SELF-SIMILARITY AS A MUSICAL TECHNIQUE

DAVID M. ETLINGER

1. Mathematical Background

Self-similarity is a highly general concept with wide-ranging applications, particularly in complexity and chaos theory [1, 2]. It also appears in the apparently far-removed domain of music [3]. Using elements of physics and mathematics as both motivation and quantitative tools, we will explore several kinds of self-similar structures in music as they occur, both in the works of the masters and in smaller etudes composed specifically to demonstrate these techniques.

Before understanding self-similarity in music, it is of course helpful to understand self-similarity in general. Rather than attempting to formulate a precise definition, we will use a more intuitive description: given some sort of object or system (the details depend heavily on the domain we are working in), we say there is *self-similarity* when the structure is the same (up to certain transformations we'll discuss) at different levels or scales. Nonmusical examples abound in nature. For example, consider a mountainous skyline. From far away, there is a certain level of jaggedness determined by the resolution of your eye. As you move closer, you essentially zoom in on a certain segment—this is a *change of scale*. Although certain peaks will move out of the field of vision, new features that were too small to see before will emerge. Even at close distances, rocks and crags will prevent the mountains from appearing smooth. Thus, through a wide range of scales, the mountain range will look essentially the same.

Even in this apparently simplistic example, there are several interesting features. First, note that there is a certain range of scales within which self-similarity holds. From too far away, the mountains will not be discernible at all. Zoom in too far, and we pass from the domain of rocks to the domain of molecules and atoms. These have their own scaling features, but they are not the same as those of the mountain range. This limited range of scales is the *domain* in which self-similarity functions. Next, note that as we move in closer, peaks leave our field of vision at approximately the same rate as smaller features resolve. In other words, the spectrum of scales is *continuous*.[1]

The final point to note from this example is that the self-similarity is not exact. That is, the jaggedness at one scale looks similar to that at another, but it is not precisely the same. To quantitatively determine self-similarity (which we will not attempt to do here), we would need to come up with appropriate tolerances. For example, although the skyline will not be identical at different scales, it seems quite reasonable to think that many statistical properties will in fact be the same at all

[1] For an example of a *discrete* spectrum, consider the layers of an onion. There are only a certain number of layers, which we can index by integers $1 \dots n$, and hence there are a finite number of scales. In practice, most examples from nature have a continuous spectrum. Many musical examples exhibit discrete spectra, so we will discuss this at more length later.

levels—the likelihood that the skyline will change direction at any given point, the slope of each segment, etc. The operations which relate non-identical structures in this way are called *similarity transformations* or *transforms*. Scaling is one such transform; others include *translation* (shifting the structure's coordinates without changing its shape); *reflection* (flipping the structure about an axis); *rotation* (turning the structure about a pivot point); and *shearing* (scaling different dimensions by different amounts) [4]. Many more complicated transforms exist, but we will not try to define them. When we can turn one structure into another by applying a transform, we say they are in the same *equivalence class*.[2] By allowing transforms that apply significant changes, we are using the idea of *fuzzy invariance*: the notion that the similarity can be very loose, as long as it falls within set tolerance limits [5, 6].

There are many other examples of self-similarity in nature. Trees exhibit this property: the tree as a whole looks quite similar to any main branch and its subsidiaries; these smaller branches in turn have even smaller branches, and so on. Shorelines are self-similar in almost the same way as mountain ranges. Snowflakes are a beautiful example of this: looking at any section of the outer boundary reveals smaller features of exactly the same sort as the larger ones. In general, almost all naturally occuring boundaries or areas exhibit some sort of self-similarity.

There are also quite a number of examples from mathematics. Since they are often generated by some sort of algorithm or iterated procedure, they have many features in common with the musical examples we will look at later. Here we will examine two classes of these objects, beginning with the *Julia* and *Mandelbrot sets*. Julia sets are formed from a relatively simple procedure. Consider any function f whose domain is the complex plane.[3] Although any function will technically give a Julia set, interesting results only tend to appear when the function is nonlinear. For each value x of the complex plane, consider the sequence $x, f(x), f(f(x)), f(f(f(x))), \ldots$. If these values grow without bound, then x is not in the Julia set. If they remain within a certain range no matter how many times we iterate the procedure, then x is in the Julia set. Thus we divide the complex plane into two sets. Further distinctions can be made by using different colors to represent different rates of divergence (i.e., how quickly or slowly the results grow for elements not in the Julia set, according to some appropriate measure). Examples of these sets are shown in figures 1 and 2, both with $f(x) = x^2 - \mu$ [7]. The self-similarity of these sets lies on the boundary. No matter how much we zoom in on any given segment, there is still an infinite amount of jaggedness, with the same mathematical properties of any other scale.

