Appears in "Proceedings of the Fifth International Conference on Genetic Algorithms", Ed: S. Forrest, Morgan Kaufmann (San Mateo), 1993.

A Study in Set Recombination

Nicholas J. Radcliffe & Felicity A. W. George
Edinburgh Parallel Computing Centre
University of Edinburgh
King's Buildings
EH9 3JZ
Scotland

Abstract

A family of problems for which the solution is a fixed size set is studied, using fitness functions with varying degrees of epistasis. An empirical comparison between a traditional crossover operator with a binary representation and a penalty function, and the representation-independent Random Assorting Recombination Operator (RAR) is performed. RAR is found to perform marginally better in all cases. Since RAR is a parameterised operator, a study of the effect of varying its parameter, which can control any trade-off between respect and assortment, is also presented.

1 Introduction

Throughout the short history of genetic algorithms there has been a creative tension between two quite distinct schools of thought on the subjects of representations and operators. The fundamental disagreement concerns the relative merits of simple, low cardinality representations used with standard recombination operators—typically "n-point" and more recently uniform crossover[1]—and more complex, problem-specific recombination operators applied to more diverse representations. The "traditional" ("binary is best") view derives primarily from the observations about implicit[2] parallelism and schema processing by Holland (1975), who used a simple counting argument to suggest that low cardinality representations should be more powerful. While it is probably fair to say that this has been the dominant view amongst workers in genetic algorithms (as opposed to Schwefel's *Evolutionstrategie* school) there have always been approaches based on non-standard, high-cardinality representations together with "custom" genetic operators. While these approaches have arguably had less theoretical under-pinning, it is far from experimentally established that they are inferior. Davis (1991), for example, has argued forcefully that on "real-world" problems he invariably gains superior results with "natural" representations and custom operators, while Goldberg (1990c) has defended the use of binary representations and traditional operators for problems in real parameter optimisation with the development of his "theory of virtual alphabets". (It should be stressed that there are dozens, if not hundreds of papers which discuss and form part of this debate, and there is no attempt here to undertake the Stachanovite task of cataloguing or reviewing them all.)

In addition to the countless empirical studies which pertain to the debate, there have been a few attempts to re-examine and generalise schema analysis and the notion of implicit parallelism, notably those by Goldberg & Lingle (1985), Wright (1990), Antonisse (1989), Radcliffe (1990, 1991), Vose (1991), Vose & Liepins (1991) and Eshelman & Schaffer (1992). In particular, Antonisse argued for a re-interpretation of schemata under which the counting argument actually suggests that higher cardinality representations will exhibit greater implicit parallelism than their lower cardinality counterparts, while Radcliffe has developed a comprehensive generalisation of schema analysis called forma analysis. The latter stresses the key rôle of the fitness variance of schemata (and their generalisation, *formae*) in determining the ability of the genetic algorithm to exploit performance correlations in the search space and provides a framework for constructing problem-specific versions of various "universal", representation-independent operators (Radcliffe, 1993).

Forma analysis has been applied, at a theoretical level, to problems for which the solution is a set of fixed size (Radcliffe, 1992). This paper presents a simple empirical comparison between the performance of the specialised representation and recombination operator suggested by forma analysis and that of a traditional binary representation with a conventional crossover operator, over a range of problems ranging from simple non-epistatic problems through to harder problems with parameterised degrees of epistasis. The constraint of fixed set size will be seen to necessitate the use of some mechanism to restrict (or encourage) the set to be of fixed size, and for this reason a penalty function has been used with the traditional operator for this study.

[1] See, for example, Syswerda (1989) for a description of these
[2] *née intrinsic*

2 Approaches to Set Recombination

2.1 Statement of problem

The problems which will be considered in this paper all consist of searches for a particular subset $S^\star$ of size n from a larger universal set $\mathcal{E}$ of size N ($N > n$), and in each case the search space will be considered to be the set of all subsets of $\mathcal{E}$ of the given fixed size n. In particular, the universal set, $\mathcal{E}$, will always be considered to be the set of integers modulo 120

$$\mathcal{E} \triangleq \mathbb{Z}_{120} \triangleq \{0, 1, \ldots, 119\} \tag{1}$$

and the subset sought will in each case be of size 60 and contain the lowest 60 numbers, so that the optimum $S^\star$ will be

$$S^\star = \mathbb{Z}_{60} = \{0, 1, \ldots, 59\} \subset \mathbb{Z}_{120}. \tag{2}$$

Various fitness (utility) functions will be used to control the degree of epistasis in the problem.

2.2 Traditional approaches

Regardless of any detailed aspects of the fitness function used, this problem immediately poses a problem if a traditional linear string representation is to be used. The size of the given search space—the set of all subsets of size n drawn from $\mathcal{E}$—is

$$\binom{N}{n} \triangleq \frac{N!}{n!(N-n)!}, \tag{3}$$

which in the present case, with $n = 60$ and $N = 120$ is approximately 10^{35}. The most naïve approach might be to use a string of length n, each of whose characters could take on any of the N values in $\mathcal{E}$. This, however, would given enormous scope for repetition of elements and would enlarge the search space to size n^N, some 10^{213} with the suggested values.

