University of Wollongong

Research Online

Faculty of Engineering and Information
Sciences - Papers: Part A

Faculty of Engineering and Information
Sciences

1999

Simulated annealing: Searching for an optimal temperature schedule

Harry Cohn
University of Melbourne

Mark James Fielding
University of Wollongong, mfieldin@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/eispapers>

Part of the [Engineering Commons](#), and the [Science and Technology Studies Commons](#)

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Simulated annealing: Searching for an optimal temperature schedule

Abstract

A sizable part of the theoretical literature on simulated annealing deals with a property called convergence, which asserts that the simulated annealing chain is in the set of global minimum states of the objective function with probability tending to 1. However, in practice, the convergent algorithms are considered too slow, whereas a number of nonconvergent ones are usually preferred. We attempt a detailed analysis of various temperature schedules. Examples will be given of when it is both practically and theoretically justified to use boiling, fixed temperature, or even fast cooling schedules which have a small probability of reaching global minima. Applications to traveling salesman problems of various sizes are also given.

Keywords

annealing, optimal, temperature, simulated, schedule, searching

Disciplines

Engineering | Science and Technology Studies

Publication Details

Cohn, H. & Fielding, M. James. (1999). Simulated annealing: Searching for an optimal temperature schedule. *SIAM Journal On Optimization*, 9 (3), 779-802.

SIMULATED ANNEALING: SEARCHING FOR AN OPTIMAL TEMPERATURE SCHEDULE*

HARRY COHN[†] AND MARK FIELDING[†]

Abstract. A sizable part of the theoretical literature on simulated annealing deals with a property called convergence, which asserts that the simulated annealing chain is in the set of global minimum states of the objective function with probability tending to 1. However, in practice, the convergent algorithms are considered too slow, whereas a number of nonconvergent ones are usually preferred. We attempt a detailed analysis of various temperature schedules. Examples will be given of when it is both practically and theoretically justified to use boiling, fixed temperature, or even fast cooling schedules which have a small probability of reaching global minima. Applications to traveling salesman problems of various sizes are also given.

Key words. simulated annealing, temperature, cooling, Markov chain, convergence, inhomogeneous chain, fundamental matrix, time to absorption

AMS subject classifications. 60J05, 60K35

PII. S1052623497329683

1. Introduction and summary. Suppose that a function f is defined on a finite (but large) set of states S . The aim of simulated annealing (SA) is to find a state x such that $f(x) = \min_{y \in S} f(y)$. Because, for some large S , such an aim is not in general feasible in a reasonable time frame, we may confine ourselves to finding a near optimal x , i.e., a state x for which $f(x)$ is close to $\min_{y \in S} f(y)$.

For each state x in S , define a set $N(x)$, called the set of neighbors of x . Write $\mathcal{N}$ for the family of neighborhoods $\{N(x), x \in S\}$.

A neighbor choosing matrix $\mathbf{G}$ with entries $G(x, y)$ is defined for each x and y in S such that $G(x, y) > 0$ if and only if $y \in N(x)$. The matrix $\mathbf{G}$ is called a *generation* matrix.

Define

$$P_T(x, y) = \begin{cases} G(x, y) \exp(-[f(y) - f(x)]^+/T) & \text{if } y \neq x, \\ 1 - \sum_{z \neq x} P_T(x, z) & \text{if } y = x, \end{cases}$$

with $a^+ = \max(a, 0)$.

The parameter T is called *temperature*. Write $\mathbf{P}_n$ for the transition probability corresponding to $T = T_n$, where T_n is the temperature at time n . The sequence $\{T_n\}$ is called a *temperature schedule*. If $\lim_{n \rightarrow \infty} T_n = 0$ we say that $\{T_n\}$ is a *cooling schedule*; we say it is a *fixed temperature schedule* if $T_n = T$ for all n .

An initial probability distribution and the sequence of one-step transition probabilities $\{\mathbf{P}_n\}$ define an inhomogeneous Markov chain $\{X_n\}$. This chain will be called an *SA chain*. The SA chain is the basis of the SA algorithm. It originates from an idea that goes back to the paper by Metropolis et al. [24] and was followed up by many other contributors (see, e.g., Aarts and van Laarhoven [2], Aarts and Korst [3], Chiang and Chow [5], Connolly [9], Connors and Kumar [10], Gelfand and Mitter [11],

*Received by the editors October 30, 1997; accepted for publication (in revised form) July 1, 1998; published electronically June 30, 1999.

<http://www.siam.org/journals/siopt/9-3/32968.html>

[†]Department of Mathematics and Statistics, University of Melbourne, Parkville 3052, Victoria, Australia (harry@ms.unimelb.edu.au).

Geman and Geman [12], Hajek [15], Hwang and Sheu [16], Romeo and Sangiovanni-Vincentelli [28]). Recently, Niemi and Pokarowski [26] and Niemi [27] have clarified the asymptotic behavior of the SA chain by relating it to the theory of the tail events (see Cohn [6], [7], [8]).

Write $P^{(m,n)}(x, y) = P(X_n = y | X_m = x)$ for $m < n$. We shall say that y is *reachable from* x if there exist an integer p and states $x = x_0, x_1, x_2, \dots, x_p = y$ such that $x_{k+1} \in N(x_k)$ for $0 \leq k < p$. It is easy to see that if y is reachable from x then there must be a number p such that $P^{(m, m+p)}(x, y) > 0$ for any m .

We assume that $(S, \mathcal{N})$ is irreducible, i.e., that any state x is reachable from any state y .

We shall say that state y is *reachable at height* h from state x if h is the smallest number such that $x = y$ and $f(x) \leq h$ or if there is a sequence of states $x = x_0, x_1, \dots, x_p = y$ for some $p \geq 1$ such that $x_{k+1} \in N(x_k)$ for $0 \leq k < p$ and $f(x_k) \leq h$ for $0 \leq k \leq p$.

State x is said to be a *local minimum* if no state y with $f(y) < f(x)$ is reachable from x at height $f(x)$. The *depth* of x , $d(x)$, is defined to be ∞ if x is a global minimum; otherwise it is the smallest number h , $h > 0$, such that some y with $f(y) < f(x)$ can be reached from x at height $f(x) + h$.

We assume that y is reachable from x at height h if and only if x is reachable from y at height h . This assumption is called *weak reversibility*.

Write S^* for the global minimum set of states, i.e., the set of states x with $f(x) = \min_{y \in S} f(y)$. We say that the SA chain (or algorithm) is *convergent* if

$$(1.1) \quad \lim_{n \rightarrow \infty} P(X_n \in S^*) = 1.$$

Hajek [15] identified the smallest value of c for which an SA chain with cooling schedule of the form

$$T_n = \frac{c}{\log(n + n_0)}$$

is convergent. (Here n_0 is a positive integer.) It was proved in [15] that the SA algorithm is convergent if and only if $c \geq d^*$, where d^* is the largest depth of the local minima, which are not global minima.

Thus $T_n = d^*/\log(n + n_0)$ gives the fastest logarithmic-type cooling schedule leading to convergence. Such a cooling schedule is called *canonical*, and d^* is said to be the *canonical constant*. It is important to stress that it would be wrong to assume that a canonical cooling schedule necessarily reaches global minimum faster than other schedules.

A convergent chain obtains optimality in the long run even if we adopt a *memoryless algorithm*, i.e., an algorithm that does not recall the past values of the chain. An algorithm that stores the best solution of all iterations will be called a *memory algorithm*.

The aim of this paper is to study the behavior of the SA algorithm in terms of temperature schedules. It turns out that the key critical points for the limit behavior of the SA chain occur in the range of logarithmic cooling schedules. We shall describe a number of optimality criteria corresponding to various situations. Then we study some theoretical properties of algorithms that are used in practice. It turns out that there is no theoretical reason why some temperature schedules that are attached to nonconvergent SA chains should be overruled. Examples are given to illustrate each case.

2. Which optimality? An algorithm needs to specify a stopping rule (time), i.e., a time when the process is terminated and a decision is adopted. This stopping rule, denoted by τ , may or may not depend on the values taken by the chain up to the stopping time and may be random. It is also a function of the temperature schedule and other parameters of the algorithm as well as the optimality criterion adopted for the problem.

