

Specifying spectra for musical scales

William A. Sethares^{a)}

Department of Electrical and Computer Engineering, University of Wisconsin, Madison,
Wisconsin 53706-1691

(Received 5 July 1996; accepted for publication 3 June 1997)

The sensory consonance and dissonance of musical intervals is dependent on the spectrum of the tones. The dissonance curve gives a measure of this perception over a range of intervals, and a musical scale is said to be *related* to a sound with a given spectrum if minima of the dissonance curve occur at the scale steps. While it is straightforward to calculate the dissonance curve for a given sound, it is not obvious how to find related spectra for a given scale. This paper introduces a "symbolic method" for constructing related spectra that is applicable to scales built from a small number of successive intervals. The method is applied to specify related spectra for several different tetrachordal scales, including the well-known Pythagorean scale. Mathematical properties of the symbolic system are investigated, and the strengths and weaknesses of the approach are discussed.

© 1997 Acoustical Society of America. [S0001-4966(97)05509-4]

PACS numbers: 43.75.Bc, 43.75.Cd [WJS]

INTRODUCTION

The motion from consonance to dissonance and back again is a standard feature of most Western music, and several attempts have been made to explain, define, and quantify the terms "consonance" and "dissonance." For example, Tenney¹ provides a historical overview that identifies five separate uses of the terms, and a longstanding debate surrounds the "reductionist" explanations of Plomp and Levelt² and Terhardt,³ and the proponents of "cultural conditioning" such as Cazden.⁴ One of the most successful of the reductionist approaches is called tonal or *sensory* dissonance, in which the sensory dissonance between pairs of sine waves is determined from psychoacoustic experiments. The sensory dissonance of more complex sounds is then defined to be the sum of the dissonances between all simultaneously sounding sine wave partials. Because sensory dissonance depends on the partials, sounds with different spectra may function differently. For instance, an interval may be quite consonant when played with one sound, but quite dissonant when performed with another.

The dissonance curve $D_F(r)$ is a function that describes how the sensory dissonance of a sound with spectrum F varies when played at different intervals r . Figure 1, for instance, shows a plot of the dissonance curve for a sound with six harmonic partials over a range of intervals slightly larger than an octave. The minima of this curve occur at the simple integer ratios of the Just Intonation scale, reinforcing the familiar notion that the most consonant (least dissonant) intervals for sounds with harmonic spectra are those with small integer ratios. The top axis shows the steps of the 12-tone equal tempered scale, which can be viewed as approximating many of these just ratios. Techniques for drawing dissonance curves are described in detail in Ref. 5 and a computer program is given in Ref. 6. These are based on an explicit parametrization of the perceptual data gathered by Plomp and Levelt² (and replicated in Ref. 7). These are in

turn closely related to Helmholtz's⁸ "beat theory" of dissonance in which the beating between higher harmonics causes a roughness that is perceived as dissonance.

Nonharmonic sounds can have dissonance curves that differ dramatically from Fig. 1, indicating that intervals with the most sensory consonance depend strongly on the structure of the partials of the sound. A dramatic example of this is provided in the "Tones and tuning with stretched partials" selection of the Auditory Demonstrations recording by Houtsma *et al.*,⁹ in which stretched sounds appear more consonant when played in the corresponding stretched octaves than when played in "real" octaves. To talk about this kind of effect more generally, a spectrum and a scale are said to be *related* if the dissonance curve has minima at the scale steps. Thus Fig. 1 shows that the Just Intonation scale and harmonic sounds are related. It is easy to find the related scale for a given spectrum simply by drawing the dissonance curve. But the inverse problem of finding a spectrum that is related to a given scale is not as straightforward. This paper focuses on certain classes of scales (such as tetrachordal scales) which are defined by only a few different successive intervals, and presents an algorithm for constructing families of spectra related to these scales.

This is important because related spectra can provide the composer and/or performer with additional flexibility in terms of controlling the consonance and dissonance of a given piece. For example, the Pythagorean tuning is often criticized because its major third is sharp compared to the equal tempered third, which is itself sharper than the just third. This excessive sharpness is heard as a roughness or beating, and is especially noticeable in slow, sustained passages. Using a related spectrum that is specifically crafted for use in the Pythagorean tuning, however, can ameliorate much of this roughness. The composer or performer thus has the option of exploiting a smoother, more consonant third than is available when using unrelated spectra.

The next section reviews previous approaches to the spectrum selection problem, and recalls the principle of co-

^{a)}Electronic mail: sethares@ece.wisc.edu

1452 PAX
SETHARES, W.A.
1996

C₂

Specifying spectra for musical scales

William A. Sethares^{a)}

Department of Electrical and Computer Engineering, University of Wisconsin, Madison,
Wisconsin 53706-1691

(Received 5 July 1996; accepted for publication 3 June 1997)

The sensory consonance and dissonance of musical intervals is dependent on the spectrum of the tones. The dissonance curve gives a measure of this perception over a range of intervals, and a musical scale is said to be *related* to a sound with a given spectrum if minima of the dissonance curve occur at the scale steps. While it is straightforward to calculate the dissonance curve for a given sound, it is not obvious how to find related spectra for a given scale. This paper introduces a “symbolic method” for constructing related spectra that is applicable to scales built from a small number of successive intervals. The method is applied to specify related spectra for several different tetrachordal scales, including the well-known Pythagorean scale. Mathematical properties of the symbolic system are investigated, and the strengths and weaknesses of the approach are discussed.

© 1997 Acoustical Society of America. [S0001-4966(97)05509-4]

PACS numbers: 43.75.Bc, 43.75.Cd [WJS]

INTRODUCTION

The motion from consonance to dissonance and back again is a standard feature of most Western music, and several attempts have been made to explain, define, and quantify the terms “consonance” and “dissonance.” For example, Tenney¹ provides a historical overview that identifies five separate uses of the terms, and a longstanding debate surrounds the “reductionist” explanations of Plomp and Levelt² and Terhardt,³ and the proponents of “cultural conditioning” such as Cazden.⁴ One of the most successful of the reductionist approaches is called tonal or *sensory* dissonance, in which the sensory dissonance between pairs of sine waves is determined from psychoacoustic experiments. The sensory dissonance of more complex sounds is then defined to be the sum of the dissonances between all simultaneously sounding sine wave partials. Because sensory dissonance depends on the partials, sounds with different spectra may function differently. For instance, an interval may be quite consonant when played with one sound, but quite dissonant when performed with another.

The dissonance curve $D_F(r)$ is a function that describes how the sensory dissonance of a sound with spectrum F varies when played at different intervals r . Figure 1, for instance, shows a plot of the dissonance curve for a sound with six harmonic partials over a range of intervals slightly larger than an octave. The minima of this curve occur at the simple integer ratios of the Just Intonation scale, reinforcing the familiar notion that the most consonant (least dissonant) intervals for sounds with harmonic spectra are those with small integer ratios. The top axis shows the steps of the 12-tone equal tempered scale, which can be viewed as approximating many of these just ratios. Techniques for drawing dissonance curves are described in detail in Ref. 5 and a computer program is given in Ref. 6. These are based on an explicit parametrization of the perceptual data gathered by Plomp and Levelt² (and replicated in Ref. 7). These are in

turn closely related to Helmholtz’⁸ “beat theory” of dissonance in which the beating between higher harmonics causes a roughness that is perceived as dissonance.

