Dominance Measures for Multi-Objective Simulated Annealing

Kevin I. Smith
K.I.Smith@exeter.ac.uk

Richard M. Everson
R.M.Everson@exeter.ac.uk

Jonathan E. Fieldsend
J.E.Fieldsend@exeter.ac.uk

Department of Computer Science
University of Exeter
Exeter, EX4 4QF. UK.

***Abstract*— Simulated annealing (SA) is a provably convergent optimiser for single-objective (SO) problems. Previously proposed MO extensions have mostly taken the form of an SO SA optimising a composite function of the objectives. We propose an MO SA utilising the relative dominance of a solution as the system energy for optimisation, eliminating problems associated with composite objective functions. We also propose a method for choosing perturbation scalings promoting search both towards and across the Pareto front.**

We illustrate the SA's performance on standard test problems. The new SA is shown to promote rapid convergence to the true Pareto front with a good coverage of points across it.

I. Introduction

A popular algorithm for solving single-objective optimisation problems (those in which the user cares only about a single dependant variable of the system) is simulated annealing (SA) [1]. Geman & Geman [2] provided a proof that simulated annealing, if annealed sufficiently slowly, converges to the global optimum. Although the required cooling rate is infeasibly slow for most purposes, simulated annealing often gives excellent results when run with a faster cooling schedule. It is frequently the case in optimisation problems, however, that there are several objectives of the system which the user is interested in optimising simultaneously. Clearly, simultaneous optimisation of several objectives is usually impossible and the curve (for two objectives) or surface (for three or more objectives) that describes the trade-off between objectives is known as the Pareto-front. Although there are several well developed genetic algorithms and evolutionary schemes to address such multi-objective problems (see, [3], [4] for recent reviews), simulated annealing does not, in its usual formulation, provide a method for optimising more than a single objective. Simulated annealing has been adapted to multi-objective problems by combining the objectives into a single objective function [5], [6], [7], [8], [9]; however, these methods either damage the proof of convergence, or are limited (potentially severely) in their ability to fully explore the trade-off surface.

We propose a modified simulated annealing algorithm which maps the optimisation of multiple objectives to a single-objective optimisation using the true trade-off surface, maintaining the convergence properties of the single-objective annealer while encouraging exploration of the full trade-off surface. A method of practical implementation is also described, which overcomes the limitation that the true trade-off surface is unavailable for most real-world problems, using the available non-dominated data points from the current optimisation.

In this paper we start by briefly discussing methods that combine objectives into a single composite objective. In section III we describe our dominance-based SA algorithm and, in section IV, methods for improving the quality of the optimisation energy measure when the available data points are few are described. Choosing an efficient scale for perturbations is an important component of scalar SA algorithms. The issue is further complicated in multi-objective algorithms because a perturbation may not only move the current state closer to or further from the Pareto front, but also transversally. In section V we describe a method for setting the scale of perturbations and other run-time parameters. Results showing that the algorithm converges on a range of test problems are given in section VI and we draw conclusions in section VII.

II. Background

A. Dominance and Pareto Optimality

In multi-objective optimisation we attempt to simultaneously maximise or minimise D objectives, y_i, which are functions of P variable parameters or decision variables, $\mathbf{x} = (x_1, x_2, \dots, x_P)$:

$$y_i = f_i(\mathbf{x}), \qquad i = 1, \dots, D \tag{1}$$

Without loss of generality, we assume that the objectives are to be minimised, so that the multi-objective optimisation problem may be expressed as:

$$\text{Minimise} \quad \mathbf{y} = \mathbf{f}(\mathbf{x}) \equiv (f_1(\mathbf{x}), \dots, f_D(\mathbf{x})) \tag{2}$$

The idea of dominance is generally used to compare two solutions $\mathbf{a}$ and $\mathbf{b}$. If $\mathbf{f}(\mathbf{a})$ is no worse for all objectives than $\mathbf{f}(\mathbf{b})$ and wholly better for at least one objective it is said that $\mathbf{a}$ *dominates* $\mathbf{b}$, written $\mathbf{a} \prec \mathbf{b}$. Thus $\mathbf{a} \prec \mathbf{b}$ iff:

$$\begin{aligned} f_i(\mathbf{a}) \leq f_i(\mathbf{b}) \quad \forall i = 1, \dots, D \quad \text{and} \\ f_i(\mathbf{a}) < f_i(\mathbf{b}) \quad \text{for at least one } i. \end{aligned} \tag{3}$$

Clearly the dominates relation is not a total order and two solutions are *mutually non-dominating* if neither dominates the other. A set F of solutions is said to be a non-dominating set if no element of the set dominates any other:

$$\mathbf{a} \not\prec \mathbf{b} \;\forall\; \mathbf{a}, \mathbf{b} \in F \tag{4}$$

A solution is said to be globally non-dominated, or Pareto-optimal, if no other feasible solution dominates it. The set of all Pareto-optimal solutions is known as the Pareto-optimal front, or the Pareto set, $\mathcal{P}$; solutions in the Pareto set represent the possible optimal trade-offs between competing objectives. A human operator can select a solution once this set has been revealed. When using heuristic algorithms, the non-dominated set produced by one, or several, runs will only be an approximation to the true Pareto front, thus some care with terminology is required, and in this paper the set produced by such an algorithm is referred to as the archive of the estimated Pareto front, which we denote by F.