It is usually assumed that an optimal (near optimal) algorithm is one that

(i) with probability 1 reaches the global (near global) minimum in finite time, and

(ii) is faster than other algorithms.

Both desiderata need to be qualified. It will turn out that (i) may have a different meaning for memoryless algorithms than for those with memory. Besides, we shall further see that (i) may be done away with if we consider multiple run algorithms. As far as (ii) is concerned, there are a number of principles for optimality based on τ , the most common one being searching for the algorithm that attains the minimum of $E(\tau)$. However, such a criterion would exclude algorithms with $P(\tau = \infty) > 0$. The choice of the algorithm and the corresponding optimality criterion may be problem dependent. A discussion of various criteria of optimality follows.

2.1. Convergent algorithms. A convergent chain, defined by (1.1), ensures that the global minimum is eventually reached if a sufficiently large number of iterations is allowed. If convergence fails for a cooling schedule $\{T_n\}$, there is a positive probability that the SA chain will never reach a global minimum state (see Hajek [15]). Some authors seem to assume that convergence is a necessary property of a successful algorithm. As previously mentioned, the canonical schedule, despite being the fastest tending to 0, is not necessarily optimal, i.e., the fastest in reaching a global minimum.

In fact, convergence is not a necessary attribute of a successful algorithm either. It may be necessary in relation to a memoryless algorithm. For memory algorithms, there is no a priori reason for a cooling schedule with property (1.1) to be preferred to temperature schedules that do not satisfy (1.1).

By the same token, one should not a priori eliminate convergent algorithms from the search of optimal schedules.

2.2. Regular algorithms. For memory algorithms, (2.1) is replaced by the less restrictive requirement that S^* be reached with probability 1. In such a case, a criterion for optimality should depend only on how early S^* is reached.

Suppose that, for any given temperature schedule $\{T_n\}$, we define the stopping time τ to be the first n such that X_n hits S^* . An algorithm for which

$$(2.1) \quad P(\tau < \infty) = 1$$

is said to be *regular* and *defective* otherwise.

For memory algorithms, τ is the variable that should be optimized. The relevant sequence of random variables is $\{\min(f(X_1), \dots, f(X_n))\}$, and (2.1) is equivalent to

$$(2.2) \quad \lim_{n \rightarrow \infty} P(\min(f(X_1), \dots, f(X_n)) = \min_{x \in S} f(x)) = 1.$$

Obviously (2.2), or equivalently (2.1), is a convergence property, and it is easy to see that it holds for a much larger class of temperature schedules than the ones satisfying (1.1). For example, all the chains corresponding to fixed temperature schedules satisfy

it, as they are ergodic Markov chains with stationary transition probabilities. It is well known that for such chains all states are recurrent and (2.1) holds. On the other hand, if $\{\pi_x\}$ is the stationary probability distribution, then $\lim_{n \rightarrow \infty} P(X_n = x) = \pi_x > 0$ for all x and therefore

$$\lim_{n \rightarrow \infty} P(X_n \in S^*) = 1 - \sum_{x \notin S^*} \pi_x < 1,$$

so that (1.1) fails.

We shall see later an example where it may be optimal to boil, i.e., to let T_n tend to ∞ . This case corresponds to a Markov chain with stationary transition probabilities given by the generation matrix.

We shall show that there are problems where a fixed optimal temperature may be identified.

For regular algorithms, a criterion for τ^* to be optimal is

$$E(\tau^*) = \min_{\tau \in \mathcal{T}} E(\tau),$$

where $\mathcal{T}$ is the class of stopping times attached to all temperature schedules. This criterion is often used in operations research.

A number of papers have pointed out the potential usefulness of memory algorithms (see, e.g., Kirkpatrick [20], Gelfand and Mitter [11] for cooling schedules and Connolly [9] for a fixed temperature algorithm).

We shall study the properties of τ for fixed temperature schedules in a later section.

2.3. Defective algorithms. It may seem natural to consider optimality with respect to the class of all temperature schedules defining a regular algorithm. However, a closer examination does not justify such a criterion. We also need to consider temperature schedules that may correspond to defective algorithms, even if $P(\tau < \infty)$ is not even close to 1. In fact, such algorithms are the ones mostly used in practice.

Suppose that we want to allow for a fixed number of iterations N and choose the algorithm that performs the best within N iterations.

Consider the case when the numbers N and p are suitably chosen such that, for a stopping time τ corresponding to some temperature schedule,

$$P(\tau \leq N) \geq p.$$

Define $\mathcal{T}$ to be the class of all regular and defective τ , where τ is the first hitting time of S^* (or near optimal states). An optimality criterion for such a case will be satisfied by a stopping time τ^* in $\mathcal{T}$ such that

$$P(\tau^* \leq N) = \sup_{\tau \in \mathcal{T}} P(\tau \leq N).$$

In fact, we may achieve a property close to (2.2) in terms of some number, say k , of independent runs of size N . Indeed, if $\{X_n^{(i)}, n = 1, \dots, N\}$ is the i th run with $i = 1, \dots, k$, then

$$(2.3) \quad P\left(\min_{i \in \{1, \dots, k\}} \min \left(f\left(X_1^{(i)}\right), \dots, f\left(X_N^{(i)}\right)\right) = \min_{x \in S} f(x)\right) \geq 1 - (1 - p)^k.$$

By suitably choosing k such that the right-hand side of (2.3) is as large as desired, and adopting the stopping rule $\tau_N = \min(\tau, N)$, we may ensure both the quality of the algorithm and a limitation on the number of iterations.

2.4. Near optimality. If optimality requires too many iterations, near optimality may be a suitable alternative. The latter is in fact the case with most algorithms used in practice. In fact, many feasible algorithms will only provide a near optimal solution.

In some cases, getting a near minimum, say, within two percent of the global minimum, can be achieved with a drastic reduction in the number of iterations required for finding a global minimum state. An improvement from two percent to, say, one percent may result in a huge increase in the number of iterations, which is not always practical.

3. Fixed temperature schedules. If a simulated annealing chain is run with a fixed temperature, then, under minor conditions on the generation matrix, the “best state so far” will, with probability 1, ultimately become a global minimum, and the expected time and variance of the time until reaching a global minimum will be finite. This follows from the classical theory of finite Markov chains.

We will see that for some small and medium-size problems, the fixed temperature schedules seem to work better than the simulated algorithms based on a cooling schedule.

Define a Markov chain $\{X_n^*\}$ with one absorbing state representing the states of S^* lumped together. The states outside S^* , as well as the transition probabilities among themselves, remain unchanged. Clearly, the first time the chain $\{X_n^*\}$ reaches S^* is a stopping time, say, τ . Such a case is well known in the theory of Markov chains (see, e.g., Kemeny and Snell [18, Theorem 3.5.3]). Denote by Q the transition matrix corresponding to the states outside S^* . The matrix $N = (I - Q)^{-1}$ is called *fundamental*. The square matrix I , called identity matrix, has the diagonal entries equal to 1 and is 0 elsewhere. Let A be an arbitrary finite matrix. The matrix A_{sq} is formed from A by squaring all entries. We denote by ξ a column matrix having all components equal to 1.

The following result is extracted from Theorem 3.5.4 of [18].

THEOREM 3.1. (i) *The i th component of $N\xi$ is the mean number of steps needed to reach S^* given that the chain starts in i .*

(ii) *The i th component of $(2N - I)N\xi - (N\xi)_{sq}$ is the variance of the same function.*

3.1. An optimal boiling schedule example. In Hajek [15], a small problem instance, shown here as Figure 3.1, is given for SA consisting of 26 states. The chain is used by Hajek to illustrate convergent schedules. Ironically, it turns out that boiling to ∞ is the optimal temperature schedule.

Shown is the neighborhood structure as well as the cost associated with each state. The states have been numbered arbitrarily to give the state space $S = \{1, 2, \dots, 26\}$. The relationship $y \in N(x)$ is represented by an arrow from x to y . So the set of neighbors of state 9, for example, is $N(9) = \{8, 10, 13\}$, and for state 3, we get $N(3) = \{2\}$.

