

# Survival of the Sickest: A Site-Specific Recombination Operator for Accelerated Function Optimization

Stephen Drake    Phil Husbands

School of Cognitive and Computing Sciences  
University of Sussex  
Brighton, UK, BN1 9QH  
Email: {stevedr, philh}@cogs.susx.ac.uk

Cognitive Science Research Paper No. 530  
March 2001

## Abstract

This paper describes experiments with a new crossover operator which is based on the mechanism of biological site-specific recombination. By using hill-climbing to gather additional information about the fitness landscape, it increases the constructive power of crossover. The nature of the operator calls for a somewhat unusual selection strategy which, in contrast with traditional methods, automatically selects a relatively unfit member of the population to be a parent. When applied to difficult continuous-variable function optimization problems, the operator is seen to perform better than standard one-point crossover in terms of quality of solutions found and its speed in finding them (as measured by the number of function evaluations carried out).

## 1 Introduction

A considerable body of literature has amassed which attempts to explain crossover's place among evolutionary algorithms, and research has yielded increasingly sophisticated analyses since the time when conventional wisdom decreed that recombination was the primary engine of optimization while mutation was regarded as being somehow 'less powerful'. However, a complete and accurate analysis is still proving elusive, with much of the research either posing more questions than it answers (questions such as: is crossover redundant? Is crossover simply macromutation?) or contradictory, in the case of Holland's Schema Theorem versus Goldberg's Building Blocks Hypothesis [12]. There appears to be scant work which marries investigation of such theoretical bases as those mentioned above to the efficacy of crossover as observed empirically (other than evaluating its performance under variations of parameter values (for example, [4, 10]) or in recombination strategy itself, such as in [11, 13]). Moreover, it can be illustrated (see section 5) that, very often, standard one-point crossover constructs better solutions for only a relatively short length of the time; it very quickly becomes redundant, and a GA can then only rely on mutation to explore different areas of the search space. Thus the immediate motivation for this paper was to improve the performance of crossover by increasing the length of time for which it constructs better solutions. An 'intelligent' operator is proposed, which incorporates a generic – that is, problem non-specific – mechanism inspired by biological site-specific recombination. Biological site-specific recombination works by mediating proteins that bind specific target sequences in a strand of DNA, and catalyse recombination at those positions [6].

## 2 An Artificial Site-Specific Recombination Operator

This operator incorporates a hill-climbing technique into its mechanism, and can so be viewed as a hybrid. Hybrid approaches have traditionally involved employing problem-specific search techniques, or relegating the additional technique to a supplementary capacity by bolting it onto the end of the GA: for example, applying some local search to the fittest individual (or, indeed, every individual) at the end of a generation, or some hill-climbing at the end of a run. However, in this site-specific recombination (S-SR) operator, hill-climbing plays a much more important rôle, and can be seen as imbuing the operator with its 'intelligence'.

## 2.1 Allele Bounds

The structure of an individual in the population has been modified to incorporate a higher and a lower ‘bound’ for every gene of its genotype, as illustrated in figure 1. (These are initialised respectively to be the higher and lower values of the evaluation function’s variable range.) During the random generation of a population, and after every mating cycle, an approximation to random-mutation hill-climbing (RMHC) is applied to a specified percentage of an individual’s genes. RMHC mutates random loci in the current best genotype until a maximum number of evaluations have been performed (or until the optimum string has been found), only retaining mutations which lead to an equal or higher fitness [9]. This strategy was adapted to accommodate real-valued genotypes and the concept of allele bounds.




Figure 1: allele bounds

The values of a gene’s allele bounds are altered according to the effects of hill-climbing on that gene: if incrementing a gene’s value by a constant `STEP_SIZE` brings about a fitness increase, its lower bound is set to the gene’s new value; conversely, if there occurs a decrease in fitness, the gene, together with its upper bound, is reset to its previous value. (The above operates in reverse if hill-climbing decrements a gene’s value; incremental hill-climbing is applied first; if this effects a fitness decrease, the gene’s value is reset and decremental hill-climbing is applied.) The algorithm is represented by the following pseudo-code:

```

REPEAT N TIMES
  Select random Gene in CurrentGenotype
  Increment Gene by StepSize
  IF CurrentGenotype is fitter OR of equal fitness THEN
    Lower bound of Gene = Gene
  ELSE IF CurrentGenotype is less fit THEN
    Reset Gene to original value
    Upper bound of Gene = Gene
    Decrement Gene by StepSize
  IF CurrentGenotype is fitter OR of equal fitness THEN
    Upper bound of Gene = Gene
  ELSE IF CurrentGenotype is less fit THEN
    Reset Gene to original value
    Lower bound of Gene = Gene
  END IF
END IF
END REPEAT

