

# Review of **The Annealing Algorithm**

## by Otten and van Ginnekin

Daniel R. Greening  
University of California, Los Angeles

Simulated annealing is a technique used to find good approximate solutions for  $NP$ -hard problems. The algorithm begins with a feasible starting state  $s \in S$ , with cost  $\epsilon(s)$ , then generates a proposed state change  $s' \in S$  from the current state. It evaluates the new state's cost,  $\epsilon(s')$ , and decides either to set the current state to the new state, or to reject the proposed state. The state change is accepted with probability

$$p(\langle s, s' \rangle) = \begin{cases} 1 & \text{if } \epsilon(s') < \epsilon(s) \\ e^{(\epsilon(s) - \epsilon(s'))/t} & \text{otherwise} \end{cases},$$

where  $t$  is a control parameter.

The probabilities of generating new states from old form a Markov state transition matrix,  $G$ . Likewise, one can express state acceptance probabilities in another transition matrix  $A$ . If  $G$  is weakly ergodic ( $A$  is ergodic by definition), repeated transitions  $GA$  from any starting state will converge to a stationary Boltzmann distribution.

The Boltzmann distribution, originally discovered in thermodynamics, describes the energy distributions in a particle system held at constant temperature  $t$ . As a result, intuition derived from physics often applies to simulated annealing.

Simulated annealing has proved better than other techniques for solving several industrial design problems, such as circuit layout and network optimization. Unfortunately, as might be expected of any  $NP$ -hard problem solver, it is often slow—large circuit layout problems have been known to take days.

An army of researchers have set about attacking the execution time of simulated annealing: some proceed using trial and error, without a theoretic base; others develop elaborate theories, without reference to actual problems or resource constraints. The most successful embrace both practice and theory.

It was a pleasure to find that *The Annealing Algorithm* supplies its reader with a collection of interesting theoretical results, and shows how these results apply in practice. In Chapter 9 we find a proof that simulated annealing converges at a rate bounded proportionally by the second largest eigenvalue of the

transition matrix  $GA$ . In Chapter 10, we discover that the “smoothness” of the move space is related directly to that eigenvalue. Example implementations with “smooth” and “rough” move spaces for the traveling salesman problem are shown.

Similar useful facts appear throughout the monograph. In Chapter 5, we find that the Boltzmann distribution maximizes state-space exploration (entropy) for a fixed average cost—this provides an intuitive feel for why simulated annealing works. Chapter 8 shows us that if  $E_\infty \gg \sigma_\infty^2$ , where  $E_\infty$  and  $\sigma_\infty$  are the cost mean and standard deviation over all states, the starting temperature  $t$  need not be greater than  $\sigma_\infty^2$ —this makes choosing a starting temperature easy.

Chapter 7 shows us that an annealing run can be divided into two regions—a high temperature region  $T \geq t$ , where  $E_t \approx E_\infty - \sigma_\infty^2/t$ , and a low temperature region  $t \leq T$ , where  $E_t \approx E_\infty - (\frac{2}{T} - 2)\sigma_\infty^2/T$ —this can help us determine how close we are to equilibrium.

The authors refrain from extravagant claims about simulated annealing. They clearly describe the limitations of their work, and point out where they make mathematical approximations. This is a rigorous work.

*The Annealing Algorithm* has some flaws, perhaps endemic to monographs. In the preface, it claims to be “at the same time an introduction into annealing and its applications, a compendium for the theoretical background of annealing, a basis for further research, and a report on the progress made in developing a multi-purpose annealing routine.” All claims were correct, save one: *The Annealing Algorithm* is **not** an introduction to annealing. For that, a better text exists [1].

Many difficult mathematical preliminaries appear in Chapters 2–6: with few motivating comments, dragging myself through this dry presentation seemed like crawling through a vast desert. Finally, Chapters 7–10 bring welcome relief. Problems at the end of each chapter, particularly in Chapters 2–6, would have exercised and deepened my knowledge: none were included. Occasionally the authors invoke equations derived in preceding chapters, but fail to provide references.

The typeface, Computer Modern San Serif, was a poor choice: it brings confusion between  $I$ ,  $l$ , and 1. Formulaic variables were not italicized, causing more confusion. Symbols undergo metamorphosis as you read: at one moment  $\nu$  is the valley of two states, at another moment  $v$ ;  $\mu$  and  $\bar{x}$  are the same mean;  $\sigma$  and  $s$  are the same standard deviation;  $\Phi$  is at one point the normal distribution, at another the space conductance. This cryptography distracted and slowed me—it will trap other unwary readers.

Facts and techniques gleaned from *The Annealing Algorithm* will benefit my own work. I recommend it highly for colleagues versed in simulated annealing, particularly those who lean toward theory. The lack of motivation and notational confusion may irritate you, as it did me, but the wealth of information contained within these pages provides more than ample compensation for its difficult style.

## References

- [1] P.J.M. van Laarhoven and E.H.L. Aarts. *Simulated Annealing: Theory and Applications*. D. Reidel Publishing Company, Boston, 1987.