The Annals of Applied Probability
1995, Vol. 5, No. 4, 1087–1094

PROBABILISTIC SEARCH WITH OVERRIDES

BY BENNETT L. FOX[1] AND GEORGE W. HEINE

University of Colorado

Consider a time-inhomogeneous Markov chain which converges in probability to a subset S_0 of its state space. Override the standard move mechanism up to a random transition time, almost surely finite, but not necessarily a stopping time. Under weak conditions, the modified process converges in probability to the same set S_0. Two examples of independent interest illustrate this result.

1. Introduction. Under various conditions, the (nonhomogeneous) Markov chain $(X(1), X(2), \dots)$ of simulated annealing [see Kirkpatrick, Gelatt and Vecchi (1983) or Černy (1985)] converges in probability to the set S_0 of global optimizers. Suppose that we modify the search strategy on a random but finite number of steps. In particular, let the search strategy have its standard move mechanism overridden up to a random time N, assumed almost surely finite. Just after N, the standard move mechanism takes over. We show that the current state of the modified process converges in probability to the set S_0 of optimal states if the current state of the original process does; that is, the Markov chain associated with the modified process converges in probability to S_0 whenever this is true in the unmodified process. Our main result states this assertion precisely. Interestingly, our result holds under certain conditions even when N is not a stopping time. The result is not surprising; in fact, it would be surprising if it were false. However, the proof, for N not a stopping time, seems to require more subtlety than one might expect.

An obvious goal of overrides, especially with N a stopping time, is to shorten the number of moves in the chain X to first hit the set S_0 of optimal states. The examples in Sections 3 and 4 show uses for overrides when N is not a stopping time. In Section 4, the setting is noisy objective functions. There, the goal is not so much to shorten the time to first hit S_0 but to hit S_0 repeatedly, to refine as quickly as possible our estimate of the objective-function value on S_0.

The set S_0 can be an arbitrary measurable subset of S, not just the set of global optimizers, and the Markov chain associated with the process does not

Received November 1994; revised April 1995.

[1]Work partially supported by a grant from the Air Force Office of Scientific Research and the Office of Naval Research, Contract F49620-90-C0033.

AMS 1991 *subject classifications*. 60J05, 60J10, 60K10, 90C10, 90C11, 90C26.

Key words and phrases. Simulated annealing, optimization, loop-skipping, coupling, noisy objective function.

have to arise from simulated annealing. We phrase our theorems in terms of generic probabilistic search, but our results are not limited to that setting.

The remainder of this paper is organized as follows. Section 2 gives notation and definitions and states our main result and a slight generalization. Sections 3 and 4 give examples where N is not a stopping time. Proofs of our main result and a generalization are in Sections 5 and 6.

2. The basic theorem. We are given a state space S. A probabilistic search algorithm on S is a (generally nonstationary) Markov chain $(X(1), X(2), \dots)$, where the probability distribution of $X(k+1)$ given $X(k)$ is computed according to some (generally time-dependent) rule, which we call the *standard move mechanism*. We modify the chain so that the rule can be overridden up to and including transition number N, assumed a.s. finite. The result is a new Markov chain $Y(\cdot)$, also generally nonstationary.

Let S_0 be an arbitrary measurable subset of S, not necessarily finite. Both our theorems require that the unmodified chain $X(\cdot)$ converge in probability to S_0 from any starting state at any initial time. More precisely,

$$\lim_{k\to\infty} \inf_{x\in s} \Pr(X(k) \in S_0 \mid X(n) = x) = 1, \tag{1}$$

for all positive integers n. Condition (1) holds when S is a compact subset of finite-dimensional Euclidean space and the density for proposed moves is uniformly bounded away from zero over S, along with certain other weak conditions detailed by Bélisle (1992). In fact, examination of his proof shows that, for finite S, it is enough that, from every state in S, there is a path to S_0 on which the objective function is nonincreasing. However, this condition is stronger than is generally assumed in finite-state simulated annealing. When S is finite, we can replace (1) by the simpler condition

$$\lim_{k\to\infty} \Pr(X(k) \in S_0 \mid X(n) = x) = 1, \tag{2}$$

for all states x in S and all positive integers n. Various sets of conditions which imply (2) when S is finite are given by Chiang and Chow (1988), Connors and Kumar (1989), Hájek (1988) and Tsitsiklis (1989). Thus, (1) and (2) are compatible with sufficient conditions in the simulated-annealing literature.

