J Heuristics (2008) 14: 627–654
DOI 10.1007/s10732-007-9058-7

# Simulated annealing in the presence of noise

**Jürgen Branke · Stephan Meisel · Christian Schmidt**

Received: 10 January 2006 / Revised: 19 March 2007 / Accepted: 26 March 2007 /
Published online: 25 October 2007


**Abstract** In many practical optimization problems, evaluation of a solution is subject to noise, e.g., due to stochastic simulations or measuring errors. Therefore, heuristics are needed that are capable of handling such noise. This paper first reviews the state-of-the-art in applying simulated annealing to noisy optimization problems. Then, two new algorithmic variants are proposed: an improved version of stochastic annealing that allows for arbitrary annealing schedules, and a new approach called *simulated annealing in noisy environments* (SANE). The latter integrates ideas from statistical sequential selection in order to reduce the number of samples required for making an acceptance decision with sufficient statistical confidence. Finally, SANE is shown to significantly outperform other state-of-the-art simulated annealing techniques on a stochastic travelling salesperson problem.

**Keywords** Simulated annealing · Uncertainty · Noise · Sequential sampling

## 1 Introduction

Many real-world optimization problems are noisy. In such problems, a solution's quality cannot be determined accurately, but has to be estimated via the realization of a random variable. The sources of noise can be manifold, including optimization

J. Branke (✉) · C. Schmidt
Institute AIFB, University of Karlsruhe, 76128 Karlsruhe, Germany
e-mail: branke@aifb.uni-karlsruhe.de

C. Schmidt
e-mail: csc@aifb.uni-karlsruhe.de

S. Meisel
Dept. of Management Studies, Univ. of Braunschweig, 38106 Braunschweig, Germany
e-mail: stephan.meisel@tu-braunschweig.de

based on randomized simulations, measurement errors, stochastic sampling, and interaction with users. In this paper, we examine the use of simulated annealing (SA) for noisy optimization problems, and propose some modifications to make the algorithm more efficient and more flexible.

SA is a randomized search heuristic inspired by thermodynamics and the process of annealing in crystallography or metallurgy. In short, it can be regarded as a stochastic hill-climber, where the probability to accept a move depends on the fitness difference between current and candidate solution, and a temperature parameter. The temperature is decreased (annealed) over the run, making it less and less likely to accept moves to inferior solutions. A more detailed discussion of simulated annealing is provided in Sect. 2.

Generally, noise is considered a major challenge for optimization. It influences the perceived quality difference between current and candidate solution, and thereby changes the acceptance probabilities of the algorithm. If, for example, the candidate solution is evaluated better than it actually is, and/or the current solution is evaluated worse than its true quality, the probability to move from current to candidate solution is much higher than the annealing schedule suggests. With decreasing signal-to-noise ratio, the true fitness values have less and less influence, up to the point when search is reduced to a random walk in the search space.

For most noisy optimization problems, the uncertainty in fitness evaluation can be reduced by sampling a solution's fitness several times and using the average as estimate for the true mean fitness. Sampling $n$ times reduces the variance of a random variable's estimated mean by a factor of $n$, but on the other hand increases the computation time also by a factor of $n$. Thus, there is a generally perceived trade-off: either one can use relatively exact estimates but only run the algorithm for a small number of iterations (because a single estimate requires many evaluations), or one can let the algorithm work with relatively crude quality estimates, but allow for more iterations (as each estimate requires less effort).

In this paper, we propose a new SA variant that can deal more efficiently with noise. It relies on two basic observations.

1. SA is a randomized search algorithm, which deliberately introduces randomness in the acceptance step by accepting a move only with a certain probability. In particular in the early phases of the run, this kind of uncertainty is considered helpful in order to avoid local optima. One idea of this paper is to accept the noise in the optimization problem and (at least partially) use it to replace the deliberate randomness in the algorithm. In other words, noise is partially neutralized by making the SA algorithm more deterministic.
2. When the signal-to-noise ratio is sufficiently large, the noise has little influence on the acceptance probability. Therefore, we propose to use a sequential sampling scheme: Generate a few samples, estimate the signal-to-noise ratio, and then resample until the decision whether to accept the new candidate solution can be made with sufficient confidence.

By integrating these ideas (and some others) into SA, the algorithm can be made much more efficient.

The paper is structured as follows. In the next section, we describe the fundamental principles of simulated annealing that are important for this paper. Related work is

reviewed in Sect. 3. Section 4 presents in more detail the two methods which form the foundation for our new approach. We present our new SA variants in Sect. 5. In Sect. 6, the new algorithm is empirically compared with several previous SA variants for noisy optimization problems on a simple stochastic travelling salesman problem. The paper is concluded with a summary and some ideas for future work.

## 2 Simulated annealing and the thermodynamical equilibrium

### 2.1 Background from thermodynamics

As mentioned above, SA has been motivated by physical annealing processes in crystallography or metallurgy, where a solid is melted and then cooled down slowly in order to obtain perfect molecular structures, which correspond to a state of minimum energy. According to statistical mechanics, a system is said to be in a thermodynamical equilibrium, if its probability of being in a given state $x$ follows the Boltzmann-Gibbs-Distribution:

$$P(x) \propto e^{-\beta E(x)} \tag{1}$$

where $E(x)$ is the energy level of state $x$, and $\beta = 1/(c_B T)$ is a factor depending on the temperature, $T$, and the Bolzmann constant, $c_B$. A sufficient condition for equilibrium is the so-called detailed balance condition,

$$P_{tr}(x_1 \to x_2) P(x_1) = P_{tr}(x_2 \to x_1) P(x_2) \quad \forall x_1, x_2 \tag{2}$$

with $P_{tr}(x_1 \to x_2)$ denoting the probability that the state changes from $x_1$ to $x_2$. Small transformations yield

$$\frac{P_{tr}(x_1 \to x_2)}{P_{tr}(x_2 \to x_1)} = \frac{P(x_2)}{P(x_1)} = \frac{e^{-\beta E(x_2)}}{e^{-\beta E(x_1)}} = e^{-\beta \Delta E} \tag{3}$$

with

$$\Delta E = E(x_2) - E(x_1). \tag{4}$$

As has been noted first by Kirkpatrick et al. (1983), the process of physical annealing can be regarded to correspond to the process of optimization. According to this analogy, a solution in the search space corresponds to a molecular structure of the material, and the solution's quality to the energy level of the molecular structure. Using the above results from thermodynamics, removing the Boltzmann constant, and substituting the transition probability $P_{tr}$ by the product of the probability to attempt a move, $P_m$, and the probability to accept a move, $P_a$, we get

$$\frac{P_a(\Delta E) P_m(x_1 \to x_2)}{P_a(-\Delta E) P_m(x_2 \to x_1)} = e^{-\Delta E/T}. \tag{5}$$

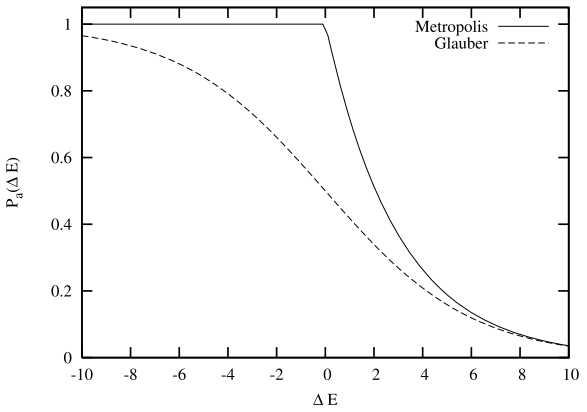


**Fig. 1** Acceptance probabilities depending on energy difference according to the Metropolis and Glauber criteria

Assuming that $P_m(x_1 \to x_2) = P_m(x_2 \to x_1)$, this simplifies to

$$\frac{P_a(\Delta E)}{P_a(-\Delta E)} = e^{-\Delta E/T}. \tag{6}$$

There are different possible acceptance probability functions $P_a$ satisfying Eq. 6. Metropolis et al. (1953) suggested the so-called *Metropolis* criterion, which is defined as

$$P_a^{\text{Metropolis}}(\Delta E) = \begin{cases} 1 & : \quad \Delta E \leq 0, \\ e^{-\Delta E/T} & : \quad \Delta E > 0. \end{cases} \tag{7}$$

Another valid acceptance probability function is based on the work of Glauber (1963):

$$P_a^{\text{Glauber}}(\Delta E) = \frac{1}{1 + e^{\Delta E/T}}. \tag{8}$$

The two acceptance criteria are visualized in Fig. 1 for $T = 3$.

## 2.2 Simulated annealing

SA is an iterative randomized search heuristic based on the above physical principles. The basic structure is described in Algorithm 1. The algorithm keeps track of a current solution $x_c$. In every iteration, a new candidate solution $x_n$ is generated in the neighborhood of the current solution, $N(x_c)$. Based on the energy difference $\Delta E = E(x_n) - E(x_c)$ and the temperature $T$, and according to the acceptance probability $P_a(\Delta E)$, either the candidate solution replaces the current solution, or is rejected. During the search, the temperature, and thus the probability to accept worse candidate solutions, is gradually reduced.