B. Simulated Annealing

Simulated annealing is the computational analogue of slowly cooling a metal so that it adopts a low-energy, crystalline state. At high temperatures particles are free to move around, whereas as the temperature is lowered they are increasingly confined due to the high energy cost of movement. It is physically appealing to call the function to be minimised the *energy*, $E(\mathbf{x})$, of the state $\mathbf{x}$, and to introduce a parameter T, the computational temperature which is lowered throughout the simulation according to an annealing schedule. At each T the SA algorithm aims to draw samples from the equilibrium distribution $\pi_T(\mathbf{x}) \propto \exp\{-E(\mathbf{x})/T\}$. As $T \to 0$ more and more of the probability mass of π_T is concentrated in the region of the global minimum of E, so that any sample from π_T will almost surely lie at the minimum of E.

Sampling from the equilibrium distribution is usually achieved by Metropolis-Hastings sampling, which involves making proposals $\mathbf{x}'$ that are accepted with probability

$$A = \min\left(1, \exp\{-\delta E(\mathbf{x}', \mathbf{x})/T\}\right) \tag{5}$$

where

$$\delta E(\mathbf{x}', \mathbf{x}) \equiv E(\mathbf{x}') - E(\mathbf{x}) \tag{6}$$

Intuitively, when T is high perturbations from $\mathbf{x}$ to $\mathbf{x}'$ which increase the energy are likely to be accepted (in addition to perturbations which decrease the energy, which are always accepted) and the samples can explore the state space, but as T is reduced only perturbations leading to small increases in E are accepted, so that only limited exploration is possible as the system settles on (hopefully) the global minimum. The algorithm is summarised in Algorithm 1: during each of the K epochs, the computational temperature is fixed at T_k and L_k samples are drawn from π_{T_k} before the temperature is lowered in the next epoch. Each sample is a perturbation (mutation) of the current state from a proposal density (line 3); the perturbed state $\mathbf{x}'$ is accepted with probability given by (5), as shown in lines 4-8. In the work reported here we perturb each element of $\mathbf{x}$ singly, drawing the perturbations from a Laplacian distribution, which has tails that decay relatively slowly, thus ensuring that there is a high probability of exploring regions distant from the current solutions.

Algorithm 1 Simulated annealing

Inputs:
$\{L_k\}_{k=1}^{K}$ *Sequence of epoch durations*
$\{T_k\}_{k=1}^{K}$ *Sequence temperatures,* $T_{k+1} < T_k$
$\mathbf{x}$ *Initial feasible solution*

1: for $k := 1, \ldots, K$
2: for $i := 1, \ldots, L_k$
3: $\mathbf{x}' := \mathit{perturb}(\mathbf{x})$
4: $\delta E := E(\mathbf{x}') - E(\mathbf{x})$
5: $u := \mathtt{rand}(0, 1)$
6: if $u < \min(1, \exp(-\delta E/T_k))$
7: $\mathbf{x} := \mathbf{x}'$
8: end
9: end
10: end

As already alluded to, convergence is guaranteed if and only if the cooling schedule is sufficiently gradual [2], but experience has shown SA to be a very effective optimisation technique even with relatively rapid cooling schedules.

C. Multi-Objective SA with Composite Objective Functions

An attractive approach to multi-objective simulated annealing (MOSA), adopted by several investigators [10], [11], [6], [7], [8], [9], is to combine the objectives into a weighted sum:

$$E(\mathbf{x}) = \sum_{i=1}^{D} w_i f_i(\mathbf{x}) \tag{7}$$

The composite objective is then used as the energy to be minimised in a scalar SA optimiser. An equivalent alternative [5] is to sum $\log f_i(\mathbf{x})$, and others (e.g., [11], [7] have investigated a number of non-linear and stochastic composite energies.

It is clear that simulated annealing with a composite energy (7) will converge to points on the Pareto optimal front where the objectives have ratios given by w_i^{-1}, if such points exist. However, it is unclear how to choose the weights in advance, indeed, one of the principal advantages of multi-objective optimisation is that the relative importance of the objectives can be decided with the Pareto front on hand. Perhaps more importantly, parts of the front are inaccessible with fixed weights [12]. Recognising this, investigators have proposed a variety of schemes for adapting the w_i during the annealing process to encourage exploration along the front.

It is natural to keep an archive, F, of all the non-dominated solutions found so far, and this archive may be utilised to further exploration by periodically restarting the annealer from a randomly chosen element of F [9].

A proposal $\mathbf{x}'$ in scalar SA is either better or worse than the current state $\mathbf{x}$ depending on the sign of $\delta E(\mathbf{x}', \mathbf{x})$ (we

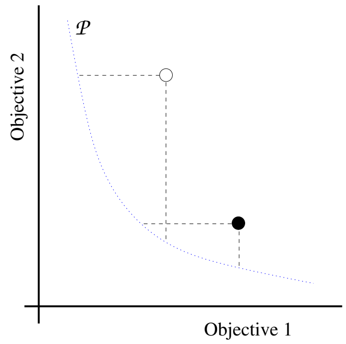

Fig. 1. Energy from area of the true Pareto front $\mathcal{P}$ dominating a solution. Solutions are marked by circles and lines indicate the regions of $\mathcal{P}$ dominating each one.

