ELSEVIER

Signal Processing 79 (1999) 117–128

SIGNAL PROCESSING

www.elsevier.nl/locate/sigpro

Adaptive simulated annealing for optimization in signal processing applications

S. Chen[a,*], B.L. Luk[b]

[a]*Department of Electronics and Computer Science, University of Southampton, Highfield, Southampton SO17 1BJ, UK*
[b]*Department of Electrical and Electronic Engineering, University of Portsmouth, Anglesea Building, Portsmouth PO1 3DJ, UK*

Received 12 June 1998; received in revised form 27 May 1999

Abstract

Many signal processing applications pose optimization problems with multimodal and nonsmooth cost functions. Gradient methods are ineffective in these situations. The adaptive simulated annealing (ASA) offers a viable optimization tool for tackling these difficult nonlinear optimization problems. Three applications, maximum likelihood (ML) joint channel and data estimation, infinite-impulse-response (IIR) filter design and evaluation of minimum symbol-error-rate (MSER) decision feedback equalizer (DFE), are used to demonstrate the effectiveness of the ASA. © 1999 Elsevier Science B.V. All rights reserved.

Zusammenfassung

Viele Anwendungen der Signalverarbeitung werfen Optimierungsprobleme mit multimodalen und nichtglatten Kostenfunktionen auf. Gradientenmethoden sind in diesem Fall uneffektiv. Das "adaptive simulated annealing" (ASA) ist ein funktionsfähiges Optimierungswerkzeug zur Lösung dieser schweren nichtlinearen Optimierungsprobleme. Drei Anwendungen, gemeinsame Maximum Likelihood (ML) Kanal- und Datenschätzung, Filterdesign mit unendlichen Impulsantwort (IIR) und Berechnungen eines entscheidungsrückgekoppelten Entzerrers, der die Symbolfehlerrate minimiert, werden verwendet, um die Effektivität des ASA zu zeigen. © 1999 Elsevier Science B.V. All rights reserved.

Résumé

De nombreuses applications de traitement de signaux posent des problèmes d'optimisation avec des fonctions de coût multimodales et non lisses. Des méthodes de gradient sont inefficaces dans ces situations. Le recuit simulé adaptatif offre un outil d'optimisation viable pour traiter ces problèmes complexes d'optimisation non linéaire. Trois applications, l'estimation conjointe de canal et de données par le maximum de vraisemblance, la conception de filtres à réponse impulsionnelle infinie et l'évaluation d'égaliseurs à rétroaction de la décision à taux d'erreur par symbole minimal sont utilisées pour démontrer l'efficacité du recuit simulé adaptatif. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Simulated annealing; Global optimization; Blind equalization; IIR filter; Decision feedback equalizer

* Corresponding author.

PII: S0165-1684(99)00084-5

1. Introduction

Optimization problems with multimodal and/or nonsmooth cost functions are commonly encountered in signal processing applications. Conventional gradient-based algorithms are ineffective in these applications due to the problem of local minima or the difficulty in calculating gradients. Optimization methods that require no gradient and can achieve a global optimal solution offer considerable advantages in solving these difficult optimization problems. Two best-known classes of such global optimization methods are the genetic algorithm (GA) [8,10,12] and the simulated annealing (SA) [7,16,31].

The GA and SA belong to a class of so-called guided random search methods. The underlying mechanisms for guiding optimization search process are, however, very different for the two methods. The GA is population-based, and evolves a population-based solution according to the principles of the evolution of species in nature. The SA, on the other hand, evolves a single solution in the parameter space with certain guiding principles that imitate the random behaviour of molecules during the annealing process.

The GA seems to have attracted the main attention in the application to signal processing problems (e.g. [1,19,30]). The SA by contrast has not received similar interests. The SA represents a global optimization technique with some striking positive and negative features. An attractive feature of SA is that it is very easy to program and the algorithm typically has few parameters that require tuning. Furthermore, its statistical guarantee of convergence should make SA very appealing. However, a serious drawback of SA is that it is often very slow. In many diverse applications, a standard SA algorithm often requires much more number of cost-function evaluations to converge, compared with a carefully designed and tuned GA.

However, an improved version of SA, referred to as the ASA [2,13–15,25], provides significant improvement in convergence speed over standard versions of SA and maintains all the advantages of standard SA algorithms. This paper aims to introduce the ASA to the signal processing community. To illustrate its simplicity and versatility, we apply the ASA to three signal processing applications, ML joint channel and data estimation, IIR filter design and evaluation of the MSER DFE. Our study demonstrates that the ASA offers a viable alternative to the GA as a global optimization tool for solving diverse signal processing problems.

2. Adaptive simulated annealing

The ASA, also known as the very fast simulated reannealing, is a very efficient version of SA. Detailed analysis of the algorithm can be found in [2,13–15,25]. Many signal processing applications pose the following general optimization problem:

$$\min_{\boldsymbol{w}\in\mathscr{W}} J(\boldsymbol{w}), \tag{1}$$

where $\boldsymbol{w} = [w_1 \dots w_n]^{\mathrm{T}}$ is the n-dimensional parameter vector to be optimized,

$$\mathscr{W} \triangleq \{\boldsymbol{w}: (L_i \leqslant w_i \leqslant U_i,\ 1 \leqslant i \leqslant n) \cap (\alpha_j \leqslant g_j(\boldsymbol{w}) \leqslant \beta_j,\ 1 \leqslant j \leqslant m)\} \tag{2}$$

is the feasible set of $\boldsymbol{w}$, L_i and U_i are the lower and upper bounds of w_i, respectively, and $\alpha_j \leqslant g_j(\boldsymbol{w}) \leqslant \beta_j$, $1 \leqslant j \leqslant m$, are the m inequality constraints. The cost function $J(\boldsymbol{w})$ can be multimodal and nonsmooth. The ASA is a global optimization scheme for solving this kind of constrained optimization problems.

2.1. Search guiding mechanisms

The ASA evolves a single point $\boldsymbol{w}$ in the parameter or state space $\mathscr{W}$. The seemingly random search is guided by certain underlying probability distributions. An elegant discussion on how the general SA algorithm works can be found in [25]. Specifically, the general SA algorithm is described by three functions.