The epoch of the last override must be almost surely finite, but we do not require it to be a stopping time. It is enough that the override mechanism stops eventually and has no further influence on the destiny of the process. Denote by $(Y(k),\ k = 1, 2, \dots)$ the Markov chain of moves with overrides possible, and by N the epoch of the last override. We assume that there exists a family $\{Z(n, \cdot)\colon n \geq 0\}$ of Markov chains such that the following hold:

(i) The chains $Z(n, \cdot)$ and $Y(\cdot)$ coincide exactly up to time n; that is, $Z(n, k, \omega) = Y(k, \omega)$ for $k \leq n$ and every ω in the probability space.

(ii) After time n, the transition probabilities of $Z(n, \cdot)$ and $X(\cdot)$ are the same; that is,

$$\Pr(Z(n, k+1) = y \mid Z(n, k) = x) = \Pr(X(k+1) = y \mid X(k) = x),$$

for $k > n$ and $x, y \in S$.

(iii) There exists a random time N, a.s. finite, such that $Z(n, k, \omega) = Y(k, \omega)$ for all k on the set $\{N(\omega) \leq n\}$.

Here is the intuition for this structure. We start with the unmodified process and define a sequence of processes in which random overrides are allowed, but only up to a deterministic time n. Each process $Z(n, \cdot)$ couples with $Y(\cdot)$ up to time n and behaves stochastically like $X(\cdot)$ after time n. We can characterize $Y(j)$ as $\lim_{n \to \infty} Z(n, j)$. The limit is attained for some finite n, which can be chosen uniformly in j.

Here is our main result.

THEOREM 1. *Assume* (1) *and conditions* (i)–(iii). *Then* (1) *also holds for the chain* Y.

For reasons explained in Section 4, it is sometimes desirable to replace (ii) with some weaker condition. The following is one such:

(ii′) There exist constants $\eta_1, \eta_2, \ldots$ such that $\sum_{k=1}^{\infty} \eta_k < \infty$, and, for all x, y in S and all $k > n$ and for all n large enough,

$$\Pr(Z(n, k+1) = y \mid Z(n, k) = x) \geq (1 - \eta_k) \Pr(X(k+1) = y \mid X(k) = x).$$

It is not usually clear how to check that (ii′) holds.

THEOREM 2. *The conclusion of Theorem* 1 *continues to hold if we assume only* (i), (ii′) *and* (iii).

We prove Theorem 1 in Section 5, and Theorem 2 in Section 6.

3. Example 1 (Self-loop skipping). Fox (1993) introduces an algorithm QUICKER (stated at the end of this section). It prunes the sequence of states visited of self-loops of the form $x \to x \to \cdots \to x$, retaining only the first state in each self-loop, while preserving the convergence property (1) of the simulated annealing algorithm. Although QUICKER does require the evaluation of the objective function at every neighbor of the current state, Fox (1993, 1995a) for finite S and Fox (1995b, c) for general S give a rationale that the procedure might nevertheless be desirable. Here we use QUICKER, rather than the naive algorithm, as the standard move mechanism. Fox (1995a) and Heine (1994) show that its use can lead to a significant decrease in computer time over the naive algorithm. They make this notion precise and give sufficient conditions for it to occur.

To illustrate our theorem, we truncate QUICKER to QUICKER-j, stated below. Truncated QUICKER differs from the original version in that the former is forced to exit after at most a user-specified number j of iterations of its inner loop (in the original version, that number is unbounded stochasti-

cally). Heine (1994) shows how to pick j to make N, here the last time QUICKER takes more than j iterations, finite almost surely. In fact, Heine (1995) shows that under broad conditions we never need pick j larger than 2.

We cannot say that QUICKER-j shortens the time to first hit S_0. However, QUICKER-j shortens computer time per accepted move and is essential to synchronize parallel processors with SIMD architecture executing (stochastically) independent runs of simulated annealing. See Fox (1993) for details. Theorem 1 shows that (1) is preserved under QUICKER-j.

