

CHAPTER 10

THE BIMODAL FEATURE AT LARGE POPULATION SIZES AND HIGH SELECTION PRESSURE: IMPLICATIONS FOR DIRECTED EVOLUTION

Martin Oates¹, David Corne² and Douglas Kell³

¹ *Evosolve Ltd, Stowmarket, UK, moates@btinternet.com* ² *Dept. of Computer Science, University of Reading, UK, d.w.corne@reading.ac.uk* ³ *Institute of Biological Sciences, University of Aberystwyth, UK, dbk@aber.ac.uk*

Abstract.: The ‘bi’ and ‘higher modal features’ are aspects of Evolutionary Algorithm (EA) behaviour which are revealed, for a wide range of conditions, when extensive parametric studies are done to explore convergence time over a wide range of mutation rates. The bimodal feature indicates optimal mutation rates in terms of convergence time, which often correspond to optimal mutation rates in terms of final solution quality. The significance of the bimodal feature lies in parameter setting issues, and it is of interest to see how it varies with parameters and EA designs. Previous work shows that it appears in a wide range of conditions, but attenuates (the local optimum in convergence time becomes less apparent) with larger population sizes and low selection pressure. This chapter extends exploration of the bimodal feature into EAs with much larger population sizes, and show that under sufficiently high selection pressure it ‘returns’. It is interesting to note that these observations apply directly in the emerging field of ‘Directed Evolution’ for novel bio-molecules, in which large parallel populations undergo evolutionary search, with solution quality and number of generations being vital to optimise. This has potentially highly significant consequences for setting of mutation rates in Directed Evolution and high selection pressure large-scale parallel EAs in general.

10.1 Introduction

Much experimental and theoretical work has been done examining optimum parameter settings for Evolutionary Algorithms when applied to a very wide range of problems such as combinatorial and function optimisation, for example ^{4, 5, 10, 12, 15, 32, 18, 26}. These parameters have included, amongst others, population size, selection pressure, mutation rate, crossover rate and crossover operator. Previous work by the authors has focussed on optimising the search process for an industrial application (such as automated web load balancing ^{20, 21, 27}) with an emphasis on the repeatability, speed and accuracy of the search). In general this application was facilitated by the use of small, embedded controllers where sequential processing has been the norm and thus the inherent parallelisation of EAs to allow concurrent fitness evaluation has not been readily exploitable. This has provided results indicating that small populations running with steady state algorithms with Tournament ⁶ style selection pressure and traditionally high mutation rates tend to produce good results in a minimum number of evaluations. However, in direct contrast to this, biological studies in ‘Directed Evolution’ ^{1, 29}, where bacteriological samples are bred to improve a desired characteristic such as toxin immunity etc, are interested in getting reliable results in a minimum number of generations, and where the use of large populations with parallel evaluation is commonplace ^{2, 7, 36}. Much work has also been done by others on parallel EAs, where fitness evaluation is carried out across a cluster of processors, utilised by a central (or sometimes distributed) Evolutionary Algorithm controller. These configurations lend themselves more naturally (though not exclusively) to Generational style EAs using a form of Elitist ‘Breeder’ ¹⁷ style selection strategy.

In support of a biological study utilising ‘Directed Evolution’ targeted at rapid, novel enzyme development the authors are part of a team now examining ^{8, 9, 27} the performance characteristics of some of these large population, minimum generation EAs to attempt to optimise the bacteriological and virological studies being carried out by the biologists. In these biological studies it is possible to have populations of several

thousand members, derived from a considerably smaller elitist breeding pool, with parallel evaluation. Each generational evaluation cycle may take hours or days to complete regardless of the population size. To give an example of where Directed Evolution may be applied, consider the early development of an epidemic of a new strain of bacteriological or virological threat. Here infection rates typically rise exponentially, and thus it is crucial to cut the number of such evaluation cycles to a minimum in the search for an effective vaccine. Therefore it is critical to find control parameters which deliver good results in a minimum number of generations, with far less regard to the actual number of evaluations carried out. Perhaps not surprisingly, existing work in bacterial strain improvement traditionally uses mutation rates focussed around the reciprocal of the chromosome length ($1/L$), based on work commonly attributed to Baltz⁷. This mutation rate is typically induced by exposure to radiation or specific chemicals, the latter of which can also be used for 'targetted mutation' at particular loci and/or alleles. However emerging evidence from more recent studies suggests considerably higher mutation rates can prove more effective^{3,36}.

This chapter presents some of the initial results from this new study showing that with very large population sizes and very high selection pressure, higher than traditional mutation rates deliver improved results on a range of standard test problems. The chapter begins with a background summary of relevant previous work leading to the experiments carried out to date. A discussion of these new results is provided together with initial conclusions and plans for future work.

10.2 Background

Figure 10.1 shows the mean performance profile (averaged over 50 runs) for a steady state, 3 way single tournament EA using uniform crossover³⁰ and 'New Random Allele' mutation at a specified rate per gene on Watson's H-IFF problem (described later). The algorithm has a population size of 20 and each run is allowed 1,000,000 evaluations. The graph shows cyclical and phasic behaviour in the number of evaluations

used to first find the best solution found, the standard deviation in this value and the fitness of the best solution found. This has been explained in ²⁶, where it was shown that the performance of the algorithm over the range of mutation rates examined passes through 3 distinct phases, which repeat at least 3 times. In the first cycle of the first phase, the algorithm is starting to exploit the low level of mutation available to it, predominantly occurring as single point mutations. As mutation rates rise, these mutations occur with increasing frequency allowing the algorithm to utilise an increasing number of evaluations, until a point is reached where the usefulness of single point mutations is exhausted. As mutation rates increase, this point is reached earlier in the run, and hence the number of evaluation used falls. This is the second 'phase' of the performance profile and is further characterised by the flattening of the best found fitness plot and the reduction in the standard deviation of the number of evaluations used. At this point, mutation rates are still too low for the occurrence of 2 point mutations within the same chromosome to have any significant affect. However, as mutation rates increase, a point is reached where the likely occurrence of 2 point mutation becomes significant and this occasionally allows the algorithm to break though the 'fitness barrier' ¹⁴ surrounding the local optimum it has become stuck in. Hence the number of evaluations used becomes erratic (shown by a sudden, marked increase in its standard deviation. This is now the third 'phase' of performance behaviour. As mutation rates continue to increase, 2 point mutations become commonplace and the algorithm reverts back to its original phase behaviour exploiting increasing occurrences of these and the 3 phase cycle repeats until the usefulness of 2 point mutations is exhausted. The cycle is shown to repeat at least one further time, before excessive mutation rates cause the algorithm to deteriorate into random search. This explanation is described and analysed in far more detail in ²⁶ and is shown to exist in a range of multi-modal problems in ^{24, 25}. In mono-modal problems, only a single cycle of these 3 phases is usually observed as would be expected by such a hypothesis, as no 'fitness barriers' exist which require specific types of mutation to breach.

