a novel counterpoint of intermedial foci is given substance, which provides experiential evidence of an expanded view of reality. Intricately interwoven methodological ontologies, involving the interplay between each contributing constituent, facilitate unique aesthetical and formal stratifications of structures that are intended to raise the level of consciousness of the audience through exposure to perspectives that are unlikely to be found elsewhere.

CHAPTER 7. CONCLUSIONS

7.1 Conclusions

At this point it should be clear that, pairing holophons and pseudo holograms, diffused and projected in 3D/360 venues, creates an evolutionary progression of counterpoint in the field of visual music composition that approaches a multidimensional spatiotemporal experience. Intricately interwoven methodological ontologies involving generative audio algorithms, aleatoric visual procedures, and high-performance computing techniques pose a confluence of process that facilitates unique aesthetical and formal stratifications of structure in support of this approach. Further, the interplay between the composer's sensibilities, audio and visual indeterminacy, and high-performance computing, necessitates a novel quadrilateral and reciprocative feedback loop actively engaging each constituent. Due in part to lengthy rendering cycles, this work expands the traditional tone row methodology of musical composition, wherein thematic elements are constrained and then developed to the greatest feasible extent, to foster a novel counterpoint of intermedial foci.

The complex matrices of processes involved in the creation of these holochoric visual music compositions are rich with creative opportunity as well as fraught with creative challenges. Yet, these challenges and opportunities can act as the fuel driving a long-term endeavor such as the creation of the "Antitheses" project. As demonstrated in this document as well as in experiencing the compositions, the project required over two years of planning, building, configuring, rendering, editing, creating, listening, seeing, experiencing, intuiting, and much more before approaching fruition. As an artist the continuous stream of the results of this investment are worthy of the time, energy, dedication, and love required. The pursuit of the state of mind required to move in this manner becomes an appreciated lifestyle.

It was surprising how much was learned throughout this project. As compared to the initial work, that being produced now is obvious proof of this assessment. It is said that we learn by correction, not by perfection, which is extremely apropos here. It was only through the praxis of creating these visual music compositions that questions, mistakes, and challenges have arisen and forced solutions in order to proceed… and there have been many of each! In working with the visual, the auditory, the high-performance computing, and, perhaps most importantly, with creating their symbiosis, a great expanse was traversed. Additionally, great vistas lie ahead… the journey has only just begun.

What was documented here are, primarily, the results of this study… the processes that eventually formed a unified approach to the composition of the material created. It is clear that the refinement of process and attitude over the course of the first ~15 months of the project was extensive. That which now seems simple, such as the most effective resolution for rendering the images in Maya, or how to integrate convolution reverb into the Csound orchestra were major obstacles in the beginning. However, first conceiving of what might successfully contribute toward holochory, for instance, and then testing, evaluating, and ultimately implementing it continues as an extensive and protruded process. Further refinement, that occurs during the process of creating each new scene, is constantly occurring and will continue as an ongoing process.

The contemporary multimedia artist must be ready, willing, and even excited to acquire a modicum of expertise in numerous diverse and hybridized areas of endeavor. For it is their confluence that summons heretofore unimaginable creative expressions into being. Technologically, socially, psychologically, and personally, we human beings are currently advancing at an exponential rate (Kurzweil). Our capabilities in many areas have perhaps far outpaced their implementation. It is hoped that this document provides a humble example of a fraction of that which is currently within our grasp and that it inspires others to extend beyond their imagined limits.

In refining any creative process, the goal is in achieving its transparency. Though developing, creating, and apprehending the required methodologies is obviously necessary, it is only when this knowledge becomes second nature that creativity can flourish. In the throes of the process, while in the timeless state of mind that engenders creativity, any technical diversion is not only unwelcome but also counterproductive. Maintaining an efficient and highly functional working environment is an absolute necessity. In my experience, the act of creating, which is an actively engaged, vital, and quiet mind, engenders diving ever more deeply into the depths of my being. This can only be viewed as a journey. Ultimately, it is this journey that shines in comparison to most anything in the world of form. As such, the resultant composition is a byproduct… or perhaps a souvenir. In this case the journey is the destination.

One last thought is that this document is not intended as a recipe. If the process could be exactly reproduced the results would nonetheless be astoundingly diverse. Such is the open-ended nature of the process. Instead, it is intended to document a personal pathway through an unexplored territory and to leave signposts as potential directional markers. Certainly, there are techniques and pipelines described herein that may be of some benefit to those following in a similar direction. Yet, here it is proposed that forging one’s own path through a previously unexplored forest that is the most fascinating journey imaginable. The destination is simultaneously a point of arrival and a point of departure. For though the past may inform us and the future may inspire us, it is the present that empowers us.

