Open Research Online
The Open University's repository of research publications
and other research outputs

Scratching the scale labyrinth

Book Section

For guidance on citations see FAQs.

Version: Accepted Manuscript

oro.open.ac.uk

Scratching the Scale Labyrinth

Andrew J. Milne[1], Martin Carlé[2], William A. Sethares[3], Thomas Noll[4], and
Simon Holland[1]

[1] The Open University, Milton Keynes, MK7 6AA, UK
[2] Humboldt-Universität zu Berlin, Germany
[3] University of Wisconsin-Madison, WI 53706, USA
[4] Escola Superior de Música de Catalunya, 08013 Barcelona, Spain

Abstract. In this paper, we introduce a new approach to computer-aided microtonal improvisation by combining methods for (1) interactive scale navigation, (2) real-time manipulation of musical patterns and (3) dynamical timbre adaption in solidarity with the respective scales. On the basis of the theory of well-formed scales we offer a visualization of the underlying combinatorial ramifications in terms of a *scale labyrinth*. This involves the selection of generic well-formed scales on a binary tree (based on the Stern-Brocot tree) as well as the choice of specific tunings through the specification of the sizes of a period (pseudo-octave) and a generator (pseudo-fifth), whose limits are constrained by the actual position on the tree. We also introduce a method to enable transformations among the modes of a chosen scale (generalized and refined "diatonic" and "chromatic" transpositions). To actually explore the scales and modes through the shaping and transformation of rhythmically and melodically interesting tone patterns, we propose a playing technique called Fourier Scratching. It is based on the manipulation of the "spectra" (DFT) of playing gestures on a sphere. The coordinates of these gestures affect score and performance parameters such as scale degree, loudness, and timbre. Finally, we discuss a technique to dynamically match the timbre to the selected scale tuning.

Keywords: MOS Scales, Well-Formed Scales, Diatonic, Chromatic, Stern-Brocot Tree, Farey Sequence, Fourier Scratching

1 Introduction

The scale labyrinth is a visualization of a widely studied class of musical scales, which form a deeply structured and interconnected scale universe. The terminology applied in this paper originates from two separate strands of music-theoretical discourse. The term *moment of symmetry scale* (MOS) is rooted in the field of investigations into musical tunings, while the term *well-formed* (WF) scales is common in algebraic scale theory. The present paper attempts to merge these two traditions within an experimental paradigm. For now, however, the terms are used interchangeably.

Well-formed scales are generated by a single interval and contain exactly two different step sizes (large and small) that are maximally evenly distributed. The

case where these step sizes coincide is included as a degenerate instance. Each scale has a *valid tuning range* over which it maintains a structural identity but over which the sizes of the two steps co-vary. At any given tuning, each scale is embedded within a unique family of scales with successively larger numbers of tones (and successively smaller valid tuning ranges). Carey [1] denotes the class of all well-formed N-note scales as $WF(N, g)$ where g is the factor that converts generator order into scale step order (mod N). For example, the diatonic scale and its inverse belong to $WF(7, 2)$ while the chromatic scale and its inverse belong to $WF(12, 7)$. These relationships are represented visually in Sect. 4.

Visualizing the structure of these scales is useful from an analytic point of view, and it can also function as a graphical user interface (GUI) object for use in musical applications. An example of an analytical application is given in Figure 5 in [2], where a scale labyrinth is used to indicate MOS scale tunings that provide good approximations of just intonation. A concrete example of use in a GUI is given in Sec. 5, where the scale labyrinth allows a musician to choose, simultaneously, a scale structure (number of small and large steps) and its tuning (the sizes of its period and generator).

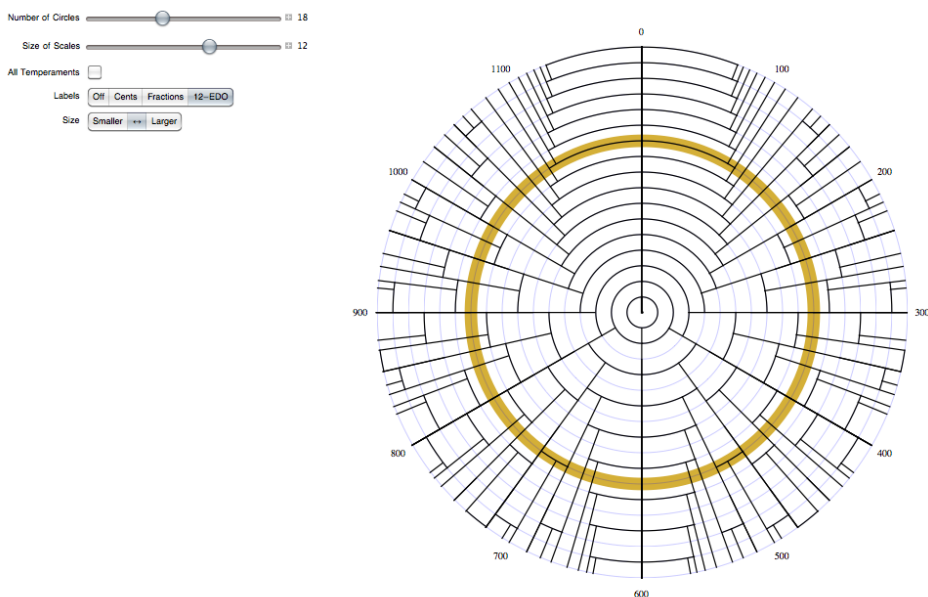**Fig. 1.** This scale labyrinth, generated interactively in Mathematica, can be adjusted in several ways to emphasize certain aspects of the tunings as described in Sect. 4. The program can be downloaded from `http://homepages.cae.wisc.edu/~sethares/ MOSLabyrinth.nbp`.

