Design, Harmony, and Voice Leading

Garry S. Sittler

University of Illinois at Urbana-Champaign
Department of Computer Science
Urbana, Illinois 61801
sittler@cs.uiuc.edu

Abstract

This paper describes an automated design system which harmonizes music. The goal is to take a given melody line and produce four part harmony according to procedures set forth by harmonization and voice leading techniques. Voice leading techniques provide some guidance about how to produce a musical arrangement. Many of these techniques, however, are represented by rules which specify certain situations which must be avoided. Thus, the principle challenge is how to determine what design choices should be made based on what cannot be done.

Introduction

The process of automated design involves choosing appropriate values for a given set of design parameters. This paper describes the application of design tools to the domain of music harmonization.

The objective of the system described is to take as input a musical melody and add to it three additional voices to produce four part harmony based on general voice leading procedures. Such techniques have evolved and have been followed in Western music from approximately 1650 to 1900 (Turek 1988).

After an overview of the domain, with an emphasis on how it relates to automated design, a solution for satisfying the various design constraints is presented and discussed.

Background

Voice leading is a specific approach to harmonization, and it is best suited for arrangement of chorale music. Four voices, or music lines, move independently, and ideally each forms its own melody. When they are brought together, however, the individual pitches combine to form chords, also known as triads. These chords should match and compliment the phrase structure of the overall song. For example, when the song ends, the listener should be able to tell from a conclusive sounding harmonization that the end has occurred and that nothing else will follow.

General voice leading rules have evolved with Western music over time. These rules come in several categories. Some, such as the available major, minor, diminished, and augmented triads, are dictated by the structure and the pitches of the twelve notes within an octave. Others are based on preferences of what sounds pleasant to the human ear. Many of these types of rules designate musical situations which are prohibited. Another category could be called heuristic knowledge. Several of these rules are based on what a composer might do, and they are also good methods for avoiding the prohibited situations. Each category is further explained and exemplified below.

Each of the seven notes in a scale has associated with it a triad which is formed with the two notes forming intervals of a third and a fifth above it. The possible triads for a major scale are as follows:

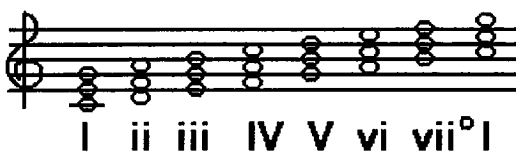

Thus, a specific note in a melody could be harmonized with three possible chords. Consider the note *E*. It could form the third of the *I* (tonic) chord, the root of the *iii* (mediant) chord, or the fifth of the *vi* (submediant) chord. One of the tasks of harmonization and of the system, therefore, is to decide for a given note which chord it will be harmonized with.

Voice leading techniques define several situations to avoid, even though these situations appear to be functionally correct with respect to the chord choices above. Certain note sequences between voices end up sounding "hollow" or displeasing. Since they do fit with a given harmonization, they are sometimes elusive and difficult to recognize, especially in written form. A typical example is allowing any two moving voices to form intervals of either a perfect fifth or an octave in consecutive beats, as illustrated by the following:

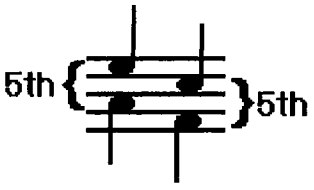

This situation, known as parallel fifths, should strictly be avoided.

Rules deemed heuristic are generally higher level because they incorporate several rules of the type described above. Heuristics generally provide good direction toward arranging a piece of music. One very useful heuristic to follow is to strive for contrary motion, especially between the outer (soprano and bass) voices, since they are the most easily audible. As mentioned earlier, each voice should strive to form its own melody, so a moving bass line helps fulfill this requirement. Furthermore, conforming to this practice avoids the parallel problem, since this can occur only when voices move in the same direction.

In addition to considering individual notes, it is necessary to take into account the overall musical phrases. A phrase is much like a sentence in language in that it expresses a musical thought. Some phrases, especially those in the middle of songs, act as questions by finishing with an open-ended sound. Others are more conclusive and serve to answer the questions. Each phrase should have an overall character and provide a feeling of forward motion leading toward its end. Certain chord sequences, called progressions, which provide this forward motion, are recognizable.

Especially important are the ends of phrases. Like sentence endings, phrases should be "punctuated." Short harmonization sequences called cadences have been identified which nicely mark phrase ends. For example, a sequence ending with the *V* (dominant) chord followed by the *I* (tonic) chord, provides a conclusive sound. Thus, this cadence, called an authentic cadence, usually appears at the end of a song.

Thus, the goal is to take a soprano voice, the melody, and produce three additional musical lines which both conform to the overall phrase structure and follow voice leading principles with respect to all of the lines.

Implementation

Based on rules and techniques governing harmonization and voice leading, I have identified what I feel to be the main subtasks and abstraction levels at which to view a given melody and the state of its harmonization progress. This a problem where in order to make efficient progress, some commitments must be chosen at appropriate places. So these levels essentially are defining a priority for making commitments. The abstraction levels considered, from top down, are as follows:

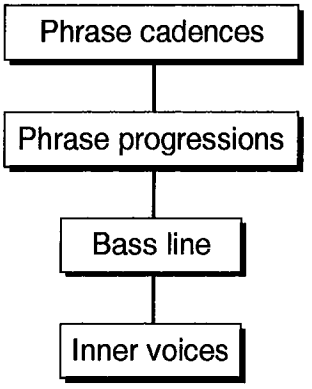

The top level decisions resolve the cadences and chords with which to end each phrase. From that point the overall phrase harmonizations are set forth by taking into consideration forward harmonic movement. The bass line is closely related to the chord choices, since after selecting a chord, only three notes are available. It is not, however, tightly constrained, since the bass can move rather freely in both steps and leaps within its range of almost two octaves. Also considered in selecting each bass note is its direction of motion with respect to both the previous bass note and the soprano melody. This is where contrary motion is highly preferred to avoid parallel fifths and octaves.