Nonharmonic sounds can have dissonance curves that differ dramatically from Fig. 1, indicating that intervals with the most sensory consonance depend strongly on the structure of the partials of the sound. A dramatic example of this is provided in the “Tones and tuning with stretched partials” selection of the Auditory Demonstrations recording by Houtsma *et al.*,⁹ in which stretched sounds appear more consonant when played in the corresponding stretched octaves than when played in “real” octaves. To talk about this kind of effect more generally, a spectrum and a scale are said to be *related* if the dissonance curve has minima at the scale steps. Thus Fig. 1 shows that the Just Intonation scale and harmonic sounds are related. It is easy to find the related scale for a given spectrum simply by drawing the dissonance curve. But the inverse problem of finding a spectrum that is related to a given scale is not as straightforward. This paper focuses on certain classes of scales (such as tetrachordal scales) which are defined by only a few different successive intervals, and presents an algorithm for constructing families of spectra related to these scales.

This is important because related spectra can provide the composer and/or performer with additional flexibility in terms of controlling the consonance and dissonance of a given piece. For example, the Pythagorean tuning is often criticized because its major third is sharp compared to the equal tempered third, which is itself sharper than the just third. This excessive sharpness is heard as a roughness or beating, and is especially noticeable in slow, sustained passages. Using a related spectrum that is specifically crafted for use in the Pythagorean tuning, however, can ameliorate much of this roughness. The composer or performer thus has the option of exploiting a smoother, more consonant third than is available when using unrelated spectra.

The next section reviews previous approaches to the spectrum selection problem, and recalls the principle of co-

^{a)}Electronic mail: sethares@ece.wisc.edu




FIG. 1. Dissonance curve for a spectrum with fundamental at 500 Hz and six harmonic partials has minima at 1.0, 1.14, 1.2, 1.25, 1.33, 1.5, 1.67, 1.78, and 2.0, which are shown by the tick marks on the frequency axis. Observe that many of these coincide exactly with steps of the Just Intonation scale, and coincide approximately with 12-tet scale steps, which are shown above for comparison.

inciding partials, which can be used to transform the problem into simpler form. A symbolic system is then introduced along with a method of constructing related spectra. Several examples are given in detail, and related spectra are found for a Pythagorean scale and for a diatonic tetrachordal scale. A simple pair of examples then shows that it is not always possible to find such related spectra. The symbolic system is investigated in the Appendix, where several mathematical properties are revealed.

I. GENERAL TECHNIQUES

The problem of finding spectra for a specified scale was stated in Ref. 5 in terms of a constrained optimization problem that can sometimes be solved via iterative techniques such as the genetic algorithm¹⁰ or simulated annealing.¹¹ Though these approaches are very general, the problem is high dimensional (on the order of the number of partials in the desired spectrum), the algorithms run slowly (overnight, or worse), and they are not guaranteed to find optimal solutions (except “asymptotically”). Moreover, even when a good spectrum is found for a given scale, the technique gives no insight into the solution of other closely related spectrum selection problems. There must be a better way.

Several general properties of dissonance curves are given in Ref. 6. The fourth of these is the key to simplifying the spectrum selection problem:

Property of Coinciding Partial: Up to half of the $2n(n-1)$ minima of a dissonance curve occur at interval ratios r for which $r = f_i/f_j$, where f_i and f_j are partials of F .

In essence, whenever the j th partial of the lower tone coincides with the i th partial of the upper tone, there is a potential minimum of the dissonance curve. The minima corresponding to such intervals r typically “look like” the minima that occur at the simple integer ratios in Fig. 1. The other half of the potential minima are caused by more widely spaced partials that do not interact in a significant way. For instance, in Fig. 1, only the very shallow minimum at 1.78 is of this kind. Most musical tones are quite complex, with numerous partials, and the majority of minima are caused by

Pythagorean Diatonic Scale

ratio	cents	
1/1	0	
9/8	204	a=9/8
81/64	408	a
4/3	498	b=256/243
3/2	702	a
27/16	906	a
243/128	1110	a
2/1	1200	b

FIG. 2. Pythagorean major scale has intervals $a=9/8$ between all major seconds and $b=256/243$ between all minor seconds. It is laid out here in the “key” of C.

coinciding partials. This property of coinciding partials is the key to solving the spectrum selection problem in certain situations. The simplest case is for equal temperaments.

A. Spectra for equal temperaments

The ratio between successive scale steps in the 12-tone equal-tempered (abbreviated 12-tet) scale is the 12th root of 2, $\sqrt[12]{2}$, or about 1.0595. Similarly, m -tet has a ratio of $s = \sqrt[m]{2}$ between successive steps. Consider spectra for which successive partials are ratios of powers of s . Each partial of such a spectrum, when transposed into the same octave as the fundamental, lies on a note of the scale. Such a spectrum is said to be *induced* by the m -tet scale.

Induced spectra are good candidate solutions to the spectrum selection problem since the ratio between any pair of partials in an induced spectrum is s^k for some integer k . By the property of coinciding partials, the dissonance curve will tend to have minima precisely at steps of the scale. Thus such spectra will have low dissonance at scale steps, and many of the scale steps will be minima.

This insight can be exploited in two ways. First, it can be used to reduce the search space of the optimization routine. Instead of searching over all frequencies in a bounded region, the search need only be conducted over induced spectra. More straightforwardly, the spectrum selection problem for equal tempered scales can be solved by careful choice of induced spectra. In Ref. 5, this method was used to find spectra related to 10-tet, and other equal temperaments are equally straightforward. Unfortunately, it is not so clear how to proceed when confronted with nonequal tunings.




FIG. 3. Tetrachordal scales divide the octave into two 4:3 tetrachords separated by an interval of 9:8. The tetrachords are each divided into three intervals to form a seven note scale, which is labeled in the key of C.

B. Pythagorean scales and the tetrachord

To see why spectrum selection is more difficult for non-equal tunings, consider the Pythagorean diatonic scale, which is shown in Fig. 2 mapped to the “key” of C. This scale is created¹² from a series of pure 3/2 fifths (translated back into the original octave whenever necessary), and all seven of the fifths are pure. An interesting structural feature is that there are only two successive intervals, a “whole step” of $a=9/8$ and a “half step” of $b=256/243$. This whole step is 4 cents larger than the equal tempered version, while the half step is 10 cents smaller than in 12-tet.

In attempting to mimic the “induced spectrum” idea of the previous section, it is natural to attempt to place the partials at scale steps. Unfortunately, the intervals between scale steps are not necessarily scale steps themselves. For instance, if one partial occurred at the seventh ($f_i=243/128$) and the other at the fourth ($f_j=4/3$), then a minimum of the dissonance curve might occur at $r=f_i/f_j=a^3=729/512$, which is not a scale step. Similarly, the ratio between a partial at 4/3 and another at 81/64 is $256/243=b$, which again is not a scale step.


The Pythagorean scale is one example of a large class of scales based on “tetrachords”¹³ which were advocated by a number of ancient theorists such as Archytas, Aristoxenus, Didymus, Eratosthenes, and Ptolemy.¹⁴ A tetrachord is an interval of a pure fourth (a ratio of 4/3) that is divided into three subintervals. Combining two tetrachords around a central interval of 9/8 forms a seven tone scale spanning the octave. For instance, Fig. 3 shows two tetrachords divided into intervals r, s, t and r', s', t' . When $r=r', s=s',$ and $t=t'$, the scale is called an equal-tetrachordal scale. The Pythagorean scale is the special equal-tetrachordal scale where $r=r'=s=s'=9/8$. A modern treatment of tetrachords and tetrachordal scales is available in Ref. 15.

