The Geometry of Musical Chords

The Geometry of Musical Chords

Dmitri Tymoczko

A musical chord can be represented as a point in a geometrical space called an orbifold. Line segments represent mappings from the notes of one chord to those of another. Composers in a wide range of styles have exploited the non-Euclidean geometry of these spaces, typically by using short line segments between structurally similar chords. Such line segments exist only when chords are nearly symmetrical under translation, reflection, or permutation. Paradigmatically consonant and dissonant chords possess different near-symmetries and suggest different musical uses.

Western music lies at the intersection of two seemingly independent disciplines: harmony and counterpoint. Harmony delimits the acceptable chords (simultaneously occurring notes) and chord sequences. Counterpoint (or voice leading) is the technique of connecting the individual notes in a series of chords so as to form simultaneous melodies. Chords are usually connected so that these lines (or voices) move independently (not all in the same direction by the same amount), efficiently (by short distances), and without voice crossings (along nonintersecting paths) (Fig. 1, A to C). These features facilitate musical performance, engage explicit aesthetic norms (1, 2), and enable listeners to distinguish multiple simultaneous melodies (3).

How is it that Western music can satisfy harmonic and contrapuntal constraints at once? What determines whether two chords can be connected by efficient voice leading? Musicians have been investigating these questions for almost three centuries. The circle of fifths (fig. S1), first published in 1728 (4), depicts efficient voice leadings among the 12 major scales. The Tonnetz (fig. S2), originating with Euler in 1739, represents efficient voice leadings among the 24 major and minor triads (2, 5). Recent work (5–13) investigates efficient voice leading in a variety of special cases. Despite tantalizing hints (6–10), however, no theory has articulated general principles explaining when and why efficient voice leading is possible. This report provides such a theory, describing geometrical spaces in which points represent chords and line segments represent voice leadings between their endpoints. These spaces show us precisely how harmony and counterpoint are related.

Human pitch perception is both logarithmic and periodic: Frequencies f and $2f$ are heard to be separated by a single distance (the octave) and to possess the same quality or chroma. To model the logarithmic aspect of pitch perception, I associate a pitch's fundamental frequency f with a real number p according to the equation

$$p = 69 + 12 \log_2(f/440) \qquad (1)$$

The result is a linear space (pitch space) in which octaves have size 12, semitones (the distance between adjacent keys on a piano) have size 1, and middle C is assigned the number 60. Distance in this space reflects physical distance on keyboard instruments, orthographical distance in Western musical notation, and musical distance as measured in psychological experiments (14, 15).

Musically, the chroma of a note is often more important than its octave. It is therefore useful to identify all pitches p and $p + 12$. The result is a circular quotient space (pitch-class space) that mathematicians call $\mathbb{R}/12\mathbb{Z}$ (fig. S3). (For a glossary of terms and symbols, see table S1.) Points in this space (pitch classes) provide numerical alternatives to the familiar letter-names of Western music theory: C = 0, C♯/D♭ = 1, D = 2, D quarter-tone sharp = 2.5, etc. Western music typically uses only a discrete lattice of points in this space. Here I consider the more general, continuous case. This is because the symmetrical chords that influence voice-leading behavior need not lie on the discrete lattice.

The content of a collection of notes is often more important than their order. Chords can therefore be modeled as multisets of either pitches or pitch classes. ("Chord" will henceforth refer to a multiset of pitch classes unless otherwise noted.) The musical term "transposition" is synonymous with the mathematical term "translation" and is represented by addition in pitch or pitch-class space. Transpositionally related chords are the same up to translation; thus, the C major chord, {C, E, G} or {0, 4, 7}, is transpositionally related to the F major chord, {F, A, C} or {5, 9, 0}, because {5, 9, 0} ≡ {0 + 5, 4 + 5, 7 + 5} modulo 12ℤ. The musical term "inversion" is synonymous with the mathematical term "reflection" and corresponds to subtraction from a constant value. Inversionally related chords are the same up to reflection; thus, the C major chord is inversionally related to the C minor chord {C, E♭, G}, or {0, 3, 7}, because {0, 3, 7} ≡ {7 − 7, 7 − 4, 7 − 0} modulo 12ℤ. Musically, transposition and inversion are important because they preserve the character of a chord: Transpositionally related chords sound extremely similar, inversionally related chords fairly so (movie S1).

