Available online at www.sciencedirect.com
SCIENCE DIRECT®
International Journal of Solids and Structures 42 (2005) 203–237
INTERNATIONAL JOURNAL OF SOLIDS and STRUCTURES
www.elsevier.com/locate/ijsolstr

Improved global–local simulated annealing formulation for solving non-smooth engineering optimization problems

K. Genovese [a], L. Lamberti [b,*], C. Pappalettere [b,1]

[a] *Dipartimento di Ingegneria Fisica e dell'Ambiente, Università degli Studi della Basilicata Viale dell'Ateneo Lucano 10, 85100 Potenza, Italy*

[b] *Dipartimento di Ingegneria Meccanica e Gestionale, Politecnico di Bari, Viale Japigia 182, 70126 Bari, Italy*

Received 30 September 2003; received in revised form 19 July 2004
Available online 11 September 2004

Abstract

This paper is concerned with a novel optimization algorithm that implements an enhanced formulation of simulated annealing (SA). The new algorithm is denoted as ISA (improved simulated annealing) in the rest of the paper. ISA includes a two-level random search: "global annealing" where all design variables are perturbed simultaneously and "local annealing" where design variables are perturbed one at a time.

The improvement with respect to classical SA is in the fact that trial designs are generated always taking care to choose directions along which the cost function may improve. To this purpose, cost function sensitivities are computed in order to properly choose the size of each random perturbation. In addition, the optimization problem is linearized about the current design point if the optimizer ends up in an infeasible region or there is no significant reduction in cost even though the cost function gradient is not close to zero. The linearization is controlled by a trust region model. The optimization algorithm continuously shifts from global to local annealing based on the current best record at the beginning of each cooling cycle. Finally, the cooling schedule is automatically adjusted within ISA based on the convergence behavior.

In this work, the ISA algorithm is successfully utilized to solve complicated optimization problems which exhibit non-smooth/non-convex behavior: (i) the large-scale (200 design variables and 3500 constraints) weight minimization of a 200 bar truss under five independent loading conditions; (ii) the configuration optimization of a cantilevered bar truss with 45 elements and 81 design variables; (iii) an example of reverse engineering where in-plane elastic properties of an eight-ply woven composite laminate are to be determined.

The performance of ISA is compared to that of classical SA, gradient based optimization codes recently published in literature and commercial software. The results obtained in this study indicate that ISA is a very efficient optimization code. In fact, ISA was much faster than classical SA. The present code allowed about 300 kg weight saving in the

* Corresponding author. Tel.: +39 080 596 2774; fax: +39 080 596 2777.
E-mail address: lamberti@poliba.it (L. Lamberti).

[1] President of the Italian Society of Stress Analysis (AIAS).

0020-7683/$ - see front matter
doi:10.1016/j.ijsolstr.2004.07.015

200 bar truss case and about 80 kg in the cantilevered bar truss case. In addition, the residual error on elastic constants in the material identification problem was less than 3%.

Keywords: Simulated annealing (SA); Optimization; Reverse engineering

1. Introduction

Engineering optimization problems have often to deal with non-smooth and/or non-convex design spaces. Consequently, many optimization methods may become not attractive since they get stuck in local optima or are computationally very expensive. In general, gradient based optimizers (GBOs) are capable to reach an optimum design rather quickly but there is no guarantee on that the attained design will really be the true global optimum. On the other hand, non-gradient based optimizers (NGBOs) allow us to perform the optimum search in a zone of design space significantly larger than in the GBO case but at a computational cost that is often unaffordable. Hence, a reasonable compromise is to use an optimization algorithm, which explores large fractions of design space but uses also gradient information in order to speed up the design process.

Amongst global optimization techniques, *simulated annealing* (SA) is perhaps the method with the widest variety of applications in problems taken from different disciplines. The reader can refer to the recent reviews included in the works of Yu Chen and Su (2002) and Hasancebi and Erbatur (2002). As is quoted in Rao (1996), simulated annealing is a stochastic technique to find a global minimizer for non-linear optimization problems. The basic idea of the method is to generate randomly a trial point and evaluate the problem functions. If the trial point is infeasible, it is rejected and a new trial point is evaluated. If the trial point is feasible and the cost function is smaller than the current best record, then the point is accepted and the best record is updated. If the trial point is feasible but the cost function is higher than the best value, then the point is accepted or rejected based on a probabilistic criterion, which estimates if design may improve in the next function evaluations. In order to compute probability, a parameter called the *temperature* is utilized. In the optimization problem, temperature can be a target value (estimated) for the cost function corresponding to a global minimizer. Initially, a larger target value is selected. As the trials progress, the target value is reduced based on a *cooling schedule*. The acceptance probability steadily decreases to zero as the temperature is reduced.

Researchers proposed many different schemes for generating randomly the trial points. A comprehensive review of those strategies can be found in Blachut (2003). In general, any random generation mechanisms used in simulated annealing may be classified in two main groups: the *1-directional search* where the design variables are perturbed one at a time and the *multi-directional search* where all design variables are perturbed simultaneously. The latter approach allows to increase the convergence speed but may fail in finding the global optimum. On the other hand, the former approach may result in too large computational times. For this reason, the present authors carried out a trade study on two non-convex optimization problems (Banana function with 50 design variables and weight minimization of a 10 bar truss) in order to choose a random generation mechanism which ensured enough robustness of the SA optimizer and a reasonable convergence speed. It was found that 1-directional search was superior over multi-directional search. More details on the comparison of different annealing schemes are given in Appendix A. Indeed, the trade study finally served to write an optimization code—based on "classical" simulated annealing—robust enough to be used as comparison basis with a new SA algorithm developed in this research.

It is apparent that the main deficiencies of classical SA are the large number of trials which result in just marginal improvements in design and the fact that infeasible designs are rejected with obvious waste of

computational resources. The latter fact is a serious drawback in optimization problems where each generation of a new trial design involves finite element analysis. In order to overcome these limitations, the present authors developed a new optimization algorithm based on an improved formulation of simulated annealing. The new algorithm is denoted as ISA (improved simulated annealing) in the rest of this paper.

The ISA algorithm combines the global optimization capability of SA with the search for feasible directions. ISA includes a two-level random search: ''global annealing'' where all design variables are perturbed simultaneously and ''local annealing'' where design variables are perturbed one at a time. The optimization algorithm presented here continuously shifts from global to local annealing based on the nominal design at the beginning of each cooling cycle. Local annealing requires constraint linearization. The accuracy of approximation is controlled by a trust region model. Since local annealing is performed only if global annealing fails, constraint gradients are evaluated not at each cooling cycle. Therefore, the ISA optimizer is basically a SA algorithm with infrequent gradient calculation. The mean features of ISA can be summarized as follows.

If the design at the beginning of a cooling cycle is feasible, a new trial design point P_{TR} is defined by perturbing all the optimization variables simultaneously. Each new trial point P_{TR} defines a new trial search direction $\mathbf{S}_{TR}$ whose components are the perturbations given to each design variable. The $\mathbf{S}_{TR}$ direction is chosen so that it is always a descent direction. This is done because the total change in cost is the sum over changes in cost when design variables are changed one at a time. As is clear, this strategy serves to reduce the number of trial designs generated in the optimization process.

If the optimizer enters in an infeasible region, the design is perturbed by taking movements along directions where the constraint violation may be reduced. In order to steer the design back to a feasible domain, constraint functions are linearized and trial designs are evaluated in the approximate model. The trial design which violates linearized constraints the least is taken as the starting point for a new search. This strategy serves to reduce the computational cost of optimization when trial points are infeasible. The reliability and accuracy of the approximate model are ensured by a trust region model.

Moreover, ISA activates improvement routines based on a quadratic approximation of cost function and constraints when a trial point is better than the current best record but violates the constraints of the optimization problem. Finally, the cooling schedule is adaptively changed by ISA during the optimization process based on the improvement in design obtained in the current annealing cycle.

In order to check on the performance of the ISA algorithm, complicated optimization problems have been solved in this study. The first test case is the weight minimization of a 200 bar truss under five independent loading conditions. This large-scale problem (200 design variables and 3500 constraints) is such that GBOs could not find designs consistent with the amount of design freedom included in the optimization process (see the work of Lamberti and Pappalettere, 2003a). The second test case is the configuration optimization of a cantilevered bar truss with 45 members and 81 design variables. The complication is in the fact that the optimization variables belong to two well distinct design spaces: size and configuration.

The last test case is a typical example of reverse engineering. The in-plane elastic properties of an eight-ply composite laminate (woven fiberglass-epoxy) used as substrate for printed circuit boards are to be determined. To this purpose, the reverse engineering problem of in-plane material characterization is transformed into an optimization problem where the goal is to minimize the difference Ψ between the displacements computed by means of finite element analyses and the displacements measured experimentally by means of a powerful non-contact optical technique known as *phase shifting electronic speckle pattern interferometry* (PS-ESPI) (Cloud, 1998; Creath, 1985). Although this problem includes only four design variables and there are no severe constraints, the optimization process is complicated by the fact that the cost function is highly non-linear and the sensitivity values may even be significantly different when design variables are smaller or greater than their optimum value.

The relative merits of the ISA algorithm, state-of-art GBOs recently published in literature, classical SA and commercial software have been compared. Results indicate that ISA proved itself to be very efficient.

In fact, the present code allowed to save about 300 kg in the 200 bar truss case with respect to the best solution quoted in literature (i.e., the LESLP algorithm described in Lamberti and Pappalettere, 2003a). In addition, ISA improved the design of the cantilevered bar truss by about 80 kg. Finally, the residual error on the in-plane displacement field found in the material identification problem was less than 3%.

The present paper is divided in five sections and an appendix. After the Introduction section, we describe the different steps of the ISA optimization algorithm in Section 2. Section 3 presents the two test cases; the section includes also a summary of the PS-ESPI optical technique and the description of the experimental set-up used in the material identification procedure. Section 4 analyzes the results obtained in the optimization runs and discusses the relative merits of ISA and the other optimization codes considered in this study. Finally, Section 5 summarizes the work presented in this paper and reports the main findings of the study. The appendix discusses the relative merits of classical annealing schemes.

2. The ISA algorithm

A non-linear optimization problem may be formulated as follows:

$$\begin{cases} \min W(x_1, x_2, \ldots, x_N), \\ G_k(x_1, x_2, \ldots, x_N) \leqslant 0, \\ x_j^{\mathrm{l}} \leqslant x_j \leqslant x_j^{\mathrm{u}}, \end{cases} \quad \begin{cases} j = 1, \ldots, N, \\ k = 1, \ldots, \mathrm{NC}, \end{cases} \tag{1}$$

where

- $(x_1, x_2, \ldots, x_N)$ are the N design variables;
- $W(x_1, x_2, \ldots, x_N)$ is the objective function;
- $G_k(x_1, x_2, \ldots, x_N)$ are the NC inequality constraint functions;
- x_j^{l} and x_j^{u} are the lower and upper bounds of the jth design variable.

The pseudo code of the ISA algorithm is now provided. Each step is described in detail in order to make potential users able to code the ISA algorithm on computers. The flow chart of ISA is shown in Fig. 1.

1. Start the optimization process. Choose the initial design vector $\mathbf{X}_0(x_{1,0}, x_{2,0}, \ldots, x_{N,0})$ and set it as the current best record $\mathbf{X}_{\mathrm{OPT}}$. Define the corresponding point P_{OPT} in the design space.

In general, the initial design $\mathbf{X}_0$ should be feasible and pretty far away from constraint domain boundaries in order to explore a zone of design space approximately centered about $\mathbf{X}_0$ with no risk to generate infeasible points that might bias the optimum design search since the very beginning of the optimization process. This strategy serves also to reduce the number of constraint evaluations eventually done in the infeasible region when design variables are perturbed one by one.

Set the K counter of cooling cycles as $K = 1$ and choose the limit number of cooling cycles as $K_{\mathrm{LIM}} = 100$.

Set the I_{GLOB} counter of the global annealing cycles eventually performed by ISA in a cooling cycle as $I_{\mathrm{GLOB}} = 0$. Set the I_{SEC} counter of the inner loops eventually performed by ISA in a cooling cycle as $I_{\mathrm{SEC}} = 0$. Choose the limit number of global annealing cycles and inner loops as $I_{\mathrm{LIM}} = 5$. The reason why the number of global annealing cycles is taken equal to the number of inner loops in each cooling cycle is the following. Each global annealing cycle requires N function evaluations in order to compute gradients. Hence, ISA will require $\mathrm{NFV} = I_{\mathrm{LIM}} \times N$ function evaluations in one cooling cycle if local annealing is never used. However, such value of NFV is also the number of function evaluations required by classical implementations of simulated annealing in order to complete I_{LIM} secondary cycles (i.e., inner loops). Since

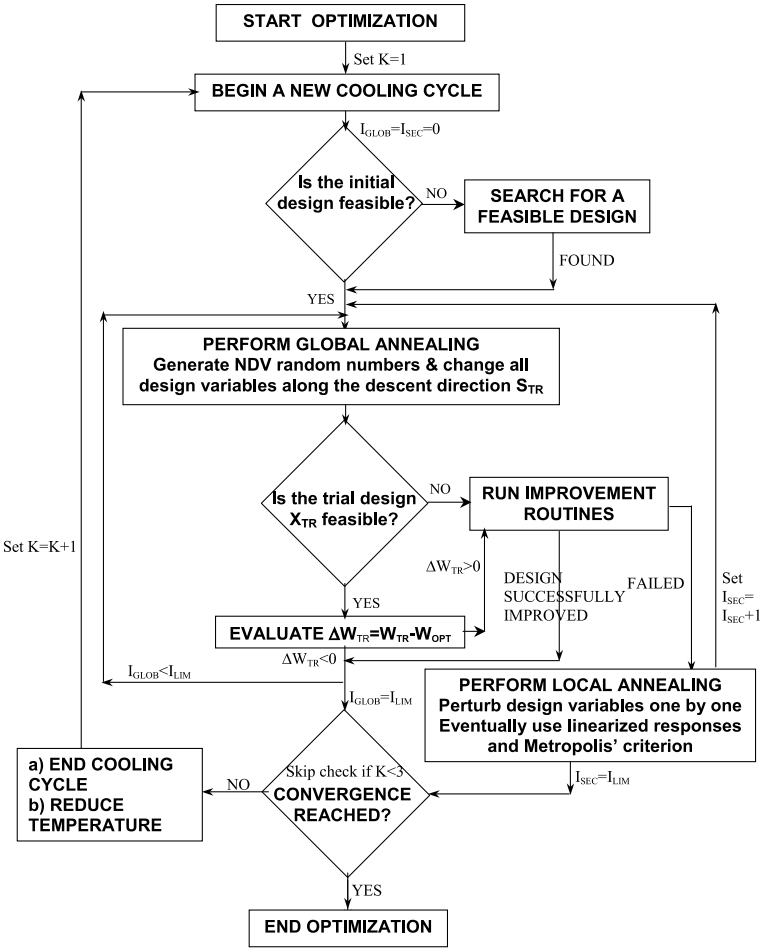

Fig. 1. Flow chart of the ISA algorithm.

the task of calculating gradients is less expensive than completing a single inner loop, to limit global annealing cycles to I_{LIM} is a rather conservative strategy in terms of CPU time.

2. If the optimization problem is constrained and the starting design is feasible, set the initial temperature T_0 as about 10% of the initial cost. If there are "soft" non-linear constraints or the optimization problem is unconstrained, use a very large value of T_0. The two strategies are practically equivalent because the cost of a "very feasible" initial design (i.e., located very far away from the constraint domain boundaries) is usually much more than 10 times as large as the final optimum cost.

If the starting point violates constraints, set the initial temperature T_0 as 10 times the cost corresponding to the first feasible trial design $P_{\mathrm{FEA},1}$ generated by ISA. The rationale behind this strategy is the following. In order to steer the design back to a feasible region, ISA perturbs the nominal design by taking movements along a series of directions whose components depend from the constraint gradients (see Step 7). If some of these directions are not descent, the cost $W_{\mathrm{FEA},1}$ computed at the $P_{\mathrm{FEA},1}$ point will be larger than the cost W_0 computed at the starting point. In such a case, ISA will try to reduce $W_{\mathrm{FEA},1}$ in order to minimize the cost function $W(\mathbf{X})$. If the $P_{\mathrm{FEA},1}$ point is pretty far away from the constraint domain boundary, the $W_{\mathrm{FEA},1}$ cost will be much higher than the final optimum cost and hence ISA will easily reduce $W(\mathbf{X})$ by performing

global annealing cycles (see Step 4). Conversely, if the $P_{\mathrm{FEA},1}$ point is close to constraint boundaries, the $W_{\mathrm{FEA},1}$ may be comparable to the minimum cost. Hence, ISA will have to evaluate an acceptance/rejection probability function in order not to get stuck in local minima. Since this is done in the very early stages of the optimization process, the temperature T_0 should be high enough to ensure exploration of a sufficiently large zone of the design space.