This approach would almost certainly, after a moment's thought, be replaced by a binary string of length N, with one element corresponding to each member of the universal set, a one indicating inclusion of the element in the subset, and a zero indicating its exclusion. If the number of ones in the string were kept fixed at n, this would be a faithful representation, with exactly one binary string corresponding to each element of the search space. Unfortunately this is not guaranteed with standard crossover operators. In practice, one of two standard approaches would probably be adopted. The more common would be to add a (negative) penalty (Richardson *et al.,* 1989) to the fitness function to reduce the performance of solutions which contain too many or too few elements (and thus fall outside the true search space, into an embedding space). While technically the size of the search space available using this approach would rise to 2^N (some 10^{36} with the given values), in practice the hope would be that the penalty function would restrict the search to solutions of sizes close to n. There is then a trade-off between exploring illegal solutions in the embedding space and a possible gain in navigability of the underlying search space. This forms the "traditional" approach which is examined in this paper, trenchant criticism from, for example, Michalewicz & Janikow (1991) notwithstanding.

The alternative "traditional" approach would be to employ a post-recombinative repair function using some form of directed mutation to produce a child with the correct number of elements (cf. Harp *et al.,* 1989). This approach will not be considered directly in this paper, but the careful reader will be able to see that the approach of forma analysis taken here can be interpreted as being related to this technique.

2.3 Forma analysis

The non-traditional approach considered here is to use the *random assorting recombination* operator (RAR) developed and justified in the general context by Radcliffe (1993) and specialised for application to set problems in Radcliffe (1992). A *forma* is an equivalence class of some equivalence relation defined over the search space (or equivalently a partition of the search space). The approach motivated by *forma analysis* is to choose equivalence relations which induce formae that the experimenter believes group together solutions having similar fitness. Such equivalence classes (formae) will not only obey the schema theorem (p. 102 of Holland, 1975, with suitable changes to the disruption coefficients) but will tend to result in relatively accurate fitness estimators.

The formae (generalised schemata) which motivate the approach taken in this paper specify subsets of elements which their members must include, and elements which they must exclude. For example, a particular forma for sets of size 3 chosen from $\mathbb{Z}_6$ might consist of those sets which include 0 and exclude 1 and 3

$$0\bar{1}\square\bar{3}\square\square = \big\{\{0,2,4\}, \{0,2,5\}, \{0,4,5\}\big\}, \tag{4}$$

where $\square$ is a (redundant) "don't care" character.

Radcliffe has characterised recombination operators according to two key criteria with respect to a given set of formae, called *respect* and *assortment.* A recombination operator is said to *respect* a given set of formae if children generated by the operator are always members of all those formae to which both parents belong. In the present context, respect can readily be seen to amount to the requirement that all those elements common to the two parents be present in all children that a recombination operator can produce. Thus given parents $\{0,2,4\}$ and $\{0,1,5\}$ respect allows only solutions chosen from

$$\Sigma\big(\{0,2,4\}, \{0,1,5\}\big) \equiv \big\{\{0,1,2\}, \{0,1,5\}, \{0,2,4\}, \{0,2,5\}\big\} \tag{5}$$

which is called their *similarity set.* (Informally, one can think of respect as requiring that if both parents have blue eyes and curly hair, all children produced by recombination must share these properties.)

In contrast, assortment requires that a recombination operator be capable of generating a child in the intersection of two formae ξ_1 and ξ_2 provided that one parent is in the first forma, the other is in the second, and that the intersection is non-empty. Using the same example for the present case, the first parent $\{0, 2, 4\}$ is a member of the forma $\square\square 2\square 4\square$, and the second, $\{0, 1, 5\}$, is a member of $\square 1\square\square\square\square$, so assortment requires that a recombination operator be capable of generating a child in the (non-empty) intersection

$$\square\square 2\square 4\square \cap \square 1\square\square\square\square = \square 12\square 4\square. \qquad (6)$$

(Informally, assortment requires that if one parent has blue eyes, and the other has curly hair, then it must be possible for recombination to produce a blue-eyed, curly-haired child, provided that these characteristics are compatible.)

Observe, however, that in the current example no member of the similarity set (equation 6) is included in the forma $\square 12\square 4\square$, so that no recombination operator can achieve both respect and assortment for these formae. Formae which cannot be respected and assorted simultaneously are termed *non-separable*. The occurrence of non-separable formae is not restricted to set problems, but is in fact quite wide-spread. For example, the natural formae for the travelling salesrep problem and neural network topology optimisation are all non-separable. It is therefore important to find a sensible approach for such problems.

2.4 Respect vs. assortment

Respect restricts the range of children which a recombination operator may produce so that features[3] common to the two parents are guaranteed to be transmitted to their children. It is a weaker form of a stricter notion of *gene transmission*, which requires that each of a child's alleles is inherited from at least one of its parents.[4] This seems like such a natural property, encompassing some of the need to exploit information already gathered, that it may be hard to see a reason for ever using non-respectful recombination operators.

Similarly, however, assortment merely requires that when the parents contain all the genetic material required to construct a given child, that the recombination be powerful enough to construct the relevant child. Again, it may be hard to see how non-assorting recombination operators can be used usefully.

One way to address this apparent *impasse* is to build a parameterised operator which allows controlled sacrifice of respect in pursuit of assortment. The *random assorting recombination* operator proposed by Radcliffe (1993, 1992) achieves exactly this.

[3] expressed as membership of particular formae

[4] It is important to emphasize that only recombination is subject to this restriction: mutation is included explicitly to violate respect.