2.1.1. Generating probability density function

$$G(w_i^{\mathrm{old}}, w_i^{\mathrm{new}}, T_{i,\mathrm{gen}};\ 1 \leqslant i \leqslant n). \tag{3}$$

This determines how a new state $\boldsymbol{w}^{\mathrm{new}}$ is created, and from what neighbourhood and probability distributions it is generated, given the current state

$\boldsymbol{w}^{\text{old}}$. The generating "temperatures" $T_{i,\text{gen}}$ describe the widths or scales of the generating distribution along each dimension w_i of the state space.

Often a cost function has different sensitivities along different dimensions of the state space. Ideally, the generating distribution used to search a steeper and more sensitive dimension should have a narrower width than that of the distribution used in searching a dimension less sensitive to change. The ASA adopts a so-called reannealing scheme to periodically re-scale $T_{i,\text{gen}}$, so that they optimally adapt to the current status of the cost function. This is an important mechanism, which not only speeds up the search process but also makes the optimization process robust to different problems.

2.1.2. Acceptance function

$$P_{\text{accept}}(J(\boldsymbol{w}^{\text{old}}), J(\boldsymbol{w}^{\text{new}}), T_{\text{accept}}). \tag{4}$$

This gives the probability of $\boldsymbol{w}^{\text{new}}$ being accepted. The acceptance temperature determines the frequency of accepting new states of poorer quality.

Probability of acceptance is very high at very high temperature T_{accept}, and it becomes smaller as T_{accept} is reduced. At every acceptance temperature, there is a finite probability of accepting the new state. This produces occasionally uphill move, enables the algorithm to escape from local minima, and allows a more effective search of the state space to find a global minimum. The ASA also periodically adapts T_{accept} to best suit the status of the cost function. This helps to improve convergence speed and robustness.

2.1.3. Reduce temperatures or annealing schedule

$$\begin{aligned} &T_{\text{accept}}(k_a) \to T_{\text{accept}}(k_a + 1), \\ &T_{i,\text{gen}}(k_i) \to T_{i,\text{gen}}(k_i + 1), \quad 1 \leqslant i \leqslant n, \end{aligned} \tag{5}$$

where k_a and k_i are some annealing time indexes. The reduction of temperatures should be sufficiently gradual in order to ensure that the algorithm finds a global minimum.

This mechanism is based on the observations of the physical annealing process. When the metal is cooled from a high temperature, if the cooling is sufficiently slow, the atoms line themselves up and form a crystal, which is the state of minimum energy in the system. The slow convergence of many SA algorithms is rooted at this slow annealing process. The ASA, however, can employ a very fast annealing schedule, as it has self adaptation ability to re-scale temperatures.

2.2. Algorithm implementation

Although there are many possible realizations of the ASA, an implementation is illustrated in Fig. 1, and this algorithm is detailed here. How the ASA realizes the above three functions will also become clear during the description.

(i) In the initialization, an initial $\boldsymbol{w} \in \mathscr{W}$ is randomly generated, the initial temperature

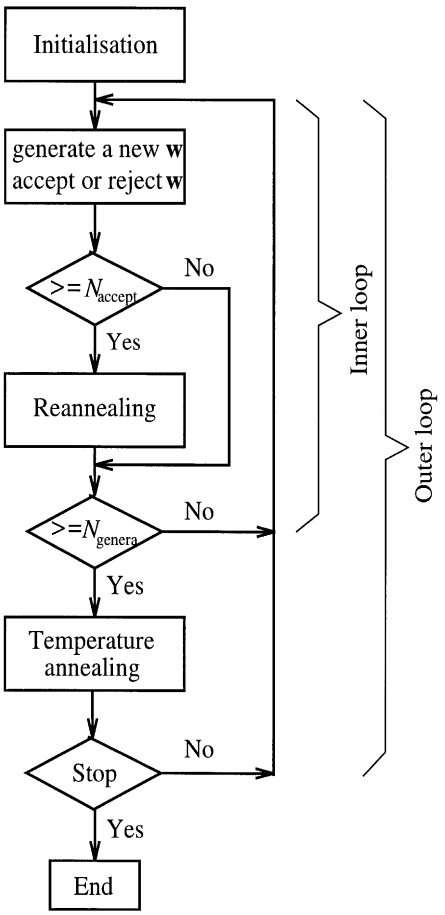

Fig. 1. Flow chart of the adaptive simulated annealing.

of the acceptance probability function, $T_{\text{accept}}(0)$, is set to $J(\boldsymbol{w})$, and the initial temperatures of the parameter generating probability functions, $T_{i,\text{gen}}(0)$, $1 \leqslant i \leqslant n$, are set to 1.0. A user-defined control parameter c in annealing is given, and the annealing times, k_i for $1 \leqslant i \leqslant n$ and k_a, are all set to 0.

(ii) The algorithm generates a new point in the parameter space with

$$w_i^{\text{new}} = w_i^{\text{old}} + q_i(U_i - L_i) \quad \text{for } 1 \leqslant i \leqslant n$$

and

$$\boldsymbol{w}^{\text{new}} \in \mathscr{W}, \tag{6}$$

where q_i is calculated as

$$q_i = \text{sgn}\left(v_i - \frac{1}{2}\right) T_{i,\text{gen}}(k_i) \times \left(\left(1 + \frac{1}{T_{i,\text{gen}}(k_i)}\right)^{|2v_i - 1|} - 1\right) \tag{7}$$

and v_i a uniformly distributed random variable in [0,1]. Notice that if a generated $\boldsymbol{w}^{\text{new}}$ is not in $\mathscr{W}$, it is simply discarded and a new point is tried again until $\boldsymbol{w}^{\text{new}} \in \mathscr{W}$.

The value of the cost function $J(\boldsymbol{w}^{\text{new}})$ is then evaluated and the acceptance probability function of $\boldsymbol{w}^{\text{new}}$ is given by

$$P_{\text{accept}} = \frac{1}{1 + \exp((J(\boldsymbol{w}^{\text{new}}) - J(\boldsymbol{w}^{\text{old}}))/T_{\text{accept}}(k_a))}. \tag{8}$$

A uniform random variable P_{unif} is generated in [0,1]. If $P_{\text{unif}} \leqslant P_{\text{accept}}$, $\boldsymbol{w}^{\text{new}}$ is accepted; otherwise it is rejected.

