MULTIOBJECTIVE OPTIMISATION OF BICYCLE FRAMES USING SIMULATED ANNEALING

Apichart Suppapitnarm, Keith A. Seffen, Geoffrey T. Parks, Andy M. Connor and P. John Clarkson

Engineering Design Centre, Department of Engineering
University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, UK

E-mail: as314@eng.cam.ac.uk, keith.seffen@umist.ac.uk, gtp@eng.cam.ac.uk, amc50@eng.cam.ac.uk, pjc10@eng.cam.ac.uk

ABSTRACT

This paper describes a novel implementation of the Simulated Annealing algorithm designed to explore the trade-off surface in multiobjective optimisation problems, in which an appropriate annealing temperature is determined and controlled for each individual objective under consideration. The algorithm maintains and updates an archive record of the non-dominated solutions encountered during search. Thus, the final archive represents the trade-off surface between the objectives and enables the designer to make an informed decision when choosing the best overall solution. The algorithm's performance is illustrated by considering the multiobjective optimisation of bicycle frames subject to multiple loading conditions. The results obtained illustrate some important differences between the structural performance of men's and women's bicycles.

INTRODUCTION

Many, perhaps most, real-world design problems are, in fact, multiobjective optimisation problems in which the designer seeks to optimise simultaneously several conflicting performance attributes of the design. Evolutionary Algorithms are well suited to multiobjective optimisation and a number of different multiobjective Genetic Algorithms (GAs) have been developed (Fonseca and Fleming, 1995). Multiobjective GAs provide the means the explore the trade-off surface between competing objectives in a single optimisation and are therefore a very attractive tool for the designer, particularly when there is little prior knowledge available about the problem to be solved. However, GAs are not always easy to implement as, in order to perform well, they can require carefully chosen representations and operators.

Simulated Annealing (SA) (Kirkpatrick *et al.*, 1983) is an alternative stochastic optimisation method which has traditionally been used for single objective optimisation. However, Engrand (1997) has recently proposed a multiobjective variant on the SA algorithm.

In this paper we describe a multiobjective SA (MOSA) optimisation algorithm developed from Engrand's variant and use it to investigate the multiobjective performance of bicycle frames subject to multiple loading cases.

The frame is the structural heart of a bicycle and in everyday use is subject to many different loading conditions. The design of bicycle frames can be regarded as a structural optimisation problem where the designer is normally concerned with producing a lightweight bicycle frame that is stiff enough to satisfy several structural performance requirements. It is readily apparent that the optimisation of bicycle frames is a multiobjective optimisation problem. Despite the recent introduction (and imminent banning) of monocoque frames for racing bicycles, the conventional tubular diamond frame continues to dominate current bicycle designs. Decades of real world experience have demonstrated the satisfactory performance attributes of this design.

SIMULATED ANNEALING

The Simulated Annealing algorithm is a mathematical analogy of the physical annealing processes in solids (Kirkpatrick *et al.*, 1983). The algorithm, which is by nature a minimisation algorithm, employs a random search that not only accepts changes that decrease the objective function, f, but also some changes that increase it. The latter are accepted with a probability:

$$P_T = \exp\left(-\frac{\delta f}{T}\right) \qquad (1)$$

where δf is the change in f and T is a control parameter, which, by analogy with the original application, is known as the "system temperature" irrespective of the objective function involved. T is

lowered periodically according to an *annealing schedule*. When T is high almost all generated solutions are accepted, irrespective of whether or not they have improved the objective ($P_T \approx 1$). When T is low only the changes that do improve the objective ($\delta f < 0$ and $P_T > 1$) are accepted. Thus, a typical SA search will initially accept large deteriorations (when T is high) and, as the search continues and T decreases, fewer and fewer deteriorations will be accepted. It is this interesting feature of SA that makes it capable of escaping from local optima and gives it a significant advantage over many other optimisation methods.

MULTIOBJECTIVE SIMULATED ANNEALING

Principle of Archiving

To extend the SA algorithm to problems with several objectives, it is first necessary to introduce the concept of "archiving" the competing solutions. The principle of archiving in SA was first proposed by Engrand (1997) and can simply be illustrated as in Figure 1, for a two-objective problem. Here it is assumed that, for each objective, an improvement in performance corresponds to a reduction in the value of each objective.

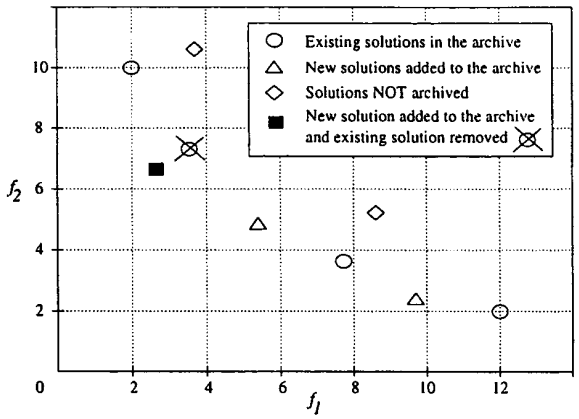

Figure 1. Illustration of archive evolution

At a given iteration, the current solution is compared with the solutions already in the archive. If this new solution dominates any existing solutions in the archive (i.e. has better values for *every* objective), those solutions are removed and the new solution is added. If the new solution is dominated by any members of the archive, it is not archived. If the new solution neither dominates nor is dominated by any members of the archive, suggesting an improvement in one objective but a worsening of another, it is added to the archive. It is this non-dominated set of archived solutions that eventually forms a trade-off surface between each of the competing objectives