3. Evaluate non-linear constraints at the current best record $\mathbf{X}_{\mathrm{OPT}}$. If constraints are satisfied execute Step 4. Conversely, if constraints are violated execute Step 7.
4. *Global annealing. All design variables are perturbed simultaneously.*

Set $I_{\mathrm{GLOB}} = I_{\mathrm{GLOB}} + 1$. Evaluate the gradient $\overline{\nabla} W(\mathbf{X}_{\mathrm{OPT}})$ of the cost function $W(\mathbf{X})$ at the P_{OPT} point. Perturb randomly each design variable x_j $(j = 1, \ldots, N)$ so that $(\partial W/\partial x_j)\,\Delta x_j < 0$. Each movement Δx_j is calculated as follows:

$$\begin{aligned} &\partial W/\partial x_j > 0 \;\Rightarrow\; \Delta x_j = -(x_{\mathrm{OPT},j} - x_j^{\mathrm{l}}) \cdot N_{\mathrm{RND},j} \cdot \gamma_K \cdot (1 + \mu_j), \\ &\partial W/\partial x_j < 0 \;\Rightarrow\; \Delta x_j = (x_j^{\mathrm{u}} - x_{\mathrm{OPT},j}) \cdot N_{\mathrm{RND},j} \cdot \gamma_K \cdot (1 + \mu_j), \end{aligned} \qquad (j = 1, \ldots, N), \tag{2}$$

where each weighting coefficient μ_j is defined as $|\partial W/\partial x_j|/\|\overline{\nabla} W(\mathbf{X}_{\mathrm{OPT}})\|$. The purpose of μ_j is to adjust the Δx_j movement based on the contribution that the jth sensitivity gives to the magnitude of cost function gradient. Design variables are changed following their order sequentially.

Define the descent direction $\mathbf{S}_{\mathrm{TR}}(\Delta x_1, \Delta x_2, \ldots, \Delta x_N)$ and the corresponding trial point of the design space $P_{\mathrm{TR}}(x_{\mathrm{OPT},1} + \Delta x_1, x_{\mathrm{OPT},2} + \Delta x_2, \ldots, x_{\mathrm{OPT},N} + \Delta x_N)$. Let $\mathbf{X}_{\mathrm{TR}}(x_{\mathrm{TR},1}, x_{\mathrm{TR},2}, \ldots, x_{\mathrm{TR},N})$ be the trial design vector containing the co-ordinates (i.e., the design variables) of the P_{TR} point. Hence, the design variables are "temporary" updated as follows:

$$x_{\mathrm{TR},j} = x_{\mathrm{OPT},j} + \Delta x_j \quad (j = 1, \ldots, N). \tag{3}$$

The $N_{\mathrm{RND},j}$ parameter in expression (2) is a random number chosen in the interval $(0, 1)$. The γ_K factor is computed as follows:

$$\gamma_K = \max\left[T_K/T_0, \alpha_K = \max(0.01, 0.2 \cdot (0.9)^{K-1})\right], \tag{4}$$

where the α_K parameter is a knockdown factor introduced by Huang and Arora (1997) in order to keep all the optimization variables inside the range defined by their lower and upper bounds. If it happens that $x_{\mathrm{TR},j} < x_j^{\mathrm{l}}$ or $x_{\mathrm{TR},j} > x_j^{\mathrm{u}}$, the Δx_j movement in expression (3) is reset to $\Delta x_j = \Delta x_j/(1 + \mu_j)$.

The γ_K factor is initially equal to 1 and is shrunk as the optimization progresses.

5. Evaluate cost function and constraints at the trial point P_{TR} generated in Step 4. Let W_{TR} denote the corresponding value of the cost function. Calculate $\Delta W_{\mathrm{TR}} = W_{\mathrm{TR}} - W_{\mathrm{OPT}}$. Based on the ΔW_{TR} value and on whether constraints are satisfied or not, different scenarios may occur.
   - 5A. If it holds $\Delta W_{\mathrm{TR}} < 0$ and the P_{TR} point is feasible, accept the design $\mathbf{X}_{\mathrm{TR}}$ as the new optimum. Hence, set $\mathbf{X}_{\mathrm{OPT}} \equiv \mathbf{X}_{\mathrm{TR}}$ and $W_{\mathrm{OPT}} = W_{\mathrm{TR}}$. If $I_{\mathrm{GLOB}} < I_{\mathrm{LIM}}$, go to Step 4. Otherwise, if $I_{\mathrm{GLOB}} = I_{\mathrm{LIM}}$, go to Step 9.
   - 5B. If it holds $\Delta W_{\mathrm{TR}} < 0$ and the P_{TR} point is infeasible, ISA builds a linear approximation of the optimization problem about the current best record P_{OPT}. The intersection point I_{TR} between the $\mathbf{S}_{\mathrm{TR}}$ direction and the linearized constraint domain boundary is found. Non-linear constraints are evaluated at I_{TR}. Quadratic approximations of the cost function and $\mathrm{NC}_{\mathrm{act}}$ active constraints are built by interpolating the responses gathered at P_{OPT}, P_{TR} and I_{TR}. Therefore, no second order sensitivities are needed. Once the quadratic model is set, ISA solves a set of $\mathrm{NC}_{\mathrm{act}}$ quadratic equations in order to find the steps Δs_k for which the constraint violation vanishes in the approximate model. The quadratic equations derive from the equalities $G_{\mathrm{quad}}^k(\Delta s_j) = 0$ $(k = 1, \ldots, \mathrm{NC}_{\mathrm{act}})$. The smallest step amongst the Δs_j steps is taken as the trial step Δs_{TR}. The trial step Δs_{TR} serves to define the new trial design point F_{TR} on the $\mathbf{S}_{\mathrm{TR}}$ direction at which the real cost function W is evaluated.

If the cost function improves ($W_{TR} < W_{OPT}$), ISA evaluates non-linear constraints at P_{TR}. If P_{TR} is feasible, accept it as the new current best record. Hence, execute Step 4 or Step 9 if it holds $I_{GLOB} < I_{LIM}$ or $I_{GLOB} = I_{LIM}$, respectively.

Conversely, if the improvement routine fails ($W_{TR} > W_{OPT}$), execute Step 6.

5C. If it holds $\Delta W_{TR} > 0$, linearize constraints about P_{OPT} in order to search the optimum design in a zone limited by the current best record and the constrain domain boundary.

- If the P_{TR} point is feasible, define a new trial point $P_{TR,NEW}$ as the symmetric of the old P_{TR} point about the P_{OPT} point. Evaluate the cost function at the $P_{TR,NEW}$ point.
  - If the new increment $\Delta W_{TR} = W_{TR,NEW} - W_{OPT}$ turns negative, evaluate non-linear constraints at $P_{TR,NEW}$. If the new trial point $P_{TR,NEW}$ is feasible, accept it as the new current best record and execute Step 4. Conversely, if $P_{TR,NEW}$ is infeasible, go to Step 5B.
  - If the new ΔW_{TR} increment is yet greater than 0, execute Step 6.
- If the P_{TR} point is infeasible, find the intersection point I_{TR} between the $\mathbf{S}_{TR}$ direction and the linearized constraint domain boundary. Redefine the trial point P_{TR} as the symmetric of the I_{TR} point about the P_{OPT} point. Evaluate the cost function W at the P_{TR} point.
  - If the new ΔW_{TR} increment turns negative, evaluate non-linear constraints. If the new trial point P_{TR} is feasible, accept it as the new current best record and execute Step 4. If the new trial point P_{TR} is still infeasible, go to Step 5B.
  - If the new ΔW_{TR} increment is yet greater than 0, execute Step 6.

As is clear, ISA attempts in Step 5C to perturb the nominal design by moving along a descent direction which pushes the optimizer away from the constraint domain boundary where there may be local minima.

6. *Local annealing. Design variables are perturbed one at a time.*

Set $I_{SEC} = I_{SEC} + 1$. Perturb the design variables one by one in order to move away from constraint boundaries and to escape from local minima. Use linearized constraints in order to reduce the number of exact analyses and to save computational time. No re-linearization is required here since Step 6 is executed only if global annealing (i.e., all the optimization variables are perturbed simultaneously) failed in Steps 5A–C where ISA builds the linearized sub-problem about the current best record P_{OPT}.

For each design variable, generate another random number ρ_j in the interval $(0, 1)$,

$$\begin{aligned} \rho_j > 0.5 \;\Rightarrow\; x_{TR,j} &= x_{OPT,j} + (x_j^u - x_j^l)\cdot\rho_j\cdot\gamma_K, \\ \rho_j < 0.5 \;\Rightarrow\; x_{TR,j} &= x_{OPT,j} - (x_j^u - x_j^l)\cdot\rho_j\cdot\gamma_K, \end{aligned} \tag{5}$$

where γ_K is again computed with expression (4). As it is done in Step 4, design variables are changed following their order sequentially.

As one can see, cost function sensitivities are not used in expression (5) for defining the new trial values $x_{TR,j}$. The reason for this strategy is that computing the gradient of cost function is useless. In fact, the ISA algorithm executes a local annealing search each time the movements taken along the descent directions $\mathbf{S}_{TR}$ in a global annealing cycle lead to violate constraints or push ISA towards non-convex regions of the design space. In such cases, ISA will recover the constraint violation eventually at the cost of some penalty in weight or it will get back to a convex region of the design space. The former fact actually implies $\mathbf{S}_{TR}^T \overline{\nabla} W(\mathbf{X}_{OPT}) > 0$. In the latter case, linear approximation may be not effective in non-convex regions.

Now, let us denote as $P_{TR,j}$ each new trial design generated. Let $\mathbf{X}_j(x_{OPT,1}, x_{OPT,2}, \ldots, x_{TR,j}, \ldots, x_{OPT,N})$ be the corresponding design vector. Compute cost function $W(\mathbf{X}_j)$ at each $P_{TR,j}$. Compute the change in cost function $\Delta W_j = W(\mathbf{X}_j) - W_{OPT}$.

An $\mathbf{X}_j$ design is immediately rejected if violates linearized constraints. This strategy is justified by the informal argument that if linearized constraints are violated, non-linear constraints also will be very likely to be violated.

If it occurs $\Delta W_j < 0$ and the linearized constraints are satisfied, store the $\mathbf{X}_j$ vector and the W_j cost in the Π_1 database of "temporary" best records. The notation "temporary" indicates that each $\mathbf{X}_j$ might violate the non-linear constraints. In fact, as the W_j cost decreases, the corresponding $\mathbf{X}_j$ design may get too close to the boundary of linearized constraint domain and finally end up infeasible. For this reason, ISA chooses from the Π_1 database the two designs $\mathbf{X}_{j,\mathrm{SMALL}}$ and $\mathbf{X}_{j,\mathrm{LARGE}}$ that correspond to the smallest and the largest $|\Delta W_j|$ changes, respectively. Then, ISA evaluates the non-linear constraints at those two points. If the $\mathbf{X}_{j,\mathrm{LARGE}}$ design is feasible, ISA takes it as the new best record. Conversely, the $\mathbf{X}_{j,\mathrm{SMALL}}$ design is taken as the new best record.

If it occurs $\Delta W_j > 0$ and the linearized constraints are satisfied, define Metropolis' probability function in the following way:

$$P(\Delta W_j) = \exp\left\{\frac{-\Delta W_j}{\left(\sum_{r=1}^{\mathrm{NDW}} \Delta W_r/\mathrm{NDW}\right)\cdot T_K}\right\}. \tag{6}$$

The NDW parameter in expression (6) represents the number of trial points at which the cost function resulted larger than the current best records throughout the optimization process. The increases in cost are expressed by the ΔW_r terms. The $\sum_{r=1,\mathrm{NDW}} \Delta W_r/\mathrm{NDW}$ ratio accounts for the general formation of all the previous candidate designs and serves to normalize the probability function with respect to the change in cost function.

Each design $\mathbf{X}_j$ is provisionally accepted or certainly rejected based on the Metropolis' criterion:

$$\begin{aligned} &P(\Delta W_j) > \min(\mathrm{NRD}_j, \rho_j) \quad \Rightarrow \quad \text{provisionally accepted,} \\ &P(\Delta W_j) < \min(\mathrm{NRD}_j, \rho_j) \quad \Rightarrow \quad \text{certainly rejected.} \end{aligned} \tag{7}$$

The $\mathbf{X}_j$ designs provisionally accepted are included in the Π_2 database. If there are no trial designs that yield reductions in cost, ISA extracts from the Π_2 database the design $\mathbf{X}_j^{\mathrm{BEST}}$ for which the cost function value is the least and sets this as the current best record. In simple words, ISA minimizes the increase in cost if the local annealing could not improve design.

It should be noted that ISA compares design acceptance probability to two random numbers and not just to one random number as, instead, is done in classical SA. The present strategy is certainly conservative since ISA shifts from global annealing to local annealing when the former search mechanism fails.

Once Step 6 is completed, go to Step 8.

7. *Search for a feasible design or for a design which violates constraints the least.*

Set $I_{\mathrm{GLOB}} = I_{\mathrm{GLOB}} + 1$. This step is executed if the optimization process or a cooling cycle begins from an infeasible design point. The infeasible starting design is yet denoted as $\mathbf{X}_{\mathrm{OPT}}$ following the notation used in Step 3. The rationale behind this step is that ISA tries to change the current design by moving along directions where the cost function improves and the constraints get less violated than at the cooling cycle starting point.

- Define the gradient vector of each of the NCV violated constraints. That is, define NCV directions in fashion of $\boldsymbol{\delta}_{G,V} = -\nabla G_V(\mathbf{X}_{\mathrm{OPT}})$.
- If it holds $\boldsymbol{\delta}_{G,V}^{\mathrm{T}} \nabla W(\mathbf{X}_{\mathrm{OPT}}) < 0$, solve the linear system formed by the $\boldsymbol{\delta}_{G,V}$ directions and the linearized constraints $G_{\mathrm{LIN},V}$ that are violated. Find hence the steps $\Delta G_{\mathrm{LIN},V}$.
- Solve the linear system formed by the $\boldsymbol{\delta}_{G,V}$ directions and the linearized constraints that are satisfied. Find hence the steps $\Delta G_{\mathrm{INIT},V}$.
- Include the $\Delta G_{\mathrm{LIN},V}$ and $G_{\mathrm{INIT},V}$ steps as components of the $\mathbf{S}_{\mathrm{VIOL}}$ direction. Combine the $\Delta G_{\mathrm{LIN},V}$ and $\Delta G_{\mathrm{INIT},V}$ steps in order to define the ΔG_{VIOL} step along the $\mathbf{S}_{\mathrm{VIOL}}$ direction. As is clear, the $\mathbf{S}_{\mathrm{VIOL}}$ vector is a direction is a vector where each component reduces the constraint violation for one or more constraints. In addition some of the $\mathbf{S}_{\mathrm{VIOL}}$ components are such that the double goal of reducing

constraint violation and reducing cost function is achieved. If it holds $\mathbf{S}_{\mathrm{VIOL}}^{\mathrm{T}}\nabla W(\mathbf{X}_{\mathrm{OPT}}) < 0$, the effect of cost function reduction is larger than the effect of constraint violation. Conversely, the optimizer has to pay some weight penalty in order to get in a feasible zone of the design space.