GLOSSARY OF TERMS

The definitions that follow are derived from the understandings achieved through the research described in this document. Though many may not be strictly textbook definitions, this glossary is intended to define terms used in this document in the context of the research in order to augment clarity of intent.

Aleatory. Forms of artistic expression derived in some manner from quasi-random processes.

Ambisonics. A spatialization system that uses a weighted sum of all loudspeakers in a sound system to create sound fields in specified locations. It allows for specific localization of sound objects in a specific listening environment.

Antitheses. In this study antitheses are perspectival views, heightened states of awareness, that are diametrically opposed to our normal everyday states of awareness. They go against the grain of commonly accepted perspectives and act as openings into non-linear and yet corporeal levels of consciousness.

Beowulf Cluster. A group of three or more commercial-grade computers configured to work as one. They consist of a head node, that directs the computational process, and two or more compute nodes that carry out the computational tasks.

Binocular. The use of two visual images to produce the perception of depth.

Central Processing Unit (CPU). The main computational engine in a computer through which all computational tasks are accomplished. Many CPUs are subdivided into several separate cores in order to provide parallel processing, which increases the computational breadth of a single processor.

Choreography. The pre-determined, or not, movements of dancers. In this study it is the relationships between audio and visual objects from a movement perspective.

Complexity. A structural organization of which the origins of the constituent patterns are not obvious. This might often involve numerous constituents as opposed to few. It could be characterized by the evolution of simple algorithms that, through reiteration, evolve in a manner that leaves their source nearly incomprehensible.

Composition. A completed work of art with constituents organized in some manner. Though often thought of from a musical perspective, in this study it is considered more broadly visual art, musical art, and computer science. Also, the act of creating a work of such organized constituents.

Convolution. A mathematical process whereby one function is mapped onto another. In this study, the reverberation signature of an impulse response, recorded in a particularly reverberant space, is mapped onto a second sound sample. This applies very realistic reverberation characteristics to the sound sample it is applied to. Implementing this technique, a very realistic distance algorithm may be achieved.

Domain Name Service (DNS). A server that translates computer network names to computer network addresses upon request. A web browser commonly asks a DNS server for the TCP IP address for a given domain name in order to connect to a web server.

Dynamic Host Configuration Protocol (DHCP). A network system that automatically assigns network addresses to hosts that request it. There are numerous other functions that DHCP can provide however, this is the one pertinent to this study.

Extended Reality (XR). An umbrella under which the terms Virtual Reality, Augmented Reality, and Mixed Reality are included.

Head Mounted Display (HMD). A device worn on the head, covering the eyes, that consists of two projection lenses, one for each eye. In this case the left- and right-eye images from a stereoscopic video are presented in the appropriate lens, turning off and on in an alternating fashion, thus providing a pseudo binocular image to the brain. Though often used for gaming applications, here HMDs are used as mediums through which the stereoscopic visual music videos may be presented.

High-Density Loudspeaker Array (HDLA). A large-scale sound system that is precisely designed to project sound in a performance venue. It often consists of 40 or more loudspeakers.

High-Performance Computing. See Super-Computing

Holochory. Any medium that involves wave interactions. In this study the term is used in a broad sense to include three-dimensional auditory and visual objects that appear to occupy the same three-dimensional space as the experiencer.

Holography. Three-dimensional visual objects created by wave interference patterns of light. In this study, the term is generalized to include any virtual three-dimensionally perceived visual object that occupies the same three-dimensional space as the viewer.

Holophony. Three-dimensional auditory objects created by wave interference patterns of air. In this study, the term is generalized to include any virtual three-dimensionally perceived auditory object that occupies the same three-dimensional space as the listener.

Indeterminacy. Circumstances created or found that are conducive to unanticipated outcomes.

Immateriality. That which lacks a fixation on the physical and the logical.

Interaxial Separation. The distance between the left- and right-eye cameras, which is the same as the distance between the eyes in stereoscopic imagining. This distance is crucial, in tandem with the object sizes and their distances from the camera, for achieving viable stereoscopic imaging.

Local Area Network (LAN). A computer networking system, based upon one of several addressing schemes, that allows several computers to exchange data.

Open Sound Control (OSC). A protocol for data exchange between computers and audio applications and equipment. It is a very efficient and accurate system that occurs over a local area network.

Practice-Based Research (PBR). A form of scholarly research that considers an artifact produced by novel processes to be the evidence of the research. Often a written document forms a contextualization of the artifact and the processes used in its production.

Practice-Led Research (PLR). A form of scholarly research that considers the novel processes involved in the production of an artifact be the evidence of the research. For it, the written dissertation documents the research and an artifact may of not be included.

