American Review of Mathematics and Statistics
December 2023, Vol. 11, No. 2, pp. 28-41
ISSN: 2374-2348 (Print), 2374-2356 (Online)

Published by American Research Institute for Policy Development
DOI: 10.15640/arms.v11n2a3
URL: https://doi.org/10.15640/arms.v11n2a3

VARIATIONS ON A TUNING ALGORITHM

WILL TURNER[1]

Abstract

In previous work we have presented a tuning algorithm for music written on a stave. The main idea is that tuning the music written on a stave in this way preserves a body of the consonances of the piece. This body of consonances does not consist of all the consonances of the piece, but rather a subset of the set of all consonances determined by a specific algorithm. The consonances form the arrows of a certain quiver whose vertices are the notes of the piece and whose underlying graph is a tree. In this paper we follow up on some of the possibilities of this algorithm, by presenting variations on it. The variations each have a musical purpose. For example, when we select certain tunings, applying our algorithm gives music which in part lies outside the human audible range; a variation corrects for this. To give another example, for certain tunings, the resulting music has big leaps between notes; a variation moderates this.

Key words: consonance, tuning, quiver, variations

1. Introduction

The tuning of music represented on a stave has been the subject of some study (see eg. [2, 5]). Usually the intervals of an octave, fifth, and major third are tuned to frequency ratios of $2, 3, 5$ respectively, or to some approximation of those. In further work [6] we have developed tunings in which these intervals are tuned to frequency ratios $\zeta_2, \zeta_3, \zeta_5 \in \mathbb{R}$ that might differ significantly from $2, 3, 5$. The new tunings preserve a body of the consonances of the piece. The consonances form a quiver whose vertices are the notes of the piece, and picks out consonances between notes that are close in the score, as specified by a certain algorithm. Here we present variations on this algorithm, which solve various issues and follow up on ideas suggested by it.

In the second section of the paper we describe the original algorithm, which we regard as a 'theme' which is subject to 'variations' in the rest of the paper.

Of particular interest to us are frequency ratios $\zeta_2, \zeta_3, \zeta_5$ which are small integers. This is because the harmonics of a vibrating string or similar are given by integral frequency ratios [2]. One issue that arises is that, when choosing integral elements for $\zeta_2, \zeta_3, \zeta_5$ we frequently obtain a piece in which the pitch drift is so large, the piece becomes inaudible in places. This is addressed in the third section of the paper. Another issue that can arise is successive leaps in the piece become quite large. This is addressed, by a variety of methods, in section 4 of the paper.

[1] Institute of Mathematics, University of Aberdeen, Aberdeen, AB24 3UE, Scotland, UK. w.turner@abdn.ac.uk.

A tuning behaves somewhat like a key. One can move between tunings, analogously to moving through keys. One way to do this is described in the fifth section of the paper.
The quiver associated to the Bach Invention we tune is connected in the original algorithm. However, with a little tweaking, it has two connected components, an even component and an odd component. We describe the effect of this in the sixth section.
In section 7 we connect successive notes of the quiver with a single harmonic. In section 8 we connect consecutive notes in each of the two parts with glissandi.
In section 9, our final variation, we discuss the problem of going backwards, and constructing a piece of music from a quiver. Our technique is to think of a piece of music as defined by a sequence of constraints on a large set of potential compositions. We label these constraints as either 'rational' or 'arbitrary'. Our strategy is to create a piece of music with few arbitrary constraints.
There are sound files to accompany this paper [7].
It is perhaps worth mentioning some differences between our mathematical approach to a musical score and that of A.I. (see eg. [4]). We 'train' our quiver on only one score, rather than a large body of work. And our new tunings often sound quite different from the music we are representing.

