CHAPTER 5

Recursive Pitch Structures

Introduction

Part one of this thesis explored, in a cursory manner, the non-musical realms of natural science, cognitive science, and linguistics. Part two focuses exclusively on recursive elements of music production and perception, and assumes more expertise on the part of the reader.

Pitch – along with rhythm, timbre, and amplitude – is one of the fundamental components of sensory perception of music. In this chapter, we will investigate recursive aspects of pitch structures, including scales, chords, and melodies. Two things are worth mentioning before we begin: first, this chapter (like the entirety of part two) will focus almost exclusively on the Western system of musical organization, although intermittent mention will be made of other traditions, and the concepts discussed are intended to be understood as universals. Second, disentangling the parameters listed above is sometimes problematic – where does pitch end and timbre begin? – and there will therefore be some overlap between chapters 5, 6 and 7.

Fractal frequencies

Almost every musical culture utilizes some variety of scale in melodic construction, and the vast majority of these assume octave equivalence – meaning that scalar repetition is bounded within the space of an octave.[1] This is the case for all the diatonic rotations, as well as the acoustic scale, octatonic scale, whole-tone scale, and every other scale predominant in Western Music. The ratio of the frequencies in a perfect octave is, of course, 2:1 – meaning that for any pitch with a frequency x, the

1 Edward M. Burns, "Intervals, Scales, and Music," in *The Psychology of Music* 2nd Ed., Diana Deutsch, ed. (San Diego: Academic Press, 1999)

frequency of that same pitch one octave higher will be $2x$, and the octave above that $4x$, and so on. In perceiving stacked octaves like those in figure 5.1, our ears must traverse ever-widening segments of the frequency spectrum – a more accurate depiction of our *perception* of this sound can be seen in figure 5.2.

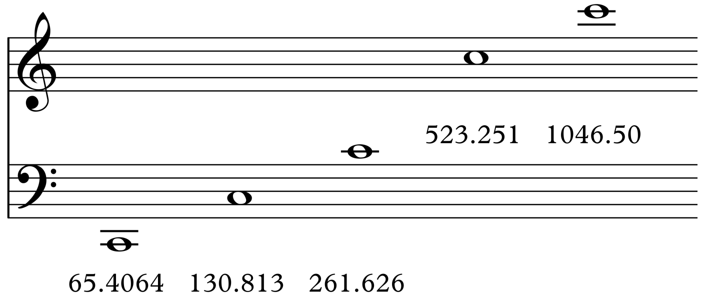

Figure 5.1 – Stacked octaves with frequencies in Hz.[2] *Traditional notation shows the octaves as approximately equidistant.*

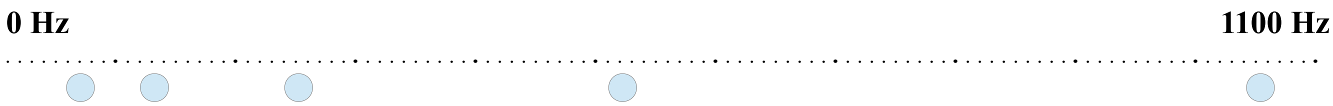

Figure 5.2 – In reality, the frequency of each octave is double that of the previous one. Each small dot represents 10 Hz.[3]

When considered this way, it is apparent that the term “octave equivalence” is another way of describing self-similarity. Each octave can be related to the last, but is greater in magnitude. Furthermore, any pattern of frequencies bounded by the octave – which is to say, practically any scale – will also display this self-similarity. Consider examples 5.3 and 5.4, which depict an octave-bounded

2 These frequencies assume equal temperament and A4 = 440 Hz.

3 This pattern can be extended into infinity in either direction, although the human capacity for perception extends only to frequencies between approximately 20 and 16,000 Hz. Searching for ever-lower octaves is interesting – consider that the next iterations of a C octave would be around 33 Hz, then 17.5 Hz, then 8.75, 4.375, etc. Eventually, differences of just a fraction of a cycle per second would (theoretically) cause the pitch to shift by many octaves. We can continue to halve these values indefinitely, in a scenario reminiscent of Zeno's paradox. The smallest vibration on record was measured in femtohertz – Hertz on a scale of 10-15. The “pitch,” emitted by a black hole 250,000,000 light-years from earth, was a B-Flat 57 octaves below middle C. (http://science.nasa.gov/science-news/science-at-nasa/2003/09sep_blackholesounds, accessed 29 January 2016).

Carnatic rāgam in several octaves, first in standard Western notation, and then according to frequency. Although the frequency scale is doubled with every iteration, the structures remain identical.