Now we state QUICKER-j, noting that QUICKER-∞ is QUICKER. Let

$$\alpha(x, k) = \Pr(Y(k+1) \neq x \mid Y(k) = x)$$

be the probability of accepting a proposed move to a new state y when at state x at step k, and let $G(x, k)$ be the trial number of the first success in iid Bernoulli trials, each with success probability $\alpha(x, k)$. The point of QUICKER is to generate

$$L = \arg \min_{l \geq i} \{Y(l) \neq Y(i)\} - 1,$$

the transition number on which Y next moves to a different state $Y(L+1)$, when it entered the current state x on move i. Given L, we generate $Y(L+1)$ with NEXT using the temperature at simulated time L; this is routine. Fox (1993, 1995a) proves that L is stochastically the same as what the naive algorithm would produce. Below, *Exit* means go to the next state. Inputs are the parameter j defined above, the current state x, the current simulated time t, the cooling schedule, and the neighborhood of x.

ALGORITHM (QUICKER-j).

Set $k \leftarrow t$
[*Set* $i \leftarrow 0$]
Until exit, *repeat*
  [*Set* $i \leftarrow i + 1$]
  Generate a geometric variate $G(x, k)$
  Set $L \leftarrow k + G(x, k) - 1$
  Set $V \leftarrow$ uniform (0, 1) r.v.
  If $V < \alpha(x, L)/\alpha(x, k)$ [or $i > j$] *then*
    Execute NEXT(x, L)
    Reset simulated time to $L + 1$
    Exit
  Else
    Set $k \leftarrow L + 1$
End

With the boxed statements deleted, QUICKER-j reduces to QUICKER.

4. Example 2 (Noisy objective function). Consider simulated annealing when the objective-function value at a state x must be estimated by a measurement $H(x)$ with expected value $f(x)$. Although f is unknown, we can observe $H(x)$ whenever we visit x. Perform simulated annealing with tentative-move and acceptance probabilities formed dynamically from some estimate $\hat{f}(\cdot)$, based on history of observed H-values. For simplicity, assume that S is finite. Other than the conditions below on $f(\cdot)$ and $\hat{f}(\cdot)$, we introduce no additional restrictions.

Until further notice, assume that there are only a finite number of possible values for $f(x)$ at any state x (e.g., a subset of the computer-representable numbers such as integer multiples of 0.001 between 0 and 1). We show at the end of this section how this requirement can be relaxed by using Theorem 2 in place of Theorem 1. In effect, until then we accept finite precision in our knowledge of $f(\cdot)$. We assume that the estimator $\hat{f}(\cdot)$ is consistent and that every state is visited infinitely often. [See Fox and Heine (1995) for a way to get a consistent estimator.] Hence, from some time N, a.s. finite, onward, the vector of $\hat{f}$-estimates coincides with the vector of true f-values. In this setting, the standard move mechanism is simulated annealing using the true f-values in the computation of the move probabilities. The modification uses the dynamically updated estimates $\hat{f}$. Our theorem says that if simulated annealing converges in probability to S_0 with $f(\cdot)$, it also converges with the consistent estimator $\hat{f}(\cdot)$.

Fox (1995b) shows how to use the setup of this section to embed an experimental-design problem in simulated annealing.

4.1. *Previous literature.* Yan and Mukai (1992) have previously considered such situations. They consider only the case where the matrix of proposed-move probabilities is symmetric. The probability of accepting a proposed move is itself a random variable, typically estimated by multiple measurements of the objective-function value of the tentative next state. The number of measurements required for each move required is large and grows without limit as the algorithm proceeds. No account is taken of measurements at previous visits.

Gelfand and Mitter (1989) study an approach similar to ours. They assume that the tentative moves are independent of objective-function values, that measurement errors are Gaussian and state-independent, and that the maximum noise variance at any state gets small sufficiently fast that its product with the logarithm of the probability of the move which increases objective-function value goes to zero. Like our condition (ii′) of Section 2, this last condition seems hard to check.

4.2. *Generalizing.* Neither of the above studies limit $f(s)$ to a finite set of values. As indicated above, we do not consider this an important restriction.

However, we indicate briefly how to remove it. The price is to replace (ii) with the weaker condition (ii'). Roughly speaking, (ii') says that the ratio of transition probabilities in the modified and unmodified chains must approach 1 sufficiently fast. Although something like (ii') seems to be needed, perhaps there are other, more easily verified conditions. If so, we are not aware of them.