**Algorithm 1** Simulated annealing

$T \leftarrow T_0$
Generate initial solution $x_c$
**repeat**
  **repeat**
    Generate $x_n \in N(x_c)$ // generate new solution
    $\Delta E \leftarrow (E(x_n) - E(x_c))$
    $u \leftarrow U[0:1]$ // draw a random number
    **if** $u < P_a(\Delta E)$
      $x_c \leftarrow x_n$ // accept new solution
    **end if**
  **until** time to reduce temperature
  $T \leftarrow r(T)$
**until** termination condition met

There are a number of possible design choices, such as:

1. The *starting solution*, which is usually generated at random.
2. The *neighborhood*, which is usually assumed to be symmetric (i.e., the probability to attempt a move from $x_1$ to $x_2$ is equal to the probability to attempt a move from $x_2$ to $x_1$). For a discussion of the effect of neighborhood structure in SA, see e.g., Cheh et al. (1991).
3. The *acceptance probability* $P_a(\Delta E)$. Here, usually the Metropolis criterion is used. In principle, it has been shown that the higher the probability to accept a move, the faster the algorithm's convergence to the thermodynamical equilibrium (Hastings 1970; Hertz et al. 1991). In this sense, the Metropolis criterion is preferable to the Glauber criterion.
4. The *annealing schedule*, consisting of an initial temperature, and rules to govern the change of the temperature during the search. An often used rule here is geometric annealing, i.e., the multiplication of the temperature with a factor $0 < \alpha < 1$, or $r(T) = \alpha T$. Such a schedule does not in general guarantee convergence to the optimum, but is very useful in practice because it allows convergence in reasonable time.
5. The *termination criterion*, e.g., the number of iterations.

There is a wealth of publications on the above design choices and how to set them appropriately. However, the details of these design choices are not of primary relevance to this paper. We henceforth assume that appropriate design choices have been made, and only attempt to implement these design choices most efficiently in a noisy environment. For this reason, the reader is referred to the literature for details on how to make the above design decisions (e.g., van Laarhoven and Aarts 1987; Rees and Ball 1987; Nourani and Andresen 1998).

### 2.3 Simulated annealing disturbed by noise

Obviously, noisy evaluations affect the determination of the energy difference $\Delta E$. If evaluations are disturbed by a Gaussian noise with mean zero and variance $\sigma_E^2$, then

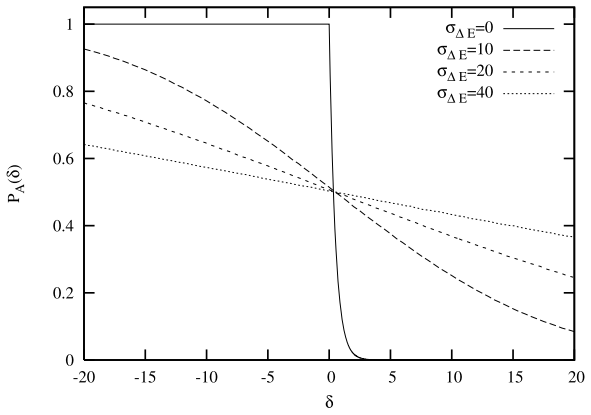


**Fig. 2** Actual acceptance probabilities for different levels of noise and $T = 0.5$

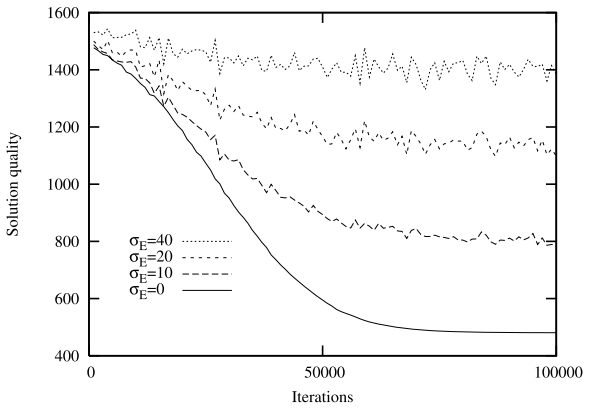


**Fig. 3** Convergence curves for different levels of noise

the observed energy difference is also a random variable

$$\Delta E \sim N(\delta, \sigma^2_{\Delta E}), \tag{9}$$

where $\delta$ is the true (undisturbed) energy difference and $\sigma^2_{\Delta E} = 2\sigma^2_E$ assuming the disturbances are independent. Then, the *true* acceptance probability, $P_A(\delta)$ can be calculated as

$$P_A(\delta) = \int_{x=-\infty}^{x=\infty} P_a(x) f_{\Delta E}(x|\delta) dx, \tag{10}$$

where $f_{\Delta E}(x|\delta)$ is the pdf of the observed (estimated) energy difference $\hat{\delta}$ given a true difference $\delta$. Figure 2 visualizes the effect of noise for the Metropolis acceptance criterion.

Figure 3 compares the convergence curves for SA on a TSP problem with different levels of noise, showing the true (undisturbed) quality of the current solution over iterations. For the plot, the curves have been averaged over 300 independent runs. We observe that noise has a deteriorating effect on the performance of SA. Specifically,

1. The stronger the noise, the more rugged is the convergence curve. Noise causes the algorithm to often regard an inferior solution as best and leads to more random decisions.
2. The stronger the noise, the slower the convergence.
3. The stronger the noise, the worse the final solution quality given a fixed number of iterations.

## 3 Related work

This section surveys previous work on applying SA to noisy environments. Three main categories are identified and discussed in turn: papers that deal mostly with convergence properties, papers that adapt the number of samples per solution based on some statistical analysis, and papers that modify the acceptance function in an attempt to maintain thermodynamical equilibrium. Some additional approaches are summarized at the end of this section.

### 3.1 Convergence and classical approaches

There are several papers looking at convergence of SA in noisy environments. All of them show that SA is able to converge to the optimum under certain conditions, usually including sufficient time and decreasing noise over time (e.g., due to an increasing number of evaluations per solution).

Kushner (1987) was the first to examine the behavior of SA when the value of the objective function can only be sampled via the Monte Carlo method. He presents a general approach by specifying a system which represents a combination of a stochastic approximation method and SA. Based on the theory of large deviations and assuming Gaussian noise he gives a full asymptotic analysis of the system under suitable conditions.

Gelfand and Mitter (1989) consider discrete search spaces and use Markov chains to prove that SA converges in probability to the globally optimal solutions, assuming that the standard deviation of the noise decreases linearly with the temperature (i.e., $\sigma_E(k) = o(T_k)$), which means that for constant noise, the number of samples per solution has to increase quadratically with the inverse of the temperature.

Gutjahr and Pflug (1996) look at the same setting as Gelfand and Mitter (1989). They specify an annealing schedule for which they prove convergence in probability if the standard deviation decreases at least inversely proportional to the number of iterations, i.e., $\sigma_E(k) = O(k^{-\gamma})$ where $\gamma > 1$. Also, they generalize the convergence proof to any noise distribution that is symmetric and more peaked around zero than the normal distribution.

Fox and Heine (1995) show that SA converges in probability if errors in acceptance decisions only occur during some initial transition time, and argue that this is true also in a noisy environment when each solution is visited and re-evaluated infinitely often, and the average over all evaluations so far is used as fitness estimate.

The SA variant proposed in Alrefaei and Andradóttir (1999) works different from the standard method. At termination, instead of returning the last solution visited, it

returns either the solution visited most often throughout the run or (with better performance) the solution with the best average fitness over all samples taken. Because of this, the algorithm doesn't require convergence of the underlying Markov chain and the authors suggest to use a constant temperature instead of an annealing process. They prove that, under the condition of decreasing noise, the procedure almost surely converges to the set of optimal solutions. The algorithm is compared to approaches of Gutjahr and Pflug (1996), Gelfand and Mitter (1989), Fox and Heine (1995) and seems to perform best on the test problems considered.

Alkhamis et al. (1999) also provide a modified SA algorithm and show its convergence in probability under suitable conditions. They take into account the uncertainty of evaluation by suggesting to base the acceptance decisions on the lower bound of a confidence interval instead of $\Delta E$. That way, the higher the uncertainty about the true quality difference, the lower the probability to accept a new solution. The authors also provide a convergence proof which is closely related to that of Gutjahr and Pflug (1996) but slightly more general. Test cases seem to demonstrate that the new SA variant based on confidence intervals outperforms the approach of Gutjahr and Pflug (1996).

Alkhamis and Ahmed (2004) is more or less a combination of two other approaches. It uses a constant temperature, returns the solution visited most often as in Alrefaei and Andradóttir (1999), and also uses the lower bound of a confidence interval to make acceptance decisions as in Alkhamis et al. (1999).

Prudius and Andradottir (2005) present two algorithms that are also based on existing approaches (Alrefaei and Andradóttir 1999; Fox and Heine 1995) and prove that they almost surely converge to the optimal solution. Both algorithms require an annealing schedule with decreasing temperature and return the solution showing the best average fitness with respect to all the observations throughout the run as final solution. The main difference between these two approaches concerns the decision about the acceptance of candidate solutions. Either it is based on all the samples available so far or on only the samples collected in the current iteration.

A drawback of the approaches of Alrefaei and Andradóttir (1999), Alkhamis and Ahmed (2004), Fox and Heine (1995) and Prudius and Andradottir (2005) is that they only work on discrete search spaces and require a lot of memory, having to store fitness values for *all* visited solutions.

### 3.2 Sample as long as necessary

The approaches above either store all evaluated solutions so far, or increase the number of evaluations over time according to a predefined schedule. The approaches discussed in this subsection attempt to improve fitness estimates as much as needed by repeatedly sampling a solution's quality only until some stopping criterion is met.

Bulgak and Sanders (1988) use SA to optimize buffer sizes in an automatic assembly system. The basic idea is to evaluate the *candidate* solution until the energy difference is statistically significant. As a criterion for statistical significance Bulgak and Sanders consult a confidence interval for $E(x_c)$. For the annealing they use the schedule proposed in Hajek (1988) but add that the convergence proof associated with this schedule for deterministic problems does not apply to the given stochastic system.