```

## 2.2 Selection and Recombination

Mating was restricted to small neighbourhoods of individuals as detailed in section 3.2. In contrast to traditional methods of selection, the least fit member of a neighbourhood is selected as the first parent. (Preliminary experiments were carried out, of which each set of 50 runs consecutively selected the next fittest neighbour to be the first parent. Results showed that the worse the first parent is, the more effective the S-SR operator proves to be – unsurprisingly, given the nature of the operator, since there is more scope for improvement in an unfit individual.)

A set of allele bounds determines the general direction in the search space which an individual takes by defining the range in which the bounds' gene may find a profitable point. Moreover, if an increase in a gene's value due to hill-climbing effects an increase in the individual's fitness then presumably moving in the immediately opposite direction would prove detrimental to the individual's fitness. The S-SR operator looks for a mate which contains genes whose alleles lie between the corresponding gene's allele bounds in the first parent. The neighbour containing the most such 'desirable' genes is selected as a mate and one offspring is produced which consists of all these desirable genes and the remaining genes of the first parent.

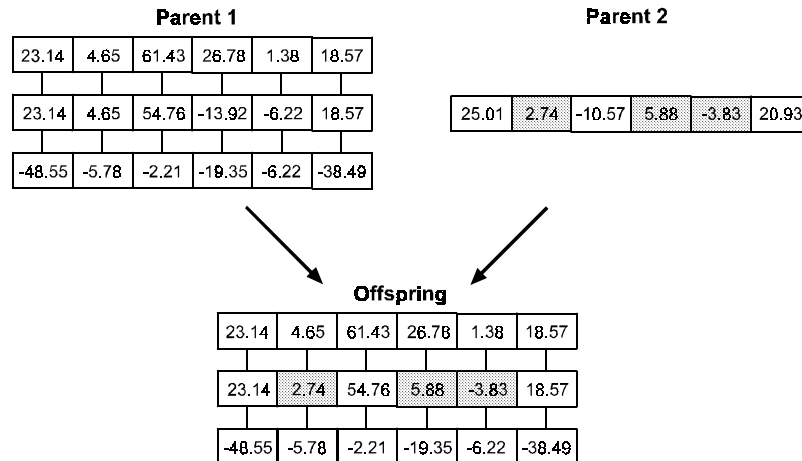


Figure 2: site-specific recombination, as directed by allele bounds

Thus the pseudocode for S-SR selection and crossover is as follows:

**SELECTION:**

```

Neighbourhood = array of six vicinal population members
MostDesirableNo is the greatest number of desirable genes found in any
neighbour so far = 0
CurrentFave is the current neighbour with the greatest number of desirable
genes
CurrentDesirableNo is the current number of desirable genes found in
current neighbour
Sort Neighbourhood in order of increasing fitness
Parent1 = Neighbourhood[1]
FOR i from 2 to NeighbourhoodSize DO
  CurrentDesirableNo = 0
  FOR j from 1 to GenotypeLength DO
    IF jth gene of Neighbourhood[i] > jth lower allele bound AND < jth
upper allele bound of Parent1 THEN
      CurrentDesirableNo = CurrentDesirableNo + 1
      Append value of j to Neighbourhood[i]'s DesirableGenesList
    END IF
  END FOR
  IF CurrentDesirableNo ≥ MostDesirableNo THEN
    MostDesirableNo = CurrentDesirableNo
    CurrentFave = Neighbourhood[i]
  END IF
END FOR
Parent2 = CurrentFave

```

**CROSSOVER:**

```

Child = Parent1
FOR i from 1 to MostDesirableNo DO
  j = value of ith element in Parent2's DesirableGenesList
  ith gene of Child = jth gene of Parent2
END FOR

```

### 3 Experiments

#### 3.1 Functions Used

In order to assess the performance of S-SR in relation to standard one-point crossover, various optimization problems were used, as represented by the following functions.