There are six local minima, states 1, 2, 10, 12, 17, and 26, and the set of global minima is $S^* = \{1, 2, 26\}$. It is easy to check that this chain is weakly reversible.

We shall assume that the generation matrix is given by $G(x, y) = 1/|N(x)|$ for all $x \in S$ and $y \in N(x)$.

We note that this example is only trivial in size. It does, however, allow us to examine the application of Markov chain theory to SA in a way not plausible for practical problems. That is the explicit examination of the transition matrix. It may also raise interesting questions about the behavior of SA in real-life problems.

FIG. 3.1. A 26-state example given in Hajek [15].

TABLE 3.1

Performance of fixed temperature schedules for a 26-state SA chain. Mean and standard deviation of time to hitting a global minimum are given.

T	$E(\tau)$	$SD(\tau)$
0	∞	∞
1	2964.04	2949.57
2	250.77	252.25
3	129.00	126.93
4	97.87	94.48
5	84.73	80.73
10	67.02	62.16
50	58.67	53.48
100	57.91	52.70
∞	57.20	51.97

TABLE 3.2

Performance of logarithmic cooling schedules for a 26-state SA chain. Estimates of mean and standard deviation of time until hitting a global minimum are given. The values at $c = \infty$ follow from the calculations with fixed temperature.

c	$E(\text{time})$	$SD(\text{time})$
6	∞	∞
8	1026.74	3464.46
10	259.78	499.17
50	64.45	61.70
80	60.39	56.06
100	60.44	55.74
200	58.72	53.84
500	57.81	53.58
1000	57.78	52.50
∞	57.20	51.97

To get from the local minimum state 12 to a state of lower cost, it is necessary to climb at least five units, so the depth of state 12 is equal to 5. Similarly, the depth of state 10 is 2, the depth of state 17 is 6, and the depth of the globally minimal states is defined as being infinite. The cups associated with the local minima that are not global minima are $\{10\}$, $\{11, 12\}$, and $\{14, 15, 16, 17, 18\}$.

The generation matrix for Hajek's 26-state example is irreducible, and as a result, the homogeneous SA Markov chain is also irreducible. Thus all states are recurrent. For Hajek's example, we investigate how the value of the temperature influences the time it takes a homogeneous SA chain to reach a global minimum. To this end we consider the Markov chain formed by making all global minima absorbing states. For $x = 1, 2, 26$, we set $G(x, y) = 1$ for $y = x$ and 0 otherwise. Given $\mathbf{Q}$, we can go on to calculate the fundamental matrix $\mathbf{N} = (\mathbf{I} - \mathbf{Q})^{-1}$. This can then be used to calculate the mean and variance of the time it takes a homogeneous SA Markov chain to reach a global minimum given by Theorem 3.1.

Performing these calculations for Hajek's example, we find, if we start the SA chain in, say, state 13, that the mean time until absorption is equal to

$$\begin{aligned} & (8 + 64\delta + 22\delta^2 + 190\delta^3 + 85\delta^4 + 243\delta^5 + 318\delta^6 + 180\delta^7 \\ & + 600\delta^8 + 107\delta^9 + 632\delta^{10} + 101\delta^{11} + 391\delta^{12} + 127\delta^{13} \\ & + 135\delta^{14} + 118\delta^{15} + 24\delta^{16} + 65\delta^{17} + 2\delta^{18} + 18\delta^{19} + 2\delta^{21}) \\ & / (\delta^6 (4 + 16\delta^2 + 21\delta^4 + 13\delta^6 + \delta^7 + 3\delta^8 + \delta^9 + \delta^{11})) \end{aligned}$$

with $\delta = \exp(-1/T)$, and the variance is equal to

$$\begin{aligned} & (64 + 896\delta + 4288\delta^2 + 5888\delta^3 + 26380\delta^4 + 25352\delta^5 + 82400\delta^6 \\ & + 90128\delta^7 + 168711\delta^8 + 236270\delta^9 + 267229\delta^{10} + 444772\delta^{11} \\ & + 387603\delta^{12} + 616478\delta^{13} + 557847\delta^{14} + 658450\delta^{15} + 747564\delta^{16} \\ & + 582882\delta^{17} + 849119\delta^{18} + 481773\delta^{19} + 779677\delta^{20} + 418840\delta^{21} \\ & + 571970\delta^{22} + 376007\delta^{23} + 340216\delta^{24} + 305693\delta^{25} + 174206\delta^{26} \\ & + 202119\delta^{27} + 86565\delta^{28} + 102605\delta^{29} + 45346\delta^{30} + 38468\delta^{31} \\ & + 23053\delta^{32} + 10177\delta^{33} + 9792\delta^{34} + 1782\delta^{35} + 3092\delta^{36} \\ & + 184\delta^{37} + 654\delta^{38} + 8\delta^{39} + 80\delta^{40} + 4\delta^{42}) \\ & / (\delta^{12} (4 + 16\delta^2 + 21\delta^4 + 13\delta^6 + \delta^7 + 3\delta^8 + \delta^9 + \delta^{11})^2). \end{aligned}$$

We see from Table 3.1 that, for Hajek's example, SA with a fixed temperature will find a global minimum more quickly, on average, for larger temperatures. The optimum strategy based on $E(\tau)$ is to take $T = \infty$, i.e., to adopt the "boiling" schedule which corresponds to a Markov chain with transition matrix given by the generation matrix. This strategy is the one that accepts all moves with probability 1.

Shown in Table 3.2 are the results from simulations performed for cooling schedules of a logarithmic type, including the canonical cooling schedule. Ten thousand runs were performed at each value of c . Again, it is apparent that the optimal strategy is to adopt boiling.

4. A state classification and eventual traps. As in the homogeneous case, the states of an inhomogeneous Markov chain may be classified as positive, null, recurrent, or transient. However, some of the definitions used for the homogeneous chains do not seem to carry over, whereas other definitions for inhomogeneous chains

which reduce to the classical ones are available. A state classification for finite and countable inhomogeneous chains is given in [8]. We shall adapt it here to the particular case of a SA chain. Also, the atomic sets of the tail σ -field (see [8]) admit in this case a neat representation in terms of some sets, which we shall call *eventual traps*.

A state x will be said to be *null* if $\lim_{n \rightarrow \infty} P(X_n = x) = 0$ and *positive* if $\lim_{n \rightarrow \infty} P(X_n = x) > 0$. Such a classification is not a dichotomy for inhomogeneous chains, but in the case of an SA chain, it is (see [26]). Let $\{A_n\}$ be a sequence of events. Write $\{A_n \text{ i.o.}\} = \cap_{n=1}^{\infty} \cup_{m=n}^{\infty} A_m$, where i.o. stands for *infinitely often* and $\{A_n \text{ ult.}\} = \cup_{n=1}^{\infty} \cap_{m=n}^{\infty} A_m$, where ult. stands for *ultimately*. We say that $\lim_{n \rightarrow \infty} A_n = A$ almost surely (a.s.) if $P(\{A_n \text{ ult.}\}) = P(\{A_n \text{ i.o.}\})$ and A is an event differing from $\{A_n \text{ ult.}\}$ only by a set of probability 0. We say that a state x is *recurrent* if

$$P(X_n = x \text{ i.o.}) > 0,$$

and *transient* otherwise. A positive state x is always recurrent and is called *positive recurrent*. Null states may be transient or recurrent. A null state which is recurrent will be called *null recurrent*. These definitions were given in [8].

We say that A is a recurrent class if it contains only recurrent states and for any $x \in A$

$$\{X_n = x \text{ i.o.}\} = \{X_n \in A \text{ i.o.}\} \text{ a.s.}$$

We say that the recurrent class A is an *eventual trap* if

- (i) $\lim_{n \rightarrow \infty} P(X_n \in A) > 0$ and
- (ii) $P(X_n = x \text{ i.o.}) = P(X_n \in A \text{ i.o.}) = P(X_n \in A \text{ ult.})$ for any $x \in A$.

