
Optimally directed shape generation by shape annealing

J Cagan

Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA

W J Mitchell†

Graduate School of Design, Harvard University, Cambridge, MA 02138, USA

Received 1 September 1991; in revised form 24 February 1992

Abstract. The new design technique of *shape annealing* is introduced. Shape annealing is a variation of the simulated annealing stochastic optimization technique. It produces optimally directed shapes within the language specified by a shape grammar. It does so by controlling the selection and application of shape rules such that superior shapes evolve as the algorithm progresses.

1 Introduction

Shape grammars, as defined by Stiny (1980), have successfully been employed to generate designs for many different types of artifacts—villas in the style of Palladio (Stiny and Mitchell, 1978), Mughul gardens (Stiny and Mitchell, 1980), prairie houses in the style of Frank Lloyd Wright (Koning and Eizenberg, 1981), Greek meander patterns (Knight, 1986), and suburban Queen Anne houses (Flemming, 1987), to name just a few. A shape grammar derives designs in the language which it specifies by recursive application of shape transformation rules to some starting shape. In general, at any step in the derivation of a design, there is a choice to be made among different possibilities for rule application. A particular design in the language results from a particular sequence of choices: it is derived by taking a distinctive path through the state-action tree of the grammar, from the starting shape to one of the terminal nodes.

The literature of shape grammars and their applications contains little discussion of *how* these choices among alternative rule applications might appropriately be made. Indeed, the issue is usually avoided, and grammars are routinely demonstrated by exhaustively enumerating possible derivations or by randomly sampling them. Here we propose a different approach—specifically, the use of optimization criteria to control the derivation of shapes. We formulate the design problem as an optimization problem with a quantifiable objective function and a set of design constraints. We then utilize a variation of the stochastic optimization technique known as simulated annealing, which we call *shape annealing*, to control choice among alternative rule applications as shapes are derived. The result is an *optimally directed* design solution.

Cagan and Agogino (1991) define optimally directed design as an approach to design which attempts to determine optimal regions of the design space by directing search toward improving the objectives and eliminating suboptimal or dominated regions. The result is not necessarily a globally optimal numerical solution, but rather insight into how the design might be improved relative to the objectives, and a superior—that is, optimally directed—design solution.

Simulated annealing procedures for production of such optimally directed designs generate random perturbations of a given solution state. They absolutely

† Present address: School of Architecture and Planning, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

accept the new solution if it is better, or accept it with a certain probability if it is worse. Past applications of simulated annealing have been to determine the optimal solution of continuous or ordered discrete design variables within a *fixed* design configuration: the form of the solution was known but specific values needed to be determined. Our shape annealing procedure, by contrast, uses a variation of simulated annealing not to determine the values of variables in a *given* design configuration, but rather to *generate* the configuration itself. Application of this procedure results in evolution of designs: the fittest (most optimal) configurations survive. It requires one or more design criteria (objective functions) with a set of design constraints, a starting shape, and a set of shape transformation rules.

In the next section we summarize the principles of simulated annealing. Our shape annealing procedure is then described. Finally, shape annealing is demonstrated with a simple shape grammar which can generate a countably infinite set of different shapes.

2 Simulated annealing

Simulated annealing is a stochastic optimization technique which has been demonstrated to solve continuous (for example, Cagan and Kurfess, 1991; Jain et al, 1990; van Laarhoven and Aarts, 1987) or ordered discrete (for example, Jain and Agogino, 1990; Kirkpatrick et al, 1983) optimization problems of fixed structure. Kirkpatrick et al (1983) developed the simulating annealing algorithm based on Metropolis's Monte-Carlo technique (Metropolis et al, 1953). The idea is analogous to the annealing of metals. A cooling schedule is defined giving a temperature reduction over a certain number of iterations. Temperature is a gradient variable with no physical meaning; the variable is called *temperature* to maintain the analogy with metal annealing. At high temperatures selection of a solution point is quite random whereas at lower temperatures the solution is more stable; the metal annealing analogy is that at high temperatures the molecules are at a highly random state whereas at lower temperatures they reach a stable minimum energy state.