The first three, to be minimized, are the last in a suite of five functions originally constructed by De Jong [3] and which were intended to represent common difficulties among optimization problems in an isolated manner.

F1 De Jong's F3 has a single optimal value of 0, and is defined by

$$\sum_{i=1}^5 \text{integer}(x_i)$$

for  $-5.12 \leq x_i \leq 5.12$ .

F2 De Jong's F4 is 'noisy': random Gaussian noise is added to its value every time it is evaluated, and is defined by

$$\sum_{i=1}^{30} ix_i^4 + \text{Gauss}(0,1)$$

for  $-1.28 \leq x_i \leq 1.28$

F3 De Jong's F5 has a global minimum of 0.002 - although there are many suboptimal minima - and is defined by

$$0.002 + \sum_{j=1}^{25} \frac{1}{j + \sum_{i=1}^2 (x_i - a_{ij})^6}$$

for  $-65.536 \leq x_i \leq 65.536$

The following two functions, to be maximized, are taken from a set due to Baluja [1]. Both have a global maximum of 100,000 at the origin, where the variable ranges are  $-2.56 \leq x_i \leq 2.56$ . A small constant  $C = 1e-5$  is added to the denominator of the functions to avoid division by zero.

F4 Trial solutions to this function are affected by high epistasis; that is, there exists a high degree of inter-dependence between loci: the variables in the first portions of the solution string have a large influence on the quality of the rest of the solution – small changes in their values can cause large changes in the evaluation of the solution [1].

$$\frac{1}{C + \left| |y_1| + \sum_{i=2}^{100} |y_i| \right|}$$

where  $y_I = x_I$

and  $y_i = x_i + \sin(y_{i-1})$

F5 In contrast to the previous function, trial solutions to this function are affected by low epistasis; that is, there exists a low degree of inter-dependence between loci.

$$\frac{1}{C + \sum_{i=1}^{100} |0.024(i+1) - x_i|}$$

F6 To simulate a multi-peak problem the following constrained function was defined by Keane [5].

$$\frac{\left| \sum_{i=1}^n \cos^4(x_i) - 2 \prod_{i=1}^n \cos^2(x_i) \right|}{\sqrt{\sum_{i=1}^n ix_i^2}} \text{ for } 0 < x_i < 10, i = 1, \dots, n$$

for  $0 < x_i < 10, i = 1, \dots, n$

subject to  $\prod_{i=1}^n x_i > 0.75$

and  $\sum_{i=1}^n x_i < \frac{15n}{2}$

Details of the penalty function used with the above function can be found in [5].

F7 This is the most complex in a set of multimodal functions  $q$  constructed by Corana et al. [2]. These functions, to be minimized, are obtained by defining a regular, rectangular grid in a space  $R^n$  and a set of open, non-overlapping, rectangular subdomains, each centred around the node of the grid. The definition of the function  $q_n(\mathbf{x})$  of  $n$  variables is a rectangular subdomain of  $R^n$  centred at the origin and including several nodes of the grid. The function  $q_n$  is a paraboloid with axes to the coordinate directions except inside the open subdomains mentioned above, where it is constant with a value lower than the lowest value of  $q_n$  computed on the boundary of each subdomain. These subdomains are like a set of ‘holes’ representing local minima of  $q_n$  and introducing strong discontinuities in the test function. The total number of local minima of the test function  $q_n(\mathbf{x})$  is  $10^{5n-1}$ , and for the purposes of this comparison,  $n = 10$ . The absolute minimum lies on the origin and has value 0. (A comprehensive definition of the function can be found in [2].)

### 3.2 The Genetic Algorithm

A geographically distributed GA was used, with the population spread across a two-dimensional toroidal grid of size  $15 \times 15$ , each cell of which contained a single individual. Mating was restricted to small groups of individuals which were generated as follows:

- Select a grid cell at random.
- Build a neighbourhood of six individuals around the current cell by, for each neighbour, generating  $x$ - and  $y$ -distances from the current cell dependent upon a binomial approximation to a Gaussian distribution where  $n = 4$  and  $p = 0.85$ . Thus, for individuals in cells at consecutive distances away from the current cell, the probabilities of being selected are 0.52, 0.37, 0.1 and 0.01. (The direction of the distances – up, down, left or right – are chosen at random.)
- Rank the neighbourhood members according to their fitness.