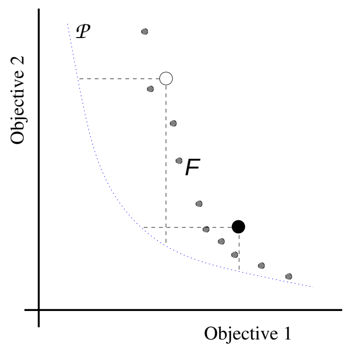

Fig. 2. Energy from proportion of the estimated Pareto front F dominating points dominating a solution. Elements of F are shown as small grey circles, solutions are shown as larger open or filled circles.

ignore the rare possibility that $\delta E = 0$). In multi-objective SA, however, $\mathbf{x}'$ may dominate $\mathbf{x}$ or $\mathbf{x}'$ may be dominated by $\mathbf{x}$ or they may be mutually non-dominating. Energies such as (7) may lead to $\mathbf{x}'$ being unconditionally accepted ($\delta E < 0$) even though $\mathbf{x}' \not\prec \mathbf{x}$, because a large negative energy change from one objective may outweigh small positive changes on the other objectives. Nam & Park [7] therefore modify the acceptance rule so that proposals are accepted with probability given by (5) if they are dominated by $\mathbf{x}$, but unconditionally accepted if $\mathbf{x}' \prec \mathbf{x}$ or if $\mathbf{x}'$ and $\mathbf{x}$ are mutually non-dominating. In a similar vein, Suppapitnarm *et al.* [9] promote exploration along the front by unconditionally accept proposals that are not dominated by any member of F, otherwise using (5).

Although it is clear that the assurance of a convergence proof can be provided for a multi-objective simulated annealer using a scalar objective function and fixed weights (7), such annealers are fundamentally limited in their coverage of the Pareto front. On the other hand, it is difficult to see how proofs of convergence might be obtained with the heuristic modifications designed to promote exploration transversal to the front. With this in mind, we investigate the efficacy of an energy function based on the defining notion of dominance.

III. A Dominance Based Energy Function

In single objective optimisation problems the energy $E(\mathbf{x})$, the quantity to be minimised, is an absolute measure of the quality of any solution $\mathbf{x}$; the optimum is the $\mathbf{x}$ with the lowest energy. However, in the multi-objective case optimum solutions are only defined in relation to each other: the Pareto front is the set of solutions that dominate all other solutions. We can compare the relative quality of $\mathbf{x}$ and $\mathbf{x}'$ with the dominance relation, but note that it gives essentially only three values of quality—better, worse, equal—in contrast to the energy difference in uni-objective problems which usually gives a continuum.

If the true Pareto front $\mathcal{P}$ were available, we could define a simple energy of $\mathbf{x}$ as the measure of the front that dominates $\mathbf{x}$. Let $\mathcal{P}_{\mathbf{x}}$ be the portion of $\mathcal{P}$ that dominates $\mathbf{x}$:

$$\mathcal{P}_{\mathbf{x}} = \{\mathbf{y} \in \mathcal{P} \,|\, \mathbf{y} \prec \mathbf{x}\} \tag{8}$$

Then we define

$$E(\mathbf{x}) = \mu(\mathcal{P}_{\mathbf{x}}) \tag{9}$$

where μ is a measure defined on $\mathcal{P}$. We shall be principally interested in finite sets approximating $\mathcal{P}$ and so shall take $\mu(\mathcal{P}_{\mathbf{x}})$ to simply be the cardinality of $\mathcal{P}_{\mathbf{x}}$. If $\mathcal{P}$ is a continuous set, we can take μ to be the Lebesgue measure (informally, the length, area or volume for 2, 3 or 4 objectives). As illustrated in Fig. 1, this energy has the properties we desire: if $\mathbf{x} \in \mathcal{P}$ then $E(\mathbf{x}) = 0$, and solutions more distant from the front are in general dominated by a greater proportion of $\mathcal{P}$ and so have a higher energy; in Fig. 1 the solution marked by an open circle has a greater energy than the one marked by a filled circle.

Clearly this formulation of an energy does not rely on an *a priori* weighting of the objectives and the assurances of convergence [2] for uni-objective SA continue to hold in this case. Since all solutions lying on the front have equal minimum energy, we would anticipate that a simulated annealer using this energy would, having reached the front, perform a random walk exploration of the front.

We note that Fleischer [13] has proposed an alternative measure of a non-dominated set, which may be loosely characterised as being based on the volume dominated by the set rather than the area of the dominating set.

Unfortunately, the true Pareto front $\mathcal{P}$ is unavailable during the course of an optimisation. We therefore propose to use an energy defined in terms of the current estimate of the Pareto front, F, which is the set of mutually non-dominating solutions found thus far in the annealing. Denote by $\tilde{F}$ the union of the F, the current solution $\mathbf{x}$ and the proposed solution $\mathbf{x}'$, that is $\tilde{F} = F \cup \mathbf{x} \cup \mathbf{x}'$. Then, in a similar manner to (8), let $\tilde{F}_{\mathbf{x}}$ be the elements of $\tilde{F}$ that dominate $\mathbf{x}$:

$$\tilde{F}_{\mathbf{x}} = \{\mathbf{y} \in \tilde{F} \,|\, \mathbf{y} \prec \mathbf{x}\} \tag{10}$$

so that we obtain an energy difference between the current and proposed solutions of

$$\delta E(\mathbf{x}, \mathbf{x}') = \frac{1}{|\tilde{F}|}\left(|\tilde{F}_{\mathbf{x}}| - |\tilde{F}_{\mathbf{x}'}|\right) \tag{11}$$

Division by $|\tilde{F}|$ ensures that δE is always less than 1 and provides some robustness against fluctuations in the number of solutions in F. If $\tilde{F}$ is a non-dominating set the energy difference between any two of its elements is zero. Note also that $\delta E(\mathbf{x}, \mathbf{x}') = -\delta E(\mathbf{x}', \mathbf{x})$.

The inclusion of the current solution and the proposal in $\tilde{F}$ ensures that $\delta E(\mathbf{x}, \mathbf{x}') < 0$ if $\mathbf{x}' \prec \mathbf{x}$, which ensures that proposals that move the estimated front towards the true front are always accepted.