- Reduce the ΔG_{VIOL} step to $G_{\mathrm{FIN,VIOL}}$ by means of a trust region model along the $\mathbf{S}_{\mathrm{VIOL}}$ direction. Let $G_{\mathrm{FIN,VIOL}}$ be the corresponding point of the design space. That is,

$$\min_{k=1,\ldots,\mathrm{NCV}}\left[\frac{G_k(\mathbf{X}_{\mathrm{OPT}}) - G_k(\mathbf{X}_{\mathrm{OPT}} + \boldsymbol{\delta}\mathbf{X}_{\mathrm{VIOL}})}{\boldsymbol{\delta}\mathbf{X}_{\mathrm{VIOL}}^{\mathrm{T}}\overline{\overline{\nabla}}G_k(\mathbf{X}_{\mathrm{OPT}})}\right] \geqslant 0.75, \tag{8}$$

where $\|\boldsymbol{\delta}\mathbf{X}_{\mathrm{VIOL}}\|$is the unknown step size along $\mathbf{S}_{\mathrm{VIOL}}$ determined by solving the non-linear equation (8). The 0.75 value in the RHS of expression (8) is derived from classical implementation of trust region models (see, for instance, Wujek and Renaud, 1998).

- If $G_{\mathrm{FIN,VIOL}}$ satisfies the linearized constraints set it as the new trial point P_{TR}. If non-linear constraints are satisfied at P_{TR}, choose P_{TR} as the current best record and execute from Step 3 onward.
- If $G_{\mathrm{FIN,VIOL}}$ violates the linearized constraints, re-linearize constraints about $G_{\mathrm{FIN,VIOL}}$. Rename the $G_{\mathrm{FIN,VIOL}}$ point as the new trial point P_{TR}. Repeat this step until some new $G_{\mathrm{FIN,VIOL}}$ point satisfying linearized constraints is found.

The rationale behind the strategy described above is the following. Classical SA deals with violated constraints by generating new random designs until constraints are satisfied. However, such a strategy may result in increasing the cost function too much. From the stand point of the number of required exact analyses, the strategy implemented in ISA performs N function evaluations in order to find the gradients of non-linear functions. Although classical simulated annealing also requires up to N exact analyses in each inner loop (one local annealing cycle in ISA or the same as a "secondary cooling cycle" in classical SA), it does not provide the designer with any information on the direction along which the constraint violation reduces the most, on which variables govern the constraint violation reduction rate, and on the relationship between the amount of constraint violation and reduction in cost function.

8. If it occurs $I_{\mathrm{SEC}} = I_{\mathrm{LIM}}$, go to Step 9. Conversely, set $I_{\mathrm{SEC}} = I_{\mathrm{SEC}} + 1$ and repeat from Step 3 onward.

It is to be noticed that ISA may execute less than I_{LIM} local annealing cycles in a cooling cycle (in fact, ISA tries to perform global annealing each time a local annealing cycle improves the previous best record) while in classical SA the optimizer is forced to perform all the I_{LIM} secondary cooling cycles included in a primary cooling cycle. This fact contributes to reducing significantly the CPU time required in the optimization process.

9. If $K > 3$ check for convergence according to the following criterion (9):

$$\begin{aligned}&\max\left\{\max\left[\frac{|W_{\mathrm{OPT},K} - W_{\mathrm{OPT},K-1}|}{W_{\mathrm{OPT},K}};\frac{|W_{\mathrm{OPT},K-1} - W_{\mathrm{OPT},K-2}|}{W_{\mathrm{OPT},K-1}};\frac{|W_{\mathrm{OPT},K-2} - W_{\mathrm{OPT},K-3}|}{W_{\mathrm{OPT},K-2}}\right];\right.\\&\left.\mathrm{Max}\left[\frac{\|\mathbf{X}_{\mathrm{OPT},K} - \mathbf{X}_{\mathrm{OPT},K-1}\|}{\|\mathbf{X}_{\mathrm{OPT},K}\|};\frac{\|\mathbf{X}_{\mathrm{OPT},K-1} - \mathbf{X}_{\mathrm{OPT},K-2}\|}{\|\mathbf{X}_{\mathrm{OPT},K-1}\|};\frac{\|\mathbf{X}_{\mathrm{OPT},K-2} - \mathbf{X}_{\mathrm{OPT},K-3}\|}{\|\mathbf{X}_{\mathrm{OPT},K-2}\|}\right]\right\} \leqslant \varepsilon_{\mathrm{CONV}},\end{aligned} \tag{9}$$

where $W_{\mathrm{OPT},K}$ and $\mathbf{X}_{\mathrm{OPT},K}$, respectively, denote the best record and the corresponding design vector found in the Kth cooling cycle.

The $\varepsilon_{\mathrm{CONV}}$ parameter is set to 10^{-5} in order to avoid premature convergence if the last four cooling cycles resulted in marginal improvements in design.

If the convergence criterion (9) is satisfied go to Step 11.

10. If $K < 3$ or the convergence criterion (9) is not satisfied:

- Reset K as $K = K + 1$;
- Reset I_{SEC} as $I_{\mathrm{SEC}} = 0$ and I_{GLOB} as $I_{\mathrm{GLOB}} = 0$;

- Reduce temperature T_K in fashion of $T_K = \beta_K T_{K-1}$ where the parameter β_K is chosen as:

$$\beta_K = \max\left[\frac{0.95}{\left(1+\frac{N_{\mathrm{REJE}}}{N_{\mathrm{TRIA}}}\right)}; \left(1-\frac{W_{\mathrm{FIN},K-1}}{W_{\mathrm{INIT},K-1}}\right)\right]. \tag{10}$$

The $W_{\mathrm{FIN},K-1}$ and $W_{\mathrm{INIT},K-1}$ terms respectively indicate the cost function values at the beginning and at the end of the current annealing cycle. N_{REJE} is the number of trial designs rejected by ISA out of the total number of trial designs N_{TRIA} generated in the current cooling cycle. The N_{REJE} number includes each trial point which does not yield immediate improvement in design. For instance, if $I_{\mathrm{LIM}} = 5$ global annealing cycles are performed within the current cooling cycle and ISA executes Step 5B in each global annealing cycle (that is, there are five infeasible trial points at which the cost function decreases), it yields $N_{\mathrm{REJE}} = 5$ and $N_{\mathrm{TRIA}} = I_{\mathrm{LIM}} + N_{\mathrm{REJE}} = 10$.

In general, it is suggested in literature to use a constant temperature reduction factor chosen ranging between 0.9 and 0.99. However, unlike many referenced SA algorithms (see, for instance, Rao, 1996; Huang and Arora, 1997; Yu Chen and Su, 2002), ISA utilizes a variable temperature reduction factor that accounts for two effects: the percentage of trial points that immediately yield improvements in design and the trend of the cost function. The former effect is captured by a term whose upper limit is the 0.95 value that is about the average between 0.9 and 0.99; this upper limit is eventually reduced by means of a knockdown factor λ_K defined as the inverse of $(1 + N_{\mathrm{REJE}}/N_{\mathrm{TRIA}})$: the λ_K factor is obviously equal to 1 if $N_{\mathrm{REJE}} = 0$. The effect of cost function is captured in (10) by computing the relative change in cost attained in the current cooling cycle.

The rationale behind expression (10) is the following. If the cost function decreased much in the current cooling cycle, the temperature can be kept high since the optimizer is still exploring a zone where the cost function gradient is negative. Hence, new descent directions can be found easily and the risk to reject candidate designs is very low. In simple words, the criterion (10) is driven by the $(1 - W_{\mathrm{FIN},K-1}/W_{\mathrm{INIT},K-1})$ term. For instance, if Step 5B is executed one time within each global annealing cycle ($N_{\mathrm{REJE}} = 5$; $N_{\mathrm{TRIA}} = 10$), one gets 0.633 for the $0.95\lambda_K$ product. This implies that a decrease in cost function larger than 63.4% (very likely to occur in the early cooling cycles if the initial design is far enough from the constraint domain boundaries) makes the $(1 - W_{\mathrm{FIN},K-1}/W_{\mathrm{INIT},K-1})$ term predominant over the $0.95\lambda_K$ product.

Conversely, if the cost function improved marginally or even increased, the optimizer entered in a zone where there are few descent directions or a descent direction cannot even be defined. Hence, the temperature should be reduced in order to reject many trial points because they certainly will not yield significant improvements in design. In simple words, since the $(1 - W_{\mathrm{FIN},K-1}/W_{\mathrm{INIT},K-1})$ term may get close to zero, the criterion (10) is driven by the $0.95\lambda_K$ product: the larger the number of trial designs rejected in the current design cycle, the higher will be the reduction in temperature. This strategy is not too conservative. In fact, assuming that the 30% of total trials were rejected (that is, a very high rejection rate), one gets a value of about 0.75 for the $0.95\lambda_K$ product.

- Repeat from Step 3 onward.

11. End the optimization process.

3. Test cases

From the description of ISA contained in the previous section, it could seem that the present algorithm will be able to solve only optimization problems where explicit gradient information are available. This may appear as a limitation with respect to classical simulated annealing which is generally considered a versatile optimization tool since it does not involve gradient computation and hence does not include the explicit formulation of the optimization problem as a pre-requisite. However, the present authors point out that

ISA performance is either insensitive to how the optimization problem is formulated and how the gradients are determined. In order to prove this statement, we chose a set of four optimization problems where not all of the gradients can be determined explicitly since their computation involves finite element analysis.

The test cases used for testing the computational efficiency of the ISA algorithm are now described.

3.1. Weight minimization of bar truss structures

The weight minimization problem for a bar truss structure comprised of NOD nodes and NEL elements may be formulated as follows:

$$\begin{cases} \min W = \rho g \sum_{j=1}^{\mathrm{NEL}} l_j x_j, & \\ u^{\mathrm{l}}_{(x,y,z),k} \leqslant u_{(x,y,z),k,\mathrm{ilc}} \leqslant u^{\mathrm{u}}_{(x,y,z),k}, & \begin{cases} i = 1,\ldots,N, \\ j = 1,\mathrm{NEL}, \\ k = 1,\mathrm{NOD}, \\ \mathrm{ilc} = 1,\mathrm{NLC}, \end{cases} \\ \sigma^{\mathrm{l}}_j \leqslant \sigma_{j,\mathrm{ilc}} \leqslant \sigma^{\mathrm{u}}_j, & \\ x^{\mathrm{l}}_i \leqslant x_i \leqslant x^{\mathrm{u}}_i, & \end{cases} \tag{11}$$

where

- x_j is the cross sectional area of the jth element of the structure included as sizing variable in the optimization process.
- x_j is the length of the jth element of the structure.
- g is the gravity acceleration value (9.81 m/s^2); ρ is the material density.
- NLC the number of independent loading conditions acting on the structure.
- $u_{(x,y,z),k,\mathrm{ilc}}$ are the displacements of kth node in the directions x, y, z, with the lower and upper bounds $u^{\mathrm{l}}_{(x,y,z),k}$ and $u^{\mathrm{u}}_{(x,y,z),k}$.
- $\sigma_{j,\mathrm{ilc}}$ is the stress on the jth element, with the lower and upper bounds σ^{l}_j (compressive) and σ^{u}_j (tensile).
- The ilc subscript indicates that displacement and stress constraints are relative to the ilc-th loading condition. The constraints on stresses and displacements are put in a dimensionless form.

The planar 200 bar truss structure shown in Fig. 2 is to be designed under five independent loading conditions and with constraints on nodal displacements and member stresses. The structure has 77 nodes. The optimization includes 200 design variables (cross sectional area of truss members) and 3500 non-linear constraints. The Young modulus of the material is 2.069×10^{11} N/m^2 while the density is 7833.413 kg/m^3. The lower bound of the cross sectional area is set to 0.1 in.^{2} (0.00064516 m^2). The displacements of the free nodes must be less than 0.5 in. (0.0127 m). The allowable stress (the same in tension and compression) is set to 30,000 psi (21.092 kgf/mm^2). The optimization was carried out also with the more severe stress limit of 10,000 psi (7.031 kgf/mm^2). More details on geometry and loading conditions are given in the work of Venkayya (1978).

Lamberti and Pappalettere (2003a) considered different variable linkages for the structure in Fig. 2 and carried out optimization runs including 96, 105 and 200 design variables. Their optimization algorithm—LESLP, where the acronym stands for Linearization Error Sequential Linear Programming—implemented a sophisticated formulation of Sequential Linear Programming where the move limit domain included as many candidate designs as possible. Results indicated that LESLP was superior over other referenced optimization algorithms and commercial software. However, the designs optimized with LESLP and the other codes were not consistent with the amount of design freedom included in the optimization. In particular, the structure optimized with 200 design variables was about 230 kg heavier than its counterpart optimized with only 96 design variables. It was seen that the stress constraints were not active in the 200 design variable case (about 30% of the 30,000 psi allowable limit) while they were active in the 96 design variable case. This occurred probably because the design space includes a zone where there is a sharp change in

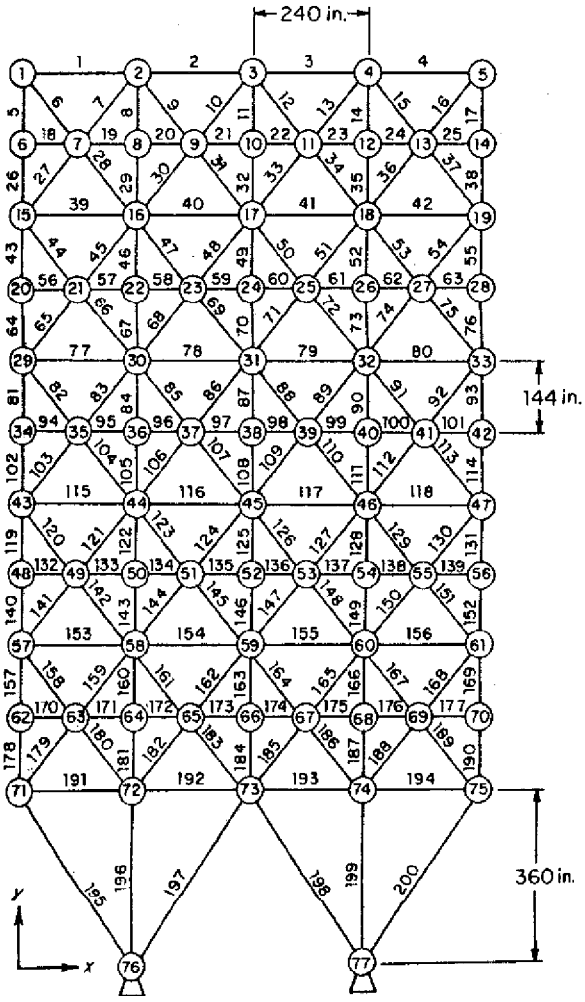

Fig. 2. Schematic of the planar 200 bar truss structure.

stress value if the design is perturbed so to keep displacement constraints active. The very large stress gradients resulted in very tight move limits and hence limited the zone where the optimum is searched. In order to overcome this problem, the planar structure was re-optimized here with the ISA algorithm including 200 design variables. Interestingly, since the cost function of the truss problem is linear, the $\overline{\nabla} W$ gradient vector required in the ISA global annealing cycles can be immediately determined and does not change through the entire optimization process.

The sizing optimization problem (11) may become a configuration optimization problem if the co-ordinates of structure nodes are also included as optimization variables. Let $x_{j1,2}$, $y_{j1,2}$, $z_{j1,2}$ be the co-ordinates of the nodes limiting the generic ith element of the structure. It follows:

$$\min W = \rho g \sum_{j=1}^{\mathrm{NEL}} x_j \sqrt{(x_{j1} - x_{j2})^2 + (y_{j1} - y_{j2})^2 + (z_{j1} - z_{j2})^2}. \tag{12}$$

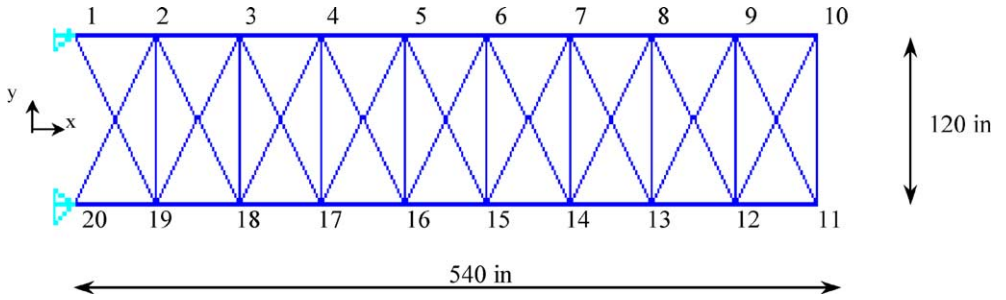

Fig. 3. Schematic of the cantilevered bar truss structure.