Praxis. Practice as distinguished from theory.

PXE Boot. In this study, the PXE methodology allows a server, such as a cluster compute node, to obtain a boot image from a server that is broadcasting it. Here it allows the head node of a cluster to push a hard drive image to a new compute-node that is joining the cluster.

Quasi-random. Randomness is thought to be impossible. Therefore, it could instead be described as a phenomenon that consists of a pattern that is not apprehended… thus the term quasi-random.

Reciprocation. A back and forth exchange between one of more constituents of a system.

Schedulers. In computer science schedulers direct computational tasks in an optimal manner. This is one of the main functions of the head node of a cluster or high-performance computing system.

Secure File Transfer Protocol (SFTP). Allows for file transfer to and from a remote computer via console commands.

Secure Shell (SSH). A protocol enabling a secure remote console session. Using SSH one can log into to one computer remotely to another and operate it as if working directly on it.

Stereopsis. The perception of distance or depth using two ears or two eyes as understood and calculated by the brain.

Stereoscopic. Two visual images, one depicting the angular perspective of the right eye and the other from the angular perspective of the left, that are combined to provide the perception of a depth dimension.

Super-Computing. A computing system that performs at the highest rate currently possible for computers. Today that usually consists of parallel computing using multiple CPUs or multiple CPU cores.

TCP IP. A set of network protocols that allow computers to share a network. The addressing schema, unique to a particular group of clients, offers shared address ranges that separate computers on one address range from those on others, thus restricting access to computers on specific ranges.

Vector-Based Amplitude Panning (VBAP). An inverse triangulation methodology that uses three speakers to specifically place the perception of the location of a sound between them using amplitude modulation.

Virtual Reality (VR). Artificially derived experiences that are usually created by digital means.

Visual Music. A somewhat ambiguous term that includes compositions in which light is experienced as sound and/or sound begin experienced as light. In this study the term includes light and sound being used as a singular compositional medium.

VR Theater. A performance of a movie that employs a head-mounted display as opposed to a typical projection screen.

Wave Field Synthesis. A spatialization technique involving a large number of loudspeakers to specifically create wavefront interaction to create the perception of a sound existing in specific locations.

Zero Parallax. The distance from a stereoscopic camera rig to an object at which the angular overlap between left- and right-eye images converge.

REFERENCES

"3.3. Rocks 6: Install and Configure Your Frontend." *Rocks 6: Install and Configure Your Frontend*, http://central-7-0-x86-64.rocksclusters.org/roll-documentation/base/7.0/install-frontend.html.

Abbado, Adriano. *Visual Music Masters: Abstract Explorations: History and Contemporary Research*. Skira, 2017.

Abbasi, et al. "Studying the Recent Improvements in Holograms for Three-Dimensional Display." *International Journal of Optics*, Hindawi, 28 Aug. 2014, www.hindawi.com/journals/ijo/2014/519012/.

Barrera, Taurin (2014). *Digital Synesthesia: Exploring Interactive Electroacoustic Environments*. Unpublished manuscript.

Barrett, Natasha. "A Musical Journey towards Permanent High-Density Loudspeaker Arrays." *Computer Music Journal*, vol. 40, no. 4, 2016, pp. 35–46., doi:10.1162/comj_a_00381.

Barrett, N. "Ambisonics spatialisation and spatial ontology in an acousmatic context". In *Proc. from the Electroacoustic Music Studios conference.* http://www.ems-network.org/ems10/ 2010.

Barrett, Natasha. "Spatio-Musical Composition Strategies." *Organised Sound*, vol. 7, no. 3, 2002, pp. 313–323., doi:10.1017/S1355771802003114.

Barrett, Natasha. "The Perception, Evaluation and Creative Application of High Order Ambisonics in Contemporary Music Practice." Ircam Composer in Research Report 2012.

Benyon, Margaret. "On the Second Decade of Holography as Art and My Recent Holograms." *Leonardo*, vol. 15, no. 2, 1982, p. 89., doi:10.2307/1574541.

Bolt, B. (2006) Materializing pedagogies. Working Papers in Art and Design 4 Retrieved May 2019 from URL http://sitem.herts.ac.uk/artdes_research/ papers/wpades/ vol4/bbfull.html ISSN 1466-4917.

"Book Reviews." *Music, Sound, and the Moving Image*, vol. 4, no. 1, 2010, pp. 109–129. doi:10.3828/msmi.2010.6.

Boulanger, Richard. *The Csound Book: Perspectives in Software Synthesis, Sound Design, Signal Processing and Programming*. MIT Press, 2000.