If the estimators $\hat{f}(\cdot)$ are consistent, but not necessarily restricted to a finite range, then at a random but a.s. finite time N, $\|\hat{f}(x) - f(x)\| < \xi$, for all x, for any prescribed $\xi > 0$. Although N is not a stopping time, our Theorem 2 shows that this does not matter. We need to show that knowing the objective-function values to within an arbitrary ξ suffices. In most formulations of simulated annealing, this is quite straightforward. For example, Heine (1994) shows how convergence is implied by a theorem in Hájek (1988), assuming the conditions in that paper.

5. Proof of Theorem 1. Let $\varepsilon > 0$ be arbitrary. Let $\phi(m)$ be the smallest integer larger than m that satisfies

$$\inf_{x \in S} \Pr(X(t) \in S_0 \mid X(m) = x) > 1 - \varepsilon, \tag{3}$$

for all $t > \phi(m)$. We now show that the function $\phi(\cdot)$ is nondecreasing. Select $n > \phi(m+1)$. Then

$$\inf_{y \in S} \Pr(X(n) \in S_0 \mid X(m+1) = y) > 1 - \varepsilon. \tag{4}$$

Applying the Chapman–Kolmogorov equations, we have, for any x in S,

$$\begin{aligned} &\Pr(X(n) \in S_0 \mid X(m) = x) \\ &\quad = \sum_{y \in S} \Pr(X(n) \in S_0 \mid X(m+1) = y) \Pr(X(m+1) = y \mid X(m) = x) \\ &\quad \geq \inf_{y \in S} \Pr(X(n) \in S_0 \mid X(m+1) = y); \end{aligned}$$

this, together with (4), implies that $n > \phi(m)$.

Since $\phi(\cdot)$ is nondecreasing and grows without limit, there is, for $k > \phi(0)$, a unique integer $\phi^*(k)$ such that

$$\phi(\phi^*(k)) < k \leq \phi(\phi^*(k) + 1).$$

Note that the function $\phi^*(\cdot)$ is also nondecreasing and grows without limit.

Take any $k > \phi(0)$; write ν for $\phi^*(k)$. By condition (ii),

$$\begin{aligned} \Pr(Z(\phi^*(k), k) \in S_0) &\geq \inf_{x \in S} \Pr(Z(\nu, k) \in S_0 \mid Z(\nu, \nu) = x) \\ &= \inf_{x \in S} \Pr(X(k) \in S_0 \mid X(\nu) = x) \\ &> 1 - \varepsilon. \end{aligned} \tag{5}$$

To finish the proof, we use a variant of the coupling inequality [see, e.g., Asmussen (1987), page 143]. By conditions (i) and (iii), the chains $Z(\nu, \cdot)$ and $Y(\cdot)$ coincide exactly for all time on the set $\{N < \nu\}$. Thus

$$\begin{aligned} &\left|\Pr(Z(\nu,k) \in S_0) - \Pr(Y(k) \in S_0)\right| \\ &\quad = \left|\Pr(Z(\nu,k) \in S_0, N \geq \nu) - \Pr(Y(k) \in S_0, N \geq \nu)\right| \\ &\quad \leq \Pr(N \geq \nu). \end{aligned} \tag{6}$$

We have proved that, for arbitrary ε, for all large k,

$$\Pr(Y(k) \in S_0) > 1 - \varepsilon - \Pr(N \geq \phi^*(k)).$$

As remarked above, the function $\phi^*(\cdot)$ is nondecreasing and grows without limit; by condition (iii), this completes the proof. □

6. Proof of Theorem 2. Note first that (ii′) relates the j-step transition probabilities in $Z(n, \cdot)$ and $X(\cdot)$. In fact, for all n large enough,

$$\begin{aligned} &\Pr(Z(n,k+2) = y \mid Z(n,k) = x) \\ &\quad \geq (1 - \eta_k)(1 - \eta_k + 1) \\ &\qquad \times \sum_{z \in S} \Pr(X(k+2) = y \mid X(k+1) = z) \Pr(X(k+1) = z \mid X(k) = x) \\ &\quad = (1 - \eta_k)(1 - \eta_{k+1}) \Pr(X(k+2) = y \mid X(k) = x) \\ &\quad \geq \exp(-\eta_k - \eta_{k+1}) \Pr(X(k+2) = y \mid X(k) = x), \end{aligned}$$

and, in general,

$$\begin{aligned} &\Pr(Z(n,k+j) = y \mid Z(n,k) = x) \\ &\qquad \geq \exp\left(-\sum_{i=k}^{k+j-1} \eta_i\right) \Pr(X(k+j) = y \mid X(k) = x). \end{aligned}$$