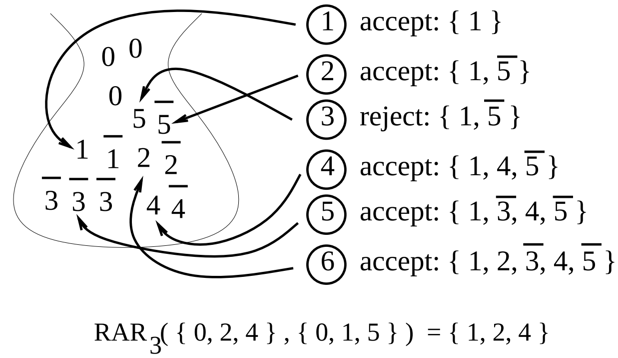

Figure 1: An example of how RAR can produce a child in the required intersection. Bars denote excluded elements.

2.5 Random assorting recombination

The *random assorting recombination* operator takes a parameter w and is denoted RAR_w. In its simplest conception this is parameter is an integer, and the operation of the operator may be described as follows:

Given a positive integer parameter w, and two parents A and B, each of size n,

1. Place w copies of each element common to the two parents in a bag $\mathcal{G}$, together with w copies of "barred elements", which are present in neither parent. Thus if their only common element is 0, and they commonly exclude only 3, and if $w = 3$, then three copies of 0 and $\bar{3}$ would be placed in the bag.
2. Place one copy of each element in only one of the parents, together with one copy of its barred counterpart.
3. Repeatedly draw from the bag, in a random order, barred and normal elements. Normal elements are included in the child, whereas barred elements are excluded, in both cases subject to the primacy of earlier decisions (i.e. an element previously included cannot later be excluded, and vice-versa), and a requirement that elements not be included which would cause violation of the fixed set size.
4. This process continues until either the child is fully specified, or the bag is empty. Should the bag empty before the child is fully specified, remaining elements and barred elements are assigned at random subject to the constraints listed above.

An example of the operation of RAR_3 is given in figure 1. It is easy to see that

- For all positive values of $w < \infty$, RAR_w assorts the set formae, because it is possible to construct any child which includes only elements available in its parents.
- Lower values of w lead to more assortment, in the sense that it becomes more probable that arbitrary mixtures of the parents genes will be generated.

- Higher values of w lead to less violation of respect and less assortment. If w were notionally set to ∞, it is clear that there would be full respect. In this case RAR_∞ actually reduces to a simpler operator called *random respectful recombination*, (R^3) which includes all elements common to the two parents, excludes all those which neither contains, and randomly chooses the remaining elements. Clearly R^3 does not assort the set formae, since these are non-separable.

It should be stressed that the description above is explanatory only, and that more efficient implementations than that described are possible. Moreover, since in practice w is just the probability of selecting elements common to the two parents relative to those unique to one parent, there is no difficulty about allowing w to take on any positive value, i.e. the restriction to integers is unnecessary.

For a description of the operation of RAR in the general context, and a more detailed motivation for it, the reader is referred to Radcliffe (1993).

3 Test Problems

3.1 Non-epistatic problem

The simplest problem which will be considered is completely non-epistatic, and simply awards one point to a solution for each element which it contains which is also in the optimum $S^\star$. (In all cases the universal set is $\mathbb{Z}_N$ with $N = 120$, and the optimum sought is $\mathbb{Z}_n$, with $n = 60$.) Thus this fitness function, which will be described formally as

$$\mu_1 : \mathbb{P}(\mathbb{Z}_N) \longrightarrow \mathbb{Z} \tag{7}$$

(where $\mathbb{P}$ denotes "the set of all subsets of") is given by

$$\mu_1(A) = |A \cap \mathbb{Z}_n|. \tag{8}$$

The penalty function which will be applied when using the binary representation simply deducts one for every surplus or missing element, and is given by

$$P : \mathbb{P}(\mathbb{Z}_N) \longrightarrow \mathbb{Z} \tag{9}$$

with

$$P(A) \triangleq -\big||A| - n\big|. \tag{10}$$

Clearly μ_1 represents a completely non-epistatic problem, and is in some ways analogous to the common "one-max" problem (or "counting ones" problem, e.g. Vose & Liepins (1991)).

3.2 Epistatic problems

Probably the most widely discussed epistatic problems in the literature on genetic algorithms are the deceptive problems introduced by Goldberg (1990a, 1990b), together with the epistatic members of De Jong's test suite (DeJong, 1975). None of these can be applied directly since they are all defined with respect to arbitrary binary strings,

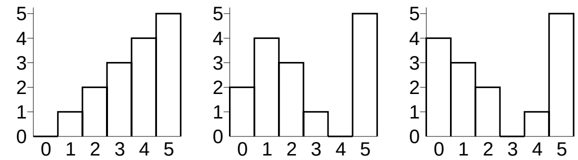

Figure 2: The left figure shows the credit assigned for getting k elements of a sub-range of length 5 correct for a non-epistatic problem. The central and right figures show two possible arrangements in the epistatic case. Notice that getting all elements correct always results in the maximum score.

though it would be possible to build analogues of various deceptive problems. This, however, is not the aim here: rather the intention is to perform a simple investigation of the effect of epistasis on the performance of the operators under consideration. Instead, therefore, a novel set of functions of sets with varying degrees of random epistasis will be introduced and explored. The basic idea will be to introduce a grouping of the elements into equal-length sub-ranges, such as 0–4, 5–9 and so forth. The credit assigned for getting k of these elements correct in the non-epistatic case would be k, as shown in figure 2. In the epistatic version, the credit for getting k elements correct will be shuffled randomly with respect to the non-epistatic case, subject only to the constraint that getting all elements correct will always result in a maximum score (figure 2). Clearly the longer the sub-blocks considered, the more epistasis will be allowed, and in general the fitness function μ_k $(k > 0)$, will be defined on the basis of subranges of length k, this definition coinciding with the earlier one for μ_1. Note that the random shuffling is *independently* chosen for each sub-range of length k.