The selection of parent individuals differed in strategy depending on which method of crossover was to be applied: in anticipation of site-specific recombination, the least fit member of the neighbourhood was chosen automatically as the first parent, and a mate selected according to the selection strategy described in section 2.2. However, if conditions precluded S-SR (that is, if none of the first parent’s neighbours contained desirable genes, or if the number of desirable genes in a neighbour equalled the length of the genotype) standard one-point crossover was used. In this instance, parents were selected according to a linear selection function favouring the fitter individuals. After mating, in both cases, the fittest child replaced a member of the neighbourhood selected according to the inverse of this function.

### 3.3 Mutation

Mutation was applied after crossover with a probability of 0.01, and the experiments were carried out essentially using two different methods with both one-point crossover and S-SR. Method 1 incorporated a fairly standard method of real-valued mutation: in 90% of mutations, a gene was selected randomly, and mutated according to a uniform distribution centred on the current value and of width 10% of the gene range. In the remaining 10%, the selected gene was simply assigned a random value from within its entire range. Method 2 granted genotypes a greater chance of exploring more remote areas of the search space by assigning a randomly selected gene a random value within its range (while resetting the gene's allele bounds to those parameters respectively for S-SR) 100% of the time. However, resetting the allele bounds could mislead S-SR. Recall that the neighbour with the most desirable genes is chosen to be the second parent; with a gene's allele bounds reset, S-SR would now identify any value for that gene's equivalent in a potential mate as desirable. Therefore, it is possible that a neighbour with several desirable genes – plus one misleadingly desirable gene – would be chosen in favour of a neighbour with one less desirable gene, but which ultimately would have proved more profitable. To help rectify this, a third mutation method, 2.1, for S-SR, acted in the same way as the second, but, in addition, it was ensured that when hill-climbing was subsequently applied, it would be applied to the mutated gene first and foremost, thus updating its allele bounds once again in anticipation of a later S-SR operation.

### 3.4 Hill-Climbing

A major concern while designing the S-SR operator was the potentially enormous expense incurred due to repeated function evaluations during the hill-climbing stages: for example, using one-point crossover with genotypes of any length for a run of 300 generations, the total number of evaluations would be 67500; however, using S-SR with genotypes of length 100 (and assuming hill-climbing is applied to every gene), the number of evaluations for a single run potentially could range from 6817500 to 13567500. Therefore in order to reduce this expense, hill-climbing was only applied to specified percentages of randomly selected genes; the percentages used were 50%, 30%, 10%, 5% and 1%.

In order that the comparison was fair in terms of the amount of function evaluations, all runs were set to last the number of evaluations equal to the minimum number of evaluations possible in a run of 300 generations using S-SR, with hill-climbing applied to 50% of a genotype, i.e.  $((1 + 50) \times 225) \times 300 = 3442500$  for a genotype of length 100;  $((1 + 25) \times 225) \times 300 = 1755000$  for a genotype of length 50, and so on.


### 3.5 Local Search

In addition to comparing S-SR with simple one-point crossover, a subsequent set of runs were carried out in order to compare S-SR performance with a GA using one-point crossover which was also boosted by local search. The method of local search was shown to perform well in studies such as [7, 8] and comprised 200 mutations applied to the current fittest individual after every 225 breeding cycles. (A breeding cycle includes hill-climbing during S-SR runs.)

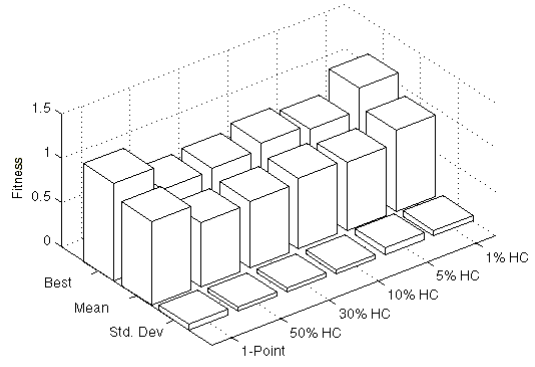
## 4 Results

De Jong published his suite of test functions in 1975 [3], and they have been used extensively ever since, becoming a standard test bed. However, it is interesting to note how the increase in computing power and GA efficiency has rendered trivial what once represented common difficulties among optimization problems: while comparing the one-point and site-specific operators, both implementations invariably found the optimum during the first few generations (if not during the initialisation of the population) for all of the three De Jong problems used. Therefore this section will concentrate only on results obtained for the other four functions.