A voice leading between two multisets {x_1, x_2, …, x_m} and {y_1, y_2, …, y_n} is a multiset of ordered pairs (x_i, y_j), such that every member of each multiset is in some pair. A trivial voice leading contains only pairs of the form (x, x). The notation (x_1, x_2, …, x_n) → (y_1, y_2, …, y_n) identifies the voice leading that associates the corresponding items in each list. Thus, the voice leading (C, C, E, G) → (B, D, F, G) associates C with B, C with D, E with F, and G with G. Music theorists have proposed numerous ways of measuring voice-leading size. Rather than adopting one, I will require only that a measure satisfy a few constraints reflecting widely acknowledged features of Western music (16). These constraints make it possible to identify, in polynomial time, a minimal voice leading
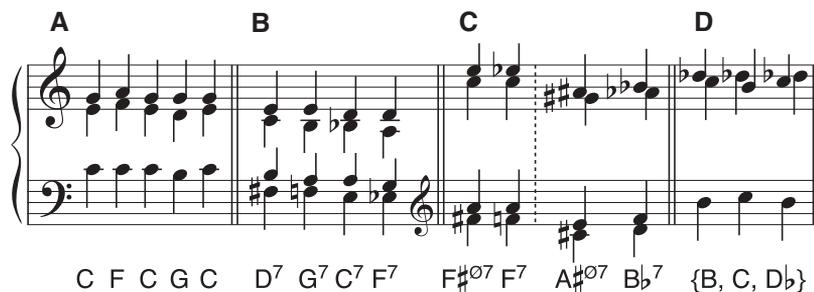
Fig. 1. Efficient voice leading between transpositionally and inversionally related chords. These progressions exploit three near-symmetries: transposition (**A** and **B**), inversion (**C**), and permutation (**D**). Sources: a common 18th-century upper-voice voice-leading pattern (A), a common jazz-piano voice-leading pattern, which omits chord roots and fifths and adds ninths and thirteenths (B), Wagner's *Parsifal* (C), Debussy's *Prélude à l'après-midi d'un faune* (C), and contemporary atonal voice leading in the style of Ligeti and Lutoslawski (D) (soundfile S1). Chord labels refer to unordered sets of pitch classes and do not indicate register.

Department of Music, Princeton University, Princeton, NJ 08544, USA, and Radcliffe Institute for Advanced Study, 34 Concord Avenue, Cambridge, MA 02138, USA. E-mail: dmitri@princeton.edu

(not necessarily bijective) between arbitrary chords (*16*). Every music-theoretical measure of voice-leading size satisfies these constraints.

I now describe the geometry of musical chords. An ordered sequence of *n* pitches can be represented as a point in $\mathbb{R}^n$ (fig. S4). Directed line segments in this space represent voice leadings. A measure of voice-leading size assigns lengths to these line segments. I will use quotient spaces to model the way listeners abstract from octave and order information. To model an ordered sequence of *n* pitch classes, form the quotient space $(\mathbb{R}/12\mathbb{Z})^n$, also known as the *n*-torus $\mathbb{T}^n$. To model unordered *n*-note chords of pitch classes, identify all points $(x_1, x_2, \ldots x_n)$ and $(x_{\sigma(1)}, x_{\sigma(2)}, \ldots x_{\sigma(n)})$, where σ is any permutation. The result is the global-quotient orbifold $\mathbb{T}^n/\mathcal{S}_n$ (*17, 18*), the *n*-torus $\mathbb{T}^n$ modulo the symmetric group $\mathcal{S}_n$. It contains singularities at which the local topology is not that of $\mathbb{R}^n$.

Figure 2 shows the orbifold $\mathbb{T}^2/\mathcal{S}_2$, the space of unordered pairs of pitch classes. It is a Möbius strip, a square whose left edge is given a half twist and identified with its right. The orbifold is singular at its top and bottom edges, which act like mirrors (*18*). Any bijective voice leading between pairs of pitches or pairs of pitch classes can be associated with a path on Fig. 2 (movie S2). Measures of voice-leading size determine these paths' lengths. They are the images of line segments in the parent spaces $\mathbb{T}^n$ and $\mathbb{R}^n$, and are either line segments in the orbifold or "reflected" line segments that bounce off its singular edges. For example, the voice leading (C, D♭) → (D♭, C) reflects off the orbifold's upper mirror boundary (Fig. 2).