4 The Experimental Framework

All experiments were conducted using the Reproductive Plan Language (RPL) Interpreter described in Russo (1991) and Jones (1992). Serial experiments used a non-generational reproductive plan with a population size of 100, a mutation rate (definitely exchanging a member of the set with a non-member after crossover in the case of the set representation, and definitely flipping a bit in the case of the binary representation) of 0.01, recombination with probability 1.0, binary tournament selection with parameter 0.6 (no replacement), certain replacement of worst individual, RAR with various weights, and uniform crossover with parameter half. In all cases the all solutions in initial population are of the correct size (60), i.e. when the penalty function is used, the penalty is zero for all members of the initial population.

It is important to understand the reasons for choosing uniform crossover as the "traditional" operator to compare against. Eshelman *et al.* (1989) have analysed crossover operators and shown that the differences between n-point crossover and uniform crossover may be fully described by considering two factors, which they term "positional" and "distributional" bias. Uniform crossover has no positional bias because the probability of any set of alleles being transferred from a parent to its child is independent of their location on the genome (i.e. it is a function of the number of genes considered only). In contrast, the probability that a group of alleles will be transmitted *en masse* by n-point crossover operators is strongly dependent on their location on the genome. The epistatic functions introduced in this paper have strong positional dependence which n-point crossover could exploit, but such exploitation is not of interest in this paper, for similar functions which did not group together the sub-ranges considered could equally easily be constructed. Similarly, a variant of RAR which had positional bias could easily be constructed if exploitation of the such dependence of the epistatic functions were desired.

An operator is said to exhibit *distributional bias* if the probability distribution for the quantity of material taken from one parent is non-uniform. One-point crossover has no distributional bias (assuming that the cross point is chosen uniformly) while the distribution for uniform crossover is binomial. Spears & DeJong (1991) have provided evidence that biasing the operators to take more genetic material from one parent than the other tends to improve their performance, but again this is not the purpose of this study.

Thus uniform crossover has been chosen for this study because its positional and distributional characteristics are the same as those of RAR.

5 Experimental Results

Two sets of experiments were performed using the non-epistatic evaluation function μ_1 and five epistatic variations. The degree of epistasis was increased by increasing the length of the blocks whose values were randomly permuted through the values 2, 3, 4, 6, and 10. It should be noted that the degree of epistasis for even the first of these values is quite high, and it is unclear whether typical "real-world" problems would be likely to exhibit the degree of epistasis represented by the harder problems studied here (though one might expect other sorts of difficulties to arise in real-world situations).

Graphs of all results are shown in figures 3 and 4. All results are averaged over twenty runs on the different instances of randomly-generated epistatic functions, and show the best solution in the population at the given generation. The error bars indicate standard deviations. In the case of the results with uniform crossover, the fitness includes the penalty, but in all cases the best solution in the final population was of the correct size (60).

For the first set of experiments (figure 3) the performance of RAR_2 with the 'natural' weight of 2 is compared with uniform crossover with a penalty function. Although in most cases the standard deviations overlap, in every case from the non-epistatic to the most epistatic fitness function, RAR marginally out-performs uniform crossover with the penalty function, in terms of the quality of its solutions and usually in its speed also.

The second set of experiments compares the performance of RAR_w for weights $w = 4$, $w = 2$, $w = 1$, $w = 0.5$, $w = 0.25$ and $w = \infty$ ($\equiv \mathrm{R}^3$). Again the graphs include the performance of uniform crossover for reference, but this time standard deviations are not shown because of the density of information on the graphs. The following general observations can be made:

- The lower values of w generally achieve the best ultimate performance, but they are relatively slow to achieve this.
- Conversely, the higher values of w (especially as represented by $\mathrm{R}^3 \equiv \mathrm{RAR}_\infty$) tend to perform better earlier in the runs.

This is much as would be expected: the higher weights lead to more respect and thus greater exploitation at the expense of less thorough exploration of the space. For the problems studied the ultimate benefits of the lower weights and greater exploration can take as much as an order of magnitude longer to be achieved than the fast gains from respect.

6 Conclusions

It should be emphasized that the experiments performed are far from exhaustive or definitive. A simple penalty function was used with uniform crossover, and the epistatic test functions used for set optimisation are contrived and all of a single class. Nevertheless, the following conclusions may fairly be drawn. First, the RAR operator, in its simplest form, clearly works fairly effectively for the fixed-size set problems studied. Since previous work published on forma analysis has been largely theoretical, this in itself is significant. Obviously it would be desirable to study a wider range of set problems, particularly "real-world" examples.

Secondly, although the gains are not dramatic, RAR_2 out-performs uniform crossover with a simple penalty function on a family of problems in which the solution is a fixed-size set, using fitness functions which vary from highly-epistatic to non-epistatic. This therefore provides some evidence that forma analysis is capable of leading to effective genetic search in problem domains not well suited to conventional binary representations. Since RAR is a universal operator, in the sense that it can be used in any domain provided only that appropriate formae have been defined over the search space, this is potentially significant. Preliminary work in other problem domains such as the TSP and a radar sensing problem, not yet published, have also been encouraging and detailed studies will appear.