The scale labyrinth is related to a Stern-Brocot tree (Fig. 4), which is a systematic enumeration of the rational numbers. In the case of the labyrinth,
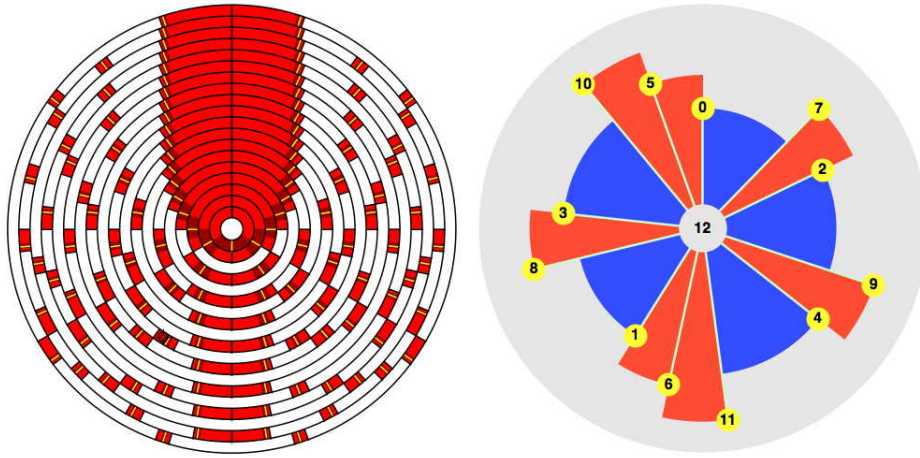
Fig. 2. This scale labyrinth can be used directly to control the tuning of an interactive algorithmic music application, as described in Sect. 5. See also http://www.youtube.com/watch?v=OBUUcjLbCIk for a short movie.

Fig. 3. The scale labyrinths of Figures 1 and 2 have a remarkable symmetry and are reminiscent of labyrinths of Gothic architecture, such as the famous Chartres cathedral labyrinth shown here (original photograph by Mich De Mey `http://www.flickr.com/photos/dumbo/2555996059/`).

the tree has been bent into a circular shape so that 0/1 and 1/1 lie at the same location. In the scale labyrinth, each nested circle corresponds to a fraction with denominator equal to the radius of the circle; in the Stern-Brocot tree, each row contains fractions with different denominators. The Stern-Brocot tree has been used by Erv Wilson to illustrate the structure of MOS scales using scale tree diagrams [3]. More recently, Holmes [4] produced a version, rotated by 90 degrees, where the left-right position of each fraction was determined by its denominator. Inspired by this diagram, Milne [2] produced a circular version to visually emphasize the period (octave) equivalence so that rotation by 360 degrees corresponds to the period.
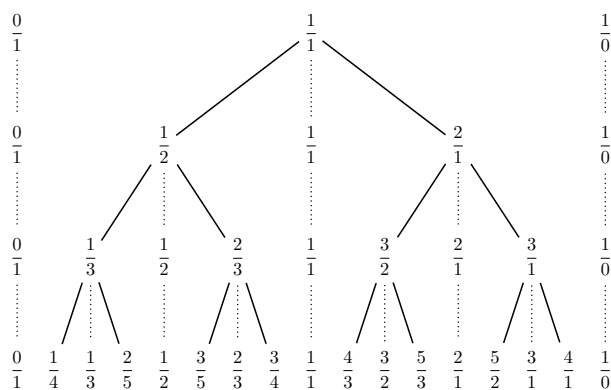

Fig. 4. The Stern-Brocot tree.

2   Some Properties of MOS/Well-formed Scales

There are several alternative but equivalent definitions: an MOS/well-formed scale is a *generated scale* containing exactly two step sizes that are distributed with maximal evenness.[5] A generated scale is produced by repeatedly adding a *generator* interval (typically a perfect fifth) and then reducing all such intervals by repeatedly subtracting a *period* interval (typically the octave) so all intervals are smaller than the period. The number of times the generator can be stacked so as to produce just two evenly distributed step sizes depends on the ratio of the generator and period. For example, if the generator is 702 cents and the period is 1200 cents (a generator/period ratio of 0.585), the two-note scale C-G belongs to scale class $WF(2,1)$, the 3-tone scale C-D-G belongs to $WF(3,2)$, the 5-tone scale C-D-E-G-A belongs to $WF(5,2)$, the 7-tone scale C-D-E-F♯-G-A-B belongs to $WF(7,2)$, the 12-tone chromatic scale belongs to $WF(12,7)$, $WF(17,12)$ is a 17-tone scale, and so forth [5]. Different generator/period ratios

---

[5] That is, the distribution of the two steps forms a Christoffel word.

require different numbers of tones to produce well-formed scales. For example, when the generator is 316 cents (a just intonation "minor third") and the period is 1200 cents, the following numbers of tones produce well-formed scales: 2, 3, 4, 7, 11, 15, 19, and so forth. Whenever the size of the period is not explicitly mentioned, it is assumed to be 1200 cents.

Each MOS scale can be characterized by its number of large (L) and small (s) steps, so the familiar (anhemitonic) pentatonic scale can be characterized with the signature 2L, 3s, the familiar diatonic scale by the signature 5L, 2s. For any given generator/period ratio, the number of tones required to form an MOS scale and the range of tunings over which any such scale maintains its identity (i.e., its number of large and small steps is invariant), is given by the Stern-Brocot tree (as discussed in the next section).

Well-formed scales have a number of properties that are thought to give them aesthetic value. For example: every scale span (generic interval size) occurs in exactly two interval sizes (Myhill's property [6]). The two scale step sizes are evenly distributed throughout the period. Within the period, every scale degree has a unique pattern of intervals surrounding it [7], which helps support tonal functionality. When transposed by the generator, the resulting scale shares all but one tone, facilitating modulation [7]. Collectively, these features suggest a good compromise between the excessive simplicity of equal step scales and the complexity of completely irregular scales [8].