Generalizing to higher dimensions is straightforward. To construct the orbifold $\mathbb{T}^n/\mathcal{S}_n$, take an *n*-dimensional prism whose base is an (*n* − 1) simplex, twist the base so as to cyclically permute its vertices, and identify it with the opposite face (figs. S5 and S6) (*16*). The boundaries of the orbifold are singular, acting as mirrors and containing chords with duplicate pitch classes. Chords that divide the octave evenly lie at the center of the orbifold and are surrounded by the familiar sonorities of Western tonality. Voice leadings parallel to the height coordinate of the prism act as transpositions. A free computer program written by the author allows readers to explore these spaces (*19*).

In many Western styles, it is desirable to find efficient, independent voice leadings between transpositionally or inversionally related chords. The progressions in Fig. 1 are all of this type (movie S3). A chord can participate in such progressions only if it is nearly symmetrical under transposition, permutation, or inversion (*16*). I conclude by describing these symmetries, explaining how they are embodied in the orbifolds' geometry, and showing how they have been exploited by Western composers.

A chord is transpositionally symmetrical (T-symmetrical) if it either divides the octave into

equal parts or is the union of equally sized subsets that do so (*20*). Nearly T-symmetrical chords are close to these T-symmetrical chords. Both types of chord can be linked to at least some of their transpositions by efficient bijective voice leadings. As one moves toward the center of the orbifold, chords become increasingly T-symmetrical and can be linked to their transpositions by increasingly efficient bijective voice leading. The perfectly even chord at the center of the orbifold can be linked to all of its transpositions by the smallest possible bijective voice leading; a related result covers discrete pitch-class spaces (*16*). Efficient voice leadings between perfectly T-symmetrical chords are typically not independent. Thus, composers

have reason to prefer near T-symmetry to exact T-symmetry.

It follows that the acoustically consonant chords of traditional Western music can be connected by efficient voice leading. Acoustic consonance is incompletely understood; however, theorists have long agreed that chords approximating the first few consecutive elements of the harmonic series are particularly consonant, at least when played with harmonic tones (*21*). Because elements *n* to 2*n* of the harmonic series evenly divide an octave in frequency space, they divide the octave nearly evenly in log-frequency space. These chords are therefore clustered near the center of the orbifolds (Table 1) and can typically be linked by efficient, independent
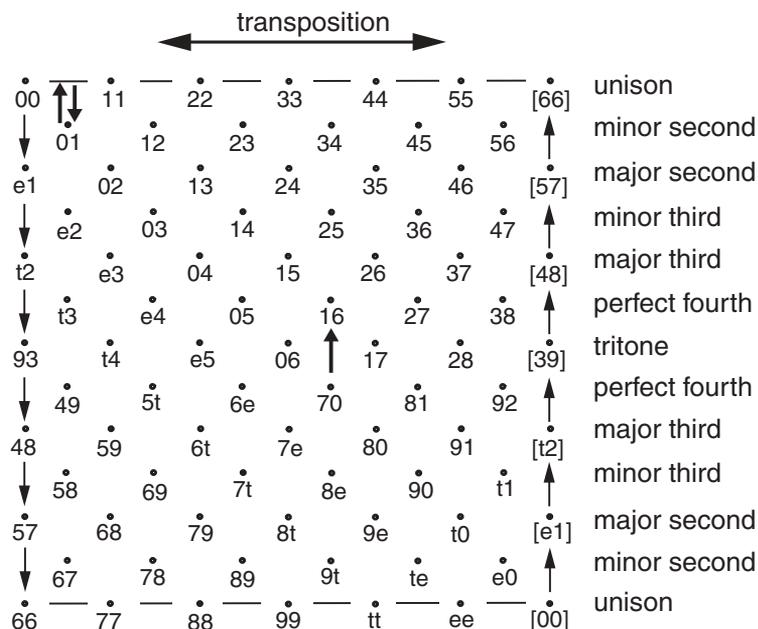
Fig. 2. The orbifold $\mathbb{T}^2/\mathcal{S}_2$. C = 0, C♯ = 1, etc., with B♭ = t, and B = e. The left edge is given a half twist and identified with the right. The voice leadings (C, D♭) → (D♭, C) and (C, G) → (C♯, F♯) are shown; the first reflects off the singular boundary.