A benefit of this energy measure is that it encourages exploration of sparsely populated regions of the front. Imagine two proposals, each dominated by some solutions in F; for example, the solutions illustrated by the filled and unfilled circles in Fig. 2. The solution that is dominated by fewer elements (the unfilled circle) has the lower energy and would therefore be more likely to be accepted as a proposal.

Defining the energy in this manner thus, unlike some proposed multi-objective enhancements to simulated annealing discussed in section II-C, provides a single energy function encouraging both convergence to and coverage of the Pareto front without requiring other modifications to the single-objective simulated annealing algorithm (beyond the obvious storage of an archive of the estimated Pareto front). In particular no additional rules are required for cases in which the current and proposed solutions are mutually non-dominating.

Convergence to the true Pareto front is no longer an immediate consequence of Geman & Geman's work [2], because the energy based on F is only an approximation to (9). However, Greening [14] offers proof of convergence, albeit more slowly, even when the energy contains errors. Current work is investigating the application of this work to MOSA and in section VI we offer empirical evidence of the convergence.

An energy function based on (11) is straightforward to calculate; counting the number of elements of $\tilde{F}$ that dominate $\mathbf{x}$ and $\mathbf{x}'$ can be achieved in logarithmic time [15]. Our proposed multi-objective algorithm closely follows the standard SA algorithm (Algorithm 1), the only addition that is necessary is to maintain an archive, F of the current estimate of the Pareto front and to calculate the energy difference using (11).

IV. Increasing Energy Resolution

As mentioned earlier, the true Pareto-optimal front of solutions is, in general, unavailable to us. While using the archive of the estimated Pareto front F provides an estimate of solution energy, when F is small the resolution in the energies can be very coarse. In fact, the difference in energy between two solutions is an integer multiple of $1/|\tilde{F}|$ between 0 and 1. Since the acceptance criterion (5) for new solutions is determined by the difference in energy $\delta E(\mathbf{x}, \mathbf{x}')$ between the current solution and the proposed solution, low resolution of the energies leads to a low resolution in acceptance probabilities. At low computational temperatures it will become increasingly likely that, for small archives, this granularity will make it almost impossible for even slightly detrimental moves to be made. This is undesirable as, at its most severe, this effect will reduce the algorithm to behaviour similar to a greedy search optimiser.

This problem is alleviated by using a larger set for energy calculations. There are a couple of straightforward, but ultimately inadequate, methods for artificially increasing the size of F which we now briefly discuss before describing a method using the attainment surface.

A. Conditional Removal of Dominated Points

An obvious method for increasing the size of the archive is to not delete solutions known to be dominated if deleting them would reduce $|F|$ size below some predefined minimum. The existence of old solutions in F, may lead to desirable proposals (not non-dominated solutions) being rejected. In addition the old solutions may bias the search away from regions of the front that were previously well populated.

A further disadvantage of this method is that no solutions are introduced which have not been generated by the algorithm, so dominated solutions kept in F will not necessarily be of any benefit. In the worst case they could be sufficiently far from the current estimation of the Pareto front that they are dominated by all the non-dominated solutions, and will therefore increase the energy resolution at the expense of range.

B. Linear Interpolation

Another apparently suitable method of augmenting F is linear interpolation between the solutions in F. In this method when the archive is smaller than some predefined size new points in objective space are generated on the simplices defined by an element of F and its $D-1$ nearest neighbours. This overcomes the limitations of the previous method: Since new solutions are generated 'on' the current estimated Pareto front, the problems which could occur with using old, dominated elements of F in the energy calculations are avoided. The interpolated points generated can also be evenly distributed between the current estimated Pareto-optimal solutions, which is beneficial as it does not deter the algorithm from exploring any region of the estimated front which is not already densely populated. The principal disadvantage of this method is that proposals may be dominated by an interpolated point, but not by any of the real elements of F, meaning that the proposal may erroneously be disregarded.

C. Attainment Surface Sampling

Consideration of the previous two methods of augmenting the estimated Pareto front suggests that the augmenting points should have the following properties.

1) The artificial points must be sufficiently close to the current estimation of the Pareto front that they can affect

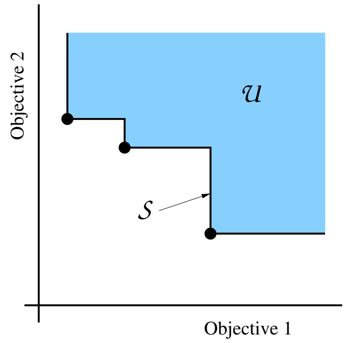

Fig. 3. Attainment surface $\mathcal{S}_F$ is the boundary of the region, $\mathcal{U}$ dominated by the non-dominated set F, whose elements are marked as dots.

the energy of solutions generated near to the current estimated Pareto front.

2) They must be evenly distributed across the currently estimated Pareto front so as to not discourage the algorithm from accepting proposals in poorly populated regions of the front.
3) They must not dominate any proposal which is not dominated by any member of F, so that potential entrants to the archive are not discarded. A consequence of this is that they must all be dominated by at least one member of F.