A related and even more studied set is the Mandelbrot set, shown in figure 3 [7]. It is formed as follows: let $f(x) = x^2 - u$, where u is a complex number. A

[2]More precisely: suppose we have a type of transformation T that acts on a set $\{a, b, c, \ldots\}$ (for example, we would call scaling a type of transformation, and scaling by 200% a particular example of that type). Write $a \sim b$ if there is some transformation of type T that turns a into b. We say $\sim$ is an *equivalence relation* if the following three properties hold: *reflexivity* ($a \sim a$); *commutativity* ($a \sim b$ implies $b \sim a$); *transitivity* ($a \sim b$ and $b \sim c$ implies $a \sim c$). The *equivalence class* of a, denoted $[a]$, is the set $\{b | a \sim b\}$. Equivalence classes are often used in music theory, particularly 20th century music theories which are based on group theory. The concept of a pitch-class, for instance, is an equivalence class of pitches related by octaves. There are (theoretically) unlimited pitches, but only 12 pitch-classes.

[3]That is, we require $f(x)$ to be defined for all values x of the form $a + bi$ where a and b are real numbers and $i = \sqrt{-1}$.

point a in the complex plane is in the Mandelbrot set when the Julia set generated by $f(x) = x^2 - a$ is connected (that is, the set is one contiguous piece rather than several separate parts). If the Julia set is disconnected, the point is not in the Mandelbrot set. The Mandelbrot set is self-similar in much the same way as the Julia sets, but it has even more interesting properties. For example, there are an infinite number of smaller copies of the original set, and of course each of these copies contains an infinite number of further copies, etc. There is a vast amount of literature on the geometry and mathematics of Mandelbrot and Julia sets, but we will not examine them any further [1].

The other example from mathematics that we will study is the *Cantor set* [8]. This is a relatively simple instance of a much larger class; it will not have as many interesting properties, but it is quite straightforward to understand. The procedure is as follows: consider the interval of real numbers from 0 to 1 inclusive, denoted $[0, 1]$. We begin by removing the open[4] middle third, leaving us with $[0, 1/3] \cup [2/3, 1]$. We then remove the open middle third from each remaining interval. The procedure is repeated an infinite number of times, on each interval that remains from the previous step. The set of numbers that is left after this infinite procedure is the Cantor set. This set has extreme self-similarity in that if we look at any single closed interval that remains at any step, it will eventually look just like the whole thing. For example, after the first step one of the intervals is $[0, 1/3]$. We can view this as a scaled-down version of the interval $[0, 1]$.[5] Since the removal of middle thirds continues an infinite number of times, the interval $[0, 1/3]$ will ultimately be made to look just like the interval $[0, 1]$. One thing to note about this example is that, unlike many of our previous cases, the similarity is exact other than scale (in the case of the mountain range, for instance, the similarity was statistical rather than exact). This is because of the straightforward nature of our algorithm.

2. Self-Similarity in Existing Music

Self-similarity abounds in the canonical works of Western art music, present in different forms at all time periods. However, it becomes immediately clear when examining these works that classifying the self-similar features quantitatively is not at all trivial. This is logical given the historical nature of composition: until the 20^{th} century, mathematics was often not an explicit compositional determiner (with exceptions during the pre-Baroque period, some of which are discussed below). At best (for our purposes), the mathematical regularities of music were byproducts of the prevailing aesthetics of form and clarity; at worst, there was an active aesthetic imperative to disrupt the expected relations and orders by setting up regular patterns and then breaking them. In other words, composition was a qualitative rather than quantitative activity. Consequently, the self-similarity of these works is often "irregular," less-than-exact, or requiring some subjective interpretation to bring it to light. Even when the 20^{th} century saw mathematics reach new prominence in music, it was often applied toward goals other than self-similar structures. Nonetheless, there is more than ample material to work with, as long as we confine

[4] *Open* in this context means the endpoints are not included; *closed* means the endpoints are part of the set.

[5] The fact that $[0, 1/3]$ and $[0, 1]$ have virtually the same properties at different scales is a consequence of the real number axioms that we will not prove here.

ourselves to somewhat (not overly) qualitative descriptions and do not press too hard for the exact transforms at work.