We use the term eventual trap as distinct from trap to emphasize that a Markov chain reaching A may have a positive probability of escaping from A at all times, but as $n \rightarrow \infty$ such an escape becomes less and less likely and the chain must end up in an eventual trap with probability 1.

Remark. For an SA chain, it turns out that if A is an eventual trap with $\lim_{n \rightarrow \infty} P(X_n \in A) < 1$, then A^c , the complementary set to A , must contain at least one eventual trap.

A result of one of the present authors (see [8] and the references therein) describing the tail σ -field of a finite inhomogeneous Markov chain leads to the assertion that a chain of SA type has a finite number of disjoint eventual traps $A_1, \dots, A_t$ such that

$$\lim_{n \rightarrow \infty} P(X_n \in \cup_{i=1}^t A_i) = 1.$$

Obviously, the number of eventual traps does not exceed the cardinality of S .

5. Weak and strong ergodicity: Conditional convergence. Write $P^{(m,n)}(x, y) = P(X_n = y | X_m = x)$ for $m < n$. We shall say that $\{X_n\}$ is *weakly ergodic* if for any m, x, y , and z ,

$$\lim_{n \rightarrow \infty} (P^{(m,n)}(x, z) - P^{(m,n)}(y, z)) = 0.$$

A sufficient condition for weak ergodicity is the existence of some constant u such that

$$(5.1) \quad \sum_{k=1}^{\infty} \min_{x,y} P^{(k,k+u)}(x, y) = \infty.$$

However, it is easy to see that (5.1) requires that for all x

$$\sum_{n=1}^{\infty} P(X_n = x) = \infty.$$

In general, the above property is not necessary for weak ergodicity.

We say that $\{X_n\}$ is *strongly ergodic* if there is a probability distribution $\pi = (\pi_1, \dots, \pi_s)$ on S such that for any x, y , and $m \geq 1$,

$$(5.2) \quad \lim_{n \rightarrow \infty} P^{(m,n)}(x, y) = \pi_y.$$

It is easy to see that strong ergodicity implies weak ergodicity. For properties of weakly and strongly ergodic chains, see Seneta [29].

We say that $\{X_n\}$ is *conditionally convergent* if for some numbers $\pi_{x,y}^{(m)}$ and any $m \geq 1$,

$$(5.3) \quad \lim_{n \rightarrow \infty} P^{(m,n)}(x, y) = \pi_{x,y}^{(m)}.$$

In the literature of inhomogeneous Markov chains, such chains are known as *convergent* (see Mukherjea [25] and Cohn [7]), but that term has been used before in relation to property (1.1), so *conditional convergence* will be used for (5.3). For weakly ergodic chains, conditional convergence is equivalent to strong convergence, as defined in (5.2).

6. Slow cooling schedules. Let us write $\underline{f} = \min_{\{x \in S\}} f(x)$ and denote by $\bar{d}$ the number with the property

$$\sum_{k=1}^{\infty} \exp(-\bar{d}/T_k) = \infty$$

and

$$\sum_{k=1}^{\infty} \exp(-d/T_k) < \infty$$

for any $d > \bar{d}$.

If $\bar{d} > 0$ we shall say that $\{T_n\}$ is a slow cooling schedule. If $\bar{d} = 0$ we say that $\{T_n\}$ is a fast cooling schedule.

The following result is extracted from Niemiro and Pokarowski [26] and Niemiro [27].

THEOREM 6.1. *Suppose that $\{X_n\}$ is an SA chain with $\bar{d} > 0$. Then*

- (i) *there exist t recurrent classes $A_1, \dots, A_t$ which are eventual traps;*
- (ii) *the chain is conditionally convergent;*
- (iii) *if $x \in A_i$ then $y \in A_i$ if and only if y is reachable from x at height lower than or equal to $\underline{f}(A_i) + \bar{d}$, where $\underline{f}(A_i) = \min_{z \in A_i} f(z)$;*
- (iv) *if $\bar{S}$ is the set consisting of deepest states of $A_1, \dots, A_t$, then $\lim_{n \rightarrow \infty} P(X_n \in \bar{S}) = 1$;*
- (v) *if $x \in \bar{S}$, then $\lim_{n \rightarrow \infty} P(X_n = x) > 0$.*

Notice in particular that a convergent chain may admit either only one eventual trap (the case of a weakly ergodic chain) or several eventual traps, each of them containing some global minimum states on which the whole probability mass will eventually concentrate.

On the other hand, it is easy to see that by decreasing $\bar{d}_1$ to, say, $\bar{d}_2$, the number of eventual traps does not decrease, because any eventual trap for $\bar{d}_2$ is either an eventual trap for $\bar{d}_1$ or belongs to a partition of an eventual trap for $\bar{d}_1$. Thus, if $\bar{d}$ is smaller than d^* , the SA algorithm does not converge because the set of eventual traps will necessarily include some that do not contain any global minima states. Such eventual traps attract the chain to local minima. Since each of the limit probabilities $\lim_{n \rightarrow \infty} P(X_n \in A_i)$ is positive, property (2.1) also fails for chains of this kind. Indeed, notice that $\lim_{n \rightarrow \infty} P(X_n \in S^*) = 1 - \sum_{i \in \Lambda_{\bar{d}}} \lim_{n \rightarrow \infty} P(X_n \in A_i)$, where $\{A_i : i \in \Lambda_{\bar{d}}\}$ is the collection of eventual traps corresponding to $\bar{d}$ which do not contain global minima states. Clearly, $\lim_{n \rightarrow \infty} P(X_n \in S^*)$ becomes smaller as $\bar{d}$ decreases and, as a result, the number of eventual traps increases. If $\bar{d}$ is sufficiently small, then any local minima states may form the bottom of some eventual trap. This is the reason why some heuristics cooling faster than logarithmic are not convergent. Such algorithms may end up in a local minimum. We shall describe a number of such algorithms later in the paper.

THEOREM 6.2. *A convergent chain is weakly ergodic if and only if one of the following two statements holds:*

- (i) *There is only one global minimum state.*
- (ii) *If x and y are two global minima states, then x is reachable from y at height smaller than or equal to $\underline{f} + \bar{d}$.*

Proof. It is easy to see that if (i) holds, the only global minima state, say x , is in a recurrent class A which is an eventual trap. However, A is the only eventual trap since $\lim_{n \rightarrow \infty} P(X_n \in A) \geq \lim_{n \rightarrow \infty} P(X_n = x) = 1$. Thus $\{X_n\}$ has a trivial tail σ -field, which implies weak ergodicity (see [6]).

To prove (ii) notice that by Theorem 6.1 (iii) all the global minima states must be in one recurrent class which is the unique eventual trap. This completes the proof.

COROLLARY 6.3. *A convergent chain is not weakly ergodic if and only if there exist two global minima states x and y such that y is reachable from x at height higher than $\underline{f} + \bar{d}$.*

THEOREM 6.4. *A weakly ergodic SA chain corresponding to a cooling schedule is convergent and strongly ergodic.*

Proof. Any weakly ergodic chain has a trivial tail σ -field and therefore could not admit more than one eventual trap. However, the only cooling schedules that are not convergent are the ones that admit several eventual traps, with at least one having no global minimum states. This proves convergence. Strong ergodicity follows from Theorem 6.1 and weak ergodicity.

Remark. If we do not confine ourselves to cooling schedules, then weak ergodicity may not imply convergence, as we have seen in the case of fixed temperature schedules.

To summarize the above results on convergence, we conclude that

- (i) the canonical cooling schedule may result in a chain that is not weakly ergodic;
- (ii) the canonical constant d^* is the cutoff point for $\bar{d}$ below which the process exhibits a *phase transition*, with its class of eventual traps increasing to include some local minima traps.