Western theory recognizes the first five fifth-generated well-formed scales: authentic division of the octave, tetractys, the pentatonic, diatonic and chromatic.[6] But there are a number of MOS scales that, due to their microtonal intervals, are unfamiliar and may be difficult to play on standard instruments. Interestingly, a number of such well-formed scales contain numerous intervals and chords that approximate consonant just intonation intervals and chords as effectively as the familiar diatonic scale [9]. And, as discussed in Sect. 6, it is possible to use synthesizers capable of spectral retuning to minimize sensory dissonance for any MOS scale at any tuning [10]. It seems, therefore, that with novel musical controllers and synthesizers, the musical possibilities of MOS structures may become more accessible to musicians and composers. This paper discusses one such application.

Some intriguing features of scales with the MOS/well-formed structure are:
Co-prime step numbers: In every MOS, the number of small steps and number of large steps (in each period) is always co-prime. For example, the pentatonic scale is 2L, 3s; the diatonic is 5L, 2s. There is no MOS scale with, for example, 2 large steps and 4 small steps (within the period).
Inverse scales: Every MOS scale has an "inverse" form, where the number of large and small steps swaps. For example, the diatonic scale (5L, 2s) has an inverse, the *anti-diatonic* scale, which is 2L, 5s.
Landmark equal tunings: As the tuning of generator changes, the sizes of the small and large steps co-vary. For example, when the generator is 700 cents, the

---

[6] Strictly speaking, the chromatic scale in 12-TET is a *degenerate* well-formed scale because its two step sizes are identical [5].

diatonic scale's large steps (major seconds) are 200 cents, its small steps (minor seconds) are 100 cents; when the generator is 710 cents, the large steps are 220 cents, the small are 50 cents (in all cases, $5 \times$ large step size $+ 2 \times$ small step size = period size).

The co-varying step sizes produce three "landmark" tunings: a) the tuning where the small and large steps become the same size (this tuning marks the transition between an MOS scale and its inverse), b) the tuning at which the small steps of the MOS shrink to zero, c) the tuning at which the small steps of the inverse MOS scale shrink to zero. All three landmark tunings are equal temperaments: the cardinality of a) is the number of L steps plus the number of s steps, the cardinality of b) is the number of L steps, and the cardinality of c) is the number of s steps. For example, the diatonic 5L, 2s scale meets its inverse (the anti-diatonic scale) at 7-TET (685.714 cents), where the large and small steps become identically sized; the diatonic is also bounded at 5-TET (720 cents), where its two small steps shrink to zero size; the anti-diatonic is also bounded at 2-TET (600 cents), where its five small steps shrink to zero.

Embeddings: Every MOS scale with signature pL, qs is embedded in a family of MOS scales. The lowest cardinality embedding scale has $2p + q$ steps. For example, the lowest cardinality MOS scale that embeds the 5L, 2s diatonic scale has $2 \times 5 + 2 = 12$ tones; the lowest cardinality scale that embeds the 2L, 5s anti-diatonic has $2 \times 2 + 5 = 9$ tones. A method to determine the tuning of this embedding scale using the Stern-Brocot tree is given in Sect. 3.

Coherence within a well-defined tuning range: A scale is *coherent* [11] or *proper* [12] if there is a monotonic relationship between that scale's generic interval sizes and its specific interval sizes. This requires, for example, that every fifth be larger than every fourth, which are larger than every third, which are larger than every second, which in turn are larger than the unison. Well-formed scales are coherent over the tuning range within which the ratio of L to s (Blackwood's R [13]) is less than 2. This is precisely the tuning at which the lowest cardinality embedding scale is equally tuned. For instance, the 5L, 2s diatonic scale is coherent between 4/7 and 7/12. The anti-diatonic scale 2L, 5s is coherent between 5/9 and 4/7.

3   The Stern-Brocot Tree

The Stern-Brocot tree [14] [15] is a systematic enumeration of the rational numbers independently discovered in the 19th century by the mathematician Moritz Stern and the watchmaker Achille Brocot. Stern's focus was mathematical whereas Brocot's focus was the specification of gear ratios for clock design. The tree provides a method to iteratively generate all rational numbers, in reduced form, exactly once. In Fig. 4, note that the top row consists of the three rationals 0/1, 1/1, and 1/0; the next row down consists of the mediants of each adjacent pair of the above row, which gives the rationals 1/2 and 2/1 (the *mediant* of a/b and c/d is $(a + c)/(b + d)$, where all fractions are reduced). The iterative process of populating each new row with the mediants of the adjacent pairs above produces the Stern-Brocot tree, and all fractions are ordered, by size,

from left to right. The following paragraphs show the correspondences between the Stern-Brocot tree and the scale properties described in the previous section.

In the Stern-Brocot tree, each fraction can be thought of as representing a generator/period tuning ratio. Because these are co-prime ratios n/d, they generate an equal division of the period of cardinality d. For instance, 4/7 can represent a generator of 685.714 cents and a period of 1200 cents, and thus generates 7-TET. With this in mind, the tree can be used to provide the precise landmark tunings, range of coherence, and embeddings of any well-formed scale. Fig. 5 provides an illustration of the methods described below.
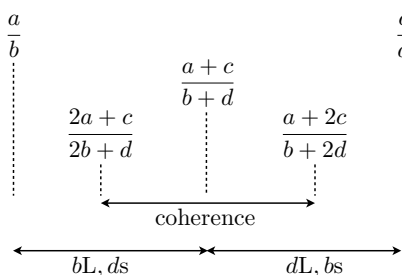
$$\frac{a}{b} \qquad\qquad\qquad \frac{c}{d}$$

$$\frac{a+c}{b+d}$$

$$\frac{2a+c}{2b+d} \qquad\qquad \frac{a+2c}{b+2d}$$

coherence

$$b\mathrm{L}, d\mathrm{s} \qquad\qquad d\mathrm{L}, b\mathrm{s}$$

Fig. 5. A well-formed scale's landmark tunings, its range of tuning coherence, and its lowest cardinality embedding can be read directly off the Stern-Brocot tree.