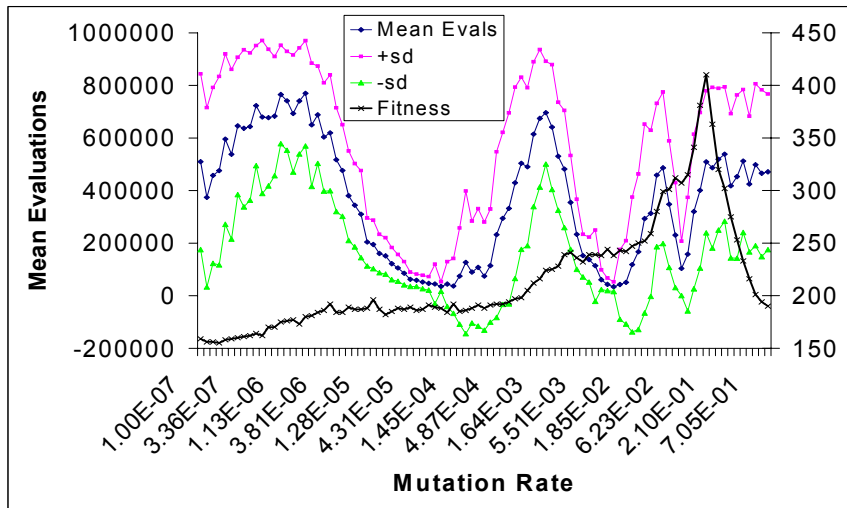


Figure 10.1 H-IFF 64 Performance Profile at 1 Million evaluations




Figure 10.2 Royal Stair 50-1 Performance Profile with 3 way Tournament selection Population Size 100

Figure 10.2 shows the performance of a steady state Evolutionary algorithm on an instance of the Royal Staircase problem (length 50, block size 1), showing the co-incidence in troughs of minimum error, minimum evaluations used and minimum coefficient of variation (standard deviation divided by the mean ie a minimum in the normalised process variability). The results are again the average of 50 runs of the algorithm, each time with a population size of 100, uniform crossover at a probability of 1.0 and New Random Allele replacement mutation at the indicated rate per gene. Each run was allowed 20,000 evaluations, reporting the first evaluation number at which the best result in the run was first seen. The algorithm employed 3 way, single Tournament selection. As can be seen, in the trough of optimum performance (at mutation rate of around 2.5%), the algorithm requires around 5,000 evaluations to find the global optimum.

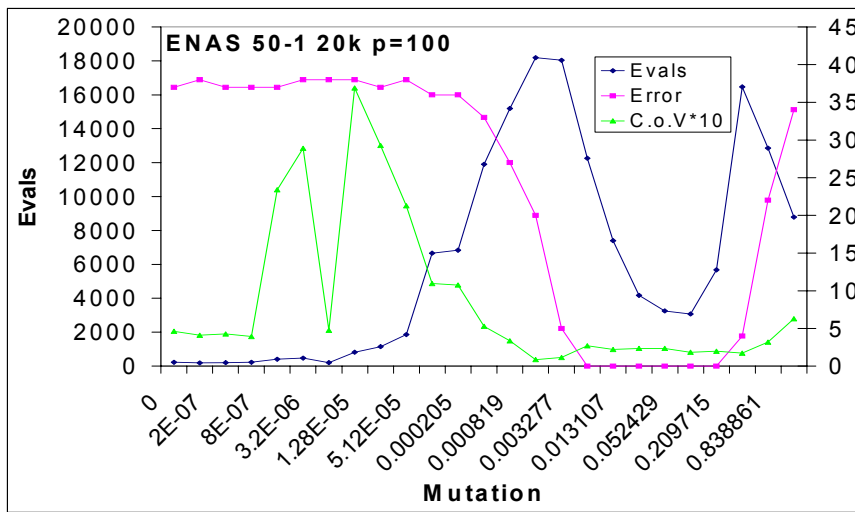


Figure 10.3 Royal Stair 50-1 Performance Profile with 8 way Tournament selection, Population Size 100




Figure 10.4 Evaluations used, RS50-1, 50% Elitism

The experiment was then repeated with the selection pressure increased to a single 8 way Tournament, where 8 members of the population are chosen at random and ranked. The first and second best are used as parents to produce a child which replaces the 8th ranked member of the Tournament back in the original population. This increased selection pressure can be seen in Figure 10.3 to have 3 predominant effects on the performance profile. Firstly, the number of evaluations used in the trough of optimum performance was seen to fall from around 5,000 evaluations (3 way Tournament) to below 3,000 evaluations (8 way Tournament). Secondly, the average error of ‘best solutions found’ at low mutation rates was seen to deteriorate with higher selection pressure, and finally the average number of evaluations used at these mutation rates was seen to fall. Neither of these last two effects are surprising as

the increased selection pressure is clearly causing earlier premature convergence from which the algorithm cannot escape due to lack of mutation. These results are a subset of previously published results in ²², ²³ wherein these effects are shown over a wider range of population sizes, problem instances and algorithm designs.

