

MaxAnnealing: A Tool for Algorithmic Composition Based on Simulated Annealing

Fernando Iazzetta

Center for New Music and Audio Technologies
University of California at Berkeley
1750 Arch St. - Berkeley, CA - 94709

also from

PUC-SP - Communication and Semiotic Program
fernando@CNMAT.berkeley.edu

Fabio Kon

Department of Computer Science
Institute of Mathematics and Statistics
University of São Paulo
Cx. Postal 66281 - 05389-970
05389-970 - São Paulo (SP) Brazil
kon@ime.usp.br

Abstract

Musical composition can be roughly viewed as a search for the best solution among a finite - although huge - universe of possibilities. Some of the algorithmic compositional techniques try to simulate the act of composing doing this search automatically. However, this approach has two major problems. The first one is the hardness of depicting aesthetic concepts through mathematical rules. The second problem is the low efficiency of the exhaustive search among all possible solutions.

The "Simulated Annealing" algorithm - first proposed in [1] - presents very good results on finding the optimal solution for many combinatorial problems efficiently (in polynomial time). In this paper we present an adaptation of this algorithm to the problem of algorithmic composition. We then discuss some possibilities regarding goal functions to this algorithm and describe MaxAnnealing, a tool designed to help composers and musicologists study the possibility of defining aesthetic concepts through mathematical rules. The system is implemented in the MAX programming environment.

1 Introduction

Musical composition involves many interrelated parameters which interact with each other creating a complex structure of musical possibilities and constraints. As Minsk has said "the problem of creating a good piece of music is a problem of finding a structure that satisfies a lot of different constraints" [2]. To deal with this complexity the composer uses different strategies, some of which are very well defined and can easily be formalized while others remain so flexible and context dependent that resist to any kind of formalization. These strategies include intuition, chance, adaptation, and trial and error based choices. However when the compositional task is transferred to a computer, these strategies are not always available since they are hard to implement as a computer program.

Usually, the compositional task consists of defining a musical goal to be reached by a compositional project. In most cases, not only the compositional project is determined by the goal, but also the latter is sensitive to the constraints imposed by the former. Thus, music arises from the balance of the composer's intentions (goal) and the compelling processes of organizing musical ideas (project). In other words, composition is a constant exercise of adaptation and interaction between project and goal.

However, in the case of computer generated compositions, sometimes it is hard to come up with an efficient project which can be translated into a compositional algorithm. Although the composer can have a clear idea of his musical intentions, he may be incapable of formalizing these ideas in order to build a program. Actually, for certain kinds of musical problems it is hard, or even impossible, to delineate an algorithm to solve them using the contemporary Computer Science tools. Usually, it can be due to the diversity of elements involved in the problem, to the complexity of parameters that affect the problem, or even to the lack of knowledge about some aspects of the system.

In his piece *Protocol* for solo piano, Charles Ames [3] proposes a compositional method that goes beyond the traditional processes strictly based on random selection or rigid determinism used in previous computer music programs. Ames formalizes his ideas in a "protocol" which is a "collection of tests where each test has been ranked according to one's preferences [...] Then by having the computer evaluate a substantial repertory of alternatives, one can direct it to search for the alternative best fitting these criteria. If there is a choice passing all tests, then systematic evaluation must find it; otherwise, the search will provide the best of imperfect choices" [3, p. 215]. The caveat of this method is that the computer must evaluate all in an enormous range of possible musical configurations to find the best one.