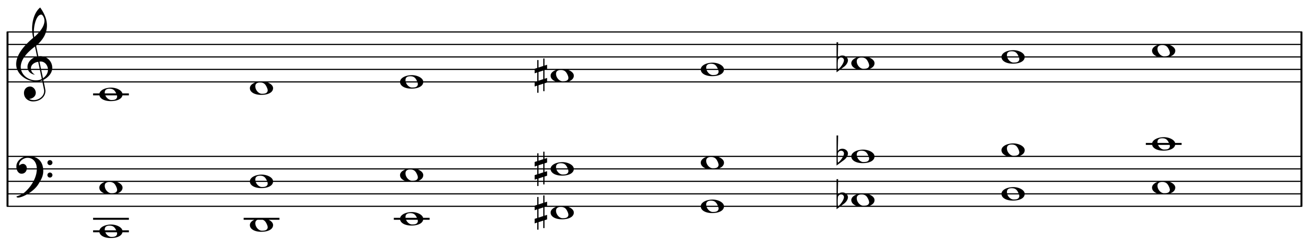

Figure 5.3 – Latangi rāgam, shadjam at C^4

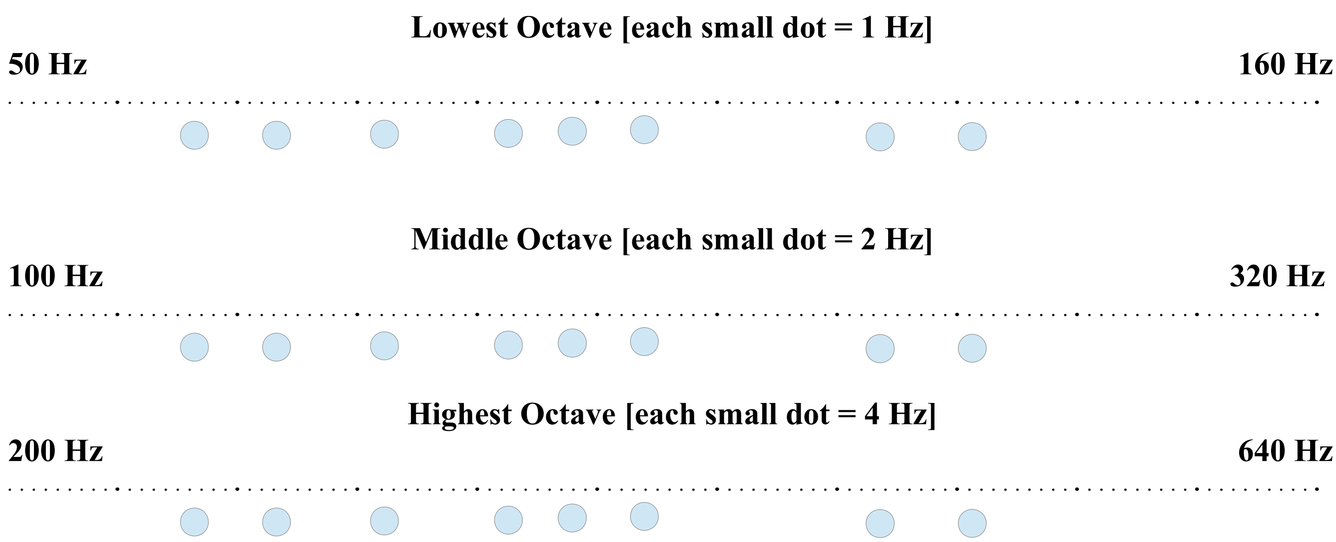

Figure 5.4 - Latangi rāgam in frequency scale notation. Notice that the frequency scale is condensed with each successive iteration – if arrayed in real space, these self-similar patterns would form a recursive wedge. The frequencies of the pitches in Hz are:

65.406, 73.416, 82.407, 92.499, 97.999, 103.826, 123.471, 130.813 (Lowest Octave)

130.813, 146.832, 164.814, 184.997, 195.998, 207.652, 246.942, 261.626 (Middle Octave)

261.626, 293.665, 329.628, 369.994, 391.995, 415.305, 493.883, 523.251 (Highest Octave)

4 S. Bahgyalekshmy, *Ragas in Carnatic Music* (Jaipur: CBH Publishers, 2006).

As arranged above,the frequencies in each column are self similar, increasing at a fixed proportion as one moves from top to bottom. This forms a recursive wedge. In this sense, a scale spanning multiple octaves resembles the periodic table of the elements (figure 1.1).

All of this demonstrates that scales – or any repeating division of the octave – are not just logarithmic, but fractal in nature. The self-similar array of pitches we see in each example in figure 5.4 would theoretically continue to repeat infinitely in either direction, even as the frequency scale approached zero (by being halved over and over again) or raced off to infinity (by being doubled). In this respect, octave-bounded scales are very similar to the *Cantor Ternary Set*, a fractal collection which can be represented visually by drawing a line, then removing the middle third, then doing so again, and so on.