Consider two fractions in the same row of the tree, $\frac{a}{b}$ and $\frac{c}{d}$, and their mediant $\frac{a+c}{b+d}$. The interval between $\frac{a}{b}$ and $\frac{a+c}{b+d}$ corresponds to the valid (generator/period) tuning range of an MOS scale with b large steps and d small. There must be d small steps because the boundary tuning at $\frac{a}{b}$ has b tones, which implies that d small steps have shrunk to zero size. Conversely, the interval between $\frac{a+c}{b+d}$ and $\frac{c}{d}$ corresponds to the valid tuning range of the inverse MOS scale, which has d large steps and b small—the boundary tuning at $\frac{c}{d}$ has d tones, so b small steps have shrunk to zero size. The tuning at precisely $\frac{a+c}{b+d}$ is $(b+d)$-TET, so this is the tuning at which the "large" and "small" steps become equivalent in size and the MOS meets its inverse.

For example, the fractions $\frac{1}{2}$ and $\frac{3}{5}$ have a mediant $\frac{4}{7}$, so the interval between $\frac{1}{2}$ and $\frac{4}{7}$ is the valid tuning range of an MOS with 2 large and 5 small tones (the anti-diatonic), while the interval between $\frac{4}{7}$ and $\frac{3}{5}$ is the valid tuning range of its inverse—the diatonic 5L, 2s. At precisely $\frac{4}{7}$, the large and small steps are equally sized and the two scales meet. Each triple of fractions made from two adjacent fractions and the fraction between them on the row below corresponds, therefore, to the three landmark tunings: The central value gives the generator/period ratio where the MOS scale meets its inverse, the outer values give the tunings at which the small steps of the MOS, or its inverse shrink to zero.

The lowest cardinality scales within which any MOS is embedded can be identified by the mediant of its two boundary tunings. As above, the boundary tunings of bL, ds are $\frac{a}{b}$ and $\frac{a+c}{b+d}$, and their mediant is $\frac{2a+c}{2b+d}$. This means the

embedding scale contains $2b + d$ tones, with either $b + d$ large steps and b small, or b large and $b + d$ small, depending on the tuning. For example, the diatonic scale (whose boundaries are 4/7 and 3/5) is embedded within their mediant, which is 7/12 (which represents the scale 7L, 5s or its inverse 5L, 7s). Similarly the anti-diatonic whose boundary tunings are 1/2 and 4/7 is embedded within their mediant, which is 5/9 (7L, 2s, or 2L, 7s).

The tuning range over which bL, ds, and its inverse dL, bs, are coherent is bounded by the tunings at which their respective lowest cardinality embedding scales are equally tuned. That is, at $\frac{2a+c}{2b+d}$ and $\frac{a+2c}{b+2d}$. For instance, the diatonic has boundary tunings of 4/7 and 3/5, and their mediant (embedding scale) is 7/12, so the range over which the diatonic scale is coherent is 4/7 to 7/12; similarly, the anti-diatonic has boundary tunings at 1/2 and 4/7 with a mediant (embedding scale) of 5/9, so the range over which it is coherent is 5/9 to 4/7, which can be gleaned from [5].

4   Reading the Interactive Labyrinths

Depending on its purpose, different visualizations of the scale labyrinth may be preferred. Figure 1 displays several kinds of information that may be useful in a detailed analysis while the simpler Fig. 2 may be more appropriate as an interface element for the purpose of choosing a scale and tuning. In Fig. 1, the angle indicates the ratio between the generator and the period—the top of the circle represents a generator/period ratio of zero, the bottom of the circle a ratio of 1/2, the 11 o'clock position, a ratio of 11/12. For example, the familiar 12-tone equal tempered diatonic scale, which requires a generator of 700 cents (7 semitones) and a period of 1200 cents (12 semitones), can be found at the location $700/1200 = 7/12 = 0.583$. Figure 1 is labeled in cents (assuming a period of 1200) but, in the interactive version, fractions can be displayed instead.

Note that the structure is left-right symmetric, because a scale created using any generator is identical to the scale generated by the complement of the generator within the period (e.g., precisely the same scale is produced by generators of 697 cents and $1200 - 697 = 503$ cents). Each ring corresponds to the set of MOS scales that contain the numbers of notes indicated by its integer label (the dot at the center of the circle is ring 1). Each radial line (spoke) extending inwards from the edge of the circle represents an equal temperament scale. For each spoke, there is some circle that the spoke touches but does not cross. This circle gives the number of notes in the corresponding equal step scale. The angle of the spoke gives the tuning of its generator relative to the period. For instance, at 700 cents, there is a spoke that extends from the edge of the circle to the 12th ring, which indicates that this tuning produces 12-TET. Similarly, at 685.714 cents ($685.714/1200 = 4/7$) there is a spoke extending to the 7th ring, illustrating that this tuning produces 7-TET.

As described in Sect. 2, as the generator/period tuning ratio changes, the small and large steps co-vary across three landmark tunings, which are visually prominent in the labyrinth. Using the 7-note diatonic and anti-diatonic scales as

an example, focus on the seventh ring as in Fig. 6, which zooms in about the 700 cent region. There are two lines which cross (rather than merely touch) this part of the 7-ring, those at 600 and 720 cents. These delimit an arc which corresponds to the possible tunings of a 7-note MOS scale, (in this case the diatonic) and its inverse (in this case the anti-diatonic). The 7-TET spoke at 685.7 cents is the only spoke which meets (but does not cross) this arc. This spoke represents the point at which the sizes of the small and large steps equalize before reversing roles, and so marks the tuning at which the diatonic and anti-diatonic meet.
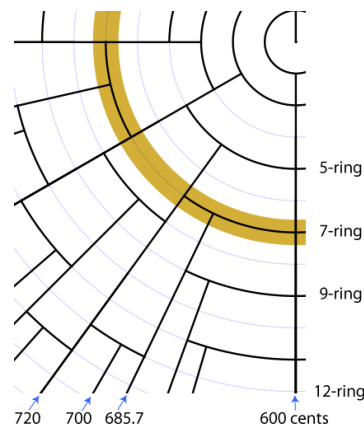

Fig. 6. This is a zoom into the scale labyrinth of Figure 1. Extra annotations have been added to help clarify the discussion in this section.

