

# Efficiently Identifying Pareto Solutions when Objective Values Change

Jonathan E. Fieldsend<sup>\*</sup>  
Computer Science  
University of Exeter, UK  
J.E.Fieldsend@exeter.ac.uk

Richard M. Everson  
Computer Science  
University of Exeter, UK  
R.M.Everson@exeter.ac.uk

## ABSTRACT

In many multi-objective problems the objective values assigned to a particular design can change during the course of an optimisation. This may be due to dynamic changes in the problem itself, or updates to estimated objectives in noisy problems. In these situations, designs which are non-dominated at one time step may become dominated later not just because a new and better solution has been found, but because the existing solution's performance has degraded. Likewise, a dominated solution may later be identified as non-dominated because its objectives have comparatively improved. We propose management algorithms based on recording single "guardian dominators" for each solution which allow rapid discovery and updating of the non-dominated subset of solutions evaluated by an optimiser. We examine the computational complexity of our proposed approach, and compare the performance of different ways of selecting the guardian dominators.

## Categories and Subject Descriptors

E.1 [Data]: Data Structures—*Graphs and networks*; G.1.6 [Mathematics of Computing]: Numerical Analysis—*Optimization, Global optimization*

## General Terms

Algorithms, Experimentation, Performance, Theory

## Keywords

Multi-objective optimisation, dynamic problems, uncertainty

## 1. INTRODUCTION

A typical assumption when performing an optimisation is that the evaluation of a solution does not vary unless the solution itself is modified. A number of data structures

<sup>\*</sup>Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.  
Copyright 2014 ACM 978-1-4503-2662-9/14/07 ...\$15.00.  
<http://dx.doi.org/10.1145/2576768.2598279>.

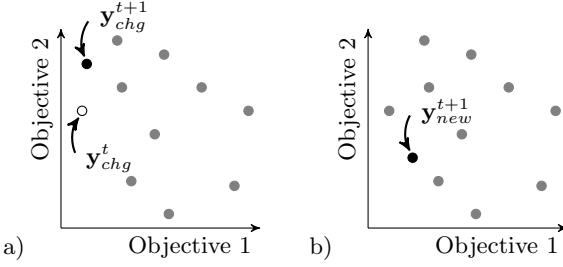
exist for efficiently maintaining and querying sets of multi-objective solutions in this situation (see [1] for a recent review). In the worst case, the number of domination comparisons required before a solution can be classified as non-dominated is equal to the size of the current non-dominated set. However, this assumes that a domination relationship between two solutions at time  $t$  will persist at all future time steps. If this does not hold it becomes necessary to track dominated as well as non-dominated solutions as the search progresses if we want to guarantee access at a particular time step to the best estimate of the non-dominated subset of solutions visited by an optimiser. How to efficiently maintain an elite archive when the assigned objectives are susceptible to change has not been widely addressed in the literature. We present here data structures and algorithms to facilitate the efficient identification of the non-dominated subset of solutions, when the *assigned* objectives to previously evaluated solutions may vary over time. This work is based on our recent technical report [5].

Assigned objectives may vary in a number of circumstances. In dynamic optimisation problems the problem itself changes over time, meaning that the objective vectors of all solutions evaluated at any time may vary if reassessed at a later time [10]. In noisy problems repeated evaluation of the same solution leads to different values for the objectives [8, 7]. Finally, an objective evaluation may be updated or refined at some later point due to the receipt of new information [4]. In all these cases, the objective vector  $\mathbf{y}$  associated with the design vector  $\mathbf{x}$  at some time  $t$ , may not be the same as the objective vector associated with  $\mathbf{x}$  at  $t + 1$ .

We now briefly review the ideas of dominance and Pareto optimality in Section 2 before discussing in Section 3 the computational complexity of managing a set where at each time step a previously evaluated solution may have its objectives changed, or a brand new solution may be evaluated. We introduce a data structure and algorithms to efficiently maintain the set of evaluated solutions and identify the non-dominated subset in Section 4. In Section 5 we provide empirical assessments of different algorithms. The paper concludes with a discussion in Section 6.

## 2. MULTI-OBJECTIVE OPTIMISATION

In multi-objective optimisation problems the optimiser seeks to simultaneously extremise  $D$  objectives,  $f_d(\mathbf{x}), d = 1, \dots, D$ , where each objective depends upon an  $L$ -dimensional solution vector  $\mathbf{x}$  of parameters or design variables. The parameters may also be subject to equality and inequality constraints which, for simplicity, we assume can be eval-



**Figure 1:** Two possible transitions from  $Y^t$  to  $Y^{t+1}$ .  
a: a single member of  $Y^t$  has had its value shifted from  $y_{chg}^t$  to  $y_{chg}^{t+1}$ . b: a new solution has been evaluated with objective vector  $y_{new}^{t+1}$ .

ated precisely. When the objectives are to be minimised the multi-objective optimisation problem may be expressed as:  $\text{minimise } \mathbf{y} = \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_D(\mathbf{x}))$  with  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^L$ . The feasible domain  $\mathcal{X}$  is defined by any design constraints. When there is more than one objective to be minimised, solutions may exist for which performance on one objective cannot be improved without reducing performance on at least one other. Such solutions are *Pareto optimal*.

The notion of *dominance* may be used to make Pareto optimality clearer. Assuming, without loss of generality, that the goal is to minimise the objectives, and given two *objective vectors*,  $\mathbf{v}$  and  $\mathbf{u}$ , then  $\mathbf{u}$  is said to *dominate*  $\mathbf{v}$  iff  $\mathbf{u}_d \leq \mathbf{v}_d, \forall d = 1, \dots, D$  and  $\mathbf{u} \neq \mathbf{v}$ . This is denoted as  $\mathbf{u} \prec \mathbf{v}$ .<sup>1</sup> Furthermore, we denote by  $U \prec \mathbf{v}$  the situation where at least one member of the set  $U$  dominates  $\mathbf{v}$ . A set of objective vectors is said to be a *mutually non-dominating set* if no member of the set is dominated by any other member. A solution to a multi-objective optimisation problem is thus *Pareto optimal* if it is not dominated by any other feasible solution. The set of all Pareto optimal solutions is the Pareto set, and its image in objective space is the Pareto front. We denote the estimated Pareto front (the non-dominated subset of objective vectors associated with the solutions evaluated by an optimiser) by  $E$ .

Here we address problems in which the objectives associated with a solution may *vary* from one time step to the next. As such we effectively have a solution whose associated objective vector  $\mathbf{y}$  is dependent on the time  $t$  at which it is evaluated; we indicate this with  $\mathbf{y}^t$ . We denote by  $Y^t$  the set of all objective vectors that have been evaluated at time  $t$ . An illustration is provided in Figure 1. In the left panel a single member  $\mathbf{y}_{chg}^t$  of  $Y^t$  has a new objective value at time  $t+1$ , namely  $\mathbf{y}_{chg}^{t+1}$ . This new objective value is associated with the same solution as  $\mathbf{y}_{chg}^t$ . Also in this illustration all other  $\mathbf{y}_i \in Y^t, i \neq chg$  are unaltered so that  $\mathbf{y}_i^t = \mathbf{y}_i^{t+1}$ , and  $|Y^t| = |Y^{t+1}|$ . In the right panel a new solution has been evaluated, resulting in  $\mathbf{y}_{new}^{t+1}$ , meaning that  $|Y^{t+1}| = |Y^t| + 1$ .