Solution Acceptance

In analogy to single objective optimisation, an overall probability of solution acceptance for multiple objectives is defined by:

$$P_{all} = \prod_{i=1}^{N} \exp\left\{ -\frac{\{f_i(X_{n+1}) - f_i(X_n)\}}{T_i} \right\} \qquad (2)$$

in which T_i are assigned system temperatures for each objective f_i. Thus, the overall probability is the product of a series of individual acceptance probabilities for each objective, and each objective has its own associated temperature. This obviates the need to scale the objectives carefully with respect to each other, as required by the method proposed by Engrand (1997), as long as appropriate temperatures can be determined automatically.

Any move in which at least objective is decreased is potentially a move onto the trade-off surface. The acceptance probability of such a move depends on the relative changes of all the objectives (and the current temperatures). To avoid the clearly undesirable possibility of moves onto the trade-off surface not being accepted, because they fail the probabilistic test, each new solution generated is first submitted as a candidate for archiving. If the solution is archived, then it is automatically accepted. If it is not archived, then it is accepted with the probability given by Equation (2). This solution acceptance logic overcomes the other major weakness of the algorithm proposed by Engrand (1997). Thus the overall structure of our MOSA algorithm is as shown in Figure 2.

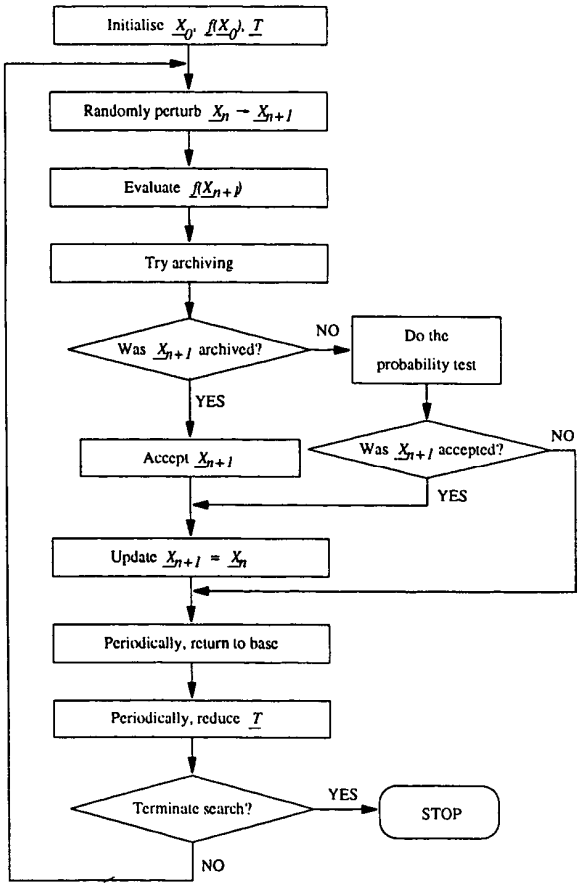

Figure 2. Structure of our MOSA algorithm

Annealing Schedule

In the initial part of the search, all the temperatures are set to ∞, so that all feasible perturbations are accepted. After a predetermined number of trials, the temperatures are set to appropriate values using White's formulation (White, 1984):

$$T_i = \sigma_i \tag{3}$$

where σ_i is the standard deviation of the variation in f_i observed. Periodically thereafter the temperatures are lowered according to the formula:

$$T_i^{'} = \alpha_i T_i \tag{4}$$

where α_i is determined using the formulation of Huang *et al.* (1986):

$$\alpha_i = \max\left[0.5, \exp\left(-\frac{0.7T_i}{\sigma_i}\right)\right] \tag{5}$$

Return to Base

In a traditional single objective SA implementation the "return to base" option retrieves the best solution found and continues the search from there. In our MOSA implementation when a return to base occurs a solution is retrieved from the archive (which contains the best solutions found). The selection of solutions from the archive favours the most isolated and extreme solutions on the trade-off surface, in order to try to ensure a uniformly distributed exposure of solutions on the trade-off surface.

In the implementation used in this study, the i^{th} return-to-base is made after N_{Bi} iterations where

$$N_{Bi} = 0.9N_{Bi-1}; \quad N_{Bi} \geq 10, \tag{6}$$

with $N_{B1} = 1000$.

The return-to-base solution is selected randomly from a candidate list of A_{Ii} solutions from the archive, with

$$A_{Ii} = 0.9^{i-1} A_S; \quad A_{Ii} \geq 10, \tag{7}$$

where A_S is the number of solutions in the archive. The candidate list contains the M extreme solutions (M is the number of objectives), i.e. those at the ends of the trade-off, and the $A_{Ii} - M$ most isolated solutions, where the degree of isolation of a solution j is defined by

$$I_j = \sum_{\substack{i=1 \\ i \neq j}}^{A_S} \sum_{k=1}^{M} \left(\frac{[f_k(i) - f_k(j)]}{[f_{k\max} - f_{k\min}]} \right)^2 \tag{8}$$

in which $f_{k\max}$ and $f_{k\min}$ are the maximum and minimum values of the k^{th} objective function.