Another problem involving an extensive search of compositional solutions is presented by Schottstaedt in his "Automata Counterpoint" [4]. Schottstaedt implements a program that generates five species of counterpoint based on the rules exposed by J.J. Fux in the *Gradus ad Parnassum* (1725). Although the rules which govern the counterpoint are very clear and well defined, there are many different solutions for the same counterpoint problem and the computational time to check each possible solution becomes a significant constraint. "If we make an exhaustive search of every possible branch of a short (10-note) first-species problem, we have 16 raised to the 10th power possible solutions (there are 16 ways to move from the current note to the next note). Even if we could check each branch in a nanosecond, an exhaustive search in this extremely simple case would take 1,000 sec (about 20 minutes)" [4, p. 203]. The solution Schottstaedt presented to this case was to start with the best first result found for each interval, that is, the one to which the program assigns the smallest penalty. But as Schottstaedt recognizes, "the first such solution may not be very good. By accepting the smallest local penalty we risk falling into a bad overall pathway" [4, p. 203].

Some computational tools have been applied in music in order to solve this kind of problem where it is necessary to find the best configuration in a large space of possibilities. Examples are the back-propagation training algorithms used to weight the connections in a neural network, or the genetic algorithms which, inspired in the selective processes that occur in a natural environment, apply successive transformations to a system that evolves toward its environmental fitness. Those tools are based on the search of the best solution of a problem through iterative processes. In these processes, a possible solution is generated and evaluated by a particular function. Then, the program generates a new possible answer which again is evaluated. The results of these successive "guesses" are compared in order to direct the search for the best result. After doing a number of these iterations the system can reach a solution which is sufficiently close to the optimal.

2 The Simulated Annealing Algorithm

Simulated annealing (SA) is a probabilistic algorithm used in the search for optimum solution first described in [1]. It has been developed after the analogy with a Condensed Matter Physics thermal process named annealing and can be used to find near optimal configurations in very large and complex systems.

The annealing process consists on heating a piece of metal until it reaches a temperature slightly above its critical homogenization temperature and then carefully decreasing the temperature until the molecules are arranged in a way so that the metal reaches its thermodynamic equilibrium. The thermodynamic equilibrium is the state in which the molecules form a structure strictly organized and the energy of the system is minimal. If the heating and cooling processes are not correctly done, the metal does not reach the thermodynamic equilibrium.

As a combinatorial optimization process, the purpose of the SA is to find the minimum or maximum values of a cost function for a specific system. The cost function or goal function measures how good a specific configuration is. The SA starts with an initial structure S which can be randomly generated

and has a method for modifying this structure generating a neighbor structure. For each step in this process, the quality of the new structure is determined and, if the neighbor solution S' is better than the current solution, then S' becomes the current solution. But if the new solution does not represent an improvement in the goal function, it still can be accepted with probability

$$e^{\frac{Quality(S') - Quality(S)}{T}}$$

This condition, known as Metropolis criterion, helps the SA algorithm to escape from local minima. Unlike traditional local search algorithms, SA can make occasional moves in the search space which can decrease the value given by the goal function $Quality()$. The probability of acceptance of the new structure is greater if the difference between the cost of S and S' is small, even if it would represent a decrease in quality. Also, the probability of acceptance of a structure that decreases quality gets smaller as the temperature T decreases providing that the algorithm gets stabilized under a certain temperature.

It is possible to demonstrate that if the goal function and temperature lowering functions meet some constraints, the SA ends its running in polynomial time and finds the optimal solution with probability almost one [6].

3 Applying SA to Musical Composition

Figure 1 presents a version of the algorithm applied to the musical composition problem.

```

Procedure Simulated Annealing
Begin
  S ← random initial song
  T ← 1 /* initial temperature */
  L ← L0
  While (Quality is increasing) do
    Repeat L times
      S' ← Neighbor(S)
      If Quality(S') > Quality(S) Then S ← S'
      Else S ← S' with probability  $e^{\frac{Quality(S') - Quality(S)}{T}}$ 
    T ← T × 0.9 /* decreases the temperature */
    L ← reduce(L)
  X ← S, the final song given by the algorithm
End.