Ahmed and Alkhamis (2002) combine SA with the ranking and selection procedure for solving discrete simulation-based optimization problems. Their algorithm divides the sampling for each solution into two stages. The number of samples taken in the second stage depends on two parameters as well as on the mean solution quality that results from the first stage. The method resembles the proposal of Alrefaei and Andradóttir (1999) in that it returns the solution visited most often throughout the optimization run. Ahmed and Alkhamis assume Gaussian noise with unknown variance and prove that their procedure converges almost surely to the global optimal solution. They test their method on two problems but provide no comparisons with other approaches.

### 3.3 Adapt acceptance function, maintain equilibrium

Another interesting idea used in several papers is to modify the acceptance function such that it fulfills the detailed balance equation (Eq. 2) despite the noise.

Kennedy and Kuti (1985) allow unknown variance and arbitrary noise distribution. They derive a parameterized acceptance function which is different from the original Metropolis function but which, depending on the value of the parameter, fulfills the detailed balance equation. Two kinds of difficulties are associated with this approach: On the one hand an appropriate parameter value must be chosen which depends on the (a priori usually unknown) range of $\Delta E$. On the other hand, the wider this range, the lower the resulting acceptance rate and the less efficient the optimization runs will be. This means that in a practical context the method usually implies a tradeoff between runtime and adherence to the detailed balance equation.

Ceperley and Dewing (1999) develop an acceptance function under the assumption of normally distributed noise. They argue that this method is more efficient than Kennedy and Kuti (1985) by providing better acceptance rates especially at higher noise levels. Ceperley and Dewing specify an acceptance function for known variance and also derive a similar function allowing for estimated variances. In the absence of noise ($\sigma^2_{\Delta E} = 0$), the given function equals the Metropolis criterion. But the larger the variance and the lower the temperature, the more the resulting acceptance probabilities differ from Metropolis and provide worse acceptance rates. Bowler (2001) provides an alternative derivation of the result of Ceperley and Dewing (1999). More details on this approach are provided in Sect. 4.2.

Fink (1998) presents a simple yet powerful threshold acceptance criterion called *stochastic annealing*. It assumes normally-distributed noise and requires a known variance in order to control the temperature. Stochastic annealing doesn't meet the thermal equilibrium exactly, but results in acceptance probabilities very close to the Glauber criterion. This approach is also discussed in Ball et al. (2003a) and Bowler (2001), along with the method from Ceperley and Dewing (1999) and yet another acceptance criterion for arbitrary noise and unknown variance which is simpler than Kennedy and Kuti (1985), but suffers from similar difficulties. In both of these papers, only stochastic annealing is actually applied to an optimization problem, suggesting that it is considered superior in practice by these authors. In Ball et al. (2003b), stochastic annealing has been applied to the probabilistic travelling salesman problem. Stochastic annealing is also discussed in more detail in Sect. 4.1.

### 3.4 Others

Haddock and Mittenthal (1992) investigate the feasibility of using SA in conjunction with a simulation model to determine the optimal system parameters. They apply the Glauber acceptance function and consider the single simulation results as point estimates for the true solution qualities. Unfortunately the considered problem is of small size and no theoretical treatment is given.

Painton and Diwekar (1995) extend the evaluation function by adding a penalty term depending on the number of samples taken. Penalty is increased with decreasing temperature, thus in the long run favoring solutions evaluated based on many samples. The number of samples used for comparison is decided randomly for each tournament and previously evaluated samples are not re-used. Painton and Diewekar base their analysis on Latin Hypercube Sampling instead of common Monte Carlo sampling and analyze the effect of different noise distributions.

Wang and Zhang (2006) assume Gaussian noise with unknown variance and use a hypothesis test to figure out whether the current solution and the candidate solution are significantly different or not. If they are not, Wang and Zhang propose to discard the candidate solution and generate a new one. Otherwise they apply the Metropolis criterion with the estimated mean solution qualities.

## 4 Base procedures

The methods we propose in this paper are partially based on the approaches of Fink (1998) and Ceperley and Dewing (1999), thus these are presented in detail in this section.

### 4.1 Stochastic annealing

This procedure has originally been proposed in Fink (1998). It is based on the observation that a cumulative normal distribution with appropriately chosen variance looks very similar to the Glauber acceptance criterion. Thus, Fink suggests to simply accept the new candidate solution if the observed (mean) fitness difference $\hat{\delta}$ is smaller than zero, i.e.

$$P_a(\hat{\delta}) = \begin{cases} 1 & : \quad \hat{\delta} \leq 0, \\ 0 & : \quad \hat{\delta} > 0. \end{cases} \tag{11}$$

Due to the noise, sometimes $\hat{\delta} > 0$ is observed even if the candidate solution is better ($\delta \leq 0$), and vice versa. The true acceptance probability can be calculated as

$$P_A(\delta) = \text{Prob}(\Delta E \leq 0). \tag{12}$$

Assuming a normally distributed observed fitness difference $\Delta E \sim N(\delta, \sigma^2_{\Delta E})$, this results in

$$P_A(\delta) = \Phi\left(\frac{-\delta}{\sigma_{\Delta E}}\right). \tag{13}$$

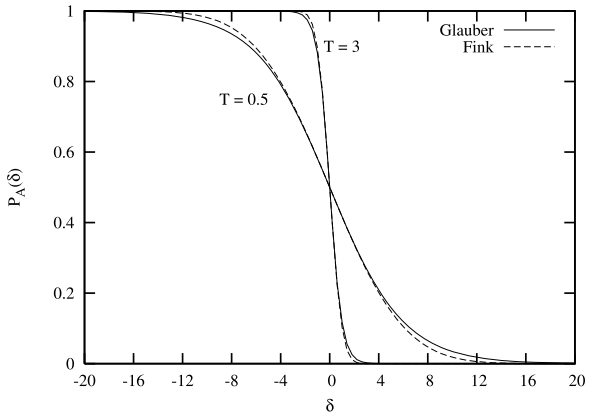


**Fig. 4** Acceptance probabilities for different temperatures and $\sigma^2_{\Delta E} = 50$

Note that the standard deviation of the observed fitness difference $\sigma_{\Delta E}$ can be reduced by increasing the number of samples, $n$. Fink derives a correlation between the temperature $T$ and the number $n$ of samples per solution.

$$\frac{1}{T} \approx \sqrt{\frac{8n}{\pi \sigma^2_{\Delta E}}}. \tag{14}$$

Figure 4 compares the true acceptance probabilities obtained according to the Glauber function without noise, and for stochastic annealing with noise and using a number of samples according to Eq. 14. As can be seen, they match quite closely, with stochastic annealing accepting the candidate solution slightly less often if it is worse, and slightly more often if it is better than the current solution.

Despite its simplicity, the approach is quite powerful. Basically, instead of trying to reduce the noise in evaluation, it reduces the uncertainty inherent in the stochastic acceptance criterion of SA by making acceptance deterministic. Then it utilizes this noise in evaluation to generate the desired level of uncertainty for the true acceptance probability. As a result, a certain level of noise can be tolerated without practically affecting the algorithm.

Nevertheless, the approach has a number of drawbacks which we want to address in this paper:

1. Since a solution has to be evaluated at least once ($n \geq 1$), there is a maximum temperature that can be implemented with this method.
2. Since the number of samples, $n$, is a discrete variable, so are the temperatures that can be implemented by this method. For example, when switching from $n = 1$ to $n = 2$ in the above example, the temperature is changed from $T = 4.43$ to $T = 3.13$. The annealing schedule is thus more governed by the granularity of $n$ rather than by the designer's choices. Fink (1998) propose to use $n_k \propto k^2$, i.e., $T_k \propto 1/k$ as annealing schedule, with $k$ denoting the stage of the algorithm. Compared to the most common geometric annealing schedule $T_k \propto \alpha^k$, this results in a faster reduction of temperature in the early stages, and a slower reduction later on, see Fig. 5. In Ball et al. (2003b), the authors mention as future work that instead of

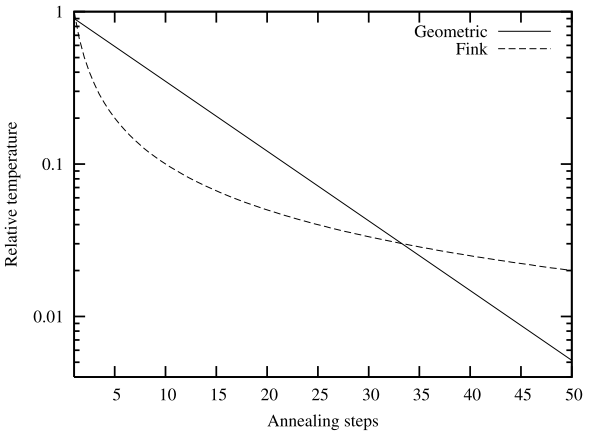


**Fig. 5** Typical geometric annealing schedule ($T = 0.9^{(k-1)}$) and typical schedule as proposed by Fink (1998) ($T = 1/k$)

the used schedule to increase $n$, "smoother annealing can be obtained by directly controlling the effective temperature", but do not provide any means to do so.

3. As the temperature decreases, the number of samples required increases drastically. In the above example, for a temperature of $T = 0.1$, the approach suggests to draw $n = 1963$ samples per solution, i.e., 3926 samples in each iteration.