II. A SYMBOLIC SYSTEM

This section presents a symbolic system that uses the desired scale to define an operation that generates “strings” representing spectra, i.e., sets of partials. Admissible strings have all ratios between all partials equal to some interval in the scale, and thus are likely to be related spectra, via the property of coinciding partials.

A. Basic definitions

A desired scale S can be specified either in terms of a set of intervals $(s_0, s_1, s_2, \dots, s_m)$ with respect to some fundamental frequency f or by the successive ratios $r_i = s_i/s_{i-1}$.



For instance, for the Pythagorean diatonic scale of Fig. 2,

$$S = (1, 9/8, 81/64, 4/3, 3/2, 27/16, 243/128, 2/1),$$

and r_i is either $a=9/8$ or $b=256/243$ for all i . The intervals $s_i \in S$ are called the *scale intervals*.

A spectrum F is defined by a set of partials with frequencies at $(f_1, f_2, \dots, f_n)$. The property of coinciding partials suggests that related spectra can be constructed by ensuring that the ratios of the partials are equal to scale steps. The following definitions distinguish the situation where all ratios of all partials are equal to some scale step, from the situation where all scale steps occur as a ratio of some pair of partials.

Definition: If for each i and j there is a k such that $f_i/f_j = s_k$, then the spectrum is called *complementary* to the scale.

Definition: If for each k there is at least one pair of i and j such that $s_k = f_i/f_j$, then the spectrum is called *complete* with respect to the scale.

If a spectrum is both complete and complementary, then it is called *perfect* with respect to the given scale. Of course, scales and spectra need not be perfect in order to sound good or to be playable, and many scales have no perfect spectra at all. Nonetheless, when perfect spectra exist, they are ideal candidates.

B. An example

The simplest nonequal scales are those with only a small number of different successive ratios. For example, one scale generated by two intervals a and b has scale intervals

$$\begin{aligned} s_0 &= 1, \quad s_1 = a, \quad s_2 = ab, \quad s_3 = a^2b, \quad s_4 = a^2b^2, \\ s_5 &= a^3b^2, \quad \text{and} \quad s_6 = a^3b^3 = 2, \end{aligned} \quad (1)$$

where a and b are any two numbers such that $a^3b^3=2$. For this scale,

$$r_1 = a, \quad r_2 = b, \quad r_3 = a, \quad r_4 = b, \quad r_5 = a, \quad \text{and} \quad r_6 = b.$$

To see how it might be possible to build up a perfect spectrum for this scale, suppose that the first partial is selected arbitrarily at f_1 . Then f_2 must be

$$af_1, \quad abf_1, \quad a^2bf_1, \quad a^2b^2f_1, \quad a^3b^2f_1, \quad \text{or} \quad 2f_1 \quad (2)$$

since any other interval will cause f_2/f_1 to be outside the scale intervals. Suppose, for instance, that $f_2 = a^2bf_1$ is selected. Then f_3 must be chosen so that f_3/f_1 and f_3/f_2 are both scale intervals. The former condition implies that f_3 must be one of the intervals in (2) while the latter restricts f_3 even further. For instance, $f_3 = a^3b^2f_1$ is possible since $a^3b^2f_1/a^2bf_1 = ab$ is one of the scale intervals in (1). But $f_3 = a^3b^3f_1$ is not possible since $a^3b^3f_1/a^2bf_1 = ab^2$ is not one of the scale intervals. Clearly, building complementary

TABLE I. $\oplus$ -table for the scale defined in (1).

$\oplus$	(0,0)	(1,0)	(1,1)	(2,1)	(2,2)	(3,2)
(0,0)	(0,0)	(1,0)	(1,1)	(2,1)	(2,2)	(3,2)
(1,0)	(1,0)	*	(2,1)	*	(3,2)	*
(1,1)	(1,1)	(2,1)	(2,2)	(3,2)	(0,0)	(1,0)
(2,1)	(2,1)	*	(3,2)	*	(1,0)	*
(2,2)	(2,2)	(3,2)	(0,0)	(1,0)	(1,1)	(2,1)
(3,2)	(3,2)	*	(1,0)	*	(2,1)	*

spectra for nonequal scales requires more care than in the equal tempered case where partials can always be chosen to be scale steps. For some scales, no complementary spectra may exist. For some, no complete spectra may exist.

C. Symbolic computation of related spectra

This process of building spectra rapidly becomes complex. A symbolic table called the $\oplus$ -table (pronounced “Oh-plus-table”) simplifies and organizes the choices of possible partials at each step. The easiest way to introduce this is to continue with the example of the previous section.

Let the scalar intervals in (1) be written (1,0), (1,1), (2,1), (2,2), (3,2), and (3,3), where the first number is the exponent of a and the second is the exponent of b . Since the scale is generated by a repeating pattern, i.e., it is assumed to repeat at each octave, (3,3) is equated with (0,0). Basing the scale on the octave is not necessary, but it simplifies the discussion. The $\oplus$ -Table I represents the relationships between all the scale intervals. The table shows, for instance, that the interval a^2b combined with the interval ab gives the scale interval a^3b^2 , which is notated $(2,1) \oplus (1,1) = (3,2)$.

The asterisk indicates that the given product is not permissible since it would result in intervals that are not scalar intervals. Thus $a^2b = (2,1)$ cannot be $\oplus$ -added to $a = (1,0)$ since together they form the interval a^3b which is not an interval of the scale. Observe that the “octave” has been exploited whenever the product is greater than 2. For instance, $(1,1) \oplus (3,2) = (4,3)$. When reduced back into the octave, (4,3) becomes (1,0) as indicated in the table, expressing the fact that $a^4b^3/a^3b^3 = a^1b^0$. At first glance this may appear to be some kind of algebraic structure such as a group or a monad.¹⁶ However, algebraic structures require closure, i.e., that operations on members of the set give answers that remain within the set. The presence of the asterisks indicates that $\oplus$ does not define a closed operator.

TABLE II. A perfect spectrum for the scale (1).

i	1	2	3	4	5	6	7	k
t_i	(3,3)	(5,5)	(6,6)	(9,8)	(10,9)	(11,10)	(13,12)	
s_i	(0,0)	(2,2)	(0,0)	(3,2)	(1,0)	(2,1)	(1,0)	
$r_{i,k}$		(2,2)	(1,1)	(3,2)	(1,1)	(1,1)	(2,2)	1
			(0,0)	(1,0)	(1,0)	(2,2)	(0,0)	2
				(3,2)	(2,1)	(2,1)	(1,1)	3
					(1,0)	(3,2)	(1,0)	4
						(2,1)	(2,1)	5
							(1,0)	6

TABLE III. $\oplus$ -table for the Pythagorean scale defined in (4).

$\oplus$	(0,0)	(1,0)	(2,0)	(2,1)	(3,1)	(4,1)	(5,1)
(0,0)	(0,0)	(1,0)	(2,0)	(2,1)	(3,1)	(4,1)	(5,1)
(1,0)	(1,0)	(2,0)	*	(3,1)	(4,1)	(5,1)	*
(2,0)	(2,0)	*	*	(4,1)	(5,1)	*	*
(2,1)	(2,1)	(3,1)	(4,1)	*	(0,0)	(1,0)	(2,0)
(3,1)	(3,1)	(4,1)	(5,1)	(0,0)	(1,0)	(2,0)	*
(4,1)	(4,1)	(5,1)	*	(1,0)	(2,0)	*	*
(5,1)	(5,1)	*	*	(2,0)	*	*	*

D. Construction of spectra

The $\oplus$ -Table I was constructed from the scale steps given in Eq. (1); other scales S define analogous tables. This section shows how to use such $\oplus$ -tables to construct spectra related to a given scale.