### 3. UPDATING ON A SINGLE TIME STEP

Consider a general set  $Y^t$ . At time  $t+1$  either a single solution has had its associated objective vector in  $Y^t$  changed

<sup>1</sup>As we are only concerned with the objective space representation of solutions here, we restrict our dominance notation to this space.

or a new solution has been evaluated for the first time, which is added to  $Y^t$  to form  $Y^{t+1}$ . In a situation where more than one member of  $Y^t$  has its value changed or there is more than a single new solution added to  $Y^t$ , each individual modification to the set may be viewed as a distinct time step and the order these are processed is immaterial to the final states of the solutions and the elite set.

The non dominated members  $E^t$  of  $Y^t$  may be selected via the `nondom` function defined below:

$$A = \text{nondom}(Z) = \{\mathbf{z} \in Z \mid \nexists \mathbf{v} \in Z, \mathbf{v} \prec \mathbf{z}\}. \quad (1)$$

The estimated Pareto front at time  $t$  may therefore be identified as  $E^t = \text{nondom}(Y^t)$ .  $E^t$  is also known as the elite set or archive. It is also useful to define a function which returns the subset of a population whose members dominate an objective vector  $\mathbf{u}$ :

$$B = \text{dom\_members}(\mathbf{u}, Z) = \{\mathbf{z} \in Z \mid \mathbf{z} \prec \mathbf{u}\}. \quad (2)$$

If  $\mathbf{v}^t \in Y^t$  and  $\text{dom\_members}(\mathbf{v}^t, Y^t) = \emptyset$ , then  $\mathbf{v}^t \in E^t$ . Conversely if  $\text{dom\_members}(\mathbf{v}^t, Y^t) \neq \emptyset$ , then  $\mathbf{v}^t \notin E^t$ .  $E^t$  may thus be alternatively defined as:

$$E^t = \{\mathbf{y}^t \in Y^t \mid \text{dom\_members}(\mathbf{y}^t, Y^t) = \emptyset\}. \quad (3)$$

This may seem a slightly convoluted route to defining the non-dominated set, however it is a useful formulation for determining which solutions need to be compared to  $E^t$  for possible entry into  $E^{t+1}$

#### 3.1 Changing a dominated member of $Y^t$

Consider the effect of changing the objective vector of a dominated member of  $Y^t$ ,  $\mathbf{y}_{chg}^t$ . If  $\mathbf{y}_{chg}^t \notin E^t$ , then the change to  $\mathbf{y}_{chg}^{t+1}$  (from  $\mathbf{y}_{chg}^t$ ) means that it should now enter  $E^{t+1}$  if  $\text{dom\_members}(\mathbf{y}_{chg}^{t+1}, E^t) = \emptyset$ . However, there are no other members of  $Y^{t+1}$  that can enter  $E^{t+1}$  as a result of the movement of  $\mathbf{y}_{chg}^t$  to  $\mathbf{y}_{chg}^{t+1}$ , because if they were dominated by members of  $E^t$ , they will still be dominated by the same members at  $t+1$ . That is, because  $\mathbf{y}_{chg}^t \notin E^t$ , then there is at least one  $\mathbf{e}^t \in E^t$  which dominates  $\mathbf{y}_{chg}^t$ . At time  $t+1$  this  $\mathbf{e}^t$  therefore also dominates *all* solutions that  $\mathbf{y}_{chg}^t$  dominated at time  $t$ .

Existing members of  $E^t$  will require removal from  $E^{t+1}$  if they are dominated by  $\mathbf{y}_{chg}^{t+1}$ , so if  $\mathbf{y}_{chg}^{t+1}$  enters  $E^{t+1}$ , then each  $\mathbf{e}_i^t \in E^t$  should be compared to  $\mathbf{y}_{chg}^{t+1}$ , and if  $\mathbf{y}_{chg}^{t+1} \prec \mathbf{e}_i^t$  then  $\mathbf{e}_i^t$  cannot be a member of  $E^{t+1}$ . As long as the domination comparison results are stored when  $\mathbf{y}_{chg}^{t+1}$  is compared to  $E^t$  for entry into  $E^{t+1}$ , this stage will not require any additional domination calculations. Consequently a maximum of  $|E^t|$  domination comparisons are required when the objective vector of a dominated member of  $Y^t$  is changed.

#### 3.2 Changing a non-dominated member of $Y^t$

If prior to varying its objective values  $\mathbf{y}_{chg}^t \in E^t$ , then its change in location may mean it should not be in  $E^{t+1}$ . However, even if it does enter  $E^{t+1}$ , as illustrated in Figure 1a, solutions in  $Y^t$  that it dominated may enter  $E^{t+1}$  because the change in location to  $\mathbf{y}_{chg}^{t+1}$  from  $\mathbf{y}_{chg}^t$  may mean solutions in  $Y^t$  that were previously dominated by  $\mathbf{y}_{chg}^t$  are no longer dominated in  $Y^{t+1}$ .

If  $\text{dom\_members}(\mathbf{y}_{chg}^{t+1}, E^t) = \emptyset$ , then  $\mathbf{y}_{chg}^{t+1}$  will enter  $E^{t+1}$ ; otherwise it will be excluded and  $E^{t+1}$  is initially set as

$E^t \setminus \{\mathbf{y}_{chg}^t\}$ . If  $\mathbf{y}_{chg}^{t+1}$  does enter  $E^{t+1}$ , then all other  $\mathbf{e}_i^t \in E^t$  must be assessed to discover whether they are dominated by  $\mathbf{y}_{chg}^{t+1}$ . If  $\mathbf{y}_{chg}^{t+1} \prec \mathbf{e}_i^t$ , then  $\mathbf{e}_i^t$  will not be a member of  $E^{t+1}$ . The members  $\mathbf{y}_i^t$  of  $Y^t$  which are dominated by  $\mathbf{y}_{chg}^t$  need to be compared to  $\mathbf{y}_{chg}^{t+1}$  only if they are dominated exclusively by  $\mathbf{y}_{chg}^t$  at time  $t$ ; that is  $\text{dom\_members}(\mathbf{y}_i^t, Y^t) = \{\mathbf{y}_{chg}^t\}$ . This is because if a  $\mathbf{y}_i^t$  is dominated by members of  $Y^t$  other than  $\mathbf{y}_{chg}^t$ , then it will still be dominated by them at  $t+1$  since these members are unchanged.

We denote by  $K$  the number of solutions in  $Y^t$  which are dominated *exclusively* by  $\mathbf{y}_{chg}^t$ . In practise  $K$  is typically small (often zero); in order to be exclusively dominated by a single member of  $Y^t$  it is necessary (but not sufficient) for a solution to lie in the level 2 Pareto shell, namely those solutions that are non-dominated in the set  $Y^t \setminus E^t$  (see [2]).