Figure 5.5 – The Cantor Ternary Set. Every iteration is derived by removing the middle third of each line in the iteration before.

The Cantor Ternary Set, like an octave-bounded pitch array, demonstrates a binary logarithm. The number of lines in each successive iteration increases in the pattern 1*n*, 2*n*, 4*n*, 8*n*... – the powers of two. This same pattern dictates the frequencies of each successive octave in figure 5.1 (65.4064, 130.813 [= **2** x 65.4064], 261.626 [= **4** x 65.4064], 523.251 [= **8** x 65.4064])... Moreover, the Cantor Set displays self-similarity at each hierarchical level, just the same way as the Latangi rāgam in figure

5.4. The fractal dimension of the Cantor Ternary Set is ≈ 0.631.[5] We can hypothesize that the rāgam (or a major scale, or any other octave-bounded repeating array) would likewise possess a fractal dimension, although this dimension would vary based on the number of constituents in the array and the precise dimensions at which they divide the octave.

Bounded scale patterns which contain smaller repeating elements (including all of Messiaen's "modes of limited transposition") possess yet another recursive level of self-similarity. In these instances, each patterned sub-array can also be classified as self-similar at a different proportion than the larger array. Consider Messiaen's third mode, which is bounded by the octave and contains three four-pitch sub-arrays which follow the pattern M2, m2, m2.[6]

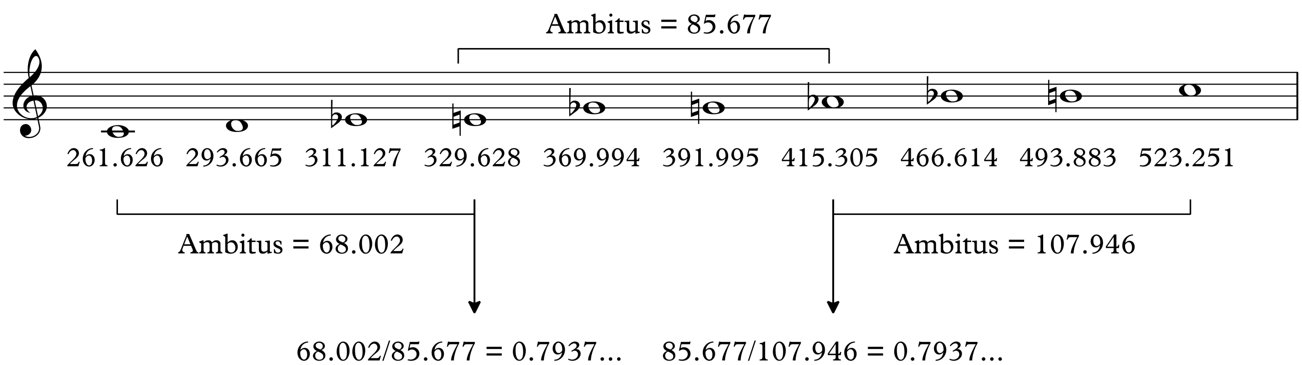

Figure 5.6 – Messiaen's third mode with frequencies in Hz. The large interval pattern, bounded by the octave, would be self-similar at the ratio of 1:2, or .500, in subsequent iterations. The smaller pattern, bounded by the interval of a M3, is self-similar at the ratio of 4:5, or .800 (because we are assuming equal temperament, the frequencies only approximate this ratio).

5 Mandelbrot, *Fractal Geometry of Nature*.
6 Olivier Messiaen, *The Technique of My Musical Language* (Paris: Alphonse Leduc, 1956).

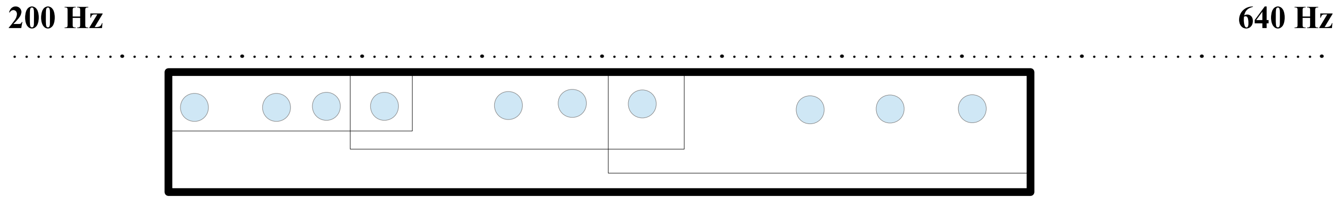

Figure 5.7 – Messiaen's third mode in frequency scale notation. Each repeating small array is enclosed in a thin box. Each small dot = 4 Hz.

In short, all bounded scales can be considered fractal when realized in different octaves, and scales comprised of smaller repeating patterns demonstrate additional levels of recursion. It's worth mentioning that, while all scales predominant in Western music (and all of Messiaen's modes) presume a twelve-fold division of the octave, any repeating division would demonstrate these same properties.