At 720 cents, the spoke marking the left edge of the arc extends inwards to touch the 5th ring. This tuning marks 5-TET, the point at which the 2 small scale steps shrink to zero. Similarly, the spoke at the right edge of the arc (at 600 cents) extends inwards to the 2nd ring. This marks 2-TET, the point at which 5 small scale steps shrink to zero. Thus, this MOS scale has 5 large and 2 small steps (5L, 2s), with the numbers reversed for the inverse scale. This simple procedure of following the spokes inward gives the structure of the MOS sale and its inverse. Conversely, following outward the spokes that delimit the scale and its inverse shows which MOS scales 7-TET is embedded inside. In this case, 7-TET is embedded in MOS scales of size 9 and 12. Similarly, any arc and its associated three spokes correspond to a MOS scale and its three landmarks. Thus the labyrinth can be used to investigate visually the structure, inverse, embedding scales, associated equal tunings, and valid tuning range of any MOS scale.

Figure 2 uses a slightly different visualization. In this version, each scale's valid tuning range is indicated with a thick colored band; this enables the scales to be quickly spotted and easily clicked upon when used in a GUI. The landmark tunings are now indicated by the boundaries of each scale arc (as before), and

by the radial segment inside it, which marks the location at which the inverses meet. The thicker arcs also allow for the tuning range of coherence to be clearly indicated with a darker shading.

The scale labyrinth in Fig. 1 is drawn interactively in a Mathematica program that allows the user to control how the information is presented. The basic size of the labyrinth (how many concentric circles it contains) is controlled by the top slider, while the second slider moves the golden ring to highlight scales of the specified size. The circumference of the circle can be labelled in cents or fractions, and in absolute terms or with respect to the chosen scale size. While Fig. 1 shows scales up to 18 tones, the interactive application can display scales of any size.

A large variety of tunings that provide approximations to (temperaments of) just intonation—such as meantone, srutal, magic, hanson, etc.—are built in and detailed information about them is shown by clicking on the names. When displaying individual temperaments, red lines are superimposed whose angles show their 5-limit TOP (Tenney optimal) tunings (such tunings minimize the maximum error of all possible 5-limit just intonation intervals [9]). The number of the ring they extend to indicates the lowest cardinality scale for which every scale degree is a member of a major or minor triad. For instance, in the pentatonic scale C, D, E, G, A, there is no major or minor triad that contains tone D. Adding another fifth (to make the Guidonian hexachord) gives C, D, E, G, A, B, and now every tone is a member of at least one major or minor triad. This scale, however, is not well-formed. The lowest cardinality well-formed scale, all of whose tones are a member of a major or minor triad is, in this case, the seven-tone diatonic scale. So, in general, it may be musically useful to choose an MOS close to a location where a red line passes through its arc.

The concept of *scale*, as we use it, presupposes periodicity. Under this perspective all rotations of a scale are equivalent. We use the concept of *mode* in order to grasp all the combinatorial refinements, which emerge from the variation of a fundamental domain for this period.

The modes of a well-formed scale have a remarkable property of parsimony, by which one mode can be transformed into another by the replacement of a single tone either by a pseudo-octave (= period) or by the augmented prime (= difference interval between large and small step). These transformations are modal refinements of diatonic and chromatic transposition. Every well-formed scale has an associated universe of modes which is freely generated from two basic and commuting transformations and which is therefore isomorphic to the free commutative group $\mathbb{Z}^2$ of rank 2. One transformation preserves the origin of a mode and yields a new finalis (lowest tone) a second higher. The other transformation preserves the finalis of a mode and yields a new origin a pseudo-fifth (generator) sharp-wards. The details of the interactive navigation through the modes of a well-formed scale are based on a mathematical approach given in [22] [23]. Every mode occupies a fundamental frame in $\mathbb{Z}^2$ whose borders correspond to the augmented prime (horizontal dimension) and the pseudo-octave (vertical dimension). The zig-zag trajectories within each frame represent the

step interval patterns and the folding patterns. Minimal changes of the mode can be described as vertical or horizontal shifts of the fundamental domain. Any modal change can be decomposed into such minimal changes. The height-width representation therefore offers an effective navigation method with the modal universe of a well-formed scale.
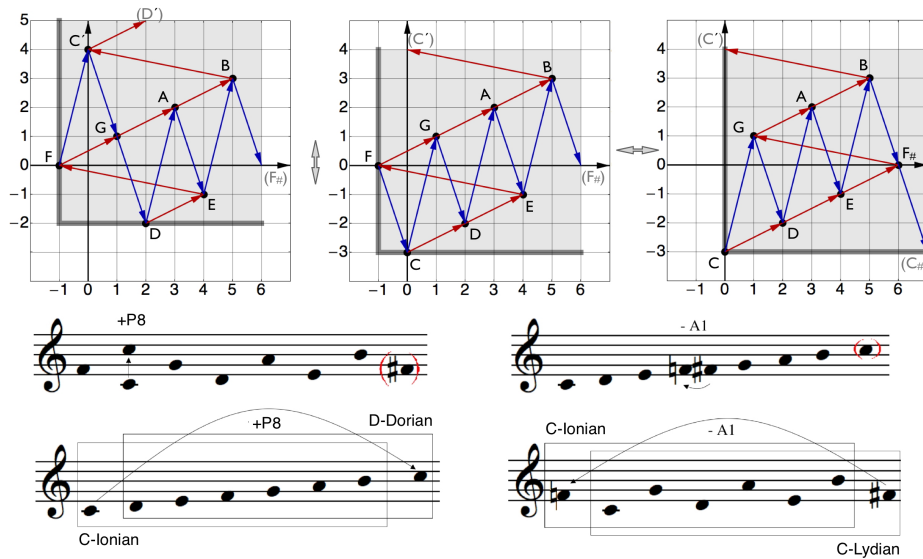