Probably the oldest use of the technique is the *cantus firmus* of medieval and renaissance polyphony. The definition of this term varied somewhat throughout its usage; it originally denoted a given voice over which keyboard players would improvise additional lines. It then came to be synonymous with a Gregorian chant melody. The usage we are concerned with gained prominence in 14^{th} and 15^{th} century English and French music. It referred to a procedure in polyphonic writing whereby a voice, usually the tenor (next-to-lowest), would be given a melody from a preexisting composition, often in note values much longer than the other voices. This melodic material would be used as a basis for the rest of the composition, often integrated in highly ornamented ways. For example, it was quite common for other voices to present the cantus firmus using much shorter note values. In other cases it would be transformed through the addition of ornamental tones, altered rhythmic figuration, or even some manipulation of the melodic intervals. The cantus firmus melody could be repeated many times during a given composition, with as few or as many of these alterations as composers saw fit (assuming that they always followed the established rules of "good" compositional technique). For example, consider the Psalm "Miserere mei Deus" by Josquin dez Prez, for five voices [9]. The cantus firmus is an eight-note melody (unusually, it is not borrowed but original to Josquin). In the first section of the piece, it is presented eight times, once starting on every degree of the mode. In the second section, the melody is reversed, in shorter note values, and again on every degree of the mode. The third section resembles the first. Voices not presenting the cantus firmus often refer to it melodically nonetheless, quoting fragments or prominent contours.

Canons and fugues are old techniques (canonic material can be found dating at least to the 13^{th} century; fugues since the 14^{th}), both widely used (everyone from Bach to Stravinsky has written fugues). Both terms are used with many technical and historical conditions, none of which are relevant here. We can adequately say that *canonic* imitation is when the same material is presented exactly in two or more voices. The usual exception to exact repetition is to allow a scaling of the rhythmic values. *Fugal* imitation is when the imitating voices transpose the original material, do not necessarily follow it exactly, depart at some point and begin new material, etc. In strict fugues, there is a specified pattern of transforms. The material is first presented in the home key (tonic), and then in the key a perfect fifth above (dominant). Further imitations repeat this pattern, which often serves as a basis for looser imitative forms. As with all genres he tackled, J.S. Bach was a master of both canons and fugues. One set of canons in particular stands out: the 14 canons of the Goldberg Variations (not usually included in the Goldberg Variations, they were found in Bach's handwritten copy in 1974) [10]. These canons are all based on the first eight notes of the aria ground which serves as the basis for the rest of the Variations. They increase sequentially in complexity, and culminate in a quadruple canon involving multiple rhythmic scaling. These canons were all written as *puzzle canons*: only one voice was presented, along with clues as to how the other voices should be performed. It was up to the would-be performer to figure out exactly how to realize the different voices; incorrect solutions would violate the rules of counterpoint. The fugue reached its peak in Bach's monumental *Art of Fugue*. This work systematically developed the fugue, increasing in complexity by adding voices, rhythmic diminutions and augmentations, retrogrades, etc. It culminates in an (unfinished) quadruple fugue, presenting four intricate and interwoven themes simultaneously. After this achievement, composers rarely created simply “fugues” or “canons.” Rather, they incorporated these ideas and techniques into other works. Sometimes, a particular section would in fact be a canon or fugue. More frequently, the idea of interweaving the same material into different musical lines was used in a looser fashion, not meeting the technical criteria of fugue or canon but conveying a similar impression. Ultimately, these ideas became subsumed in the much more abstract notion of *motivic parallelism*.

Motivic parallelism refers to the notion that surface melodic fragments can be replicated at deeper structural levels, and indeed serve as the foundation for entire works. This is quite an involved claim; let us examine it more closely. First, we should briefly discuss structural levels. It is universally accepted that, in tonal Western art music, not all notes carry equal importance. Some serve as focal points and signify important moments of tension and release; others are surface tones that merely connect other notes to form more pleasing melodies. And of course there are all degrees in between. Furthermore, it is almost as universally accepted that, just as there are melodies and harmonies at the immediate note-to-note level, there are governing lines and tonal areas at successively “deeper” (i.e., large-scale) levels. These “fundamental” structural features are then elaborated by more active surface motion. Of course, the difficulty comes in determining what is and is not structural, but that is not our present concern. The important thing is that sometimes a prominent melody or contour from the very surface level will reappear as a large-scale line at a deeper level. When this happens, it seems reasonable to conclude that this is both intentional and relevant; it is also reasonable to look for more of the same. If there is a sufficient degree of this self-similarity, then we can justifiably call it a governing element of the piece’s design. In a sense, the piece can be said to be built entirely from and around this motive.