Scales in which every constituent is equidistant from every other (in traditional notation) are special cases, because they are recursive at every measurable level. In the whole-tone scale, for example, the fundamental sub-array consists of a single interval (M2), such that every adjacent scale degree can be said to be self-similar to the next at a ratio of roughly 8:9. But likewise, every collection of two adjacent degrees (bounded by a M3) is self-similar to every other collection of two at a ratio of approximately 4:5, and every collection of three degrees (bounded by a tritone) is self-similar to every other collection of three at a ratio of approximately 5:7. And so on.

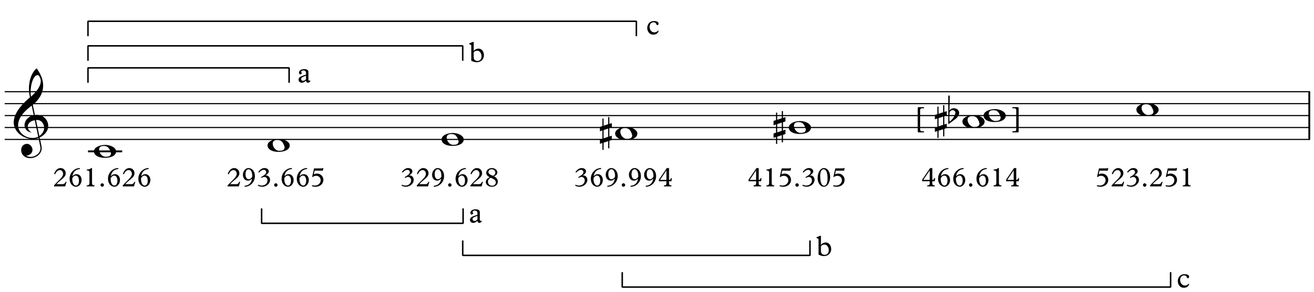

Figure 5.8 – The whole-tone scale. The bracketed segments marked "a" are demonstrably self-similar, as are those marked "b" and "c."

Therefore, when we, as listeners, are confronted with a construction based on the whole-tone scale (or the chromatic scale, or any other scale comprised of a single interval), we are navigating a large number of recursive layers which cannot be reliably separated from one another. The perceived ambiguity of the scale is very much like the disorienting nature of this sentence from chapter 4:

(1) Abbey called her first cousin Jim's lawyer Bobby's grand-nephew Sam's brother Jeff's mother Karen's friend Tom.

Just as we might be hard-pressed to define Sam's relationship with Abbey, we cannot determine whether a constituent of the whole-tone scale is one-third of a three pitch array, one-half of a two-pitch array, etc. The many layers of recursive nesting preclude immediate comprehension. But unlike the characters mentioned in sentence (1), the relationships between whom can be untangled if one is suitably patient, the degrees in the whole tone scale also demonstrate *individual* self-similarity – as if every person in (1) was the same person. This pushes the level of ambiguity even further:

(2) Abbey called her first cousin Abbey's lawyer Abbey's grand-nephew Abbey's brother Abbey's mother Abbey's friend Abbey.

Special recursive qualities of the diatonic and pentatonic collections

The diatonic and anhemitonic pentatonic[7] collections are arguably the two most important scale arrays in the Western tradition. The former is the structural basis for all Western modal and tonal music – its rotations include the major and minor scales, as well as all the diatonic modes. The latter is not only common in art and popular music, but pervasive in the folk traditions of Europe, America, Indonesia, China, and many other cultures.

7 The designation *anhemitonic* denotes a pentatonic array which contains no half-steps. In this paper, I intend it to refer to the "major" pentatonic collection (M2 M2 m3 M2) and its rotations exclusively.

These scales share two important characteristics. First, they are "well-formed,"[8] meaning that they are generated by a single interval. In other words, the degrees of the pentatonic scale can be found by reiterating the perfect fifth (or its inversion) five times:

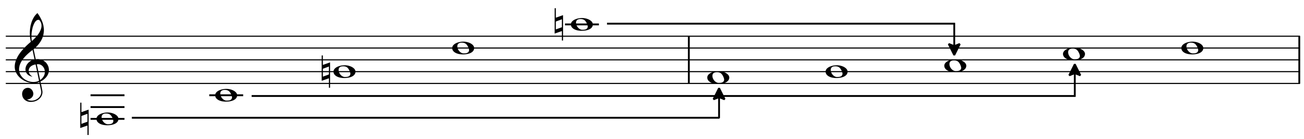

Figure 5.9 – The well-formedness of the pentatonic array.