### 4.2 Acceptance criterion by Ceperley and Dewing

Ceperley and Dewing (1999) proposed the following acceptance criterion:

$$P_a(\hat{\delta}) = \begin{cases} 1 & : \quad \hat{\delta} \leq -\frac{1}{2}\sigma^2_{\Delta E}/T, \\ e^{-(\hat{\delta}/T + \frac{1}{2}\sigma^2_{\Delta E}/T^2)} & : \quad \hat{\delta} > -\frac{1}{2}\sigma^2_{\Delta E}/T. \end{cases} \tag{15}$$

This acceptance criterion almost fulfills the detailed balance equation (Eq. 2) despite the noise, and without requiring additional evaluations. For $\sigma^2_{\Delta E} = 0$ it is equivalent to the Metropolis criterion. Figure 6 visualizes the true acceptance probability for the CD acceptance criterion and different levels of noise. The main drawback of this acceptance criterion is that for low temperatures, a new solution is only accepted if it is perceived as much better than the current solution, see Fig. 7. Figure 8 demonstrates the same issue by plotting the probability to accept a new solution if it is perceived as equally good (dashed line) or even better by 1 unit (solid line) depending on the temperature. As can be seen, the probability to accept the new solution becomes negligible as the temperature falls below 1.5, even if the new solution has a perceived better solution quality. As has been discussed in Sect. 2.2, the resulting extremely low acceptance probabilities slow down the algorithm reaching an equilibrium state, up to the point where it basically comes to a complete halt.

The time to reach thermodynamical equilibrium depends on the time between accepted moves. The expected number of evaluations between accepted moves, $E(m)$, can be calculated as the number of samples per iteration divided by the true accep-

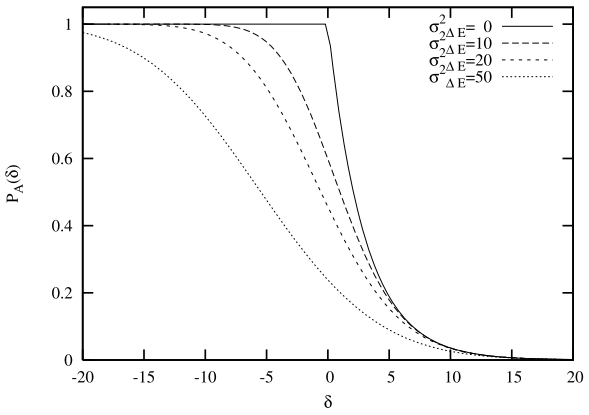


**Fig. 6** Acceptance probability function $P_A(\delta)$ for the approach of Ceperley and Dewing, for $T = 3$ and different levels of noise

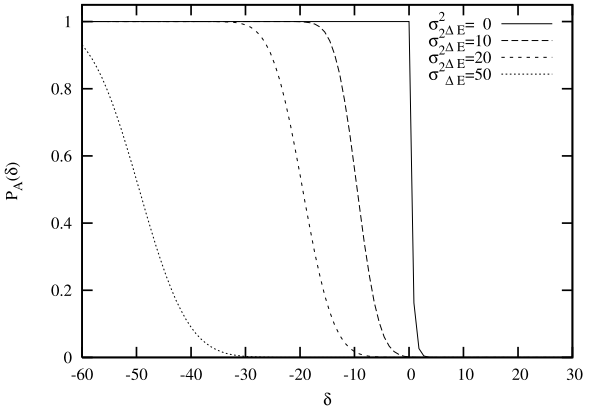


**Fig. 7** Acceptance probability function $P_A(\delta)$ for the approach of Ceperley and Dewing, for $T = 0.5$ and different levels of noise

tance probability. For CD, this can be calculated as Ceperley and Dewing (1999)

$$E^{CD}(m) = 4\left(e^{-(\delta/T)}\mathrm{erfc}\left(\frac{\sigma^2_{\Delta E}/(2T^2) - \delta/T}{\sqrt{2\sigma^2_{\Delta E}/T^2}}\right) + \mathrm{erfc}\left(\frac{\sigma^2_{\Delta E}/(2T^2) + \delta/T}{\sqrt{2\sigma^2_{\Delta E}/T^2}}\right)\right)^{-1}, \tag{16}$$

where $\mathrm{erfc}(x)$ is the complementary error function. For stochastic annealing, combining Eqs. 14 and 13 leads to

$$E^{\text{stoch. ann.}}(m) = \frac{\pi\sigma^2_{\Delta E}}{4T^2\Phi(-\delta/T\sqrt{\pi/8})}. \tag{17}$$

Figure 9 shows the number of samples required to accept a move to an equally good solution ($\delta = 0$) depending on the ratio $\sigma^2/T^2$. Except for a small range around $\rho = 6$, CD requires clearly more function evaluations to accept a move (note that $\sigma^2/T^2 \geq 8/\pi$ for stochastic annealing to be applicable).

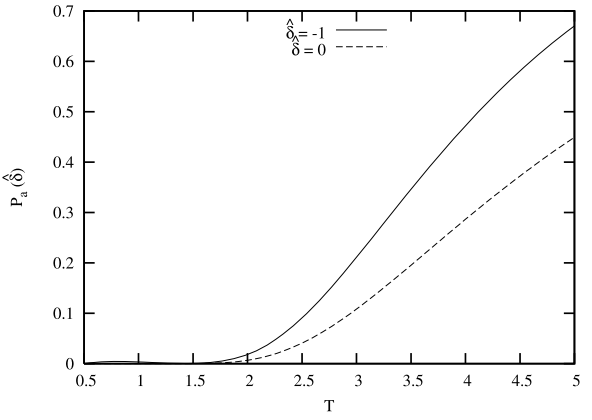


**Fig. 8** Probability to accept the new solution depending on the temperature and the perceived quality difference $\hat{\delta}$. $\sigma^2_{\Delta E} = 40$

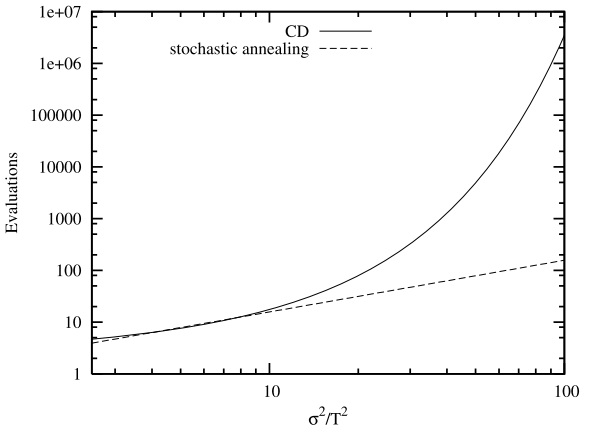


**Fig. 9** Expected number of samples to required to accept a move to an equally good solution ($\delta = 0$), depending on the ratio $\sigma^2/T^2$

## 5 New variants of simulated annealing

In this section, we present two new SA variants. First, in the next subsection we propose some extensions of stochastic annealing that allow for arbitrary annealing schedules. Then, we present our new simulated annealing for noisy environments (SANE) which combines the method of Ceperley and Dewing for higher temperatures with a threshold-based sequential sampling scheme for lower temperatures.

### 5.1 Improved stochastic annealing

As has been noted above, stochastic annealing has some restrictions regarding the annealing schedule used.

The maximal temperature in stochastic annealing is reached when $n = 1$ and is defined by $1/T_{\max} = \sqrt{8/(\pi\sigma^2_{\Delta E})}$. In order to allow higher temperatures, we suggest here to artificially increase the noise in the evaluation such that the resulting temperature corresponds to the desired temperature. We achieve this by adding a random

variable to the fitness difference and using $\Delta E' = \Delta E + H$ for acceptance, where $H$ is distributed according to

$$H \sim N(0, \sigma_H^2), \quad \text{and} \tag{18}$$

$$\sigma_H^2 = \frac{8T^2}{\pi} - \sigma_{\Delta E}^2. \tag{19}$$

The variance of this random variable, $\sigma_H^2$, has been derived by solving Eq. 14 with $n = 1$ for $\sigma_{\Delta E}^2$ (the required variance for the given temperature) and then subtracting the actual variance.

In principle, this method of artificially adding noise could also be used to generate arbitrary temperature levels that lie between those defined by two adjacent sample numbers $n$ and $n + 1$. However, the method only allows us to artificially *increase* the temperature, never to *decrease* it. For a desired temperature between the levels defined by $n$ and $n + 1$, this would mean that always $n + 1$ samples have to be used. In the following, we present a more efficient method to reach arbitrary temperatures that would correspond to a non-integer $n$ according to Eq. 14.

In particular, we propose to use $\lfloor n \rfloor$ samples with probability $p$, and $\lceil n \rceil$ samples with probability $(1 - p)$. The actual acceptance probability in stochastic annealing, depending on the number of samples, is

$$P_A(\delta) = \Phi\left(\frac{-\delta\sqrt{n}}{\sigma_{\Delta E}}\right). \tag{20}$$

Hence, for arbitrary (non-integer) $n$ derived from Eq. 14, we would like to choose $p$ such that

$$\Phi\left(\frac{-\delta\sqrt{n}}{\sigma_{\Delta E}}\right) = p\Phi\left(\frac{-\delta\sqrt{\lfloor n \rfloor}}{\sigma_{\Delta E}}\right) + (1-p)\Phi\left(\frac{-\delta\sqrt{\lceil n \rceil}}{\sigma_{\Delta E}}\right) \tag{21}$$

or

$$p = \frac{\Phi\left(\frac{-\delta\sqrt{n}}{\sigma_{\Delta E}}\right) - \Phi\left(\frac{-\delta\sqrt{\lceil n \rceil}}{\sigma_{\Delta E}}\right)}{\Phi\left(\frac{-\delta\sqrt{\lfloor n \rfloor}}{\sigma_{\Delta E}}\right) - \Phi\left(\frac{-\delta\sqrt{\lceil n \rceil}}{\sigma_{\Delta E}}\right)}. \tag{22}$$

This equation would require to know the true difference between the two solutions, $\delta$, which is not known. However, we know that we have to sample at least $\lfloor n \rfloor$ times, and then we can estimate $\delta$ by the observed fitness difference $\hat{\delta}$. As Fig. 10 demonstrates, even for small $n$ this approximation is very good.