LEMMA 6.5. *Suppose that $\sum_{n=1}^{\infty} P_n(x, y) = \infty$, where $\liminf_{n \rightarrow \infty} P(X_n = x) > 0$. Then $P(X_n = x, X_{n+1} = y \text{ i.o.}) > 0$.*

Proof. Write $A_n = \{X_n = x, X_{n+1} = y\}$. We shall show that a divergent part of the Borel–Cantelli-type lemma holds for the events $\{A_n\}$. Write $\mathcal{F}_n$ for the σ -field generated by $X_1, \dots, X_n$. The Markov property of $\{X_n\}$ yields

$$(6.1) \quad P(A_n | \mathcal{F}_n) = P_n(x, y) 1_{\{X_n = x\}}.$$

According to the Borel–Cantelli–Levy lemma,

$$(6.2) \quad P(A_n \text{ i.o.}) > 0 \text{ if and only if } P\left(\sum_{n=1}^{\infty} P(A_n | \mathcal{F}_n) = \infty\right) > 0.$$

Consider now the random variable

$$(6.3) \quad Y_n = \frac{\sum_{k=1}^n P_k(x, y) 1_{\{X_k=x\}}}{\sum_{k=1}^n P_k(x, y)}.$$

Notice that the denominator in (6.3) tends to ∞ as $n \rightarrow \infty$. Since $0 \leq Y_n \leq 1$, for $\{Y_n\}$ to converge in probability to 0, it is necessary that $E(Y_n) \rightarrow 0$. However, this is not the case as $\liminf_{n \rightarrow \infty} E(Y_n) \geq \liminf_{n \rightarrow \infty} P(X_n = x) > 0$, which implies $P(\sum_{k=1}^{\infty} P_k(x, y) 1_{\{X_k=x\}} = \infty) > 0$; the proof is concluded on account of (6.1) and (6.2).

This lemma provides a criterion of recurrence for a state y which is reachable in one step from a positively recurrent state.

THEOREM 6.6. *If $\{X_n\}$ is convergent, then*

- (i) $\{f(X_n)\}$ converges in probability to $\underline{f}$;
- (ii) $\{f(X_n)\}$ converges a.s. to $\underline{f}$ if and only if

$$(6.4) \quad \sum_{n=1}^{\infty} \exp\left(-\frac{d(x)}{T_n}\right) < \infty$$

for any state x with $x \notin S^*$ and $x \in N(y)$, where y is a global minimum state.

Proof. Since f is constant on S^* , (i) follows from the definition of convergence (1.1).

To prove (ii), notice that (6.4) implies

$$\sum_{n=1}^{\infty} P(\{X_n \in S^*\} \cap \{X_{n+1} \notin S^*\}) < \infty.$$

By a Borel–Cantelli-type lemma given by Barndorff-Nielsen [4], the above implies that $P(\{X_n \in S^* \text{ ult.}\}) = 1$. Thus all states outside S^* are transient, which proves the first implication of (ii).

Assume now that (6.4) fails. Thus there exists a state x with $f(x) > \underline{f}$, $x \in N(y)$, where y is a global minimum state, and

$$\sum_{n=1}^{\infty} \exp\left(-\frac{d(x)}{T_n}\right) = \infty.$$

According to Theorem 6.1, all states of S^* are positive. Thus we can use Lemma 6.5 to conclude that x is recurrent. However, in this case, $P(\{f(X_n) \geq f(x) > \underline{f} \text{ i.o.}\}) > 0$, contradicting the almost sure convergence of $\{f(X_n)\}$ to $\underline{f}$. This completes the proof of (ii).

7. Critical points for the SA chains. Next we shall identify a number of critical points for the constant c of an SA chain with logarithmic temperature schedule.

THEOREM 7.1. *Suppose that the SA chain $\{X_n\}$ admits a cooling schedule $\{T_n = c/\log(n_0 + n)\}$ for some constant c .*

1. Define c_0 to be the smallest $h > 0$ such that

$$\exists x \in M, y \in N(x) : d(x) < h \text{ and } f(x) < f(y) \leq f(x) + h,$$

where M is the set of all local minima, including global minima. Then c_0 is the smallest c such that null recurrent states exist. For $c < c_0$, the SA chain assumes only positive recurrent and transient states, and its collection of eventual traps is maximal in number.

2. Define

$$c_1 = \min_{x \in M, x \notin S^*} d(x),$$

where M is the set of all local minima states. c_1 is the smallest c such that the number of recurrent classes that are eventual traps decreases.

3. Define

$$c_2 = d^* = \max_{x \in M, x \notin S^*} d(x).$$

Then c_2 is the smallest c such that the algorithm is convergent. It is also the smallest c such that all local minima that are not global minima are null states, or the smallest c such that the only positive recurrent states are global minima states.

4. Define c_3 to be the smallest c such that a null recurrent local minimum exists. c_3 is the smallest h such that there exist a local or global minimum state x and a local minimum state y with $d(x) > h$ and y reachable at height $f(x) + h$ from x .

5. Define

$$h^* = \max_{x, y \in S^*} \{h : y \text{ is reachable from } x \text{ at height } \underline{f} + h\}$$

and $c_4 = \max\{d^*, h^*\}$. Then c_4 is the smallest c for which weak ergodicity occurs.

6. Define

$$c_5 = \max_{x \in S} f(x) - \min_{y \in S} f(y).$$

Then c_5 is the smallest c for which all states are recurrent.

7. Define $c_6 = +\infty$ in the case when the transition probabilities of the SA chain do not depend on the temperature. Then c_6 is the only c for which all states are positive recurrent.

Proof. Notice first that by simple manipulations we deduce that for any c_i with $i \in \{1, \dots, c_6\}$ we have for $\alpha \geq c_i$

$$\sum_{k=1}^{\infty} \exp(-\alpha/T_k) = \infty$$

and

$$\sum_{k=1}^{\infty} \exp(-\alpha/T_k) < \infty$$

for $\alpha < c_i$.

FIG. 7.1. $0 \leq c < 1$. All local and global minima are positive states. It is a defective algorithm. No null recurrent states exist. The chain may freeze in any (connected set of) local minima including global minima. There are four eventual traps. This also is the case for all fast cooling schedules.

FIG. 7.2. $1 \leq c < 2$. A null recurrent state first occurs. One local minimum is rendered transient. If the SA chain becomes trapped in the eventual trap containing the null recurrent state, then, strictly speaking, the chain will never freeze. That is, the null recurrent state will be visited infinitely often. Such visits will, however, become less and less frequent and further apart. There are three eventual traps.

To prove point 1 we take into account that for $c = c_0$ the positive states will remain the same as for $c < c_0$ but, according to Lemma 6.5, the set of null recurrent states will increase.

To prove point 2, notice that for $c = c_1$, at least two eventual traps for $c < c_1$ become merged in one eventual trap. This follows from Theorem 6.1(iii).

Point 3 is also a consequence of Theorem 6.1, because for $c = c_2$ the eventual traps containing global minima must contain all local minima as well. As the probability mass concentrates in the bottom states of a recurrent eventual trap the chain must be convergent.

Remark. It has turned out that the critical points identified above belong to a logarithmic cooling schedule. For cooling schedules that go faster to 0 than a logarithmic one, we can easily see that the SA chain behavior is the one described for $c < c_0$. For temperature schedules that are slower than logarithmic, the SA chain

FIG. 7.5. $4 \leq c < 5$. The chain is weakly ergodic. All global minima are contained in the single eventual trap.

FIG. 7.6. $c \geq 5$. There is one eventual trap, incorporating the entire state space. All states are recurrent, but only global minima are positive recurrent.

behavior is as in the case $c \geq c_5$. For temperature schedules with a subsequence of $\{T_n\}$ bounded away from 0, we get a weakly ergodic $\{X_n\}$ with all states positive recurrent.

We shall consider now an example of an SA chain with 15 states to illustrate the asymptotic behavior of slow cooling schedules described above. The example is shown in Figures 7.1–7.7, where the properties of states are depicted at various values of c . The dotted graphs delineate the eventual traps. Marked in black are the positive recurrent, in gray the null recurrent, and in white the transient states. For this example we get $c_0 = c_1 = 1$, $c_2 = 2$, $c_3 = 3$, $c_4 = 4$, and $c_5 = 5$.