Let S be a set of scale intervals with unit of repetition or “octave” s^* . Let $T = [S, s^* + S, 2s^* + S, 3s^* + S, \dots]$ be a concatenation of S and all its octaves. (The symbol “+” is used here in the sense of vector addition.) Each element of $s \in S$ represents an equivalence class $s + ns^*$ of elements in T .

Example: For the scale of the previous section,

$$S = [(0,0), (1,0), (1,1), (2,1), (2,2), (3,2)]$$

with octave $s^* = (3,3)$. Then

$$s^* + S = [(3,3), (4,3), (4,4), (5,4), (5,5), (6,5)],$$

$$2s^* + S = [(6,6), (7,6), (7,7), (8,7), (8,8), (9,8)],$$

etc., and T is a concatenation of these.

The procedure for constructing spectra can now be stated.

Symbolic Spectrum Construction:

(1) Choose $t_1 \in T$ and let $s_1 \in S$ be the corresponding representative of its equivalence class.

(2) For $i = 2, 3, \dots$, choose $t_i \in T$ with corresponding $s_i \in S$ so that there are $r_{i,i-j}$ with

$$s_i = s_j \oplus r_{i,i-j} \quad (3)$$

for $j = 1, 2, \dots, i-1$.

The result of this procedure is a string of t_i which defines a set of partials. By construction, the spectrum built from these partials is complementary to the given scale. If, in addition, all of the scale steps appear among either the s or the r , then the spectrum is complete, and hence perfect.

TABLE IV. A perfect spectrum for the Pythagorean scale (4).

i	1	2	3	4	5	6	7	k
t_i	(5,2)	(8,3)	(10,4)	(12,4)	(14,5)	(15,5)	(17,6)	
s_i	(0,0)	(3,1)	(0,0)	(2,0)	(4,1)	(5,1)	(2,0)	
$r_{i,k}$		(3,1)	(2,1)	(2,0)	(2,1)	(1,0)	(2,1)	1
			(0,0)	(4,1)	(4,1)	(3,1)	(3,1)	2
				(2,0)	(1,0)	(5,1)	(0,0)	3
					(4,1)	(2,0)	(2,0)	4
						(5,1)	(4,1)	5
							(2,0)	6




FIG. 4. Dissonance curve for the spectrum specially designed for play in the Pythagorean diatonic scale has minima at all the specified scale steps. Two extra “broad” minima are not caused by coinciding partials.

Equation (3) expresses the desire to have all of the intervals between all of the partials be scale intervals. A set of s_j are given (which are defined by previous choices of the t_j). Solving this requires finding a single s_i such that Eq. (3) is well defined for all j up to $i-1$. This can be done by searching all the columns s_j for an element s_i in common. If found, then the corresponding value of $r_{i,i-j}$ is given in the leftmost column. Whether this step is solvable for a particular i, j pair depends on the structure of the table and on the particular choices already made for previous s_i . Solution techniques for (3) are discussed at length in the Appendix.

It is probably easiest to understand the procedure by working through an example. One spectrum related to the scale (1) is given in Table II. This shows the choice of t_i , the corresponding scale steps s_i (which are the t_i reduced back into the octave), and the $r_{i,k}$ that complete Eq. (3). Since all the s_i and $r_{i,k}$ are scale steps, this spectrum is complementary. Since all scale steps can be found among the s_i or $r_{i,k}$, the spectrum is complete. Hence the spectrum of Table II is perfect for this scale. To translate the table into frequencies for the partials, recall that the elements t_i express the powers of a and b times an unspecified fundamental f . Thus the first partial is $f_1 = a^3 b^3 f$, the second is $f_2 = a^5 b^5 f$, etc.

III. PERFECT SPECTRA FOR PYTHAGOREAN SCALES

The Pythagorean diatonic scale of Fig. 2 is constructed from two intervals a and b in the order a, a, b, a, a, a, b . Thus the scale steps are given by

$$\begin{array}{cccccc} 1 & a & a^2 & a^2 b & a^3 b \\ (0,0) & (1,0) & (2,0) & (2,1) & (3,1) \\ & a^4 b & a^5 b & & a^5 b^2 = 2 \\ & (4,1) & (5,1) & \text{and} & (5,2) = (0,0). \end{array} \quad (4)$$

Typically, $a^2 b$ is a pure fourth. Along with the condition that $a^5 b^2 = 2$, this uniquely specifies $a = 9/8$ and $b = 256/243$, and so the scale contains two equal tetrachords separated by the standard interval $9/8$. The $\oplus$ -table for this Pythagorean scale is shown in Table III. These exact values are not necessary for the construction of the perfect spectra that follow, and it is not necessary that (5,2) be an exact octave; any “pseudo-octave”¹⁷ or interval of repetition will do.

Spectra can be assembled by following the procedure for symbolic spectrum construction, and one such spectrum is given in Table IV. Observe that all of the s_i and $r_{i,k}$ are scale steps, and that all seven scale steps are present among the s_i and the $r_{i,k}$. Hence this spectrum is perfect for the Pythagorean scale (4). Assuming the standard values for a and b , this spectrum has its partials at

$$f, 2f, 3f, 4f, \frac{81}{16} = 5.0625f, \frac{27}{4} = 6.75f, \quad (5)$$

$$\frac{243}{32} \approx 7.594f, \quad \text{and} \quad \frac{81}{8} = 10.125f.$$

The first several partials are harmonic, and this is the “closest” perfect Pythagorean spectrum to harmonicity. For example, there are no suitable partials between $(12,4) \approx 5$ and $(14,5) = 6.75$, and thus no way to closely approximate the sixth harmonic partial $6f$. It is easy to check that $(13,4)$ and $(14,4)$ are not scale steps, and that $(13,5) = (3,1)$ forms the interval ab with $(12,4)$. Since ab is not a scale step, $(13,5)$ cannot occur in a complementary spectrum. [However, $(13,5) = 6$ can be used if $(12,4)$ is replaced by $(11,4) = 9/2$. This would then sacrifice the accuracy of the fifth harmonic in order to increase the accuracy of the sixth. Trade-offs such as this are common.]

The dissonance curve for the Pythagorean spectrum (5) is shown in Fig. 4, under the assumption that the amplitude of the i th partial is 0.9^i . As expected from the principle of coinciding partials, this curve has minima that align with the scale steps. Thus there are significant minima at the just fourth and fifths, and at the Pythagorean third $81/64$ and the Pythagorean sixth $27/16$, rather than at the just thirds and sixths as in the harmonic dissonance curve from Fig. 1. This spectrum will not exhibit rough beating when its thirds or sixths are played in long sustained passages in the Pythagorean tuning. There are also two extra minimum which are very shallow and broad, and are not due to coinciding partials. The exact location and depth of these minima changes significantly as the amplitude of the partials are changed. As is usual for such extra minima, they are only barely distinguishable from the surrounding regions of the curve. Thus perfect spectra, as constructed by the symbolic procedure, do give dissonance curves with minima that correspond closely with scale steps of the desired scale.

IV. SPECTRUM FOR A DIATONIC TETRACHORD

A more general diatonic tetrachordal scale is constructed from three intervals a , b , and c in the order a, a, b, c, a, a, b . The scale steps are given by

$$\begin{array}{cccccc} 1 & a & ab & a^2 b & a^2 bc \\ (0,0,0) & (1,0,0) & (1,1,0) & (2,1,0) & (2,1,1) \\ & a^3 bc & a^3 b^2 c & & a^4 b^2 c = 2 \\ & (3,1,1) & (3,2,1) & \text{and} & (4,2,1) = (0,0,0). \end{array} \quad (6)$$

As before, $a^2 b$ is a pure fourth that defines the tetrachord. The new interval c is typically given by the interval remain-

TABLE V. $\oplus$ -table for the tetrachordal scale defined in (6).