In the worst case  $|E^t| + K - 1$  domination comparisons are required before the membership of  $E^{t+1}$  is determined. That is,  $\mathbf{y}_{chg}^{t+1}$  will, at worst, need to be compared with all members of  $E^t$  bar itself ( $|E^t| - 1$  comparisons) and the  $K$  solutions dominated only by  $\mathbf{y}_{chg}^t$  must be compared to  $\mathbf{y}_{chg}^{t+1}$ , requiring  $K$  domination comparisons. If a solution is not dominated by  $\mathbf{y}_{chg}^{t+1}$  it may enter  $E^{t+1}$  directly without any other domination comparisons required because if  $\mathbf{y}_{chg}^t$  was the *only* dominating member in  $Y^t$ , then by definition no other members of  $E^t$  can dominate it. This comparison may also result in  $\mathbf{y}_{chg}^{t+1}$  being removed from  $E^{t+1}$ .

### 3.3 Evaluating a new location

Finally, a completely new location,  $\mathbf{y}_{new}^{t+1}$ , may be suggested due to evaluation of a new solution. In this case determination of membership of  $E^{t+1}$  is made by domination comparison of this proposal with the members of  $E^t$ . From this we can see the computational complexity of determining the elite set of solutions when adding a new location into  $Y^{t+1}$  is in the worst case  $|E^t|$  domination comparisons.

### 3.4 Multiple sequential updates

The cost of determining the elite set  $E^{t+1}$  of  $Y^{t+1}$  from a single changed member of  $Y^t$ , when all domination relations between members of  $Y^t$  have been recorded, is the sum of two components:  $|E^t|$  domination comparisons plus an additional  $K - 1$  domination comparisons when  $\mathbf{y}_{chg}^t$  is non-dominated. This complexity depends on the dominance relations between elements of  $Y^t$  being stored, which in itself is computationally expensive.

However, even if all the domination relations between members of  $Y^t$  are stored and the  $|E^t|$  or  $|E^t| + K - 1$  domination comparisons necessary to determine the elite set  $E^{t+1}$  when a member of  $Y^t$  is changed are computed, this fails to update all the domination relations between elements of  $Y^{t+1}$ . In particular, the domination relations between elements of  $Y^t$  which are dominated by more than one element of  $Y^t$  are not updated. For instance, the updating scheme described in Section 3.3 relies on knowing *which* members of  $Y^t$  a solution is dominated by. We also need to compare a new solution  $\mathbf{y}_{new}^{t+1}$  against  $Y^t \setminus E^t$  to keep track of this information. Specifically, if  $\mathbf{y}_{new}^{t+1}$  enters  $E^{t+1}$ , then the members of  $Y^t$  that it dominates will need to be identified (requiring at worst  $|Y^{t+1} \setminus E^{t+1}|$  domination comparisons). If  $\mathbf{y}_{new}^{t+1}$  is dominated by  $E_t$ , then we can utilise the observation that  $\mathbf{y}_{new}^{t+1}$  can at most dominate the subset of  $Y^{t+1}$  which is dom-

inated by *every* member of  $\text{dom\_members}(\mathbf{y}_{new}^{t+1}, E^{t+1})$ , but in the worst case this too is also  $Y^{t+1} \setminus E^{t+1}$ .

We think of two solutions being *linked* if one dominates the other. Maintaining all domination links between members of  $Y^t$  as new solutions enter and existing solutions change is exorbitant for any  $|Y^t|$  larger than a modest size. As such, it is impractical to use the approach outlined in Section 3 directly. However, efficient updates are possible by recording the information for just a *single* solution that dominates each member of  $Y^t \setminus E^t$ . The storage requirement of  $\mathcal{O}(|Y^t|)$  is vastly reduced in comparison with the  $\mathcal{O}(|Y^t|^2)$  needed to store all links, and the computational cost of maintaining these single links is also far superior to maintaining all domination links. We will now outline this proposed approach.

## 4. GUARDIAN DOMINATORS

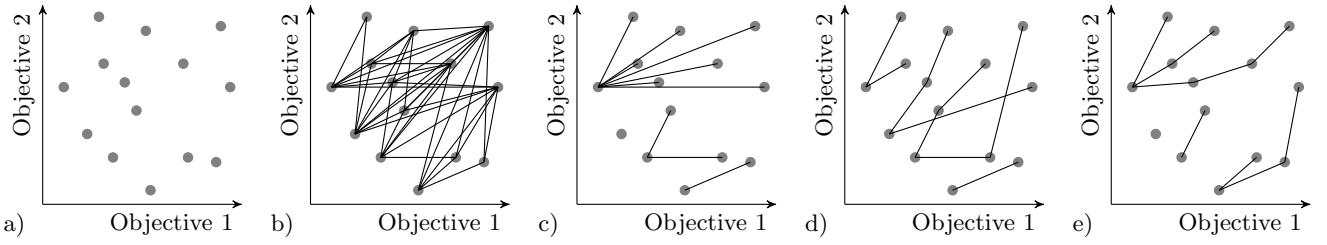
Instead of keeping track of *all* domination relations between members of  $Y^t$ , here we suggest that for each member  $\mathbf{y}_i^t$  of  $Y^t \setminus E^t$  a *single* other solution in  $Y^t$  is recorded which dominates it: its *guardian dominator*. As we show below, tracking these links as search progresses is computationally inexpensive, and permits  $Y^t$  and  $E^t$  to be cheaply updated.

The guardian dominator of the  $i$ th member of  $Y^t$  is labelled  $\mathbf{y}_i^{t,*}$  for convenience (therefore  $\mathbf{y}_i^{t,*} \prec \mathbf{y}_i^t$ ). By definition there is no solution in  $Y^t$  which dominates any members of  $E^t$ , and as such the elements of  $E^t$  are not assigned guardian dominators. Figure 2 illustrates the ways in which these single relationships are maintained for  $D = 2$  objectives. Figure 2b shows the 43 domination relationships between the 14 solutions plotted in Figure 2a; domination links are shown as edges between each dominating and dominated solution. Figure 2c shows the same set, but with only domination links between members of  $Y^t \setminus E^t$  and their guardian dominators, which are selected from  $E^t$ , making 10 edges in total as  $|Y^t \setminus E^t| = 10$ . Figure 2d again shows  $Y^t$ , but this time with the  $\mathbf{y}_i^{t,*}$  selected arbitrarily from the set of all potential candidates, not just those in  $E^t$ , which again results in 10 edges. Finally, Figure 2e shows the links between each element of  $Y^t \setminus E^t$  and their guardian dominator, with the  $\mathbf{y}_i^{t,*}$  assigned as the *closest* dominating member of  $Y^t$ . Here we have used the Euclidean distance to determine the proximity of solutions, but other distance measures might be more useful in other situations.

As can be seen in Figure 2, a substantial number of domination relations can exist between the elements of a general set of data. The minimum number is  $|Y^t \setminus E^t|$ , which corresponds to each element of  $Y^t \setminus E^t$  being dominated by a single element of  $E^t$ ; that is, there are two Pareto shells, with each element of  $E^t$  dominating non-intersecting subsets of the second shell. The maximum is  $\binom{|Y^t|}{2}$ , which corresponds to the situation where the number of Pareto shells equals the number of elements in  $Y^t$  so that the element forming the first shell dominates  $|Y^t| - 1$  solutions, the next shell element dominates  $|Y^t| - 2$  solutions, and so on.

The minimum number of guardian dominators and corresponding domination links needed for a set  $Y^t$  is  $|Y^t \setminus E^t|$ , because each member of  $Y^t \setminus E^t$  must be assigned a guardian.