Cabrera, Andrés, et al. "The Evolution of Spatial Audio in the AlloSphere." *Computer Music Journal*, vol. 40, no. 4, 2016, pp. 47–61., doi:10.1162/comj_a_00382.

Candy, L. (2006). *Practice-Based Research: A Guide*. Unpublished material.

Candy, L. and Edmonds, E. Practice-Based Research in the Creative Arts: Foundations and Futures from the Front Line. *Leonardo 51*, 1 (2018), 63–69.

Candy, L. and Edmonds, E. The Role of the Artefact and Frameworks for Practice-Based Research. *The*

Routledge Companion to Research in the Arts.

Davies, Char (1998). *Changing Space: Virtual Reality as an Arena of Embodied Being.*, www.immersence.com/publications/char/1998-CD-Virtual_Dimension.html.

Carpentier, Thibaut, et al. "Holophonic Sound in IRCAMs Concert Hall: Technological and Aesthetic Practices." *Computer Music Journal*, vol. 40, no. 4, 2016, pp. 14–34., doi:10.1162/comj_a_00383.

Chaudhari, Aunsh, et al. "Transforming the World Using Holograms." *International Journal of Computer Applications*, vol. 130, no. 1, 2015, pp. 30–32., doi:10.5120/ijca2015906867.

Chion, Michel, and Ben Brewster. "Quiet Revolution... And Rigid Stagnation." *October*, vol. 58, 1991, p. 69., doi:10.2307/778798.

Clarke, Rachel, and Claudia Hart. "The Real-Fake." *Metaverse Creativity*, vol. 3, no. 1, 2013, pp. 28–39., doi:10.1386/mvcr.3.1-2.28_1.

Cope, Nick, and Tim Howle. "Electroacoustic Moviemaking: a Creative Media Practice Research Collaboration Case Study." *Studies in Australasian Cinema*, vol. 12, no. 2-3, 2018, pp. 116–134., doi:10.1080/17503175.2018.1539539.

Csound Manual - Frames Version, www.csounds.com/manual/html/indexframes.html.

Derakhshani, Dariush. *Introducing Autodesk Maya 2014*. Sybex, 2013.

Dewitt, Tom. "visual music: Searching for an Aesthetic." *Leonardo*, vol. 20, no. 2, 1987, p. 115., doi:10.2307/1578326.

Edmonds, Ernest, and Linda Candy. "Relating Theory, Practice and Evaluation in Practitioner Research." *Leonardo*, vol. 43, no. 5, 2010, pp. 470–476., doi:10.1162/leon_a_00040.

Einstein, Albert. "I Think 99 Times and Find Nothing. I Stop Thinking, Swim in... at QuoteTab." *QuoteTab*, www.quotetab.com/quote/by-albert-einstein/i-think-99-times-and-find-nothing-i-stop-thinking-swim-in-the-silence-and-the.

Eno, Brian (2015). *The Ship*. Unpublished material.

Eno, Brian (2018). *The Working Backwards*. Unpublished material.

Evans, Brian. "Computer Music Journal-Volume 29, Number 4, Winter 2005." *Project MUSE*, muse.jhu.edu/issue/10291.

"Film, a Sound Art." *Choice Reviews Online*, vol. 47, no. 04, 2009, doi:10.5860/choice.47-1896.

Foucault, Michel. *The Order of Things: an Archaeology of the Human Sciences*. Routledge, 2010.

Garavaglia, Javier Alejandro. "Creating Multiple Spatial Settings with 'Granular Spatialisation' in the High-Density Loudspeaker Array of the Cube Concert Hall." *Computer Music Journal*, vol. 40, no. 4, 2016, pp. 79–90., doi:10.1162/comj_a_00384.

Gaver, William. "What Should We Expect from Research through Design?" *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems - CHI 12*, 2012, doi:10.1145/2207676.2208538.

Gorglione, Nancy. “Forms of Light: A Personal History in Holography.” *Leonardo*, vol. 25, no. 5, 1992, p. 473., doi:10.2307/1575759.

Greene, Brian. *The Fabric of the Cosmos ;Brian Greene: Space, Time, and the Texture of Reality*. Vintage Books, 2005.

Heintz, Joachim, et al. *CSound: :*Floss Manuals, 2015.

Hollerweger, Florian. *An Introduction to Higher-Order Ambisonic - Iki.fi*. decoy.iki.fi/dsound/ambisonic/motherlode/source/HOA_intro.pdf.

Howe, Alita, “Holophony | A New Sound System.” *Multimedia Information and Technology*, vol. 32, no1, 2006, pp. 16-18.

Illényi, A. “Holophonics: a Spread-out of the Basic Ideas on Holography into Audio-Acoustics.” *Holography*, May 1990, doi:10.1117/12.2283686.