Thus, with the above two ideas we are able to implement arbitrary annealing schedules based on the ideas of stochastic annealing, which allows us to use annealing schedules that have been shown to work well.

### 5.2 A new simulated annealing for noisy environments (SANE)

In this section, we propose a new SA variant that improves previous approaches in several aspects. The algorithm combines ideas from stochastic annealing and the

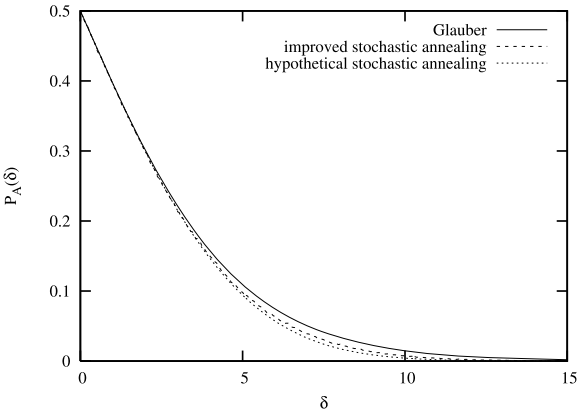


**Fig. 10** Acceptance probability for Glauber with $T = 2.38$, our improved stochastic annealing and the hypothetical curve for stochastic annealing according to Eq. 20 and $n = 1.39$

method of Ceperley and Dewing, and additionally integrates some ideas from statistical sampling. The resulting method is called simulated annealing for noisy environments (SANE).

Although the idea of artificially adding noise as presented in the previous section allows us to implement also higher temperatures within stochastic annealing, the acceptance criterion by Ceperley and Dewing (1999) seems better for high temperatures because it features higher acceptance probabilities (it is more similar to Metropolis than to Glauber) and thus is preferable according to the discussion in Sect. 2. Also, it doesn't unnecessarily remove information by adding artificial noise. Thus we propose here to use Eq. 15 as acceptance criterion for temperatures higher than the temperature equivalent to a single sample in stochastic annealing, and to use ideas from stochastic annealing for lower temperatures.

In order to reduce the number of samples required for low temperatures, we propose to additionally integrate a sequential statistical selection procedure. Selection procedures aim at efficiently selecting the better of two or more solutions with a given error probability by sequentially taking one sample after the other until there is sufficient confidence for correct selection. Sequential sampling allows for reaching a given error probability level (probability to erroneously perceive the worse solution as better) with less samples on average. It can stop sampling early if the estimated signal-to-noise ratio is high and a decision seems obvious, and spend more samples if it is low and a decision seems difficult. A comparison of several state-of-the-art selection procedures can be found in Branke et al. (2005). Although we also tested some of the above mentioned state-of-the-art procedures, in the following we will propose a new sequential sampling procedure for two reasons. First, the error probability of sequential procedures may be difficult to control (Branke et al. 2005), and second, stochastic annealing requires the probability for selecting the inferior solution (which is the error probability for a selection procedure) to depend on the true difference $\delta$.

We aim at the Glauber acceptance probability and, as does stochastic annealing, always select the solution perceived as better (threshold criterion). Our idea is to start with a single sample, and estimate the error probability. If the estimated error probability is below the desired probability to accept the inferior solution at the estimated

energy difference according to the Glauber criterion, we stop. Otherwise, we take another sample and repeat the process. The process is described in more detail in Algorithm 2. The overall structure of SANE is summarized as Algorithm 3.

The resulting true acceptance probability $P_A$ does not exactly match the intended Glauber criterion, but comes very close (top panel of Fig. 11). In fact, it is almost identical to the original acceptance probability resulting from stochastic annealing. However, it drastically reduces the average number of samples $E[N]$ required, see bottom panel of Fig. 11. This means that for identical annealing schedules one can expect SANE to behave mostly identical to stochastic annealing, independent of the problem at hand or the parameter settings, while requiring only a fraction of the samples. The savings are larger for smaller temperature, larger $\delta$ and higher noise. Note that SANE also allows to naturally obtain arbitrary temperatures, the original limitation of stochastic annealing is resolved.

Comparing SANE with other approaches from the literature, we believe the following intuitive reasoning is a strong indication that SANE is at least competitive

**Algorithm 2** Sequential sampling

$n = n_0 = 1$
Draw $n_0 = 1$ sample from $\Delta E$ and estimate $\delta$ by $\hat{\delta}$
$P_{err}(\hat{\delta}) = \Phi\left(\frac{-|\hat{\delta}|\sqrt{n}}{\sigma_{\Delta E}}\right)$
**while** $P_{err} > P_a^{\text{Glauber}}(|\hat{\delta}|)$ **do**
  Draw another sample ($n \leftarrow n + 1$)
  Update $\hat{\delta}$ and $P_{err}$
**end while**
Accept better solution

**Algorithm 3** Simulated annealing for noisy environments (SANE)

$T \leftarrow T_0$
Generate initial solution $x_c$
**repeat**
  **repeat**
    Generate $x_n \in N(x_c)$  // generate new solution
    **if** $T \geq 1/\sqrt{8/\pi\sigma_{\Delta E}^2}$
      Sample $\Delta E$ once
      Use Ceperley and Dewing (Eq. 15) to decide about acceptance
    **else**
      Use Alg. 2 to decide about acceptance
    **end if**
  **until** time to reduce temperature
  $T \leftarrow r(T)$
**until** termination condition met

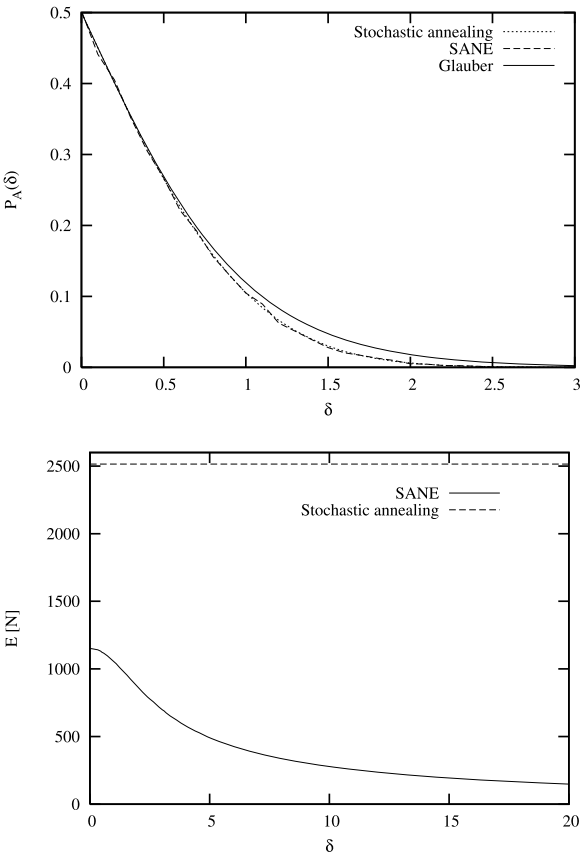


**Fig. 11** Comparison of acceptance probability function $P_A(\delta)$ (*top*) and expected number of required samples (*bottom*) for SANE and stochastic annealing for $\sigma_E = 20$ and $T = 0.5$. Note that in the *top panel*, the lines for stochastic annealing and SANE are almost identical

with previously proposed approaches. Generally, the different approaches to SA in the presence of noise differ in at least three fundamental aspects:

1. The shape of the true acceptance probability function $P_A(\delta)$.
2. The pressure towards good solutions (selection pressure).
3. The number of samples needed for a given selection pressure.

Assuming it is best to shape the acceptance function to adhere to the detailed balance equation, SANE, stochastic annealing, and the approaches by Ceperley and Dewing (1999) and Kennedy and Kuti (1985) are probably the only approaches whose adherence to the detailed balance equation is practically unaffected by noise. The resulting true acceptance probability function of SANE lies between Metropolis and Glauber, thus providing a relatively high acceptance rate. Furthermore, SANE simply removes the effects of noise, thus behaving more or less as SA in deterministic environments. This may be an advantage because the user can rely on experience with deterministic problems when tuning parameters.

Regarding the method's effect on the selection pressure, this can be regulated by choosing an appropriate annealing schedule. A reduced selection pressure can be counterbalanced by decreased temperatures and vice versa. So, if one is free to

choose the optimal annealing schedule for each method, these differences will cancel out. However, SANE can achieve a certain selection pressure with particularly few samples, because

1. it counteracts the effects of noise partly by making the selection step more deterministic (always selecting the seemingly better solution). This seems intuitively more efficient than reducing the noise by additional sampling, and then re-introducing randomness in the selection step.
2. it uses ideas from sequential sampling to minimize the effort to distinguish between two solutions.

In the next section, we will provide empirical evidence substantiating the effectiveness of SANE.

## 6 Empirical evaluation

This section evaluates the proposed methods empirically. In particular, the tested approaches are:

1. Standard SA with Glauber criterion on the deterministic problem
2. The method of Ceperley and Dewing, subsequently denoted as "CD"
3. The confidence-interval based method proposed in Alkhamis et al. (1999), subsequently denoted as "AAT"
4. Stochastic annealing with the improvements presented in Sect. 5.1 to allow arbitrary annealing schedules
5. SANE as proposed in Sect. 5.2.