In the third test case of this study, the cantilevered truss structure shown in Fig. 3 is to be designed with constraints on nodal displacements, member stresses and critical buckling loads. The structure has 45 elements and 20 nodes. It is built by nine repeating modules each of which contains five elements. In addition, the nodal co-ordinates of the eighteen free nodes are taken as configuration variables. Therefore, 81 design variables are considered. Vertical forces are applied to the structure: respectively, 150,000 lbf (68,038.856 kgf) at nodes 9 and 10 acting downward, and 50,000 lbf (22,679.619 kgf) at nodes 11 and 12 acting upward. The lower bound of the cross sectional areas is 0.1 in.^{2}. The allowable tensile stress is 25,000 psi (17.577 kgf/mm^2), the stress limit in compression accounts also for buckling loads (see Dhingra and Lee, 1994). The displacements of the free nodes must be less than 2 in. (0.0508 m). As far as it concerns the computation of the $\overline{\nabla} W$ gradient vector, it should be noticed that cost function sensitivities with respect to configuration variables can be easily obtained in fashion of closed form expressions while sizing variable sensitivities are obtained in the same way as in the 200 bar truss case.

Since the truss structures shown in Figs. 2 and 3 are statically undetermined, structural response sensitivities (i.e., constraint gradients) are obviously not available in explicit form.

All optimizations were started from both the lower bound and upper bound of cross sectional areas (the latter was set to 100 in.^{2}: i.e., 0.64516 m^2). ISA was compared to a classical SA algorithm where the optimization variables are perturbed one by one and global annealing is not performed (see Appendix A). Finally, the structures were optimized also with the TRLP-DOT algorithm where an enhanced version of LESLP (indicated as LSTRLP in Lamberti and Pappalettere, 2004) is combined with the commercial optimizer DOT®, 1995. In particular, DOT solved the approximate sub-problems with a feasible direction based routine thus saving considerable amounts of CPU time with respect to the Simplex solver used by LESLP and LSTRLP.

3.2. Composite material characterization

An eight-ply woven reinforced fiberglass-epoxy composite laminate utilized as substrate for printed circuit boards was characterized in terms of in-plane mechanical behavior. The laminate is shown in Fig. 4. A 46 mm long, 13 mm tall and 1.2 mm thick specimen was cut from a slice of material. The in-plane properties assumed as target values in the optimization process were known from standard mechanical tests: $E_x = 25$ GPa; $E_y = 22$ GPa; $G_{xy} = 5$ GPa; $\nu_{xy} = 0.28$. In order to find the values of in-plane elastic constants of the laminate, we minimized the difference between numerical data obtained by means of FEM analysis and experimental data measured by means of phase-shifting electronic speckle pattern interferometry (PS-ESPI).

PS-ESPI is based on the fact that two beams originated from a coherent light source (laser) produce an interferometric pattern when they hit on a surface. The light intensity in each point of the specimen surface is

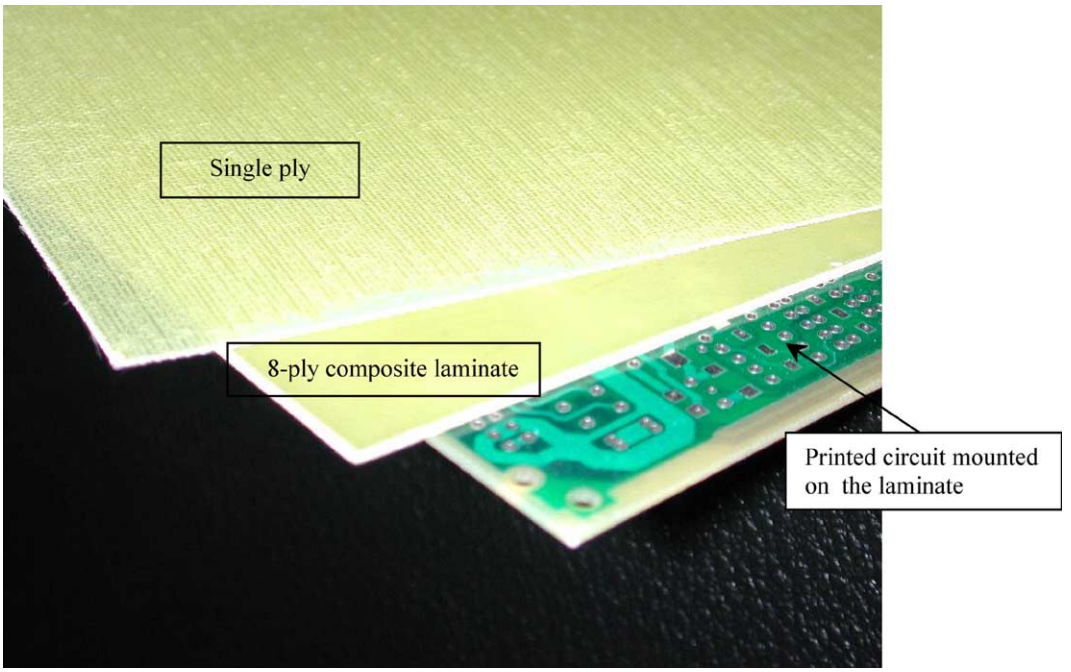

Fig. 4. Reinforced fiberglass-epoxy laminate characterized by means of ISA and PS-ESPI.

an harmonic function $I(x,y,\Phi(x,y))$ of the pixel coordinate (x,y). In order to recover the phase $\Phi(x,y)$ of the $I(x,y,\Phi(x,y))$ light intensity function, N different interferometric patterns are taken at two different exposures: the reference configuration (i.e., no load is applied) and the loaded one. This is equivalent to introduce N shifts in phase in order to span at least the 2π angle for which the harmonic function $I(x,y)$ stays the same (i.e., for an harmonic function it obviously holds: $I(x,y,\Phi(x,y)+2\pi)=I(x,y,\Phi(x,y))$). From each set of N acquisitions done in the reference or in the loaded configuration, it is possible to get the distribution of phase $\Phi(x,y)$ for each point of the specimen surface in fashion of a fringe pattern. Finally, the unknown displacements are easily computed as they are proportional to the $\Delta\Phi(x,y)$ phase difference between the two different exposures. For each pixel (x,y), the corresponding $u(x,y)$ displacement is then:

$$u(x,y)=\frac{\Delta\Phi(x,y)}{4\pi}\cdot\frac{\lambda}{\sin\theta}, \tag{13}$$

where λ and θ, respectively, are the laser light wave length and the illumination angle (see the schematic of the experimental set-up shown in Fig. 5).

A very common strategy in PS-ESPI is to use the *four-phase* technique where $N=4$ and phase shifts are hence chosen spaced by 90° (i.e., $2\pi/N$). In such a case, the phase difference $\Delta\Phi(x,y)$ is given by the expression:

$$\Delta\Phi(x,y)=\arctan\left(\frac{I_{4,\text{loaded}}(x,y)-I_{2,\text{loaded}}(x,y)}{I_{1,\text{loaded}}(x,y)-I_{3,\text{loaded}}(x,y)}\right)-\arctan\left(\frac{I_{4,\text{reference}}(x,y)-I_{2,\text{reference}}(x,y)}{I_{1,\text{reference}}(x,y)-I_{3,\text{reference}}(x,y)}\right), \tag{14}$$

where the $I_{1,\ldots}(x,y)$, $I_{2,\ldots}(x,y)$, $I_{3,\ldots}(x,y)$ and $I_{4,\ldots}(x,y)$ are the light intensity values recorded in the four different acquisitions done for the "loaded" or "reference" configurations.

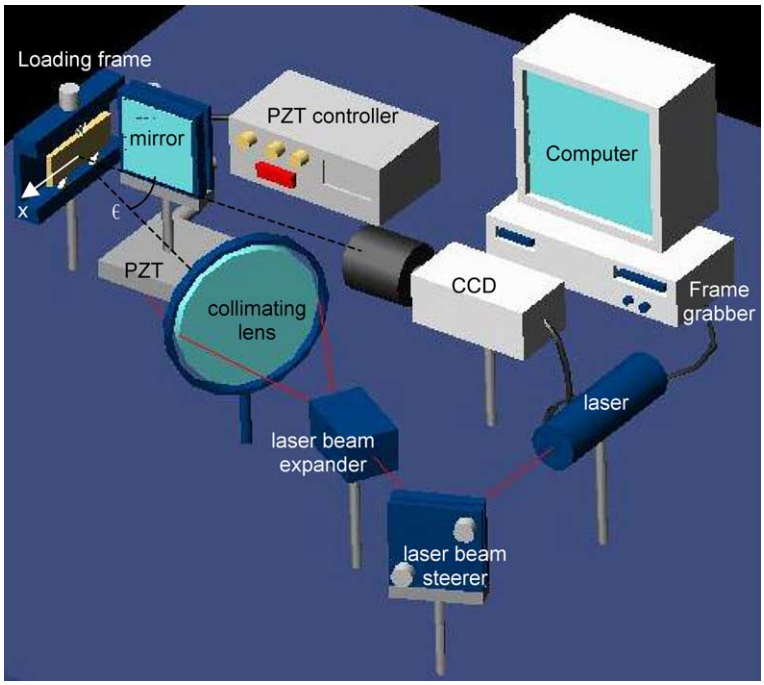

Fig. 5. Schematic of the PS-ESPI set-up utilized in the material characterization (test case 4).

Fig. 5 shows the typical PS-ESPI experimental set-up based on the Leendertz's double illumination interferometer (see Leendertz, 1970). A 35 mW He–Ne laser ($\lambda = 632.8$ nm) provides the coherent light source. A closed loop controlled piezoelectric transducer (PZT) is used as phase shifter. The intensity distributions of the combined light beams are recorded by a B/W CCD camera (795×596 pixel sensors). The images are then digitized by means of an 8-bit frame grabber. The laser beam is expanded first, filtered then and collimated finally. In order to preserve coherence, the double illumination is obtained by reflecting a certain fraction of the laser beam onto a mirror (mounted on the PZT device) which is orthogonal to the surface of the specimen. The illumination angle θ made by the laser beams with the direction of observation is 20°. More details on the optical set-up used in this research can be found in Genovese et al. (2004).

In order to determine the elastic constants of the laminate shown in Fig. 4, the following optimization problem can be formulated:

$$\begin{cases} \min \left[\Psi(E_x, E_y, G_{xy}, v_{xy}) = \sqrt{\frac{1}{N_{\text{pix}}} \sum_{j=1}^{N_{\text{pix}}} \left(\frac{u^j_{x,y,\text{TOT}} - \overline{u^j_{x,y,\text{TOT}}}}{\overline{u^j_{x,y,\text{TOT}}}} \right)^2} \right], \\ E_x > E_y, \\ 1 - v_{xy}(E_y/E_x) > 0, \\ E^{\text{l}}_x \leqslant E_x \leqslant E^{\text{u}}_x, \\ E^{\text{l}}_y \leqslant E_y \leqslant E^{\text{u}}_y, \\ G^{\text{l}}_{xy} \leqslant G_{xy} \leqslant G^{\text{u}}_{xy}, \\ v^{\text{l}}_{xy} \leqslant v_{xy} \leqslant v^{\text{u}}_{xy}, \end{cases} \tag{15}$$

where the error function Ψ is to be minimized and the unknown elastic constants E_x, E_y, G_{xy} and v_{xy} of the laminate are included as optimization variables. The two additional constrains in expression (15) ensure positive definiteness of the $[Q]$ stiffness matrix.

Let $u^j_{x,y,\mathrm{TOT}}$ and $\overline{u^j_{x,y,\mathrm{TOT}}}$ denote the values of in-plane displacement in correspondence of the jth model node and image pixel determined, respectively, by means of finite elements and by means of PS-ESPI. The "x", "y" and "TOT" subscripts respectively indicate that horizontal, vertical or total $\left(\text{i.e., } \sqrt{u_x^2 + u_y^2}\right)$ displacements are considered. It is intended that the $\overline{u^j_{x,y,\mathrm{TOT}}}$ displacements are target values since experimental measurements do not require to know the values of elastic constants a priori. Conversely, the values of elastic constants must be specified as input in the numerical analysis carried out with finite elements in order to calculate the $u^j_{x,y,\mathrm{TOT}}$ displacements. In addition, let N_{pix} denote the total number of points (image pixels/FEM nodes) at which the experimental and numerical results are compared. In order to preserve the correspondence between the pixels of the recorded images and the nodes of the FEM model, the FEM model of the specimen is built by setting the element size equal to the pixel size.

Although the error function Ψ depends on only four variables, it is certainly a highly non-smooth function. In order to support this statement, a sensitivity analysis was carried out for two different load cases: three-point-bending and shear load cases (FEM models are shown in Fig. 6). The expected "optimum design" (i.e., the laminate elastic constants to be determined) was perturbed by changing one elastic constant at a time. It was found that the "optimum" is clearly and univocally located in correspondence of the material property target values. However, the shape and the slope of the Ψ function changed significantly depending on the fact that each elastic constant was increased or reduced. This fact introduced significant non-smoothness in the optimization problem (15) and thus justified the use of simulated annealing.

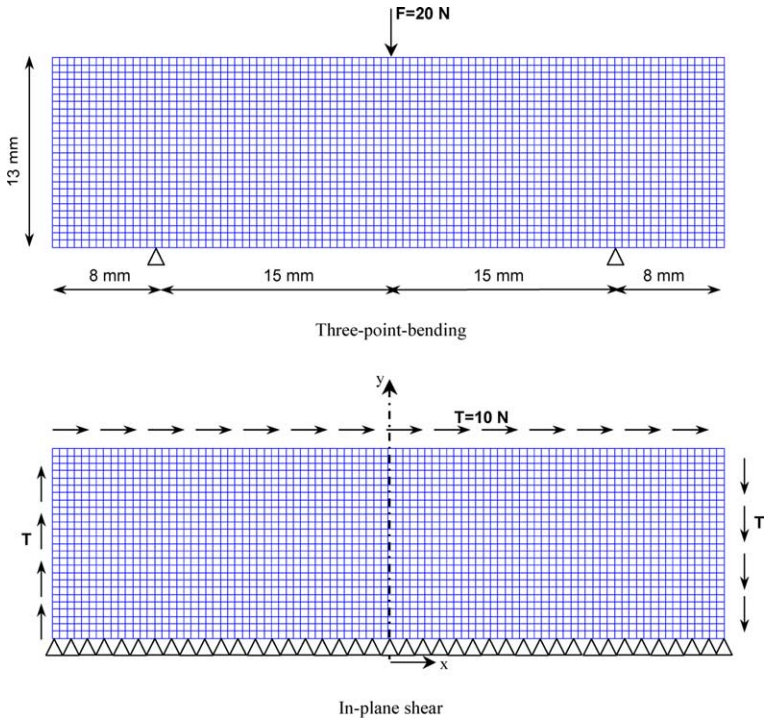

Fig. 6. Finite element models used in the numerical trials of material property identification.

As far as it concerns determination of gradients, it should be noticed that the scenario in problem (15) is the exact reversal of truss design problems. In fact, constraint gradients can be now determined explicitly while computation of cost function sensitivities involves finite element analysis.

The suitability of the ISA optimizer for material characterization problems was evaluated through a set of eight cases where the target displacements are generated numerically by FEM analyses. The commercial finite element code ANSYS® (2003) was utilized to perform structural analyses. The minimization problem (15) was also solved by means of the optimization routine available in ANSYS. The specimen was modeled with PLANE42 elements each of which includes four nodes and two degrees of freedom per node. Two different loading conditions were considered: three-point-bending and in-plane shear (see Fig. 6).

The eight numerical trials are grouped as follows: Runs 1–3 and 7 where the composite specimen experienced three-point-bending; Runs 4–6 and 8 where the composite specimen experienced in-plane shear. The definition strategy of the Ψ error function changed in the different optimization runs: the errors on the u_x displacements were utilized in Runs 1–2 and 4–5 while the errors on the u_y displacements were utilized in Runs 3 and 6. Finally, the total displacements $u_{\mathrm{TOT}} = \sqrt{u_x^2 + u_y^2}$ were utilized in Runs 7–8.