$\oplus$	(0,0,0)	(1,0,0)	(1,1,0)	(2,1,0)	(2,1,1)	(3,1,1)	(3,2,1)
(0,0,0)	(0,0,0)	(1,0,0)	(1,1,0)	(2,1,0)	(2,1,1)	(3,1,1)	(3,2,1)
(1,0,0)	(1,0,0)	*	(2,1,0)	*	(3,1,1)	*	(0,0,0)
(1,1,0)	(1,1,0)	(2,1,0)	*	*	(3,2,1)	(0,0,0)	*
(2,1,0)	(2,1,0)	*	*	*	(0,0,0)	(1,0,0)	(1,1,0)
(2,1,1)	(2,1,1)	(3,1,1)	(3,2,1)	(0,0,0)	*	*	*
(3,1,1)	(3,1,1)	*	(0,0,0)	(1,0,0)	*	*	(2,1,1)
(3,2,1)	(3,2,1)	(0,0,0)	*	(1,1,0)	*	(2,1,1)	*

ing when two tetrachords are joined, and so $c=9/8$. There are no standard values for a and b . Rather, many different combinations have been explored over the years. The $\oplus$ -table for this diatonic tetrachordal scale is given in Table V. As before, it is not necessary that (4,2,1) be an exact octave, though it must define the intervals at which the scale repeats.

Spectra can be constructed by following the symbolic spectrum construction procedure, and one such spectrum is given in Table VI. Observe that all of the s_i and $r_{i,k}$ are scale steps and that all seven scale steps are present among the s_i or $r_{i,k}$. Hence this spectrum is perfect for the specified tetrachordal scale (6).

In order to draw the dissonance curve, it is necessary to pick particular values for the parameters a , b , and c . As mentioned above, $c=9/8$ is the usual difference between two tetrachords and the octave. Somewhat arbitrarily, let $b=10/9$, which, combined with the condition that $a^2b=4/3$ (i.e., forms a tetrachord) implies that $a=\sqrt{6/5}$. With these values, the spectrum defined in Table VI is

$$f, 2f, 3f, 4f, 6.57f, 8f, 12f, \text{ and } 16f,$$

and the resulting dissonance curve is given in Fig. 5 when the amplitude of the i th partial is 0.9^i . Minima occur at all scale steps except the first, the interval a . While this may seem like a flaw, it is really quite normal for very small intervals (like the major second) to fail to be consonant; the Pythagorean spectrum of the previous section was quite atypical in this respect. Again, although a few broad minima occur, they are fairly undistinguished from the surrounding intervals. Thus the symbolic method of spectrum construction has again found a spectrum that is well suited to the desired scale.

TABLE VI. A perfect spectrum for the tetrachordal scale (6).

i	1	2	3	4	5	6	7	k
t_i	(4,2,1)	(6,3,2)	(8,4,2)	(11,5,3)	(12,6,3)	(14,7,4)	(16,8,4)	
s_i	(0,0,0)	(2,1,1)	(0,0,0)	(3,1,1)	(0,0,0)	(2,1,1)	(0,0,0)	
$r_{i,k}$		(2,1,1)	(2,1,0)	(3,1,1)	(1,1,0)	(2,1,1)	(2,1,0)	1
			(0,0,0)	(1,0,0)	(0,0,0)	(3,2,1)	(0,0,0)	2
				(3,1,1)	(2,1,0)	(2,1,1)	(1,1,0)	3
					(0,0,0)	(0,0,0)	(0,0,0)	4
						(2,1,1)	(2,1,0)	5
							(0,0,0)	6

V. WHEN PERFECTION IS IMPOSSIBLE

The above examples may lull the unsuspecting into a belief that perfect spectra are possible for any scale. Unfortunately, this is not so. Consider first a simple scale built from three arbitrary intervals a , b , and c in the order a, b, c, a . The scale steps are

$$\begin{array}{ccccccc} 1 & a & ab & abc & & a^2bc=2 & \\ (0,0,0) & (1,0,0) & (1,1,0) & (1,1,1) & \text{and} & (2,1,1)=(0,0,0) & \end{array} \quad (7)$$

As suggested by the notation, (2,1,1) serves as the basic unit of repetition which would likely be the octave. The $\oplus$ -table for this scale is given in Table VII.

The difficulty with this scale is that the element (1,1,0) cannot be combined with any other. The symbolic construction procedure requires at each step that the s_i be expressible as a $\oplus$ -sum of s_j and some $r_{i,k}$. But it is clear that the operation does not allow (1,1,0) as a product with any element [other than the identity (0,0,0)] due to the column of asterisks. In other words, if the interval (1,1,0) ever appears as a partial in the spectrum or as one of the $r_{i,k}$, then the construction process must halt since no more complementary partials can be added. In this particular example, it is possible to create a perfect spectrum by having the element (1,1,0) appear only as the very last partial. However, such a strategy will not work if there are two columns of asterisks.

An extreme example for which no perfect spectrum is possible is a scale defined by four different intervals a , b , c , and d taken in alphabetical order. The scale steps are




FIG. 5. The dissonance curve for the spectrum related to the diatonic tetrachord with $a^2=6/5$, $b=10/9$, and $c=9/8$, has minima at all scale steps except for the first. The broad minima at 1.16, 1.41, and 1.71 are not caused by coinciding partials.

$$\begin{array}{ccccccc} 1 & a & ab & abc & & abcd=2 & \\ (0,0,0,0) & (1,0,0,0) & (1,1,0,0) & (1,1,1,0) & \text{and} & (1,1,1,1) & = (0,0,0,0) \end{array} \quad (8)$$

As suggested by the notation, (1,1,1,1) serves as the basic unit of repetition which would likely be the octave. The $\oplus$ -table for this scale is given in Table VIII.

Partials of a complementary spectrum for this scale can only have intervals that are multiples of the octave (1,1,1,1) due to the preponderance of disallowed asterisk entries in the $\oplus$ -table. The only possible complementary spectrum is (0,0,0,0) f , (1,1,1,1) f , (2,2,2,2) f , etc., which is clearly not complete, and hence not perfect. Thus a given scale may or may not have perfect spectra, depending on the number and placement of the asterisk entries in the table.

VI. DISCUSSIONS AND CONCLUSIONS

A. Related versus perfect spectra

Do not confuse the idea of a spectrum related to a given scale with the notion of a perfect (complete and complementary) spectrum for the scale. The former is based directly on a psychoacoustic measure of the perceived sensory dissonance of the sound, while the latter is a construction based on the coincidence of partials within the spectrum. The latter is best viewed as an approximation and simplification of the former, in the sense that it leads to a tractable system for spectrum determination via the principle of coinciding partials.

Some scale intervals that appear in the spectrum (i.e., among the s_i or the $r_{i,k}$ of Tables II, IV, or VI) may not be minima of the dissonance curve. For instance, the tetrachordal spectrum of Sec. IV does not have a minimum at the first scale step even though the spectrum is complete. Alternatively, some minima may occur in the dissonance curve that are not explicitly ratios of partials. Three such minima occur in Fig. 5; they are the broad kind of minima that are due to wide spacing between certain pairs of partials.