Our focus is to gain insights into the noise handling mechanism of the different algorithms, rather than a particular application. For this reason, we use a rather simple test problem, namely a standard travelling salesperson (TSP) problem disturbed by adding a Gaussian random noise to the tour length whenever a tour is evaluated. As underlying TSP instances, we use the following variety of data sets from TSPLIB (2002): `eil51`, a randomly generated TSP problem with 51 cities; `bier127`, a TSP instance based on real-world data, trying to find the shortest tour through the 127 beer gardens in Augsburg, and `lin105`, a very regular drilling hole problem. The standard deviation of the noise is assumed to be known and is specified in the experiments below.

On the algorithm side, we use a standard permutation representation with swap neighborhood. All results reported below are averaged over 300 runs with different random seeds. Statistical significance is determined using a two-sided t-Test with significance level 0.05. We use common random numbers to sharpen contrasts between the different procedures. Unless specified otherwise, we report on the true (undisturbed) quality of the current solution. Since for all tested methods the computational overhead can be considered small compared to the time required for a single evaluation, running time is equated to the number of evaluations performed.

We split the analysis into two subsections. The first subsection compares SANE with related methods like CD and stochastic annealing. Then, in Subsect. 6.2, we compare SANE with AAT on three test problem instances and with a large number of different parameter settings.

### 6.1 Comparison with related methods

Stochastic annealing, CD, and SANE all follow the same basic idea of removing the effects of noise on the detailed balance equation while at the same time adhering to a given annealing schedule. For a given annealing schedule, the ratio of acceptance probabilities as defined by Eq. 6 is the same for all three algorithms, thus they should behave more or less identical except for the number of samples they need to reach a thermodynamical equilibrium. From the discussion of the methods in the previous sections, we may expect that SANE requires significantly fewer samples per iteration than stochastic annealing, and CD is expected to stagnate at lower temperatures due to the extremely low move acceptance probabilities.

For the above reasons, it seems safe to assume that annealing schedule or test problem have little impact on the qualitative results, and we restrict our experiments in this section to a single TSP instance (`eil51`). The annealing schedule is chosen based on a few preliminary tests on the deterministic problem, assuming that the methods try to remove the effect of noise, and that an annealing schedule that works well on the deterministic problem should also work well with these methods. It is a simple geometric sequence, starting with initial temperature $T_0 = 50$ and then reducing temperature every 10 iterations according to $T_{k+1} = \alpha T_k$, with $\alpha = 0.95^{(1/100)}$. All tested procedures use the same parameter settings as far as possible. The observed differences can be attributed to the way the different methods handle the noise.

Figure 12 compares the performance over the number of iterations on the problem with a noise level $\sigma_E = 40$. The figure also shows the performance of SA with Glauber criterion and without noise for comparison. In the early stages of the run, the differences are small. SANE and CD perform identical in the first iterations (as SANE uses CD during this period), and both are slightly better than stochastic annealing. However, after a while CD stagnates, because the probability to accept another solution, even if better, decreases dramatically with decreasing temperature. In the later stages, stochastic annealing and SANE show a similarly good convergence. Despite the noise, stochastic annealing and SANE are rather close to (and even slightly better than) SA in the deterministic case, which shows that their mechanism to handle the noise is effective. In the stochastic case, however, the curves are more rugged,

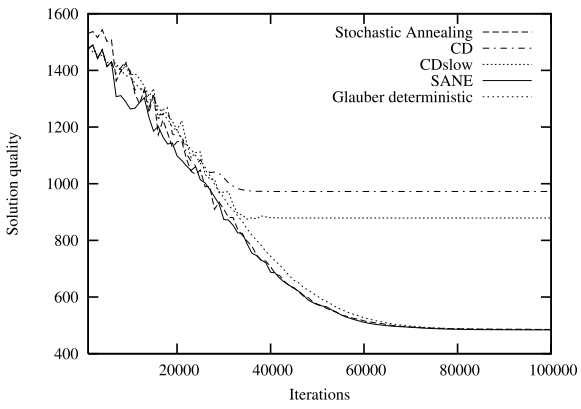


**Fig. 12** True quality of current solution depending on the iteration for $\sigma_E = 40$

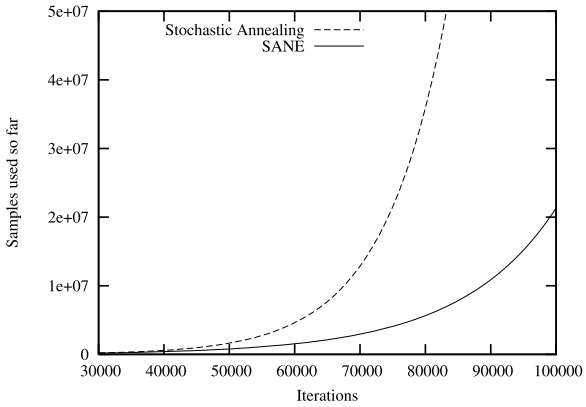


**Fig. 13** Number of samples used per iteration for $\sigma_E = 40$

because the suggested approaches can only reproduce the *expected* acceptance probabilities, while exhibiting a larger variance. The difference in solution quality obtained after 100,000 iterations of stochastic annealing, SANE, and the Glauber criterion in the deterministic case are pairwise statistically insignificant.

However, solution quality is not the only performance criterion. The three methods (CD, improved stochastic annealing, and SANE) use different numbers of fitness evaluations in each iteration. Figure 13 shows the number of fitness evaluations up to a given iteration. Clearly, stochastic annealing requires many more samples than our new method SANE. CD uses only one evaluation per iteration.

To see whether the low quality of solutions obtained by CD is only due to it using so few evaluations, we tested CD again but this time adapted the number of iterations on each annealing stage such that the total number of evaluations is equivalent to the number of evaluations used by stochastic annealing. In other words, we adjust the temperature not based on the number of *iterations*, but based on the number of *evaluations*, and use the same annealing schedule as for stochastic annealing. The result is labeled as "CDslow" in the plots. As can be seen in Fig. 12, even if CD uses the same large number of evaluations as stochastic annealing, it does not reach the same performance level as the other two methods.

If the performance is plotted based on the number of samples instead of the number of iterations (Fig. 14), the advantage of SANE is evident and the quality differences at the end of the run are statistically significant. For lower noise levels, the results look similar, only CD stagnates later and is competitive for a longer time, see Fig. 15. Again, SANE significantly outperforms stochastic annealing.

### 6.2 Comparison with Alkhamis et al.

In this subsection, we compare the best of the three SA variants examined above (SANE), with the approach by Alkhamis et al. (1999) (denoted as AAT), the most recent of the approaches from Sect. 3.1 that do not store all visited solutions. While AAT has proved asymptotic convergence to the optimum, it does not simply remove the noise, but results in a modified true acceptance probability function.

The comparison of fundamentally different approaches such as SANE and AAT is more challenging, as the result may be strongly dependent on the parameter settings

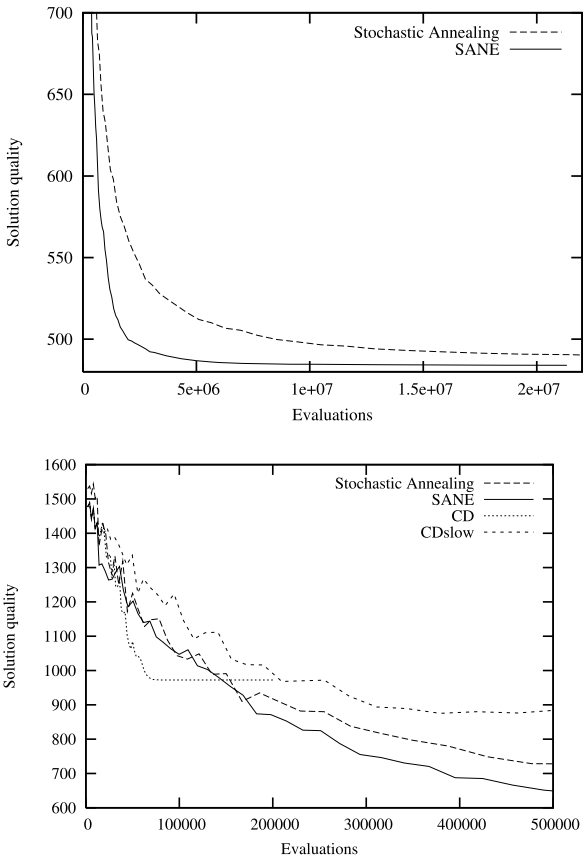


**Fig. 14** True quality of current solution depending on the number of evaluations so far, for $\sigma_E = 40$. *Bottom panel* shows only the first 500,000 evaluations on a larger scale

or problem instance. For a fair comparison, we attempt to fine-tune the parameters of each approach. While SANE only has the annealing schedule (initial temperature and $\alpha$) as parameters, AAT additionally requires to define the number of samples $N_k$ in the $k$-th annealing stage, which makes this method more difficult to fine-tune. In the empirical investigation of Alkhamis et al. (1999), a linearly increasing number of samples is suggested, namely $N_k = 5 + \lfloor k/300 \rfloor$, where $k$ is the annealing stage. We use the same underlying formula, multiplied by an additional parameter $\beta$ for tuning.

The initial parameter tuning is again based on the noisy `ei151` TSP. We aim for the best possible final result with a total of $10^6$ samples for $\sigma_E = 10$ and $5 \cdot 10^6$ samples for $\sigma_E = 40$. Initially, we did a full factorial experimental design with $T_0 \in \{10, 30, 50\}$, $\alpha \in \{0.99, 0.995, 0.999\}$, and $\beta \in \{1, 5, 10, 20\}$. These results showed that SANE could probably benefit from higher $\alpha$, while AAT could benefit from lower $\alpha$ and higher $\beta$, and we did some additional tests in these directions. Tables 1 and 2 for SANE and Tables 3 and 4 for AAT illustrate how the different approaches react on different parameter settings.