And the degrees of the diatonic scale can be found by reiterating the same interval class seven times:

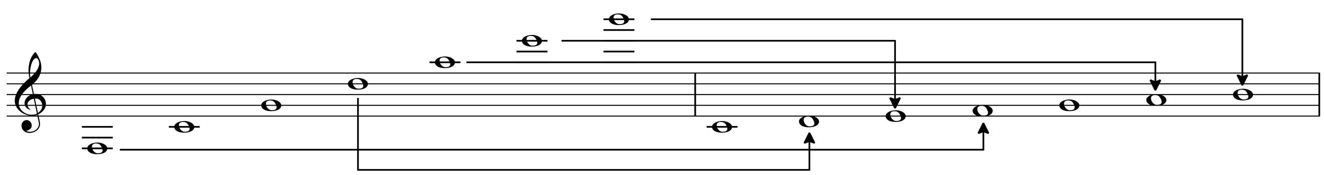

Figure 5.10 – The well-formedness of the diatonic array.

It is possible to show this property arithmetically. An equation for finding the frequency F of a scale degree *n* in the pentatonic or diatonic collections is $Fn = 3/2\,[F(n-1)]$, where $0 < n \leq 5$ (pentatonic scale) or $0 < n \leq 7$ (diatonic scale) and *n* is a whole number. Because the output of one iteration of this equation becomes the input for the next, the equation forms a feedback loop, and is therefore recursive. The property of octave equivalency (which is to say, self-similarity – see above) is then used to condense the results of the function into a single octave as shown in figures 5.9 and 5.10.

Some well-formed scales can be shown to possess *Myhill's Property*,[9] meaning that every generic scale interval (that is, span between scale degrees) is found in exactly two different specific

8 Norman Carey and David Clampitt, "Aspects of Well-Formed Scales," *Music Theory Spectrum* 11/2 (Autumn 1989): 187-206.

9 Named for British mathematician John Myhill (1923-87), although he had little or nothing to do with its formulation. Myhill's property was actually defined by Myhill's colleagues George Clough and Gerald Myerson. Timothy Johnson, *Foundations of Diatonic Theory: A Mathematically-Based Approach to Musical Fundamentals* (Lanham, MD: Scarecrow Press, 2008).

sizes. For example, generic interval 1 in the pentatonic scale can be either a M2 (for example, F-Natural to G-Natural in figure 5.9) or a m3 (A-Natural to C-Natural in figure 5.9) – but those are the only two possibilities that exist. The precise qualities of intervals available in the pentatonic and diatonic arrays are listed in figure 5.11.

Pentatonic Scale

Generic Interval	Specific Interval (smaller)	Specific Interval (larger)
1	M2	m3
2	M3	P4
3	P5	m6
4	M6	m7

Diatonic Scale

Generic Interval	Specific Interval (smaller)	Specific Interval (larger)
1	m2	M2
2	m3	M3
3	P4	A4
4	d5	P5
5	m6	M6
6	m7	M7

Figure 5.11 – List of generic and specific intervals in the diatonic and anhemitonic pentatonic arrays. Because every generic interval has exactly two possible specific realizations, these scales are said to possess Myhill's Property.

In each of these cases, every larger interval can be divided into fixed combinations of "step intervals" (that is, intervals of generic class 1). So every M6 in the pentatonic scale contains a certain number of M2s and a certain number of m3s. Every M6 in the diatonic scale contains a certain number of m2s and a certain number of M2s. And so on. This way of dividing intervals is far from new – the theorist Prosdocimus Beldemandis (d. 1428) describes it in a treatise from 1412.[10] The results when

10 De'Beldomandi, *Contrapunctus,* ed. and tr. Jan Herlinger (Lincoln, NE: University of Nebraska Press, 1984).

applied to the pentatonic and diatonic scales are as follows:

Pentatonic Scale

Generic Interval	Specific (Smaller)	no. of M2s	no. of m3s	Specific (Larger)	no. of M2s	No. of m3s
1	M2	1	0	m3	0	1
2	M3	2	0	P4	1	1
3	P5	2	1	m6	1	2
4	M6	3	1	m7	2	2

Diatonic Scale

Generic Interval	Specific (Smaller)	no. of m2s	no. of M2s	Specific (Larger)	no. of m2s	No. of M2s
1	m2	1	0	M2	0	1
2	m3	1	1	M3	0	2
3	P4	1	2	A4	0	3
4	d5	2	2	P5	1	3
5	m6	2	3	M6	1	4
6	m7	2	4	M7	1	5

Figure 5.12 – The intervals of the pentatonic and diatonic scales, expressed as combinations of step intervals.

