

The Sound of Space-Filling Curves*

Herman Haverkort
Eindhoven University of Technology, the Netherlands
cs.herman@haverkort.net

3 August 2017

Abstract

This paper presents an approach for representing space-filling curves by sound, aiming to add a new way of perceiving their beautiful properties. In contrast to previous approaches, the representation is such that geometric similarity transformations between parts of the curve carry over to auditory similarity transformations between parts of the sound track. This allows us to sonify space-filling curves, in some cases in up to at least five dimensions, in such a way that some of their geometric properties can be heard. The results direct attention to the question whether space-filling curves exhibit a structure that is similar to music. I show how previous findings on the power spectrum of pitch fluctuations in music would suggest that musically sounding results are more easily obtained with three- or four-dimensional space-filling curves, whereas two-dimensional curves may be considered less musical. Conversely, assessing the musical quality of examples of sound tracks of space-filling curves of different dimensions may contribute to confirming or rejecting previous hypotheses on the relation between power spectra and music. Proof-of-concept sound tracks of two-, three-, four- and five-dimensional curves are provided on-line [11].

Introduction

A space-filling curve in d dimensions is a curve that is so crinkly that it completely fills a higher-dimensional space. More formally, it can be understood as a continuous, surjective mapping from $[0, 1]$ to a subset of \mathbb{R}^d that has volume (d -dimensional Lebesgue measure) greater than zero—for example a d -dimensional cube. Space-filling curves can be described, constructed, and analysed in many different ways and have many applications in computer software [2, 12, 30]; the function of a space-filling curve typically lies in providing a way to traverse elements in a two- or higher-dimensional space in such a way that consecutive elements in the traversal tend to lie very close to each other in space.

For this paper, the following approach to describing space-filling curves suffices (see, for example, Ventrella [32]). Consider a *plane-filling* curve, that is, a two-dimensional space-filling curve. We consider line segments to have a direction (forward or backward) and an orientation (left or right). A plane-filling curve is defined by a so-called *generator* that describes how to replace a left-forward line segment from $(0, 0)$ to $(1, 0)$ by a sequence of smaller line segments. From this generator, replacements for scaled, rotated, translated, reflected and/or reversed line segments can be deduced by applying the corresponding transformations. If we apply the generator recursively to its constituting line segments, to an infinite recursion depth, then we obtain a fractal curve—see Figure 1 for an example. If the sum of the squared lengths of the line segments in the generator is 1, then the fractal dimension of the resulting curve is 2. Hence, if the curve does not overlap itself too much, the curve must be a plane-filling curve, filling a region R of non-zero area in the plane. The boundary of the region R could be a simple shape such as a square, or it could be a fractal itself (with fractal dimension less than 2), possibly self-intersecting. The concept generalizes naturally to higher dimensions.

There is a natural correspondence between plane-filling curves and tessellations with reptiles (repeating tiles [10]). Let the segments of a generator be indexed s_1, \dots, s_n . When the generator is applied

*Original date: 1 March 2017. Cite this paper as: Herman Haverkort: The Sound of Space-Filling Curves. *Proc. Bridges 2017*, pages 399–402. Extended manuscript downloaded from <http://spacefillingcurves.net> on (today's date).

recursively, each segment s_i expands to a curve filling a two-dimensional object R_i in the plane. Because the same generator is applied to each segment of the generator, the corresponding planar objects R_1, \dots, R_n are all similar to each other, and together they tile R , which also has the same shape. In practice, when designing space-filling curves, it is often easier to first design and describe the recursive tessellation, and then the curve—this approach also underlies my work on higher-dimensional Hilbert curves [12]. However, there are also many curves with simple generators whose complicated tile shapes are not well-understood: Ventrella shows many examples [32].

Goal of this work. There is something intriguing about the fact that a curve can twist so much that it can really visit every point within a two- or higher-dimensional space: before Peano [28], this was not even known to be possible. There is also something intriguing about the recursive nature of such curves. The goal of the present work is to create a non-visual way of representing space-filling curves that adds to the perception of their beauty. Sound may be a good choice for this, since space-filling curves seem to have a structure similar to music, in which small motifs string to form themes that are repeated and subjected to various transformations (we will discuss this in detail further on), in turn forming larger sections. Ideally, we will be able to recognize some of a curve’s geometric properties by ear. Therefore, we are looking to create auditory renderings (or sonifications) that meet the following conditions:

- they represent the space-filling curves accurately, that is, the curves can be deduced from the sound or the musical scores (at least in theory), and geometric similarity between different sections of the curve corresponds to recognizable similarity between sections of the sound track;
- the sound is pleasant, that is, listening to the sound, possibly while watching a figure or an animation that shows the graphical representation of the curve, must be an entertaining experience, not an unsettling one.