Given ε, define $\phi(m)$ as in the proof of Theorem 2. Given δ, let M be large enough that $\exp(-\sum_{j=M-1}^{\infty} \eta_j) > 1 - \delta$. Proceed as in the proof of Theorem 1, except choose $k \geq \phi(M)$. This implies that $\phi^*(k) \geq M - 1$, so that

$$\begin{aligned} \Pr(Z(\phi^*(k),k) \in S_0) &\geq \inf_x \Pr(Z(\nu,k) \in S_0 \mid Z(\nu,\nu) = x) \\ &\geq \exp\left(-\sum_{j=\nu}^{k-1} \eta_j\right) \inf_x \Pr(X(k) \in S_0 \mid X(\nu) = x) \\ &\geq \exp\left(-\sum_{j=M-1}^{k-1} \eta_j\right)(1 - \varepsilon) \\ &> (1 - \delta)(1 - \varepsilon). \end{aligned}$$

Since δ and ε can be chosen arbitrarily small, the proof now goes through as before. □

Acknowledgments. We thank Paul Glasserman, Peter Glynn and the referees for helpful comments.

REFERENCES

ASMUSSEN, S. (1987). *Applied Probability and Queues*. Wiley, Chichester.

BÉLISLE, C. J. P. (1992). Convergence theorems for a class of simulated annealing algorithms on $\Re^d$. *J. Appl. Probab.* **29** 885–895.

ČERNY, V. (1985). A thermodynamical approach to the travelling salesman problem: an efficient simulation algorithm. *J. Optim. Theory Appl.* **45** 41–51.

CHIANG, T.-S. and CHOW, Y. (1988). On the convergence rate of annealing processes. *SIAM J. Control Optim.* **26** 1454–1470.

CONNORS, D. P. and KUMAR, P. R. (1989). Simulated annealing type Markov chains and their order balance equations. *SIAM J. Control Optim.* **27** 1440–1461.

FOX, B. L. (1993). Integrating and acclerating tabu search, simulated annealing, and genetic algorithms. *Ann. Oper. Res.* **41** 47–67.

FOX, B. L. (1995a). Faster simulated annealing. *SIAM J. Optim.* **5** 488–505.

FOX, B. L. (1995b). Simulated annealing: folklore, facts, and directions. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing. Lecture Notes in Statist.* **106** 17–48. Springer, Berlin.

FOX, B. L. (1995c). Uprooting simulated annealing and genetic algorithms. *SIAG/OPT: Views and News* **7** 5–9.

FOX, B. L. and HEINE, G. W. (1995). Stable estimators for self-adjusting simulations. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing. Lecture Notes in Statist.* **106** 205–212. Springer, Berlin.

GELFAND, S. B. and MITTER, S. K. (1989). Simulated annealing with noisy or imprecise measurements. *J. Optim. Theory Appl.* **69** 49–62.

HÁJEK, B. (1988). Cooling schedules for optimal annealing. *Math. Oper. Res.* **13** 311–329.

HEINE, G. W. (1994). Smart simulated annealing. Ph.D. dissertation, Dept. Mathematics, Univ. Colorado, Denver.

HEINE, G. W. (1995). Convergence properties for a loop-skipping algorithm for simulated annealing. Technical report, Univ. Colorado, Denver.

KIRKPATRICK, S., GELATT, C. D. and VECCHI, M. P. (1983). Optimization by simulated annealing. *Science* **220** 621–680.

TSITSIKLIS, J. (1989). Markov chains with rare transitions and simulated annealing. *Math. Oper. Res.* **14** 70–90.

YAN, D. I. and MUKAI, H. (1992). Stochastic discrete optimization. *SIAM J. Control Optim.* **30** 594–612.

DEPARTMENT OF MATHEMATICS
CAMPUS BOX 170
UNIVERSITY OF COLORADO
P.O. BOX 173364
DENVER, COLORADO 80217-3364