Differences for both one-point and S-SR (see figures 3 and 4) between the standard mutation type, method 1, and the less constrained method 2 showed that on average the latter yielded slightly better results. (Figure 5 shows that, for S-SR, method 2.1 performed best.) Only on F7, and using S-SR, did method 1 yield substantially better results.




(a) F4



(b) F5

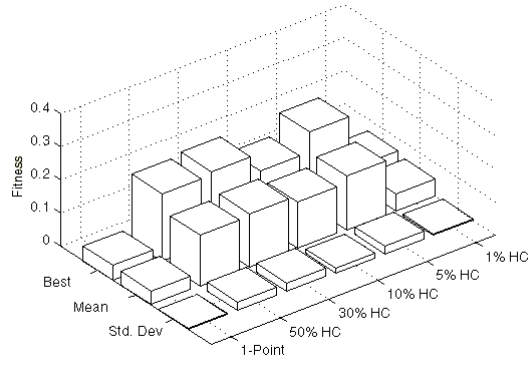


(c) F6




(d) F7


Figure 3: results for runs using mutation method 1




(a) F4



(b) F5



(c) F6



(d) F7

Figure 4: results for runs using mutation method 2

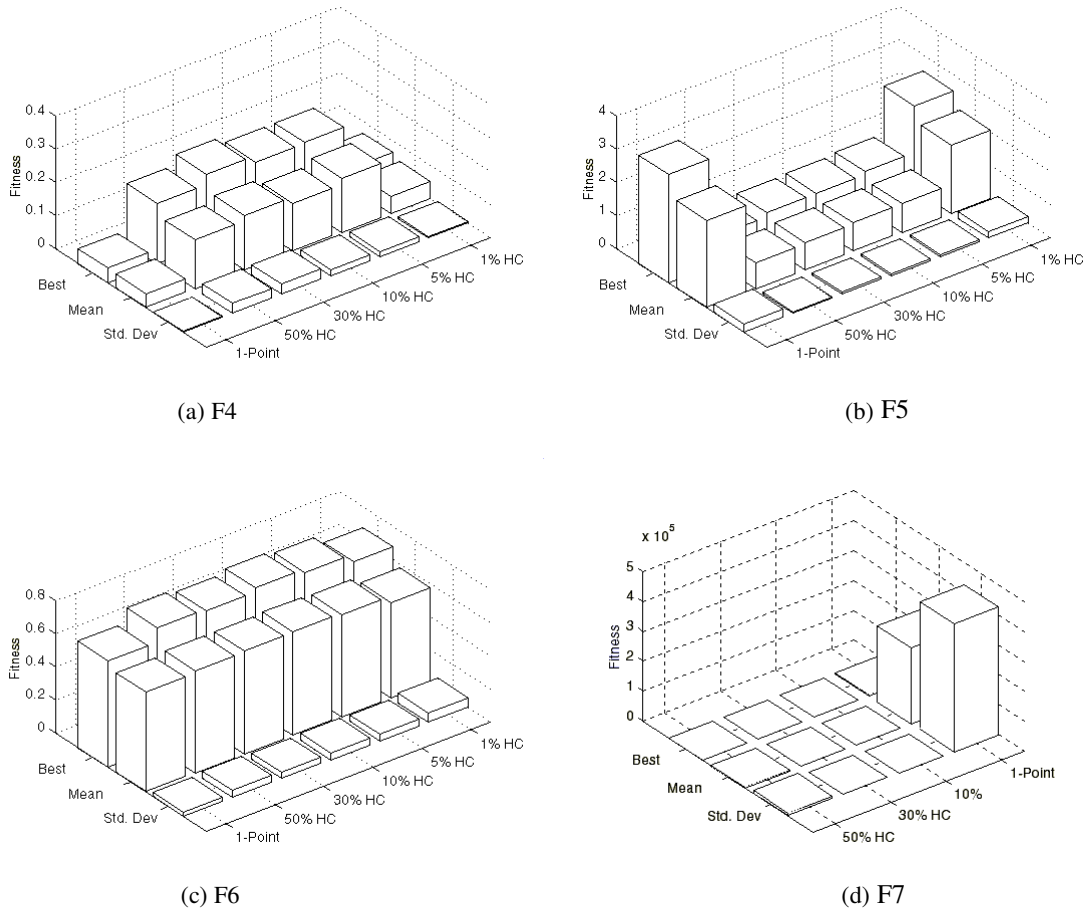


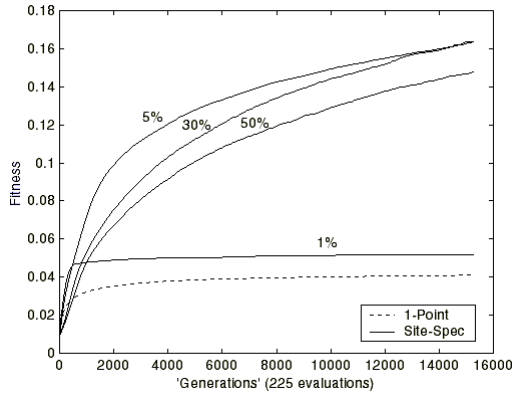
Figure 5: results for runs with S-SR using mutation method 2.1

In comparison with one-point crossover, S-SR – with the appropriate amount of hill-climbing and, typically, mutation type 2.1 – performed better on all functions except F5: the best mean achieved by S-SR was 2.06 as opposed to 2.63 by one-point. In comparison with alternative strategies, S-SR improved on results detailed in [1] with regard to quality of solutions and the speed with which they were found for F4. On F5, S-SR only performed worse than a multiple-restart stochastic hill-climbing method (MRSH) when incorporating Gray encoding. On F6, S-SR proved competitive with results detailed in [5] which were achieved by a GA which incorporated elitism and niching. On F7, S-SR achieved significantly better results than are given in [2].


Save for the occasional anomaly, a fortuitous pattern was observed with regard to the amount of hill-climbing applied during S-SR runs: in general, performance – in terms of both the quality of solutions and the speed with which they were found – improved as the amount of hill-climbing applied, and therefore the number of function evaluations, was reduced (although such evidence for runs with F6 is more ambiguous). This is illustrated in figure 6; the percentage of genes to which hill-climbing was applied in a particular run is included where convenient.

The addition of local search improved significantly the performance of runs using one-point crossover, except on F4. However, S-SR by itself still beat these runs for two of the functions, F4 and F7; adding local search to S-SR improved its performance even more and S-SR then still only performed worse than one-point with local search on one function, F5. Figure 7 illustrates runs using local search with