The best found solutions are highlighted in bold. On the problem with $\sigma_E = 10$, the best found parameter settings for SANE are $T_0 = 10$ and $\alpha = 0.9999$, yielding

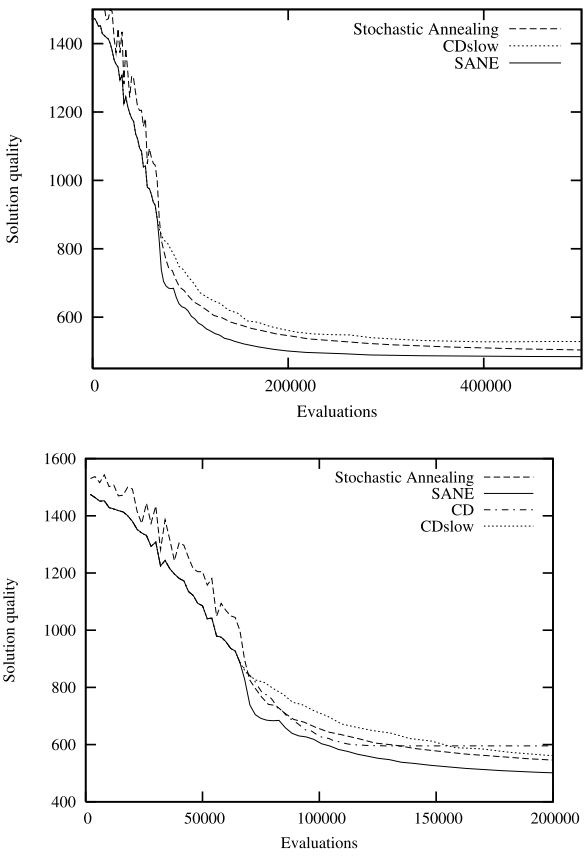


**Fig. 15** True quality of current solution depending on the number of evaluations so far, for $\sigma_E = 10$. *Bottom panel* shows only the first 200,000 evaluations on a larger scale

**Table 1** Performance of SANE on `eil51` with $\sigma_E = 10$ depending on different parameter settings

| | $T_0$ | 10 | 30 | 50 |
|---|---|---|---|---|
| | 0.99 | 587.5 | 577.0 | 575.7 |
| | 0.995 | 554.0 | 538.2 | 546.4 |
| $\alpha$ | 0.999 | 501.7 | 499.3 | 498.4 |
| | 0.9997 | 474.6 | 474.4 | 475.2 |
| | 0.9999 | **467.8** | 470.4 | 477.4 |

an average tour length of 467.8. For AAT, the best parameter settings are $T_0 = 10$, $\alpha = 0.98$ and $\beta = 10$ with an average tour length of 522.7. SANE clearly outperforms AAT on this problem, and the performance difference is significant. While SANE requires a relatively high setting of $\alpha$ to avoid very low temperatures and corresponding huge sample sizes in the later iterations, AAT performs better with low $\alpha$ values, quickly reducing the temperature and effectively turning the algorithm into a hill-climbing variant.

Figure 16 compares the convergence of AAT and SANE for the respective best found parameter settings. SANE outperforms AAT throughout the run.

**Table 2** Performance of SANE on `eil51` with $\sigma_E = 40$ depending on different parameter settings

| | $T_0$ | 10 | 30 | 50 |
|---|---|---|---|---|
| | 0.99 | 593.5 | 582.1 | 584.4 |
| | 0.995 | 558.2 | 541.2 | 547.4 |
| $\alpha$ | 0.999 | 503.5 | 499.4 | 501.0 |
| | 0.9997 | **481.8** | 482.5 | 483.3 |
| | 0.9999 | 497.1 | 512.2 | 515.1 |

**Table 3** Performance of AAT on `eil51` with $\sigma_E = 10$ depending on different parameter settings

| | $\beta$ | 5 | | | 10 | | | 20 | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | $T_0$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| | 0.96 | 527.1 | 528.0 | 527.5 | 524.3 | 526.0 | 523.8 | 532.6 | 532.8 | 534.6 |
| | 0.98 | 530.9 | 530.7 | 529.5 | **522.7** | 523.5 | 525.1 | 534.1 | 533.9 | 534.8 |
| $\alpha$ | 0.99 | 528.2 | 530.9 | 531.4 | 526.2 | 526.5 | 530.3 | 530.7 | 536.2 | 536.2 |
| | 0.995 | 530.5 | 533.0 | 532.9 | 525.2 | 533.6 | 533.7 | 533.3 | 551.1 | 551.2 |
| | 0.999 | 577.7 | 700.7 | 832.8 | 605.4 | 830.6 | 1,014.7 | 650.3 | 968.6 | 1,176.4 |

**Table 4** Performance of AAT on `eil51` with $\sigma_E = 40$ depending on different parameter settings

| | $\beta$ | 30 | | | 200 | | | 250 | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | $T_0$ | 10 | 30 | 50 | 10 | 30 | 50 | 10 | 30 | 50 |
| | 0.96 | 642.6 | 646.4 | 642.5 | **573.1** | 575.9 | 573.2 | 577.3 | 576.0 | 576.1 |
| | 0.98 | 645.8 | 644.7 | 644.7 | 574.9 | 576.4 | 576.2 | 576.1 | 579.9 | 580.7 |
| $\alpha$ | 0.99 | 644.9 | 645.1 | 645.0 | 578.3 | 582.1 | 583.1 | 579.1 | 582.3 | 588.4 |
| | 0.995 | 642.6 | 646.8 | 642.8 | 584.0 | 631.1 | 624.6 | 588.0 | 649.5 | 648.2 |
| | 0.999 | 693.2 | 803.9 | 914.3 | 762.0 | 1,124.5 | 1,314.8 | 775.5 | 1,153.7 | 1,321.4 |

For the stronger noise of $\sigma_E = 40$, the picture looks similar. Best parameter settings for SANE are $T_0 = 10$ and $\alpha = 0.9997$ with an average tour length of 481.75, while for AAT, they are $T_0 = 10$, $\alpha = 0.96$, and $\beta = 200$ with an average tour length of 573.1. Both approaches seem to prefer slightly lower $\alpha$ to counteract the otherwise slower convergence due to the higher noise. For AAT, the higher noise also requires a much higher setting for $\beta$ in order to reduce the effect of the stronger noise. SANE, on the other hand, adjusts the number of samples automatically.

To see whether the results also carry over to other problem instances, we additionally tested each approach on `bier127` and `lin105`. In practice, fine-tuning algorithm parameters for every problem instance is not realistic. Thus, in the following, we attempt to generate comparable environments for the different problem instances, allowing us to re-use the insights on the best found parameters from the `eil51` tests. The scaling of the three problem instances is quite different. We use the average distance between any two cities for scaling, which is 31.18, 1155.14 and 4875.09 for `eil51`, `lin105`, and `bier127`, respectively. To obtain simi-

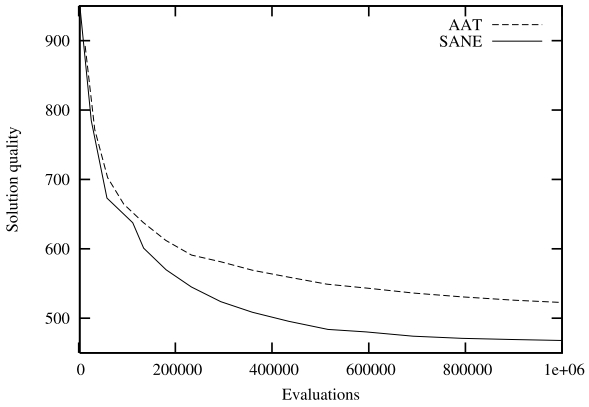


**Fig. 16** Quality of current solution depending on the number of evaluations, for `eil51`, $\sigma_E = 10$, and best found parameter settings

**Table 5** Performance of SANE on `lin105` with $\sigma_E = 262$ depending on different parameter settings

| | $T_0$ | 185.25 | 370.5 | 741.0 |
|---|---|---|---|---|
| | 0.9997 | 24,128 | 21,891 | 21,583 |
| $\alpha$ | 0.9999 | 21,763 | **20,186** | 20,360 |
| | 0.99995 | 21,576 | 20,898 | 22,606 |

**Table 6** Performance of SANE on `bier127` with $\sigma_E = 1105.6$ depending on different parameter settings

| | $T_0$ | 781.75 | 1563.5 | 3127 |
|---|---|---|---|---|
| | 0.9997 | 168,587 | 159,533 | 158,316 |
| $\alpha$ | 0.9999 | 157,532 | **151,006** | 152,514 |
| | 0.99995 | 156,631 | 154,410 | 163,626 |

lar signal-to-noise ratios, we thus scale $\sigma_{\Delta E}$ accordingly, resulting in $\sigma_E = 262$ for `lin105`, and $\sigma_E = 1105.6$ for `bier127` as equivalent to $\sigma_E = 10$ on `eil51` (note that $\sigma_{\Delta E} = \sqrt{2}\sigma_E$). The annealing schedule is also adjusted by scaling the initial temperature to $T_0 = 370.5$ and $T_0 = 1563.5$ for `lin105` and `bier127`, respectively. The annealing parameter $\alpha$ and AAT's parameter $\beta$ should not be affected by the scaling.

Tables 5 and 6 for SANE and Tables 7 and 8 for AAT show the results of these derived parameter settings, and some parameter settings around it to provide insights into the sensitivity of the results.

For SANE, the derived parameter settings seem to work fine and indeed yield the best results of all tested parameter settings. This demonstrates that SANE actually allows a user to rely on experience when choosing parameter settings.