2. The theme: a tuning algorithm

Here we describe a tuning algorithm, the 'theme', which we will present variations on in the rest of the paper. This is the tuning

Let $\mathcal{M}$ denote the subgroup of $\mathbb{Q}^\times$ generated by $2, 3, 5, 7$. Consider the group homomorphism $m : \mathcal{M} \to \mathbb{Z}$ sending $2, 3, 5, 7$ to $12, 19, 28, 34$ respectively. Suppose we have a fixed natural number n with $5 \leq n \leq 10$. We define g to be the restriction of m to $\{1, 2, ..., n\} \cdot \{1, 2, ..., n\}^{-1}$. Let $\mathcal{G}$ denote the image of g. Suppose we have a fixed section s of g.
Suppose we are given a two part composition on the stave, such as a Bach Two-Part Invention. We denote one of the parts 1 and the other part 2. We have a linear order of the notes of our composition, where notes are ordered by start time, and given two notes starting at the same time we precede the note in part 2 by the note in part 1. We denote by N the number of notes of our composition.
For $1 \leq x \leq y \leq N$ define the sequence $S(x, y)$ of elements of $\{1, 2, ..., N\}$ to be

$$(x + 1, x + 2, x + 3, ..., y, x - 1, x - 2, x - 3, ..., 1).$$

We denote by $i(x, y) \in \mathbb{Z}$ the number of semitones required to ascend from the x^{th} note to the y^{th} note.

We define a quiver Q whose vertices are given by the notes of our composition, and whose arrows are labelled with elements of $\mathcal{G}$. This is the consonance structure which, when represented, defines an interpretation of our two part composition.
Our algorithm to define Q begins with a quiver with a single vertex, corresponding to the first note of the composition, and no arrows; it adds vertices and arrows successively. We run through elements y of $\{1, 2, ..., N\}$ consecutively, in standard order. For a fixed y we run through the elements x with $1 \leq x \leq y$ in reverse order. For a fixed x and y we search through $S(x, y)$ for vertices in our quiver to connect to x. If x already belongs to our quiver, we abandon our search through $S(x, y)$ straightaway. Otherwise we run through the elements z of $S(x, y)$ in sequence. If

$i(x, z) \in \mathcal{G}$ and z belongs to our quiver, we add x to our quiver, draw an arrow from x to z, labelled with $i(x, z)$, and discontinue the search through $S(x, y)$.
The underlying graph of Q is a tree, since our algorithm involves adding leaves successively. We will assume that the vertex set of Q is the set of all notes of our composition, although there do exist examples where this is not the case.
Choose positive real numbers ζ_2, ζ_3, ζ_5, ζ_7. These determine a homomorphism ζ from $\mathcal{M}$ to $\mathbb{R}$ sending p to ζ_p, for $p = 2, 3, 5, 7$. Consider the double of Q, which is the quiver obtained from Q by adjoining a single reverse arrow from v' to v for every arrow from v to v' in Q. We label the arrows of the double of Q as follows: given an arrow in our quiver labelled by i, we label the corresponding arrow in our double quiver with $\zeta(s(i))$ and the corresponding reverse arrow with $\zeta(s(i))^{-1}$.
A path in the underlying graph of Q determines a path in the double of Q, and thus a real number, via the above representation: this real number is the product of the real numbers labelling the arrows in the path. Choose an initial frequency $F_0 \in \mathbb{R}$. Every vertex v of our quiver is connected by a unique path in the underlying graph of Q from the first note of the composition, and thus multiplying the real number given by this path by F_0 determines a frequency, which gives the frequency F_v of v.

To a frequency $f \in \mathbb{R}$ we assign the function from $\mathbb{R}$ to $\mathbb{R}$ sending t to the sum $\sum_{i=1}^{n} \sin(2\pi f\zeta(i)t)/\sqrt{i}$. We call $f\zeta(i)$ the i^{th} harmonic of this function. To a vertex v of our quiver we have associated a frequency, and to a frequency we have assigned a function. We call the resulting function 'the function assigned to v'.
By construction, an arrow in Q directed from v_1 to v_2 corresponds to at least one common harmonic of the functions assigned to v_1 and v_2. Indeed, if $s(i(v_1, v_2)) = \alpha/\beta$, for $1 \leq \alpha, \beta \leq n$, then the frequencies of v_1 and v_2 differ by the factor $\zeta(\alpha/\beta)$, and the fact that ζ is a group homomorphism implies the α^{th} harmonic of the function assigned to v_1 is equal to the β^{th} harmonic of the function assigned to v_2.
We obtain a piece by playing, for every vertex v, the function assigned to v, for the duration of the note associated to v in our score. This piece is a representation of the consonance structure Q, and an interpretation of the original two part composition.