We do not require these goals to be met by fully automatic means. In particular, similar to graphical depictions that look better in a frame, we may prefer to frame the auditory rendering of a space-filling curve with supporting voices. Note that our initial goal is not (yet) to use space-filling curves to make interesting music, but rather to use sound, possibly music, to demonstrate interesting curves. However, as we will see later, this endeavour may also teach us something about the nature of music, and music may result.

Related work. Zare’s delightful composition *Fractal Miniatures* [38] includes musical impressions of two well-known plane-filling curves. Zare writes about his composition that “elements of fractal geometry are alluded to, including symmetry and their additive nature, but all of the music flows organically without any mathematical processes guiding it” [39]. Thus, Zare’s work meets our second condition in a wonderful way, but not the first condition that we require for our present discussion.




Figure 1: (a) Generator for the Meander curve from Wunderlich [37]. (b) Result of applying the generator recursively to depth two. (c) If we start from a line segment, apply the generator recursively to infinite recursion depth, and sample the resulting curve at $3/4$ of its length, this results in a sample point that lies to the left of the original line segment. (d) The sample points for all line segments of the k -th refinement level fill a square grid of 9^k points in which consecutive samples are always neighbours—as illustrated by the curve (bold) connecting the sample points.

Many composers have used fractals or self-similar constructions (see, for example, [1, 6, 18, 19, 25, 36]) to generate music, often overlaying representations on multiple scales. Composers [5, 21] and scientists [14] also made music or sound from the contours of geometric objects drawn in the plane by mapping one coordinate axis to time and the other coordinate axis to pitch. However, such an approach, where a geometric dimension is mapped to time, is not immediately suitable for us, since a space-filling curve inevitably moves back and forth in all dimensions, and sound cannot go back and forth in time. In that sense, the work by Hermann et al. [13] comes closer to our needs: they construct a sonification based on a traversal of an unrestricted curve in two-, three-, or even higher-dimensional space. However, in their work, it is not the actual curve that is sonified, but data points that surround it—the curve is used as a means to traverse a set of point data. In other words: the sonification provides a “view” *from* the curve, but not a view *of* the curve.

Johnson [18] explored the following solution, which I call *direction mapping*. Take a sketch of a curve (as in Figure 1, for example), and calculate what computational geometers know as the *turning function* θ . This function can be described as follows: imagine we traverse the curve from one end to the other, and let t be the distance travelled so far, then θ is a function of t that goes up by α whenever we turn left by an angle α , whereas $\theta(t)$ goes down by α whenever we turn right by an angle α . Thus, at any distance t from the start, the function $\theta(t)$ indicates the direction in which we are moving. Now map t to time and $\theta(t)$ to pitch. As a result, curve sections that differ from each other only by rotation are mapped to musical phrases that differ only by transposition (shift in pitch). Curve sections that differ only by translation are mapped to musical phrases that are identical, or differ by a transposition that corresponds to a multiple of 360 degrees. For two-dimensional curves, this solution clearly has its merits, as demonstrated by Johnson’s compelling compositions (for example, [15, 16, 17]). However, if we want to render three- or higher-dimensional curves, or if we want geometric translation to be represented in the sound, we need an alternative approach.

Building on an approach by Prusinkiewicz [29], which was put into practice by Nelson [25], Mason and Saffle [24] proposed a solution for plane-filling curves that would turn a sketch of a curve such as the one in Figure 2(a) into a piece of music with two voices. One voice plays the horizontal segments of the sketch in order, while the other voice plays the vertical segments in order—where each segment’s projection on the orthogonal axis determines the pitch of a note, while the segment’s length determines the duration. Since in general, in a given part of the sketch, the total length of the horizontal segments is not the same as the total length of the vertical segments, the two voices will run out of sync. As a result, if the horizontal segments of square A in Figure 2(a) are played in counterpoint (together) with the vertical segments of square A , then, by the time we approach the bottom right corner in the figure, the voices are out of sync by four time units, so that the horizontal segments of square B are played *after* the vertical segments of square B and in counterpoint with vertical segments from the next square. This might create interesting variations when the aim is to create music¹, but not when our aim is to preserve similarity between curve sections.

I have not found any descriptions of a transformation from space-filling curves into music that successfully preserves both geometric similarity and geometric differences between sections of the curve, and may be applicable to two- and higher-dimensional space-filling curves.

About this paper. Below, I first discuss *location mapping*, which is my main approach to turning space-filling curves into sound. After that, I discuss what to expect from location mapping and direction mapping (Johnson’s approach) given previously published hypotheses on the power spectrum of pitch variations in musical sound. We will see how the answer depends on the number of dimensions of the space-filling curve, which demonstrates the potential relevance of results with space-filling curves of different dimensions for confirming or rejecting such hypotheses. Finally we discuss opportunities and challenges for further work.

¹For further references on how to create music with L-systems (including space-filling curves and turtle graphics), one may consult Manousakis [23] or Pareyon [27].