On inspection of Figures 2d-e, it can be seen that if the guardian dominator  $\mathbf{y}_i^{t,*}$  of an element  $\mathbf{y}_i^t$  is itself dominated, then there is a chain of domination links which eventually reaches a member of  $E^t$ . This is inevitable, because all el-



**Figure 2:** *a:* Graph representations of a general set of 14 solutions, plotted via their objective values in  $Y^t$ . *b:* All domination relations between solutions shown via edges. *c:* Single domination relation edge plotted per solution, only members of  $E^t$  selected as guardian dominators. *d:* Single domination relation edge shown per solution, guardian dominator selected arbitrarily from valid candidates in  $Y^t$ . *e:* Single domination relation edge shown per solution, closest dominating member selected as guardian dominator.

**Algorithm 1** Methods for determining  $E^{t+1}$  and for assigning guardian dominator when a new solution is generated. Best case and worst case domination comparisons required are indicated at the right using  $\mathcal{B}$  and  $\mathcal{W}$  respectively. Best case varies for options 2.1-2.4 depending on whether option 1.1 is selected or options 1.2-1.3 to obtain  $\mathbf{e}_{dom}^t$ . It is possible to combine this approach with data structures suggested in [1] to reduce the number of domination comparisons required for line 2.

```

1: if  $E^t \not\prec \mathbf{y}_{new}^{t+1}$  then
2:    $E^{t+1} := E^t \cup \{\mathbf{y}_{new}^{t+1}\}; E^{t+1} := E^{t+1} \setminus \{\mathbf{e}^t \in E^t \mid \mathbf{y}_{new}^{t+1} \prec \mathbf{e}^t\}$  New solution  $\mathbf{y}_{new}^{t+1}$  is not dominated by any element of  $E^t$ 
3:   Assign  $\mathbf{y}_{new}^{t+1}$  as guardian dominator for  $\{\mathbf{e}^t \in E^t \mid \mathbf{y}_{new}^{t+1} \prec \mathbf{e}^t\}$  Ensure  $E^{t+1}$  is a non-dominated set
4: else
5:    $E^{t+1} := E^t$   $\mathcal{B} = \mathcal{W} = |E^t|$ 
6:   option 1.1:  $\mathbf{e}_{dom}^{t+1} :=$  the first element in  $E^{t+1}$  which dominates  $\mathbf{y}_{new}^{t+1}$   $\mathcal{B} = 1$  or  $|E^t|$ ,  $\mathcal{W} = |E^t|$ 
7:   option 1.2:  $\mathbf{e}_{dom}^{t+1} :=$  the closest element in  $E^{t+1}$  which dominates  $\mathbf{y}_{new}^{t+1}$ 
8:   option 1.3:  $\mathbf{e}_{dom}^{t+1} :=$  as the element in  $E^{t+1}$  which dominates  $\mathbf{y}_{new}^{t+1}$  and which guards the fewest others
9:   option 2.1:  $\mathbf{y}_{new}^{t+1,*} := \mathbf{e}_{dom}^{t+1}$   $\mathcal{B} = 1$  or  $|E^t|$ ,  $\mathcal{W} = |E^t|$ 
10:  option 2.2:  $\mathbf{y}_{new}^{t+1,*} :=$  the first dominating element of  $Y_{dom}^t \cup \{\mathbf{e}_{dom}^{t+1}\}$   $\mathcal{B} = 1$  or  $|E^t| + 1$ ,  $\mathcal{W} = |E^t \cup Y_{dom}^t|$ 
11:  option 2.3:  $\mathbf{y}_{new}^{t+1,*} :=$  the closest dominating element of  $Y_{dom}^t \cup \{\mathbf{e}_{dom}^{t+1}\}$   $\mathcal{B} = |Y_{dom}^t| + 1$  or  $\mathcal{B} = \mathcal{W} = |E^t \cup Y_{dom}^t|$ 
12:  option 2.4:  $\mathbf{y}_{new}^{t+1,*} :=$  the dominating element of  $Y_{dom}^t \cup \{\mathbf{e}_{dom}^{t+1}\}$  which guards the fewest  $\mathcal{B} \& \mathcal{W}$  as option 2.3
13: end if

```

ements of  $Y^t$  have a dominating member associated with them except the members of  $E^t$  themselves. As such all chains *must* end with an  $E^t$  solution. Chains may join together as they approach  $E^t$ , but it is impossible for a dominated solution not to have a direct sequence to one of the elements of  $E^t$  via the guardian dominator links. This property is extremely useful, and we exploit it in our proposed approach for maintaining  $Y^t$  and identifying  $E^t$ . We denote the subset of  $Y^t$  for which  $\mathbf{y}_i^t$  acts as the guardian dominator by  $Y_i^t$ . Note that  $\mathbf{y}_i^t$  may not be the only element in  $Y^t$  which dominates members of  $Y_i^t$ , it is merely the dominator that is recorded.

#### 4.1 Selecting and updating guardians

At first sight it appears desirable to minimise the number of solutions for which any solution is the guardian dominator; that is to prefer arrangements like Figure 2d and 2e to Figure 2c. This sort of configuration minimises the number of domination comparisons that must be made if a guardian itself is changed. However, to achieve this more domination comparisons are required than, for example, assigning the first solution dominating  $\mathbf{y}^{t+1}$  to be  $\mathbf{y}^{t+1,*}$  the guardian dominator of  $\mathbf{y}^{t+1}$ . We suggest a number of different methods for selecting and updating guardian dominators in Algorithms 1 and 2 which span the range of low domination cost for assigning guardians, but potentially high domination cost when elements of  $Y^t$  are changed, through to higher

domination cost for guardian assignment, but fewer domination comparisons when a solution is changed. It should be noted that it is generally impossible to ensure  $|Y_i^t| \leq 1$  for all members of  $Y^t$ , because if there are more solutions in Pareto shell 2 than shell 1, then some members of  $E^t$  will inevitably be the guardian dominator for more than one element of  $Y^t$ .

Algorithm 1 presents the options we consider when a new solution  $\mathbf{y}_{new}^{t+1}$  is evaluated at time  $t + 1$ . If it is not dominated by the non-dominated subset of  $Y^t$  (line 1) then it is added to  $E^{t+1}$ , and any members of  $E^t$  which  $\mathbf{y}_{new}^{t+1}$  dominates are assigned  $\mathbf{y}_{new}^{t+1}$  as their guardian dominator. On the other hand, if  $\mathbf{y}_{new}^{t+1}$  is dominated, we consider a number of options to assign its guardian. First, one of the members of  $E^t$  which dominates  $\mathbf{y}_{new}^{t+1}$  must be selected (options 1.1-1.3), and labelled as  $\mathbf{e}_{dom}^{t+1}$ . Computationally the quickest option is to assign  $\mathbf{y}_{new}^{t+1,*}$  as the first element found in  $E^t$  which dominates  $\mathbf{y}_{new}^{t+1}$  (lines 6 & 9). At worst this requires  $E^t$  domination comparisons, but in the best case just a single domination comparison is required. However, this approach may result in non-dominated members of  $Y^t$  being the guardian dominators for many solutions (c.f. Figure 2c), so that many comparisons are required should the value of one of these guardians change later.

The idea underlying option 2.2 is to promote longer chains of domination links by assigning a guardian dominator from solutions that are not in the elite archive. Rather than searching through all  $Y^{t+1}$ , a greedy assignment is made.