```

Figure 1: Simulated Annealing Algorithm

3.1 Basic Data Structure

Each configuration of the solution space is called a "song" and its structure is defined by the following C language declarations.

```

typedef struct { char pitch, velocity, start; } note;
typedef note song[MaxVoices][MaxTimeUnit];

```

Thus, we see a song as an array of `MaxVoices` rows and `MaxTimeUnit` columns where each entry is a triple (pitch, velocity, start). Pitch may contain either a MIDI pitch value (from 0 to 127) or a rest (represented by -1). Velocity contains a MIDI velocity value (from 0 to 127). Start may contain either 0 or 1 representing that there is a note or rest starting at that time unit (1) or not (0).

3.2 Neighbor Function

The job of the neighbor function in the SA algorithm is to receive a configuration in the solution space and to return another configuration in the same space being slightly different from the first one.

Our implementation of this function returns a song identical to the one it receives except for one note which is randomly added.

3.3 Goal Function

The purpose of our work is to develop an algorithmic tool to be used by composers and musicologists. Our idea is to offer a simple way for composers or musicologists with programming experience write their own goal functions in C, link them to our system, and then use the resource created as a MAX external object. Any goal function which receives a song as described above and returns an evaluation for its quality can be linked to our basic system.

For those with no programming experience, we have written a goal function which can have some of its parameters defined through the cells and tables of a friendly MAX patch. In the next section we describe the implementation of this goal function which can be taken as an example by other people interested on using our system with their own goal functions.

4 Implementation of MaxAnnealing

We have created a program called MaxAnnealing which uses the SA algorithm in order to find the optimal solution for a compositional problem. The program was implemented in the OPCODE's MAX 2.5.2 environment [5] since MAX offers a series of handful tools for manipulating all the music and graphic data needed by the program. An external MAX object called `annealing` which performs the optimization was created using the ThinkC 5.0.3 compiler.

The program has three basic modules: 1) the parameters interface, which allows the user to set the parameters that will be used for the song evaluations; 2) the simulated annealing object, which performs the search for the optimal solution; 3) the player, which receives the song produced by the annealing algorithm and plays it through a MIDI output.

The current implementation of MaxAnnealing generates a song of sixteen bars, each bar consisting of four beats, each of which having up to four subdivisions, which we call "time units". The piece is distributed by four different tracks or voices which can be assigned to four different MIDI channels. All these values can be changed through modifying the annealing object source code.

4.1 The Parameters Interface

The program starts with the user providing three general parameters for the piece. The first one is a tension curve which determines the pitch tension at each time unit of the piece. The pitch tension measures how dissonant are the simultaneous intervals that occur at each point of the music. This parameter is set by drawing a curve where each point represents the tension for each time unit in the piece as shown in Figure 2.

A second curve determines the density parameter. In this case, each point in the curve corresponds to the number of notes that should sound at each beat (group of four time units). The range can vary from 0 (no sound at all) to 16 (4 time units × 4 voices).

Finally, the user assigns weights to every pitch class of the chromatic scale. Here, a weight 0 means that the note should never occur. This brings a generic character in the pitch domain since the user can determine which scales he wish to use and create a pitch hierarchy by assigning high weights to certain pitches and low weights to others.

4.2 The Simulated Annealing Object

All data set by the user is then given to the simulated annealing module which starts the optimization process. It begins with a piece of music composed by just one arbitrary four-note chord where the voices are distributed from the highest to the lowest note through the tracks one to four (any other initial song