Converting Space-Filling Curves into Sound by Location Mapping

For ease of explanation, we first restrict the discussion to two-dimensional space-filling curves. Later we will get back to higher dimensions.

To convert a plane-filling curve into sound, we consider the curve as a measure-preserving function f from a time interval $[0, 1]$ to a region in the plane, where we choose the desired length of the sound track as our unit of time. A practical approach to achieving this is the following. For any k , let C_k be the curve obtained by applying the generator recursively, starting from a line segment of unit length, to a recursion depth of k . Let $l(k, i)$ be the sum of the squared lengths of the first $i - 1$ segments of C_k . To calculate $f(t)$ for a given time t , we find a k and an i such that $l(k, i) = t$ (or sufficiently close to it for our purposes); the i -th vertex of C_k is the required point $f(t)$. (A more thorough approach can be found, for example, in my work on three-dimensional Hilbert curves [12].)

To convert a curve into sound, we sample the curve at regular time intervals, obtaining a series of points $f(t_1), f(t_2), \dots$, and map the coordinates of each point $f(t_i)$ to parameters of sound to be produced between time t_i and t_{i+1} . Thus, in contrast to Mason and Shaffle [24], we create sound from *point samples* that lie on the actual plane-filling curve, not from *line segments* of a sketch that approximates the curve. To implement our approach, several artistic decisions need to be made, as outlined below.

Choice of coordinate system, sampling rate and offset. Recall that we would like geometric similarity between different sections of the curve to be recognizable by ear. If rotations are by multiples of 90 degrees, this can be achieved by defining two voices, one for each coordinate in a standard Cartesian coordinate system, and traversing the curve while mapping each coordinate's values to the pitch of the corresponding voice. If rotations are by multiples of 60 degrees, we can use barycentric coordinates with respect to an equilateral triangle: this gives us three coordinates which we map to the pitch of three voices—see Figure 2 for an example. Geometric translations, reflections and rotations now correspond to what is known in music as transpositions, melody inversions and/or counterpoint inversions.

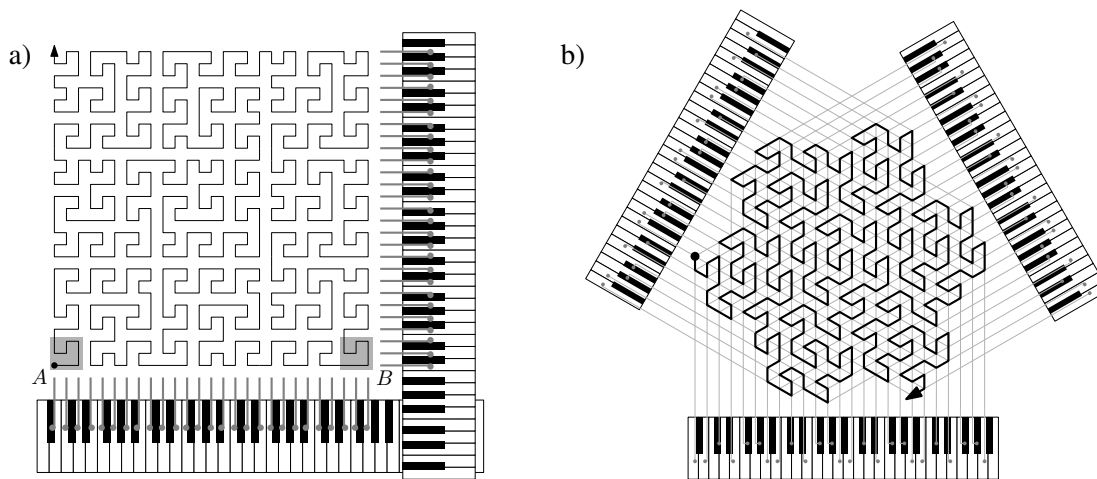


Figure 2: (a) Pitch mapping for the Meander curve [37]. (b) Pitch mapping for the Gosper flowsnake [9].

On each axis, the full range of coordinates must be mapped to pitch. One way to do this is a continuous mapping, which allows us to sample the curve at a very high rate so that we hear a continuous curve rather than a discrete approximation of a space-filling curve. This seems attractive conceptually, and indeed it facilitates listening to the sound at different levels of resolution: we may hear motifs of a few seconds, in which “notes” are “embellished” by similar motifs that are scaled down in both time and pitch range (some examples of sound tracks are available [11]). The downside is that the result is best described as “howling

winds”, “bee swarms” etc., and that my family forbids me to play this on the living room stereo.