The *attainment surface*, which has previously been used for estimated Pareto front visualisation [16] and is closely related to the attainment function [17], is an interpolating surface between the elements of F that has the requisite properties. The attainment surface, $\mathcal{S}_F$, corresponding to F is a conservative interpolation of the elements of F so that every point of $\mathcal{S}_F$ is dominated by an element of F; the attainment surface for an F comprising three two-dimensional elements is sketched in Fig. 3. More formally, the attainment surface is the boundary of the region in objective space which is dominated by elements of F. If $\mathbf{u}, \mathbf{v} \in \mathbb{R}^D$, we say that $\mathbf{u}$ *properly dominates* $\mathbf{v}$ (denoted $\mathbf{u} \lhd \mathbf{v}$) iff $u_i < v_i \ \forall i = 1, \ldots, D$. Then if

$$\mathcal{F} = \{\mathbf{y} \,|\, \mathbf{u} \prec \mathbf{y} \text{ for some } \mathbf{u} \in F\} \tag{12}$$

$$\mathcal{U} = \{\mathbf{y} \,|\, \mathbf{u} \lhd \mathbf{y} \text{ for some } \mathbf{u} \in F\} \tag{13}$$

the attainment surface is $\mathcal{S}_F = \mathcal{F} \setminus \mathcal{U}$.

We use random samples, uniformly distributed on $\mathcal{S}_F$ to interpolate F. From the definition of $\mathcal{S}_F$ it is apparent that interpolated points are dominated by an element of F, thus satisfying the third criterion. Uniform random sampling ensures that the second criterion is met, as is the first criterion because $\mathcal{S}_F$ interpolates F.

Sampling from $\mathcal{S}_F$ may be performed using Algorithm 2, which works by sampling a point from a uniform distribution on the axis-parallel hyper-rectangle bounding F and then restricting the one coordinate so that the point is dominated by an element of F. Determining whether an element of F dominates $\mathbf{v}$ on line 8 may be efficiently implemented using a binary searches of the lists L_i, in which case the algorithm requires $\mathcal{O}(|F| \log(|F|))$ time for the generation of each sample. Fig. 4 illustrates the sampled attainment surface for a set of ten 3-dimensional points; 10000 samples are shown for visualisation. In the experiments reported in section VI F was augmented with 100 samples from $\mathcal{S}_F$ before calculating the energy of the proposal.

Algorithm 2 Sampling a point from the attainment surface

Inputs:
$\{L_i\}_{i=1}^D$ *Elements of F, sorted by increasing coordinate*

Generate a random point, $\mathbf{v}$
1: for $i := 1, \ldots, D$
2: $\quad v_i := \texttt{rand}(\min(L_i), \max(L_i))$
3: end
4: $d := \texttt{randint}(1, D)$
Find smallest v_d s.t. $\mathbf{v}$ is dominated by a $\mathbf{y} \in F$
5: for $i = 1, \ldots, |F|$
6: $\quad \mathbf{u} = L_{d,i}$
7: $\quad v_d := u_d$
8: $\quad$ if $F \prec \mathbf{v}$
9: $\qquad$ return $\mathbf{v}$
10: $\quad$ end
11: end

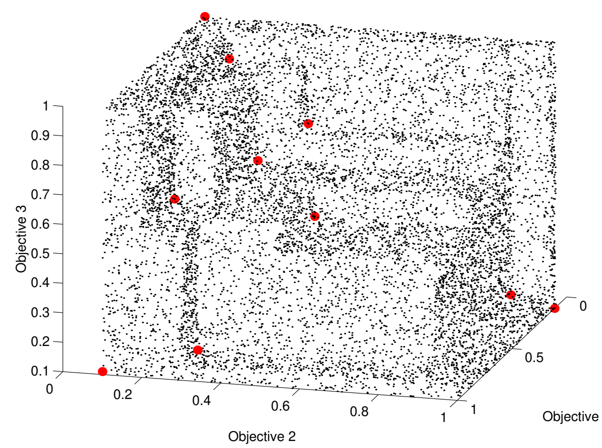

Fig. 4. 10000 samples from the attainment surface for an archive of 10 points, which are marked with heavy dots.

V. Realtime Algorithm Parameter Optimisation

The performance of this algorithm, in common with other simulated annealing systems, depends upon parameters for the initial temperature, the annealing schedule and the size of perturbations made to solutions when generating new proposals. Here we give details of methods which permit automatic setting of the initial temperate, and which adjust the scale of perturbations made to maximise the quality of proposed solutions.

A. Initial Temperatures

If this initial temperature of the system is set too high, all proposed solutions will be accepted, irrespective of their relative energies, and if set too low proposals with a higher energy than the current solution will never be accepted, transforming the algorithm into a greedy search. As a reasonable starting point we set the initial temperature to achieve an initial acceptance rate of approximately 50% on derogatory proposals. This initial temperature, T_0, can be easily calculated by using a short 'burn-in' period during which time all solutions are accepted and setting the temperature equal to the average positive change of energy divided by $\ln(2)$.

Here we adjust the temperate according to $T_k = \beta^k T_0$, where β is chosen so that the final temperature is 10^{-5}.

B. Perturbation Scalings

For simplicity a proposal is generated from $\mathbf{x}$ by perturbing only one parameter or decision variable of $\mathbf{x}$. The parameter to be perturbed is chosen at random and perturbed with a random variable drawn from a Laplacian distribution, $p(\epsilon) \propto e^{-|\sigma\epsilon|}$, where the scaling factor σ sets magnitude of the perturbation.

We maintain two sets of scaling factors, since the perturbations generating moves to a non-dominated proposal within a front (traversals) may potentially be very different from those required to locate a front closer to $\mathcal{P}$. We maintain a scaling factor for each dimension of parameter space for each of location perturbations and traversal perturbations, and adjust these independently to give the greatest probability of such a move being generated. When perturbing a solution, it is chosen randomly with equal probability whether the location scaling set will be used, or the traversal scaling set. These scalings are initially set large enough to sample from the entire feasible space. The scalings are adjusted throughout the optimisation, whenever a suitably large statistic set is available to reliably calculate an appropriate scaling factor.