See Suppapitnarm (1998) for more details of this feature of the algorithm.

OPTIMISATION OF BICYCLE FRAMES

Despite the development and introduction of new composite materials in the design of the frames of racing bicycles, carbon steel has long been the most widely used material for bicycle frames and the diamond tubular frame shape has remained virtually unchanged for most of this century. It might be concluded that over this lengthy period of development and refinement the optimum bicycle frame design must have been identified. It is nevertheless interesting to see if we can improve the design of this topology further.

As the "best" frame usually means "light weight" and "sufficiently stiff" to support all loading conditions, with an adequate safety factor to cope with unforeseen variation, we consider the optimisation of steel bicycle frames with multiple loading cases and multiple objectives. We restrict our study to two-dimensional loading cases in order to simplify the structural analysis required. The aim of this multiobjective optimisation study is to explore the trade-off between the competing objectives. This will inevitably necessitate the evaluation of many thousands of designs, and thus computational efficiency is essential.

4 Common Loading Cases

For the purposes of this study, the four loading cases considered by Mendis (1996) – starting, speeding, rolling and braking – are used to evaluate the structural performance of different frame designs. The loads applied at the three application points for each case are given in Table 1. These application points and the definition of the x and y directions are shown in Figure 3. Following the model in Mendis (1996), node O_1 is assumed fixed in all directions while node O_2 is allowed to move in the x direction only. These are kinematically and force constrained as the wheels can "splay" apart under loads that are, in fact, applied by inertia forces. The nodes (shown as solid circles) represent the ends of each beam element within the structure which are free to displace and rotate in plane.

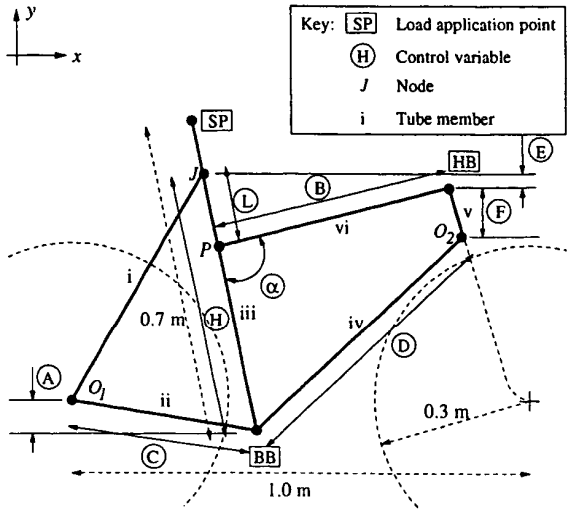

Figure 3. Bicycle frame representation

Loading case	Application point	F_x (N)	F_y (N)	M (Nm)
Starting	Handle bar [HB]	200	576	-48.8
	Seat post [SP]	0	0	0
	Bottom bracket [BB]	4930	-2465	0
Speeding	Handle bar [HB]	100	54	-12.7
	Seat post [SP]	20	-450	-31
	Bottom bracket [BB]	850	-425	0
Rolling	Handle bar [HB]	0	-900	45
	Seat post [SP]	0	-900	54
	Bottom bracket [BB]	0	-1200	0
Braking	Handle bar [HB]	-800	-700	115
	Seat post [SP]	-300	-400	36
	Bottom bracket [BB]	-300	-600	0

Table 1. Applied forces on a bicycle frame based on an average Australian rider (Mendis, 1996)

Control Variables

We assume that the frame is constructed from uniform tubular beam elements of circular cross-section with constant thickness made from isotropic material (Young's Modulus = 207 GPa, Poisson's ratio = 0.3). 21 control variables are therefore required to define completely the dimensions of the circular members making up the frame:

- 9 variables defining the frame geometry (and hence tube lengths) as shown in Figure 3;
- 6 variables defining the radii of each of the six members (i – vi);
- 6 variables defining the thicknesses of each of the six members.

3 Objectives

Three objective functions are considered in the optimisation process:

1. Minimisation of the mass of the frame;
2. Minimisation of average frame deflection (the frame deflection being defined as the sum of the displacements of points *J*, BB, HB and O_2) in all 4 loading cases (in search of the frame with the best average stiffness);
3. Minimisation of the standard deviation (SD) of the deflections in these loading cases (in search of the frame with the most consistent performance).

Constraints

The wheel base and wheel radii are fixed at 1 m and 0.3 m respectively. It is assumed that the length of the seat tube (H) plus that of the seat post is constrained to be 0.7 m. These dimensions are estimated based on the size of an average rider. The angle of the head tube is constrained such that its projection passes behind the centre of the front wheel and in front of the contact patch, for well established stability reasons (Whitt and Wilson, 1982).

In addition to these geometric and ergonomic constraints, the maximum stresses (σ_{Ti}) in each bar and the displacements (Δ_i) at all free nodes are constrained as follows:

$$\begin{aligned} &|\sigma_{Ti}| \le 100 \text{ MPa}, \quad i = 1,6 \\ &|\Delta_i|_{x,y} \le 10 \text{ mm} \end{aligned} \tag{9}$$

The stress at the joint between the seat post and the seat tube length (*J*) is also constrained as in Equation (9). This is found to be a particularly important constraint and is often the limiting one for lighter frames.