In all the examples of this paper, we set $n = 5$, in which case ζ_7 is redundant. In the numbered examples we use the two part composition Invention No. 9 by Bach.

3. Variation: Restraining motion

Of particular interest are the cases when all of the ζ_ps are integers.
Our choices for ζ_p in the preceding section are rather limited if we want to select integers, and we also want the piece to be entirely audible. Substantial pitch drift can occur that means it is not possible to place the piece in the audible range. To get around this, we restrain the motion in the piece to lie within audible constraints.
Let ξ be one of the ζ_ps.
In our algorithm, a frequency F_v is obtained via a sequence of frequencies

$$F_0, F_{v_1}, F_{v_2}, ..., F_{v_n} = F_v,$$

where $0, v_1, v_2, ..., v_n$ are the vertices in the unique shortest path from 0 to v. We obtain F_{v_i} from $F_{v_{i-1}}$ by multiplying by a certain $\zeta(j)^{\pm 1}$. In our restraining motion algorithm, we instead obtain F_{v_i} from $F_{v_{i-1}}$ by multiplying by our $\zeta(j)^{\pm 1}$ and dividing by ξ, if the result is > 2000; by multiplying by our $\zeta(j)^{\pm 1}$ and multiplying by ξ, if the result is < 50.

Example 1 $(\zeta_2, \zeta_3, \zeta_5) = (2, 3, 11)$, $\xi = 3$.

4. Variations: reordering

The variation example of the preceding section had big leaps between successive notes: the mean leap was 0.88 octaves. In this section we discuss various ways to reduce the mean leap size by reordering.
The first way is to introduce an additional pair of elements ± 1 to the set $\mathcal{G}$, and a frequency ratio a little greater than 1 in $\langle \zeta_2, \zeta_3, \zeta_5 \rangle$ which we define to be $s(1)$. We define $s(-1) = s(1)^{-1}$. The effect is to smooth out some of the leaps of the piece, as in the following example.

Example 2 $(\zeta_2, \zeta_3, \zeta_5) = (3, 5, 11)$ and $s(1) = 27/25$. The mean leap is 0.36 octaves. It would be 0.5 octaves if we didn't introduce the smoothing out.

The second way is to define $\tilde{s}$ to be the unique increasing function from $\mathcal{G}$ to $\mathrm{Im}(\zeta s)$, and work with $\tilde{s}$ instead of ζs when labelling the double of Q.

Example 3 $(\zeta_2, \zeta_3, \zeta_5) = (2, 3, 9)$, $\xi = 3$. The mean leap is 0.65 octaves. It would be 0.75 octaves if we didn't introduce the reordering.

The third way is to reorder in bars to reflect the original composition. Indeed, take a two part composition, interpreted as in section 2. For each bar, take the fundamental frequencies F_v of the notes $n_1, n_1 + 1, n_1 + 2, ..., n_1 + n_2 - 1$ of that bar in order, and take a permutation σ that permutes these notes so that these fundamental frequencies lie in increasing order. Take the fundamental frequencies of the notes $n_1, n_1 + 1, n_1 + 2, ..., n_1 + n_2 - 1$ of the original composition, and take a permutation τ that permutes these notes so these fundamental frequencies lie in increasing order. Apply the permutation $\tau^{-1}\sigma$ to the notes of the bar of the retuned interpretation, and reorder the frequencies of the piece correspondingly. In this way we obtain a reordered composition which, within bars at least, has fundamental frequencies that occur in the same order as those in the original composition.
We have not stated how we should permute notes in a bar so that fundamental frequencies lie in increasing order. To do this, we begin with our list of fundamental frequencies: $f_0, f_1, f_2, ..., f_{n_2-1}$. We let i run from 1 to $n_2 - 1$, and if $f_{i-1} > f_i$ we swap f_{i-1} and f_i. We then repeat this run $n_2 - 1$ times, at which point our frequencies will be in ascending order. This algorithm determines a permutation as required.

Example 4 $(\zeta_2, \zeta_3, \zeta_5) = (3, 4, 7)$, $\xi = 7$. The mean leap is 0.57 octaves. This compares with a mean of 1.10 octaves without the reordering.

We can also reorder notes in a bar so that fundamental frequencies lie in increasing order, as in the following example:

Example 5 $(\zeta_2, \zeta_3, \zeta_5) = (2, 3, 11)$, $\xi = 3$.