JMM: The Journal of Music and Meaning. *JMM: The Journal of Music and Meaning - Panayiotis A. Kokoras - Towards a Holophonic Musical Texture*, www.musicandmeaning.net/issues/showArticle.php?artID=4.5.

Kanellos, Emmanouil. “Visual Trends in Contemporary visual music Practice.” *Body, Space & Technology*, vol. 17, no. 1, 2018, pp. 22–33., doi:10.16995/bst.294.

Karlsen, Kathleen (2001). *The Color-Music Connection: Philosophical, Aesthetic and Scientific Perspectives.* Unpublished material.

King, Stephen. *Stephen King on Writing*. Scribner, 2000.

Kostuk, Raymond K. *Holography: Principles and Applications*. CRC Press, Taylor & Francis Group, 2019.

Krishnamurti, J., and D. Rajagopal. *Think on These Things*. HarperPerennial, 1997.

Kurzweil, Ray. *The Singularity Is near: When Humans Transcend Biology*. Duckworth, 2016.

Lanier, Lee. *Advanced Maya Texturing and Lighting*. SYBEX A Wiley Brand, 2015.

Laske, Otto E. *Music and Mind: an Artificial Intelligence Perspective*. Computer Music Association, 1982.

Laske, Otto Ernst, and Jerry Tabor. *Otto Laske: Navigating New Musical Horizons*. Greenwood Press, 1999.

Layer, Harold A. “Holographic and Stereoscopic Space: New Research Directions.” *Leonardo*, vol. 22, no. 3/4, 1989, p. 411., doi:10.2307/1575406.

Lazzarini, Victor, and Richard Charles. Boulanger. *The Audio Programming Book*. MIT Press, 2011.

Lyon, Eric, et al. “Genesis of the Cube: The Design and Deployment of an HDLA-Based Performance and Research Facility.” *Computer Music Journal*, vol. 40, no. 4, 2016, pp. 62–78., doi:10.1162/comj_a_00394.

Lund, C and Lund, H (Eds.) 2009 *Audio visual: on visual music and related media*. Stuttgart: Arnoldsche Art Publishers.

Marshall, Kingsley, and Rupert Loydell. “Thinking inside the Box: Brian Eno, Music, Movement and Light.” *Journal of Visual Art Practice*, vol. 16, no. 2, 2016, pp. 104–118., doi:10.1080/14702029.2016.1195073.

Meyer-Eppler, Werner. 1957. "Statistic and Psychologic Problems of Sound", translated by Alexander Goehr. *Die Reihe* 1 ("Electronic Music"): 55–61. Original German edition, 1955, as "Statistische und psychologische Klangprobleme", *Die Reihe* 1 ("Elektronische Musik"): 22–28.

Mitamura, Shunsuke. “Holographic Holography.” *Leonardo*, vol. 22, no. 3/4, 1989, p. 337., doi:10.2307/1575389.

Noor, Ahmed K., and Rifat Aras. “Potential of Multimodal and Multiuser Interaction with Virtual Holography.” *Advances in Engineering Software*, vol. 81, 2015, pp. 1–6., doi:10.1016/j.advengsoft.2014.10.004.

Orgeman, Keely, et al. *Lumia: Thomas Wilfred and the Art of Light*. Yale University Art Gallery, 2018.

Pasqual, A.m., and V. Martin. “Optimal Secondary Source Position in Exterior Spherical Acoustical Holophony.” *Journal of Sound and Vibration*, vol. 331, no. 4, 2012, pp. 785–797., doi:10.1016/j.jsv.2011.10.009.

Pepper, Andrew. “Holographic Space: A Generalised Graphic Definition.” *Leonardo*, vol. 22, no. 3/4, 1989, p. 295., doi:10.2307/1575382.

Peters, Nils, et al. “Current Technologies and Compositional Practices for Spatialization: A Qualitative and Quantitative Analysis.” *Computer Music Journal*, vol. 35, no. 1, 2011, pp. 10–27., doi:10.1162/comj_a_00037.

Preece, et al. *Interaction Design: beyond Human-Computer Interaction*. Wiley, 2019.

Prince, Stephen. “Audio-Vision: Sound on Screen Michel Chion.” *Film Quarterly*, vol. 48, no. 4, 1995, pp. 40–42., doi:10.2307/1213580.

Pulkke, Ville, *“Virtual Sound Source Positioning Using Vector Base Amplitude Panning*.” J.Audio Eng Soc., Vol. 45, No. 6, 1997 June.

Quiller, S. (2002). *Color choices: Making color sense out of color theory*. Watson-Guptill.