Fig. 7. Two types of mode transformation correspond to by vertical and horizontal shifts of a fundamental frame in a generic height and width coordinate system. The figure shows the transformation of a C-Ionian mode into the common origin mode D-Dorian (left) and the transformation of a C-Ionian mode into the common finalis mode C-Lydian (right), and vice versa.

5   Fourier Scratching as a Performance Technique

The Fourier Transform has been successfully applied in many different ways to the analysis and manipulation of musical sound. It has also been fruitfully applied to non-acoustic problems in mathematical music theory such as those in [16], [17], and [18]. The *Fourier Scratching* project attempts to transfer the established analysis - manipulation - resynthesis paradigm from the domain of sound processing to the domain of macroscopic musical events such as melody, rhythm, tuning and dynamics. It was first proposed in [19] and later realized as a prototype in the domain of rhythm (initially presented at the MCM 2009 [20], later publicly demonstrated at the *Lange Nacht der Wissenschaften in Leipzig*, and discussed at the SuperCollider symposium { SOUNDING CODE } in 2010,

[21]). The present extension to scales is also inspired by Quinn's and Amiot's findings about the Fourier properties of musical scales [17], [18].

Traditionally, playing microtonal music in different scales requires developing multiple instruments with suitable interfaces, which one then needs to learn to play. For example, think of the Cembalo Cromatico with 19 keys per octave. To build and master keyboards like this is quite a challenge, both technologically and pianistically. The present project began with the idea of improvising over the tunings of a scalar labyrinth, something that is infeasible without the help of a computer. Even with suitable software, which makes it easy to play any pattern in any scale at any speed, it is still necessary to specify every pattern in advance. But how can one improvise in any pattern in any scale at any speed?

The idea is to actuate a (virtual) playing robot in a musically sensible way. The basic behavior of the robot is to perform keystrokes with variable strength on variable positions of a keyboard with its n fingers one after the other at a given pulse rate. The finger movements occur with an unperturbed orderliness: finger 1 followed by finger 2 etc. and eventually finger n followed again by finger 1. The pulse and the number n of fingers can be changed but for the moment consider them fixed. The main paradigm of playing with the help of this robot is that the improviser can change the strength and the position of any finger at any time. The (virtual) keyboard is designed as a continuous circle, i.e. any point on this circle may have a (potentially) different sound. The keyboard layout for a finite scale can have specific key widths which are proportional to the sizes of the step intervals above each tone (see Fig. 8). In these examples the number of robot fingers coincides with the number of scale tones, but this is not mandatory.
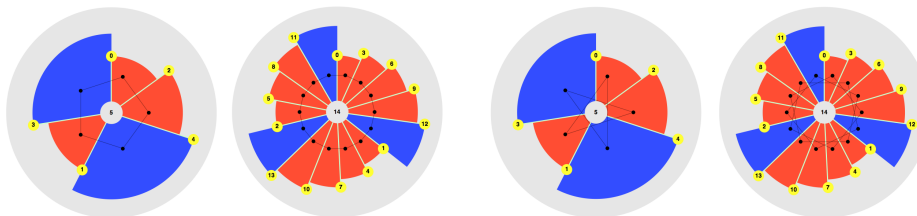

Fig. 8. The black polygons represent elementary play states with a scale being distributed on the circular keyboard (5 tones and 14 tones, respectively). In these cases the playing robot has as many fingers as there are tones in the scale and their distribution is regular. The coherence of these scale examples therefore guarantees that each tone is being played exactly once. In the Fourier picture these play states may be regarded as pure partials. In the left two examples the first Fourier coefficient satisfies $a_1 = 1$ and all others vanish. In the right two examples, $a_3 = 1$ and all other coefficients vanish.

The idea behind the Fourier Scratching technique and interface is the interactive control of musical parameters through the visualization and gestural manipulation of their Fourier coefficients. This presupposes that the musical pa-

rameters in question are encoded as complex vectors $f = (f_0, ..., f_{n-1}) \in \mathbb{C}^n$. For moderately small values of the dimension n there is no need to use the FFT to obtain realtime control, and thus the values of n are not restricted to powers of 2. The slower discrete Fourier transform $\hat{f} = DFT(f) = (a_0, ..., a_{n-1})$ can be computed quickly enough and the inverse Fourier transform can be applied after the gestural manipulation of a Fourier coefficient. The didactical incentive for the Fourier Scratching of Rhythms in the above-mentioned previous implementation is a musically faithful realization of the complex coefficients $f_k = r_k e^{it_k}$. The magnitudes r_k of the coefficients f_k encode loudness values and their phases t_k encode sound colors—given by realtime adjustable FM-sounds—audibly describing a circle. The vector $f \in \mathbb{C}^n$ represents a rhythmical loop $\mathbb{Z}_n \to \mathbb{C}$. Thereby the finite cyclic group $f : \mathbb{Z}_n$ of order n denotes a cyclic pulse of length n, which in performance triggers the values $f_k = r_k e^{it_k}$ one after another; that is, it plays percussive FM-sounds of loudness r_k and of sound color with phase-index t_k. The vector f is subject to continuous change through gestural manipulation of $\hat{f}$ by the performer and can be best described as a "traveling rhythm" in a circle of sound.