5. Variation: moving through 'keys'

We can think of a choice of ζ_ps as analogous to a key. It makes sense to smoothly change 'key' as the piece progresses.

Indeed, consider our original algorithm, together with the following variation:
Suppose we have an element $\theta_v = (\theta_{v,1}, \theta_{v,2}, \theta_{v,3}, ..., \theta_{v,n})$ of $\mathbb{R}^n_{>0}$, where $\theta_{v,1} = 1$, for every vertex v of Q.
We label the arrows of the double quiver of Q as follows. Given an arrow from v_1 to v_2 in Q, labelled with γ, we label the corresponding arrow in the double quiver with the real number $\theta_{v_2,l(s(\gamma))_1}{\theta_{v_1,l(s(\gamma))_2}}^{-1}$. We label the corresponding reverse arrow in the double quiver with the inverse of this real number.
A path in the underlying graph of Q determines a path in the double of Q, and thus a real number, via the above representation: this real number is the product of the real numbers labelling the arrows in the path. Choose an initial frequency $F_0 \in \mathbb{R}$. Every vertex v of our quiver is connected by a unique path in the underlying graph of Q from the first note of the composition, and thus multiplying the real number given by this path by F_0 determines a frequency, which gives the frequency F_v of v.

To the vertex v in Q we assign the function from $\mathbb{R}$ to $\mathbb{R}$ sending t to the sum $\sum_{i=1}^{n} \frac{1}{\sqrt{i}} \sin(2\pi F_v \theta_{v,i} t)$. We call $F_v \theta_{v,i}$ the i^{th} harmonic of this function.
By construction, an arrow in Q directed from v_1 to v_2 corresponds to at least one common harmonic of the functions assigned to v_1 and v_2. Indeed, if $l(s(i(v_1, v_2))) = (\alpha_1, \alpha_2)$, for $1 \leq \alpha_1, \alpha_2 \leq n$, then the frequencies of v_1 and v_2 differ by the factor ${\theta_{v_2,\alpha_1} \theta_{v_1,\alpha_2}}^{-1}$, which implies the α_2^{th} harmonic of the function assigned to v_1 is equal to the α_1^{th} harmonic of the function assigned to v_2.
We obtain our variation by playing, for every vertex v, the function assigned to v, for the duration of the note associated to v in our score.
It remains for us to select our $\theta_{v,j}$s. We fix three choices of ζ_p, which we call $(\zeta_2^1, \zeta_3^1, \zeta_5^1)$, $(\zeta_2^2, \zeta_3^2, \zeta_5^2)$, $(\zeta_2^3, \zeta_3^3, \zeta_5^3)$. We also fix integers $0 = a_1 < b_1 < a_2 < b_2 < a_3 < b_3 = m$, where m is the number of notes of our piece. We let $\theta_{v,j} = \zeta_j^i$ if $v \in [a_i, b_i]$, for $j = 2, 3, 5$ and $i = 1, 2, 3$. We allow $\theta_{v,j}$ to interpolate exponentially between ζ_j^i and ζ_j^{i+1} for $v \in [b_i, a_{i+1}]$, for $j = 2, 3, 5$ and $i = 1, 2$. We define $\theta_{v,1} = 1$ for all v, and $\theta_{v,4} = \theta_{v,2}^2$ for all v.

Example 6 We take $(\zeta_2^1, \zeta_3^1, \zeta_5^1) = (2, 3, 5)$, $(\zeta_2^2, \zeta_3^2, \zeta_5^2) = (2, 3, 6)$, $(\zeta_2^3, \zeta_3^3, \zeta_5^3) = (2, 4, 6)$. We take $b_1 = 60$, $a_2 = 250$, $b_2 = 310$, $a_3 = 500$, while $m = 560$.