8. Fast cooling schedules. Most of the algorithms applied to large problems are of the fast cooling type and are therefore nonconvergent. This is the case for the algorithms of Aarts and van Laarhoven [1], Kirkpatrick, Gellat, and Vecchi [19], and

FIG. 7.7. $c = +\infty$. The case where boiling is employed and all transitions are accepted. No cooling takes place. All states are positive recurrent. This also is the case for any fixed positive temperature.

Lundy and Mees [23], which satisfy the condition

$$(8.1) \quad \sum_{n=1}^{\infty} P_n(x, y) < \infty$$

for any x, y with $f(y) > f(x)$.

Consider now the Markov chain $\{X_n^A\}$ with state space A and transition probability matrix R , with entries

$$R(x, y) = \begin{cases} G(x, y) & \text{if } f(y) < f(x), \\ 0 & \text{if } f(y) > f(x), \\ 1 - \sum_{z \neq x} G(x, z) & \text{if } f(y) = f(x) \end{cases}$$

for $x, y \in A$. We shall attach such Markov chains to any (local minima) recurrent class A .

THEOREM 8.1. *If (8.1) holds, then*

- (i) $\{f(X_n)\}$ converges a.s. to a random variable W whose probability mass is concentrated on the set of local and global minima;
- (ii) all the states except for global and local minima are transient;
- (iii) $\{X_n\}$ eventually freezes in a set of states of constant objective function f ;
- (iv) if x belongs to a recurrent class A consisting of local or global minima states, then $\lim_{n \rightarrow \infty} P(X_n = x) = P(\Lambda)\pi_x$, where $\{\pi_x, x \in A\}$ is the stationary distribution of $\{X_n^A\}$, and $\Lambda = \lim_{n \rightarrow \infty} \{X_n \in A\}$ a.s.

Proof. We shall show that if A is a recurrent class consisting only of global or local minima of constant f -value, then A is an eventual trap. We shall prove first that for such A we get

$$(8.2) \quad P(\{X_n \in A \text{ ult.}\}) > 0.$$

This is equivalent to showing that

$$\lim_{n \rightarrow \infty} P(\cap_{m=n}^{\infty} \{X_m \in A\}) > 0.$$

Write

$$(8.3) \quad \alpha_n = \max_{x \in A, y \notin A} P_n(x, y).$$

By conditioning, we get

$$P(\cap_{m=n}^r \{X_m \in A\}) \geq P(\cap_{m=n}^{r-1} \{X_m \in A\})(1 - \alpha_{r-1}) \geq$$

.....

$$(8.4) \quad \geq P(\{X_m \in A\})(1 - \alpha_m) \cdots (1 - \alpha_{r-1}).$$

Letting r tend to ∞ in (8.4) and recalling that $\sum_n \alpha_n < \infty$ we get that $P(\{X_n \in A \text{ ult.}\}) > 0$ and (8.2) is proved.

Notice now that

$$(8.5) \quad \sum_{n=1}^{\infty} P(\{X_n \in A\} \cap \{X_{n+1} \notin A\}) < \infty,$$

which in conjunction with (8.2) and the Barndorff-Nielsen-Borel-Cantelli-type lemma [4] imply that A is an eventual trap. The transient states do not have an a.s. contribution in the limit. (Notice that f is constant on A but its value may differ for various eventual traps being the value of the bottom states which are global or local minima of f .) This completes the proof of (i).

Obviously, (ii) follows from (i).

It is easy to see that (iii) follows from (i) and (ii).

To prove (iv), notice that the assumption of irreducibility and accessibility of states from each other makes any chain $\{X_n^A\}$ ergodic and irreducible. Thus

$$\lim_{n \rightarrow \infty} P^{(m,n)}(y, x) = \pi_x$$

for $x, y \in A$, and

$$\lim_{n \rightarrow \infty} P^{(m,n)}(y, x) = 0$$

for $x \notin A$. But

$$P(X_n = x) = \sum_{y \in S} P(X_m = y) P^{(m,n)}(y, x).$$

Thus, if $x \in A$,

$$\lim_{n \rightarrow \infty} P(X_n = x) = \lim_{m \rightarrow \infty} P(X_m \in A) \pi_x = P(A) \pi_x,$$

and the proof is finished.

9. Some traveling salesman examples. We next investigate the relative performance of a number of cooling schedules used in applications to which we add a fixed temperature schedule. There is no claiming that the algorithms chosen are the most appropriate for the problems. The aim of the exercise is to use statistical analysis to ascertain the quality of various algorithms which appear to be problem dependent. Clearly, for some problems there is a need for faster, nonconvergent cooling than logarithmic cooling. We examine the performance of fast cooling schedules such as Aarts and van Laarhoven [1] and Lundy and Mees [23] and a basic geometric schedule as first introduced by Kirkpatrick, Gellat, and Vecchi [19]. The traveling salesman problems (TSPs) considered vary in size from 48 to 442 cities.

We consider the relative performance of these algorithms allowing a fixed number of iterations N , for an appropriately chosen N .

9.1. Application of SA to the TSP. For the TSP, we consider a path leading through all of n cities, starting in an arbitrary city and finally returning to it. A distance (or possibly time or cost) is given between each pair of cities. We consider here the symmetric TSP, where the distance is the same in either direction. The objective is to identify the path that has the smallest total distance. There are $(n-1)!/2$ possible paths.

The neighborhood structure we employ for the TSP is that generated by 2-opt moves. Consider the cities and the path of the TSP as the vertices and edges of a graph. A 2-opt move is simply the process of deleting and replacing two edges of the graph to yield a new path for the TSP. There are $n(n-3)/2$ different paths that can be created by such a move. (Note that once one edge has been deleted, if either of the neighboring edges is then deleted, it is possible only to reconstruct the original path, leaving $n-3$ edges to choose from.)

The TSP is often stated as a *benchmark* problem for testing optimization procedures. SA is often outperformed by specially tailored algorithms. The merits of SA lie in its ease of implementation and its applicability to a wide range of problems. It is our aim to use the observations of SA on TSPs to gain valuable insight into what criteria constitute an optimal temperature schedule for problems in general.

9.2. The problem instances. We have considered the six problem instances of the TSP examined in Aarts and van Laarhoven [2]. Each problem is labeled by the initials of the author(s) of the reference to it, followed by the number of cities. The problem instances are **gr48** and **gr442** from [14], **gr120** from [13], **kt57** from [17], **kroA100** from [21], and **lin318** from [22]. (We have taken **lin318** in the form of a TSP rather than a Hamiltonian circuit.)

9.3. The different schedules. Following are the rules for updating the temperature in each of the schedules considered.

Aarts: Temperature is held fixed during each loop of $R = \max_{x \in S} |N(x)|$ iterations. At the end of each loop the temperature is dropped according to the rule

$$T_{k+1} = T_k \left/ \left(1 + \frac{T_k \log(1 + \delta)}{3\sigma_k} \right) \right.,$$

where σ_k is the standard deviation of the observed values of the cost function during the k th loop of the algorithm.

Geometric: The temperature is again held fixed during each loop. We have set the length of each loop to be the same as for Aarts. At the end of each loop the

temperature is dropped according to the rule

$$T_{k+1} = \alpha T_k.$$

It is worth noting that we found that the number of iterations performed at each loop had little if any effect on the algorithm's performance, provided the value of α was adjusted appropriately.

Lundy: With Lundy's schedule the temperature is to be dropped after each iteration according to the rule

$$T_{n+1} = \frac{T_n}{1 + \beta T_n},$$

or equivalently,

$$T_n = \frac{T_0}{1 + n\beta T_0}.$$

To keep this algorithm in the same form as above, we update the temperature at the end of each loop of the same number of iterations as above. Again, we did not find this to alter the performance of the algorithm.

Logarithmic: Here again, the temperature is to be updated after each iteration, the rule for which is

$$T_n = \frac{c}{\log(n + n_0)}.$$

Again we update this temperature at the end of each loop of R iterations.

Fixed temperature: In a fixed temperature algorithm an appropriate temperature must be found. We have done so experimentally, by running a fixed temperature schedule for a range of temperatures and choosing the temperature which gives the best performance, say, the best average solution in N iterations. Connolly [9] gives a method for determining a fixed temperature by first running a fast cooling algorithm and noting the temperature at which the best solution found first occurred.