In the present application the play state of the robot with n fingers is also encoded by a complex vector $f = (f_0, ..., f_{n-1}) \in \mathbb{C}^n$. With respect to polar coordinates $f_k = r_k e^{it_k}$, r_k can be interpreted as the strength (loudness) and t_k as the play position of the finger with index k. The DFT of the play states in Fig. 8 are thus (from left to right) $(0, 1, 0, 0, 0)$, $(0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)$, $(0, 0, 0, 1, 0)$, and $(0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)$, respectively.

From a combinatorial point of view the regular distributions of Fig. 8 are special cases. In general, the position and strength of each robot finger on the circular playtable can be freely chosen at any moment. The naive control mode would be to change these parameters individually for every finger, but this requires a large set of controls. In contrast, the Fourier Scratching technique offers the ability to change the play states globally and smoothly using only a few parameters. While we cannot offer empirical evidence yet that this particular technique is musically more effective than other alternatives, it is useful to observe that the partials (as the most elementary play states) correspond to musically elementary patterns.[7] As exemplified in Fig. 8 it is precisely the family of coherent well-formed scales which will be played in generic scalar order by the first partial play state. Higher partials with indices k coprime to n generate complete generic interval cycles. Early experiments with this system give the impression that play states which are closely related in their Fourier coefficients are sensibly related by the musical ear. Navigation along the scalar hierarchy in real time can suitably be accompanied by a rising or reduction of the dimension of the play state (e.g by zero-padding the DFT of the current play state or by deleting the Fourier coefficients with minimal energy).

The play states and their Fourier Transforms can be visualized using Riemann number spheres with the vectors f and $\hat{f}$ being displayed as closed polygons with small colored balls at their points as in Fig. 9 (or refer to the demonstration

---

[7] The Nth *partial* is defined as $p_{N,k}(t) = e^{\frac{2\pi itk}{N}}$.

at `http://www.youtube.com/watch?v=-qo09XTtFMA&feature=related`). The small balls are always played in a loop, as indicated by the polygon. The equator represents the continuous circular keyboard where every finger of the robot can hit at any longitude (between 0 and 2π). Gestural control of the locations where the fingers "hit" the spherical surface of the "keyboard", selection of the "keys" (small balls), of the spheres, and activation of sonification, is currently enabled through a standard 4-axis 12-button game controller. Points which differ only in latitude from some point on the equator have the same sound and the same pitch but are played louder (north) or softer (south) than their projection to the equator. Both pitch and timbre are mapped to longitude, but pitch moves in discrete steps (as set by the choice of scale—see the sectors illustrated in Fig. 7), while timbre is continuously variable. This implies that two different pitches never have the same timbre, and that the same pitch may have more than one timbre.
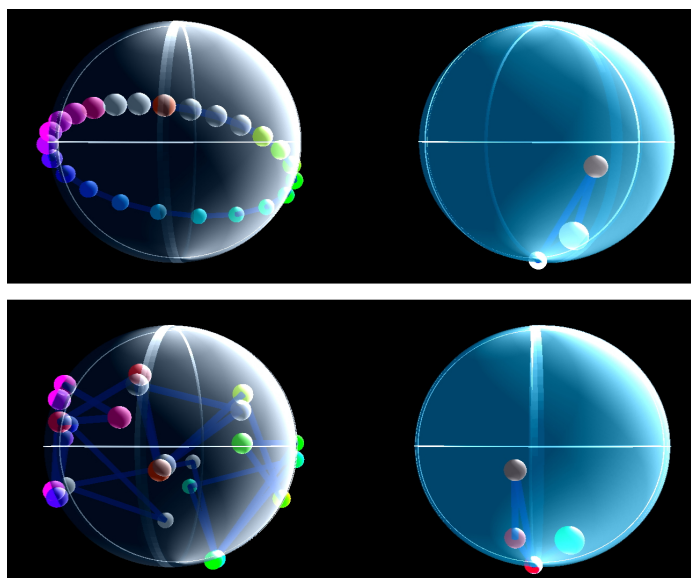

Fig. 9. The screenshots display two 24-dimensional play states with relatively simple Fourier Transforms. In each screenshot (top and bottom) there are two spheres. The left ones (passive) show the play states, while the right ones (active and selected) show the associated Fourier transforms. The screenshot on top represents a play sate with only two non-vanishing Fourier coefficients: a_0 and a_1. Here a_0 is selected (and enlarged). The screenshot at the bottom shows, how this simple play state is being "scratched" by a manipulation of the coefficient a_3, being selected (and enlarged).

The paradigm of Scratching the ascending pattern of a scale can be extended to the paradigm of Scratching a circular musical score (a loop). In addition to the play state, the score can be visualized on the surface of the sphere as a landscape

and the robot finger with index k in a certain play state $(..., r_k e^{t_k}, ...)$ plays the notes (or generalized musical events) which are positioned at the score position specified by its phase t_k and with the loudness specified by its magnitude. In this paradigm the composer of the score defines the elementary patterns.

6   Automatic Adaption of Sound Color

For tones with harmonic spectra (such as most Western instruments and the human voice), intervals close to low integer frequency ratios (e.g., the octave and perfect fifth) are typically considered to be harmonically consonant and have high melodic affinity. However, the intervals found within dynamically tuned well-formed scales can take any size and so may be quite unfamiliar and dissonant in character, or sound out-of-tune.

The inharmonic timbres produced by FM-synthesis are effective at ameliorating these issues, but there is an alternative technique—Dynamic Tonality [10]—that can be used to minimize the sensory dissonance [24] and maximize the melodic affinity and in-tuneness [25] of microtonal intervals. This is achieved by "matching" the tunings of the tones' partials (overtones or harmonics) to the underlying tuning of the scale (matching means that when a typical scale interval is played, many of the partials in one tone have the same pitches as partials in the other tone). This technique can be applied using any form of synthesis in which the tunings of partials can be precisely controlled; appropriate methods of sound synthesis include analysis-resynthesis, additive, and modal. Software synthesizers utilizing these techniques can be downloaded from `http://www.dynamictonality.com`. The pitch (relative to the fundamental) of each partial is mapped to a linear combination of the pitch heights of the period and generator of the underlying scale. This enables some of the intervals between the partials to correspond to some of the intervals found in the MOS scale, even as the underlying tuning is changed dynamically.