In practice, bicycles can fail in a large number of ways, including buckling, fatigue failure, corrosion, fast fracture etc. Any detailed examination of a candidate frame design would, of course, have to consider all possible failure mechanisms, but, for the purposes of this scoping study, it is assumed that feasible designs can be adequately screened using a suitable maximum stress constraint.

IMPLEMENTATION

The optimisation routine based on our MOSA algorithm applied to this bicycle frame problem was implemented in C++. The mass of frames was easily determined from the control variables. The deflection related objectives were calculated using a stiffness matrix, finite element approach (Paz, 1985) implemented using MATLAB® (The Mathworks, 1998).

We considered both men's and women's bicycle frame types. The stiffness matrix for men's frames (in which the distance *JP* in Figure 3 is always zero) is different from that used for women's frames. Therefore the routine requires an automatic switch between these two matrices for the MATLAB® calculations, if the optimisation is to search for both frame types simultaneously.

We examined three cases as follows:

- Only men's frame designs are generated (*JP* is always zero);

- Only women's frame designs are generated (*JP* is within some specified range);
- Both men's and women's frame designs can be generated in a single optimisation run, with a switch to the men's stiffness matrix calculation being activated when $JP < 10$ mm.

RESULTS AND DISCUSSIONS

Figure 4 shows the trade-off surfaces (projected in two-objective space) found after 20,000 iteration optimisations of both men's and women's frame designs. Note that only solutions with masses less than 10 kg are shown. It is clearly seen that virtually all the women's frames are dominated by men's ones. Men's frames have a conventional diamond (triangulated) shape, while women's frames do not due to the lowering of the top bar. As the length-to-diameter ratio of each tube is large, truss-like behaviour is not unexpected, even though the joints are rigid. Thus, the tubes in a men's frame are essentially in tension and compression, and, because of the triangulation, form an extremely strong and stiff structure in its own plane. Hence, for the same mass, men's frames will normally deflect less than women's ones.

Figure 5 shows the evolution of the trade-off surface (projected in mean deflection-mass space) for a run in which we started the optimisation from a women's frame initial solution and allowed both men's and women's frames to be generated. For ease of discussion the non-dominated solutions are classified as follows:

- Men's layout – *JP* is close to zero;
- Close-to-men's layout – solutions which have low values of *JP* (within the 10–100 mm range);
- Women's layout – solutions which have values of *JP* above 100 mm.

Initially a large proportion of the archived solutions has a women's layout, reflecting the fact that the search started from a solution with that configuration. But Figure 5 shows that, as the search progresses and the trade-off surface is advanced, fewer and fewer solutions on the trade-off surface have a women's layout. By the end of the optimisation (after 25,000 iterations) the vast majority of archived solutions have men's layouts, a number have close-to-men's layouts and only a handful have women's layouts. It is also found that the final trade-off surface is very similar to that obtained if only men's frames are allowed in the optimisation, as shown in Figure 6.