For an experience that is not only interesting but also pleasant, it seems better to sample the curve at a lower frequency (5 Hz or less) and map the coordinates of the sample points to a familiar discrete set of pitch values, such as those from twelve-tone equal temperament, that is, what we find on a standard piano. Each sample point then produces a chord of fixed duration. I advise to choose the first sample, the sampling rate and the rotation of the curve such that consecutive samples always lie on a common line of the grid implied by the Cartesian or barycentric coordinate system. With Cartesian coordinates, this ensures that the two voices are each other’s rhythmic complement and do not mask each other. When using barycentric coordinates, this ensures that with each step, one voice remains constant while the other two voices move in opposite directions, thus avoiding potentially awkward parallel fourth and fifth intervals. This approach is similar to Tom Johnson’s work *Trio* [20], which can also be understood as a traversal of a triangular grid, mapped to three-note chords via barycentric coordinates. Sometimes there are multiple solutions (see Figure 3) that will sound very different.

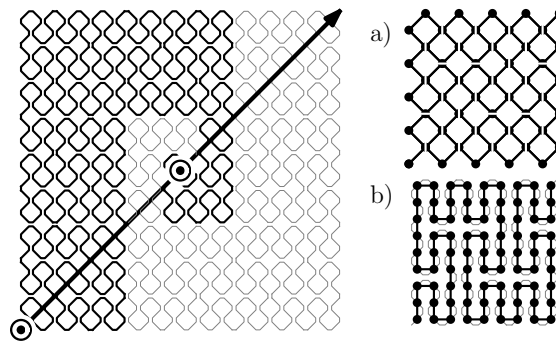


Figure 3: *Alternative sampling solutions for the Peano curve [28]. (a) Sampling at the end points of each curve section results in samples on a square grid, in which most sample points are visited twice, and in which consecutive edges of the traversal are always orthogonal to each other. (b) Sampling at the midpoints results in samples on a square grid that is rotated by 45 degrees with respect to the generating line segment, in which each grid point is visited exactly once, and the grid traversal uses vertical edges much more often than horizontal edges.*

Choice of musical scale. Note that one will get to hear all possible combinations of pitches for the two or three voices corresponding to the space filled by the curve. The selection of pitches that is used determines how many of those combinations will be dissonant and how large a range of coordinate values can be mapped. Before I implemented my approach, I actually expected that the result would inevitably be too dissonant to be practical. But to my surprise, this is not the case. Musical appeal benefits from the timely appearance and resolution of dissonances: one should have neither too many nor too few. I got the best results with simply using diatonic scales (such as the white keys of the piano) or with a hexatonic scale, obtained from the major (Ionian) diatonic scale by omitting the fourth note (for example, using all white keys except the Fs). Omitting the fourth note actually removes dissonances rather drastically: it eliminates all tritones and half of the minor seconds. Note that using these scales, as opposed to, for example, a twelve-tone or a whole-tone scale, also creates variation: translating a section of the curve will move the voices up or down in pitch and keep them recognizable, but the translation will typically change which steps in the melody are minor seconds (up or down by one piano key), which steps are major seconds (up or down by two keys) and which steps are minor thirds (up or down by three keys).

In practice, the use of diatonic scales limits the number of different pitch values that can be used effectively to about 35 (spanning five octaves), so that the number of sample points of a two-dimensional curve is limited to about 1000. To map larger integers to pitch, the sonification tool of the *On-Line Encyclopedia of*

Integer Sequences [26] takes the numbers modulo a number of discernable pitch classes. The *Audio Abacus* method by Walker et al. [35] maps large integers to sound digit by digit. Note, however, that these solutions render the mapping highly discontinuous and non-monotone: a small increment in coordinate value can result in a small increment in pitch, or in a maximum drop in pitch. In our setting, that would distort the desired correspondence between geometric and auditory congruence, so we keep to a strictly monotone mapping from single coordinates to single pitch values.

Choice of rotation, reflection, translation and direction. The choice of how to position the curve with respect to the chosen coordinate system, and in what direction to traverse the curve, determines among other things: which voices ascend and which voices descend when; which are the dominant musical motifs or rhythms (if the curve is not symmetric); and how the voices lie with respect to each other in each part of the sound track, and in particular, the harmonies with which the track starts and ends.

Choice of pace. To make sure that geometric congruence between curve sections corresponds to congruence of sound, we maintain a strict constant pace. It differs by curve what pace works best.

Choice of instrumentation. We choose instruments or sound colours for each voice so that they can be distinguished from each other.

Finish. The above choices enable us to obtain the “raw” material: a sound track that represents the plane-filling curve (some examples are available from my website [11]). However, every musician knows that merely hitting the notes will not necessarily bring out the individual voices well. To make it sound clearer and more appealing, we merge consecutive notes at the same pitch into longer notes, and we may add articulation, dynamics, and additional instruments by hand. As mentioned before, there will be many dissonances, and additional supporting voices can help in making these dissonances appear and resolve in a way that increases musical appeal, rather than reducing it.