6. Variation: two parts drifting independently

For $n = 5$, we have

$$\mathcal{G} = \{-28, -24, -19, -16, -12, -9, -7, -5, -4, 0, 4, 5, 7, 9, 12, 16, 19, 24, 28\}$$

which, in our original algorithm, we map via ζs to a set of real numbers. Let us rather use the composition of ζs and the map from

$$\mathcal{H} = \{-28, -24, -18, -16, -12, -10, -6, -4, 0, 4, 6, 10, 12, 16, 18, 24, 28\}$$

to $\mathcal{G}$ sending ± 18 to ± 19, sending ± 10 to ± 9, sending ± 6 to ± 7, and fixing α if α is even.
If we use this in an analogous way to ζs in our algorithm, we obtain a variation, in which the quiver has two distinct subquivers with no arrows between them, an even subquiver and an odd subquiver. This is because all the elements of $\mathcal{H}$ are even.

As the piece progresses, the pitch drift for the two subquivers can be different, so that even if we align the pitches of the first notes in the two subquivers at the beginning, and set $(\zeta_2, \zeta_3, \zeta_5) = (2, 3, 5)$, by the end of the piece we can be quite far from 12 tone equal temperament.

Example 7 We take $(\zeta_2, \zeta_3, \zeta_5) = (2, 3, 5)$. We set the first and fifth note of the piece to differ in frequency by a factor of 3. Consequently the 540th and the 539th note of the piece differ by 2.53 semitones.

7. Variation: Consonance structure harmonics

Example 8 As we observed in our section on the original tuning algorithm, an arrow in Q corresponds to a common harmonic of two notes n_1 and n_2. Here we play such harmonics, for the duration starting at the beginning of n_1, and ending at the end of n_2. We take $(\zeta_2, \zeta_3, \zeta_5) = (2, 3, 5)$.

8. Variation: Glissandi

Example 9 Here we connect consecutive notes in each part with glissandi. We take $(\zeta_2, \zeta_3, \zeta_5) = (3, 5, 10)$.

9. Variation: Rational constraints modelled on the tuning algorithm

To write a piece of music, a composer must choose a single element from a large set of potential compositions. To do this, a number of choices must be made, which can be thought of as subsets of the large set to which the single element must belong, or as constraints on the large set. The single element which constitutes the composition lies at the intersection of the constraining subsets. Constraints also operate in the mind of a listener to a piece of music. For example, towards the end of a piece of tonal music, a listener might be able to hum along, demonstrating the presence of constraints on the way the music is going in their mind.
In this section, we discuss some examples of compositions constructed by making a sequence of constraints on a large set of potential compositions. We assume that constraints are either *rational* constraints, or *arbitrary* constraints. Here, rational constraints are constraints which have some logic behind them, often a similarity with existing music, or an internal similarity, whilst arbitrary constraints are choices made without any rational support, beyond the fact that they are necessary to obtain a single composition after all constraints are applied.
To illustrate how this works, let us consider a simple example.

We consider a piece of music written as a sequence of frequencies and durations (we consider this restriction to be rational, because a great deal of existing music is constructed this way).
We consider a piece of music written in a single part in the frequency range 260-880 Hz(we consider this restriction to be rational, because it mimics a human soprano vocal line).

We consider music written in equal temperament (we consider this to be rational because it contains maximal translational similarities between subsets of notes).
We consider equal temperament that contains 2 as a frequency ratio (we consider this to be rational because the octave gives a very significant similarity between notes, even to the point that such notes are labelled by the same letter in the musical literature).
We consider equal temperament that approximates 3 as a frequency ratio (we consider this to be rational because 3 is the next integer ratio after 2, and the perfect fifth with frequency ratio 3/2 gives another significant similarity between notes).
We consider 12 tone equal temperament (we consider this to be rational because it satisfies the three preceding constraints).
We consider an ascending scale, or a descending scale (we consider this to be rational, because it minimises the logarithmic sum of successive leaps; here if $(f_i)_{i=1}^n$ is a sequence of frequencies the logarithmic sum of successive leaps is $\sum_{i=2}^n |\log(f_i/f_{i-1})|$).
We allow no repetitions in our scale (we consider this to be a rational way to avoid monotony).
We consider a scale of maximal length subject to the other constraints (we consider this to be rational because our piece is too short as it is).
We choose an ascending chromatic scale rather than a descending chromatic scale (we consider this to be an arbitrary choice).
We take scale frequencies $\{2^{\frac{n}{12}} \cdot 880 | -21 \leq n \leq 0\}$ (we take the initial frequency of our scale to be an arbitrary number just > 260 Hz).
We choose all notes to have the same duration (we consider this to be rational, because it maximises similarities of duration between notes).
We take notes of duration $1s$ (we consider this choice to be arbitrary).