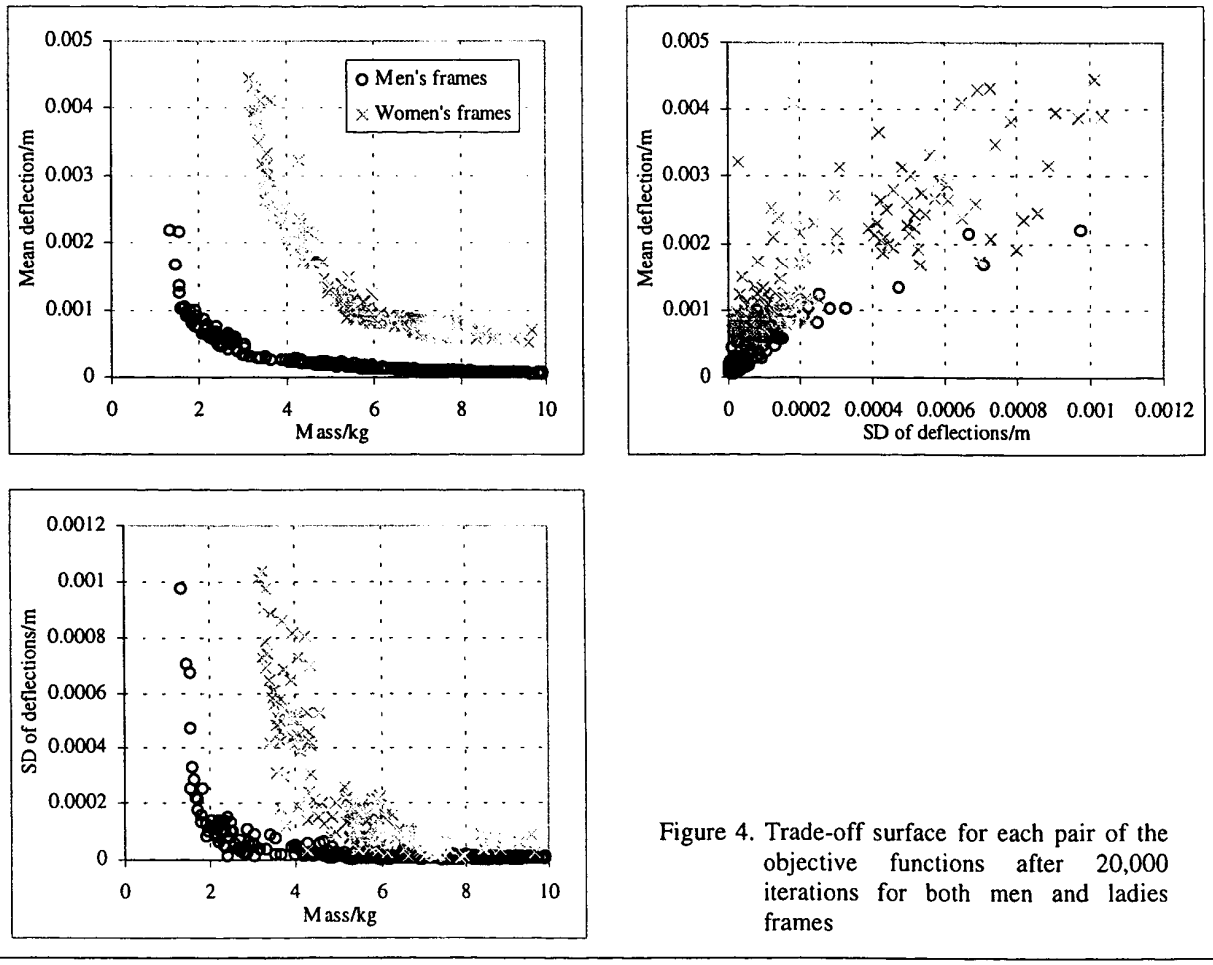

Figure 4. Trade-off surface for each pair of the objective functions after 20,000 iterations for both men and ladies frames

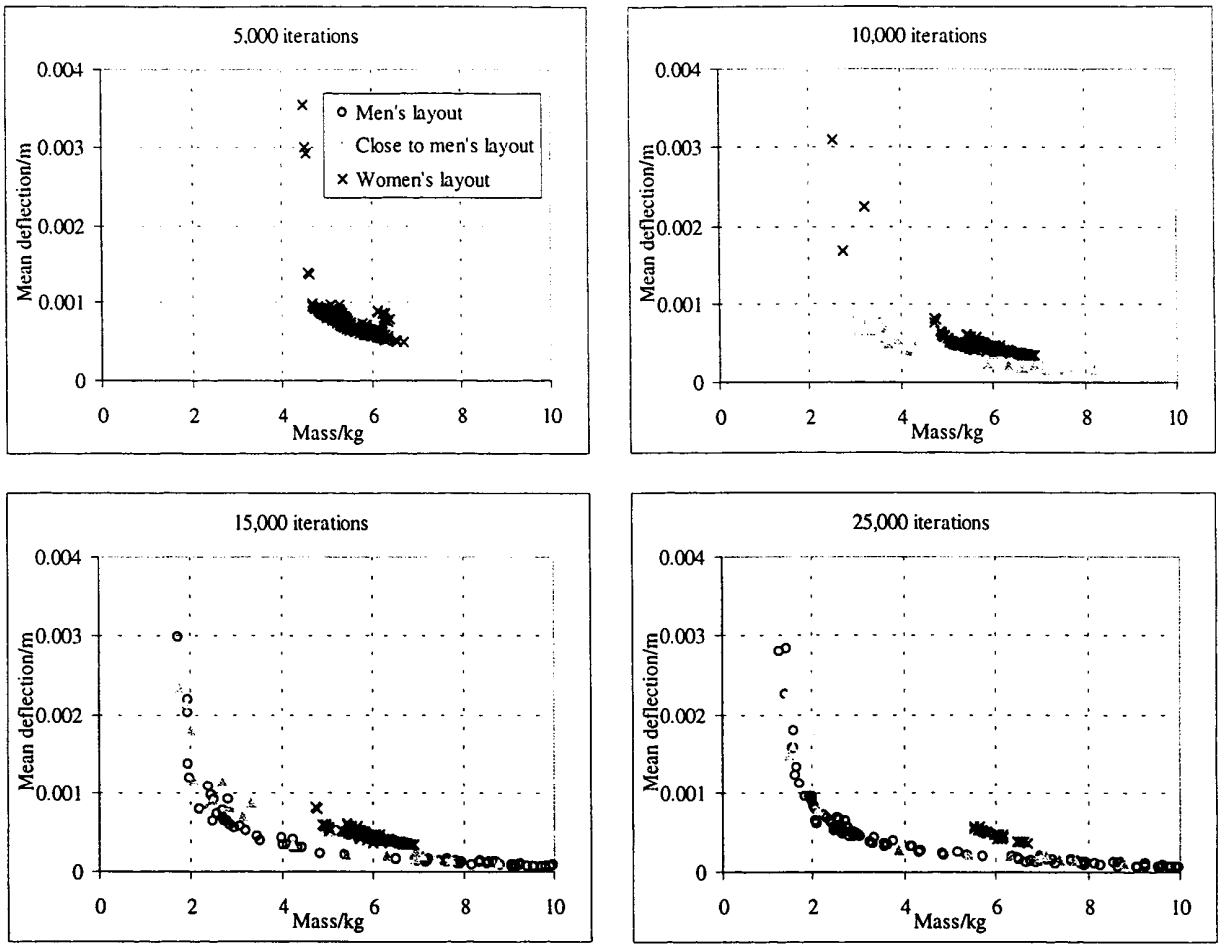

Figure 5. Evolution of the trade-off surface (projected in mean deflection-mass space) during a 25,000 iteration optimisation in which both men's and women's frame types are allowed simultaneously