both one-point crossover and S-SR (mutation type 2.1) respectively, averaged over fifty runs as before. A pattern of behaviour similar to that of the first set of runs was evident, although the inclusion of local search appeared to slow down convergence, particularly with functions F4, F5 and F7. However, with regard to F5, whereas typically solutions improved much faster using S-SR than one-point at the beginning of a run (until subsequently being overtaken) during the second set of runs, one-point invariably performed better than S-SR from the outset when local search was applied. Figure 8 illustrates results for runs using local search.







(a) F4




(a) F4




(b) F5



(b) F5



(c) F6



(c) F6

Figure 6: average performance of ordinary runs

Figure 7: average performance of runs with local search

## 5 Discussion

On examining the results, particularly those gathered for functions F4 and F5 given their opposing characteristics, it is encouraging to see that S-SR performs comparatively well on problems with high epistasis – problems which both GAs and hill-climbing traditionally find difficult. The repeated application of hill-climbing undoubtedly will assist a GA to some extent, as traditional hybrid approaches illustrate. However, regardless of any benefits reaped from hill-climbing per se, clearly, the rate of solution improvement using S-SR is very often much faster than with one-point, particularly at the beginning of a search. Such an assertion is reinforced by recalling that the reduction in the hill-climbing percentage invariably increases this speed

significantly. Nevertheless, there would appear to be distinct optimum levels of hill-climbing for the different functions: figure 6 illustrates that although reducing the amount of hill-climbing invariably increases the speed of solution improvement in the early stages of the search, sometimes runs with the least amount of hill-climbing ultimately are not the best performers. Similarly, on F4 (the first of the two Baluja functions used) the optimum hill-climbing percentage appears to be around 5%, whereas percentages of 1% and 10% elicited poor results. Conversely, on F5 (the sister Baluja function) a hill-climbing percentage of 1% elicits a performance far better than those using higher percentages. Moreover, either too low a level of hill-climbing or too high a level of hill-climbing results in an increased rate of convergence.