In conclusion, our piece is the chromatic scale with initial frequency $2^{-\frac{21}{12}} \cdot 880$ Hz and final frequency 880 Hz, and notes of duration $1s$.
The strategy was to use few arbitrary constraints.
The range of choices for our initial note frequency was $[260, 2^{-\frac{21}{12}} \cdot 880]$. If we identify such a frequency with the closest frequency to it in $\{880 \cdot 2^{\frac{n}{1200}} | n \in \mathbb{Z}\}$, then we have 12 choices (we are assuming the ear would identify notes that are this close).
A range of durations we could have chosen from is $[2^{-6}, 1]$. If we identify such a duration with the closest duration to it in $\{2^{\frac{n}{120}} | n \in \mathbb{Z}\}$, then we have 721 choices. (we are assuming the ear would identify durations that are this close).
For our human vocal range, we chose soprano. We could have just as well chosen bass, baritone, tenor, contralto or mezzo. That makes six possible choices.
We chose an ascending chromatic scale, rather than a descending one. There were two possible choices here.

Altogether, that makes $12 \cdot 721 \cdot 6 \cdot 2$ arbitrary choices (there is of course some flexibility in how we have made this calculation - for example, we might have selected a different number of possible durations or frequencies, and it is debatable how 'rational' some of our rational choices are - but the main point is that we have made far fewer arbitrary choices than if we were to generate a piece of a similar length, whilst making no rational constraints at all).
It is our intention to use the above strategy, of using rational constraints where possible, to generate a piece of music that is slightly more complicated, and departs from 12 tone equal temperament.

Example 10 Let Q be a quiver whose underlying graph is a tree, and whose edges are labelled with elements of $\{1,2,3,4,5\} \cdot \{1,2,3,4,5\}^{-1}$. We assume each vertex of our quiver has a distinguished vertex v_0, and an initial frequency F_0. We assume we have positive integers $\zeta_2, \zeta_3, \zeta_5$. We associate a note to each vertex of our quiver, just as in the tuning algorithm which is the subject of this paper.
We regard the use of a quiver in this way to be a rational way to construct a set of notes, as it formalises the use of a consonance structure. As in the pieces obtained from our eponymous tuning algorithm, we assume that notes which are close in the quiver are close in time.
We select $(\zeta_2, \zeta_3, \zeta_5) = (2, 5, 11)$. It makes sense to choose small positive integers (< 20, say) to avoid a piece containing very large leaps. So this is an arbitrary choice out of 20^3 possible choices.
Our quiver has two connected components, one for each of two movements of the piece. As is standard musical practice, we assume the two movements of the piece to be in contrast. The number two is an arbitrary choice.

We assume that each connected component is obtained by gluing together a number of similar subquivers (it is standard to generate music by similarity). To mark out a contrast between the two, for connected component 1, we assume the subquivers are all stars with g_1 vertices, whilst for connected component 2, we assume the subquivers are all lines of length g_2 (stars of diameter 2 and lines of diameter g_2 have the biggest difference in diameter for connected trees with $g_2+1 \geq 3$ vertices). For connected component 2, we glue together lines into a long line. Again for contrast we obtain connected component 1 by connecting our g_1 vertex star $S = S^0$ to $g_1 - 1$ copies of S by gluing its leaves to their centres, to form a graph S', then connecting S' to $(g_1 - 1)^2$ copies of S by gluing its leaves to their centres, to form a graph S'', etc. up to S^h which is connected component 1.

Above, g_1 and g_2 are arbitrary choices which we assume to be at most 20, estimating a typical motif to have between 1 and 20 notes.

For connected component 1, there is one obvious way to arrange our notes so that adjacent vertices in the quiver sound simultaneously. It goes as follows: Label the edges of S, and order the leaves of S linearly. Extend the labelling of S, compatible with all the gluings, across S^h. We play the note corresponding to the central vertex of S^0 for the full duration f of the piece. Play the notes corresponding to the leaves of S^0 with duration $f/(g_1 - 1)$, with order given by our linear order of the leaves of S. Play the notes corresponding to the leaves of S^1 with duration $f/(g_1 - 1)^2$, so that they sound alongside the leaves of S_0 that they are attached to, with order given by our linear order of the leaves of S, etc.