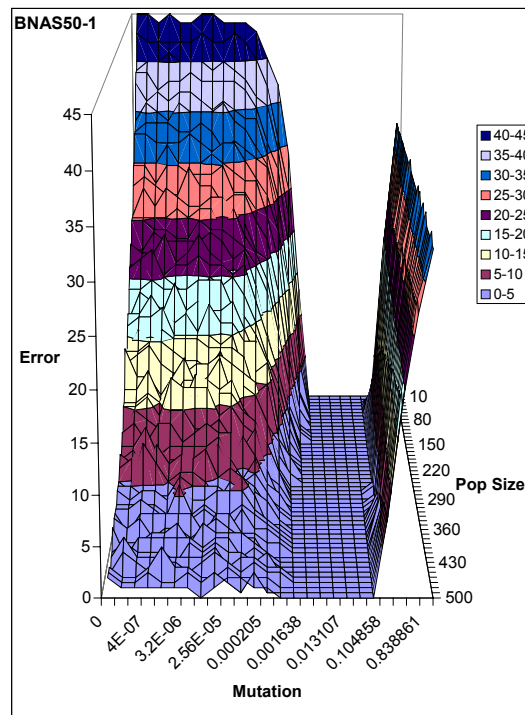


Figure 10.5 Errors, RS50-1, 50% Elitism

These results show the aforementioned effects on a steady state algorithm, whilst Figures 10.4, 10.5, 10.6 and 10.7 show similar effects on Generational Breeder style algorithms incorporating 50% at 10% elitism respectively. Here the algorithm ranks the entire population and then discards the lower performing half (or 90%). The surviving members of the population are then randomly selected in pairs as parents (using uniform crossover at probability 1.0 followed by per gene mutation) to restore the population to original size.

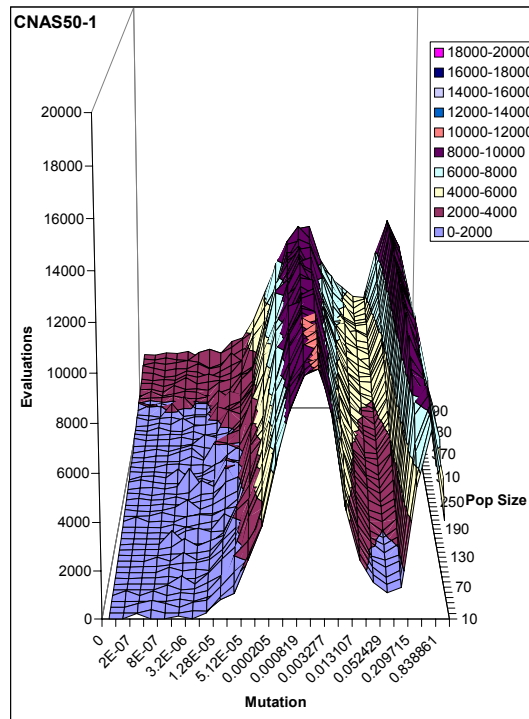


Figure 10.6 Evaluations used, RS50-1, 10% Elitism

Population sizes from 10 through to 500 (in steps of 10) have been trialled with mutation rates ranging from 1 E^{-7} to 0.83 per gene. In all cases the results are the average of 50 runs each of which is allowed 20,000 evaluations. What can clearly be seen from a baseline in Figure 10.4 is that whilst the bimodal performance profile is clearly apparent at low population sizes, it is attenuated by increased population size. Figure 10.5 shows that at low population sizes, only a specific sub-range of mutation rates can deliver good performance (zero error from the optimum fitness value), whilst as population size increases, performance at these lower mutation rates improves, until by a population size of 500, adequate performance is just beginning to be delivered. Figure 10.6 shows the contrast where selection pressure is increased by only allowing the top 10% of each generation to breed. Here, as in the case of the

steady state algorithm, the 3 effects of increased selection pressure can clearly be seen : reduced evaluations needed at optimum mutation rates; more rapid premature convergence at low mutation rates; with this convergence on poorer solutions. What is also important however, is the clear continuation of the bimodal performance profile into higher population sizes. Whilst at low population sizes the affect is attenuated with respect to the lower selection pressure case, the effect is still clear in the population size 500 case, which was not true for the low selection pressure example.




Figure 10.7 Errors, RS50-1, 10% Elitism

Figures 10.8 and 10.9 contrast with Figures 10.4 and 10.5, showing results where selection pressure remains at 50%, however the probability of performing crossover is reduced to 0.75 (in the case of no crossover, a

single parent is used with mutation only). Whereas in Figure 10.4 the first ridge of high evaluations is seen to be attenuated by increased population size, in Figure 10.8 this ridge remains high, but the trough of optimum performance is seen to rise. At a population size of 500 the bimodal profile is still just observable. Further experiments with the probability of crossover reduced to only 50% and 10% continue these trends and the 50% crossover results for ‘evaluations used’ and ‘errors’ are shown in Figures 10.10 and 10.11 respectively.. These results on the Royal Staircase problem are in direct support of earlier results in ^{21, 23} on the One Max problem




Figure 10.8 Evaluations used, RS50-1, 75% crossover

Thus it has been clearly shown that there exist ranges of optimal mutation rates capable of delivering highly robust performance in a

minimum of evaluations with a high degree of accuracy. These studies have shown that the bimodal effect, normally most prevalent at low population sizes, can be extended to affect algorithm performance at higher population sizes where high levels of selection pressure and reduced crossover are utilised. Further it can be seen that whilst this is generally at the expense of a greater number of evaluations required, it can also lead to a significant reduction in the number of generations required.




Figure 10.9 Errors, RS50-1, 75% crossover

10.3 Experimental Method

A natural extension of the above experiments is to investigate the performance of an algorithm with a very large population, derived via a highly elitist selection strategy, over a range a traditionally high mutation

rates. This models the situation when Directed Evolution is applied. In this section we present results from a range of initial experiments using such an algorithm utilising a population size of 10,000 members (initially randomly generated), where the next generation is entirely derived from the single fittest member of the population subjected to ‘per gene’ mutation at a specified rate. ‘New Random Allele’ rates of mutation from 1.024×10^{-4} to 0.838 have been trialled on an exponential scale where the mutation rate doubles between each experiment in the 14 case set. In each case, the algorithm is allowed 50 generations (ie 500,000 evaluations), reporting the fitness of the best solution found, and the generation number this was first found at. Each experiment is then repeated 50 times and results plotted show the mean results of these 50 runs, and the standard deviation of the number of generations used across the 50 runs.