One theorist, Rudolph Reti, developed an entire theory (Motivic or Thematic Analysis) around the assumption that motivic parallelism forms the basis of most great works of tonal music [11]. He essentially does what we have just discussed: looks for apparent connections between surface motifs and structural points, and then attempts to discern the transformations that lead from one to the other. By breaking the music into a series of motivic cells, he can describe later melodic activity as specific transforms of the original motives. Furthermore, he attempts to ascribe these same motives to the governing lines and harmony changes of the piece as a whole. For example, in *Thematic Patterns in Sonatas of Beethoven*, he views the “Pathetique Sonata” as basically stemming from six motivic cells [12]. He can derive all surface motion from these six cells. Furthermore, he makes compelling arguments that suggest that the long-term key plan and deep governing melodic motion both stem from variations of these cells. This is all the more astounding when we recall that the Sonata consists of three movements that, at first glance, are only cursorily related to each other. This is indeed self-similarity, from the scale of a few notes to an entire multi-movement work. Despite this apparent triumph, Reti’s methods have fallen somewhat out of favor. For one, they are too easily manipulated to suit the user’s agenda. There is no limit on the number of cells needed, or what they can consist of (in some analyses, Reti calls a single melodic interval-two notes-a motive), or what the allowable transforms are. The last in particular is troublesome; it allows all sorts of variations and alterations to a motive, such as changing some intervals but not others, in order to draw thematic parallels. Another reason for abandoning Reti is that many of the same insights can be gained in the context of a much more rigorous system-*Schenkerian Analysis*.

Schenker's system is a highly evolved, highly specific system of analysis. It posits the existence of universal musical skeletons around which all pieces are built. On top of this basic framework however, there is a near-infinite play of structural levels, harmonic and melodic variation, and formal organization. Motivic parallelism at different structural levels is not presupposed by this method. But when it does exist, a Schenkerian analysis will naturally bring it to light. This issue is explored in a great deal of depth in articles by Cadwallader and Pastille, and by Burkhart [13, 14]. Burkhart in particular gives a number of worked-out examples. To take just one example, Mozart's C major piano sonata, K. 545, contains the motive A-G-F-E near the beginning. The first time it occurs, it is a surface motive, meaning it is presented note-to-note. The next time it is encountered, it is a higher-level motive, meaning it is elaborated over a larger timescale, with other notes occurring between the notes of the motive. However, the motivic notes are more prominent, with more structural importance, and that allows them to be connected by the listener, despite their temporal separation. Further examples can be found in the articles.

Common in tonal music is *rhythmic* self-similarity [3]. In much of this music, there is an underlying sense of a *beat* or *pulse*. This is not just the tempo, or the rate at which the music plays. It refers to cyclical periods of strong and weak rhythmic events. For example, think of a waltz. The pulse is not just a flat stream of equivalent points: "one-two-three one-two-three." Rather, the first beat in each measure is emphasized: "ONE-two-three ONE-two-three." This does not just occur at the level of the fundamental pulse. *Hypermetric* and *submetric* emphases are also present. Measure-length units get grouped into larger periods, and faster events within a measure also possess differing strengths. Self-similarity occurs when the same pattern is present at multiple rhythmic levels. Let's examine a particular example: Bach's "Two-Part Invention 13" in a minor, the basis for example etude seven (see below). The meter is common time: quarter note gets the beat, four quarter notes per measure. Common time almost always has the same measure-level pulse. The first beat is the strongest; the third is also strong, although weaker than the first. The second and fourth beats are weak: "ONE-two-Three-four." Frequently, the quarter notes are divided into four sixteenth notes; they have the same pattern of emphasis. Or, two quarter notes are divided into four eighth notes, again with the same pattern. At the hypermetric level, the easiest divisions to hear are groupings of two, with the first measure strong and the second measure weak. It is harder to hear the measures in groups of four, so the pattern does not hold up completely, but the strong-weak-strong-weak relation is preserved. Many pieces have even larger groupings; it is not uncommon to have measures in common time divided into groups of four, then sixteen, and even higher. If the divisions get large enough, the hypermetric patterns often bleed into formal divisions, but this is somewhat more subjective, as it is hard to hear strong and weak patterns on the scale of an entire piece (though this did not stop Leonard Meyer from trying).