For AAT, this does not seem to be the case. For `lin105`, best results have been obtained with lower $\beta$, higher $\alpha$, and higher initial temperature than expected from the `eil51` tests ($\beta = 1, \alpha = 0.99$, and $T_0 = 741$ instead of $\beta = 10, \alpha = 0.98$, and $T_0 = 370.5$), while for `bier127`, we needed lower $\beta$, lower $\alpha$, and higher initial temperature ($\beta = 1, \alpha = 0.96$, $T_0 = 3127$ instead of $\beta = 10, \alpha = 0.98$, and $T_0 = 1563.5$).

**Table 7** Performance of AAT on `lin105` with $\sigma_E = 262$ depending on different parameter settings

| | $\beta$ | 1 | | | 5 | | | 10 | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | $T_0$ | 185.25 | 370.5 | 741.0 | 185.25 | 370.5 | 741.0 | 185.25 | 370.5 | 741.0 |
| | 0.96 | 29,364 | 29,541 | 29,365 | 30,401 | 30,691 | 30,666 | 31,936 | 32,139 | 32,219 |
| $\alpha$ | 0.98 | 29,417 | 29,374 | 29,533 | 30,600 | 30,458 | 30,366 | 32,122 | 32,079 | 31,839 |
| | 0.99 | 28,657 | 28,613 | **28,576** | 29,169 | 29,311 | 29,351 | 30,550 | 30,294 | 30,243 |

**Table 8** Performance of AAT on `bier127` with $\sigma_E = 1105.6$ depending on different parameter settings

| | $\beta$ | 1 | | | 5 | | | 10 | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | $T_0$ | 781.75 | 1,563.5 | 3,127 | 781.75 | 1,563.5 | 3,127 | 781.75 | 1,563.5 | 3,127 |
| | 0.96 | 196,993 | 196,655 | **196,147** | 200,003 | 200,359 | 200,659 | 206,199 | 207,253 | 205,617 |
| $\alpha$ | 0.98 | 200,263 | 197,147 | 197,560 | 200,263 | 200,015 | 200,284 | 206,293 | 206,542 | 207,326 |
| | 0.99 | 196,449 | 196,768 | 196,351 | 200,288 | 198,371 | 200,757 | 206,705 | 205,324 | 206,956 |

**Table 9** Comparison of SANE and AAT on eil51, lin105, and bier127 noisy TSP instances, with best found parameter settings (mean $\pm$ std. error)

| | SANE | AAT |
|---|---|---|
| `eil51` | $467.8 \pm 0.72$ | $522.7 \pm 1.41$ |
| `lin105` | $20,186 \pm 70.0$ | $28,576 \pm 108.8$ |
| `bier127` | $151,006 \pm 309.3$ | $196,147 \pm 462.4$ |

If we just compare the tours obtained by SANE and AAT with the parameters derived based on the `eil51` tests, the tours found by AAT are 59% longer for `lin105`, and 37% longer for `bier127`. And although we tested three times as many parameter settings for AAT as for SANE, the best obtained result is still 42% (`lin105`) and 30% (`bier127`) worse than with SANE. Table 9 summarizes the results, showing that SANE significantly outperforms AAT on all problem instances. Furthermore, SANE has a much lower standard error, i.e., it is more reliable.

## 7 Conclusion

Many practical optimization problems are noisy, and optimization algorithms able to cope with such noise are needed. There have been numerous previous attempts to adapt simulated annealing to noisy optimization problems. We provided a state-of-the-art survey of the field, proposed an improved and more flexible version of a very promising SA variant, stochastic annealing, and also proposed a new SA variant, called *simulated annealing in noisy environments* (SANE), which uses ideas from sequential sampling to drastically reduce the number of function evaluations required to obtain a desired acceptance probability.

On a noisy travelling salesperson problem, SANE was empirically compared with three other state-of-the-art approaches: the improved version of stochastic annealing, the approach of Ceperley and Dewing (1999), and the approach of Alkhamis et

al. (1999). As the results show, SANE significantly outperforms all the other tested approaches, generating better solutions with fewer required function evaluations.

There are several avenues for future work. First, all of our approaches took new samples of the old and the candidate solution in each iteration. Probably, it would be beneficial to retain information about samples taken in previous iterations, although this might introduce some unwanted correlations (a solution that survived one iteration because it was over-evaluated is more likely to survive also the next if the old evaluations are kept). Second, assuming a known variance is a limitation of our study. Hence, examining the effect of unknown variances (which have to be estimated) on all approaches, and improving variance estimates, e.g., by batching, are natural next steps. Finally, the sequential sampling procedure could be further improved by not always sampling both solutions but sampling only one at a time.

**Acknowledgements** We would like to thank the anonymous referees for their constructive comments.

## References

Ahmed, M.A., Alkhamis, T.M.: Simulation-based optimization using simulated annealing with ranking and selection. Comput. Oper. Res. **29**(4), 387–402 (2002)

Alkhamis, T.M., Ahmed, M.A.: Simulation-based optimization using simulated annealing with confidence interval. In: Ingalls, R.G. et al. (eds.) Winter Simulation Conference, pp. 514–519 (2004)

Alkhamis, T.M., Ahmed, M.A., Tuan, V.K.: Simulated annealing for discrete optimization with estimation. Eur. J. Oper. Res. **116**(3), 530–544 (1999)

Alrefaei, M.H., Andradóttir, S.: A simulated annealing algorithm with constant temperature for discrete stochastic optimization. Manag. Sci. **45**(5), 748–764 (1999)

Ball, R.C., Fink, T.M.A., Bowler, N.E.: Stochastic annealing. Phys. Rev. Lett. **91**(3), 030201 (2003a)

Ball, R.C., Fink, T.M.A., Bowler, N.E.: Characterization of the probabilistic traveling salesman problem. Phys. Rev. E **68**(3), 036703 (2003b)

Bowler, N.E.: The role of noise in optimisation and diffusion limited aggregation. Ph.D. thesis, University of Warwick (2001)

Branke, J., Chick, S., Schmidt, C.: New developments in ranking and selection: an empirical comparison of the three main approaches. In: Kuhl, N.E., Steiger, M.N., Armstrong, F.B., Joines, J.A. (eds.) Winter Simulation Conference, pp. 708–717. IEEE, New York (2005)

Bulgak, A.A., Sanders, J.L.: Integrating a modified simulated annealing algorithm with the simulation of a manufacturing system to optimize buffer sizes in automatic assembly systems. In: Winter Simulation Conference, pp. 684–690. IEEE, New York (1988)

Ceperley, D.M., Dewing, M.: The penalty method for random walks with uncertain energies. J. Chem. Phys. **110**(20), 9812–9820 (1999)

Cheh, K.M., Goldberg, J.B., Askin, R.G.: A note on the effect of neighborhood structure in simulated annealing. Comput. Oper. Res. **18**(6), 537–548 (1991)

Fink, T.M.A.: Inverse protein folding, hierarchical optimisation and tie knots. Ph.D. thesis, University of Cambridge (1998)

Fox, B.L., Heine, G.W.: Probabilistic search with overrides. Ann. Appl. Probab. **5**(4), 1087–1094 (1995)

Gelfand, S.B., Mitter, S.K.: Simulated annealing with noisy or imprecise energy measurements. J. Optim. Theory Appl. **62**(1), 49–62 (1989)

Glauber, R.J.: Time-dependent statistics of the Ising model. J. Math. Phys. **4**(2), 294–307 (1963)

Gutjahr, W.J., Pflug, G.C.: Simulated annealing for noisy cost functions. J. Glob. Optim. **8**(1), 1–13 (1996)

Haddock, J., Mittenthal, J.: Simulation optimization using simulated annealing. Comput. Ind. Eng. **20**(4), 387–395 (1992)

Hajek, B.: Cooling schedules for optimal annealing. Math. Oper. Res. **13**(2), 311–329 (1988)

Hastings, W.K.: Monte Carlo sampling methods using Markov chains and their applications. Biometrika **57**(1), 97–109 (1970)

Hertz, J., Krogh, A., Palmer, R.G.: Introduction to the Theory of Neural Computation. Addison-Wesley, Reading (1991)

Kennedy, A.D., Kuti, J.: Noise without noise: a new Monte Carlo method. Phys. Rev. Lett. **54**(23), 2473–2476 (1985)

Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. Science **220**(4598), 671–680 (1983)

Kushner, H.J.: Asymptotic global behavior for stochastic approximation and diffusions with slowly decreasing noise effects: global minimization via Monte Carlo. SIAM J. Appl. Math. **47**(1), 169–185 (1987)

Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of state calculations by fast computing machines. J. Chem. Phys. **21**(6), 1087–1092 (1953)

Nourani, Y., Andresen, B.: A comparison of simulated annealing cooling strategies. J. Phys. A: Math. Gen. **31**, 8373–8385 (1998)

Painton, L., Diwekar, U.: Stochastic annealing for synthesis under uncertainty. Eur. J. Oper. Res. **83**(3), 489–502 (1995)

Prudius, A.A., Andradottir, S.: Two simulated annealing algorithms for noisy objective functions. In: Kuhl, M.E. et al. (eds.) Winter Simulation Conference, pp. 797–802. IEEE, New York (2005)

Rees, S., Ball, R.C.: Criteria for an optimum simulated annealing schedule for problems of the travelling salesman type. J. Phys. A: Math. Gen. **20**(5), 1239–1249 (1987)

TSPLIB. http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/index.html (2002)

van Laarhoven, P.J.M., Aarts, E.H.L.: Simulated Annealing: Theory and Applications. Reidel, Dordrecht (1987)

Wang, L., Zhang, L.: Stochastic optimization using simulated annealing with hypothesis test. Appl. Math. Comput. **174**(2), 1329–1342 (2006)