(iii) After every N_{accept} acceptance points, reannealing takes place by first calculating the sensitivies

$$s_i = \left|\frac{J(\boldsymbol{w}^{\text{best}} + \boldsymbol{e}_i\delta) - J(\boldsymbol{w}^{\text{best}})}{\delta}\right|, \quad 1 \leqslant i \leqslant n, \tag{9}$$

where $\boldsymbol{w}^{\text{best}}$ is the best point found so far, δ is a small step size, the n-dimensional vector $\boldsymbol{e}_i$ has unit ith element and the rest of elements of $\boldsymbol{e}_i$ are all zeros. Let $s_{\max} = \max\{s_i,\ 1 \leqslant i \leqslant n\}$. Each parameter generating temperature $T_{i,\text{gen}}$ is scaled by a factor $s_{\max}/s_i$ and the annealing time k_i is reset,

$$T_{i,\text{gen}}(k_i) = \frac{s_{\max}}{s_i} T_{i,\text{gen}}(k_i), \qquad k_i = \left(-\frac{1}{c}\log\left(\frac{T_{i,\text{gen}}(k_i)}{T_{i,\text{gen}}(0)}\right)\right)^n. \tag{10}$$

Similarly, $T_{\text{accept}}(0)$ is reset to the value of the last accepted cost function, $T_{\text{accept}}(k_a)$ is reset to $J(\boldsymbol{w}^{\text{best}})$ and the annealing time k_a is rescaled accordingly,

$$k_a = \left(-\frac{1}{c}\log\left(\frac{T_{\text{accept}}(k_a)}{T_{\text{accept}}(0)}\right)\right)^n. \tag{11}$$

(iv) After every N_{genera} generated points, annealing takes place with

$$k_i = k_i + 1, \qquad T_{i,\text{gen}}(k_i) = T_{i,\text{gen}}(0)\exp(-ck_i^{1/n}), \quad 1 \leqslant i \leqslant n, \tag{12}$$

and

$$k_a = k_a + 1, \qquad T_{\text{accept}}(k_a) = T_{\text{accept}}(0)\exp(-ck_a^{1/n}); \tag{13}$$

otherwise, go to step (ii).

(v) The algorithm is terminated if the parameters have remained unchanged for a few successive reannealings or a preset maximum number of cost function evaluations has been reached; otherwise, go to step (ii).

As in a standard SA algorithm, this ASA contains two loops. The inner loop ensures that the parameter space is searched sufficiently at a given temperature, which is necessary to guarantee that the algorithm finds a global optimum. The ASA also uses only the value of the cost function in the optimization process and is very simple to program.

2.3. Algorithm parameter tuning

For the above ASA algorithm, most of the algorithm parameters are automatically set and "tuned", and the user only needs to assign a control

parameter c and set two values N_{accept} and N_{genera}. Obviously, the optimal values of N_{accept} and N_{genera} are problem dependent, but our experience suggests that an adequate choice for N_{accept} is in the range of tens to hundreds and an appropriate value for N_{genera} is in the range of hundreds to thousands. The annealing rate control parameter c can be determined form the chosen initial temperature, final temperature and predetermined number of annealing steps [13,14]. We have found out that a choice of c in the range 1.0 to 10.0 is often adequate.

It should be emphasized that, as the ASA has excellent self adaptation ability, the performance of the algorithm is not critically influenced by the specific chosen values of c, N_{accept} and N_{genera}. This has been observed in a variety of applications, including the three problems given in the following section and some other applications [2,3].

2.4. Statistical convergence to a global minimum

Sufficient conditions for various SA algorithms to statistically converge to a global minimum are known, and these conditions depend on the algorithm annealing schedules being "slow enough" [9,15,29]. For an SA algorithm to statistically converge to a global minimum is to statistically (i.e., requiring many trials) assure that any state in the parameter space can be sampled infinitely often in annealing time. It is informative to consider three cases, namely, the standard SA with the Gaussian generating function [9], the so-called fast SA with the Cauchy generating function [29] and the ASA with the generating function specified by Eqs. (6) and (7) [15].

A necessary and sufficient condition for the standard SA with Gaussian generating function to statistically converge to a global minimum has been proved [9]. It is required that the annealing schedule selected is not faster than

$$T_{i,\text{gen}}(k_i) = \frac{T_{i,\text{gen}}(0)}{\ln(k_i)}. \tag{14}$$

It is also known that the fast SA with Cauchy generating function can statistically converge to a global minimum if its annealing schedule is not faster than [29]

$$T_{i,\text{gen}}(k_i) = \frac{T_{i,\text{gen}}(0)}{k_i}. \tag{15}$$

The ASA, on the other hand, is statistically guaranteed to find a global minimum with an annealing schedule not faster than [15]

$$T_{i,\text{gen}}(k_i) = \frac{T_{i,\text{gen}}(0)}{\exp(ck_i^{1/n})}. \tag{16}$$

The convergence rate of an SA algorithm is determined primarily by its annealing schedule. The slow convergence of the standard SA is inherent from the annealing schedule (14). The ASA can use the much faster annealing schedule (16), which is even faster than what can be used for the fast SA. The ASA therefore has the fastest convergence rate among the three algorithms discussed here.

3. Optimization applications

The versatility of the ASA as a global optimization tool is demonstrated on three very different problems.