However, the rules of the game as I choose to play it, are that pitch of the primary voices remains completely determined by the process described above: accidentals (incidental flat or sharp notes) are allowed only in the supporting voices, and the latter should avoid masking the primary voices that represent the curve. With two or three voices fixed in this way for the full duration of the sound track, the addition of harmonizing supporting voices may constitute quite a puzzle with severe restrictions on standard techniques such as modulation. Nevertheless, the puzzles seem to be quite possible to solve (see my website for a few proof-of-concept examples [11]).

Higher-dimensional curves. The whole approach generalizes naturally to curves that can be sampled on a d -dimensional Cartesian grid, resulting in raw material with d voices. On my website [11], I provide examples of three-dimensional Hilbert curves [12] traversing a $16 \times 16 \times 16$ grid of 4096 three-note chords; a four-dimensional Hilbert curve traversing an $8 \times 8 \times 8 \times 8$ grid of 4096 four-note chords, and a five-dimensional Hilbert curve traversing a $4 \times 4 \times 4 \times 4 \times 4$ grid of 1024 five-note chords. Note that in the five-dimensional example, each voice only uses four different pitch values: none of them covers a full octave of a diatonic or hexatonic scale. The artistic decision to make thus becomes less a matter of choosing a scale, and more a matter of choosing a different set of four pitch values for each voice, such that interesting combinations result.

Is This Like Music?

In the introduction I claimed that space-filling curves have a structure similar to music, being composed of sections that are transformations of each other. Indeed, “certain long-standing compositional procedures, such as those of canon, fugue, and motivic development, depend entirely on making new musical material by systematically transforming previous musical material. In many instances, these procedures result in clearly self-similar musical structures.” [7] “Formal manipulations such as retrograde (backward motion) or

inversion (inverting the direction of intervals in a melody) are found in the music of J. S. Bach and became the basis for the twentieth century twelve-tone (dodecaphonic) serial techniques of Arnold Schoenberg” [8].

A quantitative approach to the claim is offered by Voss and Clarke [34]. One can treat the pitches of the notes as a signal and analyse its power spectrum. Intuitively, the power spectral density as a function of the frequency f indicates how much the signal varies within time windows of duration $1/f$; more precisely, the power spectral density for a given frequency f is the square of the variance of the signal, averaged over all time windows of duration $1/f$.

For many types of signals the power spectral density as a function of f is proportional to $1/f^\beta$, for some value of β . The higher β , the stronger the correlation between successive pitches. With $\beta = 0$ there is no correlation, that is, white noise. With $\beta = 2$ there is strong correlation as in Brownian motion, which is considered boring [31]. Voss and Clarke found that the β value of the pitch signal (or of other aspects of sound) in music is typically close to 1, and that $1/f^1$ noise sounds more musical than $1/f^0$ noise or $1/f^2$ noise. Su and Wu [33] confirm this for low frequencies—we can roughly interpret this as: on the scale of musical phrases or larger sections. They find that for high frequencies (such as 0.1 Hz, that is, within musical phrases or smaller sections) β values range from 1.4 to 1.8 for a number of classical works for violin. Manaris et al. [22] investigated many genres of music and find β values (aggregated over different aspects of the music) ranging from 1.05 for jazz to 1.35 for modern romantic music, with outliers 0.82 for twelve-tone music and 1.53 for punk rock. Dodge exploited these correlations creatively by overlaying $1/f$ noise on different scales to create music that has similar structure on multiple scale levels [6].

Where do my sound tracks of space-filling curves fit in this picture? A d -dimensional space-filling curve visiting n grid points per second, will, with frequency f , move from one region of n/f grid points to an adjacent region that occupies an area that is at mean distance $\Theta((n/f)^{1/d})$ from the first. Thus, the variance in coordinates and pitch within time windows of size $O(1/f)$ will be $\Theta((n/f)^{2/d})$. Note that for $d = 2$, this means that the variance is linear in the number of notes in the time window—similar to Brownian motion². In general, the power spectral density as a function of f will be proportional to the square of the variance, that is $\Theta((n/f)^{4/d})$. So each voice of a sound track representing a d -dimensional space-filling curve is like $1/f^\beta$ noise with $\beta = 4/d$. This allows us to make the following predictions.

Two-dimensional curves by location mapping. A raw sound track of a two-dimensional space-filling curve ($d = 2$) would, supposedly, be as boring as Brownian motion: when played slow, the music is too monotonous within short time frames; when played fast, the music would move too wildly in long time frames (one might say that the musical potential for repetitions of longer sections is underused). This matches my (very) subjective impression after experimenting with many two-dimensional curves (and before calculating the spectral density function): some two-dimensional curves have musical potential and can be turned into music with proper hand-made “finish”—these curves make the endeavour worthwhile—but one has to look for these nice curves among many curves whose sound tracks do the job of representing the curves, but do not sound all that interesting by themselves from a musical point of view (at least not raw). Nevertheless, should you judge my (raw) sound tracks of selected two-dimensional space-filling curves [11] to sound musical, then this indicates that there is something that can make sound appear musical even if it does not have the statistical properties of $1/f^\beta$ noise with β well below 2.