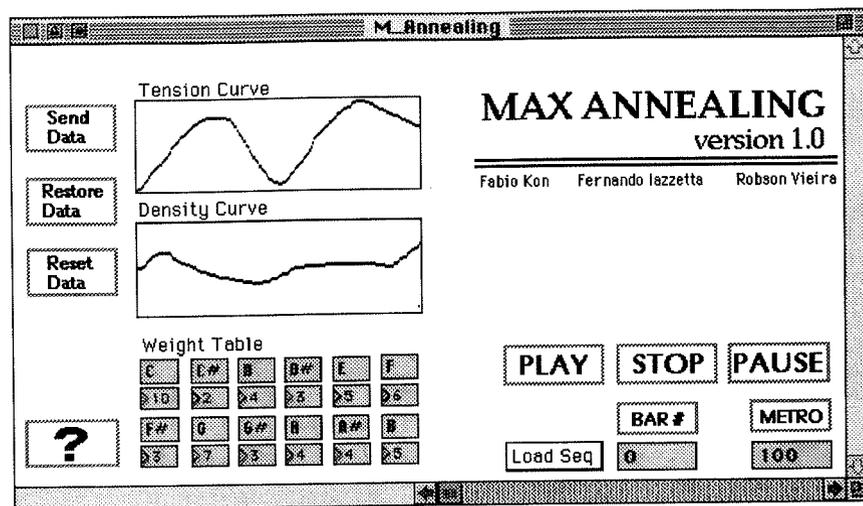


Figure 2: The MaxAnnealing Patch

would fit here). This is the first song evaluated by the algorithm. Given a song, its quality is evaluated by the goal function which is the weighted average of the five following criteria.

1. For each time unit, it calculates the tension among the intervals generated according to a previously established table of dissonance. To calculate the tension, MIDI notes are transformed in pitch classes and then the algorithm verifies all intervals that occur among the four voices for each time unit. These intervals are translated into previously established values which represent a dissonance level. The pitch tension for each time unit is calculated by adding these values.
The closer is the pitch tension to the tension determined by the user for that time unit, the higher is its evaluation. The final evaluation for the whole song is the average of its time unit evaluations.
2. For each bar, it calculates how many notes would sound and compares the result with the values established by the density curve giving a higher evaluation for the bar if the result gets closer to the value established by the user.
3. For each of the four voices, it evaluates the leaps between each note and its antecedent. Leaps close to a minor third receive higher evaluation than bigger and smaller ones. This will assure that the melodic contour of each voice will not have too many large leaps nor too many small intervals.
4. The more crossings between voices a song has, the lower is its evaluation. By doing so we try to avoid too many voice crossings.
5. Lower pitches receive higher evaluation if they have longer durations than higher pitches. This would lead to a music structure in which lower pitches will be associated to long durations and higher pitches will be associated to short durations.

As we have seen in section 2 and 3 the value given by the goal function is compared with the quality of a previous song. Then the program decides if it takes the new song as the new temporary solution, or if it keeps the last value.

4.3 The Player

This is a simple module which receives the best result obtained by the simulated annealing and translates this data so it can be played as a four-channel MIDI sequence. It presents some standard control buttons including play, pause, stop, metro, and save and load sequence.

5 Results

We have run MaxAnnealing for several times, using different parameters to test its performance. Initially, we have set the program to search through an average of 9,200 pieces of music for each new parameter configuration. This process took about 30 seconds in a shared SUN SPARCserver 1000 and one minute in a Macintosh Performa 630. Further tests have shown that, in some cases, MaxAnnealing was able to find very good solutions after a search through only 2,000 songs, which lowers considerably the necessary time to run the program. It is worth noting that even searching through 9,200 different songs, the algorithm runs almost 2,000 times faster than that which searches through all the 2^{24} possible solutions for our particular problem.

For each test we have set different parameters in order to generate specific kinds of musical output. By assigning certain weights to the notes in the pitch table we were able to generate pieces of music based on different modes and scales. The curves also provided an easy method of control over the tension and density of events which happened at each point of the song. After only a few experiments with these controls we could make satisfactory predictions about the general characteristics of the music MaxAnnealing was going to generate as the best solution.

Although our intention was only to demonstrate the validity of using the SA technique in the solution of musical problems, and despite the simplicity of the compositional rules applied, MaxAnnealing has produced some interesting musical results.

6 Conclusion and Further Work

We have introduced the use of simulated annealing algorithm as a powerful tool in the fields of musical composition and musicology. The Simulated Annealing has shown to be very effective in the search for a satisfactory solution for problems involving a large number of possible musical configurations in a very reduced time span.