3.1. ML joint channel and data estimation

Consider the digital communication channel modelled as a finite impulse response filter with an additive noise source. The received signal at sample k is given by

$$r(k) = \sum_{i=0}^{n_a-1} a_i s(k-i) + e(k), \tag{17}$$

where n_a is the channel length, a_i are the channel taps, the symbol sequence $\{s(k)\}$ is independently identically distributed with an M-PAM symbol constellation, and $e(k)$ is a Gaussian white noise. Let

$$\begin{aligned} \boldsymbol{r} &= [r(1)r(2)\dots r(N)]^{\mathrm{T}}, \\ \boldsymbol{s} &= [s(-n_a+2)\dots s(0)s(1)\dots s(N)]^{\mathrm{T}}, \\ \boldsymbol{a} &= [a_0 a_1 \dots a_{n_a-1}]^{\mathrm{T}} \end{aligned} \tag{18}$$

be the vector of N received data samples, the transmitted data sequence and the channel tap vector, respectively. The joint ML estimate of $\boldsymbol{a}$ and $\boldsymbol{s}$ is obtained by maximizing the conditional probability density function of $\boldsymbol{r}$ given $\boldsymbol{a}$ and $\boldsymbol{s}$. Equivalently, the ML solution is the minimum of the cost function

$$J_c(\boldsymbol{a},\boldsymbol{s}) = \sum_{k=1}^{N}\left(r(k) - \sum_{i=0}^{n_a-1} a_i s(k-i)\right)^2, \tag{19}$$

that is,

$$(\boldsymbol{a}^*,\boldsymbol{s}^*) = \arg\left[\min_{\boldsymbol{a},\boldsymbol{s}} J_c(\boldsymbol{a},\boldsymbol{s})\right]. \tag{20}$$

This joint ML estimate, however, is too expensive to compute except for the simplest case. In practice, suboptimal solutions are adopted for computational purpose. The algorithm based on a blind trellis search technique [27] is such an example.

The joint minimization process (20) can also be performed using an iterative loop first over the data sequences $\boldsymbol{s}$ and then over all the possible channels $\boldsymbol{a}$,

$$(\boldsymbol{a}^*,\boldsymbol{s}^*) = \arg\left[\min_{\boldsymbol{a}}\left(\min_{\boldsymbol{s}} J_c(\boldsymbol{a},\boldsymbol{s})\right)\right]. \tag{21}$$

The inner optimization can be carried out using the Viterbi algorithm (VA). The previous research has used the quantized channel algorithm [33] and the GA [1] to perform the outer optimization. In this study, we apply the ASA to perform the outer optimization. Specifically, given the channel estimate $\hat{\boldsymbol{a}}$, let the data sequence decoded by the VA be $\hat{\boldsymbol{s}}^*$. The cost function used by the ASA is the mean square error (MSE):

$$J(\hat{\boldsymbol{a}}) = \frac{1}{N} J_c(\hat{\boldsymbol{a}},\hat{\boldsymbol{s}}^*). \tag{22}$$

The search range for each channel tap is $-1.0 \leqslant a_i \leqslant 1.0$, since the channel can always be normalized.

The following numerical example was used to illustrate the combined ASA and VA approach for ML joint channel and data estimation. The channel was given by

$$\boldsymbol{a} = [0.407 \quad 0.815 \quad 0.407]^{\mathrm{T}}. \tag{23}$$

Because the true channel length $n_a = 3$ is unknown, an estimated length $\hat{n}_a = 4$ was assumed in the simulation. In practice, the performance of the algorithm is observed through the MSE. In simulation, the performance of the algorithm can also be assessed by the mean tap error (MTE), defined as $\mathrm{MTE} = \|\hat{\boldsymbol{a}} - \boldsymbol{a}\|^2$. Figs. 2 and 3 show the evolutions of the MSE and MTE with 2-PAM and 4-PAM symbols and different signal to noise ratios (SNRs), respectively. All the results were averaged over 100 different runs. Each run had a different noisy received data sequence and a different random initialization of the algorithm. No divergence was observed for any run.

It can be seen from Figs. 2 and 3 that the MSE converged to the noise floor. Compared with the results of using the GA [1], the ASA required a slightly more VA calls but obtained more accurate solution. It should be emphasized that no

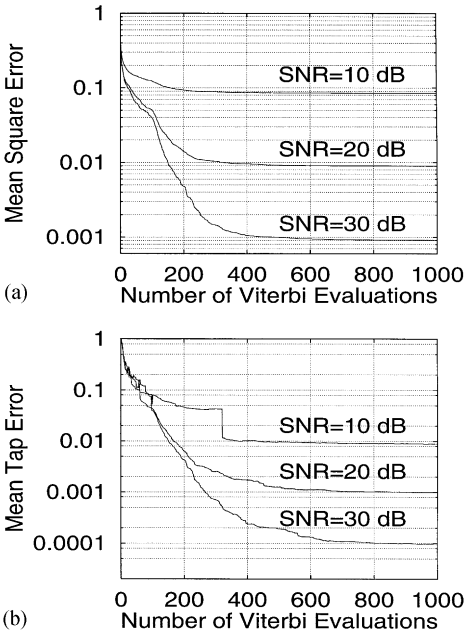

Fig. 2. (a) Mean square error and (b) mean tap error against number of VA evaluations averaged over 100 different runs. 2-PAM and data samples $N = 50$.

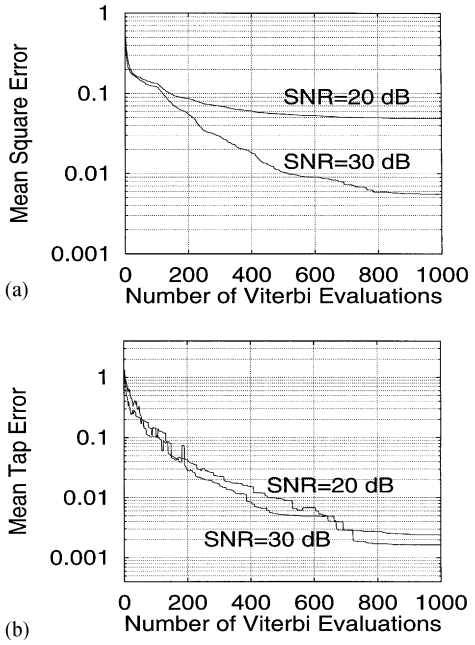

Fig. 3. (a) Mean square error and (b) mean tap error against number of VA evaluations averaged over 100 different runs. 4-PAM and data samples $N = 100$.

other existing methods, except the GA based one [1], have been observed to achieve such an accurate true ML solution of the blind joint channel and data estimation problem based on such a short data sequence.