7   Discussion

In speaking about and proposing a technique for playing music, there is an æsthetic aspect that inevitably arises. We see affinities between the Fourier Scratching techniques suggested here and minimal music, streamlined loops in DJ Culture, and even with some ideas of early serialism. Our initial motivation was to depart from the traditional dichotomy between instrument and virtuoso musician so as to allow a novel interaction between mathematical composition techniques, real-time computation, and improvised performance. Likewise, our first experiences from the more didactically oriented realization of the Fourier Scratching concept ([20] [21]) transformed into a music-æsthetic and media-theoretic challenge. It appears that the "immediate" interaction of theoretical thought and musical experience calls for a suitable environment of reflection and action, where music, mathematics and computation are "naturally" entangled. Looking towards the future, perhaps the ultimate performance interface

for this playing technique would be a large interactive sphere with multitouch and display.

8   Conclusion

In this paper, we have discussed the scale labyrinth—a visualization of a universe of interconnected well-formed scales that enables any generic well-formed scale, and its specific tuning, to be easily selected. We have shown how the selected scale can be played by the technique of Fourier Scratching, which operates on spectra of playing gestures, thereby patterning the order in which scale degrees are played, as well as their timbres and their loudnesses. We also discuss how mode transformations of the scale can be effected, and a method to ensure the timbre used is optimally matched to the underlying scale so as to maximize tonal affinity. We hope this novel system provides a set of useful constraints and parameters for performing improvisations across the valid tuning range of a given well-formed scale, across "diatonic" and "chromatic" transpositions of such a scale, and across the interconnected universe of different well-formed scales that is pictured in the labyrinth.

References

1. Carey, N.: Distribution Modulo 1 and Musical Scales. Ph.D Dissertation, University of Rochester (1998)
2. Milne, A.J., Sethares, W.A., Plamondon, J.: Tuning Continua and Keyboard Layouts. J. of Math. Music 2, 1–19 (2008)
3. Wilson, E.: Scale-tree (Peirce Sequence), `http://www.anaphoria.com/sctree.PDF` (1994)
4. Holmes, R.: 7-limit Scales, `http://doctroid.wordpress.com/pages/music/xenharmonic-scales/7-limit-scales/` (2006)
5. Carey, N., Clampitt, D.: Aspects of Well-formed Scales. Music Theory Spectrum 11, 187–206 (1989)
6. Clough, J., Engebretsen, N., Kochavi, J.: Scales, Sets, and Interval Cycles: A Taxonomy. Music Theory Spectrum 21, 74–104 (1999)
7. Balzano, G.J.: The Pitch Set as a Level of Description for Studying Musical Perception. In: Clynes, N. Music, Mind, and Brain. Plenum Press, New York (1982)
8. Carey N.: Coherence and Sameness in Well-formed and Pairwise Well-formed Scales. J. of Math. Music 2, 79–98 (2007)
9. Erlich, P.: A Middle Path Between Just Intonation and the Equal Temperaments, Part 1. Xenharmonikôn 18, 159–199 (2006)
10. Sethares, W.A., Milne, A.J., Tiedje, S., Prechtl, A., Plamondon, J.: Spectral Tools for Dynamic Tonality and Audio Morphing. Comput. Music J. 33, 71–84 (2009)
11. Balzano, G.J.: The Group-theoretic Description of 12-fold and Microtonal Pitch Systems. Comput. Music J. 4, 66–84 (1980)
12. Rothenberg, D.: A Model for Pattern Perception with Musical Applications. Math. Systems Theory 11, 199–234 (1978)
13. Blackwood, E.: Structure of Recognizable Diatonic Tunings. Princeton University Press, Princeton (1985)

14. Stern, M.A.: Über eine Zahlentheoretische. Funktion 55, 193–220 (1858)
15. Brocot, A.: Calcul des Rouages par Approximation Nouvelle Methode. Revue Chronometrique 6, 186–194 (1860)
16. Lewin, D.: Re: Intervallic Relations between Two Collections of Notes. J. of Music Theory 3, 298–301 (1959)
17. Quinn, I.: A Unified Theory of Chord Quality in Equal Temperaments. Ph.D. Dissertation, Eastman School of Music (2004)
18. Amiot, E.: David Lewin and Maximally Even Sets. J. of Math. Music 2, 157–172 (2007)
19. Amiot, E., Noll, T., Agon, C., Andreatta, M.: Fourier Oracles for Computer-Aided Improvisation. Proceedings of the ICMC: Computer Music Conference. Tulane University, New Orleans (2006)
20. Carlé, M., Hahn, S., Matern, M., and Noll, T.: Demonstration of Fourier Scratching. Presented at MCM2009, New Haven, US (2009)
21. Carlé M., Noll, T.: Fourier Scratching: { SOUNDING CODE }. SuperCollider Symposium 2010, Berlin (2010)
22. Noll, T.: Ionian Theorem. J. of Math. Music 3, 137–151 (2009)
23. Clampitt, D., Noll T.: Modes, the Height-Width Duality, and Handschin's Tone Character. Music Theory Online, 17(1) (2011)
24. Sethares, W.A.: Tuning, Timbre, Spectrum, Scale. Springer Verlag, London (2005)
25. Milne, A.J., Prechtl, A., Laney, R., Sharp, D.B.: Spectral Pitch Distance & Microtonal Melodies. Presented at ICMPC11, Seattle (2010)