Three- and four-dimensional curves by location mapping. The sound tracks of three-dimensional curves ($\beta \approx 1.3$) and four-dimensional curves ($\beta = 1.0$) would, supposedly, sound more musical. In contrast to two-dimensional curves, they should allow for a choice of pace that ensures enough variation in short time frames while keeping variation within practical limits in the long run.

Five- and higher-dimensional curves by location mapping. The melody of each single voice of a sound

²Indeed, investigating the similarity between plane-filling curves and random walks or Brownian motion has been done before, to analyse properties relevant to data structuring applications in computer science [3].

track of a higher-dimensional curve ($d > 4$, and therefore, $\beta < 1$) would sound less like ideal $1/f$ -noise and more like white noise: there may be too much variation and little coherence in short frames, and listening longer would not bring anything new. In particular, five-dimensional curves would result in melodies with a power spectrum similar to that of twelve-tone music. However, listening to the sound track of a five-dimensional curve [11], which I personally find quite pleasant, I find that power spectrum analysis of variation in pitch values is not applicable in this case. The reason is that for high-dimensional space-filling curves, one can only afford very few different pitch values per voice. For example, rendering a five-dimensional Hilbert space-filling curve on an $8 \times 8 \times 8 \times 8 \times 8$ grid would require us to play 32 768 successive chords, which is impossible to do at reasonable speed within reasonable time. Therefore, we would restrict the sampling of the curve to a $4 \times 4 \times 4 \times 4 \times 4$ grid, with only four different pitch values per note. As a result, we are not listening to melodies anymore, but really to the progression of rhythms and harmonies of the five voices. By lack of a well-defined distance measure for harmonies, we cannot easily apply the power spectrum analysis in this case.

Two-dimensional curves by direction mapping. We can also analyse the power spectrum of the pitch signal for two-dimensional curves that are mapped to one voice using the turning function. Assume the plane-filling curve consists of n sections that are congruent to each other. Let v_k be the variance of pitch within a sketch of $1/f = n^k$ segments. For a given $k \geq 1$, let μ_i be the mean pitch in the i -th section (of length n^{k-1}) of the sketch (of length n^k). Let q_i be the mean squared pitch in the i -th section. We have $v_{k-1} = q_i - \mu_i^2$ for each i , and $v_k = \frac{1}{n} \sum_{i=1}^n q_i - \left(\frac{1}{n} \sum_{i=1}^n \mu_i\right)^2$. Substituting $v_{k-1} + \mu_i^2$ for q_i , we obtain $v_k = v_{k-1} + \frac{1}{n} \sum_{i=1}^n \mu_i^2 - \left(\frac{1}{n} \sum_{i=1}^n \mu_i\right)^2$. Now, by the recursive construction of the curve, the differences in pitch (rotation) between the curve sections are independent of the level of recursion k , and thus we get $v_k = v_{k-1} + \Theta(1)$, and therefore, $v_k = \Theta(k) = \Theta(\log(1/f))$. Thus, the pitch signal resulting from the turning function approach has a power spectrum similar to $\log^2(1/f)$ noise, which is, asymptotically, more like white noise than any kind of $1/f^\beta$ noise with $\beta > 0$. The power spectrum analysis thus suggests the hypothesis that the sound track may have an unmusical lack of variation in the long run.

Readers may listen to Johnson's compositions [15, 16, 17] to judge for themselves: should this hypothesis perhaps be rejected? Admittedly, Johnson's sound tracks are minimalistic: there is much repetition, variation is sparse. In longer direction-mapped sound tracks of my own, I felt that variations in instrumentation are needed to make it bearable. Nevertheless, are Johnson's tracks not clearly recognizable as music? Is there perhaps a substantial difference between actual pitch variation and perceived pitch variation, because, in well-structured music, notes may have different weights in our perception, depending on their place in the melodic or harmonic structure?

Considerations for Further Work

Exploring the musical potential. So far, only the pitch of the leading voices on my sound tracks is determined by the space-filling curve. Other aspects of the sound, including timbre, dynamics, articulation, and accompanying voices, are deliberately unmusical (constant throughout) or created, more or less, by hand in a deliberate attempt to create music.

To judge whether the sound of space-filling curves in different dimensions can confirm or falsify the hypothesis that sound perceived as music has a power spectrum similar to $1/f$ -noise and well away from $1/f^2$ -noise, we would need to set up a controlled experiment in which the effects of curve-controlled aspects of sound are isolated from the effects of other aspects of sound. Moreover, we should isolate the effects of dimension from those of incidental properties of concrete curves. Currently I am not sure how to set up such an experiment, nor do I know what to expect from it. My preliminary impressions are that power spectrum analysis does in fact help in understanding the particular challenges of arranging or performing particular

styles of music (or space-filling curve sound tracks). However, I doubt that power spectrum analysis helps in distinguishing music from non-music, since I discovered several sound tracks that sound musical to me, at least for several minutes, while having supposedly unmusical power spectra.