3.2. IIR filter design

Consider the IIR filter with the input–output relationship governed by

$$y(k) + \sum_{i=1}^{M} b_i y(k-i) = \sum_{i=0}^{L} a_i x(k-i), \tag{24}$$

where $x(k)$ and $y(k)$ are the filter's input and output, respectively, and M ($\geqslant L$) is the filter order. The transfer function of this IIR filter is

$$H_M(z) = \frac{A(z)}{B(z)} = \frac{\sum_{i=0}^{L} a_i z^{-i}}{1 + \sum_{i=1}^{M} b_i z^{-i}}. \tag{25}$$

The IIR filter design can be formulated as an optimization problem with the cost function

$$J(\boldsymbol{w}_H) = \mathrm{E}[e^2(k)] = \mathrm{E}[(d(k) - y(k))^2], \tag{26}$$

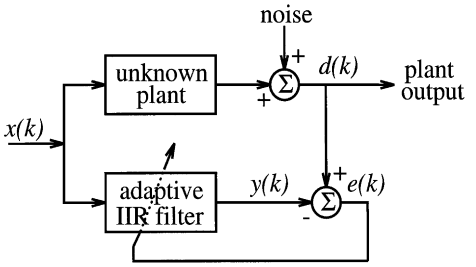

Fig. 4. Schematic of adaptive IIR filter for system identification.

where $d(k)$ is the filter's desired response, $e(k)$ is the filter's error signal and

$$\boldsymbol{w}_H = [\boldsymbol{a}^{\mathrm{T}} \quad \boldsymbol{b}^{\mathrm{T}}]^{\mathrm{T}}$$
$$= [a_0 \quad a_1 \quad \ldots \quad a_L \quad b_1 \quad \ldots \quad b_M]^{\mathrm{T}} \tag{27}$$

denotes the filter coefficient vector. The goal is to minimize the MSE (26) by adjusting $\boldsymbol{w}_H$. In practice, ensemble operation is difficult to realize, and the cost function (26) is usually substituted by the time-averaged cost function

$$J_N(\boldsymbol{w}_H) = \frac{1}{N}\sum_{k=1}^{N} (d(k) - y(k))^2. \tag{28}$$

A major concern in IIR filtering applications is that the cost function of IIR filters is generally multimodal, and a gradient-based algorithm can easily be stuck at local minima. The GA has been applied to IIR filter design (e.g. [19,20,32]) to overcome this difficulty. We demonstrate that the ASA offers an alternative to IIR filter design. To maintain the stability during optimization, we convert the direct-form coefficients b_i, $1 \leqslant i \leqslant M$, to the lattice-form reflection coefficients κ_i, $0 \leqslant i \leqslant M-1$, and make sure that all the κ_i have magnitudes less than 1. Thus the filter coefficient vector used in optimization is $\boldsymbol{w} = [a_0\, a_1 \ldots a_L\, \kappa_0 \ldots \kappa_{M-1}]^{\mathrm{T}}$. Converting the reflection coefficients back to the direct-form coefficients is straightforward [11].

System identification application, depicted in Fig. 4, is used in the experiment. The unknown plant has a transfer function $H_S(z)$, and the ASA is employed to adjust the IIR filter that is used to model the system. When the filter order M is

smaller than the system order, local minima problems can be encountered [28], and this is used to simulate a multimodal environment. Two examples were tested.

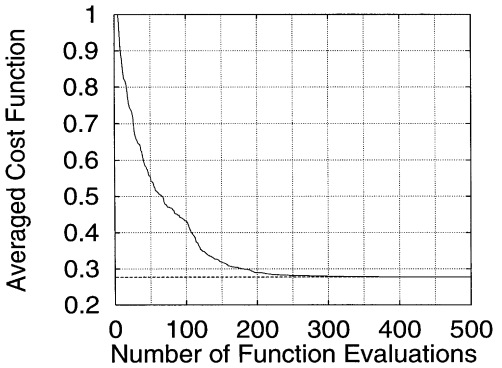

Fig. 5. Normalized cost function versus number of cost function evaluations averaged over 100 random runs of the ASA for Example 1. The dashed line indicates the global minimum.

Example 1. This example is taken from [28]. The system and filter transfer functions are respectively

$$H_S(z) = \frac{0.05 - 0.4z^{-1}}{1 - 1.1314z^{-1} + 0.25z^{-2}}$$

and (29)

$$H_M(z) = \frac{a_0}{1 + b_1 z^{-1}}.$$

The analytical MSE (26) in this case is known when the input is a white sequence and the noise is absent. The MSE has a global minimum at $\boldsymbol{w}^{\text{global}} = [-0.311 \quad -0.906]^{\text{T}}$ with the value of the normalized MSE 0.2772, and a local minimum at $\boldsymbol{w}^{\text{local}} = [0.114 \quad 0.519]^{\text{T}}$. Fig. 5 depicts the evolution of the normalized MSE averaged over 100 different runs of the ASA. Each run had a randomly chosen initial $\boldsymbol{w}$ and a random algorithm setting. Fig. 6 shows the trajectories of the filter parameter