With simulated annealing, a feasible state, s_1 , is randomly selected and the energy at that state, E_{s_1} , is evaluated. A different feasible state, s_2 , is then selected and evaluated to E_{s_2} . For objective minimization, if $E_{s_2} < E_{s_1}$, then s_2 becomes the new solution state. If $E_{s_2} \geq E_{s_1}$, then there is a probability, based on the temperature, that the new state will be accepted anyhow. Acceptance is determined by the probability calculation:

$$\Pr\{E_{s_2}\} = \exp\left(-\frac{E_{s_2}}{T}\right) / Z(T),$$

where $Z(T)$ is a normalization factor. A random number, r , between 0 and 1 is generated and compared with $\Pr\{E_{s_2}\}$. If $r < \Pr\{E_{s_2}\}$, then the new state is accepted anyhow; otherwise the old state is retained. The temperature is reduced and the process continues until convergence is reached or the temperature reaches 0.

Generally, the size of the mutation space is also reduced and asymptotes to zero. In this case, it can be proven (Lundy and Meese, 1986) that for continuous problems, if *equilibrium* (that is, convergence) is reached at each temperature and if the temperature is reduced slowly enough, then the algorithm is guaranteed to determine the global optimum. Because sufficient time cannot be guaranteed for large problems, simulated annealing is used to search only for a good solution and the exact globally optimal solution is often sacrificed for computation time.

Generally, the algorithm is run for several iterations at a given temperature until equilibrium is reached or until a certain number of iterations has occurred. The temperature is then reduced by a fixed amount and the algorithm is again run until

convergence of an iteration limit is achieved. When there is no accepted new solution at a given temperature the minimum has been found.

Simulated annealing has been used to evaluate a solution for a fixed design configuration. In the next section we extend the algorithm to *generate* the design configuration based on optimization criteria.

3 Shape annealing

In our shape annealing algorithm, simulated annealing is used to determine whether a randomly selected shape rule should be applied at a given design state. Given a current design state, an eligible rule is selected and applied to the design. If the new design does not violate any constraints, then it is sent to the Metropolis algorithm which compares the new state with the old state and, based on the temperature profile, it determines whether to accept it or not. If the rule violates a constraint, then the old state is maintained as the current. Note that a rule may not be eligible for application. The current algorithm actually selects any possible rule and then checks to verify that the rule is eligible.

By only applying additive rules, the design may rapidly converge on an inferior solution because it can work itself into a state where no rule can be applied without a constraint violation. Our algorithm can back itself out of these local solutions by reversing the previous rule. For every additive rule, there exists a subtractive opposite; if the subtractive rule is fired immediately following its companion additive rule, then the effect of the additive rule is removed.

The result of modeling a subtractive rule for every additive rule is an evolutionary design generation system. As the shape evolves, if it starts to converge on an inferior solution, the shape can backtrack and perturbate to a superior design configuration. This perturbation occurs from generating new design shapes and does not require maintaining a trail of all design decisions; the current state is all that is known about the design rather than all the mutations which had tried and failed during its evolution.

The shape annealing algorithm is given in figure 1. The number of outer loop iterations (that is, the determination of *reduction_factor*) and the number of inner loop iterations, *n*, need to be determined for the specific problem based on convergence of good solutions. Note that the size of the mutation space is not reduced at each iteration; this deviates from the normal simulated annealing approach but is required because there is no contour from which to move between neighboring configurations. In actuality, the size of the design space *increases* while the algorithm runs; however, the increase is controlled and the algorithm finds a way to generate an optimally directed design from a countably infinite number of possible configurations.

Because there is no contour from which to move between neighboring configurations (the design configurations are discrete) the simulated annealing algorithm actually acts as a sort of Monte-Carlo algorithm here. The difference from traditional Monte-Carlo approaches is that, by tweaking the parameters in the annealing schedule, the algorithm will generate shapes with different characteristics. Further, once the shape configuration has reached convergence the execution will stop, possibly preventing excessive computation time. It remains to be determined if, a priori, one can determine the appropriate parameters to produce superior solutions.

Shape grammars may be able to generate an infinite number of potential solutions. Given design constraints the number may be constrained to a finite but huge number. It may be impossible for a human or computer to generate all possible shapes in a reasonable period of time. With the shape annealing algorithm, a good solution evolves in polynomial time. It is not guaranteed to be the global minimum;

however, by letting the algorithm run a sufficient time, our experience shows the solution to be quite good. An example is presented in the next section.

```

Begin SHAPE ANNEAL
  T = 1.
  Define initial state;
  Evaluate state;
  While T > 0 Do Begin
    success = 0;
    For mutations = 1 to n Do Begin
      /* at each temperature mutate n times
      until convergence or limit reached*/

      Let temp_state = state;
      Generate rule;
      If rule is applicable
      Then Begin
        Apply rule to temp_state;
        If verify constraints of temp_state
        Then Begin
          Evaluate temp_state;
          Test temp_state with Metropolis;
          If acceptable
          Then Begin
            state = temp_state;
            success = success+1;
          End
        End
      End
    End
    If success > limit Then break;
  End
  If success = 0 Then break;
  /* no improvement...solution found*/
  T = T * reduction_factor;
End
End

```

Figure 1. Shape annealing algorithm.