1) Traversal Scaling: The traversal rescaling for a particular decision variable, x_j is performed whenever approximately 50 traversal perturbations have been made to x_j since the last rescaling.

In order to ensure wide coverage of the front we wish to to maximise the distance (in objective space) covered by the traversals to ensure the entire front is evenly covered. Generating traversals travelling a small distance will concentrate the estimated front around the point at which the current front was discovered, an effect we aim to avoid.

We seek to generate proposals on approximately the scale that has previously been successful in generating wide-ranging traversals. To achieve this, the perturbations are sorted by absolute size of perturbation in parameter space, and then trisected in order, giving three groups, one of the smallest third of perturbations, the largest third of perturbations, and the remaining perturbations. For each group the mean traversal size caused by the perturbations is calculated. The traversal size is measured as the Euclidean distance travelled in objective space when the current solution and the proposed solution are mutually non-dominating. The traversal perturbation scaling for this decision variable is then set to the average perturbation of the group which generated the largest average traversal.

This heuristic is open to the criticism that it depends upon measuring distances in objective space whose relative weighting is unknown; however, the objectives may be renormalised during optimisation so the front has approximately the same extent in objective. We emphasise that, of course, this does not affect the dominance-based energy.

2) Location Scaling: Drawing from methods widely used in evolutionary algorithms, we aim to adjust the scale of location perturbations to keep the acceptance rate for $\mathbf{x}'$ that have a higher energy than $\mathbf{x}$ to approximately $1/3$, so that exploratory proposals are made and accepted at all temperatures.

The location perturbation scaling is recalculated for each parameter for which 20 proposals for which $\delta E(\mathbf{x}, \mathbf{x}') > 0$ have been generated, after which the count is reset. Location perturbation rescaling is omitted in two cases: First, when the archive of the estimated Pareto front F has fewer than 10 members. Secondly, when the combined size of F augmented by the samples from the attainment surface when multiplied by the temperature does not exceed 1. This is because we adjust the scalings to attempt to keep the acceptance rate of derogatory moves approximately a third; when this value is too small, it becomes impossible to generate such a scaling, and so the scalings are kept at the most recent valid value.

Counting only moves generated from perturbations to a particular dimension of parameter space, the acceptance rate of derogatory moves α is the fraction of proposals to a greater energy which are accepted. If σ denotes the location perturbation scaling for a particular dimension, the new σ is set as:

$$\sigma := \begin{cases} \sigma(1 + 2(\alpha - 0.4)/0.6) & \text{if } \alpha > 0.4 \\ \sigma/(1 + 2(0.3 - \alpha)/0.3) & \text{if } \alpha < 0.3 \end{cases} \tag{14}$$

This update works because, in general, smaller perturbations in parameter space are more likely to generate small changes in objective space, resulting in smaller changes in energy.

VI. Experiments

We illustrate the performance of this annealer on some well-known test functions from the literature, namely the first three DTLZ test functions of Deb *et al.* [18]. The benefit of using these functions is that the true Pareto front, $\mathcal{P}$, is known, so we can discover how close our estimated archive F is to $\mathcal{P}$, as well as comparing our results to published work.

A. Problem description

Test problems DTLZ1, DTLZ2 and DTLZ3 are minimisation problems, and are described used here in their standard 3-objective formulation. The true front of DTLZ1 lies on the plane where the sum of all objectives equals 0.5. The fronts of DTLZ2 and DTLZ3 in contrast lie on the spherical front where the sum of the squared objectives equals 1.0.

In order to quantify the convergence of the annealer we measure two distinct properties. Firstly the median distance of the archived solutions discovered at each fold from the true

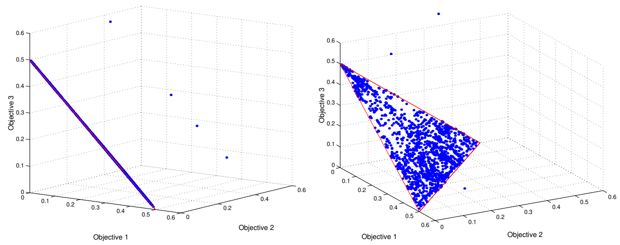

Fig. 5. The annealer archive on test problem DTLZ1 after 30000 iterations.

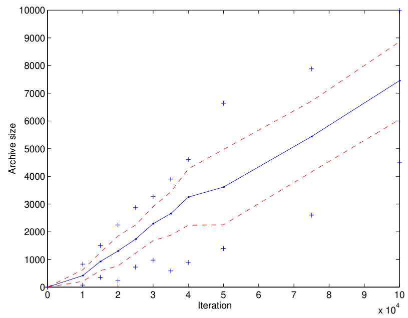

Fig. 6. Annealer archive growth on test problem DTLZ1. Mean archive size over 30 runs plotted against iteration. Dashed lines represent the mean ± 1 standard deviation; maximum and minimum sizes are marked with crosses.

front is calculated to ascertain how close on average solutions found are to the true front. Additionally we are concerned with finding solutions spread across the true Pareto front, in order to do this we also use a variant of the volume $\mathcal{V}$ measure [15]. Here $\mathcal{V}(\mathcal{P}, F)$ measures the percentage of a given objective volume which is dominated by the true front $\mathcal{P}$, but not the stored archive F. Implementing the measure in this study involves Monte Carlo sampling (10^5 samples here) of a cube bounded by the range [0,2] on each coordinate and measuring the proportion of points which are dominated by the true Pareto front, but not dominated by the annealer archive.The better covered the true Pareto front is, the smaller this measure becomes.