The inner voices make up the last level of detail, and these choices are the most constrained. After two notes are given, the main issue with the remaining lines is to avoid violating voice leading constraints. A common heuristic here is to retain common tones from the preceding chord, or if a tone is not shared between the chords, then move to the nearest note in the new chord.

The program implementation very closely matches the model outlined above. It fills in various parts of the complete song in a top down, high level (phrases) to low level (individual notes) fashion.

Input to the program consists of the specific pitch for each note, its duration, and a marker indicating whether the note ends a phrase. Once the program reads the input melody, the possible ways to harmonize each note are identified. Next, the structure of the overall melody is examined, phrase ends are identified, and from the possibilities, harmonizations are chosen for the ends of phrases. Specifically, each is ended with some appropriate cadence. The following step is to determine harmonizations for the beginning and middle parts of each phrase. This is actually done *backwards* and then forward. Starting from the phrase end, the program looks for a possible harmonic progression, or part of one, leading to the phrase cadence. It is important to understand that just choosing which chords with which to harmonize has not actually added any notes yet.

Having determined a harmonization for each melody note, the bass line is created. The primary goal is to move

the bass contrary to the melody. That is, if the melody goes up, the bass should go down, and vice versa. Of course, this can not always be accomplished because the bass may attempt to extend above its range. When it is necessary to move the two voices in similar motion, care is taken to avoid violating voice leading constraints, such as parallel fifths. Generally it is preferable for the bass note to be the root of the chosen harmonization, but chords also may appear in first inversion, in which the bass is the third of the chord.

The lowest level task is to fill in the alto and tenor voices. The heuristic of retaining common tones and moving to the nearest available note in the next chord is closely followed.

Results

To test the program I created by hand several short two-phrase melodies, each of which was intended to have one or more "appropriate" cadences between the phrases. The idea was to see whether the program recognized these and chose an appropriate harmonization.

Music which the program generates has both positive and negative characteristics. Problems appear in the lowest of the four abstraction levels shown earlier. It is not perfect with respect to all voice leading rules. The primary violations are those of note doubling and note omissions. For each harmonization, three notes are needed to form a triad, and we have four voices. Thus, some note is doubled. An interesting feature of the domain is that one of the three notes may sometimes be omitted, and two notes may be doubled, or one may even be tripled. While this is sometimes suitable, the program tends to abuse the practice, so additional care should be taken in choosing values for the inner voices. Added knowledge in the program's sections which create the inner voices could improve these results.

The operation of the program at the higher design levels appears to be highly satisfactory. One of the most favorable aspects is the production of the bass line. Contrary bass motion with respect to the melody is highly desirable, and the system accomplishes this very well. In fact, writing counterpoint such as this its own area of music study. Counterpoint music consists of two or more melodies which demonstrate independence from each other through contrary motion yet combine to form harmony (Turek 1988). The inner voices could be omitted to yield a program which produces a simple counterpoint melody based on the given melody.

For evaluation, a music instructor examined the results. When playing the music on a keyboard, the musician immediately noticed a crucial missing note in the final chord of one of the pieces. Upon further inspection the expert noticed a few other missing chord tones, but was mainly bothered by the lack of perfect resolution in the final chord. This illustrates what was mentioned earlier about sometimes being able to assign unusual note doublings. Essentially, this is a constraint which can be relaxed at times, but not always.

Another evaluation technique the musician performed was to play only the top and bottom lines of the songs. He was very satisfied with these. In fact he played them several times, listening closer each time.

Thus, the strengths and drawbacks of the current system are quite apparent. Some additional domain knowledge about voice doubling when choosing the inner voices could remedy the difficulties pointed out.

Conclusions

Music harmonization presents a suitable problem for application of automated design techniques. The domain contains typical features of design problems, such as design parameters (note and chord values), interactions between objects (notes and chords), constraints, and natural abstraction levels. In addition, central questions concern when the best time is to make commitments deciding the values for the design parameters and how to choose the values based on various rules and heuristics available.

A primary realization after implementing this system is the difficulty of identifying, representing, and succinctly describing the necessary design techniques. To get ideas of how to carry out the various subtasks, I relied mostly on my own music experience. I took some melodies and harmonized them by hand to see how I did it. But the problem is not always knowing how I arrived at a particular decision for a note of set of notes. It is not always clear whether it was due to application of some rule, to repeating a situation I had seen somewhere in the past, or to guessing. Thus, translating this knowledge into succinct rules and techniques for an automated design system was by far the most difficult part.

The main intent of the current system was to identify and investigate the issues involved in automating harmonization. Several extensions and enhancements are applicable. More music knowledge could be added to allow use of additional chords, especially seventh chords. Music also relies on non-chord tones and passing tones, which are not part of the prevalent chord, to add interest. An interesting problem would be identifying such notes in the given melody, as well as deciding when to add them.

References

Brown, D. C., and Chandrasekaran, B. "Knowledge and Control for a Mechanical Design Expert System." *IEEE,* July, 1986.

Sacerdoti, E. D. "Planning in a Hierarchy of Abstraction Spaces." *Artificial Intelligence* 5(2).

Sanjay, M., Dym, C. L., and Morjaria, M. "PRIDE: An Expert System for the Design of Paper Handling Systems." *IEEE,* July, 1986.

Stefik, M. "Planning with Constraints (MOLGEN: Part 1)." *Artificial Intelligence* 16(2), 1980.

Turek, R. *The Elements of Music: Concepts and Applications*. New York: McGraw-Hill, 1988.