“Remote High Performance Visualization of Big Data for Immersive Science.” *High Performance Computing Symposium (HPC 2018)*, 2017, doi:10.22360/springsim.2018.hpc.006.

Researchers develop projection-type holographic 3-D display technology (2016, October 14) retrieved 30 May 2019 from https://phys.org/news/2016-10-projection-type-holographic-d- technology.html

Rhoades, Michael. “Azimuth – Algorithmic Score Synthesis Techniques.” Journal SEAMUS 1st ser. 2005 (2005): 2-11.

Rhoades, Michael. “Hadronized Spectra (The LHC Sonifications).” *CEC | Canadian Electroacoustic Community*, econtact.ca/16_3/rhoades_LHCsonifications.html. 2009.

Rhoades, Michael. “Toward a Contemporary Musical Consciousness – 2013.” *Michael Rhoades' Perception Factory*, 2013, www.perceptionfactory.com/writing/toward-a-contemporary-musical-consciousness-2013/.

Richardson, Martin. “Mixed Media: Holography within Art.” *Leonardo*, vol. 20, no. 3, 1987, p. 251., doi:10.2307/1578168.

Roads, Curtis. *Composing Electronic Music: a New Aesthetic*. Oxford University Press, 2015.

Roads, Curtis. *Computer Music Tutorial*. 1996.

Roads, Curtis. *Microsound*. MIT Press, 2004.

Robinson, William (2009). *Single Command Rendering with Sun Grid Engine.* Unpublished manuscript.

Rudi, Jøran. “Computer Music Video: A Composers Perspective.” *Computer Music Journal*, vol. 29, no. 4, 2005, pp. 36–44., doi:10.1162/014892605775179937.

Ruoppa, Raine. ”*Stereoscopic Distortions as an Artistic Practice.’* 2014, Master’s Thesis, Aalto University, Department of Art.

Shaked, Natan T., et al. “Review of Three-Dimensional Holographic Imaging by Multiple-Viewpoint-Projection Based Methods.” *Applied Optics*, vol. 48, no. 34, 2009, doi:10.1364/ao.48.00h120.

Spierings, Walter, and Ana Barreto. “Dutch Holographic Laboratory.” *Leonardo*, vol. 28, no. 4, 1995, p. 273., doi:10.2307/1576187.

Spors, Sascha, et al. “Spatial Sound With Loudspeakers and Its Perception: A Review of the Current State.” *Proceedings of the IEEE*, vol. 101, no. 9, 2013, pp. 1920–1938., doi:10.1109/jproc.2013.2264784.

Stack, Microsoft Game. “Project Acoustics | Game Developers Conference 2019.” *YouTube*, YouTube, 5 Aug. 2019, www.youtube.com/watch?v=uY4G GUAQIE&feature=youtu.be&list=PLRs2lXTYCDQ3q9WF-4aO2SgCCCL6rSmA6.

Steffen, Jacob H., et al. “Framework of Affordances for Virtual Reality and Augmented Reality.” *Journal of Management Information Systems*, vol. 36, no. 3, July 2019, pp. 683–729. *EBSCOhost*, doi:10.1080/07421222.2019.1628877.

Stravinsky, Igor, et al. *Poetics of Music in the Form of Six Lessons*. Harvard Univ. Pr., 1947.

Sullivan, G.. Lecture by Graeme Sullivan at Symposium 'The Artist as Researcher' 2010, 2/2. Retrieved August 01, 2020, from https://vimeo.com/23635655

Tabor, Jerry. *Otto Laske: Navigating New Musical Horizons*. Greenwood Press, 1999.

“Three and One Chair: an Analytic Investigation into the Holo-Art.” *Taylor & Francis*, www.tandfonline.com/doi/abs/10.1080/14702029.2017.1325078.

Tricart, Celine. “Virtual Reality Filmmaking.” 2017, doi:10.4324/9781315280417.

Tolle, Eckhart. *The Power of Now: a Guide to Spiritual Enlightenment*. Hachette Australia, 2018.

Vonnegut, Kurt. *The Slaughterhouse Five*. Delacorte, 1969.

Watkins, Julie. "Composing visual music: visual music Practice at the Intersection of Technology, Audio-Visual Rhythms and Human Traces." *Body, Space & Technology*, vol. 17, no. 1, 2018, pp. 51–75., doi:10.16995/bst.296.

Whitelaw, Mitchell. "Synesthesia and Cross-Modality in Contemporary Audiovisuals." *The Senses and Society*, vol. 3, no. 3, 2008, pp. 259–276., doi:10.2752/174589308x331314

Wakunami, Koki, et al. "Projection-Type See-through Holographic Three-Dimensional Display." *Nature Communications*, vol. 7, no. 1, 2016, doi:10.1038/ncomms12954.