One way to describe these scale arrays is with a ratio describing how many small or large generic step intervals they employ out of the total number of possibilities. For example, the complete pentatonic array contains five step intervals, of which two are m3s. So the pentatonic scale [:m3:][11] is 2/5. Similarly, the diatonic scale [:m2:] is 2/7. Interestingly, whichever step interval one chooses to measure, the ratio of the smaller specific interval to the larger will *always* encompass the whole-scale ratio. Consider the following:

11 I will use this symbol to denote which interval is generating the ratio. We can imagine this to mean, "as expressed via a ratio of number of step intervals (size m3) to number of total step intervals" (which, let's face it, is a mouthful).

Pentatonic Scale [:M2:]

Generic	Specific (Smaller)		Whole-Scale Ratio		Specific (Larger)
1	1/1	>	3/5	>	0/1
2	2/2	>	3/5	>	1/2
3	2/3	>	3/5	>	1/3
4	3/4	>	3/5	>	2/4

Pentatonic Scale [:m3:]

Generic	Specific (Smaller)		Whole-Scale Ratio		Specific (Larger)
1	0/1	<	2/5	<	1/1
2	0/2	<	2/5	<	1/2
3	1/3	<	2/5	<	2/3
4	1/4	<	2/5	<	2/4

Diatonic Scale [:m2:]

Generic	Specific (Smaller)		Whole-Scale Ratio		Specific (Larger)
1	1/1	>	2/7	>	0/1
2	1/2	>	2/7	>	1/2
3	1/3	>	2/7	>	1/3
4	2/4	>	2/7	>	1/4
5	2/5	>	2/7	>	1/5
6	2/6	>	2/7	>	1/6

Diatonic Scale [:M2:]

Generic	Specific (Smaller)		Whole-Scale Ratio		Specific (Larger)
1	0/1	<	5/7	<	1/1
2	1/2	<	5/7	<	2/2
3	2/3	<	5/7	<	3/3
4	2/4	<	5/7	<	3/4
5	3/5	<	5/7	<	4/5
6	4/6	<	5/7	<	5/6

Figure 5.13 – The distribution of step intervals in the pentatonic array, the diatonic array, and every constituent interval therein.

Not only is the whole scale ratio encompassed by that of the smaller and larger intervals, it is encompassed as closely as possible. In the last table of figure 5.13, 0/1 and 1/1 are the closest approximations of 5/7 with the denominator "1." 1/2 and 2/2 are the closest with denominator "2," 2/3 and 3/3 are the closest with the denominator "3," and so on. In every case, the step makeup of the interval reflects the makeup of the scale as a whole. The larger the intervals become, the more precise the relationship becomes. In other words, *every interval contained in the diatonic and pentatonic collections can be shown to be self-similar to the collection as a whole*. Furthermore, Carey and Clampitt (1996) have shown that this self-similarity will exist for any well-formed scale possessing Myhill's property.

L-System Melody Generation

A Lindenmayer System, or *L-System*, is a set of transformation rules that recursively reiterates.[12] Every System has four parts: an *axiom* (initial state), *variables* (a vocabulary of potential results), *production rules,* and sometimes *constants* (which remain the same in every iteration). L-Systems describe a wide variety of natural phenomena – in fact, the following system (used by Lindenmayer himself to explain the growth patterns of algae) also demonstrates the reproductive pattern of honeybees, as mentioned in chapter 1, with the number of variables in each string being a Fibonacci number.[13]

Axiom: 1

Variables: 1, 0

Production Rules: (1 → 1 0) (0 → 1)

Constants: N/A

12 L-Systems are named for Hungarian Lindenmayer (1927-1989), a Hungarian botanist who originally used this system of string generation to describe the growth patterns of algae. If you think Lindenmayer is the only Hungarian botanist who will appear in this book, you are correct.

13 Grzgor Rozenberg and Arto Salomaa, *The Mathematical Theory of L-Systems* (New York: Academic Press, 1980).

$n0 = 1$
$n1 = 1\ 0$
$n2 = 1\ 0\ 1$
$n3 = 1\ 0\ 1\ 1\ 0$
$n4 = 1\ 0\ 1\ 1\ 0\ 1\ 0\ 1$
$n5 = 1\ 0\ 1\ 1\ 0\ 1\ 0\ 1\ 1\ 0\ 1\ 1\ 0$

Figure 5.14 – L-System describing the reproductive pattern of Anabaena catenula *(and honeybees).*

And so on. L-System notation can also be shown to describe various fractal figures, including the Cantor set (see above) and the Koch curve (chapter 1):

Axiom: 1

Variables: 1, 0

Production Rules: (1 → 1 0 1) (0 → 0 0 0)

Constants: N/A

$n0 = 1$
$n1 = 1\ 0\ 1$
$n2 = 1\ 0\ 1\ 0\ 0\ 0\ 1\ 0\ 1$
$n3 = 1\ 0\ 1\ 0\ 0\ 0\ 1\ 0\ 1\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 1\ 0\ 1\ 0\ 0\ 0\ 1\ 0\ 1$

Figure 5.15 – L-System describing the Cantor Set. The number of "1" markings (equivalent to solid line segments in figure 5.5) grows by powers of 2.