Notice that we have not said anything about how or why these metric pulses are present; this is a complex issue, and there is far from consensus on it.[6]

Finally, we turn our attention to the 20^{th} century. This era in music led to an explosion of innovation, and an active aesthetic of creating new compositional procedures. As a result, there was a proliferation of often unrelated methodologies, each requiring their own analytical techniques. We will focus on *twelve-tone* composition, one of the few schemes to gain a sizeable following. First developed by Schoenberg, twelve-tone takes as the fundamental unit a *twelve-tone row.* This is a specific ordering of each of the twelve pitches of the chromatic scale. Usually (but not always), we actually use *pitch classes*, meaning all octaves of a given pitch are considered equivalent (note that this partitions musical space into twelve equivalence classes). Each pitch class occurs exactly once in the row. Conventionally, we give each pitch class a number: $C = 0$, $C\# = 1 \ldots B = 11$. An example row would be: 0 4 8 2 9 11 10 3 1 7 5 6. The rows are implemented by presenting an instance of each pitch class, often melodically (placing notes one after the other in time), but also stacking several notes simultaneously. A piece that simply repeated a row over and over would not be very interesting, so there are a number of canonical transformations. The row can be *transposed* to start on a different pitch, but preserve the internal interval structure; *retrograded* by playing it in reverse order; *inverted* by keeping the number of semitones in each interval but switching ascending and descending (so 0 5 4... would become 0 7 8...); or any combination of these. Also, more than one row can be used in a piece, as long as there is some way of distinguishing them. Each instance of a row or its transforms leads to some measure of self-similarity. As an example, we can look briefly at Webern's "Variations for Piano", op. 27, mvmt. 2. This short movement exhibits a remarkable number of symmetries. The row used is: 4 5 1 3 0 2 8 9 10 6 7 11. It is used in a total of eight forms: four transposed to start on different pitches, and four additional forms retrograded and then transposed. At any given time, two different forms are being presented simultaneously (one in the left hand, one in the right). For each form, certain pitch classes must be fixed with respect to register, while others are free to occur in different octaves. The pitches that are fixed form a symmetric set about A440; in fact, virtually every aspect of the piece is centered around that pitch. Given certain constraints that are established, we can actually determine all the other row forms, given the initial form. In other words, the entire piece is generated from one row form, and is therefore self-similar up to the row transforms used.

One particular row that is quite interesting is the *Mallalieu row.* Its form is: 0 1 4 2 9 5 11 3 8 10 7 6. This row has many self-reflexive properties. Primary is the fact that we can take every second note of the row and end up with a transposition of the original (wrapping around from the end to the beginning, and adding a "dummy" 13^{th} note for counting): 1 2 5 3 10 6 0 4 9 11 8 7 is the original transposed up by one half-step. But the amazing thing is that this works for every $n = 1 \ldots 12$. For example, counting every sixth note gives 5 6 9 7 2 10 4 8 1 3 0 11. Even more

[6]There is an interesting old article by Jan LaRue that ties together issues of rhythmic emphases, metric and tonal concurrences, and quantization of patterns in music [15]. The specific problem he is addressing is phrasing and articulation in Haydn's symphonies, but today the article is much more relevant for what it tells us about early uses of computers for musical analysis. The methods used to divide the music into phrases and articulation points exhibit the then-universal ignorance of the sheer difficulty of programming. But the notion that statistical examinations of music could yield significant artistic insight was well ahead of its time.

amazing, this is the only row (up to a simple transformation: multiplying each pitch class by 7, mod 12) that has this property. A series of letters gives a number of other properties of this row, but it is clear that there is substantial potential for self-similar composition. A piece using this row, "Not Lilacs" by Robert Morris has been composed; it is also discussed in the literature [16].

The preceding discussion illustrates that opportunities for and examples of self-similarity abound in existing music. We now turn to a collection of pieces composed expressly to demonstrate some of these ideas, as well as new variants, in a simpler setting.

3. Example Etudes

The example etudes are a series of increasingly complex compositions designed to illustrate various aspects of self-similarity. Unlike most of the canonical literature we examined, the example etudes use well-defined generative algorithms, that allow us to specify the self-similarity more precisely. This does raise the aesthetic question of whether procedurally produced "music" has artistic value. The general consensus after the 20^{th} century is that algorithmic music is not *a priori* invalid. It depends on the artistic sense of the composer/programmer, as well as their ability to define procedures that lead to desired outcomes. Compositional discretion, rather than entering at the level of the notes themselves, is transferred to the algorithm design and the selection of input parameters.[7] In the case here, the etudes are for pedagogical purposes and so contain somewhat limited artistic merit. Nonetheless, there are moments of interesting music, and the methods shown here can certainly be extended to more substantial works.