In the experimental case, the Ψ error function included the u_x displacements because the Lendeertz's interferometer used in this study measures horizontal in-plane displacements. The composite laminate specimen was mounted on two supports spaced by 30 mm (see Fig. 5). The vertical load which generated the three-point-bending state was provided to the specimen by a rounded tip mounted on a sledge moved by a micrometric screw. The three-point-bending test was chosen in the experimental case of material identification because it minimizes the amount of rigid body motions (RBMs) of the specimen during the loading phase. This is very useful because RBMs may invalidate PS-ESPI measurements as they cause speckle pattern de-correlation. The vertical displacement of the rounded tip generated by an $F = 140$ N vertical load applied to the specimen was such that to preserve correlation of speckle patterns thus ensuring enough density of phase fringes. Since the load applied in the experimental case was seven times as large as the load applied in the eight numerical trial cases, four additional optimizations where the vertical load ranged from 20 to 250 N (the last value is the load at which the fringe contrast was almost lost) were run. As expected, the ISA algorithm resulted insensitive to the load amplitude. In fact, the convergence curves practically coincided after only 10 cooling cycles.

4. Results and discussion

The ISA algorithm was coded into a Fortran 90 program. The same was done for the optimizer implementing classical SA and the Sequential Linear Programming (SLP) optimizer used as basis for comparison with ISA. The optimizer based on classical SA included an 1-directional annealing search strategy (see Appendix A) and a constant rate cooling schedule. The SLP code integrated the LSTRLP algorithm developed by two of the present authors (see Lamberti and Pappalettere, 2004) and the well known commercial optimizer DOT® developed by the Vanderplaats R&D. The optimization of the truss structure was carried out on a DEC-Alpha 500 MHz Unix workstation. The material identification problem was run on a 700 MHz Pentium III personal computer in order to include the FEM solver of the ANSYS® software as a subroutine of the ISA code.

Table 1 shows the results obtained in the truss optimization problems. The table reports the number of cooling cycles, the number of global annealing cycles (within round brackets) and the number of constraint linearizations. The number of linearizations obviously coincides with the number of design cycles performed when the optimization was carried out with the Sequential Linear Programming method (this run is denoted as "TRLP-DOT" in the rest of the paper). The number of structural analyses and the CPU time required in the optimizations are also quoted in the table.

Table 1
Comparison of the numerical efficiency of different optimizers in truss design problems

Problem	Initial design	Optimization algorithm	Structural weight (kg)	Number K of cooling (global) cycles	Constraint linearizations or sub-problems	Structural analyses	CPU time (s)
Two-hundred bar truss $\sigma_{LIM} = 30{,}000$ psi	Variables at their upper bound	ISA	12,767.294	43 (21)	22	15,411	3133
		Classical SA	12,779.287	54	N/A	30,624	6107
		TRLP-DOT	12,794.060	N/A	44	17,744	3082
		LESLP[a]	13,054.851	N/A	28	11,273	4148[b]
		LESLP[a,c]	12,823.808		N/A		
	Variables at their lower bound	ISA	12,764.179	40 (15)	30	18,508	3613
		Classical SA	12,775.091	49	N/A	27,810	5542
		TRLP-DOT	12,800.971	N/A	44	17,657	3444
Two-hundred bar truss $\sigma_{LIM} = 10{,}000$ psi	Variables at their upper bound	ISA	13,102.237	37 (20)	19	12,310	3120
		Classical SA	13,105.972	49	N/A	36,421	8299
		TRLP-DOT	13,175.328	N/A	41	16,499	4874
		LESLP[a,b,d]	13,129.595	N/A	24	9697	9968[b]
	Variables at their lower bound	ISA	13,103.217	43 (20)	23	16,110	4554
		Classical SA	13,107.747	45	N/A	39,010	9050
		TRLP-DOT	13,124.309	N/A	44	17,678	5492
		LESLP[a,b]	13,120.480	N/A	20	8078	8140[b]
Cantilevered bar truss	Variables at their upper bound	ISA	3551.007	33 (17)	16	7785	60.6
		Classical SA	3610.443	52	N/A	13,238	107
		TRLP-DOT	3630.640	N/A	54	8791	198
		LSRTLP[e]	3633.748	N/A	29	4788	170[b]
	Variables at their lower bound	ISA	3581.897	42 (19)	25	11,350	83.7
		Classical SA	3663.281	53	N/A	23,840	186
		TRLP-DOT	3690.072	N/A	78	12,684	218

[a] See Lamberti and Pappalettere (2003a).
[b] This optimization was run on a computer about 30% slower than the 500 MHz DEC-Alpha used in this study.
[c] This optimization run included 96 design variables.
[d] This optimization was started from an uniform design of 98.5135 in.^{2}.
[e] See Lamberti and Pappalettere (2004).

It can be seen that ISA found the best design either when the optimization started from a feasible point (for design variables at the upper bound, the most active constraint was about 14.5% and 74% of allowable limits, respectively, for the 200 bar truss and the cantilevered bar truss) and when the starting point violated constraints very much (for design variables at the lower bound, constraint violation ranged between 7600% (nodal displacements) and 14,530% (element stresses) in the 200 bar truss case while raised to 32,000% (element stresses), 73,900% (nodal displacements) and 73,332.5% (element buckling) in the cantilevered bar truss case).

In the sizing optimization of the 200 bar truss, the structural weights found by classical SA and by the TRLP-DOT gradient based optimizer were slightly heavier than their ISA counterpart. However, the largest difference in weight resulted less than 0.5%. Remarkably, the optimized design found by ISA in the case of a 30,000 psi stress limit was about 300 kg lighter than the best solution recently quoted in literature (Lamberti and Pappalettere, 2003a). It is to be noticed from Table 1 that each optimizer used in this study designed a structure which is lighter than the 12,823.808 kg weighted structure designed by the LESLP code in the 96 design variable case (see Lamberti and Pappalettere, 2003a). This fact proves that the algorithms to which ISA was compared in this research have global optimization capabilities that allowed them to be consistent with the design freedom included in the optimization thus avoiding designs that are clearly local optima.

ISA was more efficient than LESLP also for the more severe stress limit of 10,000 psi. However, the weight improvement was less than 0.2%. Finally, TRLP-DOT found the heaviest design when the optimization started from a feasible point.

In the configuration optimization problem, the weight improvement achieved by ISA with respect to the LSTRLP algorithm recently published in literature (see Lamberti and Pappalettere, 2004) was more than 80 kg. The present code was superior over classical SA and TRLP-DOT obtaining now weight reductions of 2% and 3%, respectively. When the optimization started from an infeasible point, all of the optimization codes considered in this study exhibited a weight penalty. However, ISA was able to design a structure yet lighter than the best solution quoted in literature.

The data reported in Table 1 show also that ISA designed the very efficient structures within less cooling cycles than classical SA. This fact is very significant because SA based algorithms that perturb simultaneously all the optimization variables result, in general, faster than SA algorithms where optimization variables are perturbed one by one. However, global annealing algorithms may converge to sub-optimal designs (see the discussion in Yu Chen and Su, 2002).

The use of global annealing allowed ISA in the sizing optimization problems to save between 30% and 60% of CPU time with respect to classical SA. Reductions in computational time achieved by ISA were more significant in the configuration optimization problem. In particular, ISA was more than twice as fast as classical SA. The global annealing strategy implemented by ISA allowed to reduce significantly the number of constraint evaluations. Since the number of structural analyses quoted in the table includes also constraint function linearizations, the gradient calculation task executed by ISA in each linearization was computationally not too expensive.

The TRLP-DOT code based on Sequential Linear Programming was the fastest optimizer overall in the case of the 200 bar truss optimized with a 30,000 psi stress limit but resulted slower than ISA in the 10,000 psi stress limit case. However, the computational efficiency of TRLP-DOT is due in most part to the fact that the linear solver originally used in LESLP was replaced in the present study by the modified feasible direction solver implemented in TRLP-DOT. In fact, Table 1 shows that ISA and TRLP-DOT required almost the same number of structural analyses in all of the optimization problems. In addition, TRLP-DOT formulated and solved much more sub-problems than LESLP within about the same CPU time (this statement implies data scaling to account for differences in computer speed). However, the very large number of iterations required by TRLP-DOT in the configuration problem caused the gradient based optimizer to be about three times less fast than ISA.

In order to gather more evidence of the relative behavior of the different algorithms considered in this study, convergence curves recorded in the optimization process and constraint margins are reported, respectively, in Figs. 7–10.

Fig. 7 shows the convergence curves recorded for the 200 bar truss structure. As far as it concerns the 30,000 psi stress limit, it appears that ISA was much faster than classical SA in reducing the cost function when the optimization started from a feasible point. In addition, in the case of infeasible initial design, ISA found a feasible design already in the first cooling cycle while classical SA required seven cooling cycles. Therefore, the improvement routines implemented in ISA enhanced significantly the overall performance of the optimizer.

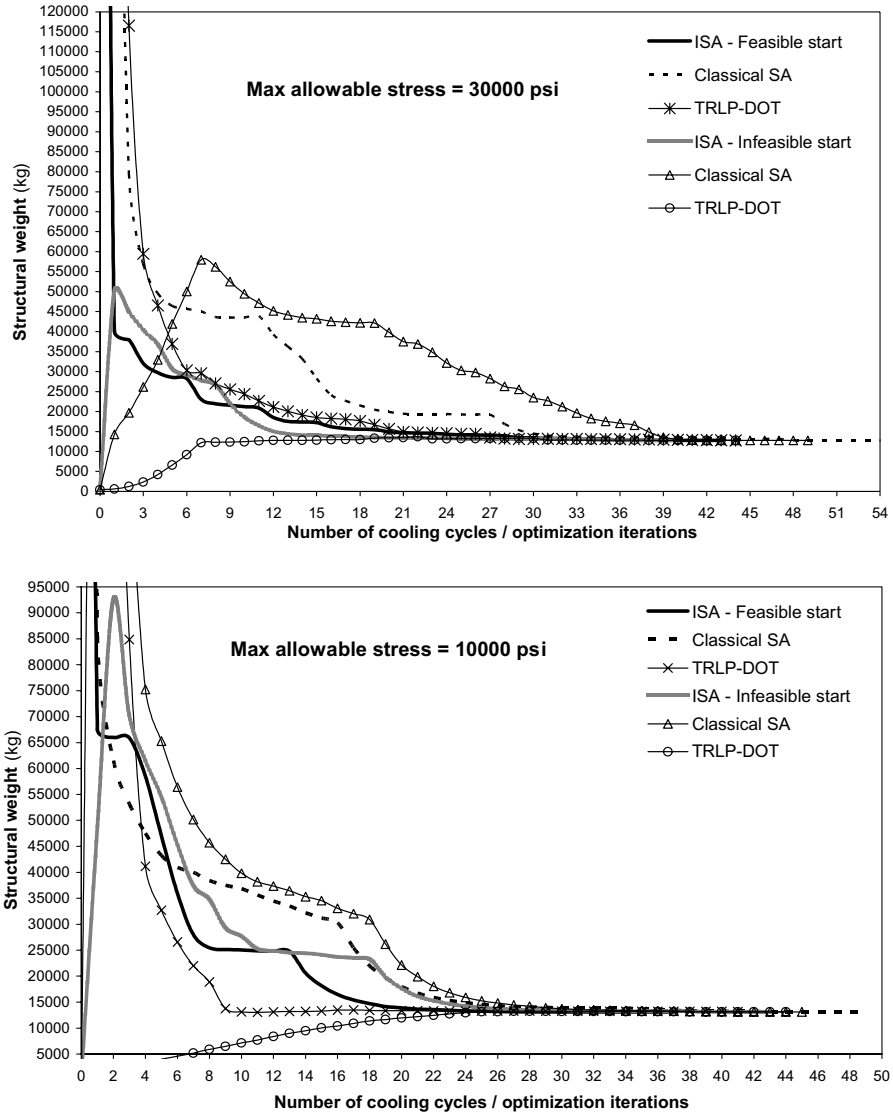

Fig. 7. Convergence curves for different algorithms in the 200 bar truss case.

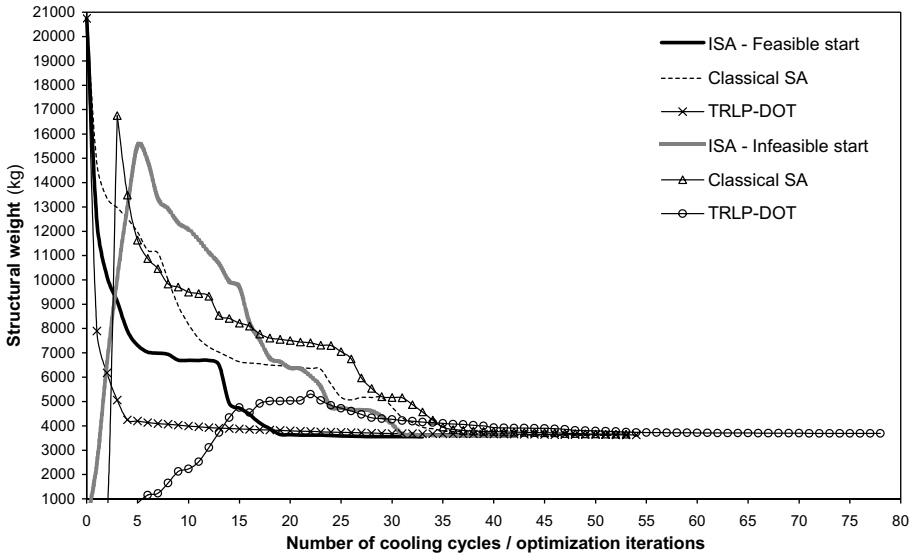

Fig. 8. Convergence curves for different algorithms in the cantilevered bar truss case.

The convergence behavior of ISA was comparable to that exhibited by TRLP-DOT. This is because ISA perturbs the design variables by moving along descent directions. When the starting design was infeasible, TRLP-DOT required about 20 iterations to enter in a feasible region because of the tight move limits imposed on the optimization variables. Although the heaviest intermediate design found by ISA in this case was about four times as large as that found by TRLP-DOT, the convergence curves of ISA and TRLP-DOT practically coincided after 21 iterations.

Similar results were found for the 10,000 psi stress limit case. ISA and TRLP-DOT convergence curves coincided after about 25 iterations in spite of the fact that ISA heaviest intermediate design was now seven times as large as that found by TRLP-DOT. Convergence behavior of classical SA resulted very sensitive to stress limit value and the optimizer was forced to increase the cost function up to about 250,000 kg (i.e., five times more than in the 30,000 psi stress limit case) in order to enter in a feasible region.

An interesting point suggested from data in Table 1 is the following. The difference in optimized weights in the 200 bar truss case is marginal (less than 0.5%). It is generally acknowledged that minor weight improvements in large-scale problems may be produced by the interaction between constraint tolerances and permitted constraint violations rather than by differences in algorithm formulations. However, the present authors want to point out that no constraint violation was permitted in this study. Therefore, any difference in weight reported in Table 1 will indicate that the corresponding optimization algorithm converged prematurely or got stuck in local minima. This statement is supported by the analysis of constraint margins plotted in Fig. 9. It can be seen that ISA was able either to recover immediately constraint violation and to keep design search process inside a feasible region throughout the optimization process. Conversely, TRLP-DOT had to resize repeatedly the move limits since many intermediate designs violated constraints more than at the previous optimization cycle or the constraint violation decreased too slowly. Further evidence of this argument is provided by the fact that classical SA, which does not use move limits found always structural weights lighter than their DOT-TRLP counterpart even though some intermediate designs were infeasible.