---

**Algorithm 2** Methods for determining  $E^{t+1}$  when objective values of a member of  $Y^t$  change, and for assigning/updating guardian dominators. Best and worst case domination comparisons required are indicated using  $\mathcal{B}$  and  $\mathcal{W}$ .

---

```

1: if  $\mathbf{y}_{chg}^t \in E^t$  then
2:   if  $\{E^t \cup Y_{chg}^t \setminus \{\mathbf{y}_{chg}^t\}\} \not\prec \mathbf{y}_{chg}^{t+1}$  then
3:      $U = \{\mathbf{e}^t \in E^t | \mathbf{y}_{chg}^{t+1} \prec \mathbf{e}^t\}$ 
4:      $E^{t+1} = \{\mathbf{y}_{chg}^{t+1}\} \cup E^t \setminus U$ 
5:     Set  $\mathbf{y}_{chg}^{t+1}$  as guardian dominator for all elements in  $U$ 
6:   else
7:      $E^{t+1} = E^t \setminus \{\mathbf{y}_{chg}^t\}$ 
8:     option 3.1:  $\mathbf{y}_{chg}^{t+1,*} :=$  first dominating element of  $E^t \cup Y_{chg}^t$ 
9:     option 3.2:  $\mathbf{y}_{chg}^{t+1,*} :=$  closest dominating element of  $E^t \cup Y_{chg}^t$ 
10:    option 3.3:  $\mathbf{y}_{chg}^{t+1,*} :=$  dominating element of  $E^t \cup Y_{chg}^t$  which guards the fewest
11: end if
12: for each  $\mathbf{y}_i^t \in Y_{chg}^t$  do
13:   if  $\{E^{t+1} \cup Y_{chg}^t\} \not\prec \mathbf{y}_i^t$  then
14:      $E^{t+1} := E^{t+1} \cup \{\mathbf{y}_i^t\}$ 
15:   else
16:     option 4.1: if  $\mathbf{y}_{chg}^{t+1} \prec \mathbf{y}_i^t$  then  $\mathbf{y}_i^{t+1,*} := \mathbf{y}_{chg}^{t+1}$ , else assign first dominating element of  $E^{t+1}$   $\mathcal{B} = 1, \mathcal{W} = |E^{t+1}|$ 
17:     option 4.2:  $\mathbf{y}_i^{t+1,*} :=$  first dominating element of  $S = Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup E^{t+1}$   $\mathcal{B} = 1, \mathcal{W} = |S|$ 
18:     option 4.3:  $\mathbf{y}_i^{t+1,*} :=$  closest dominating element of  $S = Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup E^{t+1}$   $\mathcal{B} = \mathcal{W} = |S|$ 
19:     option 4.4:  $\mathbf{y}_i^{t+1,*} :=$  dominating element of  $S = Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup E^{t+1}$  which guards the fewest  $\mathcal{B} = \mathcal{W} = |S|$ 
20:   end if
21: end for
22: else
23:   if  $E^t \not\prec \mathbf{y}_{chg}^{t+1}$  then
24:      $U = \{\mathbf{e}^t \in E^t | \mathbf{y}_{chg}^{t+1} \prec \mathbf{e}^t\}$ 
25:     Set  $\mathbf{y}_{chg}^{t+1}$  as the guardian dominator for all elements in  $U$ 
26:      $E^{t+1} = \{\mathbf{y}_{chg}^{t+1}\} \cup E^t \setminus U$ 
27:   else
28:      $E^{t+1} := E^t$ 
29:     option 5.1:  $\mathbf{y}_{chg}^{t+1,*} :=$  first dominating element of  $\mathbf{y}_{chg}^{t,*} \cup E^{t+1}$   $\mathcal{B} = 1, \mathcal{W} = |E^t| + 1$ 
30:     option 5.2:  $\mathbf{y}_{chg}^{t+1,*} :=$  first dominating element of  $S = Y_{chg}^t \cup \mathbf{y}_{chg}^{t,*} \cup E^{t+1}$   $\mathcal{B} = 1, \mathcal{W} = |S|$ 
31:     option 5.3:  $\mathbf{y}_{chg}^{t+1,*} :=$  closest dominating element of  $S = Y_{chg}^t \cup \mathbf{y}_{chg}^{t,*} \cup E^{t+1}$   $\mathcal{B} = \mathcal{W} = |S|$ 
32:     option 5.4:  $\mathbf{y}_{chg}^{t+1,*} :=$  dominating element of  $S = Y_{chg}^t \cup \mathbf{y}_{chg}^{t,*} \cup E^{t+1}$  which guards the fewest  $\mathcal{B} = \mathcal{W} = |S|$ 
33:   end if
34:   for each  $\mathbf{y}_i^t \in Y_{chg}^t$  do
35:     option 6.1: if  $\mathbf{y}_{chg}^{t+1} \prec \mathbf{y}_i^t$  then  $\mathbf{y}_i^{t+1,*} := \mathbf{y}_{chg}^{t+1}$ , else  $\mathbf{y}_i^{t+1,*} := \mathbf{y}_{chg}^{t,*}$   $\mathcal{B} = \mathcal{W} = 1$ 
36:     option 6.2:  $\mathbf{y}_i^{t+1,*} :=$  first dominating element of  $S = Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup \{\mathbf{y}_{chg}^{t,*}\}$   $\mathcal{B} = 1, \mathcal{W} = |S|$ 
37:     option 6.3:  $\mathbf{y}_i^{t+1,*} :=$  closest dominating element of  $S = Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup \{\mathbf{y}_{chg}^{t,*}\} \cup E^{t+1}$   $\mathcal{B} = \mathcal{W} = |S|$ 
38:     option 6.4:  $\mathbf{y}_i^{t+1,*} :=$  dominating element of  $S = Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup \{\mathbf{y}_{chg}^{t,*}\} \cup E^{t+1}$  which guards the fewest  $\mathcal{B} = \mathcal{W} = |S|$ 
39:   end for
40: end if

```

---

$\mathbf{y}_{new}^{t+1,*}$  is taken as the first element of  $Y^t$  which has as a guardian the selected member of the elite archive which dominates  $\mathbf{y}_{new}^{t+1}$ . Specifically  $\mathbf{y}_{new}^{t+1,*}$  is assigned to be the first element in  $Y_{dom}^t \cup \{\mathbf{e}_{dom}^{t+1}\}$  which dominates  $\mathbf{y}_{new}^{t+1}$ .

Extending the idea of promoting long chains of domination links, in option 2.3 we seek to attach a changed element to an existing chain at a location distant from  $E^t$ . The new guardian is then chosen to be the closest (using the Euclidean metric) dominating member of the subset of  $Y^t$  for which  $\mathbf{e}_{dom}^{t+1}$  was the guardian, namely  $Y_{dom}^t$ . Choosing the closest ensures that the guardian cannot dominate any of the valid alternatives in  $Y_{dom}^t \cup \{\mathbf{e}_{dom}^{t+1}\}$ .