Table 1. Common sonorities in Western tonal music. The center column lists the best equal-tempered approximation to the first *n* pitch classes of the harmonic series; the right column lists other good approximations. All divide the octave evenly or nearly evenly.

Size	Best approximation	Other approximations
2 notes	C, G	C, F♯
3 notes	C, E, G	C, E♭, G♭
		C, E♭, G
		C, E, G♯
4 notes	C, E, G, B♭	C, E♭, G♭, A
		C, E♭, G♭, B♭
		C, E♭, G, B♭
		C, E, G, B
5 notes	C, D, E, G, B♭	C, D, E, G, A
		C, D, E, G, B
6 notes	C, D, E, F♯, G, B♭	C, D, E♭, F, G, B♭
		C, D, E, F♯, G♯, B♭
7 notes	C, D, E, F♯, G, A, B♭	C, D, E, F, G, A, B♭
		C, D, E♭, F♯, G, A, B♭

voice leadings. Traditional tonal music exploits this possibility (Fig. 1, A to C, and movie S4). This central feature of Western counterpoint is made possible by composers' interest in the harmonic property of acoustic consonance.

A chord with duplicate pitch classes is permutationally symmetrical (P-symmetrical) because there is some nontrivial permutation of its notes that is a trivial voice leading. These chords lie on the singular boundaries of the orbifolds. Nearly P-symmetrical chords, such as {E, F, G♭}, are near these chords and contain several notes that are clustered close together. Efficient voice leadings permuting the clustered notes bounce off the nearby boundaries (Fig. 2 and movies S2 and S4). Such voice leadings can be independent and nontrivial. Trivial voice leadings are musically inert; therefore, as with T-symmetry, composers have reason to prefer near P-symmetry to exact P-symmetry.

Nearly P-symmetrical chords such as {B, C, D♭} are considered to be extremely dissonant. They are well-suited to static music in which voices move by small distances within an unchanging harmony (Fig. 1D). Such practices are characteristic of recent atonal composition, particularly the music of Ligeti and Lutoslawski. From the present perspective, these avant-garde techniques are closely related to those of traditional tonality: They exploit one of three fundamental symmetries permitting efficient, independent voice leading between transpositionally or inversionally related chords.

A chord is inversionally symmetrical (I-symmetrical) if it is invariant under reflection in pitch-class space. Nearly I-symmetrical chords are near these chords and can be found throughout the orbifolds (16). For example, the F♯ half-diminished seventh chord {6, 9, 0, 4} and the F dominant seventh chord {5, 9, 0, 3} are related by inversion and are very close to the I-symmetrical chord {5.5, 9, 0, 3.5}. Consequently, we can find an efficient voice leading between them, (6, 9, 0, 4) → (5, 9, 0, 3) (Fig. 1C) (16). Nearly T-symmetrical chords, such as the C major triad, and nearly P-symmetrical chords, such as {C, D♭, E♭}, can also be nearly I-symmetrical. Consequently, I-symmetry is exploited in both tonal and atonal music. It plays a salient role in the 19th century, particularly in the music of Schubert (22), Wagner (23), and Debussy (Fig. 1C).

The preceding ideas can be extended in several directions. First, one might examine in detail how composers have exploited the geometry of musical chords. Second, one could generalize the geometrical approach by considering quotient spaces that identify transpositionally and inversionally related chords (24). Third, because cyclical rhythmic patterns can also be modeled as points on $\mathbb{T}^n/S_n$, one could use these spaces to study African and other non-Western rhythms. Fourth, one could investigate how distances in the orbifolds relate to perceptual judgments of chord similarity. Finally, understanding the relation between harmony and counterpoint may suggest new techniques to contemporary composers.