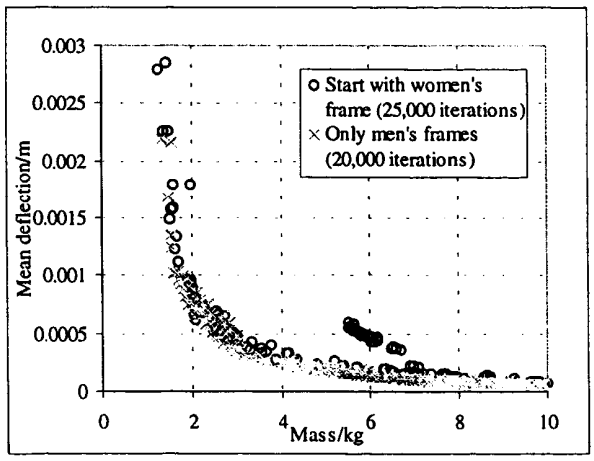

Figure 6. Comparison of trade-off surfaces obtained

Our studies show that a handful of designs with women's layouts remain on the trade-off surface shown in Figure 5 even if the optimisation is allowed to continue for a further few thousand iterations. These solutions, although less stiff than men's frame designs (for the same mass), exhibit more consistent structural performance under different loadings (resulting in lower deflection SDs – the third objective).

Figure 7 summarises the results obtained by showing some examples of different compromise frame designs identified from the trade-off surface. Note that thicknesses of the lines representing the frame members are logarithmically proportional to their second moments of area. The control variable values for these selected designs and their objective function values are given in Table 2. Examination of these solutions reveals a sensible variation in solutions from heavy and very stiff to lighter but less stiff frames.

An interesting feature of both the examples of women's frames shown (and all those archived) is that the rear stay connects to the seat tube just beneath the seat, i.e. much higher than is actually observed in practice. This feature is clearly important to the consistent performance these frames exhibit under the loading cases considered.