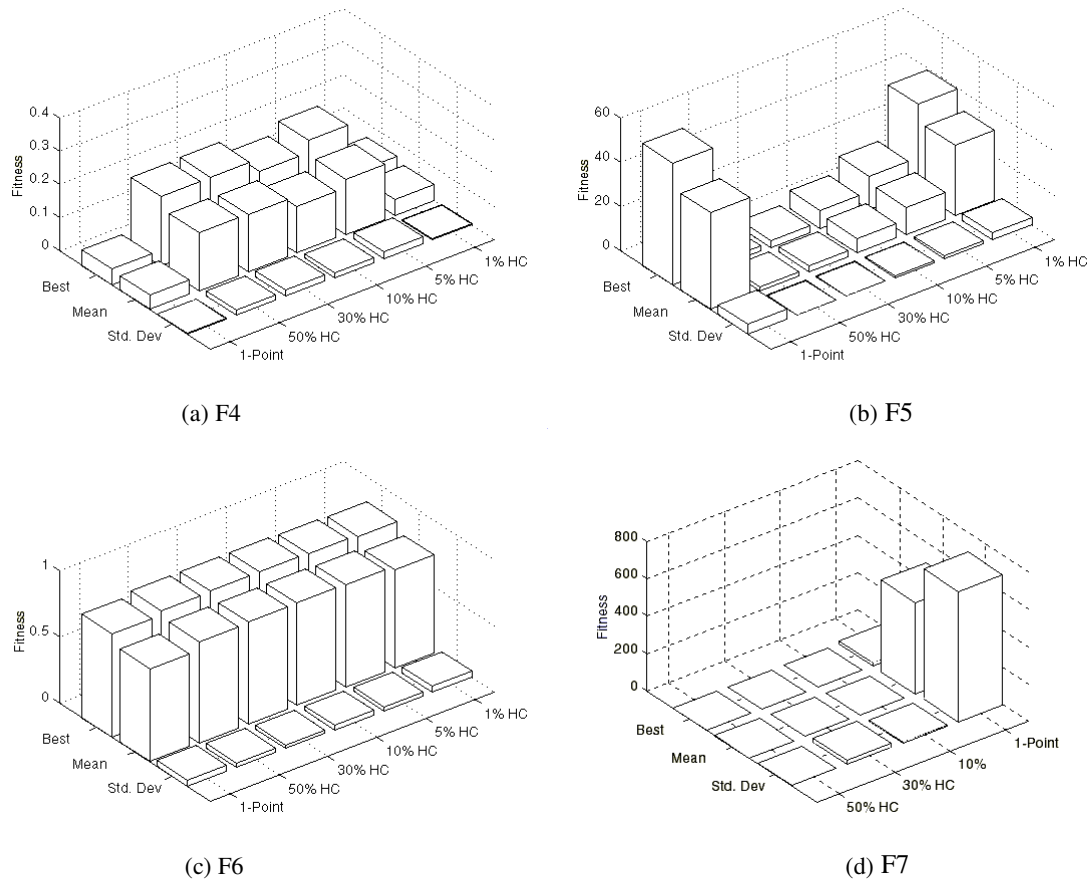


Figure 8: results for runs using local search

It would appear that, on the functions where S-SR beats one-point, the improved performance is due to there being a much greater number of constructive crossover operations, per generation, during the early stages of the search. (For the purposes of this paper, a ‘constructive crossover’ is defined as one that yields at least one offspring which is fitter than either parent; one generation is taken to equal 225 breeding cycles.) Indeed, with the exception of F6 (where the percentage of crossover operations per generation remained high throughout runs of 300 generations), when one-point is used the percentage of constructive operations falls to a negligible – if not non-existent – level in a very short time. In contrast, a relatively high percentage of constructive operations is maintained for longer when S-SR is used; furthermore, constructive operations (albeit a small number) is evident for the duration of a run. This is illustrated using F7 in figure 9. It seems likely that the S-SR selection strategy plays an important part in the success of the operator. GAs have always been based on the concept ‘survival of the fittest’, which implicitly dictates that the least fit members of a population are left to die out. Conversely, the best S-SR performance is observed when the least fit member of a neighbourhood is selected and resuscitated via the donation of genes from a fitter neighbour.

