The Sound of Space-Filling Curves

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

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 suggest that the answer depends on the number of dimensions of the space-filling curve.

Introduction

Space-filling curves are curves that are so crinkly that they completely fill a higher-dimensional space. Formally, they can be understood as continuous, surjective mappings from $[0, 1]$ to a subset of $\mathbb{R}^d$ that has d-dimensional Lebesgue measure greater than zero—for example a d-dimensional cube. In this paper, we take the following approach [12] to describing such curves. We consider line segments to have a direction (forward or backward) and an orientation (left or right). A plane-filling curve (a two-dimensional space-filling curve) is defined by a *generator* that describes how to replace a left-forward line segment from $(0, 0)$ to $(1, 0)$ by a chain 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. Under certain conditions, the curve will be plane-filling, going through every point of a region of non-zero area. The concepts and definitions generalize naturally to higher dimensions.

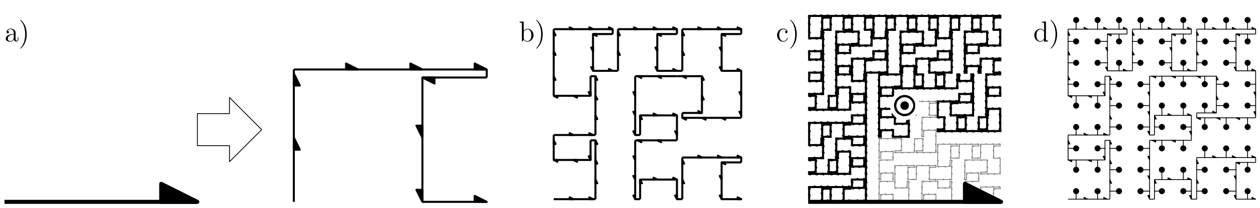

Figure 1 : *(a) Generator for the Meander curve [15]. (b) Result of applying it recursively to depth two. (c) If we start from a line segment, apply the generator recursively to infinite recursion depth, and sample the curve at 3/4 of its length, we get a sample point that lies to the left of the original line segment. (d) The samples for all segments of a given refinement level fill a square grid.*

The goal of the present work is to create a way of representing space-filling curves by sound that adds to the perception of their beauty. Our representations should be accurate and pleasant: it should be possible to reconstruct the curves from the sound (at least in theory), geometric similarity between different sections of the curve should correspond to similarity between sections of the sound track, and listening to the sound, possibly while tracing a figure or an animation of the curve, should be an entertaining experience. Note that our initial goal is not to use space-filling curves to produce sound or music that is interesting by itself, 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. Some composers have created music from the contours of geometric objects in the plane by mapping one coordinate axis to time and the other axis to pitch. However, a space-filling curve inevitably moves back and forth in all dimensions, and sound cannot go back and forth in time. Johnson [5] circumvents this problem by creating music from the sequence of left and right turns along a sketch of a plane-filling curve. This solution has its merits, but does not generalize to three- or higher-dimensional curves, and geometric translation is not represented in the sound. Building on an idea from Prusinkiewicz [11], put into practice by Nelson [10], Mason and Saffle [9] proposed to turn a sketch of a plane-filling 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. 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. 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, the voices are out of sync by four time units, so that the horizontal segments of square B are played in counterpoint with vertical segments from the next square. This conflicts with our aim to preserve similarity between curve sections, so in this paper, we will take a different approach.

Below is a brief description of my work. More details and references can be found on my website [2].

Converting Space-Filling Curves into Sound

We first restrict the discussion to two-dimensional curves. We treat a plane-filling curve as a function f from a time interval $[0, 1]$ to a region in the plane, where the desired length of the sound track is our unit of time. 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 $f(t)$. (For a more thorough approach, see other work [3].) To convert the curve into sound, we sample the curve at regular time intervals, obtaining a series of points $f(t_1), f(t_2), ...,$, 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}. To implement our approach, several artistic decisions need to be made:

Choice of coordinate system, sampling rate and offset. Recall that we would like similarity between different sections of the curve to be recognizable by ear. If rotations in the curve are by multiples of 90 degrees, we define two voices, one for each coordinate in a Cartesian coordinate system, and traverse the curve while mapping each coordinate to the pitch of the corresponding voice. If rotations are by multiples of 60 degrees, we use barycentric coordinates and map these to three voices—see Figure 2. Translations, reflections and rotations now correspond to musical transpositions and inversions. A continuous mapping from coordinates to pitch would allow us to sample at a very high rate so that we hear a continuous curve rather than a discrete approximation (examples are available [2]). However, for pleasant listening, we sample at a rate of 10 Hz or less and map coordinates to discrete pitch values, such as those from a piano. Each sample point thus 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 have one coordinate in common. Thus, with Cartesian coordinates, the two voices are each other's rhythmic complement and do not mask each other; with barycentric coordinates, we avoid (as in *Trio* [7]) that voices move in parallel.

Choice of musical scale. One will get to hear all combinations of pitches corresponding to the space filled by the curve. For musical appeal one should have neither too many nor too few dissonant combinations. I got the best results with diatonic and hexatonic scales (such as the white keys of the piano, possibly omitting F).

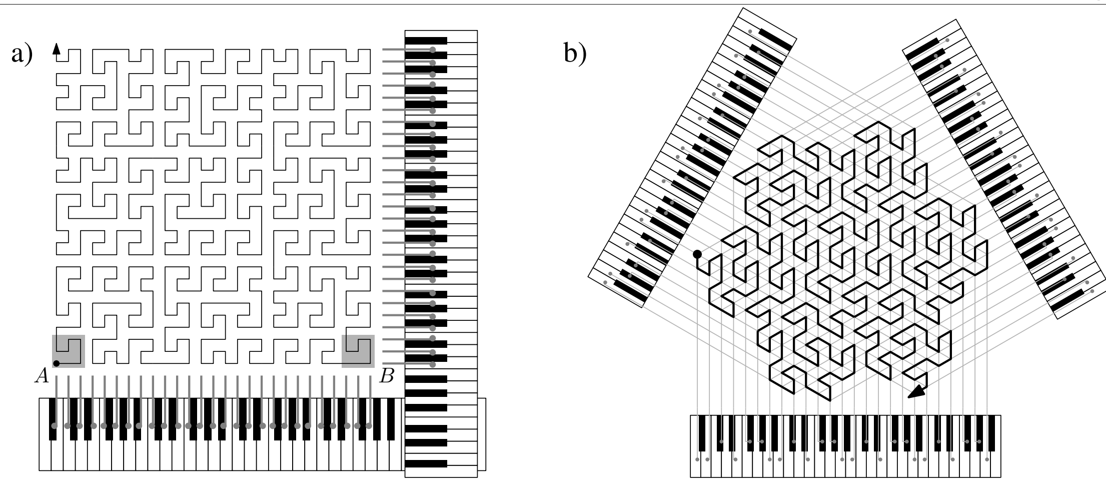

Figure 2 : *(a) Pitch mapping for the Meander curve [15]. (b) Pitch mapping for the Gosper flowsnake [1].*

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, affects which voices ascend and descend when; which motifs dominate; and how the voices lie with respect to each other in each part of the sound track; and, in particular, with what harmonies 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 so that the voices can be distinguished from each other.

Finish. The above choices result in the "raw" material: a sound track that represents the curve. To make it clearer and more appealing, we merge consecutive notes at the same pitch into longer notes, and we may add articulation and dynamics. Additional voices may help in making dissonances appear and resolve in a way that increases musical appeal (see my website for a few proof-of-concept examples [2]).

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 [2], I provide examples of three-, four- and five-dimensional curves. In the five-dimensional example, each voice only uses four different pitch values. The artistic decision to make thus becomes not a matter of choosing a scale, but rather a matter of choosing five sets of four pitch values, such that interesting combinations result.

Discussion and Considerations for Further Work

Is this like music? Voss and Clarke [14] studied the succession of pitches of notes as a signal. For many types of signals the power spectral density as a function of the frequency f is proportional to $1/f^{\beta}$, for some value of β. A β value of 0 is typical of white noise; $\beta = 2$ indicates strong correlation between successive pitches, as in Brownian motion. Voss and Clarke found that $1/f^1$ noise sounds more musical than $1/f^0$ or $1/f^2$ noise, and, as confirmed later [8, 13], that the pitch signal in music typically has a β value close to 1.