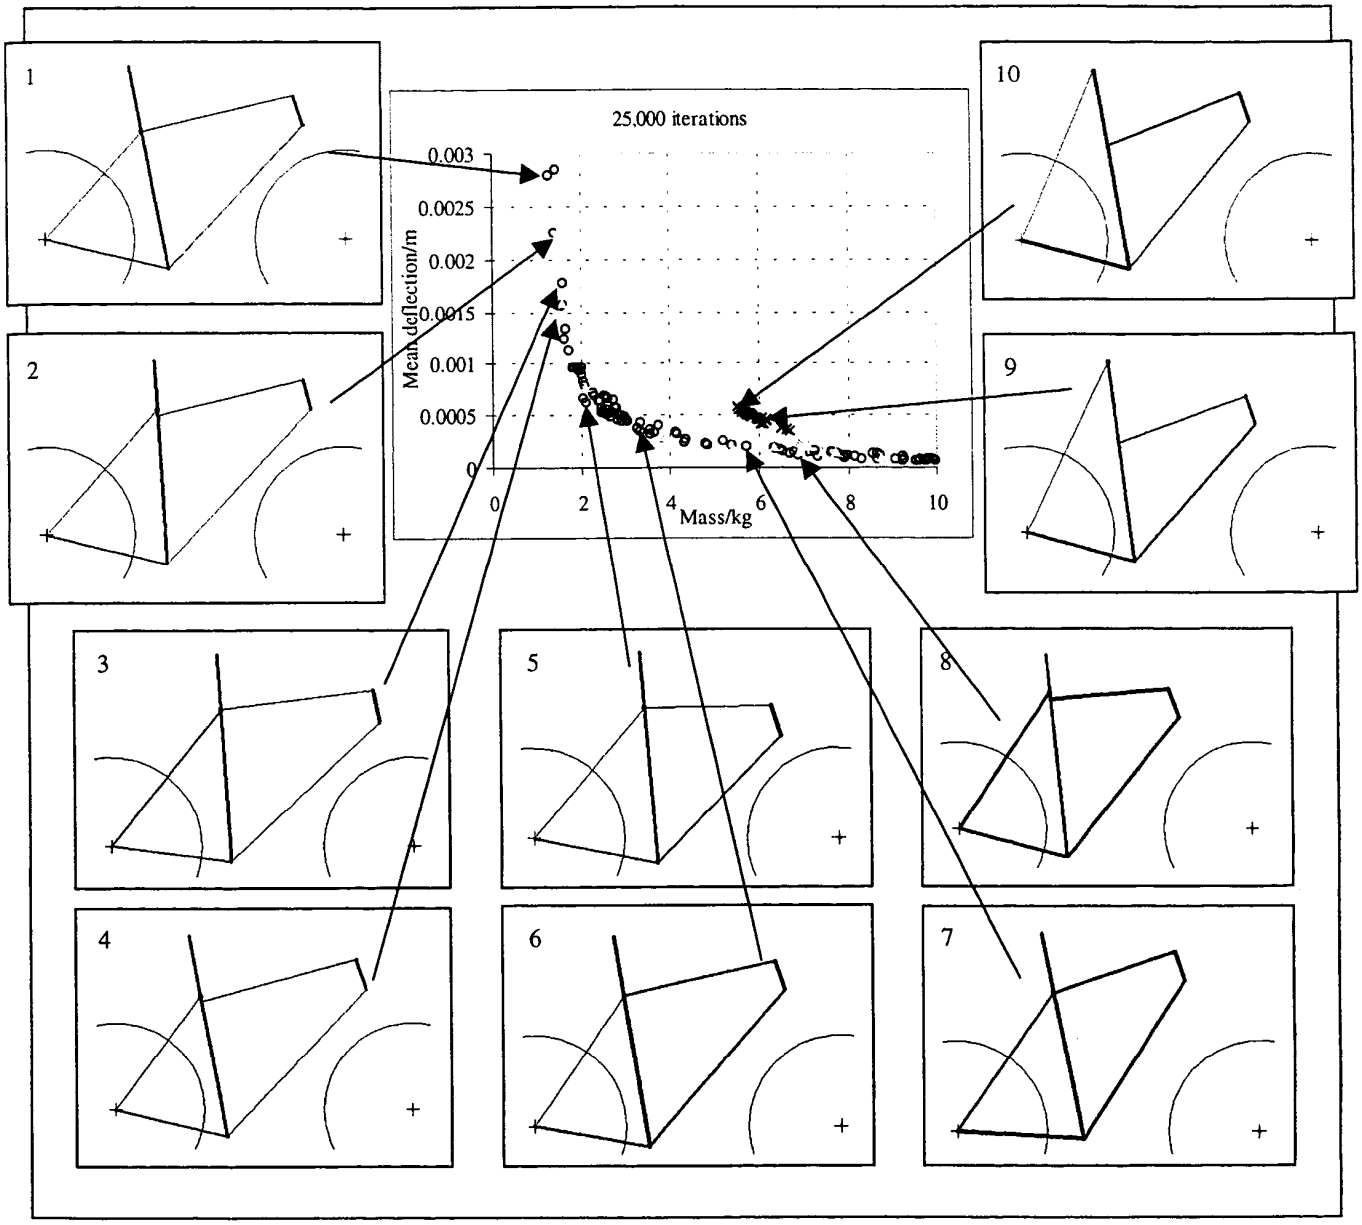

Figure 7. A selection of bicycle frame designs from the final trade-off surface in Figure 5

The examples of men's and close-to-men's frames shown all have rather similar configurations, although the heavier/stiffer frames are more compact, which, given the fixed wheel base, implies that longer forks are required. The main compression members (the seat tube and the head tube) do not vary in size significantly across the trade-off. Reductions in mass, at the cost of lower stiffness, are achieved through reductions in the size of the other members. Under the loading cases considered these members are either in tension or lightly loaded in compression. These loading cases are associated with typical road bike usage. If an off-road bicycle was being designed, the loading cases would be markedly different.

CONCLUSIONS

The foregoing example demonstrates that our MOSA algorithm can successfully explore and expose the trade-off surfaces between competing objectives. The exposure of the trade-off surface then presents the designer with a wealth of information about the range of achievable values of different (conflicting) objectives. Using this information the designer is able to make a much better informed decision about the areas of the search space worthy of closer examination.

The results of our studies presented here confirm that men's bicycle frames with a conventional diamond triangular shape are structurally superior to women's frames. Hence, in an optimisation process, if both frame