Finally, in option 2.4, the dominating member of  $Y_{dom}^t \cup \{\mathbf{e}_{dom}^{t+1}\}$  which is the currently guardian dominator for the

*Non-dominated solution being changed*

*Changed solution is still non-dominated*

*Identify dominated members of previous elite set*

$\mathcal{B} = \mathcal{W} = |E^t \cup Y_{chg}^t| - 1$   
*changed solution is now dominated*

$\mathcal{B} = 1, \mathcal{W} = |E^t \cup Y_{chg}^t|$

$\mathcal{B} = \mathcal{W} = |E^t \cup Y_{chg}^t|$

$\mathcal{B} = \mathcal{W} = |E^t \cup Y_{chg}^t|$

*Check all solutions that  $\mathbf{y}_{chg}^t$  was a guardian dominator for*

*Previously dominated solution is now non-dominated*

$\mathcal{W} = |E^{t+1} \cup Y_{chg}^t|$

*Dominated solution being changed*

*Changed solution is now non-dominated*

*Identify dominated members of previous elite set*

$\mathcal{W} = |E^t|$   
*Changed solution is still dominated*

$\mathcal{B} = 1, \mathcal{W} = |E^t|$

$\mathcal{B} = 1, \mathcal{W} = |E^{t+1}| + 1$

$\mathcal{B} = 1, \mathcal{W} = |S|$

$\mathcal{B} = \mathcal{W} = |S|$

$\mathcal{B} = \mathcal{W} = |S|$

*Check all solutions that  $\mathbf{y}_{chg}^t$  was a guardian dominator for*

$\mathcal{B} = \mathcal{W} = 1$

$\mathcal{B} = 1, \mathcal{W} = |S|$

$\mathcal{B} = \mathcal{W} = |S|$

$\mathcal{B} = \mathcal{W} = |S|$

*fewest elements of  $Y^t$  is selected as  $\mathbf{y}_{new}^{t+1,*}$ . This attempts to reduce the total number of domination links from a solution, but does not ensure that the selected guardian is at the end of a chain in  $Y_{dom}^t \cup \{\mathbf{e}_{dom}^{t+1}\}$ . Option combination 1.1 and 2.1 has the lowest best case cost, but is less effective at generating long chains. Options 2.2-2.4 require the same number of domination comparisons.*

Algorithm 2 presents the options we consider when an existing solution  $\mathbf{y}_{chg}^t$  is changed at time  $t + 1$ . If  $\mathbf{y}_{chg}^t \in E^t$  then the fact that  $\mathbf{y}_{chg}^{t+1}$  is not dominated by any element of  $E^t$  is insufficient to determine whether  $\mathbf{y}_{chg}^{t+1}$  should enter  $E^{t+1}$  because it may have moved to a position where it is dominated by a solution that it dominated at time  $t$ . Therefore it must also be compared to the elements of  $Y_{chg}^t$  to de-

termine whether it belongs to  $E^{t+1}$ . If it is not dominated by any elements of  $Y_{chg}^t$ , it joins  $E^{t+1}$ , and any elements of  $E^t$  which it now dominates do not enter  $E^{t+1}$  and are assigned  $\mathbf{y}_{chg}^{t+1}$  as their guardian dominator. If it does not enter  $E^{t+1}$  then we consider three options for assigning its guardian  $\mathbf{y}_{chg}^{t+1,*}$ . The first (line 8) assigns the first element of  $E^t \cup Y_{chg}^t$  found which dominates it, the second selects the Euclidean closest from the same set, and the third selects the dominating member from the set which currently guards the fewest members of  $Y^t$ .

Irrespective of the assignment of  $\mathbf{y}_{chg}^{t+1}$ , the solutions for which it was guardian dominator must now be checked to discover whether they are in  $E^{t+1}$ . If they are not elite, then we examine four options for assigning their guardian dominator. In the first option (line 16)  $\mathbf{y}_{chg}^{t+1}$  remains the guardian if its changed position still dominates the solution, otherwise the first dominating member of  $E^{t+1}$  is selected. The other options use  $Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup E^{t+1}$ , selecting either the first dominating member identified, the closest dominating member, or the dominating member guarding the fewest.

In the situation where  $\mathbf{y}_{chg}^t \notin E^t$ ,  $\mathbf{y}_{chg}^{t+1}$  only needs to be compared to  $E^t$  to determine if it is non-dominated or not. The options we consider for guardian dominator assignment to  $\mathbf{y}_{chg}^{t+1}$  if it is dominated are:  $\mathbf{y}_{chg}^{t,*}$  if it still dominates, otherwise the first dominating from  $E^t$ ; the first dominating element of  $Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t,*}\} \cup E^{t+1}$ ; the closest dominating member of this set; or the dominating member guarding the fewest. As in the earlier case, the guardian dominators for the elements of  $Y_{chg}^t$  may also need to be reassigned. In the first option  $\mathbf{y}_{chg}^{t+1}$  is kept as guardian dominator if its changed position still dominates the solution, if not the previous guardian dominator of  $\mathbf{y}_{chg}^{t+1}$  itself is selected, as this must dominate all members of  $Y_{chg}^t$ . The next three options use  $Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup \{\mathbf{y}_{chg}^{t,*}\} \cup E^{t+1}$ , the first selecting the first dominating member identified, the second selecting the closest dominating member, and the last selecting the dominating member which guards the fewest in  $Y^t$ .

## 5. EMPIRICAL ANALYSIS

Although we give the worst case numbers of domination comparisons needed for the different updating options, many of them are determined by the size of  $Y_{chg}^t$ , which is determined by the particular problem being optimised and how guardians were chosen earlier. We therefore illustrate how these various methods perform on synthetic data which is generated to mimic the behaviour of different convergence/search types, and on a running optimiser. We follow a general procedure in which, at alternate time steps, either the objective vector for an existing member of  $Y^t$  is changed or a new location  $\mathbf{y}_{new}^{t+1}$  is added.

In the simulation, for each member of  $Y^t$  we store an underlying ‘true’ objective location, which is never observed directly. Instead, an evaluation of a solution results in a noisy version of the true objective vector obtained by adding Gaussian noise. The  $\mathbf{y}$  of a solution is the mean of the noisy objective vector samples taken thus far. This mimics the refinement of objectives in noisy optimisation (see e.g. [9, 6]). We model the iterative generating process of  $Y^t$  in four distinct ways, based on two solution generation models, and two selection models. These mimic both how an optimiser

may progress over time, and how there may be bias in the solution selected for changing.

For solution generation the two regimes are:

1. The underlying ‘true’ objective locations of new solutions are drawn at random from a unit variance Gaussian. This emulates a search problem where the objective vector of a new solution bears little relation to fitness of the best members evaluated thus far (the likely parents), or one that is already well-converged and therefore rarely finds a better solution.
2. The underlying ‘true’ objective locations of new solutions are set as a perturbed value close to the ‘true’ objective vector of a randomly selected member of  $E^t$  (perturbed by additive multivariate isotropic Gaussian noise with  $\sigma_{mut} = 0.25$ ). This simulates a search where the evolved solutions are in the general region of objective space as previously discovered good solutions.

For selecting which  $\mathbf{y}^t$  to vary at a time step:

1.  $\mathbf{y}_{chg}^t$  is selected at random from  $Y^t$ . Any solution may have its objective vector changed.
2.  $\mathbf{y}_{chg}^t$  is selected at random from  $E^t$ . Only estimated Pareto solutions have their objective vector changed.

Combining a solution generation procedure with a selection procedure creates four distinct simulations.  $S_1$ : random new solutions, random changed solutions;  $S_2$ : random new solutions, random elite changed solutions;  $S_3$ : new solutions in vicinity of  $E^t$ , random changed solutions;  $S_4$ : new solutions in vicinity of  $E^t$ , random elite changed solutions.