4 Example

As an example application of the shape annealing algorithm, consider the simple shape grammar of half-hexagons shown in figure 2. Given this language, a countably infinite number of shapes can be generated. As designers, we desire to place as many half-hexagons as possible into a given space, without overlap, while satisfying the grammar of figure 2. For purposes of illustration, we consider all spaces to be of a constant area (25 units) and the half hexagon has a long base of one unit and a short base of one half unit. Rule 4 is given for completeness, but is actually only applied after the final solution is found.

We pose this problem as an optimization problem as:

maximize the number of pieces,

subject to:

the pieces are assembled based on shape grammar,
 no pieces can overlap,
 the shape must fit in bounds of defined space.

For this discussion, we limit the defined space to be rectangular and implement the grammar based on the center line of the half-hexagon (thus the actual rectangle will be slightly larger than 25 square units).

For implementation, the design state is stored by a trace of rule numbers; thus the detailed design is not stored but only a string of rules which create the design is stored. A new rule number is added by storing it at the end of the string. Implementation of the rule reversal occurs by randomly generating both positive and negative rule numbers. If a negative rule number is generated and the previously applied rule is the analogous positive number, then the positive rule is removed from the rule string.

Figures 3–5 show shapes generated with shape annealing for a square boundary space, and figure 6 for rectangular spaces. The fill pattern indicates which rule was applied: the light fill indicates rule 1, the medium fill indicates rule 2, and the dense fill indicates rule 3. Each figure starts with a piece in the bottom left corner.




Figure 2. Example shape grammar (from Mitchell, 1990, page 143).




Figure 3. 5×5 box with (a) 48 and (b) 45 half-hexagons.






Figure 4. (a) Transient solution with 31 half-hexagons evolving to (b) 5×5 box with 44 half-hexagons.



(a) (b)
Figure 5. (a) Transient solution with 16 half-hexagons evolving to (b) 5×5 box with 29 half-hexagons.



(a) (b) (c)
Figure 6. (a) 2×12.5 rectangle with 39 half-hexagons, (b) 10×2.5 rectangle with 33 half-hexagons, (c) 16.67×1.5 rectangle with 55 half-hexagons.

Figures 4 and 5 are of particular note. As discussed in section 3, as a shape is being generated it can work its way into a poor solution which, without rule reversal, has no further improvement. Because of the rule reversal, the shape can work its way out of the inferior state and then progress towards a better, optimally directed design. Figures 4(a) and 5(a) demonstrate solution states generated on the way to creating the final states of figures 4(b) and 5(b), respectively. Note that in figures 4(a) and 5(a) there are no other valid additive rules which can be applied from the end state; obviously these solutions are inferior designs. As the shape annealing algorithm progressed, the reversal rules became dominant and eventually the shape was able to work out of the 'corner' and found a better route with which to progress.

5 Discussion

In an examination of the shapes of figures 3–6, a few observations are relevant:

- (a) The final configuration is quite sensitive to the initial path generated. At the very beginning of the annealing process, any shapes generated will be better, based on the objective function, than the previous states. The shape quickly goes off in the initial directions generated and has no reason to backtrack to improve the beginning part of the shape. Thus, if a smart initial path is not chosen, the algorithm may not be able to densely pack portions of the space.
- (b) Initially there are a countably infinite number of shapes which can be generated; for 3 possible generative shapes at each state with a half-hexagon of 0.325 square units area and a 25 square units specified space, the worst case could be 3^{75} possible paths because up to 75 pieces could potentially fit in the space. Because there are constraints which prevent overlap and define bounds on the shape, and because the grammar is designed to prevent complete packing, this number is considerably reduced; however, the number of possible designs is combinatoric and quite large. Shape annealing is able to determine very good solutions in polynomial time.
- (c) The variance of the shape at a given temperature is sensitive to the number of possible states which remain feasible. Because there are so many available shapes at the start, there is a great variance in the shape after each initial temperature step. However, after much of the shape has been generated, it is harder to generate a shape which does not violate a constraint; thus the rate of change of the shape is smaller. The result is a shape which is more likely to be sparse near the start and dense near the end.
- (d) As mentioned in section 3, tweaking the annealing parameters can affect the characteristics of the solutions. By changing the probability of generating a rule reversal and producing a smoother temperature drop, the algorithm is able to generate solutions which spiral toward the center (and which have a better objective function). Future research will investigate the reason for the algorithm sensitivity; it may be possible to determine a priori which parameters generate better solutions.
- (e) The algorithm is computationally efficient. There is no need to maintain a trace of design decisions; only the current state and proposed mutation at each step need be stored. The algorithm itself will backtrack out of a configuration if required.
- (f) We emphasize that the designs generated by shape annealing are not guaranteed to be globally optimal. Rather, the resulting designs are generally very good solutions to a combinatoric problem. In the example problem of section 4, one could determine better solutions than those generated by the algorithm. However, we are presenting an algorithm which can control the generation of shapes *automatically*. Further, for more complicated shape grammars it is less likely that optimally directed patterns will be obvious to the human. Since we do not guarantee global optimality, the algorithm should be executed for multiple runs on each problem to obtain different solutions from which to choose.

- (g) The shapes generated by shape annealing *evolve* into their final configuration, simulating an analogy to the natural process of evolution. Although mutations may be pursued even if initially they do not improve the design, the survival of only the fittest designs remain.
- (h) The simple shape grammar used for our example does not involve the recognition and replacement of emergent subshapes. Extension of the approach to grammars that do involve emergent subshapes is an important direction for future research.

6 Conclusions

We have introduced a powerful method called *shape annealing* to control the generation of shapes with shape grammars based on optimally directed solutions to functional needs. The application of the shape annealing algorithm to a simple shape grammar has led to the evolution of a large variety of shapes which satisfy the given design requirements. Shape annealing introduces a new way of applying shape grammars to the design process.

Acknowledgements. The authors would like to thank John Corson-Rikert for providing the graphical interface to draw the figures in this paper. This work was partially sponsored by the Engineering Design Research Center at Carnegie Mellon University, an NSF-sponsored research center.

References

- Cagan J, Agogino A M, 1991, "Inducing constraint activity in innovative design" *AI EDAM: Artificial Intelligence in Engineering, Design, Analysis and Manufacturing* 5(1) 47-61
- Cagan J, Kurfess T R, 1992, "Optimal tolerance allocation over multiple manufacturing alternatives", paper presented at ASME Design Automation Conference, Scottsdale, AZ, 13-16 September
- Flemming U, 1987, "More than the sum of the parts: the grammar of Queen Anne houses" *Environment and Planning B: Planning and Design* 14 323-350
- Jain P, Agogino A M, 1990, "Theory of design: an optimization perspective" *Mechanism and Machine Theory* 25 287-303
- Jain P, Fenyves P, Richter R, 1990, "Optimal blank nesting using simulated annealing" *Proceedings of ASME Design Automation Conference: Advances in Design Automation-1988* Ed. B Ravani 2 109-116
- Kirkpatrick S, Gelatt C D Jr, Vecchi M P, 1983, "Optimization by simulated annealing" *Science* 220 671-679
- Knight T W, 1986, "Transformations of the meander motif on Greek geometric pottery" *Design Computing* 1 29-67
- Koning H, Eizenberg J, 1981, "The language of the prairie: Frank Lloyd Wright's prairie houses" *Environment and Planning B* 8 295-323
- Lundy M, Meese A, 1986, "Convergence of an annealing algorithm" *Mathematical Programming* 34 111-124
- Mitchell W J, 1990 *The Logic of Architecture* (MIT Press, Cambridge, MA)
- Metropolis N, Rosenbluth A, Rosenbluth M, Teller A, Teller E, 1953 *Journal of Chemical Physics* 21 1087-1091
- Stiny G, 1980, "Introduction to shape and shape grammars" *Environment and Planning B* 7 343-351
- Stiny G, Mitchell W J, 1978, "The Palladian grammar" *Environment and Planning B* 5 5-18
- Stiny G, Mitchell W J, 1980, "The grammar of paradise: on the generation of Mughul gardens" *Environment and Planning B* 7 209-226
- van Laarhoven P J M, Aarts E H L, 1987 *Simulated Annealing: Theory and Applications* (D Reidel, Dordrecht)