Regardless of these theoretical considerations, I believe that my work confirms that space-filling curves have musical potential. It is quite possible to entertain one’s self constructing new space-filling curves in a quest to combine geometric appeal and musical appeal with a strong relation between geometry and sound (can one hear the curve fill the space?). “Raw” two-dimensional material can be used as a starting point to be completed with more traditional musical handwork. In higher dimensions, the focus of the artistic efforts shifts more towards to the decisions that determine the raw material: designing the space-filling curve itself (already in three dimensions, there are more than ten million self-similar space-filling curves that traverse a cube octant by octant [12]), and designing a selection of pitch values for each voice. For two-dimensional curves, in addition to the approach of the present paper, Johnson’s turning-function approach is available, potentially accommodating curves that cannot be sketched on a regular grid.

Is it useful? Hetzler and Tardiff [14] mention the “rule of four” in calculus instruction, which states that mathematics should be presented numerically, graphically, analytically, and by oral or written representations. They “expect that dovetailing sonification with the Rule of Four to create a Rule of Five would be a successful pedagogical approach.” Would a musical rendering of space-filling curves also help in understanding their mathematical properties? Clearly, sonification shares some limitations with visual illustrations. The choice of sample points (discretization) and the coordinate system influences what we see or hear: an unfortunate choice may leave distinguishing aspects of a curve invisible or inaudible.

Notwithstanding these limitations, by listening to my sound tracks I discovered aspects of some curves that did not catch my attention before. On my website one can download sound tracks of a curve in which consecutive line segments of the discretization are always orthogonal to each other—a property that is very audible. In my sound tracks of the five-dimensional Butz-Hilbert curve [4], its regular structure is clearly audible, with geometric rotations appearing as voices that exchange rhythms. It would be quite a challenge to illustrate this clearly with a “five-dimensional” drawing on paper.

How to bring out the recursive structure? In my examples, the similarity between sections at the same scale (in Figure 2(a), between squares of 9 notes, or between squares of 81 notes) is easy to hear. The similarity between sections at different scales is harder to pick up. With continuous pitch mapping I found this to be possible, for some scale levels, listening really carefully, but the result will be too far avant-garde to most people’s tastes. With the discrete mapping illustrated in Figure 2, can we hear that sections of 81 time units (and the whole piece) have a structure similar to that of sections of 9 time units? We might try to apply Tom Johnson’s techniques to bring out the self-similarity in single-voiced melodies [19]. However, our task is more challenging in two ways. First, we are dealing with (at least) two- or three-voiced motifs rather than single-voiced melodies, and second, the coarser versions of the motif are stretched not only in time but also in pitch. Already for the comparatively simple case of the second movement of Johnson’s composition *Counting Keys*, Johnson remarks: “the structure is pretty obvious when you see it, though it sometimes surprises me how much difficulty people have in perceiving this from just listening”. This confirms that making not only the repetitive structure, but also the *recursive* structure of space-filling curves clearly audible, would be a very ambitious goal indeed. So far, all compositions that I consider successful in this respect, are based on a space-filling curve whose generator has only two segments.

Perhaps solutions can be found by not refining all sections of the curve to the same depth, so that we can direct the attention of the listener to different scale levels at different times—which may also bring welcome variation in pace. Ventrella shows some examples of how the graphical equivalent of such an approach can generate drawings with beautiful variations in density ([32], the $\sqrt{4}$ -square family), and we may also try to achieve this in sound.