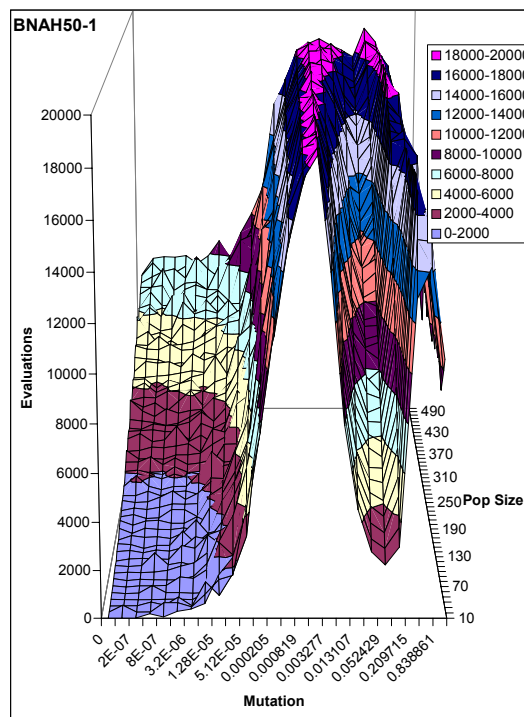


Figure 10.10 Evaluations used, RS50-1, 50% crossover

Experiments have been carried out on a range of standard test problems (Max Ones, Royal Staircase¹⁹, Kauffman NK¹⁵, H-IFF³³ etc) with only a representative sample given here for space reasons. The tunable Royal Staircase problem in this instance is a mono-modal problem with significant regions of neutral fitness plateaux. Fitness is derived by counting the number of consecutive blocks of all 1s in the chromosome starting from the left-hand side and has been extensively researched by Crutchfield and Van Nimwegen¹⁹. With this problem, with a block size set at 5, a string of 12 1's followed by any (non zero) number of 0's followed by any combination of 1's and 0's delivers a fitness of 2 out of a possible global optimum of 10. A string containing 49 1's preceded by a single 0 delivers a fitness of 0. A chromosome length of 50 was used in these experiments with block sizes set to 1, 2, 5 and 10.

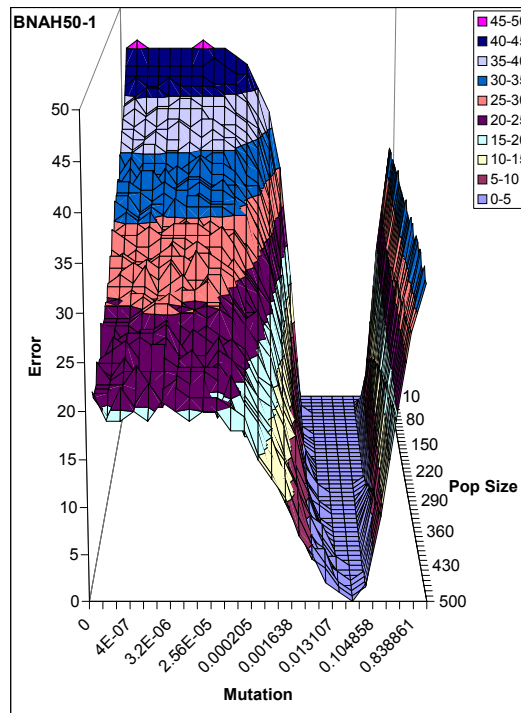


Figure 10.11 Errors, RS50-1, 50% Crossover

The tunable Kauffman NK ¹⁵ problem allows varying levels of epistatic and positional linkage to be explored. In this implementation with a chromosome length of 50 and maximum block size of 6, a randomly generated look up table is generated containing 50 rows by 64 columns. For a given block size of 1, fitness is simply derived by taking each gene individually and summing either the first or second column entry, determined by allele, over the 50 genes in the chromosome (whose locus determines the row). This is in effect a form of 'max ones' mono-modal function. However for larger block sizes, consecutive sequences of genes are used as a binary word to derive a column index into the table. Thus for a block size of 3, each gene takes part in 3 table retrievals over the 50 needed to derive overall fitness of a chromosome. Therefore any single point mutation will affect multiple aspects of the overall fitness calculation. This creates an ever ruggedised search space deteriorating to a random field as block sizes approach the length of the chromosome. Block sizes of 1, through 6 have currently been investigated.

Watson's Hierarchical If and only If problem (H-IFF) ³³ has been widely investigated by the author and others, and although first derived to explore the effects of crossover and schemata development, was critically instrumental in helping demonstrate the emergence and explanation of multi-modal algorithmic performance when subjected to varying rates of mutation. The fitness of a potential solution to this problem is the sum of weighted, aligned, decomposable blocks of either contiguous 1's or 0's. This produces a search landscape in which 2 global optima exist. one as a string of all 1's, the other of all 0's. However a single mutation away from either of these positions produces a much lower fitness. Secondary optima exist at strings of 32 contiguous 0's followed by 32 contiguous 1's (for a chromosome length of 64) and vice versa. Not surprisingly, Watson showed that hill-climbing performs extremely badly on this problem ³⁴.

Together, these test problems provide an informative and diverse set to explore many aspects of algorithm performance on combinatorial optimisation problems with low allelic range.

10.4 Results

Figure 10.12 shows the performance of the highly elite algorithm on the Royal Staircase problem with a chromosome length of 50 and block size of 1. Here it can clearly be seen that once sufficient mutation is available to the algorithm, the global optimum can be achieved 50 times out of 50 in less than 30 generations. As mutation rates increase, this number of generations required is seen to fall to a minimum of 9 at a per gene mutation rate of around 10%. Above this rate of mutation algorithm performance starts to deteriorate with a marked increase in error, number of generations needed and process unrepeatability. As the block size is increased to 2 (25 blocks thereof), this range of good mutation rates is seen to narrow (Figure 10.13), with the lower end rates no longer delivering adequate performance. By the time the block size is increased to 5, the algorithm is failing to consistently find the global optimum solution and the number of generations used is seen to be high. A slight dip is observable at the now optimum mutation rate of 20% (Figure 10.14).