Thirdly, it has been established that the weight parameter w which RAR_w uses can indeed be tailored to control the

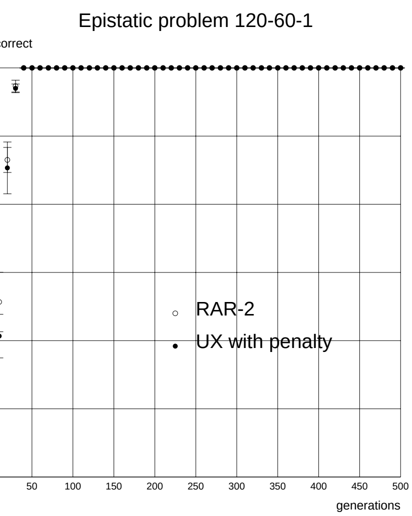
Epistatic problem 120-60-1
orrect
RAR-2
UX with penalty
50
100
150
200
250
300
350
400
450
500
generations

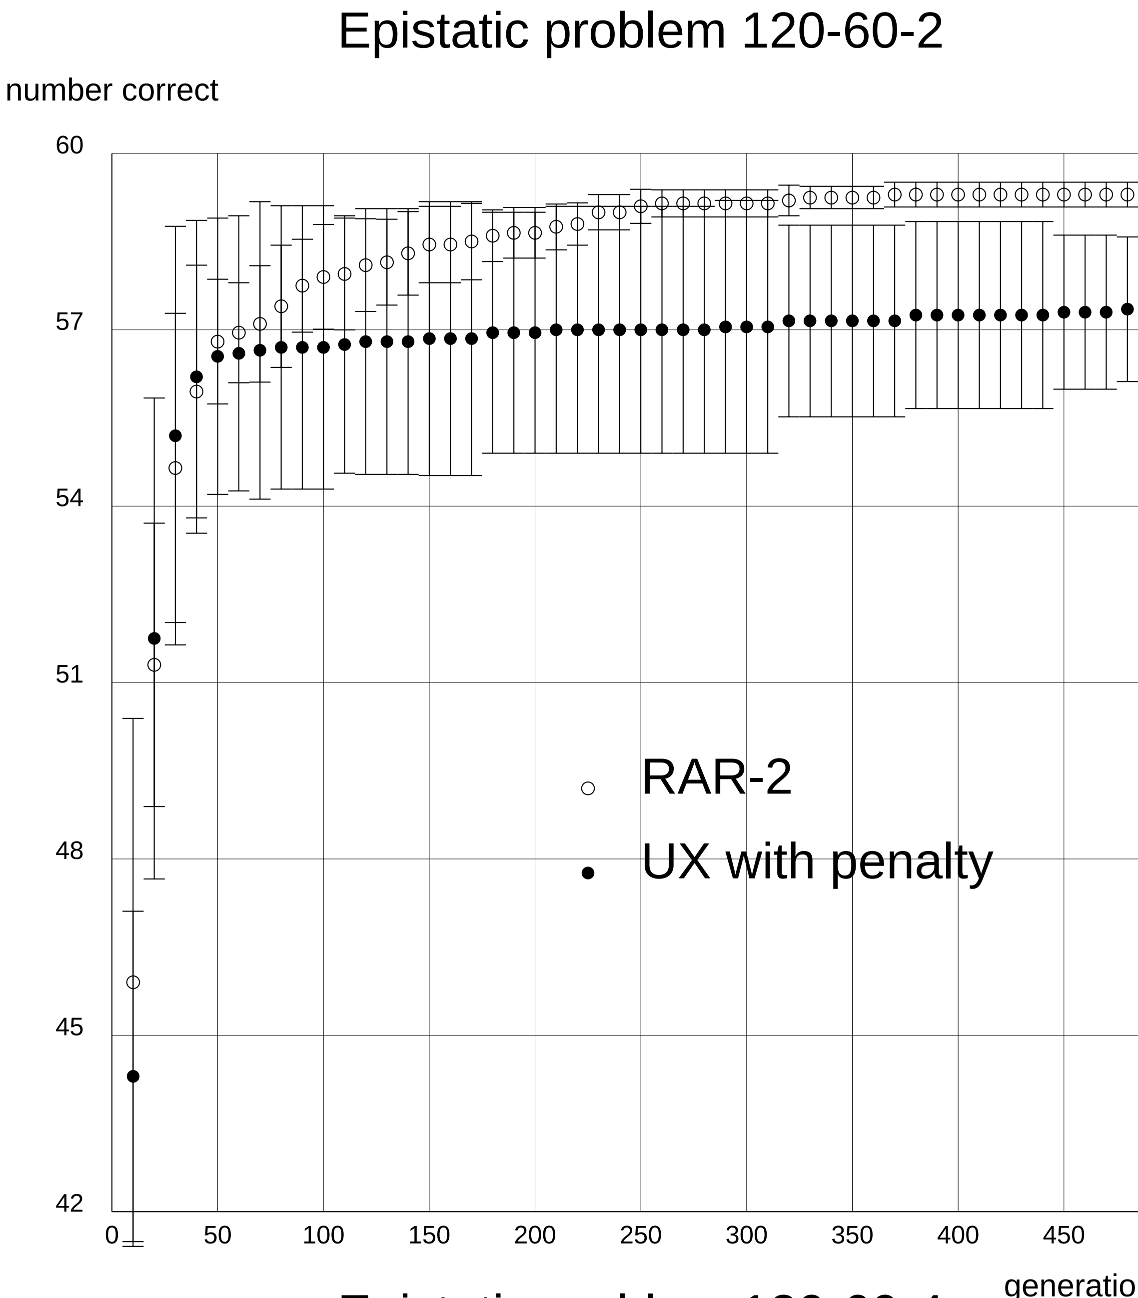
Epistatic problem 120-60-2
number correct
60
57
54
51
48
45
42
0
50
100
150
200
250
300
350
400
450
RAR-2
UX with penalty
generatio