One common way of transforming these instructions into a geometric figure is to employ *turtle graphics.* Turtle graphics make use of an imaginary moving stylus (the "turtle") which travels over a plane using L-System instructions to control its movement and orientation. When instructed, the turtle can create a line behind it as it travels. Consider the following (somewhat more advanced) L-System, which describes a shape known as the "Peano Curve:"[14]

14 Named for Giuseppe Peano (1858-1932), an Italian mathematician who was among the first to experiment with fractal geometry.

Axiom: 1

Variables: 1 0 F + -

Production Rules: (1 → 1F0F1+1+0F1F0-F-1F0F1) (Y → 0F1F0+0+1F0F1-F-0F1F0)

Constants: δ [turtle heading] = 90°

As you can see, the length of the strings grows very rapidly as this system reiterates. A potential realization of this system instructs the turtle as follows:[15]

1) For every F, move forward one step-length.
2) For every +, turn clockwise
3) For every -, turn counter-clockwise.
4) Ignore all other symbols.

For this set of turtle instructions, the following images are produced following *n1*, *n2* and *n3*.

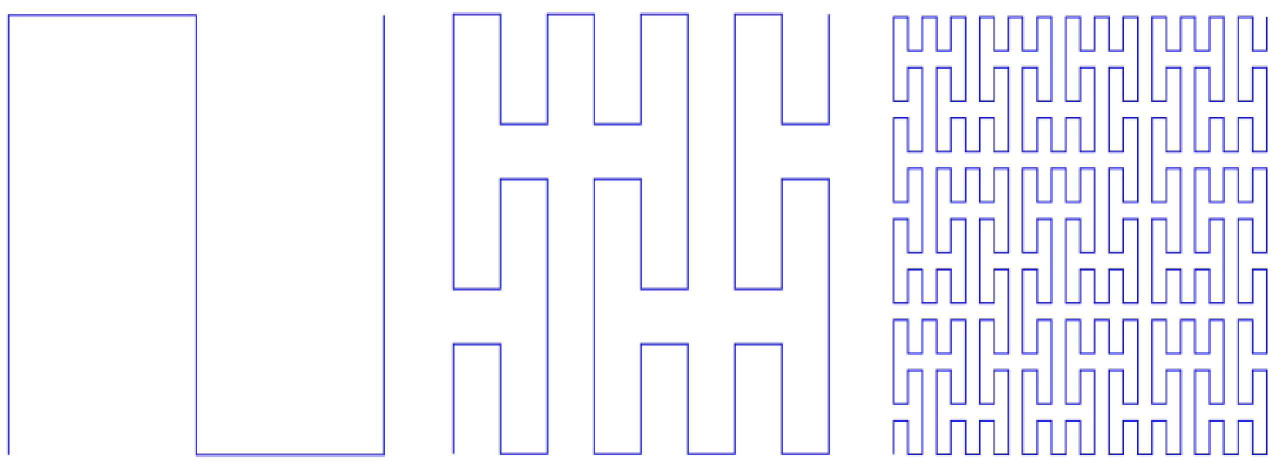

Figure 5.16 – The Peano curve after its first (left), second (middle) and third (right) post-axiom iterations.

In 1986, Mathematician and computer scientist Przemysław Prusinkiewicz (b. 1952) suggested a method for realizing L-System visualizations musically. In his system, vertical lines are interpreted as intervals within a given scale, while horizontal lines represent a given duration. So from any starting

15 Stephanie Mason and Michael Saffle, "L-Systems, Melodies, and Musical Structure," *Leonardo Music Journal* 4 (1994): 31-38.

point *x*, the Peano curve could be interpreted as follows:

1) Ascend two scale degrees.
2) State note of length 1.
3) Descend two scale degrees.
4) State note of length 1.
5) Ascend two scale degrees.

And so on. Of course, this operation can be performed beginning at either end of the figure, and the figure can be rotated to any multiple of 90°, meaning that any L-System visualization can produce eight different musical results. Because of the iterative nature of the source image, these melodies are always comprised of a repeating contour pattern. But otherwise, they can vary a great deal depending on the starting point, scale system, and orientation of the reading.

x = C♮
Basic Rhythmic Unit = ♪

Figure 5.17 – Various musical realizations of the Peano curve's second iteration. Numbers 1 and 3 read the curve in the orientation seen in figure 5.16, while numbers 2 and 4 read the curve at a rotation of 90° counter-clockwise. Numbers 1 and 2 utilize the diatonic scale, number three utilizes the chromatic scale, and number 4 utilizes the pentatonic scale.