The results for the Max Ones problem with a chromosome length of 50 are presented in Figure 10.15 and can be seen to be very similar to the block size 1 Royal Staircase results. The Max Ones problem is also representative of the Kauffman NK problem with a block size of only 1. Figure 10.16 shows results for the Kauffman NK problem with a block size of 2, now a multi-modal problem. Again we see the inability of low mutation rates to deliver good solutions, and an optimum performance at a mutation rate of around 10%. A dip is also seen in the number of evaluations used at a mutation rate of around 1.3%, but not accompanied by a similar dip in the fitness error. This is likely to be an effect similar to that seen at the beginning of this paper on the highly structured multi-modal H-IFF problem. Here we see the effect of predominant one point mutation finding certain local optima, but with insufficient 2 point mutations available, the algorithm gets trapped in these solutions. As 2 point mutations become more prevalent, at higher mutation rates, so the fitness wells surrounding these local optima can be breached and the algorithm can exploit more evaluations.

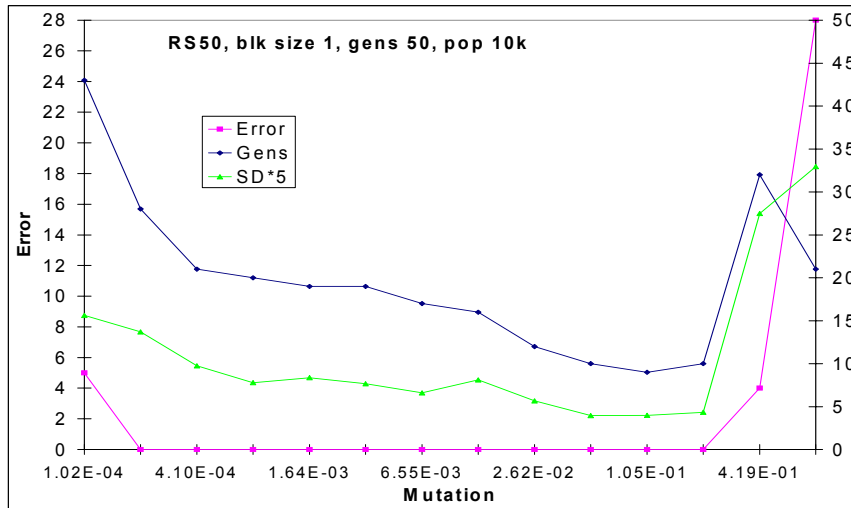


Figure 10.12 Performance Profile for highly Elite algorithm on Royal Staircase problem, block size 1




Figure 10.13 Performance Profile for highly Elite algorithm on Royal Staircase problem, block size 2

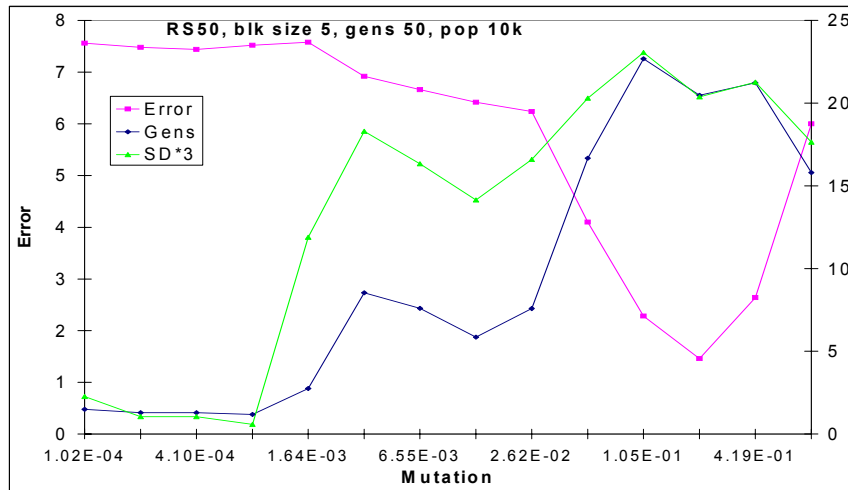


Figure 10.14 Performance Profile for highly Elite algorithm on Royal Staircase problem, block size 5

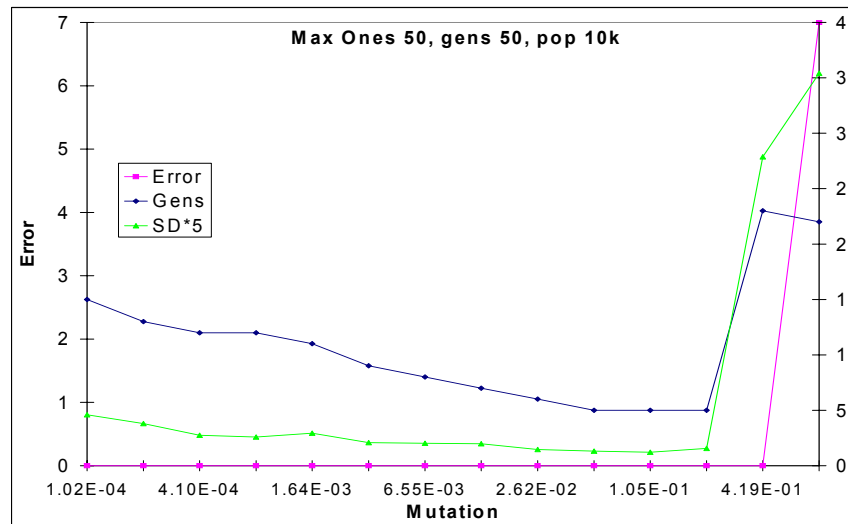


Figure 10.15 Performance Profile for highly Elite algorithm on Max Ones length 50 problem




Figure 10.16 Performance Profile for highly Elite algorithm on NK50 problem, block size 2