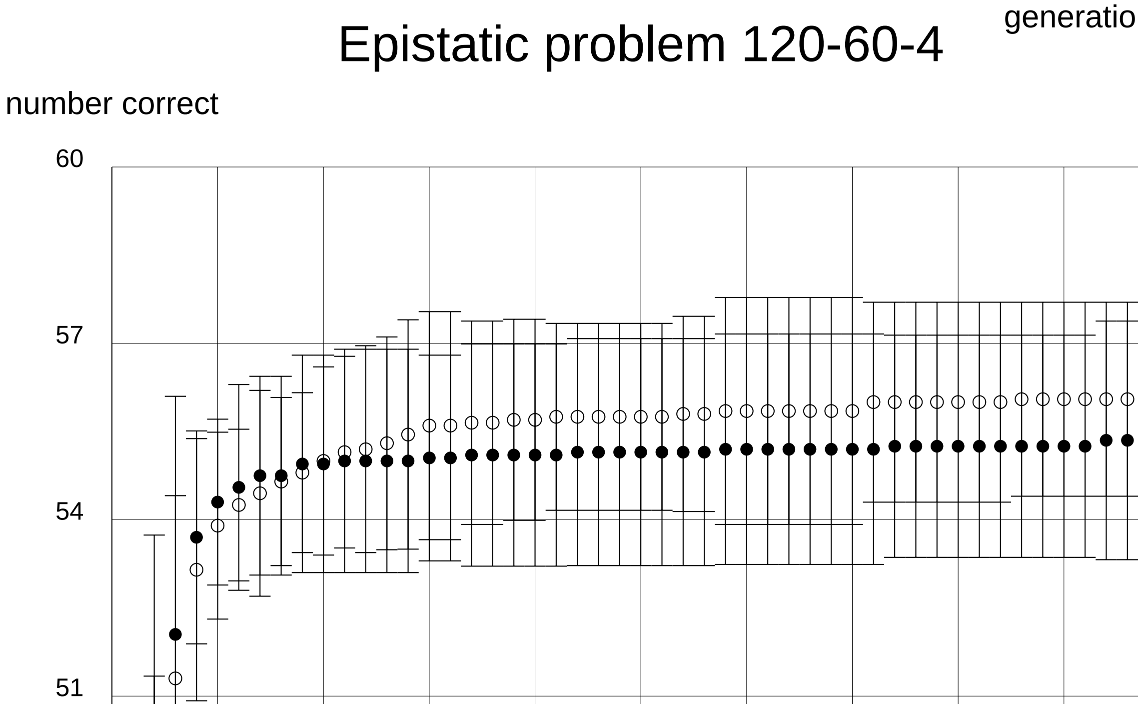
Epistatic problem 120-60-4
number correct
60
57
54
51

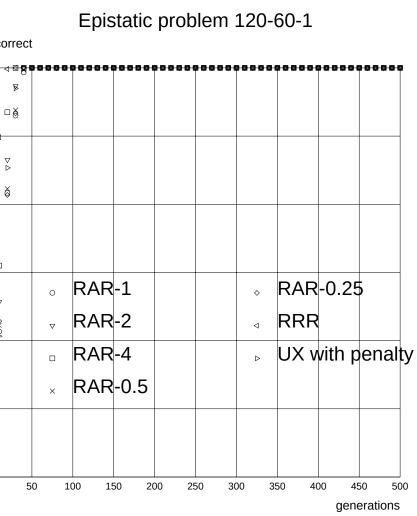

Epistatic problem 120-60-1
orrect
RAR-1
RAR-2
RAR-4
RAR-0.5
RAR-0.25
RRR
UX with penalty
50
100
150
200
250
300
350
400
450
500
generations