The notion of a perfect spectrum shows starkly that the most important feature of related spectra and scales are the coincidence of partials of a tone, a result that would not have surprised Helmholtz. Perhaps the crucial difference is that related spectra take explicit account of the amplitudes of the partials, whereas perfect spectra do not. In fact, by manipulating the amplitudes of the partials, it is possible to make various minima appear or disappear. For instance, it is possible to “fix” the problem that the tetrachordal spectrum is missing its first scale step a by increasing the amplitudes of the partials that are separated by the ratio a . Alternatively, it is often possible to remove a minimum from the dissonance curve of a perfect spectrum by decreasing the amplitudes of the partials separated by that interval. Moreover, while a minimum due to coinciding partials may be extinguished by manipulating the amplitudes, its location (the interval it forms) remains essentially fixed. In contrast, the broad type minima that are not due to coinciding partials move continuously as the amplitudes vary; they are not a fixed feature of the dissonance curve of a perfect spectrum. In choosing the

amplitudes of the partials of a perfect spectrum, care must also be taken to avoid masking one partial by another.

B. Almost perfect spectra

As the number of different intervals in a desired scale increases, it becomes more difficult to find perfect spectra; the $\oplus$ -tables become less full (i.e., have more disallowed asterisk entries) and fewer solutions to Eq. (3) exist. There are several simple modifications to the procedure that may result in spectra that are well matched to the given scale, even when perfection is impossible. One simple modification is to allow the spectrum to be incomplete. Since very small intervals are unlikely to be consonant with any reasonable amplitudes of the partials, they may be safely removed from consideration. A second simplifying strategy is to relax the requirement of complementarity, while it is certainly important that prominent scale steps occur at minima, it is not obviously harmful if some extra minima exist. Indeed, if an extra minimum occurs in the dissonance curve but never appears in the music, then its existence is transparent to the listener.

A third method of relaxing the procedure can be applied whenever the scale is specified only over an octave (or over some pseudo-octave), in which case the completeness and complementarity need only hold over each octave. For instance, a partial t_i might be chosen even though it forms a disallowed interval with a previous partial t_j , providing the two are more than an octave apart. Thus judicious relaxation of various elements of the procedure may allow specification of useful spectra even when perfect spectra are not possible.

C. The mathematician's view

From a mathematical point of view, the symbolic spectrum selection procedure raises a number of interesting issues. The $\oplus$ operation defined here is not any kind of standard mathematical operator because of the disallowed asterisk entries. Though they do not form any recognizable algebraic structure, $\oplus$ tables do have several features that would be familiar to an algebraist. For instance, the tables have an identity element, the operation $\oplus$ is commutative, and it is associative when it is well defined. These are used in the Appendix to derive a set of properties that can be used to streamline the symbolic spectrum construction procedure.

$\oplus$ -tables clearly have a significant amount of structure. For instance, any $\oplus$ -table can be viewed as a subset of the commutative group of integer m vectors $(\sigma_1, \sigma_2, \dots, \sigma_m)$ where the i th entry is taken mod n_i , from which certain elements have been removed. Can this structure be exploited? Another obvious question concerns the possibility of decomposing $\oplus$ -tables in the same kind of ways that arbitrary groups are decomposed into normal subgroups. Might such a decomposition allow the building up of spectra for larger scales in terms of spectra for simpler scales?

TABLE VII. $\oplus$ -table for the scale defined in (7).

$\oplus$	(0,0,0)	(1,0,0)	(1,1,0)	(1,1,1)
(0,0,0)	(0,0,0)	(1,0,0)	(1,1,0)	(1,1,1)
(1,0,0)	(1,0,0)	*	*	(0,0,0)
(1,1,0)	(1,1,0)	*	*	*
(1,1,1)	(1,1,1)	(0,0,0)	*	*

D. The composer's view

From the musical point of view, perfect spectra raise a number of issues. For instance, a given nonequal scale sounds different in each key because the set of intervals is slightly different. How would the use of perfect spectra influence the ability to modulate through various keys? Certain chords will become more or less consonant when played with perfect spectra than when played with harmonic spectra. What patterns of (non)harmonic motion are best suited to perfect spectra and their chords? Will perfect spectra be useful for some part of the standard repertoire, or will they be only useful for new compositions that directly exploit their strengths (and avoid their weaknesses)?

E. Informal experiments

Preliminary experiments with perfect spectra are encouraging. A simple “organ” sound with eight harmonic partials was generated via additive synthesis. A second organlike sound which is perfect for the Pythagorean scale was generated with partials specified in (5). All parameters except for the frequencies of the partials were identical. Both sound pleasant, if somewhat bland and “electronic.” The Pythagorean spectrum, though nonharmonic, gives a definite sense of pitch, and is well fused. It also has a slightly “brighter” sound, probably because the highest partials occur at somewhat higher frequencies than in the harmonic version. The Bach chorale “Aus Meine Herzens Grunde” was recorded as a standard MIDI file and comparisons were made between the piece when played in (1) 12-tet with the harmonic spectrum, (2) Pythagorean tuning (in G) with the harmonic spectrum, (3) Pythagorean tuning (in G) with the Pythagorean spectrum, and (4) 12-tet with the Pythagorean spectrum. The differences between the four versions are subtle, but clear. For instance, there are several sustained major thirds between the alto and soprano lines, as in measures 4, 10, 13, and 16. In 12-tet, beats can be readily perceived between these two voices. The beats are even more pronounced in the Pythagorean tuning due to the stretching of the thirds. However, when the perfect spectrum is played in its related scale, the beats disappear, and the clarity of the chord increases. The fourth case is only marginally distinguishable from the first two, emphasizing that the Pythagorean spectrum itself is not overly bizarre.

It is certainly not true that all music will sound greatly improved (or even much different) when employing perfect spectra. For instance, the stylistically similar Bach chorale “Als der Gütige Gott” (as was used in Ref. 9 to demonstrate the effect of stretched spectra) does not show the same effect. With this piece, there is almost no noticeable difference (other than the brightness of the spectrum) between the per-

TABLE VIII. $\oplus$ -table for the simple scale defined in (8).

$\oplus$	(0,0,0,0)	(1,0,0,0)	(1,1,0,0)	(1,1,1,0)
(0,0,0,0)	(0,0,0,0)	(1,0,0,0)	(1,1,0,0)	(1,1,1,0)
(1,0,0,0)	(1,0,0,0)	*	*	*
(1,1,0,0)	(1,1,0,0)	*	*	*
(1,1,1,0)	(1,1,1,0)	*	*	*

fect and harmonic versions. On investigation, it became clear that this is because there are no sustained major thirds in the chorale; the “held” chords are all inversions that avoid close position major thirds. Hence the piece avoids resolving to those chords that would sound most aggressive in the Pythagorean tuning.

F. Nonharmonic sounds

Exploiting nonharmonic sounds is a topic of considerable interest to the computer music community,^{18,19} and the notion of perfect spectra helps to specify classes of potentially useful nonharmonic sounds. Perfect spectra stipulate the frequencies of the partials, but leave the amplitudes free. Since spectrum is just one aspect of timbre, many different timbres may share a given spectrum. For instance, brassy timbres arise (at least in part) from a rise in the spectral energy in the higher partials, while flute timbres are dependent on a breathy puff in the attack. (For an overview of the physical correlates of timbre, see Ref. 20.) For spectra such as the Pythagorean, with partials that are detuned only a few percent from harmonic, it is likely that analogous increases in the energy of the higher partials will tend to be heard as trumpetlike, while breathy puffs of air in the attack will tend to cause the sound to appear flutelike. Thus each perfect spectrum defines a whole class of timbres which may sound as different from each other as a trumpet from a flute.