References and Notes

1. C. Masson, *Nouveau Traité des Règles pour la Composition de la Musique* (Da Capo, New York, 1967).
2. O. Hostinský, *Die Lehre von den musikalischen Klängen* (H. Dominicus, Prague, 1879).
3. D. Huron, *Mus. Percept.* **19**, 1 (2001).
4. J. D. Heinichen, *Der General-Bass in der Composition* (G. Olms, New York, 1969).
5. R. Cohn, *J. Mus. Theory* **41**, 1 (1997).
6. J. Roeder, thesis, Yale University (1984).
7. E. Agmon, *Musikometrica* **3**, 15 (1991).
8. R. Cohn, *Mus. Anal.* **15**, 9 (1996).
9. C. Callender, *Mus. Theory Online* **10** (2004) (http://mto.societymusictheory.org/issues/mto.04.10.3/mto.04.10.3.callender.pdf).
10. G. Mazzola, *The Topos of Music* (Birkhäuser, Boston, 2002).
11. R. Morris, *Mus. Theory Spectrum* **20**, 175 (1998).
12. J. Douthett, P. Steinbach, *J. Mus. Theory* **42**, 241 (1998).
13. J. Straus, *Mus. Theory Spectrum* **25**, 305 (2003).
14. F. Attneave, R. Olson, *Am. Psychol.* **84**, 147 (1971).
15. R. Shepard, *Psychol. Rev.* **89**, 305 (1982).
16. See supporting material on *Science* Online.
17. I. Satake, *Proc. Natl. Acad. Sci. U.S.A.* **42**, 359 (1956).
18. W. Thurston, *The Geometry and Topology of Three-Manifolds* (www.msri.org/publications/books/gt3m).
19. ChordGeometries 1.1 (http://music.princeton.edu/~dmitri/ChordGeometries.html).
20. R. Cohn, *J. Mus. Theory* **35**, 1 (1991).
21. W. Sethares, *Tuning, Timbre, Spectrum, Scale* (Springer, New York, 2005).
22. R. Cohn, *19th Cent. Mus.* **22**, 213 (1999).
23. B. Boretz, *Perspect. New Mus.* **11**, 146 (1972).
24. C. Callender, I. Quinn, D. Tymoczko, paper presented at the John Clough Memorial Conference, University of Chicago, 9 July 2005.
25. Thanks to D. Biss, C. Callender, E. Camp, K. Easwaran, N. Elkies, P. Godfrey Smith, R. Hall, A. Healy, I. Quinn, N. Weiner, and M. Weisberg.

A High-Brightness Source of Narrowband, Identical-Photon Pairs

James K. Thompson,[1]* Jonathan Simon,[2] Huanqian Loh,[1] Vladan Vuletić[1]

We generated narrowband pairs of nearly identical photons at a rate of 5×10^4 pairs per second from a laser-cooled atomic ensemble inside an optical cavity. A two-photon interference experiment demonstrated that the photons could be made 90% indistinguishable, a key requirement for quantum information-processing protocols. Used as a conditional single-photon source, the system operated near the fundamental limits on recovery efficiency (57%), Fourier transform–limited bandwidth, and pair-generation-rate–limited suppression of two-photon events (factor of 33 below the Poisson limit). Each photon had a spectral width of 1.1 megahertz, ideal for interacting with atomic ensembles that form the basis of proposed quantum memories and logic.

The generation of photon pairs is useful for a broad range of applications, from the fundamental [exclusion of hidden-variable formulations of quantum mechanics (1)] to the more practical [quantum cryptography (2) and quantum computation (3)]. A key parameter determining the usefulness of a particular source is its brightness, i.e., how many photon pairs per second are generated into a particular electromagnetic mode and frequency bandwidth. Parametric down-converters based on nonlinear crystals are excellent sources of photon pairs, but they are comparatively dim because their photon bandwidths range up to hundreds of GHz. However, new applications are emerging that demand large pair-generation rates into the narrow bandwidths (5 MHz) suitable for strong interaction of the photons with atoms and molecules (2, 4–7).

We report the development of a source of photon pairs with spectral brightness near fundamental physical limitations and approximately three orders of magnitude greater than the best current devices based on nonlinear crystals (8). Unlike parametric downconverters, however, the atomic ensemble can additionally act as a quantum memory and store the second photon, allowing triggered (i.e., deterministic) generation of the second photon. Triggered delays of up to 20 μs have been demonstrated (9–15), and it is expected that optical lattices hold the potential to extend the lifetime of these quantum memories to seconds (9). Lastly, proposed applications in quantum information (2, 3) rely on joint measurements of single photons for which indistinguishability is crucial for high fidelity. We observe large degrees of indistinguishability in the time-resolved interference between the two generated photons (16–19).

[1]Department of Physics, MIT–Harvard Center for Ultracold Atoms, Research Laboratory of Electronics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA. [2]Department of Physics, MIT–Harvard Center for Ultracold Atoms, Harvard University, 17 Oxford Street, Cambridge, MA 02138, USA.

*To whom correspondence should be addressed. E-mail: jkthomps@mit.edu