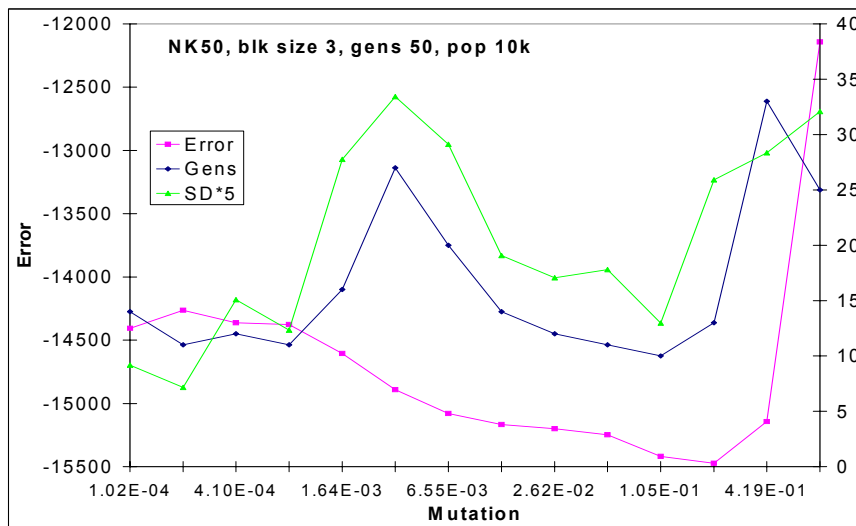


Figure 10.17 Performance Profile for highly Elite algorithm on NK50 problem, block size 3

Figure 10.17 shows results for the NK problem with a block size of 3. Again an optimum mutation rate is seen in terms of fitness error, with a local minimum in the number of generations used. The preceding multi-phasic behaviour is not apparent, but this is probably due to the small number of mutation rates sampled. Figure 10.18 shows similar results with a block size of 4, whilst Figure 10.19 shows results with a block size of 5. By this time, the algorithm is showing clear signs of insufficient generations being allowed. The multi-modal profile is severely attenuated, however there is still a clear minimum in the fitness error plot. The fact that the number of generations required is now relatively high is typical of results published for serial algorithms using both steady state and generational techniques, where the number of evaluations required is severely limited ²¹.

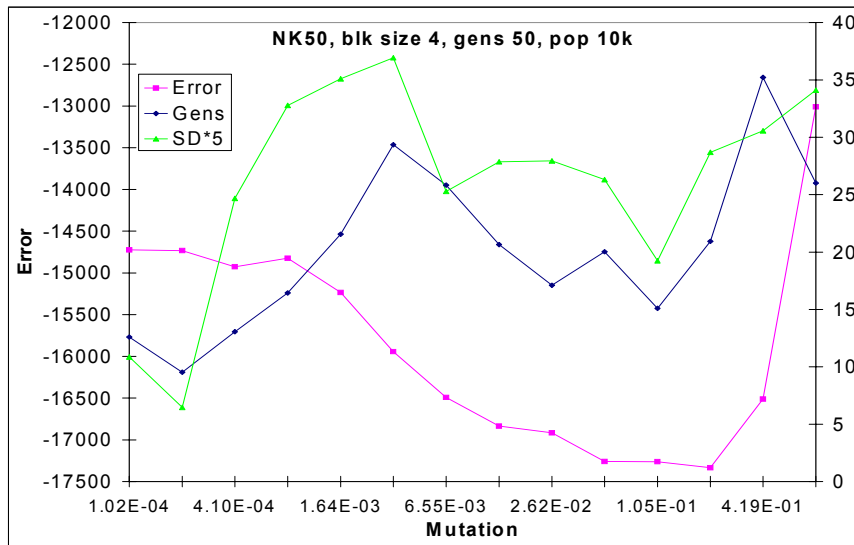


Figure 10.18 Performance Profile for highly Elite algorithm on NK50 problem, block size 4

Figure 10.20 shows the results for the highly structured multi-modal H-IFF problem with a chromosome length of 64. Here we again see a dip in

the fitness error, but accompanied by increasing numbers of generations used. This is an extension of the effect described immediately above.

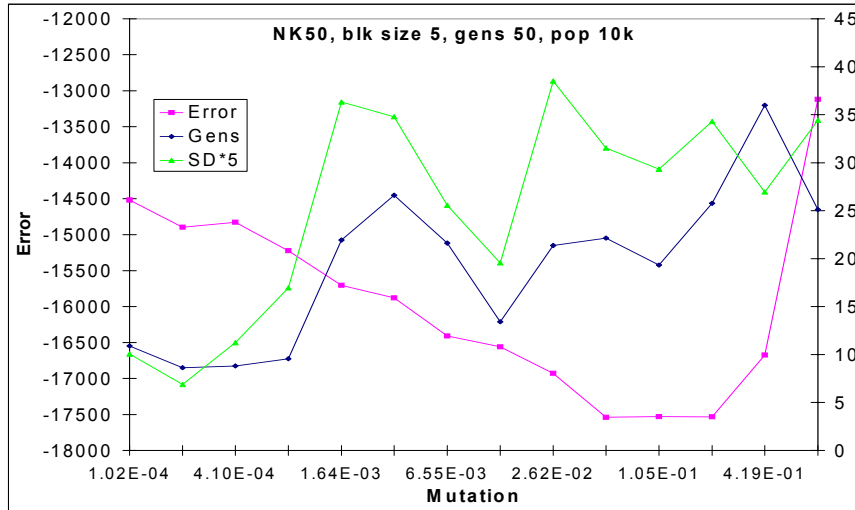


Figure 10.19 Performance Profile for highly Elite algorithm on NK50 problem, block size 5