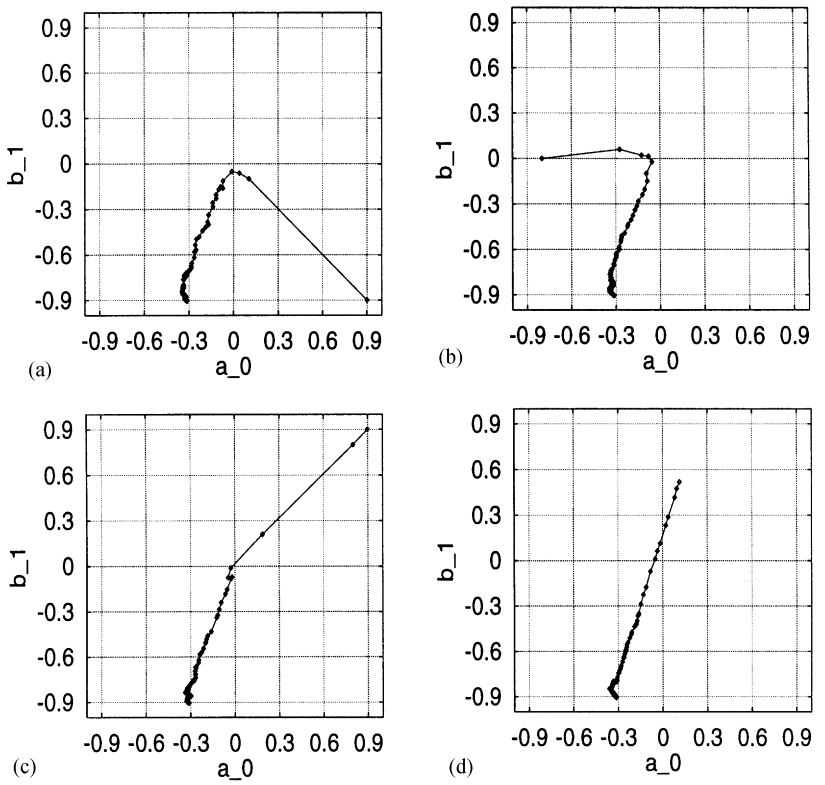

Fig. 6. Trajectories of the filter parameter vector averaged over 100 different runs of the ASA, started from the fixed initial positions: (a) $[0.9 \quad -0.9]^{\text{T}}$, (b) $[-0.8 \quad 0.0]^{\text{T}}$, (c) $[0.9 \quad 0.9]^{\text{T}}$ and (d) $[0.114 \quad 0.519]^{\text{T}}$, for Example 1.

vector averaged over 100 different runs of the ASA, started from four fixed initial positions. It can be seen that the ASA consistently found the global optimal solution.

Example 2. This is a third-order system with the transfer function given by

$$H_S(z) = \frac{-0.3 + 0.4z^{-1} - 0.5z^{-2}}{1 - 1.2z^{-1} + 0.5z^{-2} - 0.1z^{-3}}. \quad (30)$$

The system input was a uniform white sequence taking values in $[-0.5, \quad 0.5]$, and the SNR was 30 dB. The data length used in calculating the MSE (28) was $N = 2000$. When a reduced-order filter with $M = 2$ and $L = 1$ was used, the MSE was multimodal. Extensive simulation showed that the MSE had a global minimum of 0.059. The ASA consistently reached this global minimum, as shown in Fig. 7. To illustrate the multimodal nature of the cost function, Fig. 8 shows the behaviours of a standard gradient algorithm when started from two quite closed initial positions.

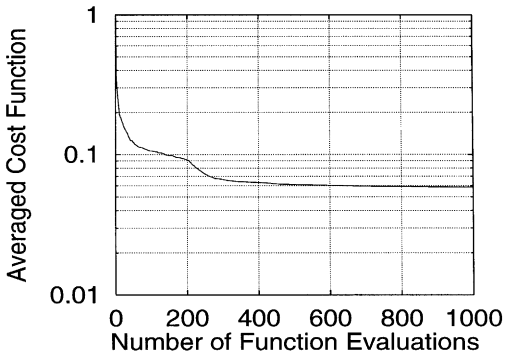

Fig. 7. Cost function versus number of cost function evaluations averaged over 100 random runs of the ASA for Example 2.

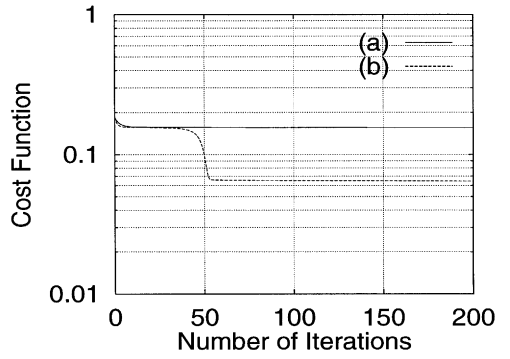

Fig. 8. Convergence behaviours of the gradient algorithm, started from the two initial conditions: (a) $[0.0 \quad 0.0 \quad 0.3 \quad 0.1]^T$ and (b) $[0.0 \quad 0.0 \quad 0.3 \quad 0.0]^T$, for Example 2.

The IIR filter design has been an active research topic and a variety of techniques have been developed. The two examples used here are well-known benchmark problems, and the GA method has been applied to them [20,30]. Compared with the results of using GAs for adaptive IIR filtering available in the literature, the efficiency of the ASA appears to be in the same order as GAs.

3.3. MSER decision feedback equalizer

Equalization is a powerful technique for combating distortion and interference in communication links [22,24] and high-density data storage systems [18,23]. The DFE, in particular, is widely used in practice as it provides a good balance between performance and complexity. A generic DFE structure is shown in Fig. 9. For the channel given by (17), the DFE produces an estimate $\hat{s}(k-d)$ of $s(k-d)$ by quantizing the filter output

$$f(\boldsymbol{r}(k), \hat{\boldsymbol{s}}_b(k)) = \boldsymbol{w}^T\boldsymbol{r}(k) + \boldsymbol{b}^T\hat{\boldsymbol{s}}_b(k), \quad (31)$$

where $\boldsymbol{w} = [w_0 \dots w_{m-1}]^T$ and $\boldsymbol{b} = [b_1 \dots b_n]^T$ are the coefficients of the feedforward and feedback filters, respectively, $\boldsymbol{r}(k) = [r(k) \dots r(k-m+1)]^T$ is the channel output vector and $\hat{\boldsymbol{s}}_b(k) = [\hat{s}(k-d-1) \dots \hat{s}(k-d-n)]^T$ is the past-detected symbol vector. It is sufficient to choose the decision

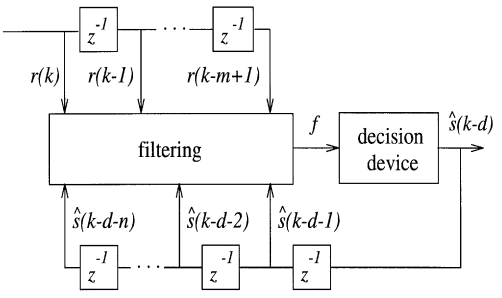

Fig. 9. Schematic of a generic DFE.

delay $d = n_a - 1$, the feedforward order $m = n_a$ and the feedback order $n = n_a - 1$ (see [4,5]).