It is assumed that the reader understands basic elements of music; specifically, note names (on the keyboard and generally), and how to read music. We will define most of the other necessary concepts as we go, but we should examine one before we begin: *chromatic vs. diatonic transposition.* Transposition means taking a section of music and raising or lowering the pitches, so that it starts and ends on different pitches while preserving the internal relations. Of the two types, chromatic transposition is more exact. It means that every note is raised or lowered by the same number of half-steps (a half-step is the distance between adjacent notes on the keyboard; from C to C#, for instance). So, if we chromatically transpose the notes C-D-E-A up by two half-steps (one whole-step), the result is D-E-F#-B. Diatonic transposition is slightly more complicated. It requires that we first identify the key and scale that the piece is in. For our purposes, we need only consider the major scale, a seven-note scale (plus the first note repeated an octave higher) with the interval pattern of whole and half steps "whole-whole-half-whole-whole-whole-half." An example is the C major scale: C-D-E-F-G-A-B-C. We give each note a number, called the *scale degree*, denoted by a number with a hat: $C = \hat{1}$, $D = \hat{2}$, etc. Diatonic transposition then means shifting each note by the same number of scale degrees. Diatonically transposing C-D-E-A by one scale degree yields D-E-F-B. Note that the number of scale degrees between each note is preserved, but not necessarily the number of half-steps.

[7] Algorithmic or procedural music is one of the most important musical developments of the computer age. It clearly involves concepts from wide-ranging areas of aesthetic theory, computer science and mathematics. For more discussion on both the technical and artistic aspects of this music, see Xenakis [17].

Fractal Melody One. This piece starts with the base melody, shown in the first four measures (Beethoven's famous "Ode to Joy" melody). Each following measure is a copy of the original melody (with repeated notes removed), transposed chromatically to start on successive notes of the original tune. In other words, stringing together the first note of each measure gives the base melody; this can be thought of as expanding or replacing each note with a copy of the melody. Since the procedure is only performed once, the result contains two hierarchical levels. The self-similarity is between these two levels; the transforms needed to relate them are time scaling and transposition (again ignoring repeated notes). Notice that while there are some interesting sounds produced, the rhythmic regularity has a "flattening" effect on the sound of the piece. Also, the exact replication of the melody causes each measure to be heard as an independent unit, obscuring the larger connections.

Fractal Melody Two. This etude builds on the same melody as the first (again replicated in the first measures). First, we made a chart containing the number of occurrences of each pitch, including duplicates:

C	2
D	4
E	5
F	2
G	2

For each measure, 15 notes were chosen at random, using a random number generator weighted by the above table, so that the likelihood of a particular note being chosen was the same as its rate of occurrence in the original. For example, the probability of choosing a D was 4/15. After this, each measure was transposed chromatically to "center" on successive pitches of the original. More precisely, we take E to be the center of the original. For each measure of the expanded tune, the random melody was first created using the notes of the original: C, D, E, F, G. Then, the entire measure was transposed up or down by the number of half-steps the corresponding note in the original called for. So the third measure corresponds to the third note of the original, F. F is one half-step above E, so after the random melody was created, the third measure was transposed up by one half-step. Much like the first piece, each measure is an expansion of a single note, giving two levels to the piece. However, instead of exact self-similarity, we now have statistical or fuzzy self-similarity. Rhythmic monotony is still an issue.

Fractal Melody Three. For this and subsequent etudes, a new melody is used: a simplification of a motive from Modest Mussorgsky's "Pictures at an Exhibition." There are three levels to this piece; each of them is shown in the score (the divisions are signalled by meter changes). First is the original melody. At the next level, each note is replaced by a measure. Rather than replicating the entire melody, for each measure five representative notes are chosen at random to stand for the original. So for example, the first expanded measure used the 1st, 5th, 6th, 8th and 12th notes of the original (A, G, E, G, A), in that order, as the basis. The second measure used the 2nd, 3rd, 4th, 7th, and 13th notes, and so on. Rather than centering each measure on the corresponding note of the original, we center them on the *tonal center* that note implies. Consider the original melody. It uses notes from the key of C major, and appears to emphasize the notes C and G. These and other musical considerations indicate that the melody is in the key of C, and we say C is the tonal center. But it starts not on C, but on A: $\hat{6}$ of C major, three half-steps below C. Since we want each measure to have a tonal center corresponding to successive notes of the original, we must first create the representative melodies, and then chromatically transpose each measure by an appropriate amount. The first measure should have a tonal center of A, three half-steps below the original tonal center of C. So we transpose the first measure down by three half-steps. The second measure should center tonally on G, five half-steps below C, so we transpose it down by five half-steps; and so on. To build the final level of the piece (what is really the actual piece; the previous levels are intermediate steps), we iterate the same procedure, applied to the second level. In other words, each note of the second level is replaced by a measure in the third. The selection of representative pitches still comes from the original melody, but now we transpose each measure to center tonally on successive notes of the second level. So the first measure of the third level centers on F#, etc. This gives three conceptual levels; again, the self-similarity is statistical rather than exact. It is more difficult to hear the large-scale connections between measures, especially at the third level, but this is not necessarily a bad thing aesthetically. It depends on how obvious we want to make the compositional process to the listener. Finally, note that rhythmic differences in the original melody (not present in "Ode to Joy") are not taken into account.