Perhaps it is unsurprising that more constructive crossovers are observed when the S-SR strategy is to take an unfit individual and make it fitter, but one possible explanation may be that fit schemata concealed within unfit individuals, which normally would be overlooked by selection, instead are being exploited. Indeed, this would seem likely if the least fit individual mates with its fittest neighbour to produce offspring fitter than either of its parents (thereby fulfilling the axiom which dictates that ‘opposites attract’!).

That constructive crossover operations are still being carried out even at the very end of a run perhaps indicates a slowing of convergence – despite the acceleration of improvement, and that, intuitively, the nature of the site-specific operator suggests an *increased* rate of convergence. Another possibility may be that S-SR makes the most of small differences between a converged population; or that, in a population of converged fitness, one individual benefits from receiving a particular gene from another individual which is of similar fitness but which inhabits a different area of the search space.




Figure 9: amount of constructive crossover operations

## 6 Conclusion

The immediate objective of designing an operator which attempts to increase the probability of affecting a constructive crossover, thereby increasing the power of a GA, has been fulfilled: it has been shown empirically that using site-specific crossover on some (currently!) difficult functions will accelerate optimization to a considerable degree as well as usually achieving fitter solutions than are reached when one-point is used. However, only continuous function optimization problems have been addressed so far – a set of problems which constitute only a small subset of those to which GAs are routinely applied. Therefore, the operator must be adapted, if necessary, to tackling such combinatorial tasks such as job-shop scheduling and the travelling salesman problem, or even neural network weight optimization. In addition, an analytical investigation of the operator should be carried out in the hope of producing a more complete explanation of the overall rôle of crossover in evolutionary algorithms.

## References

- [1] S. Baluja, *An Empirical Comparison of Seven Iterative and Evolutionary Function Optimisation Heuristics*, Internal Paper CMU-CS-95-193, School of Computer Science, Carnegie Mellon University (1995).
- [2] A. Corana, M. Marchesi, C. Martini, and S. Ridella, Minimizing Multimodal Functions of Continuous Variables with the ‘Simulated Annealing’ Algorithm, in *ACM Transactions on Mathematical Software*, Vol. 13, No. 3, September 1987, pages 262-280.
- [3] K. DeJong, *An Analysis of the Behaviour of a Class of Genetic Adaptive Systems*, PhD thesis, University of Michigan (1975).
- [4] J. J. Grefenstette, Optimization of Control Parameters for Genetic Algorithms, *IEEE Transactions on Systems, Man & Cybernetics* 16, No. 1 (1986).
- [5] A. Keane, *Experiences with Optimisers in Structural Design*. in Proceedings of the First International Conference on Adaptive Computing in Engineering Design and Control, University of Plymouth (1994).
- [6] R. F. Leach, *Genetic Recombination*, Blackwell Science Ltd. (1996).

- [7] M. McIlhagga, P. Husbands and R. Ives. A Comparison of Optimisation Techniques for Integrated Manufacturing Planning and Scheduling, in H.-M. Voigt, W. Ebeling, I. Rechenberg and H.-P. Schwefel, editors, *PPSN IV*, Springer (1996).
- [8] M. McIlhagga, P. Husbands and R. Ives, A Comparison of Search Techniques on a Wing-Box Optimisation Problem, in H.-M. Voigt, W. Ebeling, I. Rechenberg and H.-P. Schwefel, editors, *PPSN IV*, Springer (1996).
- [9] M. Mitchell, *An Introduction to Genetic Algorithms*, MIT Press (1998).
- [10] J. D. Schaffer, R. A. Caruana, L. J. Eshelman, R. Das, A Study of Control Parameters Affecting On-Line Performance of Genetic Algorithms for Function Optimization, in J. D. Schaffer, ed., *Proceedings of the Third ICGA*, Morgan Kaufman (1989).
- [11] G. Syswerda, Uniform Crossover in Genetic Algorithms, in J. D. Schaffer, ed., *Proceedings of the Third ICGA*, Morgan Kaufman (1989).
- [12] C. Thornton, The Building Block Fallacy, *Complexity International Vol. 4* (1997).
- [13] K. Vekari, C. Clack, Selective Crossover in Genetic Algorithms: An Empirical Study, in A. E. Eiben, T. Bäck, M. Schoenauer and H.-P. Schwefel, eds., *PPSN V*, Springer (1998).