A space-filling curve visiting n grid points per second moves, with frequency f, from one region of n/f grid points to an adjacent region that lies at mean distance $\Theta((n/f)^{1/d})$ from the first. Thus, the variance in coordinates and pitch within time windows of size $\Theta(1/f)$ is $\Theta((n/f)^{2/d})$, and the power spectral density as a function of f is $\Theta((n/f)^{4/d})$. So each voice of a sound track of a d-dimensional curve is like $1/f^{\beta}$ noise with $\beta = 4/d$. Therefore, a sound track of a two-dimensional curve ($d = 2$) may be as boring as Brownian motion. Should you nevertheless judge my raw sound tracks of selected two-dimensional curves [2] to sound musical, then there must be 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. A similar analysis [2] can be applied to sound tracks created with Johnson's approach [5], that is, from left and right turns. Three- and four-dimensional curves ($\beta \approx 1.3$, $\beta = 1.0$) would, supposedly, sound more musical. With even higher-dimensional curves, each voice uses only a few different pitch values. Thus, we hear a progression of rhythms and harmonies rather than melodies, and power spectrum analysis of pitch variation does not seem meaningful.

The above discussion shows that assessing how musical the sound of space-filling curves is, may have some relevance in testing Voss and Clarke's hypothesis. To really do so, we would need to set up a controlled experiment in which the effects of dimension are isolated from incidental properties of concrete curves, and in which the effects of curve-controlled aspects of sound (that is, pitch) are isolated from the effects of other aspects of sound that are not governed by the curve. Regardless of these theoretical considerations, I believe that my work confirms that space-filling curves have musical potential, and there are curves that can combine geometric appeal and musical appeal with a strong relation between geometry and sound.

Is it useful? Hetzler and Tardiff [4] mention the "rule of four" in calculus instruction (mathematics should be presented numerically, graphically, analytically, and by oral or written representations), and they propose to add sonification. Could musical rendering of space-filling curves indeed help in understanding their mathematical properties? From my website one can download sound tracks of a curve in which consecutive line segments are always orthogonal to each other—a property that is very audible and which I was not aware of before. In my sound tracks of a five-dimensional Hilbert curve, its regular structure is clearly audible. 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 is easy to hear. But can we also hear, in the example of Figure 2(a), that sections of 81 time units 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 [6]. However, our task is more challenging as we are dealing with (at least) two- or three-voiced motifs rather than single-voiced melodies.

References

[1] M. Gardner. Mathematical games, December 1976. *Scientific American*, 235(6):124–133 (1976)

[2] H. Haverkort. *The sound of space-filling curves.* http://spacefillingcurves.net, Sec. "Representations".

[3] H. Haverkort. *How many three-dimensional Hilbert curves are there?* arXiv:1610.00155 (2016)

[4] S. M. Hetzler and R. M. Tardiff. Two tools for integrating sonification into calculus instruction. *Proc. Int. Conf. Auditory Display, 2006*, p281–284.

[5] T. Johnson. *Self-similar melodies*, Editions 75, Paris, 1996, p83–92, p206–208.

[6] T. Johnson. *Self-similar structures in my music: an inventory.* Lecture at IRCAM, Paris, 14 Oct 2006.

[7] T. Johnson, *Trio*, Editions 75, Paris, 2005.

[8] 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)

[9] S. Mason and M. Saffle. L-systems, melodies and musical structure. *Leonardo Music J.* 4:31–38 (1994)

[10] G. L. Nelson. Real time transformation of musical material with fractal algortihms. *Computers and Mathematics with Applications* 32(1):109–116 (1996)

[11] P. Prusinkiewicz. Score generation with L-systems. *Proc. 1986 Int. Computer Music Conf.*, p455–457.

[12] J. Ventrella. *Brainfilling curves: a fractal bestiary.* Self-published / Eyebrain Books, 2012.

[13] Z.-Y. Su and T. Wu. Music walk, fractal geometry in music. *Physica A* 380:418–428 (2007)

[14] R. F. Voss, J. Clarke. "$^1/_f$ noise" in music: music from $^1/_f$ noise. *J. Acoust. Soc. Am.* 63:258–263 (1978)

[15] W. Wunderlich. Über Peano-Kurven. *Elemente der Mathematik* 28(1):1–24 (1973)