9.4. Method used in comparing the schedules. In an attempt to make a fair comparison of the different schedules, the following method is used.

Measure of performance: We measure the performance of each algorithm by the average best solution found in the N iterations. Results of the algorithms with regards to $P(\tau \leq N)$ are also given, where τ is taken as the time until reaching a global minimum, as well as within one and two percent of the global minimum.

To choose N : We wish to choose an N for each problem instance that is sufficiently large, but not too large, for the algorithms to find good heuristic solutions. We have chosen Aarts's algorithm to roughly determine such an N , but Lundy's or the geometric algorithm also could have been used. First, 100 runs of Aarts's algorithm are performed with the parameter setting ($\delta = 0.1$) recommended by its authors. In choosing N , we consider the number of iterations taken until first visiting the best solution found in each run. The maximum of these is taken, after removing outliers. An outlier is taken as a value more than 1.5 times the interquartile range greater than the third quartile. The initial temperature is determined experimentally to yield an initial acceptance ratio of 0.95.

Determining the parameters of the schedules: Once N has been chosen for a given problem, the parameters of Lundy's algorithm and the geometric schedule are

TABLE 9.1

The parameter settings experimentally found for five temperature schedules for various TSPs. N is the number of iterations to be allowed for each algorithm and is set according to Aarts's schedule.

Problem	N	T_0	Aarts δ	Geom. α	Lundy β	Logar. c	Fixed T
gr48	509760	2800	0.1	0.98700	2.546×10^{-7}	250	20
kt57	857223	6000	0.1	0.98920	6.173×10^{-8}	500	40
kroA100	4205532	11500	0.1	0.99220	1.196×10^{-8}	650	45
gr120	7104240	2900	0.1	0.99300	3.704×10^{-8}	150	11
lin318	102173400	11800	0.1	0.99615	1.498×10^{-9}	450	25
gr442	242935584	2420	0.1	0.99670	5.669×10^{-9}	45	2.3

determined experimentally to yield approximately the same N , when determined in the same way. The initial temperature is set as above. For fixed temperature, the optimum temperature is found experimentally by trying various temperatures and finding the one that yields on average the best solution in the N iterations. The logarithmic schedule is very slow and cooling from a high to a low temperature in the given amount of time is not possible. We therefore set $n_0 = 2$ in order to maximize the overall change in temperature, and we determine the optimal value for c in the same way as we determine the optimum fixed temperature.

Stopping the algorithms: Once parameters are chosen, the algorithms are rerun. Upon reaching N iterations the temperature is set to zero, and the algorithms are allowed to (quickly) settle in a local minimum. For the logarithmic schedule the optimum value of c is found, with this final freezing included in the algorithm. This final freezing is also included when searching for the optimal fixed temperature.

9.5. Results. Tables 9.1–9.5 show the results from running the five above-mentioned temperature schemes on the six TSP instances. One hundred runs are performed for each instance under each temperature schedule. Table 9.1 shows the number of iterations allowed for each problem instance and the parameters experimentally determined for each algorithm. Reported are the quality of final (best) solutions, iterations taken to reach these solutions, and the proportion of runs reaching global or near global minima solutions. Global minima solutions were found only for the 48, 57, and 100 cities instances.

9.6. Remarks regarding simulations. 1. From the simulations carried out, we see that it is worthwhile having a handful of algorithms available in the application of SA to a particular problem.

2. In the case of the TSP, we see that for smaller problems, the fixed and logarithmic schedules seem to perform as well as and better than the fast cooling schedules. For larger problems the fast cooling schedules seem to perform better. It appears that in such cases the schedule of Lundy and Mees outperforms the Aarts and van Laarhoven and the geometric schedules.

3. The results are likely to differ for different applications of SA. Lundy and Mees's algorithm initially cools more rapidly than the other two fast cooling schedules, and it spends more time at smaller temperatures. It may be the case, however, that the slower initial cooling of the other schedules is crucial in other applications.

4. We see that for the 48-city and 120-city TSPs, fixed temperature and the logarithmic schedule outperform the fast cooling schedules. The results suggest that it is not simply the size of the problem that is important but the structure as well. It may be the case that for applications other than the TSP, the structure of the

TABLE 9.2

A comparison of different temperature schedules, in a fixed number of iterations, for various TSPs. Mean and standard deviation given for 100 runs in each case. The solution in each run is taken as the best solution visited.

Problem	Average best solution (% above global)				
	Aarts	Geom.	Lundy	Logar.	Fixed
gr48	0.88	0.66	0.38	0.24	0.25
kt57	1.07	0.84	0.40	0.56	0.59
kroA100	0.96	0.75	0.49	0.49	0.54
gr120	1.83	1.40	1.07	0.69	0.85
lin318	1.73	1.45	1.34	2.16	2.37
gr442	1.66	1.34	1.05	2.00	2.10
Problem	Standard deviation (% above global)				
	Aarts	Geom.	Lundy	Logar.	Fixed
gr48	0.69	0.54	0.40	0.28	0.28
kt57	0.85	0.77	0.54	0.69	0.69
kroA100	0.70	0.52	0.38	0.38	0.56
gr120	0.72	0.66	0.50	0.39	0.43
lin318	0.52	0.45	0.45	0.67	0.79
gr442	0.51	0.37	0.39	0.47	0.45

TABLE 9.3

Mean and standard deviation of iterations taken until finding the best solution of each run.

Problem	Average iterations until best in run				
	Aarts	Geom.	Lundy	Logar.	Fixed
gr48	478375	440705	280476	280260	247698
kt57	809653	763375	540635	474612	494819
kroA100	4113819	3821121	2927120	2703681	2843458
gr120	6828916	6305855	4467879	4932393	4586938
lin318	100592216	95016256	72002496	99782344	96906464
gr442	237841104	226062032	160714880	222006640	206698016
Problem	Standard deviation of iterations until best				
	Aarts	Geom.	Lundy	Logar.	Fixed
gr48	15981	27161	81559	138893	145977
kt57	20622	36932	117102	226153	240173
kroA100	64690	173474	720508	1257878	1339295
gr120	114856	318401	1194903	1658310	2017348
lin318	835365	2079017	11570513	9369117	16177562
gr442	1823134	6554549	31531302	34883836	52499600

TABLE 9.4

Estimates for $P(\tau \leq N)$, where τ is the time to reaching a global minimum, for N as given in Table 9.1. A global minimum was never reached in any of the runs for the larger problems.

Problem	Proportion reaching global minimum				
	Aarts	Geom.	Lundy	Logar.	Fixed
gr48	0.05	0.17	0.35	0.30	0.34
kt57	0.04	0.06	0.28	0.31	0.30
kroA100	0.03	0.06	0.10	0.04	0.00

problem means that fixed and logarithmic schedules are suited to large problems too.

5. We do not know whether the logarithmic cooling schedule used is convergent, as we have not identified the canonical constant. Indeed, d^* is not readily available, and to get it, when feasible, may require much more extensive work than finding an optimal state. In fact, as we pointed out before, convergence is not relevant to the success of the algorithm.

TABLE 9.5

Estimates for $P(\tau \leq N)$ when τ is taken, respectively, as the time to reaching solutions with the objective function at most one percent and two percent larger than the global minimum.