The annealer was run 30 times, for iteration lengths (function evaluations) 10000, 15000, 20000, 25000 and 30000.

B. Results

Figure 5 shows the archive attained from a single run of the annealer on test problem DTLZ1 for 30000 iterations. This is an equivalent number of function evaluations as the non-dominated sorting genetic algorithm II (NSGA-II) [19], used on this problem in [18]. Although both algorithms demonstrate a similar level of convergence, the annealer generates a far higher *number* of converged solutions in comparison with those reported in [19]. Figure 6 shows how the final archive size grows with run length for DTLZ1, an average of over 2000 non-dominating points found after 30000 iterations, the majority of which have converged to the true front.

A even more striking example of the annealers performance is provided in Figure 7, where its evaluation on DTLZ3 is

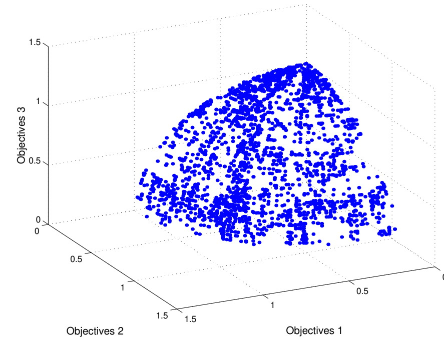

Fig. 7. The annealer archive on test problem DTLZ3 after 30000 iterations.

shown again for 30000 iterations. The front is converged to within 0.002 of the true front (as supported by the histogram of solution distances from the front in Figure 8). Deb *et al.* [19] in contrast found that the NSGA-II had still failed to converge to this problem after 50000 function evaluations.

The experimental results, averaged across runs, are provided in Table I. This table shows the median and quartile boundaries across the folds for the two error measures for each iteration. The annealer can be seen to converge very swiftly to DTLZ2, in just 10000 iterations the average distance of archive set members from the true Pareto front being less than 10^{-5}. The annealer does not converge (on average) quite so close on the other two test functions, which is probably due to some members of the estimated Pareto set being caught on one of the many local fronts that lie close to the true front for these problems. Nevertheless, by 30000 iterations the average distance of solutions from the true Pareto front is of the order 10^{-4} and 10^{-3} for DTLZ1 and DTLZ3 respectively. The $\mathcal{V}$ results evince that not only have the solutions found by the annealer converged well to the true front, but also that they have spread across it well (as also shown in Figures 5 and 7). After 30000 iterations, on average, only 0.15% of the volume in Z is dominated by $\mathcal{P}$ for DTLZ1, but not by the annealer's estimate of it. For DTLZ2 and DTLZ3 this value is the slightly higher 0.32% and 1.17% (however it must be noted that the actual volume sampled in Z for these test functions is smaller due to the shape of $\mathcal{P}$ — so relatively larger $\mathcal{V}$ values compared to DTLZ1 are unsurprising).

VII. Conclusions

We have presented an energy measure for use in multi-objective SA which is based on the fundamental notion of dominance, rather than employing a weighted combination of the objectives. Simulated annealers employing this measure were shown to have good convergence properties on the first three DTLZ test functions [18]. The simulated annealer with the dominance measure even provided an excellent estimate of the Pareto front on the DTLZ3 problem, where the multi-objective genetic algorithm NSGA-II was unable to locate the true front, even with more function evaluations.

An advantage of the dominance based energy measure is that it is not *a priori* biased towards any part of the front;

TABLE I

TEST PROBLEM RESULTS. MEDIAN AND QUARTILE BOUNDARIES FOR THE MEDIAN DISTANCE OF POINTS FROM THE TRUE FRONT, AND FOR $\mathcal{V}(\mathcal{P}, F)$.

	DTLZ1	DTLZ2	DTLZ3	DTLZ1	DTLZ2	DTLZ3
Iterations	Distance $\times 10^{-2}$	Distance $\times 10^{-6}$	Distance $\times 10^{-2}$	$\mathcal{V}$ %	$\mathcal{V}$ %	$\mathcal{V}$ %
10000	6.27 (1.68, 50.18)	5.63 (2.73, 13.77)	19.15 (7.04 ,52.26)	0.59 (0.32, 0.99)	0.66 (0.62, 0.71)	5.45 (2.60, 10.18)
15000	0.95 (0.52, 2.36)	1.89 (0.74, 4.75)	3.84 (1.52, 21.79)	0.26 (0.18, 0.35)	0.47 (0.45, 0.54)	2.44 (1.53, 6.78)
20000	0.53 (0.167, 15.93)	0.87 (0.54, 4.55)	0.96 (0.43, 2.98)	0.21 (0.14, 0.30)	0.42 (0.36, 0.45)	1.50 (1.03, 3.36)
25000	0.08 (0.06, 0.169)	0.51 (0.12, 1.190)	0.43 (0.21, 0.981)	0.17 (0.14, 0.239)	0.36 (0.33, 0.395)	1.84 (0.84, 6.512)
30000	0.05 (0.02, 0.13)	0.29 (0.15, 1.24)	0.23 (0.13, 0.54)	0.15 (0.10, 0.29)	0.32 (0.31, 0.35)	1.17 (0.65, 3.04)

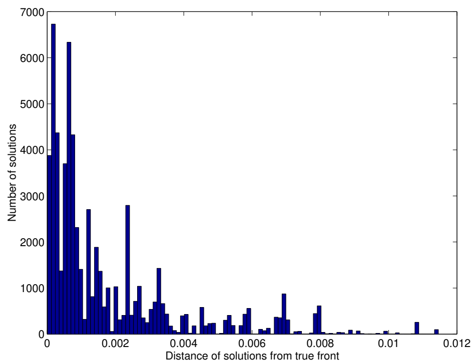

Fig. 8. Histogram of distances of archive solutions from the true front on problem DTLZ3 after 30000 iterations, for all the 30 folds. The last 5% of the right hand side of the histogram has been truncated for ease of viewing.

indeed, we have argued that it tends to promote exploration in sparsely populated regions and in practice we have shown that estimated fronts evenly and widely cover the true front.