In the configuration optimization problem, Fig. 8 shows that TRLP-DOT was apparently faster than ISA in reducing structural weight when optimization started from a feasible design. However,

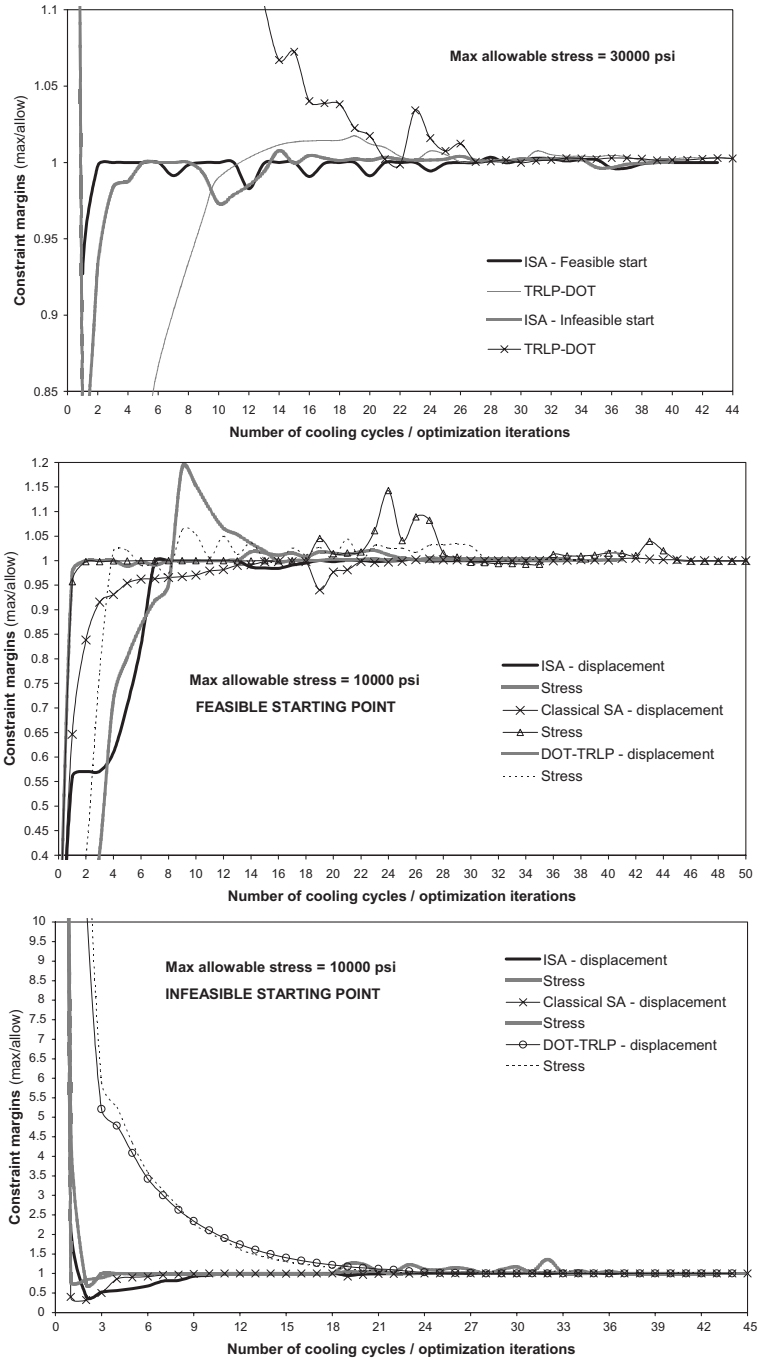

Fig. 9. Constraint margins in the 200 bar truss case.

constraint margin analysis reveals that TRLP-DOT intermediate designs violated buckling (up to 25%) and stress constraints (up to 5%) in the first 10 design cycles (see Fig. 10). Moreover, the gradient based optim-

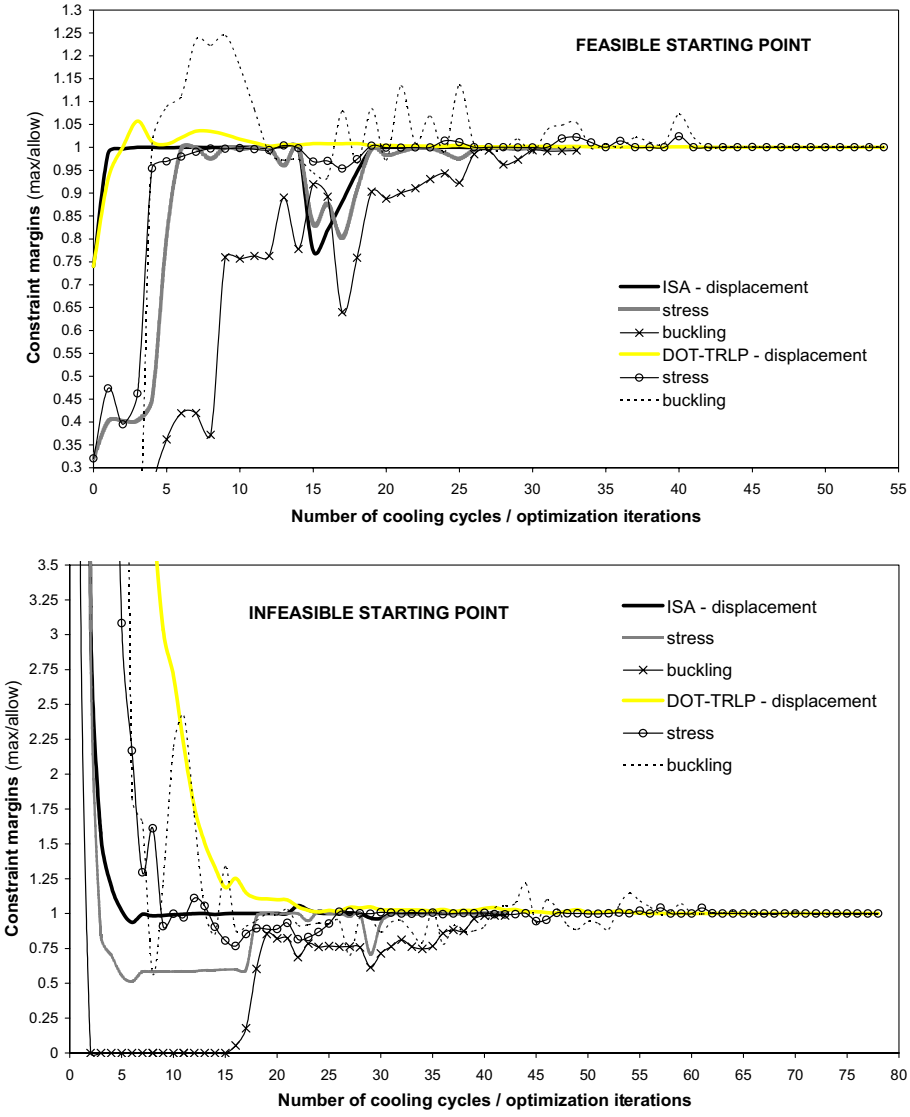

Fig. 10. Constraint margins in the cantilevered bar truss case.

izer was not able to eliminate constraint margin oscillations before than 40 iterations were completed (see again Fig. 10).

In the case of infeasible starting point, TRLP-DOT exhibited again large oscillations in buckling constraint margins (still 25% and 15% violation respectively after 44 and 54 optimization iterations) and required almost 30 iterations in order to keep intermediate designs almost feasible for some consecutive optimization cycles. Conversely, ISA recovered immediately the constraint violation and found a completely feasible (i.e., all constraint margins are strictly less than 1) intermediate design within only five cooling cycles. It should be noticed that ISA was able to "isolate" the effect of each different type of constraint in the sense that buckling and stress constraints became immediately inactive and did not return critical

before 15 design cycles. This behavior can be explained as follows. The cantilevered bar truss has been proven to have some degree of non-convexity in the configuration variable space (see Lamberti and Pappalettere, 2003a,b). However, configuration variables may greatly affect the distribution of critical buckling loads in the truss members. Since simulated annealing deals with design space non-convexity, ISA was able to immediately turn buckling constraints uncritical.

Details of optimized designs of truss structures are not reported in the paper in order to save space. However, Fig. 11 shows the distribution of the ζ_{SIZ} ratio between the optimized cross section values

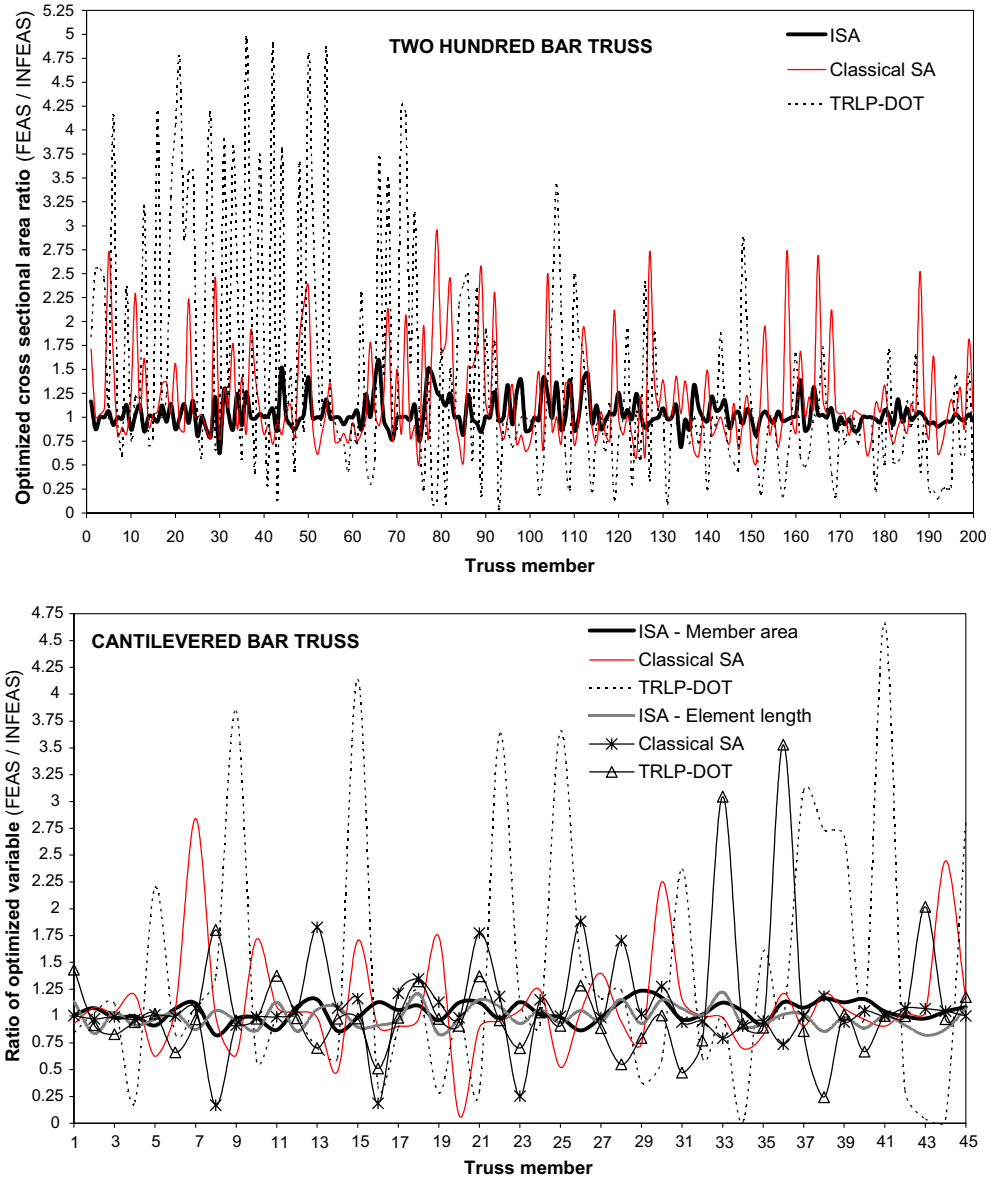

Fig. 11. Sensitivity of optimized designs to starting points.

Table 2
Comparison of numerical efficiency of ISA and ANSYS in the material identification problem when target displacements are generated numerically

Load	Run 1 Three-point-bending	Run 2	Run 3	Run 4 In-plane shear	Run 5	Run 6	Run 7 Bending	Run 8 Shear
Displacement included in ψ	u_x	u_x	u_y	u_x	u_x	u_y	u_{TOT}[a]	u_{TOT}[a]
Initial elastic constants[a] u_{TOT}[b]	$E_x = 27{,}500$ $E_y = 20{,}000$ $G_{xy} = 4000$ $v_{xy} = 0.308$	$E_x = 4000$ $E_y = 3000$ $G_{xy} = 20{,}000$ $v_{xy} = 0.05$	$E_x = 4000$ $E_y = 3000$ $G_{xy} = 20{,}000$ $v_{xy} = 0.05$	$E_x = 27{,}500$ $E_y = 20{,}000$ $G_{xy} = 4000$ $v_{xy} = 0.308$	$E_x = 4000$ $E_y = 3000$ $G_{xy} = 20{,}000$ $v_{xy} = 0.05$	$E_x = 4000$ $E_y = 3000$ $G_{xy} = 20{,}000$ $v_{xy} = 0.05$	$E_x = 4000$ $E_y = 3000$ $G_{xy} = 20{,}000$ $v_{xy} = 0.05$	$E_x = 4000$ $E_y = 3000$ $G_{xy} = 20{,}000$ $v_{xy} = 0.05$
Initial% error on displacements								
Average	12.8	89.9	287	23.0	35.4	106	297	37.4
Maximum	32.9	608	324	24.8	60.5	119	387	120
Calculated elastic constants ISA	$E_x = 24{,}998$ $E_y = 21{,}975$ $G_{xy} = 4999$ $v_{xy} = 0.280$	$E_x = 25{,}001$ $E_y = 21{,}966$ $G_{xy} = 5003$ $v_{xy} = 0.280$	$E_x = 25{,}000$ $E_y = 22{,}000$ $G_{xy} = 5000$ $v_{xy} = 0.280$	$E_x = 24{,}993$ $E_y = 21{,}999$ $G_{xy} = 5000$ $v_{xy} = 0.280$	$E_x = 24{,}993$ $E_y = 21{,}995$ $G_{xy} = 5000$ $v_{xy} = 0.280$	$E_x = 25{,}011$ $E_y = 22{,}000$ $G_{xy} = 4999$ $v_{xy} = 0.280$	$E_x = 24{,}996$ $E_y = 21{,}996$ $G_{xy} = 5002$ $v_{xy} = 0.279$	$E_x = 24{,}992$ $E_y = 21{,}987$ $G_{xy} = 5001$ $v_{xy} = 0.279$
Residual % error on displacements								
Average	0.0346	0.0349	0.00	0.000689	0.00274	0.0594	0.00379	0.00468
Maximum	0.2900	0.3120	0.00	0.001610	0.00422	0.9810	0.01740	0.00920
Calculated elastic constants ANSYS	$E_x = 25{,}035$ $E_y = 22{,}037$ $G_{xy} = 5019$ $v_{xy} = 0.281$	$E_x = 24{,}997$ $E_y = 21{,}971$ $G_{xy} = 5003$ $v_{xy} = 0.281$	$E_x = 24{,}997$ $E_y = 21{,}999$ $G_{xy} = 5003$ $v_{xy} = 0.279$	$E_x = 25{,}000$ $E_y = 21{,}995$ $G_{xy} = 5000$ $v_{xy} = 0.281$	$E_x = 25{,}076$ $E_y = 22{,}025$ $G_{xy} = 5000$ $v_{xy} = 0.280$	$E_x = 25{,}153$ $E_y = 21{,}900$ $G_{xy} = 4950$ $v_{xy} = 0.283$	$E_x = 25{,}001$ $E_y = 22{,}014$ $G_{xy} = 4997$ $v_{xy} = 0.281$	$E_x = 25{,}002$ $E_y = 22{,}042$ $G_{xy} = 4998$ $v_{xy} = 0.281$
Residual % error on displacements								
Average	0.259	0.215	0.0139	0.00363	0.0144	0.403	0.00556	0.00643
Maximum	1.785	2.140	0.0189	0.02660	0.0241	3.290	0.02241	0.01210

Target values: $E_x = 25\,\text{GPa}$; $E_y = 22\,\text{GPa}$; $G_{xy} = 5\,\text{GPa}$; $v_{xy} = 0.28$.

[a] $u_{TOT} = \sqrt{u_x^2 + u_y^2}$.

[b] $u_{TOT} = \sqrt{u_x^2 + u_y^2}$.

obtained when the optimization started from a feasible point and an infeasible point, respectively. For a given element, $\zeta_{SIZ} = 1$ will indicate that the optimized design is insensitive to the starting point. Should a design be near to the global optimum, the ζ_{SIZ} parameter will be very close to 1 for all truss members. It is apparent from the figure that ISA achieved the smallest dispersion (about 2.5%) of the ζ_{SIZ} parameter around 1 while TRLP-DOT exhibited the largest dispersion. This is consistent with the fact that TRLP-DOT terminated prematurely the optimization process. Behavior of classical SA was similar to ISA but the ζ_{SIZ} parameter exhibited a much larger number of local peaks and valleys.