We specify $h = 4$ so that the piece sounds in five parts, which is a musical choice that can be found in Bach's well-tempered clavier, for example. We take as our labellings all possible ratios of ζ_is between 1 and 3 (ratios more than this introduce very big intervals, which we would like to avoid for range reasons). These are $5/4, 2/1, 11/5, 11/4$. We need an ordering of these. The most harmonious interval appearing here is $2/1$, so we put that last in our ordering. We then choose an ordering with the smallest logarithmic sum of successive leaps: $5/4, 11/5, 11/4, 2/1$. The same sum is achieved with the ordering $5/4, 11/4, 11/5, 2/1$. Selecting between these is an arbitrary choice. The duration f is also an arbitrary choice. The frequency ratio of the highest to the lowest harmonic is $(11/4)^4 \cdot 11$. We pick an

initial frequency $F_0 = 30$ Hz so this remains in the audible range. In principle F_0 could go as low as 20 or as high as 31.8.

We next turn to connected component 2, whose quiver is a line. The labels of the edges of this line form a sequence of frequencies in $\{1, 2, 5, 4, 11\} \cdot \{1, 2, 5, 4, 11\}^{-1}$ To generate our music by similarity, we consider a pair of "motifs", which are words in $2^{\pm 1}, 5^{\pm 1}, 4^{\pm 1}, 11^{\pm 1}$, and obtain our piece by composing factorisations of these motifs. To avoid excessive pitch drift, we assume these words multiply together to give a number approximately equal to 1. The number of motifs is somewhat arbitrary, although it makes sense to use 2 since it is bigger than 1, for complexity, and small enough that the ear stands a chance to recognise the motifs. As our motifs we select $m_1 = 5^6 \cdot 11^{-2} \cdot 2^{-7} \approx 1.0088$. and $m_2 = 11^5 \cdot 2^{-8} \cdot 5^{-4} \approx 1.0066$.

We need to select some factorisations of our motifs. For m_2, we consider factorisations whose terms are 2^{-1} (6 times), $\frac{11}{4}$ (once) and $\frac{11}{5}$ (4 times). This selection is made so that successive terms are as small as possible. To keep the motifs short, we have also chosen not to extend the factorisation in length by including mutually inverse terms. For m_1, by similar logic, we consider factorisations whose terms are $\frac{5}{11}$ (twice), $\frac{5}{4}$ (four times), 2 (once).

There are a number of choices of how to compose these factorisations, as a sequence of labelled edges of our quiver (we use the obvious total order on the edges of our linear quiver). We do so as follows:

At the beginning put a factorisation of m_1, then a single term from a factorisation of m_2, then another factorisation of m_1, then a single term from our factorisation of m_2,..., until all the terms from our factorisation of m_2 are exhausted, at which point we start again with another factorisation of m_2, etc.

Our logic is to keep the copies of the terms in m_i at regular intervals to help them be recognisable by the ear. We will play the single notes connected by terms in our factorisation of m_2 simultaneously.

We want to limit the range of frequencies of our music, at least so it is audible. If we were to place no limits on the factorisations above we would have a potential range ratio of roughly $2^6 \cdot (\frac{11}{5})^2$, which together with the range ratio of harmonics of 11 gives a range ratio of $2^6 \cdot 11^3 \cdot 5^{-2} \approx 3400$, greater than the range ratio of the human ear (here the range ratio of a range r_1 Hz-r_2 Hz is r_2/r_1).

A practical way to limit factorisations is to alternate terms > 1 and terms < 1. For m_2, when we do this, odd terms are 2^{-1} whilst four even terms are $\frac{11}{5}$ and one even term is $\frac{11}{4}$. Altogether there are 5 choices of factorisation. The range of m_1 is less concerning: it is maximally only a bit more than 2 octaves. There are 105 possible factorisations for m_1, without any restrictions.

We order our factorisations lexicographically. This is a rational choice, since lexicographic ordering is based on similarity, and it is a standard principle to place similar musical elements side by side. We have $5 \cdot 11 = 55$ terms of m_2 (resp. 105 factorisations of m_1) if we take the entire sequences.