With modifications in its objective function, the system would be very useful for finding the solution of many other compositional problems. Moreover, one can conceive the utilization of the SA as a practical tool in the field of musicology, which would enable the verification of relations between the formal rules which govern a piece of music and the actual effects - in terms of auditory experience - those rules generate.

As a further step on this work, we intend to test the use of other objective functions and develop the MaxAnnealing user interface to allow the generation of more complex compositional systems, which, we believe, will lead to more interesting musical results. These developments include new configuration options to be set by the user and the introduction of more elaborated compositional constraints in the program functions.

The MaxAnnealing source code and binaries are available by anonymous FTP at ftp.ime.usp.br, directory pub/macintosh/MaxAnnealing.

7 Acknowledgments

The authors gratefully acknowledge the help provided by Robson Feichas Vieira (IME/USP) throughout the development of the computational system.

This work was supported by CNPq (process # 200124/94-3), FAPESP (process # 93/0603-1), and CNMAT.

References

- [1] Kirkpatrick S., C.D. Gelatt Jr. and M.P. Vecchi. "Optimization by simulated annealing", *Science*, 220, pp. 671-680, 1983.
- [2] Roads, Curtis, "Interview with Marvin Minsk", *Computer Music Journal*, 4 (3), 1980.
- [3] Ames, Charles, "Protocol: Motivation, Design, and Production of a Composition for Solo Piano", *Interface*, 11, pp. 213-238, 1982.
- [4] Schottstaedt, W. "Automatic Counterpoint", *Currents Directions in Computer Music Research*, Max Mathews and John Pierce (Eds.). Cambridge: The MIT Press, 1989.
- [5] Puckett, M., D. Zicarelli, *MAX - An Interactive Graphic Programming Environment*, Opcode Systems, Menlo Park, CA, 1990.
- [6] Mitra, D., Romeo, F., Sangiovanni-Vicentinelli, Alberto, "Convergence and Finite-Time Behavior of Simulated Annealing", *Advanced Applied Probability*, 18, pp. 747-771, 1986.

AREM2: a composition tool

Francisco Kröpfl, Miguel Calzón
 Laboratorio de Investigación y Producción Musical (LIPM)
 Junín 1930 - (1113) Buenos Aires, Argentina
 fk@morfo.filo.uba.ar
 mcalzon@cnea.edu.ar

Abstract

In recent years an interactive software named AREM the goal of which is to produce reinterpretation in real time of sequences input through a MIDI interface was developed. This software is based on a specific composition methodology by Francisco Kröpfl implemented by Miguel Calzón. AREM2, a non interactive program, takes the same starting point and it is oriented towards the organization of musical structures including pitch, rhythm, syntactic units and texture. Both programs are complementary and may be used simultaneously.

1. Introduction

This paper introduces AREM 2, a program that creates musical structures which are reinterpreted successively with regard to their qualities of pitch, rhythm, texture and syntax. It is based on a composition methodology designed and taught by Francisco Kröpfl (described in Kröpfl 1987).

This paper will cover three areas: firstly the relationship between AREM1 and AREM2, secondly the way in which AREM2 produces musical sequences, and thirdly how both programs are applied in different fields of music.

2. AREM 2 and its precursor AREM 1

AREM - *Algoritmos para la Reinterpretación de Estructuras Musicales* - or **ARMS** - *Algorithms for the Reinterpretation of Musical Structures* (described in Calzón 1992 and 1993) is an interactive program which allows the transformation of the structure of the musical input by a performer or composer. From now on we will refer to this program as **AREM 1**.

AREM 2, by contrast, is **not** interactive. It consists of a group of patches that produce musical sequences following patterns derived from the subjacent methodology of composition. Each patch generates information on pitch, intensity or duration of notes, while the articulation