For the cantilevered bar truss case, Fig. 11 shows also the distribution of the ζ_{LEN} ratio between the optimized lengths of each truss element for the two initial designs. It can be seen that ISA was again the most robust algorithm overall. Interestingly, element lengths optimized by ISA increased averagely by 1% when the design process started from an infeasible point. This fact is consistent with constraint margin plots of Fig. 10 and confirms that ISA was able to carry out the entire design search process inside regions where buckling constraints are not critical (in spite of the 73,332.5% initial violation) thus reducing the role played by the configuration design space. Conversely, classical SA and TRLP-DOT had to reduce significantly (between 2 and 3 times) the length of some elements, which resulted very sensitive to buckling.

Table 2 summarizes the results of the eight preliminary numerical trials of material identification carried out with the ISA code and the ANSYS program. Table 2 reports also the type of load applied to the specimen, type of displacement included in the Ψ error function and the initial guess on elastic constants.

It is apparent from Table 2 that ISA is insensitive to how the Ψ error function is defined as well as is insensitive to the loads applied to the specimen and to the initial value of elastic constants. The table shows also that ISA was superior over the ANSYS optimization module. In fact, the percentage average and maximum errors on displacements computed in correspondence of the values of elastic constants determined by ISA were always less than 0.06% and 1%, respectively. In Run 3, the present code was even able to exactly determine the values of elastic constants. Conversely, ANSYS exhibited residual errors up to 3.29% (Run 6).

Further evidence of the relative merits of ISA and ANSYS can be gathered from Fig. 12 which shows, for Runs 7–8, the plots of the Ψ error function versus the number of cooling cycles/optimization iterations. It can be seen that ISA reduced very quickly the value of Ψ while the ANSYS code still exhibited an oscil-

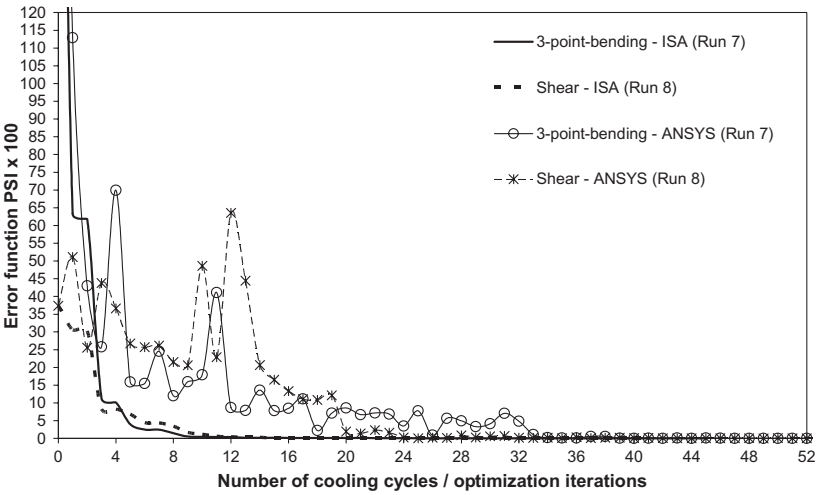

Fig. 12. Comparison of convergence curves obtained with ISA and ANSYS in the material identification problem when the ψ function includes u_{TOT} displacements.

latory behavior when ISA already had practically converged. The same relative behavior shown in Fig. 12 was seen in the other six numerical trials. This fact suggested not to use the ANSYS optimization routine in the experimental case where determination of displacements from interferometric patterns might have introduced noises in the optimization in fashion of local peaks of the error function Ψ.

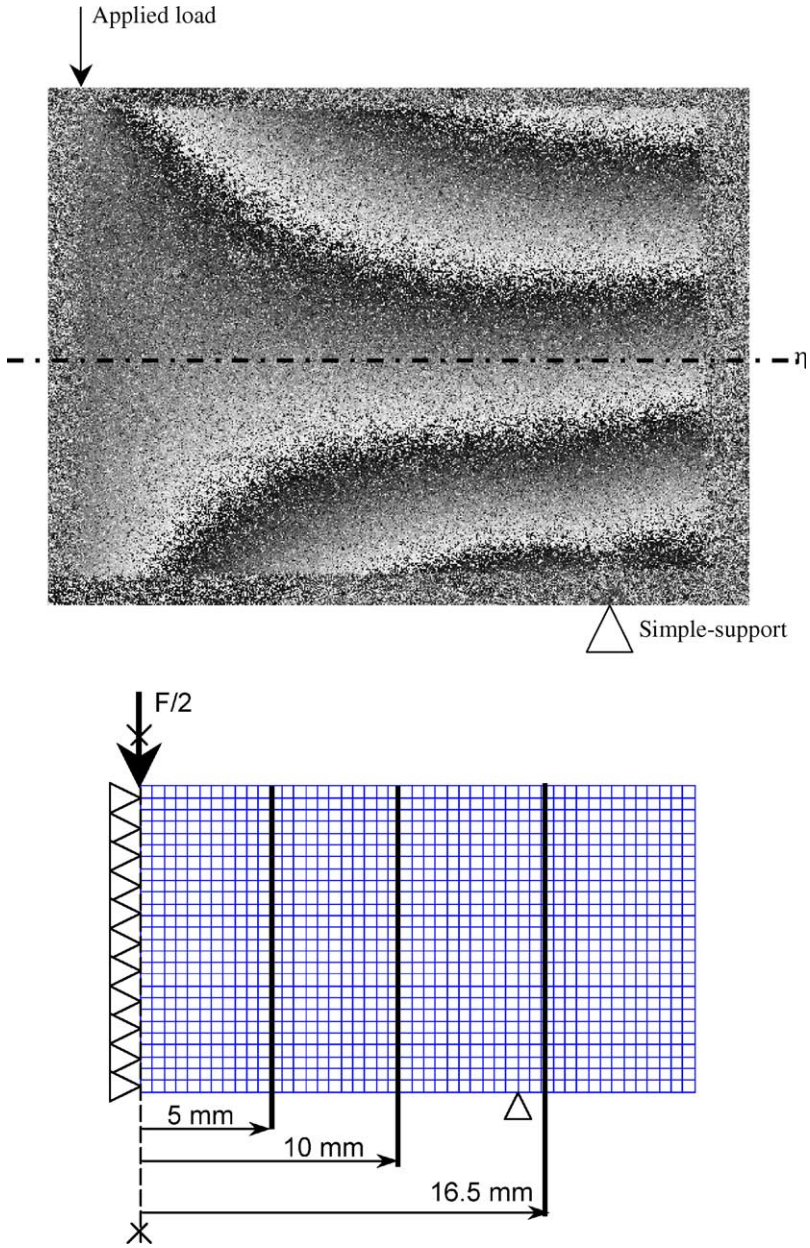

Fig. 13. Phase map obtained with PS-ESPI and schematic of FEM model, loads and locations at which experimental and numerical data are compared.

In the experimental case, ISA was not compared to any other optimization code. Besides the aforementioned argument on ANSYS, classical SA is not attractive since it implies too many FEM analyses. Finally, preliminary tests on TRLP-DOT showed that the performance (number of iterations and total CPU time) of the SLP based optimizer was sensitive to the initial guess on the material properties because of the non-smooth behavior of the sensitivities of the Ψ function. Although also ISA utilizes gradient information, the random generation process included in the present algorithm ensures enough design freedom allowing thus to by-pass quickly the non-smooth regions of design space.

Fig. 13 shows the phase distribution in the composite specimen determined by means of PS-ESPI in the experimental case. It can be seen that the phase fringes reproduced the u_x displacement field in a continuous fashion. The horizontal axis η where the ε_x strains ideally get equal to zero (bending neutral axis), the loaded zone and the simply supported zone are also indicated in the figure. In addition, the figure shows the FEM model used for structural analysis along with the locations at which numerical and experimental data were compared. Because of symmetry about the vertical axis and limitations in size of the optics used in the experiments, the region of interest considered along the x-direction was about 18 mm long.

Since the PS-ESPI set-up used in the experiments allowed us to measure horizontal displacements (see Fig. 5), the Ψ error function was built by comparing the numerical and experimental values of u_x displacements at the three locations indicated in Fig. 13. The Ψ error function was minimized in all those three cases. Finally, an additional optimization run (indicated as "ALL" in the rest of the paper) included all the nodes/pixels considered in the other three runs.

Table 3 presents the results of the material identification problem for the experimental case. Different initial guesses on material properties were made in order to introduce more uncertainty in the optimization. However, the ISA optimizer was again insensitive to the starting point and hence recovered the considerably large initial percentage errors. The values of elastic constants determined by ISA were very close to the target values. Although the average and maximum percentage errors were obviously larger than in the test cases with numerically generated target displacements (see Table 2), the accuracy achieved in the experimental case is however acceptable since the largest error was less than 3%. This residual error was certainly caused by uncertainty factors such as electronic noise mixed with interferometric patterns, overall efficiency of filtering, local de-correlation of speckle patterns, etc.

Table 3
Sensitivity of ISA to the set of experimental data used in the material identification problem

Location at which data are compared	$x = 5$ mm	$x = 10$ mm	$x = 16.5$ mm	All locations
Initial values of elastic constants[a]	$E_x = 5000$ $E_y = 3000$ $G_{xy} = 2000$ $v_{xy} = 0.1$	$E_x = 10{,}000$ $E_y = 8000$ $G_{xy} = 2000$ $v_{xy} = 0.4$	$E_x = 30{,}000$ $E_y = 10{,}000$ $G_{xy} = 8000$ $v_{xy} = 0.01$	$E_x = 3000$ $E_y = 2000$ $G_{xy} = 1000$ $v_{xy} = 0.01$
Initial % error on u_x				
Average	140.1%	146.6%	37.6%	281%
Maximum	197.4%	226.2%	148.4%	870%
Calculated values of elastic constants	$E_x = 25{,}016$ $E_y = 22{,}049$ $G_{xy} = 5002$ $v_{xy} = 0.281$	$E_x = 25{,}048$ $E_y = 21{,}989$ $G_{xy} = 5035$ $v_{xy} = 0.279$	$E_x = 25{,}031$ $E_y = 21{,}963$ $G_{xy} = 4969$ $v_{xy} = 0.279$	$E_x = 25{,}043$ $E_y = 22{,}034$ $G_{xy} = 5000$ $v_{xy} = 0.279$
Residual % error on u_x				
Average	1.351%	1.469%	1.118%	2.033%
Maximum	2.711%	2.481%	2.180%	2.775%

[a] Target values: $E_x = 25$ GPa; $E_y = 22$ GPa; $G_{xy} = 5$ GPa; $v_{xy} = 0.28$.

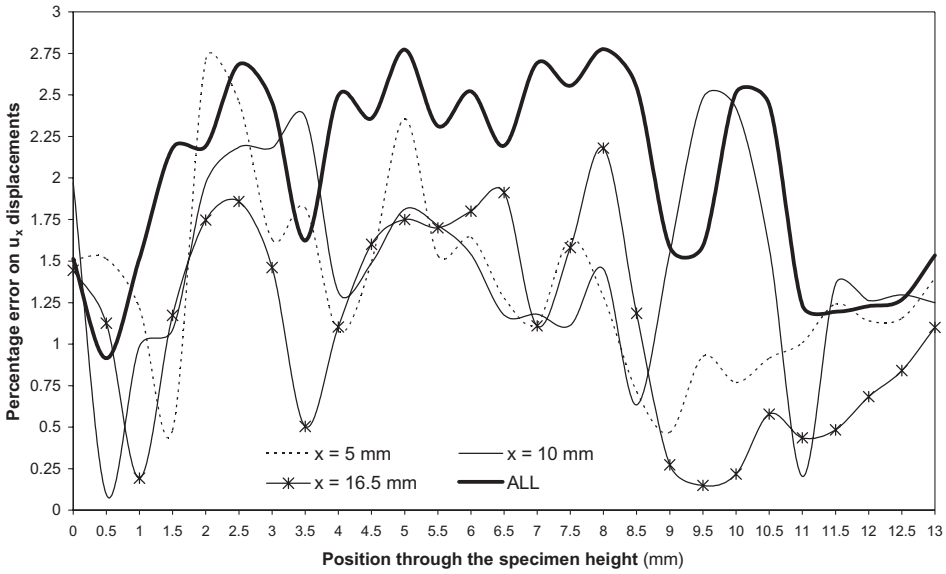

Fig. 14. Residual error on the u_x displacements in the three-point-bending experimental test.

Finally, Fig. 14 shows the percentage error on the u_x displacement computed at each FEM node/image pixel of the different control stations ("$x = 5$ mm", "$x = 10$ mm", "$x = 16.5$ mm", "ALL") when the values of elastic constants (i.e., the optimum design) found by ISA were given as input to the FEM model. It can be seen that the maximum error occurs near the middle of the specimen height: that is, where the bending neutral axis is approximately located. This happened certainly because the horizontal displacements are much smaller near the neutral axis than at the bottom and at the top of the specimen and hence the optimizer is more sensitive to sudden changes in the horizontal displacement sign which may result in local peaks of the Ψ error function.

5. Summary and conclusions

This paper described a novel optimization algorithm implementing an enhanced formulation of simulated annealing (SA). The algorithm—denoted as ISA (improved simulated annealing)—combined global and local annealing search mechanisms. In order to speed up convergence, the ISA algorithm included gradient information and defined descent directions whose components are generated randomly. In addition, in case of infeasible intermediate designs, ISA perturbed the design by moving along directions such that cost function may improve and constraints get less violated. Local annealing where optimization variables are perturbed one by one was performed each time global annealing did not yield significant improvements in design. Linear approximations controlled by a trust region model were used in order to reduce the number of exact analyses thus saving computational time. Finally, the cooling schedule implemented in ISA reduced adaptively the temperature based on the convergence behavior exhibited during the optimization process.

In this study, the ISA algorithm was tested in complicated optimization problems that exhibit non-convex and/or non-smooth behavior: (i) the weight minimization of a large-scale truss structure with 200 design variables and 3500 non-linear constraints; (ii) the configuration optimization of a cantilevered bar truss with 81 design variables; (iii) a typical example of reverse engineering where the goal was to

characterize the in-plane behavior of an eight-ply woven reinforced fiberglass-epoxy composite laminate by minimizing the difference between the displacements computed with FEM analyses and the displacements measured experimentally with a powerful optical technique referred to as phase shifting electronic speckle pattern interferometry (PS-ESPI).

It is to be noticed that the aforementioned test cases were chosen so that gradients are not available explicitly. This fact allowed to overcome the apparent limitation of ISA with respect to classical simulated annealing which does not require gradient information.

ISA was compared to a classical implementation of SA, state-of-art optimization codes and commercial software. The results indicated that ISA was the most efficient algorithm in truss optimization problems allowing weight reductions between 80 and 300 kg with respect to the best solutions recently quoted in literature. It should be noticed that the present algorithm was clearly superior over classical SA and its convergence speed resulted comparable to a gradient based optimizer implementing a sophisticated version of Sequential Linear Programming.

In order to check the suitability of ISA for material identification problems, eight trial cases with numerically generated target displacements were carried out. As expected, ISA resulted insensitive to the load type, initial guess of elastic constants and target displacements included in the Ψ error function. Moreover, ISA outperformed the optimization routine available in the well known commercial general purpose FEM software ANSYS.

Hence, we used ISA also in a real case where horizontal displacements of a 46 mm long, 13 mm tall and 1.2 mm thick specimen subjected to three-point-bending were measured by means of PS-ESPI. Three sets of points (image pixels/FEM nodes) were chosen at different locations on the specimen in order to build the error function. Four optimizations runs were carried out: one for each set of points and the last one for all of the points simultaneously.

ISA proved itself capable to accurately characterize the behavior of the 8-ply composite laminate. In fact, the residual error between the displacements measured by means of PS-ESPI and those computed by FEM analysis was less than 3% at the end of the identification process. This error was probably due to uncertainty factors inherent to PS-ESPI (electronic noise, filtering efficiency, etc.) and to numerical instabilities caused by local peaks of the Ψ error function.

The complexity and variety of the optimization problems successfully solved in this study certainly support the conclusion that ISA is a powerful optimization code. However, the present authors point out that ISA should be tested further in other optimization problems in order to draw more general conclusions.

Acknowledgment

The authors gratefully thank the Department of Mechanical and Aerospace Engineering (MAE) of the University of Florida, Gainesville (USA), for the network access to computing resources.