The Wiener or minimum MSE (MMSE) solution $(\boldsymbol{w}_{\mathrm{MMSE}}, \boldsymbol{b}_{\mathrm{MMSE}})$ is often said to be the optimal solution for the coefficients of the DFE [6]. However, the MMSE solution does not correspond to the MSER solution, the symbol error rate (SER) being the ultimate performance criterion of equalization. It can be shown that the decision feedback $\boldsymbol{b}_{\mathrm{MMSE}}^{\mathrm{T}}\hat{\boldsymbol{s}}_b(k)$ in a DFE performs a space translation which maps the DFE onto a "linear" equalizer in the translated observation space [4,5]:

$$f'(\boldsymbol{r}'(k)) = \boldsymbol{w}^{\mathrm{T}}\boldsymbol{r}'(k), \tag{32}$$

where $\boldsymbol{r}'(k)$ is the translated observation vector. This equivalent DFE is depicted in Fig. 10. In principle given the channel model, the analytic expression of the SER, $P_E(\boldsymbol{w})$, for the DFE with the weight vector $\boldsymbol{w}$ can be derived and the MSER solution $\boldsymbol{w}_{\mathrm{MSER}}$ can be obtained by minimizing $P_E(\boldsymbol{w})$. Furthermore, it becomes clear that the Wiener solution is not the MSER solution.

For the 2-PAM case, the gradient algorithm has been used to optimize $P_E(\boldsymbol{w})$ to obtain the MSER solution [4,5]. For high-order PAM channels, however, the optimization based on gradient method becomes highly complex and very costly. We propose a Monte Carlo approach based on the ASA to achieve the MSER solution for the general M-PAM channel. The SER can be approximated as

$$\hat{P}_E(\boldsymbol{w}) = \frac{1}{N}\sum_{k=1}^{N} \delta(\hat{s}(k-d) - s(k-d)), \tag{33}$$

where the indicator function

$$\delta(\hat{s}(k-d) - s(k-d)) = \begin{cases} 0, & \hat{s}(k-d) = s(k-d), \\ 1, & \hat{s}(k-d) \neq s(k-d) \end{cases} \tag{34}$$

and N is the number of training data. The MSER solution $\boldsymbol{w}_{\mathrm{MSER}}$ is obtained by minimizing $\hat{P}_E(\boldsymbol{w})$ using the ASA algorithm.

The proposed Monte Carlo approach was tested using the following example:

Channel $\boldsymbol{a} = [0.3482\ 0.8704\ 0.3482]^{\mathrm{T}}$

and (35)

4-PAM symbols.

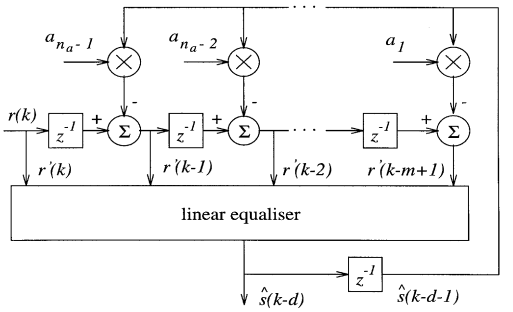

Fig. 10. Schematic of the translated DFE.

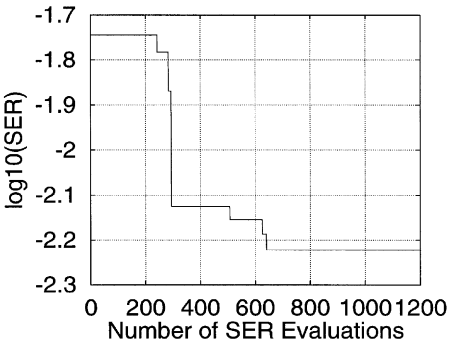

Fig. 11. Evolution of the training SER versus number of SER evaluations. SNR = 20 dB and the initial $\boldsymbol{w}$ is the Wiener solution.

For the case of SNR = 20 dB and $N = 2000$ with the initial $\boldsymbol{w}$ set to the $\boldsymbol{w}_{\mathrm{MMSE}}$, Fig. 11 depicted the evolution of the training SER. The SERs obtained by the MSER DFE and the MMSE DFE with detected symbols being fed back are compared in Fig. 12. The number of training data N for evaluating $\hat{P}_E(\boldsymbol{w})$ ranged from 100 to 100,000, depending on the SNR. The number of symbols used to estimate the SERs of the trained equalizers, shown in Fig. 12, was sufficiently large to produce 400 error counts for each SNR.

As expected, the MSER solution is superior over the MMSE solution. This Monte Carlo algorithm is block-data based, and is particularly suited for data storage systems, as in many commercial disk drives, the equalizers are trained at the factory floor and then are "frozen" before shipping. For communication links, the approach is suitable for the initial set up of the DFE. For the cases of high

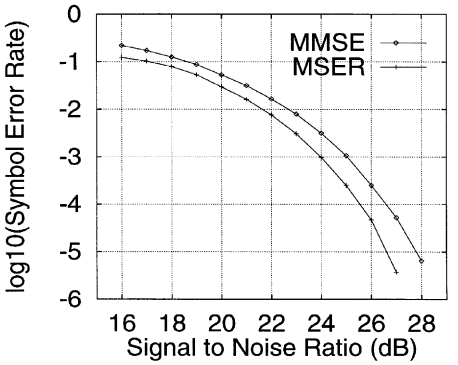

Fig. 12. Performance comparison for the MSER and MMSE DFEs with detected symbols being fed back.

SNR, the conventional Monte Carlo method will require a large number of training samples. However, this problem can easily be overcome by adopting importance sampling techniques [17,21,26] in estimating the SER.