Of course, observing that such a system can be used to generate music is of limited usefulness – almost anything can serve that function. But certain readings of L-Systems seem to compulsively restate small patterned tropes that are common in both art and folk idioms. Consider the repeating contour pattern in figure 5.17, example 2. It is comprised of two parts – a two-step scale which descends and returns (*a*), followed by another two-step descent (*b*), whereafter the pattern restarts. A huge variety of melodies can be related to this basic shape either by direct parallel, partial inversion, or complete inversion.

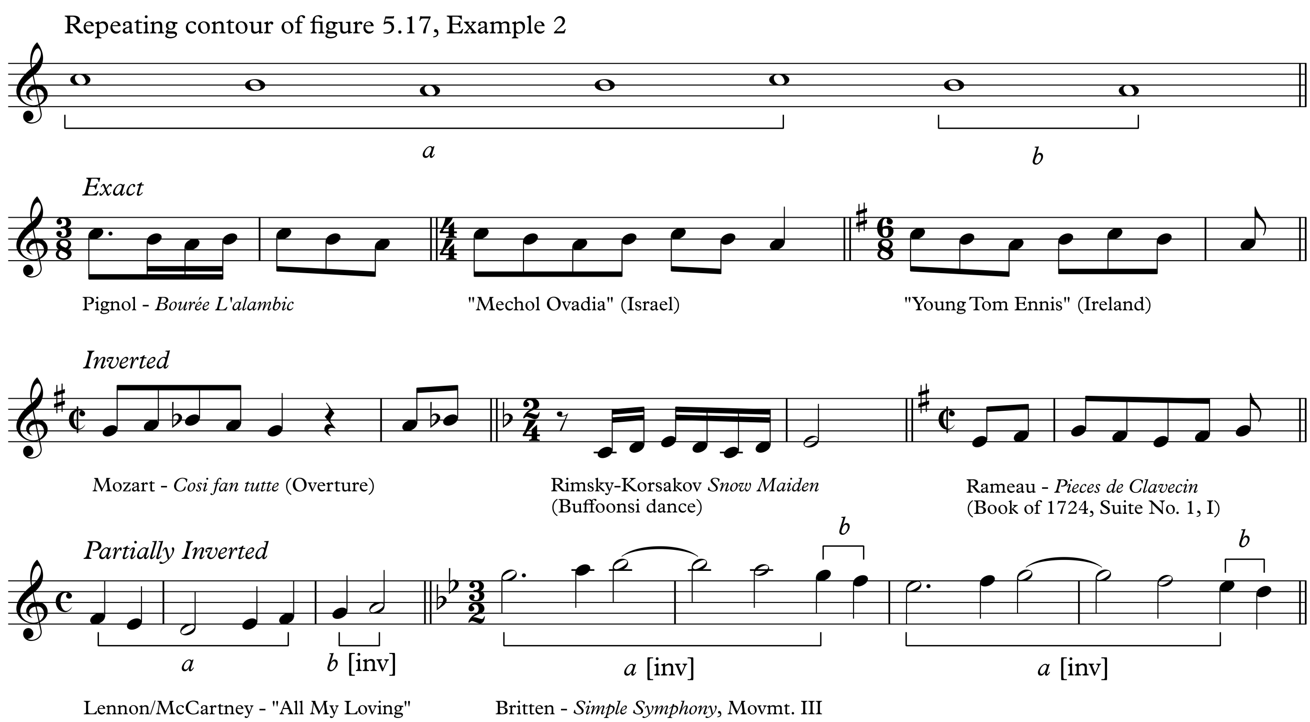

Figure 5.18 – Melodies prominently including the repeating contour pattern of Figure 5.17, example 2. This is a small sample of the total melodies available.

This is just one example. Mason and Saffle (1994) compared the melodic realization of the Peano curve in a different rotation against the entries in Denys Parsons' *Dictionary of Musical Tunes* (1975 edition) and found over two dozen direct matches. In a separate experiment reported in the same paper, they compared the quadratic Gosper curve (another L-System visualization) in every possible rotation against the same database, producing over 200 matches.[16] Based on these convincing results from a single L-System figure – and remember, there are an infinite number that are theoretically possible – the researchers hypothesized that a high percentage of melodies could be shown to possess fractal characteristics.

Whether or not you consider self-similarity discovered in this way to be perceptible, the fact remains that melodies in Western music tend to be highly patterned, and that many of these patterns can be discovered through sonic realization of fractal figures. This implies a link. The fact that these melodies are, by and large, constructed using scales which can themselves be shown to be self-similar helps to solidify this link. And as we move from the realm of melody to the realm of rhythm, we will discover further reinforcement in the recursive musical patterning of time itself.

16 These results allowed for a predetermined amount of rhythmic "stretching" of the musical L-System realization. Nonetheless, the results are compelling. Mason and Saffle, "L-Systems, Melodies, and Musical Structure."