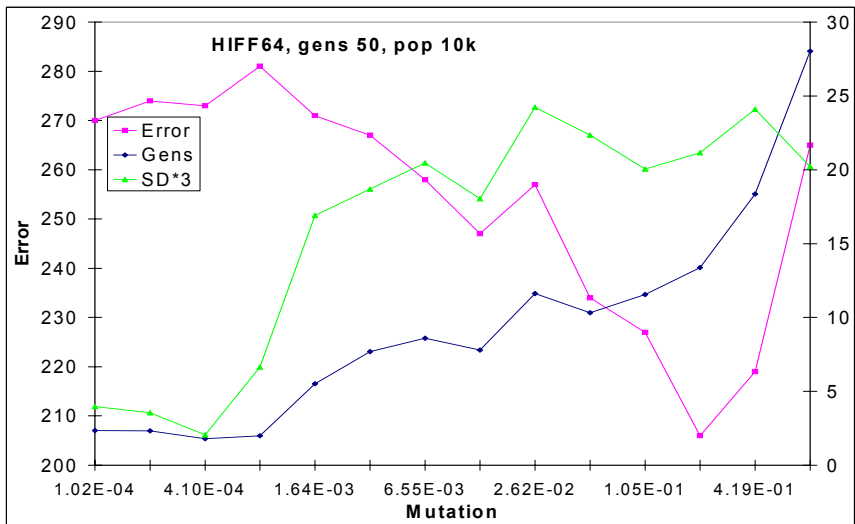


Figure 10.20 Performance Profile for highly Elite algorithm on H-IFF 64 problem

10.5 Discussion

Contrasting the behaviour of these somewhat extreme algorithms to results with more conventional parameters shows that these results are obtained at a price. Whilst the number of generations required has been seen to be small, the number of evaluations required is considerably higher than that of an algorithm tuned for sequential use. Table 1 gives comparisons between the number of evaluations and generations needed by the highly elitist algorithm described in this paper with experiments with a 50% elitist Breeder algorithm using population sizes of 100 and 500 respectively each allowed up to 20,000 evaluations. The results for each algorithm are the average over 50 runs at optimum mutation rates targeted for optimum fitness followed by minimum evaluations. As can clearly be seen, the number of evaluations is typically more than an order of magnitude greater, but the number of generations required is much reduced (typically by a factor of 3 or 4 between the Elite algorithm and the Population Size 500 algorithm, and by a similar factor again against the Population size 100 algorithm). This trade-off would be difficult to justify in terms of the massively parallel processing environment it would require for a typical EA combinatorial optimisation problem, but is a relatively insignificant price to pay in biological assay evaluations. It is also worth noting that on the NK 50-2, 50-3, 50-5 and 50-6 problems, the average fitness of best found solutions was better for the Elite algorithm than for either of the 50% Breeder algorithms. For the RS 50-1, RS 25-2 and Max 1s problems, the global optimum was consistently found by all 3 algorithms and for the RS 10-5, and NK 50-3 problems the elite algorithm average best found fitness was marginally worse. For the highly structured H-IFF problem, as would be expected with no population crossover, the elite algorithm consistently under-performed, in terms of best found fitness, relative to the other two algorithms.

The extremely high level of selection pressure utilised by this algorithm (only the single fittest being used to produce all of the next generation) draws parallels with developments along the lines of Hill Climbers and

other such, non-population based techniques. However such techniques tend to look only at single point mutations and as such contain no mechanism for escaping local or deceptive optima. The advantage of a high 'per gene' mutation rate is that over the 10,000 derivatives of the elite parent, a wide variety of differing mutation schedules is produced. Some variants will suffer only single point mutation, others 2 point, whilst potentially some could have all genes replaced by random alleles (no mutations is also a distinct, but wasted possibility). The higher the 'per gene' mutation rate, the more that multi-point mutations will dominate this distribution. During different stages of the optimisation, different types of mutation are likely to be of most use. Given the random generation of the first 10,000 evaluations, a relatively wide coverage of the search space is examined and thus low rates of mutation will allow local exploration to find a local optimum in the next generation. However, once found, in a multi-modal or deceptive search space, a considerably more disruptive mutation rate will be required to allow the search to break free of this optimum in the search for ever better optima.

10.6 Conclusions

The mutation rates seen to be effective under these circumstances on these problems appear well in excess of those traditionally used within the EA community, even higher than the now generally accepted optimal rates ($1/L$ and k/L) demonstrated by Mühlenbein¹⁸. However the absence of crossover and the extreme selection pressure used effectively invalidate such comparisons. This algorithm is indeed far closer to a traditional Evolutionary Strategy as developed by Rechenberg²⁸ than a standard Genetic Algorithm as developed by Holland¹³ and Goldberg¹¹, but once again, mutation rates and offspring sizes are larger than would be typically used in even these algorithms. Initial population generation is also significantly different. Further work is clearly required with problems of wider allelic range and other population sizes, to see at what point the trade-off becomes marginalised. Relaxation of selection pressure also requires investigation allowing the possibility of reintroduction of crossover.

On the assumption that directed evolution fitness landscapes ¹ (in contrast to those of natural evolution ³⁶) share similar properties to those problems investigated here, the reduction in number of generations required to obtain useful results becomes a highly significant advantage. This could lead, for example, to potential anti-toxins and vaccines being developed far earlier in the course of an epidemic leading to considerable reduction in suffering and potential reduction in loss of life.

Table 1. Relative reduction in number of Generations required with increased Pop Size

	50% BDR P 100			50% BDR P 500			Mono	Elite	P 10k
	Evals	Gens	SD	Evals	Gens	SD	Evals	Gens	SD
RS50-1	5050	99	23.6	14500	57	5.7	90K	9	0.80
RS25-2	6300	125	37.8	14000	55	7.1	100K	10	0.85
RS10-5	14200	283	80.7	17000	67	8.6	200K	20	6.8
MAX 1s	1750	34	4.9	4750	18	1.7	50K	5	0.24
NK50-2	6300	125	79.4	8750	34	6.4	130K	13	3.5
NK50-3	8350	166	87.3	13750	54	12.7	130K	13	5.2
NK50-4	9050	180	87.2	12250	48	10.4	210K	21	5.7
NK50-5	10600	211	103.9	17250	68	7.8	190K	19	5.9
NK50-6	11700	233	82.7	18000	71	7.2	240K	24	6.0
H-IFF64	5150	102	54.1	14000	55	10.1	130K	13	7.1

References

1. Arnold, F. (1998) Directed evolution, *Nature Biotech.*, **16**:617–618.
2. Arnold, F. (2001) Combinatorial and computational challenges for biocatalyst design. *Nature*; **409**:253–7.
3. Arnold, F. (2001) (ed.) *Evolutionary protein design*. San Diego: Academic Press.
4. Bäck, T. (1993) Optimal Mutation Rates in Genetic Search, *Proc. 5th ICGA*, pp 2 – 9.
5. Bäck, T (1994) Selective pressure in evolutionary algorithms: a characterization of selection mechanisms *Proc. 1st IEEE Conf. On Evolutionary Computation*, pp 57–62.
6. Bäck, T. (1996) *Evolutionary Algorithms in Theory and Practice*, Oxford University Press.