We initialised  $Y^1$  with a single location in objective space, whose true objective vector is  $\mathbf{0}$ , which is perturbed by additive multivariate isotropic Gaussian noise with  $\sigma_{noise} = 0.1$ . The simulations reported here used two objectives and were run for 100,000 time steps. We use the option configurations detailed in Table 1, which cover a wide range of the methods described in Algorithms 1 and 2.

Figure 3 shows how key measures of complexity develop using the different methods of guardian selection on the simulations as  $|Y^t|$  increases. The top row shows the cumulative moving average size of selected  $Y_{chg}^t$ , averaged over 50 runs. This is recorded every time step where a member of  $Y^t$  is changed (and therefore a  $Y_{chg}^t$  was used). The middle row shows the (cumulative) mean number of domination comparisons required at each time step to identify  $E^t$ . Finally, the bottom row shows the ratio between the number of domination comparisons required, and the size of  $E^t$ , thus allowing us to discount the effect of different elite population sizes in the different simulations. A few key points may be observed. Although the methods which compare fewer solutions when assigning a guardian dominator to new locations ( $C_{1-3}$ ) tend to require fewer domination comparisons on some simulations (notably  $S_1$  and  $S_3$ ), these methods require orders of magnitude more domination comparisons on  $S_2$  (randomly chosen new solutions and randomly chosen elite changed solutions). This is because the elite members tend to act as the guardian dominators for many members of  $Y^t$ , compared with the other methods, and these heavily-loaded elements are also changed relatively frequently. Therefore, this combination of methods appears to be a poor choice when an algorithm is only advancing its Pareto front estimate slowly,

Label	Option combinations	Description of guardian assignment
$C_1$	1.1, 2.1, 3.1, 4.1, 5.1, 6.1	Assigns first dominating to $\mathbf{y}_{new}^{t+1}$ , and first dominating to $\mathbf{y}_{chg}^{t+1}$
$C_2$	1.2, 2.1, 3.1, 4.1, 5.1, 6.1	Assigns closest elite to $\mathbf{y}_{new}^{t+1}$ , and first dominating to $\mathbf{y}_{chg}^{t+1}$
$C_3$	1.3, 2.1, 3.1, 4.1, 5.1, 6.1	Assigns dominating elite which guards fewest to $\mathbf{y}_{new}^{t+1}$ , and first dominating to $\mathbf{y}_{chg}^{t+1}$
$C_4$	1.1, 2.2, 3.1, 4.2, 5.2, 6.2	Assigns first dominating found – always searching through $Y_{chg}^t$ first for $\mathbf{y}_{chg}^{t+1}$
$C_5$	1.2, 2.3, 3.2, 4.3, 5.3, 6.3	Assigns closest dominating solution from comparison sets
$C_6$	1.3, 2.4, 3.3, 4.4, 5.4, 6.4	Assigns dominating solution with fewest guarded solutions from comparison sets

Table 1: Methods combinations tested empirically. Options refer to Algorithms 1 and 2.

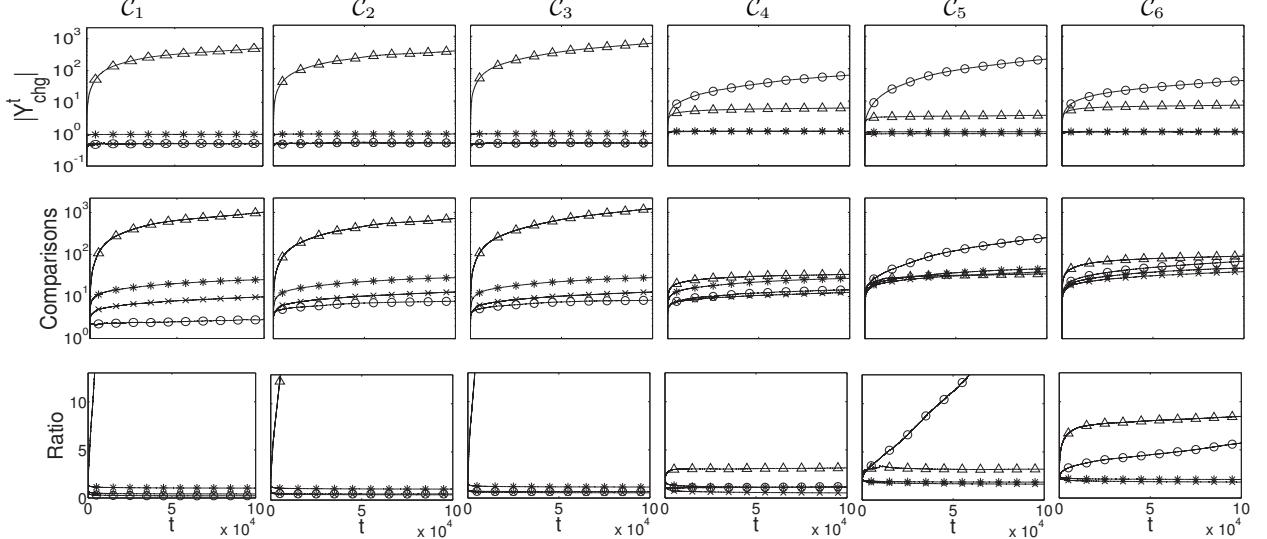


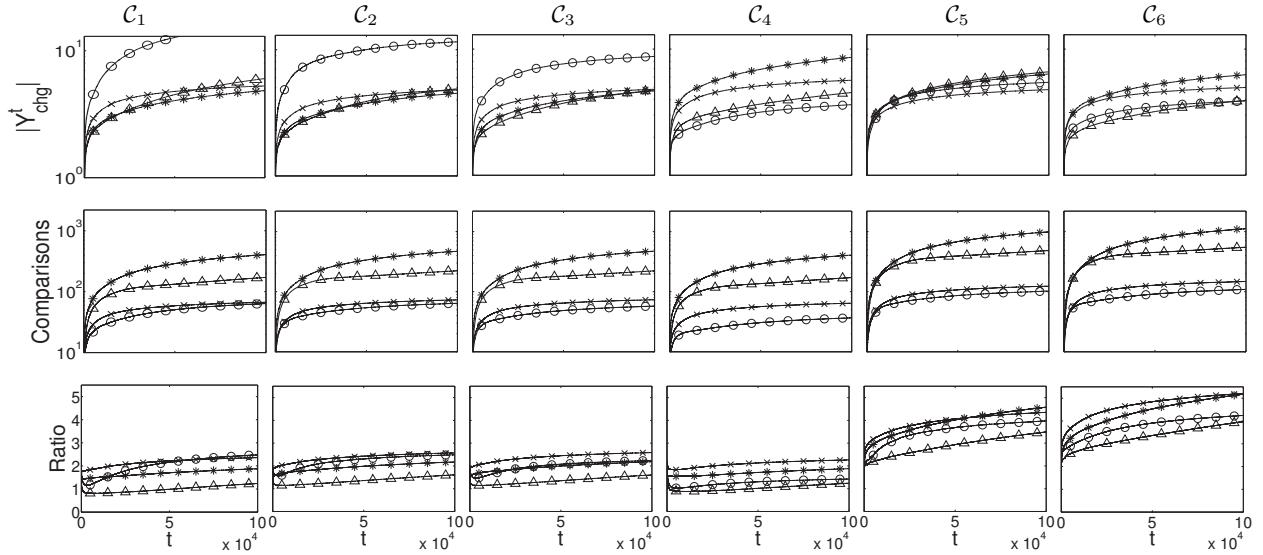
Figure 3: Growth of  $|Y_{chg}^t|$  and number of domination comparisons using different guardian selection methods. Top row (log scale):  $|Y_{chg}^t|$ . Middle row (log scale) number of domination comparisons. Bottom row: ratio of number of dominations to elite archive size. Each panel shows the cumulative moving average over 100,000 time steps, averaged over 50 runs. Circles denote simulation  $S_1$ , triangles:  $S_2$ , crosses:  $S_3$  and stars:  $S_4$ .