References

- [1] E. Amiot. Auto similar melodies. *J. Math and Music* 3(1):1–26 (2009)
- [2] M. Bader. *Space-filling curves: an introduction with applications in scientific computing*. Springer, 2013.
- [3] E. Bugnion, T. Roos, R. Wattenhofer, and P. Widmayer. Space filling curves versus random walks. In M. van Kreveld et al. (eds): *Algorithmic Foundations of Geographic Information Systems*, Lecture Notes in Computer Science 1340:199–211 (1997)
- [4] A. R. Butz. Alternative algorithm for Hilbert’s space-filling curve. *IEEE Trans. Computing* 20(4):424–426 (1971)
- [5] J. Cage, *Ryoanji*, Henmar Press, 1984.
- [6] C. Dodge. Profile: a musical fractal. *Computer Music J.* 12(3):10–14 (1988)
- [7] C. Dodge and C. R. Bahn. Musical fractals. *Byte* 11(6):185–196 (1986)
- [8] S. Dubnov and G. Surges. Delegating creativity: use of musical algorithms in machine listening and composition. Chapter 6 in Newton Lee (ed.), *Digital Da Vinci—Computers in Music*, Springer, 2014.
- [9] M. Gardner. Mathematical games—in which “monster” curves force redefinition of the word “curve”. *Scientific American*, 235(6):124–133 (1976)
- [10] S. W. Golomb. Replicating figures in the plane. *The Mathematical Gazette* 48(366):403–412 (1964)
- [11] H. Haverkort. *The sound of space-filling curves—examples*. <http://spacefillingcurves.net>, Section “Representations”.
- [12] H. Haverkort. *How many three-dimensional Hilbert curves are there?* Under review/revision for J. of Computational Geometry. arXiv:1610.00155 (2016)
- [13] T. Hermann, P. Meinicke, and H. Ritter. Principal curve sonification. In *Proc. Int. Conf. Auditory Display, 2000*, p81–86.
- [14] S. M. Hetzler and R. M. Tardiff. Two tools for integrating sonification into calculus instruction. In *Proc. Int. Conf. Auditory Display, 2006*, p281–284.
- [15] T. Johnson. *Formulas for String Quartet*, part 7. Editions 75, Paris, 1994.
- [16] T. Johnson. *Music for 88*, part 2: *Mersenne Numbers*. Editions 75, Paris, 1988. Recorded by the composer, Xi Records, 1992.
- [17] T. Johnson. *Rational Melodies*, numbers II and XIX. Editions 75, Paris, 1982. Recorded by Eberhard Blum, Hat Hut Records, 1993.
- [18] T. Johnson. *Self-similar melodies*, Editions 75, Paris, 1996, p83–92, p206–208.
- [19] T. Johnson. *Self-similar structures in my music: an inventory*. Lecture presented in the MaMuX seminar, Institut de Recherche et Coordination Acoustique/Musique, Paris, 14 Oct 2006.
- [20] T. Johnson, *Trio*, Editions 75, Paris, 2005.
- [21] A. Lucier. *Panorama*, Material Press, 1993.
- [22] B. Manaris, J. Romero, P. Machado, D. Krehbiel, T. Hirzel, W. Pharr, and R. B. Davis. Zipf’s law, music classification and aesthetics. *Computer Music J.* 29(1):55–69 (2005)
- [23] S. Manousakis. *Musical L-systems*. Master’s thesis, Royal Conservatory, Den Haag, 2006.
- [24] S. Mason and M. Saffle. L-systems, melodies and musical structure. *Leonardo Music J.* 4:31–38 (1994)
- [25] G. L. Nelson. Real time transformation of musical material with fractal algorithms. *Computers and Mathematics with Applications* 32(1):109–116 (1996)

- [26] *Listen to a Sequence in the on-line encyclopedia of integer sequences*, <http://oeis.org/play.html>, accessed 27 Feb 2017.
- [27] G. Pareyon. *On musical self-similarity—Intersemiosis as synecdoche and analogy*. Acta Semiotica Fennica XXXIX, Approaches to Musical Semiotics 13. Int. Semiotics Institute, Imatra, 2011, p443–449.
- [28] G. Peano. Sur une courbe, qui remplit toute une aire plane. *Math. Ann.* 36(1):157–160 (1890)
- [29] P. Prusinkiewicz. Score generation with L-systems. In *Proc. 1986 Int. Comp. Music Conf.*, p455–457.
- [30] H. Sagan. *Space-filling curves*. Springer, 1994.
- [31] M. Schroeder, *Fractals, chaos, power laws: minutes from an infinite paradise*. W. H. Freeman, 1991.
- [32] J. Ventrella. *Brainfilling curves: a fractal bestiary*. Self-published / Eyebrian Books, 2012.
- [33] Z.-Y. Su and T. Wu. Music walk, fractal geometry in music. *Physica A* 380:418–428 (2007)
- [34] R. F. Voss and J. Clarke. “ $1/f$ noise” in music: music from $1/f$ noise. *J. Acoustical Soc. America* 63(1):258–263 (1978)
- [35] B. N. Walker, J. Linday, and J. Godfrey. The audio abacus: representing a wide range of values with accuracy and precision. In *Proc. Int. Conf. Auditory Display, 2004*.
- [36] R. Wallin. *Fractal music—red herring or promised land? ... or “Just Another of those Boring Papers on Chaos”*. Article based on a lecture given at the Nordic Symp. for Computer Assisted Composition 1989, <http://www.rolfwallin.org>, accessed 7 Feb 2017.
- [37] W. Wunderlich. Über Peano-Kurven. *Elemente der Mathematik* 28(1):1–24 (1973)
- [38] R. Zare. *Fractal miniatures*. Self-published, 2012.
- [39] R. Zare. In the programme of the *RED NOTE New Music Festival Composition Competition—Category B (Chamber Orchestra)*, 2015.