Fractal Melody Four. "Pictures at an Exhibition" is again the base melody. For this piece, representative notes are again chosen, but this time the same five notes are used every time: we select the 1st, 3rd, 6th, 11th and 13th notes. These were chosen by a very informal musical intuition as to what the most important notes are. The first and last notes of the melody are obviously relevant; the 3rd note is C, the tonal center; the 6th note is the highest point in the melody; the 11th note is a "local maximum," falls on a strong beat, and subjectively sounds important. Yet again, each note is replaced by a melodic fragment of these five representative notes. For the first time, however, rhythmic relations are preserved: eighth notes in the original melody are replaced by five sixteenth notes; quarter notes by five eighth notes; the half note by five quarter notes. So the ratios of duration remain the same. Also for the first time, diatonic rather than chromatic transposition is used. The first measure, corresponding to A in the original, is transposed down two scale degrees; the second measure is transposed down three scale degrees; etc. This means that the resulting etude uses only the notes of the C major scale. This was motivated purely by aesthetic/compositional choice, and the desire to exhibit a variety of possible schemes.

Fractal Melody Five. This is a procedurally simple derivation from Fractal Melody Four, that nonetheless yields very interesting new features. Basically, the two levels of Melody Four are presented simultaneously, the first level scaled to take the same amount of time as the second. Here the bass plays the original melody, flute the expanded melody. There is clearly a huge jump in audible complexity when *polyphony* (multiple simultaneous voices) is introduced. For one, we must now worry about the sound produced when two notes are played together. Since one voice is derived from the other, the composer needs to consider this when choosing the original melody, as different pieces of it will be overlapped in various ways. Also, having the bass play the original helps the listener keep track of the iterative process. We know that the flute is playing a measure based on the simultaneously presented note of the original melody. This aids in perceiving long-term relations.

Fractal Melody Six. As above, Melody Six is Melody Three, reworked to include polyphony. Exactly the same naive method was used: each level in the iterative procedure is presented simultaneously, all scaled to take the same amount of time. In this case we need three voices (flute, violin and bass). Since Melody Three used chromatic transposition, the result includes many dissonant-sounding intervals, as notes from different scales are played at the same time. Again, the presence of the original tune helps in following the procedure. At this point, it is clear that simply knowing the algorithm used to generate pieces like this does not necessarily tell you what the final output will sound like; the only way to do that is to run the algorithm and listen to the output. The interrelations between different levels, especially when used polyphonically, is simply too complex for humans to guess at or entirely control. If these procedures are used in larger-scale compositions, it would probably be necessary to spend a good deal of time "tinkering" with different input melodies and algorithm choices, in order to obtain a desirable output.

Fractal Melody Seven. This melody is in some sense the reverse of the above six; instead of expanding a base melody, we "compact" or "reduce" an existing piece. The starting point is Bach's "Two-Part Invention 13" in a minor. This piece, discussed in the previous section, has very regular metric patterns which make it easy to determine the most important notes of each measure. Each hand was treated separately. Each group of four sixteenth notes or two eighth notes was replaced by a quarter note on the first pitch of the grouping (so four sixteenth notes B-E-B-D would be replaced by a single quarter note B). If the first note of such a grouping was a rest, or is tied over from the previous grouping, the third note was used (in the case of two eighth notes, the second). Recall from the above discussion regarding rhythmic emphasis that the first sixteenth of every four is the strongest, and the third the next-strongest. Taking this as given, we see that this method of replacement yields the strongest notes of each grouping. The reduction is thus in some sense a "summary" of the original; it highlights important changes while smoothing out local elaborations (which of course give it much of its aesthetic value). This method could likely be extended to other pieces; the Invention works particularly well because of its total rhythmic consistency.