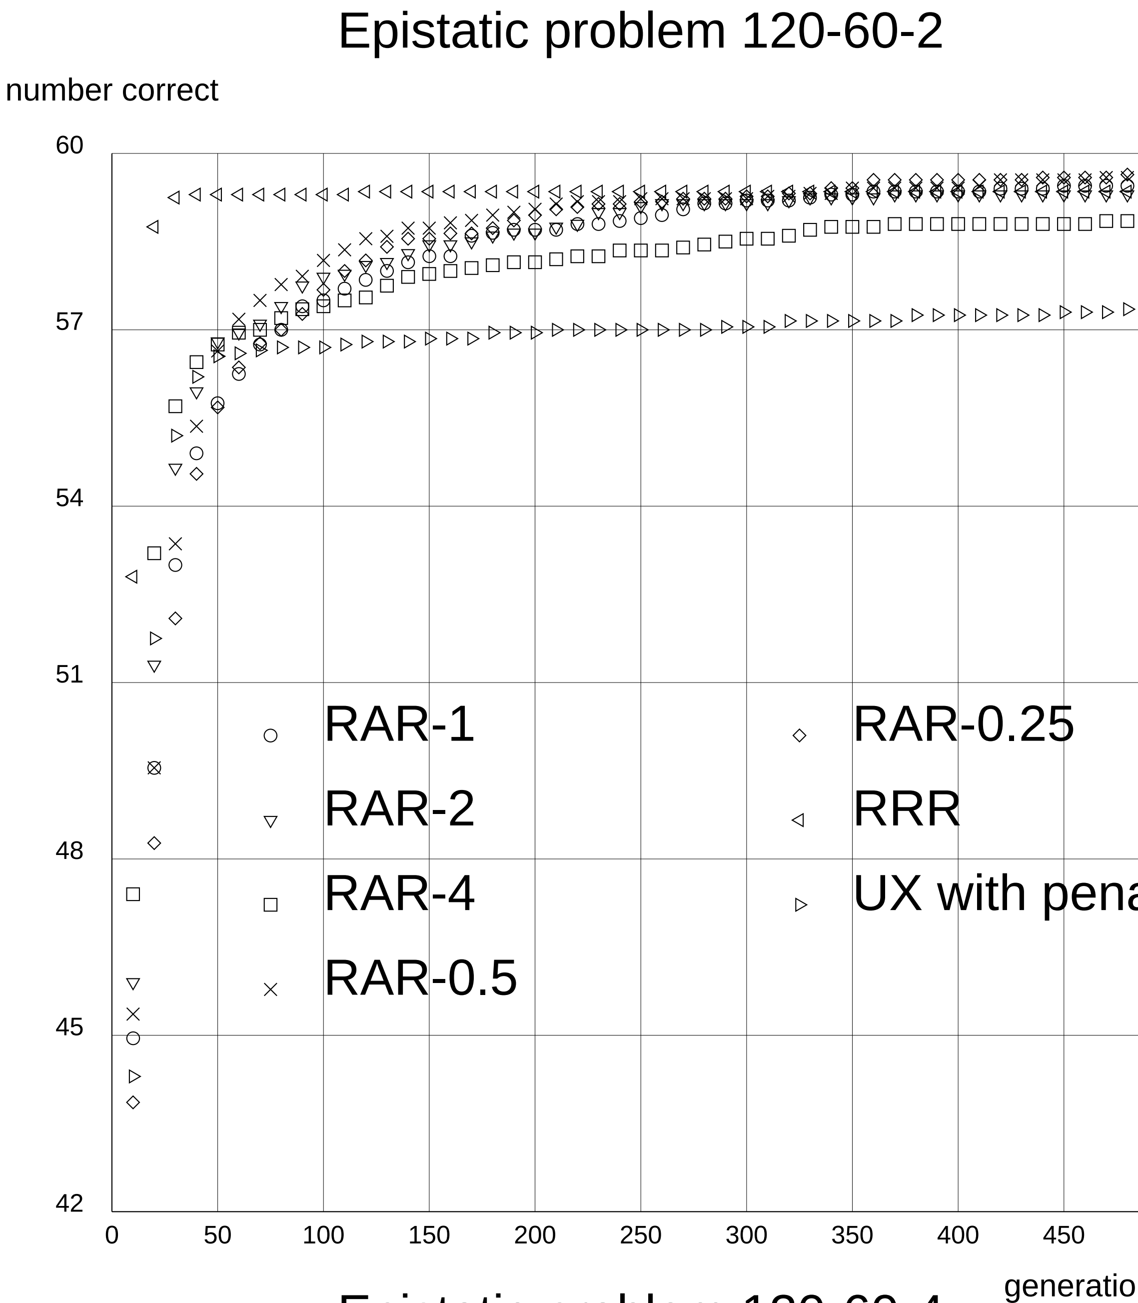

Epistatic problem 120-60-2
number correct
60
57
54
51
48
45
42
RAR-1
RAR-2
RAR-4
RAR-0.5
RAR-0.25
RRR
UX with pena
0
50
100
150
200
250
300
350
400
450
generatio

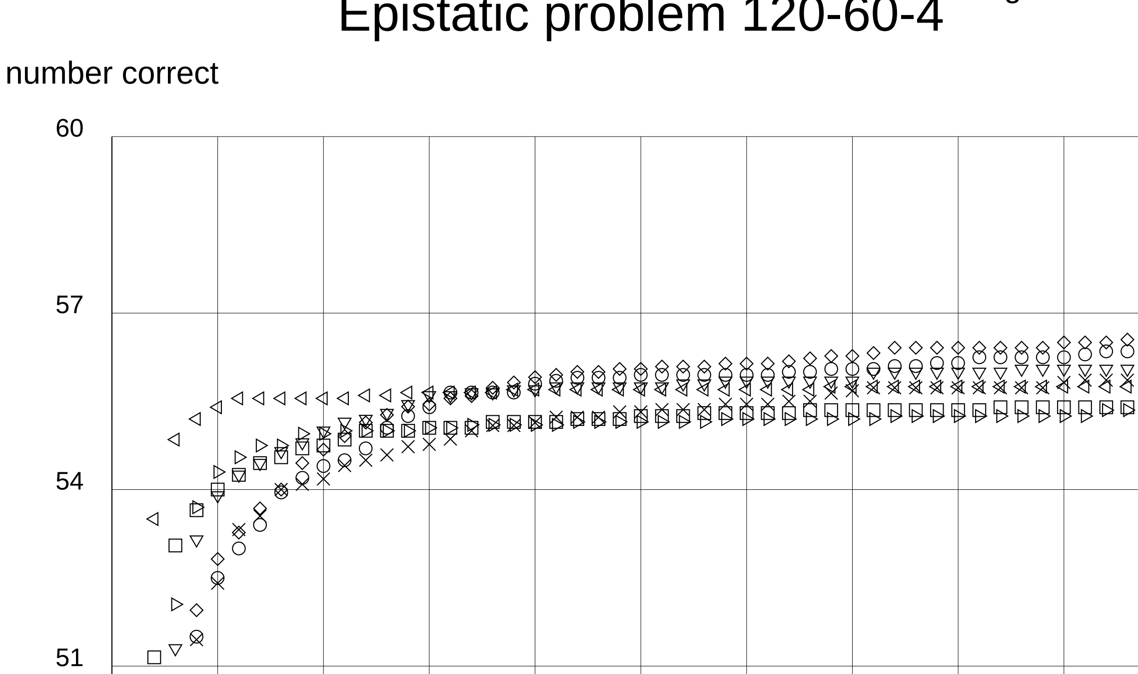

Epistatic problem 120-60-4
number correct
60
57
54
51

trade-off between respect (cf. exploitation) and assortment (cf. exploration) in problems for which these are incompatible.