Finally, the method does not give any indication of how such sounds might be generated or created. One obvious way is via additive synthesis. Another is via the technique of “spectral mapping,”²¹ which directly manipulates the partials of a sampled sound. A much more difficult question is how acoustic instruments might be given the kinds of deviations from harmonicity that are specified by perfect spectra.

APPENDIX: PROPERTIES OF $\oplus$ -TABLES

Given any set of scale intervals S , the $\oplus$ -table derived from S has the following characteristics.

Identity: The “octave” or unit of repetition s^* acts as an identity element, i.e.,

$$s^* \oplus s = s \oplus s^* = s \quad \forall s \in S.$$

Commutativity: The $\oplus$ -table is symmetric, i.e.,

$$s_1 \oplus s_2 = s_2 \oplus s_1 \quad \forall s_1, s_2 \in S. \quad (A1)$$

If one side of (A1) is undefined (is “equal” to *), then so is the other. Commutativity of $\oplus$ follows directly from the commutativity of products of powers of real numbers.

Associativity: The $\oplus$ operator is associative whenever it is well defined. Thus

$$(s_1 \oplus s_2) \oplus s_3 = s_1 \oplus (s_2 \oplus s_3) \quad \forall s_1, s_2, s_3 \in S \quad (\text{A2})$$

provided that both sides of (A2) exist.

It is indeed possible for one side of (A2) to exist but not the other.

Example: The tetrachordal scale (6) has $\oplus$ -Table V. Observe that $((2,1,1) \oplus (1,0,0)) \oplus (2,1,0)$ is well defined and equals $(1,0,0)$, but that $(2,1,1) \oplus ((1,0,0) \oplus (2,1,0))$ does not exist because $(1,0,0) \oplus (2,1,0)$ is disallowed. To further emphasize how unusual this construction is, observe that by commutativity, $(2,1,1) \oplus (1,0,0) = (1,0,0) \oplus (2,1,1)$. Substituting this in the above calculation gives $((1,0,0) \oplus (2,1,1)) \oplus (2,1,0)$ which is indeed equal to $(1,0,0) \oplus ((2,1,1) \oplus (2,1,0))$, since both sides are $(1,0,0)$.

The remaining properties of $\oplus$ -tables concern “solutions” to the equation

$$s_i = s_j \oplus r_{i,i-j} \quad (\text{A3})$$

that arises in the symbolic spectrum construction procedure. Recall that in the procedure, a set of s_j are given (which are defined by previous choices of the t_j). The goal is to find a single s_i such that Eq. (A3) is well defined for all j up to $i-1$. The properties of $\oplus$ -tables can help pinpoint viable solutions to (A3).

Theorem A.1: Suppose that $s_j \in S$ have been chosen for all $j < k$. Let $\mathbf{S}_j$ be the set of all non-* entries in the s_j column of the $\oplus$ -table. Then for all $i \geq k$, s_i must be an element of $\bigcap_{j < k} \mathbf{S}_j$.

Proof: First consider the case $i = k = 2$, with s_1 specified. Then (A3) requires choice of s_2 such that $s_2 = s_1 \oplus r_{1,1}$ for some $r_{1,1}$. Such $r_{1,1}$ will exist exactly when $s_2 \in \mathbf{S}_1$. For $i > 2$, $s_i = s_1 \oplus r_{i,i-1}$ must be solvable, which again requires that $s_i \in \mathbf{S}_1$. The general case $s_i = s_j \oplus r_{i,i-j}$ is similarly solvable exactly when $s_i \in \mathbf{S}_j$. Since this is true for every $j < k$, $s_i \in \bigcap_{j < k} \mathbf{S}_j$. $\square$

Thus when building spectra according to the procedure, the set $\mathcal{S}^k = \bigcap_{j < k} \mathbf{S}_j$ defines the allowable partials at the k th step. Clearly, $\mathcal{S}^k$ can never grow larger since $\mathcal{S}^k \supset \mathcal{S}^{k+1} \forall k$, and it may well become smaller as k increases. This demonstrates that the order in which the partials are chosen is crucial in determining whether a perfect spectrum is realizable.

The easiest way to appreciate how the theorem A.1 simplifies (and limits) the selection problem is by example.

Example: In Table I, once $s_i = (3,2)$ for some i , then for all $k > i$, s_k must be $(3,2)$, $(1,0)$, or $(2,1)$.

Example: In Table III, once $s_i = (2,0)$ has been chosen, then for all $k > i$, s_k must be either $(2,0)$, $(4,1)$, or $(5,1)$. In particular, no s_k can be the identity $(0,0)$.

Corollary A.1: Suppose that an element $\hat{s} \in S$ appears in every column of the $\oplus$ -table. Then for any choice of s_j , $j < i$, (A3) is always solvable with $s_i = \hat{s}$.

Proof: Since $\hat{s}$ is in every column of the table, $\hat{s} \in \mathbf{S}_j \forall j$ and hence $\hat{s} \in \bigcap_{j < k} \mathbf{S}_j$ for any k . $\square$

In other words, for any $s \in S$, there is always a $r \in S$ such that $\hat{s} = s \oplus r$, and so $\hat{s}$ is always permissible.

Example: In Table V, the identity $s^* = (0,0,0)$ appears in every column. Thus it is always possible to choose a partial t_i with the equivalence class s^* at any step.

Suppose, on the other hand, that an element $\bar{s} \in S$ appears nowhere in the $\oplus$ -table other than in the column and row of the identity. Then $\bar{s}$ cannot be used to define one of the s_i since $\bar{s} \notin \mathbf{S}_k$ for any k and so for any $s_i \neq s^*$, $s_i = \bar{s} + r$ has no solution. Although $\bar{s}$ cannot occur among the s_i , it is still possible that it might appear among the $r_{i,k}$. Indeed, it will need to in order to find a complete spectrum.

Example: The element $\bar{s} = (2,1)$ appears nowhere in $\oplus$ -table III defined by the Pythagorean scale. The spectrum was made complete by ensuring that $\bar{s}$ appears among the $r_{i,k}$ of Table IV.

Another property of $\oplus$ -tables is that elements are arranged in “stripes” from southwest to northeast. For instance, in Table III, a stripe of $(4,1)$ elements connects the $4,1$ entry with the $1,4$ entry. Similarly, a stripe of $(3,1)$ elements connect the $3,1$ with the $1,3$ entries, although the stripe is broken up by a *. The fact that such (possibly interrupted) stripes must exist is the content of the next theorem.

Given a m note scale S , the entries of the corresponding $\oplus$ -table can be labeled as a matrix $\{a_{j,k}\}$ for $j = 1, 2, \dots, m$ and $k = 1, 2, \dots, m$. Let P_i denote the i th stripe of the $\oplus$ -table, that is, $P_i = \{a_{j,k}\}$ for all j and k with $j+k = i+1$.

Example: For the Pythagorean $\oplus$ -table,

$$P_1 = \{(0,0)\}, \quad P_2 = \{(1,0), (1,0)\},$$

$$P_3 = \{(2,0), (2,0), (2,0)\},$$

$$P_4 = \{(2,1), *, *, (2,1)\},$$

$$P_5 = \{(3,1), (3,1), *, (3,1), (3,1)\}, \text{ etc.}$$

Theorem A.2: For each i , all non-* elements of the stripe P_i are identical.