4. Conclusion

We have successfully shown the wide applicability of the concept of self-similarity. After describing many fundamental concepts and giving examples from mathematics and nature, we examined some of the many occurrences of self-similarity in existing music. Clearly, not all of these were intended by the composer; often they were a consequence of other musical goals. Rather than being a detriment, this gives great support to the usefulness of the concept as an analytical tool. Its applicability to situations not designed with it in mind is quite desirable, as this is a necessary task of any descriptive theory. After discussion of the literature, several example etudes were presented. These pieces are obviously fairly simplistic. But it is interesting to observe the unexpected musicality of even such unsophisticated methods. Often, novel sounds are produced that do merit some attention from the listener, and do not have a clear *sonic* basis in the original melody, even though the *procedural* derivation is straightforward. This is heartening and suggests that only a slight increase in methodological complexity is needed to produce sounds that are quite compelling.

5. Acknowledgements

Although music and science are beginning to overlap in scholarly circles, they are still for the most part distinct entities. In writing this project I was fortunate to have two advisors with interests in both fields. Prof. Joseph Eberly of the University of Rochester Physics Department was instrumental in conceiving this undertaking and ensuring that it gained the proper momentum. He was the principal advisor for the quantitative areas, but his artistic perspective as a non-professional musician was absolutely crucial. Prof. David Headlam, Eastman School of Music, Theory Department balanced this with his expertise on Western music. The depth and breadth of his understanding is quite astounding. Additionally, Dr. Headlam is involved with electronic music, and was able to work productively with the technical material as well. Finally, I would like to thank Daniel DiPaolo for his assistance with Schenkerian analysis in general, and motivic self-similarity in particular.

References

[1] Mandelbrot, Benoit B. *Fractals: Form, Chance, and Dimension*. W.H. Freeman, 1977.
[2] Gleick, James. *Chaos: Making a New Science*. Penguin, 1998.
[3] Solomon, Larry. *Solomon's Music Theory and Composition Resources*. (Excellent discussion and examples of many mathematical aspects of music, including fractals and self-similarity). http://music.theory.home.att.net/
[4] Madden, Charles. *Fractals in Music: Introductory Mathematics for Musical Analysis*. High Art Press, 1999.
[5] Quinn, Ian. "Fuzzy Extensions to the Theory of Contour." *Music Theory Spectrum*. 19.2, pp.232-63, 1997.
[6] Straus, Joseph. "Voice Leading in Atonal Music." *Music Theory in Concept and Practice*. James Baker, ed. University of Rochester Press, 1997.
[7] Joyce, David E. *Mandelbrot and Julia Set Explorer*. http://aleph0.clarku.edu/~djoyce/cgi-bin/expl.cgi
[8] Weisstein, Eric W. "Cantor Set." *Mathworld*. http://mathworld.wolfram.com/
[9] Mattfield, Jacquelyn Anderson. *"Cantus Firmus" in the Liturgical Motets of Josquin des Prez*. Yale University Dissertation, 1959.
[10] Bach, J.S. *Fourteen Canons: On the First Eight Notes of the Aria Ground from the "Goldberg Variations"*. Ed. Christoph Wolff. Kassel, 1976.
[11] Cook, Nicholas. *A Guide to Musical Analysis*. Oxford, 1987.
[12] Reti, Rudolph. *Thematic Patterns in Sonatas of Beethoven*. Ed. Deryck Cooke. Macmillan, 1967.
[13] Cadwallader, Allen and William Pastille. "Schenker's High-Level Motives." *Journal of Music Theory*, 36, pp.119-148, 1992.
[14] Burkhart, Charles. "Schenker's 'Motivic Parallelisms'." *Journal of Music Theory*, 22/1, pp.145–175, 1978.
[15] LaRue, Jan. "Two Problems in Musical Analysis: The Computer Lends a Hand." *Computers in Humanistic Research: Readings and Perspectives*. Edmund A. Bowles, ed. Prentice-Hall, 1967.
[16] *In Theory Only*, 2/7, 1977. See this volume, pp.8-21, for relevant discussion by David Lewin, Milton Babbitt, Richmond Browne and Robert Morris.
[17] Xenakis, Iannis. *Formalized Music: Thought and Mathematics in Composition*. Indiana University Press, 1971.