4. Concluding remarks

The ASA is a global optimization technique having certain advantages. The algorithm is versatile and very easy to program, and has very few parameters that require tuning. In this paper, we have provided a detailed description of the ASA, including how the algorithm works and the choices of the algorithm parameters. With the aim of introducing the ASA to the signal processing community, we have applied the ASA to three different signal processing problems, ML joint channel and data estimation, IIR filter design and evaluation of the MSER DFE. Our results have demonstrated that the ASA provides a viable alternative to better known GAs for solving diverse signal processing applications with multimodal and nonsmooth cost functions.

References

[1] S. Chen, Y. Wu, Maximum likelihood joint channel and data estimation using genetic algorithms, IEEE Trans. Signal Process. 46 (5) (1998) 1469–1473.

[2] S. Chen, B.L. Luk, Y. Liu, Application of adaptive simulated annealing to blind channel identification with HOC fitting, Electronics Letters 34 (3) (1998) 234–235.

[3] S. Chen, J. Wu, R.H. Istepanian, Optimizing stability bounds of finite-precision PID controller structures, IEEE Trans. Automatic Control, 2000, to appear.

[4] S. Chen, E.S. Chng, B. Mulgrew, G. Gibson, Minimum-BER linear-combiner DFE, in: Proc. ICC'96, Dallas, Texas, 1996, Vol. 2, pp. 1173–1177.

[5] S. Chen, B. Mulgrew, E.S. Chng, G. Gibson, Space translation properties and the minimum-BER linear-combiner DFE, IEE Proc. Commun. 145 (5) (1998) 316–322.

[6] J.M. Cioffi, G.P. Dudevoir, M.V. Eyuboglu, G.D. Forney, MMSE decision-feedback equalizers and coding – Part I: Equalization results, IEEE Trans. Commun. 43 (10) (1995) 2582–2594.

[7] A. Corana, M. Marchesi, C. Martini, S. Ridella, Minimizing multimodal functions of continuous variables with the simulated annealing algorithm, ACM Trans. Mathematical Software 13 (3) (1987) 262–280.

[8] L. Davis (Ed.), Handbook of Genetic Algorithms, Van Nostrand Reinhold, 1991.

[9] S. Geman, D. Geman, Stochastic relaxation, Gibbs distribution and the Bayesian restoration in images, IEEE Trans. Pattern Analysis and Machine Intelligence 6 (1984) 721–741.

[10] D.E. Goldberg, Genetic Algorithms in Search, Optimization, and Machine Learning, Addison-Wesley, Reading, MA, 1989.

[11] A.H. Gray Jr., J.D. Markel, Digital lattice and ladder filter synthesis, IEEE Trans. Audio and Electroacoustics AU-21 (6) (1973) 491–500.

[12] J.H. Holland, Adaptation in Natural and Artificial Systems, University of Michigan Press, Ann Arbor, Michigan, 1975.

[13] L. Ingber, Simulated annealing: practice versus theory, Mathematical and Computer Modelling 18 (11) (1993) 29–57.

[14] L. Ingber, Adaptive simulated annealing (ASA): lessons learned, J. Control and Cybernetics 25 (1) (1996) 33–54.

[15] L. Ingber, B. Rosen, Genetic algorithms and very fast simulated reannealing: a comparison, Mathematical and Computer Modelling 16 (11) (1992) 87–100.

[16] S. Kirkpatrick, C.D. Gelatt Jr., M.P. Vecchi, Optimization by simulated annealing, Science 220 (4598) (1983) 671–680.

[17] D. Lu, K. Yao, Improved importance sampling technique for efficient simulation of digital communication systems, IEEE J. Selected Areas in Commun. 6 (1) (1988) 67–75.

[18] J. Moon, The role of SP in data-storage systems, IEEE Signal Processing Magazine 15 (4) (1998) 54–72.

[19] R. Nambiar, C.K.K. Tang, P. Mars, Genetic and learning automata algorithms for adaptive digital filters, in: Proc. IEEE ICASSP, 1992, Vol. IV, pp. 41–44.

[20] S.C. Ng, S.H. Leung, C.Y. Chung, A. Luk, W.H. Lau, The genetic search approach: a new learning algorithm for adaptive IIR filtering, IEEE Signal Processing Magazine 13 (6) (1996) 38–46.
[21] G. Orsak, B. Aazhang, On the theory of importance sampling applied to the analysis of detection systems, IEEE Trans. Commun. 37 (4) (1989) 332–339.
[22] J.G. Proakis, Digital Communications, 3rd edition, McGraw-Hill, New York, 1995.
[23] J.G. Proakis, Equalization techniques for high-density magnetic recording, IEEE Signal Processing Magazine 15 (4) (1998) 73–82.
[24] S.U.H. Qureshi, Adaptive equalization, in: Proc. IEEE 73 (9) (1985) 1349–1387.
[25] B.E. Rosen, Rotationally parameterized very fast simulated reannealing, IEEE Trans. Neural Networks, 1997, submitted.
[26] J.S. Sadowsky, J.A. Bucklew, On large deviations theory and aysmptotically efficient Monte Carlo estimation, IEEE Trans. Inform. Theory 36 (3) (1990) 579–588.
[27] N. Seshadri, Joint data and channel estimation using blind trellis search techniques, IEEE Trans. Commun. 42 (2/3/4) (1994) 1000–1011.
[28] J.J. Shynk, Adaptive IIR filtering, IEEE ASSP Magazine (April 1989) 4–21.
[29] H. Szu, R. Hartley, Fast simulated annealing, Physics Letters A 122 (1987) 157–162.
[30] K.S. Tang, K.F. Man, S. Kwong, Q. He, Genetic algorithms and their applications, IEEE Signal Processing Magazine 13 (6) (1996) 22–37.
[31] P.J.M. van Laarhoven, E.H.L. Aarts, Simulated Annealing: Theory and Applications, Reidel, Dordrecht, The Netherlands, 1987.
[32] P.B. Wilson, M.D. Macleod, Low implementation cost IIR digital filter design using genetic algorithms, in: Workshop on Natural Algorithms in Signal Processing, Chelmsford, Essex, 14–16 November 1993, pp. 4/1–4/8.
[33] E. Zervas, J. Proakis, V. Eyuboglu, A quantized channel approach to blind equalization, in: Proc. ICC'92, Chicago, 1992, Vol. 3, pp. 351.8.1–351.8.5.