As they are currently ordered, the notes of our realisations of our motif m_1 potentially jump up and down quite a lot. When we reorder them in time, we leave fixed the first and last frequencies of a realisation of m_1, and rearrange the intermediate frequencies in ascending order. (it is rational to have either an ascending sequence or a descending sequence, because it minimises the logarithmic sum of successive leaps, but choosing between ascending and descending is arbitrary).

We run through 55 repetitions (together with the first 55 of 105, ordered lexicographically). The choice of an interval of 55 out of 105 involves an arbitrary choice from 51.
The duration of all our notes is identical, for similarity. The value of this duration is arbitrary. Our initial frequency $F_0 = 220$ is chosen so the top harmonic of the highest note is close to the upper end of the audible range. Again, of course, this involves some arbitrary choice.

Let us discuss the problem of finding approximate relations. Suppose we are given two real numbers $r_1, r_2 \approx 1$ that are not equal to 1. Choosing $a_1, a_2 \in \mathbb{Z}$ judiciously, we can find a closer approximation $r_1^{a_1} r_2^{a_2} \approx 1$.
For example, suppose $r_1 = 5^3 \cdot 11^{-2}$, $r_2 = 2^7 \cdot 5^{-3}$. Then $r_1 r_2^{-1} = 5^6 \cdot 2^{-7} \cdot 11^{-2}$ is closer to 1. Suppose $r_1 = 11/16$, $r_2 = 11/10$. Then $r_1 r_2^4 = 11^5 \cdot 2^{-8} \cdot 5^{-4}$ is closer to 1.
Suppose $r_1 = 13 \cdot 3^{-1} \cdot 2^{-2}$, $r_2 = 3^3 \cdot 2^{-1} \cdot 13^{-1}$ Then $r_1 r_2^{-2} = 3^{-7} \cdot 13^2$ is closer to 1.
Let us be more systematic. Suppose we begin with two linearly independent elements $s_1, s_2 \in A$, where A is a subgroup of $\mathbb{Q}^\times$.
Reordering, and replacing with inverses if necessary we may assume $1 < s_2 < s_1$. For $n \geq 3$ we recursively define a relation $1 < s_n < s_{n-1}$ to be $s_n = s_{n-2} \mod \langle s_{n-1} \rangle$. By induction, the pair (s_{n-1}, s_n) is linearly independent for $n \geq 2$. The descending sequence (s_n) converges on 1, and lies in A.
(Why does this sequence converge on 1? Consider the logarithm of the sequence (s_i). This is a sequence (α_i). If $\alpha_{n-1} > \frac{1}{2}\alpha_{n-2}$ then $\alpha_n = \alpha_{n-2} - \alpha_{n-1} < \frac{1}{2}\alpha_{n-2}$. If $\alpha_{n-1} < \frac{1}{2}\alpha_{n-2}$ then $\alpha_n < \alpha_{n-1} < \frac{1}{2}\alpha_{n-2}$. Thus $\alpha_{2n+1} < \frac{1}{2^n}\alpha_1$ which implies the sequence (α_n) converges on zero.)

References

[1] J. S. Bach, *Two Part Invention No.9*, BWV 780.

[2] D. J. Benson, *Music: A Mathematical Offering*, Cambridge University Press, November 2006.

[3] A. S. Crans, T. M. Fiore, and R. Satyendra, *Musical Actions of Dihedral Groups*, The American Mathematical Monthly, vol. 116, no. 6, June 2009, pp. 479-495.

[4] G. Hadjeres, F. Pachet, and F. Nielsen, *DeepBach: a Steerable Model for Bach Chorales Generation*, Proceedings of the 34th International Conference on Machine Learning, PMLR 70, 2017, pp. 1362-1371.

[5] K. Stange, C. Wick and H. Hinrichsen, *Playing Music in Just Intonation: A Dynamically Adaptive Tuning Scheme*, Computer Music Journal, Volume 42, Issue 3, 2018, pp. 47-62.

[6] W. Turner, *On representing consonance structures*, http://homepages.abdn.ac.uk/w.turner/pages/.

[7] W. Turner, *Examples 1, 2, 3, 4, 5, 6, 7, 8, 9, 10a, 10b*, http://homepages.abdn.ac.uk/w.turner/pages/.