and where there is a bias towards changing elite solutions. Across the simulation types, the combination  $C_4$  appears to give the best overall performance, having a slightly worse ratio on simulation  $S_1$ , than  $C_{1-3}$ , but a vastly better ratio on  $S_2$ , and on all simulations the average number of comparisons per time step required is 0.5–3 times  $|E^t|$ .

In detail,  $C_4$  behaves as follows. If a new proposal is dominated it selects the guardian from the set of solutions guarded by the first member in  $E^t$  which dominates  $\mathbf{y}_{new}^{t+1}$ . If an existing elite member is changed and moved to a dominated location, its guardian is selected as the first dominating member of  $E^t \cup Y_{chg}^t$ . The new guardian of each  $Y_{chg}^t$  element in this case is selected as the first dominating member of  $Y_{chg}^t \cup \{\mathbf{y}_{chg}^{t+1}\} \cup E^{t+1}$  (otherwise, the element enters  $E^{t+1}$ ). If a dominated member of  $Y^t$  is moved to a non-dominated location, it is set as the guardian for any newly dominated elements  $E^t$ . If instead it moves to another still dominated location, its guardian dominator is kept the same (if it still dominates) or replaced with the first dominating member of  $E^{t+1}$  identified. All elements of  $Y^{t+1}$  for which  $\mathbf{y}_{chg}^t$  was guardian are compared to each other, the new location of the changed solution,  $\mathbf{y}_{chg}^{t+1}$ , and the location of the guardian dominator of  $\mathbf{y}_{chg}^t$  ( $\mathbf{y}_{chg}^{t,*}$ ). The first dominating element of this set is assigned as the guardian dominator for each previously guarded solution.

We also examined the effect of the methods in a simple multi-objective optimiser, applied to two of the DTLZ test problems [3], where objective values are subject to change. We use two and three objective versions of DTLZ1, which has multiple deceptive fronts, the majority of which are distant from the Pareto front, and DTLZ2, where random solutions are in the general vicinity of the Pareto front. We chose these two as they nicely span a range of expected population behaviours. We modified the DTLZ problems so that isotropic Gaussian noise with  $\sigma = 0.1$  contaminated the objectives. The problems were optimised with a  $(|E^t| + 1)$ -Evolution Strategy (ES), in which at each iteration (time step) a member of  $E^t$  was selected at random. This member was either mutated on a single design parameter (using Gaussian mutation with  $\sigma = 0.1$ ), creating  $\mathbf{y}_{new}^{t+1}$ , or reevaluated to update its mean objective vector estimate, resulting in  $\mathbf{y}_{chg}^{t+1}$ , using an accumulative resampling approach to improve the objective estimate [9, 6].

Figure 4 shows the empirical results of using the different methods to identify  $E^t$  for this optimiser. Note that the elite archives  $E^t$  identified by each method are identical, as is the ES’s behaviour (when started with the same random number generator seed). We are concerned solely with the cost of finding  $E^t$  at each time step. As with the synthetic data, we average results over 50 different runs for each method combination. Again, combination  $C_4$  is the best choice, having



**Figure 4:** Growth of  $|Y_{chg}^t|$  and number of domination comparisons using different guardian selection methods for noisy DTLZ problems. Top row (log scale):  $|Y_{chg}^t|$ . Middle row (log scale) number of domination comparisons. Bottom row: ratio of number of dominations to elite archive size. Each panel shows the cumulative moving average over 100,000 time steps, averaged over 50 runs. Circles: DTLZ1,  $D = 2$ . Triangles: DTLZ1,  $D = 3$ . Crosses: DTLZ2,  $D = 2$ . Stars: DTLZ2,  $D = 3$ .

better or equivalent comparison costs across problems, with the average number of comparison requires ranging from 1 to 2.2 times  $|E^t|$ .

## 6. DISCUSSION

This paper has presented a new method for rapidly discovering the non-dominated elite sets in evolving populations, considering both when new solutions are added to the population and when the objective values of solutions change during the optimisation. This latter case is particularly important for dynamic and noisy optimisation. The method relies on the assignment of a *guardian dominator* to every non-elite solution. Although a full average case complexity analysis is not available because it depends on the nature of the set being tracked, we show that guardian dominators can allow rapid location of the elite set in a changing population. Performance of the scheme depends on the way in which the population of solutions evolves and empirically we have found the different guardian assignment method perform differently in different situations. This opens up a future avenue of work on dynamically adapting the assignment method as a population evolves. Nonetheless, even a fixed assignment method vastly reduces the number of domination comparisons need to maintain the elite set and will allow noisy and dynamic problems to be efficiently handled.

We have exploited the time-cost improvements it provides in a state-of-the-art noisy optimiser [6], and example MATLAB code is available from <https://github.com/fieldsend>.

## 7. REFERENCES

- [1] N. Altwaijry and M. Menai. Data structures in multi-objective evolutionary algorithms. *J. Computer Science and Technology*, 27(6):1197–1210, 2012.
- [2] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization Problems. In *Parallel Problem Solving from Nature*, pages 849–858. Springer, 2000.
- [3] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable Multi-Objective Optimization Test Problems. In *IEEE Congress on Evolutionary Computation*, volume 1, pages 825–830, 2002.
- [4] J. Fieldsend and R. Everson. On the efficient use of uncertainty when performing expensive ROC optimisation. In *IEEE Congress on Evolutionary Computation*, pages 3984–3991, 2008.
- [5] J. Fieldsend and R. Everson. On the efficient maintenance and updating of Pareto solutions when assigned objectives values may change. Technical report, University of Exeter, UK, December 2013.
- [6] J. Fieldsend and R. Everson. The Rolling Tide Evolutionary Algorithm: A Multi-Objective Optimiser for Noisy Optimisation Problems. *IEEE Transactions on Evolutionary Computation*, in press.
- [7] C.-K. Goh and K. Tan. *Evolutionary Multi-objective Optimization in Uncertain Environments*. Springer, 2009.
- [8] J. Horn and N. Nafpliotis. Multiobjective Optimization Using the Niched Pareto Genetic Algorithm. Technical Report 93005, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, 1993.
- [9] T. Park and K. Ryu. Accumulative Sampling for Noisy Evolutionary Multi-Objective Optimization. In *Proceeding of the Genetic and Evolutionary Computation Conference*, pages 793–800, 2011.
- [10] S. Yang, Y. Ong, and Y. Jin. *Evolutionary computation in dynamic and uncertain environments*. Springer, 2007.