Determining an efficient scale on which to make proposals is more complicated in the multi-objective case than the uni-objective case, because some proposals work to advance the front, while others traverse the front. We have proposed simple heuristics to adapt the perturbation scales and future work involves applying machine learning techniques to learn the local mapping between parameter and objective space in order to more sensitively control the search direction.

Our $E(\mathbf{x})$ is a measure of a portion of the dominated set, namely $\mu(\tilde{F}_{\mathbf{x}})$; which is a close relation to Fleischer's recently proposed measure [13]; loosely, our measure deals with the area of the dominating surface—the attainment surface—while Fleischer's considers the dominated volume. It will be interesting to investigate the convergence of an annealer based on Fleischer's measure. Although a proof of convergence for simulated annealers based on our measure remains to be completed, this is an area of current work, together with the application of the annealer to large scale problems.

ACKNOWLEDGEMENTS

We gratefully acknowledge funding from Motorola [KIS] and the EPSRC [KIS & JEF].

REFERENCES

[1] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, pp. 671–680, 1983.

[2] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 6, pp. 721–741, 1984.

[3] C.A.C Coello, "A Comprehensive Survey of Evolutionary-Based Multiobjective Optimization Techniques," *Knowledge and Information Systems. An International Journal*, vol. 1, no. 3, pp. 269–308, 1999.

[4] D. Van Veldhuizen and G. Lamont, "Multiobjective Evolutionary Algorithms: Analyzing the State-of-the-Art," *Evolutionary Computation*, vol. 8, no. 2, pp. 125–147, 2000.

[5] P. Engrand, "A multi-objective approach based on simulated annealing and its application to nuclear fuel management," in *5th International Conference on Nuclear Engineering*, Nice, France, 1997, pp. 416–423.

[6] P. Czyżak and A. Jaszkiewicz, "Pareto simulated annealing – a metaheuristic technique for multiple-objective combinatorial optimization," *Journal of Multi-Criteria Decision Analysis*, vol. 7, pp. 34–47, 1998.

[7] D.K. Nam and C.H. Park, "Multiobjective simulated annealing: a comparative study to evolutionary algorithms," *Int. J. Fuzzy Systems*, vol. 2, no. 2, pp. 87–97, 2000.

[8] M. Hapke, A. Jaszkiewicz, and R. Slowinski, "Pareto simulated annealing for fuzzy multi-objective combinatorial optimization," *Journal of Heuristics*, vol. 6, no. 3, pp. 329–345, 2000.

[9] A. Suppapitnarm, K.A. Seffen, G.T. Parks, and P.J. Clarkson, "A simulated annealing algorithm for multiobjective optimization," *Engineering Optimization*, vol. 33, pp. 59–85, 2000.

[10] P. Serafini, "Simulated annealing for multiobjective optimization problems.," in *Multiple criteria decision making. Expand and enrich the domains of thinking and application*, 1994, pp. 283–292.

[11] E.L. Ulungu, J. Teghaem, Ph. Fortemps, and D. Tuyttens, "Mosa method: a tool for solving multiobjective combinatorial decision problems," *J. multi-criteria decision analysis*, vol. 8, pp. 221–236, 1999.

[12] I. Das and J. Dennis, "A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multicriteria optimization problems," *Structural Optimization*, vol. 14, no. 1, pp. 63–69, 1997.

[13] M. Fleischer, "The measure of Pareto optima: Applications to multi-objective metaheuristics," Tech. Rep. CSHCN TR 2002-17, CSHCN, University of Maryland, 2002.

[14] D. Greening, "Simulated annealing with inaccurate costs functions," in *Proceedings of the IMACS International Congress of Mathematics and Computer Science*, Trinity College, Dublin, 1993.

[15] J.E. Fieldsend, R.M. Everson, and S. Singh, "Using Unconstrained Elite Archives for Multi-Objective Optimisation," *IEEE Transactions on Evolutionary Computation*, vol. 7, no. 3, pp. 305–323, 2003.

[16] E. Zitzler, *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*, Ph.D. thesis, Swiss Federal Institute of Technology Zurich (ETH), 1999, Diss ETH No. 13398.

[17] V. Grunet da Fonseca, C.M. Fonseca, and A.O. Hall, "Inferential Performance Assessment of Stochastic Optimisers and the Attainment Function," in *First International Conference on Evolutionary Multi-Criterion Optimization*, E. Zitzler, K. Deb, L. Thiele, C.A.C. Coello, and D. Corne, Eds., pp. 213–225. Springer-Verlag. Lecture Notes in Computer Science No. 1993, 2001.

[18] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, "Scalable multi-objective optimization test problems," in *Congress on Evolutionary Computation (CEC'2002)*, 2002, vol. 1, pp. 825–830.

[19] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II," in *Proceedings of Parallel Problem Solving from Nature - PPSN VI*. 2000, pp. 849–858, Springer.