Proof: By construction, the elements s_i and $s_{i+1} \in S$ are integer vectors, and they may be ordered so that

$$s_{i+1} = s_i + e_{j,i} \forall i, \quad (\text{A4})$$

where $e_{j,i}$ is a unit vector with zeroes everywhere except for a single 1 in the j th entry. Let $\Sigma(s_i)$ represent the sum of the entries in $s_i = (\sigma_1, \sigma_2, \dots, \sigma_p)$, i.e., $\Sigma(s_i) = \sum_{j=1}^p \sigma_j$, and let Σ^* represent the sum of the entries in the element that forms the unit of repetition. Because the $\oplus$ operation adds powers of the generating intervals,

$$\Sigma(s_j \oplus s_k) = \Sigma(s_j) + \Sigma(s_k) \pmod{\Sigma^*} \quad (\text{A5})$$

whenever $s_j \oplus s_k$ is well defined. Because of the ordering, the entries in the stripe P_i can be written

$$s_j \oplus s_k, \quad s_{j-1} \oplus s_{k+1}, \quad s_{j-2} \oplus s_{k+2} \dots$$

for all positive j and k with $j+k = i+1$. Hence

$$\Sigma(s_j \oplus s_k) = \Sigma(s_{j-1} \oplus s_{k+1}) = \dots \quad (\text{A6})$$

whenever these are defined. From (A4), $\Sigma(s_j) = \Sigma(s_k)$ implies that $s_j = s_k$. Hence (A6) shows that $s_j \oplus s_k = s_{j-1} \oplus s_{k+1} = \dots$ whenever the terms are defined, and hence all well-defined elements of the stripe are identical. $\square$

This is useful because stripes define whether a given choice for the t_i (and hence s_i) is likely to lead to complete

spectra. Suppose that $\tilde{s}$ is a candidate for s_i at the i th step. Whether $\tilde{s}$ will “work” for all previous s_j (i.e., whether $\tilde{s} = s_j \oplus r$ has solutions for all s_j) depends on whether $\tilde{s}$ appears in all the corresponding S_j . Theorem A.2 pinpoints exactly where $\tilde{s}$ must appear, at the intersection of the column S_j and the stripe containing $\tilde{s}$. Thus the procedure can be implemented without conducting a search for $\tilde{s}$ among all possible columns.

A special case is when a column is “full,” i.e., when it contains no $*$ entries.

Theorem A.3: Let S_f be a full column corresponding to $s_f \in S$. Then $s_i = s_f \oplus r_i$ is solvable for all $s_i \in S$.

Proof: Since there are m entries in the column S_f and there are m different s_i , it is only necessary to show that no entries appear twice. Using the ordering (A4) of the previous proof, S_f has elements

$$s_1 \oplus s_f, \quad s_2 \oplus s_f, \dots, s_m \oplus s_f, \quad (\text{A7})$$

which are well defined by assumption. Now proceed by contradiction, and suppose that the i th and j th elements of (A7) are the same, i.e., $s_i \oplus s_f = s_j \oplus s_f$. Then

$$\Sigma(s_i \oplus s_f) = \Sigma(s_j \oplus s_f) \pmod{\Sigma^*}$$

(where Σ and Σ^* were defined in the previous proof). This implies that

$$\Sigma(s_i) + \Sigma(s_f) = \Sigma(s_j) + \Sigma(s_f) \pmod{\Sigma^*}$$

which implies that $\Sigma(s_i) = \Sigma(s_j) \pmod{\Sigma^*}$. By the same argument as in the proof of theorem A.2, this implies that $s_i = s_j$. But each s_i appears exactly once in (A7), which gives the desired contradiction. $\square$

Thus when a column is full, it must contain every element. In this case, Eq. (A3) puts no restrictions on the choice of s_i . Let $\{s_j\}$ be all the elements of S that have full columns. Then a $\oplus$ -subtable can be formed by these $\{s_j\}$ that has no illegal $*$ entries. For example, recall Table I which is generated by the scale (1). The elements (0,0), (1,1), and (2,2) have full columns and hence can be used to form a full $\oplus$ -subtable. It is easy to generate perfect spectra for such full $\oplus$ -subtables because Eq. (A3) puts no restrictions on the choice of partials for a complementary spectrum. Whether these extend to all elements of the scale, however, depends heavily on the structure of the nonfull part of the table. Finding spectra for full subtables is exactly the same as finding

spectra for equal temperaments, whose $\oplus$ -tables have no disallowed $*$ entries. In fact, full $\oplus$ -tables form a commutative group, which may explain why the equal tempered case is relatively easy to solve.

All of the above properties were stated in terms of the columns of the $\oplus$ -table. By commutativity, the properties could have been stated in terms of the corresponding rows.

¹J. Tenney, *A History of ‘Consonance’ and ‘Dissonance’* (Excelsior, New York, 1988).

²R. Plomp and W. J. M. Levelt, “Tonal consonance and critical bandwidth,” *J. Acoust. Soc. Am.* **38**, 548–560 (1965).

³E. Terhardt, “Pitch, consonance, and harmony,” *J. Acoust. Soc. Am.* **55**, 1061–1069 (1974).

⁴N. Cazden, “The definition of consonance and dissonance,” *Int. Rev. Aesthetics Soc. Music* **11**, 123–168 (1980).

⁵W. A. Sethares, “Local consonance and the relationship between timbre and scale,” *J. Acoust. Soc. Am.* **94**, 1218–1228 (1993).

⁶W. A. Sethares, “Relating tuning and timbre,” *Exp. Musical Instrum.* **IX**, No. 2 (1993).

⁷A. Kameoka and M. Kuriyagawa, “Consonance theory part I: consonance of dyads,” *J. Acoust. Soc. Am.* **45**, 1452–1459 (1969); Also, “Consonance theory part II: consonance of complex tones and its calculation method,” *J. Acoust. Soc. Am.* **45**, 1460–1469 (1969).

⁸H. Helmholtz, *On the Sensations of Tone* (Dover, New York, 1954).

⁹A. J. M. Houtsma, T. D. Rossing, and W. M. Wagenaars, *Auditory Demonstrations* (Philips compact disc No. 1126-061 and text) (Acoustical Society of America, Woodbury, NY, 1987).

¹⁰S. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison-Wesley, New York, 1989).

¹¹S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, “Optimization by simulated annealing,” *Science* **220**, 4598 (1983).

¹²T. D. Rossing, *The Science of Sound* (Addison-Wesley, Reading, MA, 1990).

¹³D. B. Doty, *Just Intonation Primer* (Just Intonation Network, San Francisco, CA, 1993).

¹⁴J. M. Barbour, *Tuning and Temperament* (Michigan State College, East Lansing, 1951).

¹⁵J. Chalmers, Jr., *Divisions of the Tetrachord* (Frog Peak Music, Hanover, NH, 1993).

¹⁶S. Lang, *Algebra* (Addison-Wesley, Reading, MA, 1965).

¹⁷F. H. Slaymaker, “Chords from tones having stretched partials,” *J. Acoust. Soc. Am.* **47**, 1469–1571 (1968).

¹⁸W. Carlos, “Tuning: at the crossroads,” *Comput. Music J. Spring*, 29–43 (1987).

¹⁹M. V. Mathews and J. R. Pierce, “The Bohlen-Pierce scale,” in *Current Directions in Computer Music Research*, edited by M. V. Mathews and J. R. Pierce (MIT, Cambridge, MA, 1991).

²⁰D. M. Green, *An Introduction to the Psychology of Hearing* (Academic, New York, 1989), 3rd ed.

²¹W. A. Sethares, “Consonance-based spectral mappings,” *Comput. Music J.* (to appear in January 1998).