Variable \ Selected design	1	2	3	4	5	6	7	8	9	10
H (m)	0.4723	0.5286	0.5158	0.4908	0.5182	0.5049	0.4983	0.5699	0.6993	0.6996
B (m)	0.5209	0.5067	0.5052	0.5404	0.4146	0.5053	0.4243	0.4026	0.4658	0.4920
C (m)	0.4232	0.4177	0.3980	0.3864	0.4171	0.3820	0.4202	0.3822	0.3811	0.3850
α (°)	92.9	101.1	93.6	95.4	86.0	93.5	97.7	89.6	103.7	102.6
A (m)	0.1000	0.1000	0.0525	0.0909	0.0828	0.0660	0.0267	0.0976	0.0997	0.0999
E (m)	-0.1205	-0.1049	-0.0624	-0.1270	-0.0044	-0.1120	-0.1375	-0.0130	0.1191	0.0773
F (m)	0.1002	0.1050	0.1087	0.1061	0.1000	0.1000	0.1000	0.1000	0.1001	0.1009
D (m)	0.6643	0.7182	0.6796	0.6875	0.5843	0.6737	0.6243	0.6148	0.6329	0.6594
d_i (m)	0.0050	0.0056	0.0058	0.0050	0.0050	0.0050	0.0161	0.0324	0.0051	0.0052
d_{ii} (m)	0.0074	0.0067	0.0069	0.0070	0.0082	0.0171	0.0406	0.0235	0.0492	0.0466
d_{iii} (m)	0.0207	0.0238	0.0262	0.0282	0.0390	0.0500	0.0461	0.0318	0.0490	0.0454
d_{iv} (m)	0.0050	0.0050	0.0050	0.0050	0.0107	0.0149	0.0302	0.0251	0.0245	0.0083
d_v (m)	0.0327	0.0373	0.0428	0.0355	0.0426	0.0500	0.0500	0.0426	0.0407	0.0459
d_{vi} (m)	0.0050	0.0050	0.0050	0.0050	0.0050	0.0174	0.0284	0.0426	0.0120	0.0141
t_i (m)	0.0020	0.0020	0.0021	0.0020	0.0021	0.0020	0.0021	0.0025	0.0022	0.0021
t_{ii} (m)	0.0025	0.0025	0.0025	0.0026	0.0025	0.0025	0.0025	0.0026	0.0025	0.0025
t_{iii} (m)	0.0024	0.0024	0.0024	0.0024	0.0024	0.0025	0.0024	0.0032	0.0031	0.0031
t_{iv} (m)	0.0017	0.0018	0.0012	0.0019	0.0016	0.0010	0.0015	0.0042	0.0017	0.0018
t_v (m)	0.0025	0.0025	0.0025	0.0025	0.0025	0.0025	0.0025	0.0039	0.0041	0.0041
t_{vi} (m)	0.0025	0.0025	0.0025	0.0024	0.0025	0.0025	0.0026	0.0022	0.0026	0.0025
L (m)	0	0.0198	0	0.0198	0	0	0	0.0235	0.2856	0.2666
Mass (kg)	1.2704	1.4821	1.5806	1.4962	2.0925	3.3072	5.7158	6.8964	6.0669	5.5334
Mean deflection (m)	2.80E-03	2.26E-03	1.78E-03	1.49E-03	6.22E-04	3.55E-04	2.00E-04	1.45E-04	4.69E-04	5.52E-04
SD of deflections (m)	1.50E-03	6.45E-04	3.92E-04	6.11E-04	1.03E-04	6.42E-05	3.07E-05	4.01E-05	1.76E-06	3.57E-06

Table 2. Control variables and objective functions of the selected designs, d = diameter, t = thickness

types can be generated simultaneously, the majority of solutions are likely to be men's. Nevertheless there are a few good frames with women's layouts. These solutions, although not as stiff, offer more consistent structural performance under a variety of loadings.

The major advantage of our MOSA algorithm over a typical multiobjective GA (MOGA) is that it is a much easier algorithm to implement, as, unlike GAs, it does not require carefully chosen representations and operators. Our MOSA algorithm has been found to give comparable performance to MOGAs in terms of computational efficiency on problems where a direct comparison is possible (Suppapitnarm *et al.*, 1999; Parks and Suppapitnarm, 1999). It therefore seems to have considerable promise as a multiobjective optimisation tool for designers.

REFERENCES

Engrand, P. (1997), "A multi-objective approach based on Simulated Annealing and its application to nuclear fuel management", *Proc. 5th International Conference on Nuclear Engineering*, Nice, May 25-29 1997, pp. 416-423.

Fonseca, C.M. and Fleming, P.J. (1995), "An overview of evolutionary algorithms in multi-objective optimisation", *Evolutionary Computation*, Vol. 3 No. 1, pp. 1-16.

Huang, M.D., Romeo, F. and Sangiovanni-Vincentelli, A. (1986), "An efficient general cooling schedule for Simulated Annealing", *Proc. IEEE International Conference of Computer Aided Design*, Santa Clara CA, November 11-13 1986, pp. 381-384.

Kirkpatrick, S., Gerlatt, C.D. Jr and Vecchi, M.P. (1983), "Optimization by Simulated Annealing", *Science*, Vol. 220, pp. 671-680.

The Mathworks Inc. (1998), "*MATLAB* ®" Version 5.1.0.421, Natick MA.

Mendis, C. (1996), "*Optimisation of a Carbon Composite Bicycle Frame*", BEng thesis, Department of Aeronautical Engineering, University of Sydney.

Parks, G.T. and Suppapitnarm, A. (1999), "Multiobjective optimization of PWR reload core designs using Simulated Annealing", to appear in *Proc. International Conference on Mathematics and Computation, Reactor Physics and Environmental Analysis in Nuclear Applications*, Madrid, September 27-30 1999.

Paz, M. (1985) "*Structural Dynamics: Theory and Computation*", Van Nostrand Reinhold, New York.

Suppapitnarm, A. (1998), "*A Simulated Annealing Algorithm for Multiobjective Design Optimization*", MPhil thesis, Engineering Department, University of Cambridge.

Suppapitnarm, A., Seffen, K.A., Parks, G.T. and Clarkson, P.J. (1999), "A Simulated Annealing algorithm for multiobjective optimisation", submitted to *Engineering Optimization*.

Whitt, F.R. and Wilson, D.G. (1982) "*Bicycling Science*", The MIT Press, Cambridge MA.

White, S.R. (1984), "Concepts of scale in Simulated Annealing", *Proc. IEEE International Conference of Computer Aided Design*, Port Chester NY, October 8-11 1984, pp. 646-651.