Walker, Evan Harris. *The Physics of Consciousness: the Quantum Minds and the Meaning of Life*. Basic Books, 2000.

Weinel, Jonathan. *Inner Sound: Altered States of Consciousness in Electronic Music and Audio-Visual Media*. Oxford University Press, 2018.

Wolfram, Stephen. *A New Kind of Science*. Wolfram Media, 2002.

Yu, Ka & Neafus, Dan & Wyatt, Ryan. (2016). Filmmaking for Fulldome: Best Practices and Guidelines for Immersive Cinema (Part I). The Planetarian. 45. 26-32,34,36,38.

Yu, Ka & Neafus, Dan & Wyatt, Ryan. (2017). Filmmaking for Fulldome: Best Practices and Guidelines for Immersive Cinema (Part II). The Planetarian. 46. 25-36.

Zicher3d-org. "zicher3d-Org/Domemaster-Stereo-Shader." *GitHub*, github.com/zicher3d-org/domemaster-stereo-shader/wiki/Maya-Domemaster3D-Tips-and-Tricks.

Curriculum Vitae

Education:

- Virginia Tech Blacksburg, Virginia Fall 2018 – Fall 2020
  Interdisciplinary Doctoral Candidate working in areas of Holophony, Holography, and High-Performance Computing. Practice-Based Research
- Virginia Tech Blacksburg, Virginia 2015 – Spring 2018
  School of Visual Arts - BFA in Creative Technologies/Honors College
  Summa Cum Laude/Honors Scholar/Certificate for Excellence in Undergraduate Research
- Indiana/Purdue University Fort Wayne, Indiana 2004 – 2015
  ~40 hours toward a Computer Science

Grants, Scholarships, Awards:

- Dean & Rosina Carter Honors Scholarship - $1500
- 2 Grants – $500 each for travel to 2016 NYCEMF
- ICAT SEAD Grant - $500 Fall 2016 (PI) for Imagined Realities Concert in the Cube at VT
- Excellence in Undergraduate Research Certificate – Spring 2017
- Honors College Enrichment Grant – $1000 for travel to 2017 NYCEMF
- ACC Innovation and Creativity Grant - $2000 Summer of 2017 for building a Super Computing Cluster for Rendering
- ICAT SEAD Grant - $500 Spring 2018 (PI) for Second Body Awareness Installation in the Cube at VT

Employment:

- Ball State University – Researcher in the Institute for Digital Intermedia Art 2019 - Present
- Budd Group/Novozymes Biologicals Facilities Maintenance Manager 2015 – 2019
- Sweetwater Sound Network Administrator 2001 – 2014
- Adaptive Technologies Network Administrator/Electrical Engineering Technician 1998 - 2001
- Votaw Electric Master Electrician 1986 – 1998

Teaching Experience:

- Guest Lecturer/Teaching for Charles Nichols' Csound class at Virginia Tech - MUS 3066 Computer Music and Multimedia - March 21st and 28th, 2019

Major Lectures:

- International Csound Conference, Berklee College of Music in Boston – 2014
- Toronto Electroacoustic Symposium – 2014
- Harvest Moon Festival, Concordia University Montreal – 2009
- SEAMUS National Conference, San Diego State University – 2004
- MapleCreek Middle School (6 talks) – 2003

Publications:

- Leonardo Music Journal Volume 30, 2020 – Exploring the Nexus of Holography and Holophony in Visual Music Composition
- CEC eContact! 16.3 2014 – Hadronized Spectra: The LHC Sonifications (Interactive version on Rhoades web site)
- Journal SEAMUS 2005 Vol. 18 no.1 – Azimuth: Algorithmic Score Synthesis Techniques (Interactive version on Rhoades web site)

Music Compositions:

- 22 – CD Projects / 150+ Electroacoustic Music Compositions

Oil Paintings:

- 400+ oil paintings

Visual Music Compositions:

- Bending Glass Walls
- Inside the Crimson Castle (An Escher Multiverse)
- Twisted Chambers (Holographic/Holophonic)
- His Dancers Three (Holographic/Holophonic)
- Petal and Branch (Holographic/Holophonic)
- Shards of Oil (Holographic/Holophonic)
- Second Body Awareness (Holographic/Holophonic)
- Ice Cube Windows (projection mapping on Moss Art Center at VT)

- Crystal Verse
- Pixelated Planes
- Cylindrical Dimensions
- Apparitions
- Reflections and Refractions
- Event Horizon
- Voice in the Forest
- Cygnus Loop

Concerts:

- NYCEMF/ICMC 2019 – Held in Manhattan, New York City
- Apparitions – 2018 EaSt Concert Series at Concordia University, Montreal.
- Reflections and Refractions – EMM 2018 at Lewis University, near Chicago.
- Crystal Verse - NYCEMF 2018 held in Manhattan, New York City.
- A Hot Jupiter - Cube Fest 2018 "Sounds in Focus" - Perform Studio at Virginia Tech.
- Apparitions" - Germany at the Festival NEUE MUSIK LÜNEBURG. Elektroakustische Musik aus New York City - 2017
- International Csound Conference 2017 in Uruguay - Second Body Experience
- DISIS Concert Spring 2017 – Supercollider Class Performance
- NYCEMF 2017 in NYC- Apparitions
- Imagined Realities 2016 - VT Cube – One man concert
- DISIS Concert Spring 2016 - Apparitions
- DISIS Concerts Fall 2016 – Pixelated Planes
- Csound International Conference 2016 in St. Petersburg Russia – Sidereal Sonata
- NYCEMF 2016 in NYC – Reflections and Refractions
- SEAMUS Conference 2015 (VT) – Amber Orbs in a White Infinity (3rd mvnt)
- NYCEMF 2015 in NYC – Amber Orbs in a White Infinity
- NYCEMF 2014 in NYC – Immersion
- Csound International Conference 2013 (Berklee College of Music) – Immersion

- Estudio de Música Electroacústica 2013 (Montevideo, Uruguay) – Voices From Another Dimension
- Three Rivers Festival Installation 2013 – The Sights and Sounds of Numbers
- SEAMUS National Conference 2011 - Espace
- SEAMUS National Conference 2011 – Event Horizon
- SEAMUS National Conference 2010 - Organic Mechanism
- Harvest Moon Festival (Montreal, Canada) - Cygnus Loop
- SEAMUS National Conference 2009 - Tunnel Maze
- SEAMUS National Conference 2008 - Anaphasics
- SEAMUS National Conference 2007 - Release!
- SEAMUS National Conference 2007 - Unperceived Dimensions
- The NYU New Music Ensemble 2007 (NYC) - Leading Turbulence
- IMMARTS' TechArt06 Exhibit (Chicago)– Release!
- SEAMUS National Conference 2006 - The Outer Edge of Possibility
- SEAMUS National Conference 2005 - Mental Aperture
- Electric Rainbow Coalition Festival 2003 - Madman's Prayer
- Synthèse 2003 (Bourges, France) - Energies @ Work
- SEAMUS National Conference 2003 - Composition Zeta
- ICMC 2001 (Cuba) - John's Door (door to the john)
- IV National Symposium for Computer Music 2000 (Cordoba, Argentina) - Beckoning the Hunter
- SEAMUS National Conference Y2K - Energies @ Work
- And Numerous Others

Installations and Exhibits:

- Second Body Awareness, a holographic/holophonic visual music installation, was presented in the Cube, in the Moss Art Center at Virginia Tech, as a part of ICAT Day - 2018.
- April 30th, 2018 – Michael's "Ice Cube Windows", a projection mapping composition, was presented as a part of a group exhibit during ICAT day - 2018

- Second Body Awareness, a holographic/holophonic visual music installation, was presented in the Cube, in the Moss Art Center at Virginia Tech - 2018
- Second Body Awareness – Installation in the Cube at VT (Holographic/Holophonic Composition) – ICAT DAY - 2018
- Second Body Awareness – Installation in the Cube at VT (Holographic/Holophonic Composition) – 2018
- Eternal Flame – Projection Mapping Installation in the Perform Studio at VT - 2018
- Virginia in the Raw - Mish Mish Gallery - 2015
- Echoes in the Cedar Canyons Installation – Fort Wayne Museum of Art – 2014
- Horizontal Blue Installation – Sweetwater Sound - 2013
- The Sights and Sounds of Numbers - Three Rivers Festival - 2013
- Artlink Exhibit at the United Arts Building - 2013
- Horizontal Blue Installation – A Taste of the Arts – 2013
- Horizontal Blue Installation – Berklee School of Music in Boston – 2013
- The Voice in the Forest – Lotus Gallery 2012
- And Others

Hobbies and Other Interests:

- Playing Racquetball
- Watching Movies
- Driving High Performance Sports Cars
- Chess, Sudoku, Checkers, Scrabble, Euchre, Rummy

PUBLICATIONS

- Leonardo Music Journal Volume 30, 2020 – Exploring the Nexus of Holography and Holophony in Visual Music Composition
- CEC eContact! 16.3 2013 - Hadronized Spectra: The LHC Sonifications
- Journal SEAMUS Volume 18, No. 1 Spring 2005 – Azimuth: Algorithmic Score Synthesis Techniques