Appendix A. Effect of design perturbation strategy in simulated annealing

This appendix presents the results of a trade study aimed to investigate on how the overall performance of simulated annealing may be sensitive to the random generation mechanism of trial designs implemented by the optimizer. Such an investigation served to choose a SA based algorithm, which could be used as a significant basis for comparison with the sophisticated algorithm ISA (improved simulated annealing) presented in the main part of this paper. In order to accomplish this task, we tested three different annealing schemes in two well-known optimization problems that have local minima: the Rosenbrock's problem and the weight minimization of a 10 bar truss under stress and displacement constraints. The optimization runs

started very far from the global minimum of each problem in order to check if the different optimizers were able to reduce the cost function quickly and to approach the optimum within not too many cooling cycles.

The Rosenbrock's valley problem is often referred to as the "Banana function problem". The analytical formulation of the optimization problem is the following:

$$\begin{cases} \min W = \sum_{i=1}^{\mathrm{NDV}-1} [100 \cdot (x_{i+1} - x_i^2)^2 + (1 - x_i)^2], \\ -2.048 \leqslant x_k \leqslant 2.048 \end{cases} \quad (k = 1, \ldots, \mathrm{NDV}), \tag{A.1}$$

where NDV is the number of optimization variables. The problem statement implies that the global minimum of the cost function is reached if all the design variables are equal to 1. Therefore, the global minimum of the cost function is 0.

The "banana problem" denotation is because the cost function W expressed by (A.1) represents in the NDV-dimensional design space an hyper-surface with a long, narrow, parabolic shaped flat valley. This valley contains the global optimum, which is hidden among many local minima. Here, the optimization run included 50 design variables. The design process started with all the optimization variables set to 2.

The 10 bar truss structure shown in Fig. A.1 is to be designed for minimum weight. This sizing optimization problems includes 10 design variables each of which is the cross sectional area of a truss member (see expression (11) in the main part of this paper). The lower bound of the design variables is 0.1 in.^{2} (0.00064516 m^2). Nodal displacements of the free nodes are constrained to be less than 2 in. (0.0508 m) while the allowable stress (the same in tension and in compression) is 25,000 psi (17.577 kgf/mm^2). Downward forces of 100,000 lbf (45,359.240 kgf) are applied at nodes 2 and 3. Previous investigations found that this optimization problem is non-convex and has a local minimum of 2302.7 kg besides the global minimum of 2295.6 kg. Here, we started the optimization from a very feasible point: design variables at the upper bound (100 in.^{2}: that is, 0.64516 m^2).

The two non-convex problems described above were hence solved by means of three annealing strategies that included different mechanisms for generating trial designs. The algorithms are briefly described in the following.

(i) The optimization variables are perturbed one at a time and the optimizer performs function evaluation each time a new trial design is generated. Each design variable is updated as follows:

$$x_{\mathrm{TR},j} = x_{\mathrm{OPT},j} + s_{\mathrm{NRD}} \cdot (x_j^{\mathrm{u}} - x_j^{\mathrm{l}}) \cdot \rho_j \cdot \gamma_K \quad (j = 1, \ldots, \mathrm{NDV}), \tag{A.2}$$

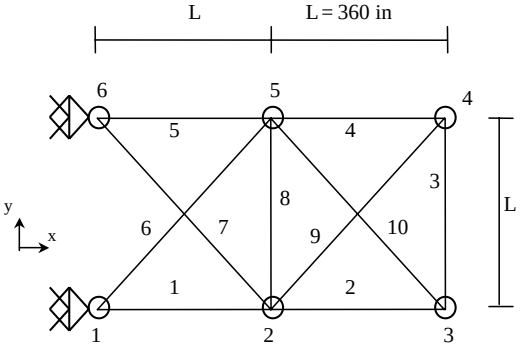

Fig. A.1. Schematic of the 10 bar truss structure.

where ρ_j is a random number in the interval $(0,1)$; s_{NRD} is set to 1 or -1 if the random number ρ_j is greater or smaller than 0.5, respectively; γ_K is a knockdown factor defined in fashion of expression (4) in the main part of this paper; K indicates the current cooling cycle.
This strategy is formally equivalent to the "local annealing" scheme included in the ISA algorithm (see description of Step 6 in the main part of this paper).

(ii) All the NDV optimization variables are perturbed simultaneously. Since the trial design vector $\mathbf{X}_{\mathrm{TR}}$ includes the optimization variables defined in fashion of expression (A.2), NDV random numbers ρ_j must to be generated in order to define each new trial design point.

(iii) All the NDV optimization variables are perturbed simultaneously. Again, definition of each new trial design $\mathbf{X}_{\mathrm{TR}}$ in fashion of expression (A.2) implies generating NDV random numbers. However, since the cost function of the problem is to be minimized, the s_{NRD} parameter for perturbing the design jth variable is now set to 1 or -1 if it holds, respectively, $\partial W/\partial x_j < 0$ or $\partial W/\partial x_j > 0$. This strategy is similar to the "global annealing" scheme included in ISA (see description of Step 4 in the main part of this paper) but the steps chosen for each design variable are not centered about the current optimal design.
Here, the $\partial W/\partial x_j$ sensitivities could be easily computed since the banana problem has an explicit cost function while the cost function (weight) of the truss is linear in sizing variables (i.e., the cross sectional areas of the truss members).

In summary, algorithm (i) is based on an 1-directional annealing search strategy while algorithms (ii) and (iii) implement multi-directional annealing search strategies. For this reason, algorithms (i), (ii) and (iii) are respectively denoted as "1-directional", "Multi-directional—no cost gradient" and "Multi-directional—with cost gradient" in the rest of the paper. In all of the algorithms mentioned above, design variables are changed following their order sequentially.

In addition, the annealing algorithms compared in this appendix include a cooling schedule with an uniform cooling rate. The temperature to be used in the Kth cooling cycle is reset as $T_K = \beta_K T_{K-1}$ by using the constant value $\beta_K = 0.95$. Finally, the stopping criterion is the same as the one used for the ISA algorithm (see expression (9) in the main part of the paper).

The convergence curves recorded in the two optimization problems for the three annealing strategies described above are shown in Fig. A.2. The initial value W_0 of the cost function is shown in the figure. The plots reported in Fig. A.2 include also the convergence curves obtained with the Sequential Quadratic Programming (SQP) optimization routine implemented in the commercial software Matlab® 6.1 (2001). This commercial optimizer was utilized in the present study because it proved to be very efficient in solving complicated structural optimization problems with multiple design spaces (see Lamberti and Pappalettere, 2004).

It can be seen from Fig. A.2 that algorithms (ii) and (iii) reduced, in the early cooling cycles, the cost function in a much faster way than algorithm (i). As expected, including gradient information in the generation of trial designs allowed to increase the cost reduction rate. However, the 1-directional search recovered the gap in convergence speed within about 55 cooling cycles in the Rosembrock's valley problem and within about 25 cooling cycles in the 10 bar truss case. This behavior is consistent with the argument that perturbing all design variables could not keep the current good searching direction also in the subsequent iterations. Conversely, the 1-directional search has a better chance to recover the good searching direction and approach the optimum point systematically. In addition, perturbing all design variables does not give any information on the effect of each single variable on constraint violation. Since, the sharp reduction in cost initially seen in the case of multi-directional algorithms resulted in approaching constraint boundaries too quickly, the search direction defined by the multi-directional random generation was certainly not conservative when the number of active constraints increased significantly within the current cooling cycle.

Fig. A.3 shows the optimized designs found by the different annealing strategies and by Matlab in the Banana problem. The standard deviation of the optimized values of design variables is also reported in

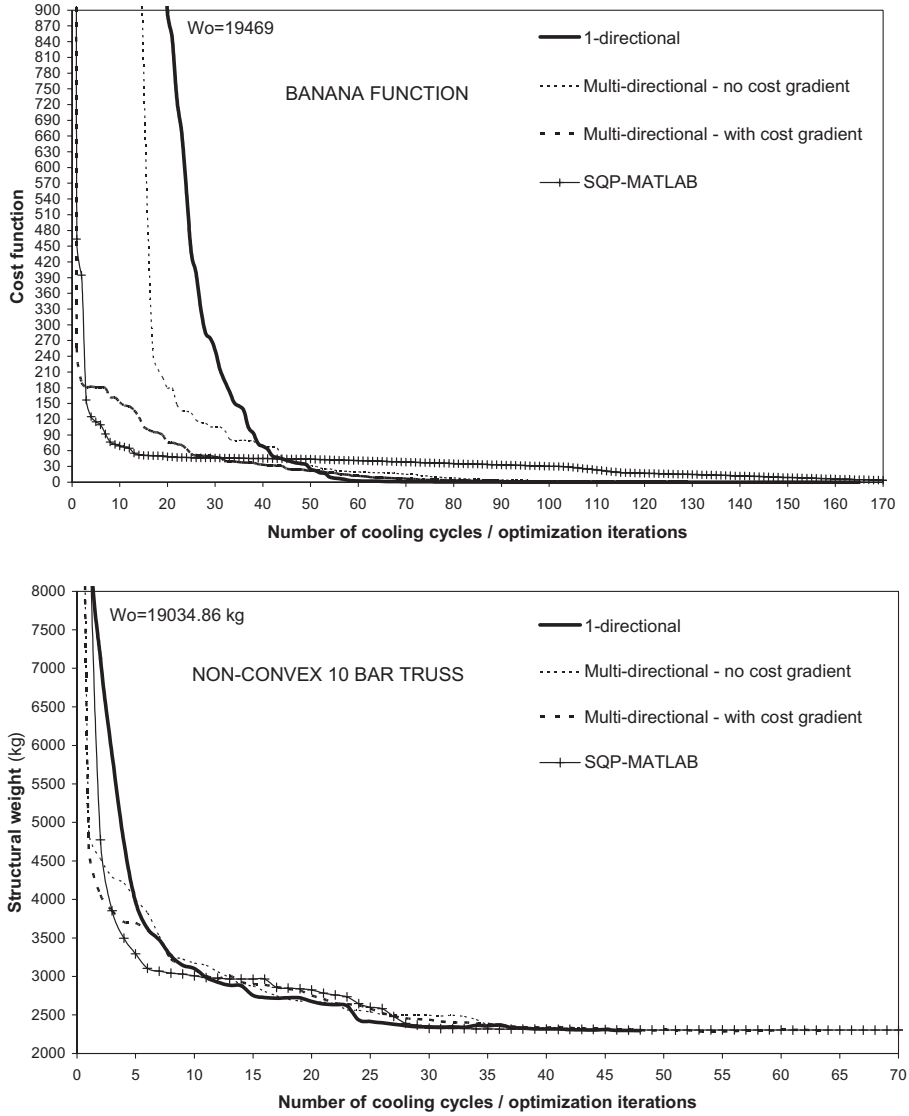

Fig. A.2. Convergence curves obtained in preliminary tests for different annealing strategies.

the figure. It is apparent that the SQP optimizer used by Matlab found the best design. This was because the cost function of the Rosenbrock's problem includes many quadratic terms (in fashion of single contributes or of product factors). The 1-directional annealing strategy found the closest design to that obtained by Matlab but the last 10 design variables were up to 0.4% larger than the optimum target value of 1. The multi-directional search strategies were not very effective and the error on design variables was up to 1%.

In the 10 bar truss case, the fact that the cost function is linear made the SQP routine implemented in Matlab miss the 2295.6 kg global optimum and to converge to the 2302.7 kg local minimum. Conversely, the annealing algorithms (i) and (iii) (i.e., "1-directional" and "multi-directional with gradients") were able to find better designs than Matlab. However, algorithm (ii) failed and converged to a design about 101 kg

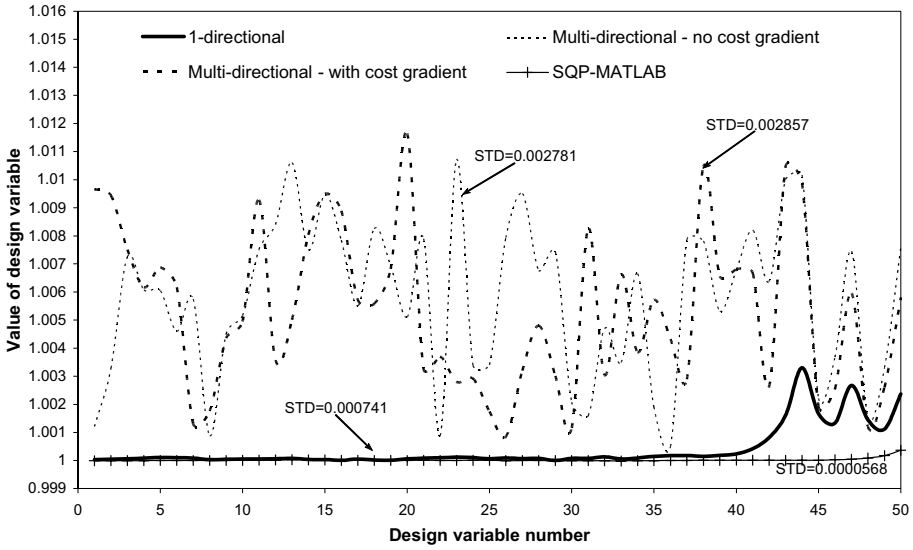

Fig. A.3. Optimized designs obtained in the banana function case.

heavier than the 2295.6 kg global optimum. This failure along with the fact that, in the Banana problem, the convergence speed of algorithm (ii) was lower than that of algorithm (iii) suggested us not to use multi-directional annealing search without including gradient information.

In view of the results presented and of the discussion developed in this appendix, it is correct to choose the 1-directional perturbation mechanism as basis for comparison with ISA. In addition, an optimization code based on simulated annealing should combine the ability of the multi-directional search to reduce sharply the cost function in the early optimization cycle and the ability of the 1-directional search to recover the good search direction as the different cooling cycles progress. Therefore, the combination of global and local annealing done by ISA appears certainly well motivated.

References

ANSYS® 7.0, 3003. User's Manual, Swanson Analysis System Inc.

Blachut, J., 2003. Optimal barreling of steel shells via simulated annealing algorithm. Comput. Struct. 81, 1941–1956.

Cloud, G.L., 1998. Optical Methods of Engineering Analysis. Cambridge University Press, New York, USA.

Creath, K., 1985. Phase-shifting speckle interferometry. Appl. Opt. 24, 3053–3058.

Dhingra, A.K., Lee, B.H., 1994. A genetic algorithm approach to single and multiobjective structural optimization with discrete-continuous variables. Int. J. Numer. Meth. Eng. 37, 4059–4080.

DOT®, 1995. Users Manual, Version 4.20. Colorado Springs, Vanderplaats R&D Inc.

Genovese, K., Lamberti, L., Pappalettere, C., 2004. A comprehensive ESPI based system for combined measurement of shape and deformation of electronic components. Opt. Laser Eng. 42, 543–562.

Hasanceby, O., Erbatur, F., 2002. Layout optimization of trusses using simulated annealing. Adv. Eng. Softw. 33, 681–696.

Huang, M.W., Arora, J.S., 1997. Optimal design with discrete variables: some numerical experiments. Int. J. Numer. Meth. Eng. 40, 165–188.

Lamberti, L., Pappalettere, C., 2003a. Move limits definition in structural optimization with Sequential Linear Programming. Part I and II. Comput. Struct. 81, 197–238.

Lamberti, L., Pappalettere, C., 2003b. A numerical code for lay-out optimization of skeletal structures with Sequential Linear Programming. Eng. Comput. 19, 101–129.

Lamberti, L., Pappalettere, C., 2004. Improved Sequential Linear Programming formulation for structural weight minimization. Comp. Meth. Appl. Mech. Eng. 193, 3493–3521.

Leendertz, J.A., 1970. Interferometric displacement measurement on scattering surface utilizing speckle effect. J. Phys. E: Sci. Instrum. 3, 214–218.

Matlab® Version 6.1, 2001. The MathWorks Inc., <http://www.mathworks.com>.

Rao, S.S., 1996. Engineering Optimization. Wiley Interscience, New York, USA.

Venkayya, V.B., 1978. Structural optimization: a review and some recommendations. Int. J. Numer. Meth. Eng. 13, 203–228.

Wujek, B.A., Renaud, J.E., 1998. New adaptive move-limit management strategy for approximate optimization. AIAA J. 36, 1911–1934.

Yu Chen, T., Su, J.J., 2002. Efficiency improvement of simulated annealing in optimal structural designs. Adv. Eng. Softw. 33, 675–680.