Problem	Proportion reaching within 1% of global				
	Aarts	Geom.	Lundy	Logar.	Fixed
gr48	0.68	0.81	0.97	1.00	1.00
kt57	0.58	0.71	0.86	0.67	0.65
kroA100	0.62	0.76	0.91	0.93	0.88
gr120	0.11	0.27	0.44	0.80	0.69
lin318	0.07	0.13	0.19	0.01	0.02
gr442	0.12	0.23	0.39	0.02	0.00
	Proportion reaching within 2% of global				
	Aarts	Geom.	Lundy	Logar.	Fixed
gr48	0.92	0.97	1.00	1.00	1.00
kt57	0.89	0.89	0.99	1.00	1.00
kroA100	0.93	0.98	0.99	1.00	0.96
gr120	0.57	0.81	0.95	0.99	0.99
lin318	0.70	0.88	0.92	0.52	0.34
gr442	0.77	0.98	1.00	0.49	0.43

6. We have seen that obtaining a global minimum is plausible for some small to medium-size problems. For the 100-city TSP, using Lundy's schedule, we get for the time until reaching a global minimum,

$$P(\tau < N) \approx 0.10,$$

for $N = 4205532$. Using $k = 50$ reruns (2.3) becomes

$$P(\min_{i \in \{1, \dots, k\}} \min(f(X_1^{(i)}), \dots, f(X_N^{(i)})) = \min_{x \in S} f(x)) \approx 0.995,$$

and with $k = 100$ we get a probability of 0.99997 of reaching the global minimum.

10. Concluding remarks. 1. We have looked at the limit behavior of the SA chain as a function of its temperature schedule. The quality of an algorithm depends on its parameters, and the temperature schedule is only one ingredient of an algorithm. However, the limit behavior of the SA chain is determined only by its temperature schedule.

2. We have a three-type classification for an algorithm: convergent, regular, and defective. Examples are provided to illustrate situations when boiling gives the optimal algorithm, when logarithmic or fixed temperature outperform a number of faster cooling schedules, or when defective algorithms are better for the problem.

3. We characterized the limit behavior of an algorithm in terms of recurrence, transience, and eventual traps. It turns out that a convergent chain may have several eventual traps or may consist of one eventual trap, as in the weakly ergodic case. A regular algorithm is not necessarily convergent. It may be weakly ergodic but not convergent. A convergent chain or a chain with a fixed temperature will exhibit a lot of changes in its objective function values, as there are usually recurrent states that are neither global nor local minima. Such changes will become less and less frequent but will not disappear. In contrast, a defective chain does not have recurrent states outside global or local minima states and will eventually have its objective function value frozen in a local or global minimum.

4. The critical points for algorithms where the asymptotic behavior changes are all in the range of logarithmic cooling schedules. There are two extreme types of behavior: the first, when each local minimum is an eventual trap, and the second,

when all states are recurrent. It may seem that the first case does not lead to a good algorithm. However, for large problems these type of schedules usually outperform the convergent and regular ones. It may also seem that the latter compares unfavorably to the canonical cooling schedule which prescribes a convergent chain with the minimal number of recurrent states. However, such an impression is also deceptive.

5. When using a memoryless algorithm for a convergent chain or a memory algorithm for a regular chain, we know that reaching global minima may be achieved with probability as large as desired if we let the chain run a sufficiently long time. However, that may not be feasible in practice, as it may require an excessively long time. In contrast, for defective algorithms we know that the probability of reaching optimality is limited, often by a small number. However, repeated independent runs may ensure a high quality for such algorithms, which are often used in practice.

Acknowledgment. The authors are thankful to the referees for a number of useful comments.

REFERENCES

- [1] E. H. L. AARTS AND P. J. M. VAN LAARHOVEN, *Statistical cooling: A general approach to combinatorial optimization problems*, Philips J. Res., 40 (1985), pp. 193–226.
- [2] E. H. L. AARTS AND P. J. M. VAN LAARHOVEN, *Simulated Annealing: Theory and Applications*, Reidel, Dordrecht, the Netherlands, 1987.
- [3] E. H. L. AARTS AND J. KORST, *Simulated Annealing and Boltzman Machines*, John Wiley, New York, 1989.
- [4] O. BARNDORFF-NIELSEN, *On the limit behaviour of extreme order statistics*, Ann. Math. Statist., 34 (1963), pp. 992–1002.
- [5] T. S. CHIANG AND Y. CHOW, *On the convergence rate of annealing processes*, SIAM J. Control Optim., 26 (1988), pp. 1455–1470.
- [6] H. COHN, *On a paper by Doeblin on non-homogeneous Markov chains*, Adv. Appl. Prob., 13 (1981), pp. 388–401.
- [7] H. COHN, *On a class of non-homogeneous Markov chains*, Math. Proc. Cambridge Philos. Soc., 92 (1982), pp. 527–534.
- [8] H. COHN, *Products of stochastic matrices and applications*, Internat. J. Math. Math. Sci., 12 (1988), pp. 209–233.
- [9] D. T. CONNOLLY, *An improved annealing scheme for the QAP*, European J. Oper. Res., 46 (1990), pp. 93–100.
- [10] D. P. CONNORS AND P. R. KUMAR, *Balance of recurrence order in time-inhomogeneous Markov chains with applications to simulated annealing*, Probab. Engrg. Inform. Sci., 2 (1988), pp. 157–184.
- [11] S. B. GELFAND AND S. K. MITTER, *Analysis of simulated annealing for optimization*, in Proc. 24th Conference on Decision and Control, Ft. Lauderdale, FL, 1985, pp. 779–786.
- [12] S. GEMAN AND D. GEMAN, *Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images*, IEEE Trans. Pattern. Anal. Machine Intelligence, 6 (1984), pp. 721–741.
- [13] M. GRÖTSCHEL, *Polyedrische Charakterisierungen Kombinatorischer Optimierungsprobleme*, Hain, Meisenheim am Glan, 1977.
- [14] M. GRÖTSCHEL AND O. HOLLAND, *Solution of large-scale symmetric traveling salesman problems*, Math. Programming, 51 (1991), pp. 141–202.
- [15] B. HAJEK, *Cooling schedules for optimal annealing*, J. Math. Oper. Res., 13 (1988), pp. 311–329.
- [16] C. R. HWANG AND S. J. SHEU, *Singular perturbed Markov chains and the exact behaviors of simulated annealing processes*, J. Theoret. Probab., 5 (1992), pp. 223–249.
- [17] R. L. KARG AND G. L. THOMPSON, *A heuristic approach to solving traveling salesman problems*, J. Management Sci., 10 (1964), pp. 225–248.
- [18] J. G. KEMENY AND J. L. SNELL, *Finite Markov Chains*, Springer-Verlag, New York, 1976.
- [19] S. KIRKPATRICK, C. D. GELLAT, AND M. P. VECCHI, *Optimization by simulated annealing*, Science, 220 (1983), pp. 671–680.

- [20] S. KIRKPATRICK, *Optimization by simulated annealing: Quantitative studies*, J. Statist. Phys., 34 (1983), pp. 975–986.
- [21] P. D. KROLAK, W. FELTS, AND G. MARBLE, *A man-machine approach toward solving the traveling salesman problem*, Comm. ACM, 14 (1971), pp. 327–334.
- [22] S. LIN AND B. W. KERNIGHAN, *An effective heuristic algorithm for the traveling salesman problem*, J. Oper. Res., 21 (1973), pp. 498–516.
- [23] M. LUNDY AND A. MEES *Convergence of an annealing algorithm*, Math. Programming, 34 (1986), pp. 111–124.
- [24] N. METROPOLIS, A. W. ROSENBLUTH, M. N. ROSENBLUTH, A. TELLER, AND E. TELLER, *Equations of state calculations by fast computing machines*, J. Chem. Phys., 21 (1953), pp. 1087–1092.
- [25] A. MUKHERJEA, *A new result on the convergence of non-homogeneous Markov chains*, Trans. Amer. Math. Soc., 90 (1981), pp. 167–182.
- [26] W. NIEMIRO, *Limit distributions of simulated annealing Markov chains*, Disc. Math. Algebra Stochastic Methods, 15 (1997), pp. 241–269.
- [27] W. NIEMIRO AND P. POKAROWSKI, *Tail events of some non-homogeneous Markov chains*, Ann. Appl. Probab., 5 (1995), pp. 261–293.
- [28] F. ROMEO AND A. SANGIOVANNI-VINCENTELLI, *Probabilistic hill climbing algorithms: Properties and applications*, Proc. Chapel Hill Conference on VLSI, H. Fuchs, ed., Computer Science Press, Rockville, MD, 1985, pp. 393–417.
- [29] E. SENETA, *Non-Negative Matrices and Markov Chains*, Springer-Verlag, New York, 1981.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.