7. Baltz, R.H.(1986) Mutation in Streptomyces, in Day, L, Queener S, (eds). *The Bacteria, Vol 9, Antibiotic-producing Streptomyces*. New York: Academic Press, 61–94.
8. Corne, D., Kell, D.B., Oates, M. On Fitness Distributions and Expected Fitness Gain of Mutation Rates in Parallel Evolutionary Algorithms, in J.J. Merelo et al. (eds.) *PPSN VII*, Springer Verlag, Berlin, 2002.
9. D.W.Corne, M.J.Oates, D.B.Kell, *Fitness Gains and Mutation Patterns: Deriving Mutation Rates by Exploiting Landscape Information*, in Foundations of Genetic Algorithms 7, Morgan Kaufmann , pp347-364
10. Deb, K., Agrawal, S.(1998) Understanding Interactions among Genetic Algorithm Parameters. in *Foundations of Genetic Algorithms*, MKP.
11. Goldberg, D. (1989) *Genetic Algorithms in Search Optimisation and Machine Learning*, Addison Wesley.
12. Goldberg, D., Deb, K., Clark, J.H. (1992) Genetic Algorithms, noise, and the sizing of populations. *Complex Systems* 6:333–362.
13. Holland, J. (1975) *Adaptation in Natural and Artificial Systems*, MIT press, Cambridge, MA.
14. Jones, T. (1995) *Evolutionary Algorithms, Fitness Landscapes and Search*, PhD Dissertation, UNM.
15. Kauffman, S.A. (1993) *The Origins of Order: Self-Organization and Selection in Evolution*, OUP.
16. Michalewicz, Z. (1996) *Genetic Algorithms + Data Structures = Evolution Programs*, Springer, 1996.
17. Mühlenbein, H., Schlierkamp-Voosen, D. (1994), The Science of Breeding and its application to the Breeder Genetic Algorithm, *Evolutionary Computation* 1:335–360.
18. Mühlenbein, H.(1993) How genetic algorithms really work: I. Mutation and Hillclimbing, in Manner, Manderick (eds), *Proc. PPSN II*, Elsevier, pp 15–25.
19. van Nimwegen, E., Crutchfield, J. (1998) Optimizing Epochal Evolutionary Search: Population-Size Independent Theory, *Comp. Meth. in Applied Mechanics and Engineering*, special issue on Evolutionary and Genetic Algorithms, Goldberg, D., Deb, K. (eds).
20. Oates, M., Corne, D. (1998) Investigating Evolutionary Approaches to Adaptive Database Management against various Quality of Service Metrics, *Proc of PPSN-V*, pp. 775-784.
21. Oates, M., Corne, D., Loader, R. (1999) Investigation of a Characteristic Bimodal Convergence-time/Mutation-rate Feature in Evolutionary Search, in *Proc. ICEC 99 Vol 3*, IEEE, pp. 2175-2182
22. Oates, M., Corne, D., Turton, B. (1999) The Effects of Selection Pressure on Parameter Choice in Evolutionary Search, in Late Breaking Papers at GECCO 99, pp.198–203.
23. Oates, M., Smedley, J., Corne, D., Loader, R. (2000), Bimodal Performance Profile of Evolutionary Search and the Effects of Crossover, in *Theoretical Aspects of Evolutionary Computing*, (Editors Kallel, Naudts, Rogers), Springer Verlag.

24. Oates, M, Corne, D, Loader, R. (2000) A Tri-Phase Multimodal Evolutionary Search Performance Profile on the 'Hierarchical If and Only If' Problem, in *Procs of GECCO 2000*.
25. Oates, M., Corne, D., Loader, R. (2000) Tri-Phase Performance Profile of Evolutionary Search on Uni- and Multi- Modal Search Spaces, in *Procs of Congress on Evolutionary Computation (CEC)*.
26. Oates, M., Corne, D. (2000) Overcoming Fitness Barriers in Multi-Modal Search Spaces, in *Foundations of Genetic Algorithms 6*, MKP.
27. Oates, M. (2001) Global Web Server Load Balancing using Evolutionary Computational Techniques, *Soft Computing*.
28. M.Oates, D.Corne, D.Kell, *The Bimodal Feature at Large Population Sizes and High Selection Pressure: Implications for Directed Evolution*, in *Procs of 4th Asia-Pacific Conf on Simulated Evolution and Learning (SEAL 02)*, Vol 1, Wang et al (Eds), ISBN 981-04-7522-5 (2002), pp81-85.
29. Rechenberg, I. (1973) *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*, Frommann-Holzboog, Stuttgart.
30. Syswerda, G. (1989), Uniform Crossover in Genetic Algorithms, in Schaffer J. (ed), *Procs of 3rd ICGA*, Morgan Kaufmann, pp. 2 – 9
31. Voigt CA, Kauffman S., Wang Z.G.. Rational evolutionary design: The theory of in vitro protein evolution. In: Arnold, F., editor. *Advances in Protein Chemistry*, Vol 55, 2001:79-160.
32. Vose M, *The Simple Genetic Algorithm; Foundations and Theory*, MIT Press, 1999
33. Watson RA, Hornby GS, and Pollack JB, Modelling Building-Block Interdependency, *PPSN-V* (1998), pp. 97-106.
34. Watson RA, Pollack JB, Hierarchically Consistent Test Problems for Genetic Algorithms, in *CEC 99*, Vol 2, IEEE, pp. 1406–1413
35. Wright S, *The roles of mutation, inbreeding, crossbreeding and selection in evolution*, *Procs of 6th Intl Conf on Genetics* (Ithaca, NY, 1932) vol 1, ed D F Jones (Menasha, WI: Brooklyn Botanical Gardens), pp. 356-366.
36. Zaccolo, M. & Gherardi, E. (1999). The effect of high-frequency random mutagenesis on in vitro protein evolution: A study on TEM-1 β -lactamase. *J. Mol. Biol.* 285, 775-783.

Published as: Oates, M. J., Corne, D. W. & Kell, D. B. (2003). The bimodal feature at large population sizes and high selection pressure: implications for directed evolution. In *Recent advances in simulated evolution and learning* (ed. K. C. Tan, M. H. Lim, X. Yao and L. Wang). Pp 215-240. World Scientific, Singapore.