Acknowledgements

This work made extensive use of the RPL Interpreter, which was developed by Claudio Russo and Graham Jones, and auxiliary routines for sets developed by Jonathan Shapcott.

References

Antonisse, 1989. Jim Antonisse. A new interpretation of schema notation that overturns the binary coding constraint. In *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann (San Mateo), 1989.

Davis, 1991. Lawrence Davis. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold (New York), 1991.

De Jong, 1975. Kenneth A. De Jong. *An Analysis of the Behaviour of a Class of Genetic Adaptive Systems*. PhD thesis, University of Michigan, 1975.

Eshelman and Schaffer, 1992. Larry J. Eshelman and David J. Schaffer. Real-coded genetic algorithms and interval schemata. In Darrell Whitley, editor, *Foundations of Genetic Algorithms 2*. Morgan Kaufmann (San Mateo, CA), 1992.

Eshelman *et al.*, 1989. Larry J. Eshelman, Richard A. Caruana, and J. David Schaffer. Biases in the crossover landscape. In *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann (San Mateo), 1989.

Goldberg and Lingle Jr, 1985. David E. Goldberg and Robert Lingle Jr. Alleles, loci and the traveling salesman problem. In *Proceedings of an International Conference on Genetic Algorithms*. Lawrence Erlbaum Associates (Hillsdale), 1985.

Goldberg, 1990a. David E. Goldberg. Genetic algorithms and walsh functions: Part I, a gentle introduction. *Complex Systems*, 3:129–152, 1990.

Goldberg, 1990b. David E. Goldberg. Genetic algorithms and walsh functions: Part II, deception and its analysis. *Complex Systems*, 3:153–171, 1990.

Goldberg, 1990c. David E. Goldberg. Real-coded genetic algorithms, virtual alphabets, and blocking. Technical Report IlliGAL Report No. 90001, Department of General Engineering, University of Illinois at Urbana-Champaign, 1990.

Harp *et al.*, 1989. Steven Alex Harp, Tariq Samad, and Aloke Guha. Towards the genetic synthesis of neural networks. In *Proceedings of the Third International Conference on Genetic Algorithms*, pages 360–369. Morgan Kaufmann (San Mateo), 1989.

Holland, 1975. John H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press (Ann Arbor), 1975.

Jones, 1992. Graham P. Jones. Parallel genetic algorithms for large travelling salesrep problems. Master's thesis, University of Edinburgh, 1992.

Michalewicz and Janikow, 1991. Zbigniew Michalewicz and Cezary Z. Janikow. Handling constraints in genetic algorithms. In *Proceedings of the Fourth International Conference on Genetic Algorithms*. Morgan Kaufmann (San Mateo), 1991.

Radcliffe, 1990. Nicholas J. Radcliffe. *Genetic Neural Networks on MIMD Computers*. PhD thesis, University of Edinburgh, 1990.

Radcliffe, 1991. Nicholas J. Radcliffe. Equivalence class analysis of genetic algorithms. *Complex Systems*, 5(2):183–205, 1991.

Radcliffe, 1992. Nicholas J. Radcliffe. Genetic set recombination. In Darrell Whitley, editor, *Foundations of Genetic Algorithms 2*. Morgan Kaufmann (San Mateo, CA), 1992.

Radcliffe, 1993. Nicholas J. Radcliffe. The algebra of genetic algorithms. *to appear in Annals of Maths and Artificial Intelligence*, 1993.

Richardson *et al.*, 1989. John T. Richardson, Mark R. Palmer, Gunar Liepins, and Mike Hilliard. Some guidelines for genetic algorithms with penalty functions. In *Proceedings of the Third International Conference on Genetic Algorithms*, pages 191–195. Morgan Kaufmann (San Mateo), 1989.

Russo, 1991. Claudio V. Russo. A general framework for implementing genetic algorithms. Technical Report EPCC–SS91–17, Edinburgh Parallel Computing Centre, University of Edinburgh, 1991.

Spears and De Jong, 1991. William M. Spears and Kenneth A. De Jong. On the virtues of parameterised uniform crossover. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 230–236. Morgan Kaufmann (San Mateo), 1991.

Syswerda, 1989. Gilbert Syswerda. Uniform crossover in genetic algorithms. In *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann (San Mateo), 1989.

Vose and Liepins, 1991. Michael D. Vose and Gunar E. Liepins. Schema disruption. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 237–243. Morgan Kaufmann (San Mateo), 1991.

Vose, 1991. Michael D. Vose. Generalizing the notion of schema in genetic algorithms. *Artificial Intelligence*, 1991.

Wright, 1990. Alden H. Wright. Genetic algorithms for real parameter optimisation. In Gregory J. E. Rawlins, editor, *Foundations of Genetic Algorithms*, pages 205–220. Morgan Kaufmann (San Mateo, CA), 1990.