
On Replacement Strategies in Steady State Evolutionary Algorithms

Jim Smith

james.smith@uwe.ac.uk

Faculty of Computing, Engineering and Mathematical Sciences,
University of the West of England, Bristol, BS16 1QY, UK

Abstract

Steady State models of Evolutionary Algorithms are widely used, yet surprisingly little attention has been paid to the effects arising from different replacement strategies. This paper explores the use of mathematical models to characterise the selection pressures arising in a selection-only environment. The first part brings together models for the behaviour of seven different replacement mechanisms and provides expressions for various proposed indicators of Evolutionary Algorithm behaviour. Some of these have been derived elsewhere, and are included for completeness, but the majority are new to this paper. These theoretical indicators are used to compare the behaviour of the different strategies. The second part of this paper examines the practical relevance of these indicators as predictors for algorithms' relative performance in terms of optimisation time and reliability. It is not the intention of this paper to come up with a "one size fits all" recommendation for choice of replacement strategy. Although some strategies may have little to recommend them, the relative ranking of others is shown to depend on the intended use of the algorithm to be implemented, as reflected in the choice of performance metrics.

Keywords

Replacement Strategies, Selection, Steady State.

1 Introduction

As the commonly used "Natural Selection" metaphor suggests, the selection mechanism in Evolutionary Algorithms (EAs) is the driving force behind the hoped-for improvement in the fitness distribution of the population. Essentially its role is to take advantage of the fitter individuals produced by the reproductive operators (recombination and mutation) and increase their relative frequency in the population, so that they are more likely to be chosen as parents during the next round of reproduction. What distinguishes EAs from many biological models is the use of a fixed population size. This enforces a split of the selection mechanism into two phases, namely parental selection and replacement strategy. For most Generational EAs (GEAs), (i.e., except those using some variant of an elitist strategy such as $(\mu + \lambda)$), the latter is simple: all members of the previous population are deleted, and if good solutions are to be preserved and propagated then this burden falls upon the reproductive operators. This greatly simplifies the task of producing analyses of the expected behaviour of the algorithm using indicators such as Takeover Time (Goldberg and Deb, 1991), Selection Intensity (Bäck, 1995; Blickle and Thiele, 1997), Genetic Drift (Rogers and Prügel-Bennett, 1999a) and Expected Loss of Diversity (Motoki, 2002). For Steady State EAs (SSEAs), the choices of replacement strategy may be split into two classes: age-based (as per GEAs) and fitness-based.

1.1 Related Work

Rogers and Prügel-Bennett (1999b) showed that an SSEA using Boltzmann selection for both parents and replacements demonstrated twice the rate of change of a GEA using the same selection method for parents. This was both in terms of the increase in the mean of the population's fitness distribution, and of the decrease of its variance. In the case of random deletion, the GEA and SSEA were shown to be (to a good approximation) equivalent in terms of the rate of increase in mean fitness, but the non-fitness-dependent term in the rate of variance decrease (Genetic Drift) was double in the SSEA. This "statistical mechanics" approach as espoused in (Prügel-Bennet and Shapiro, 1994) and subsequently developed in many other papers by those authors are valuable because they permit comparisons to be made within the context of an evolutionary algorithm with variation operators on real fitness function. However although these are important results in their own right, the techniques used are not easily adaptable to many commonly used replacement methods. This is because they either rely on the assumption that members are not picked to be replaced from the extrema of the fitness distribution, and/or have no way of taking into account the age of the solution. The former precludes modelling well known algorithms such as GENITOR (Whitley and Kauth, 1988), the latter strategies such as First-In-First-Out (FIFO).

Other approaches which model EAs with variation and selection operators acting on a specific fitness function include the "dynamic systems" approach (as described in (Vose, 1999) and further developed by Vose and co-workers) and the "statistical dynamics" approach (van Nimwegen et al., 1999; van Nimwegen and Crutchfield, 2001). Both approaches model an EA as the action of an "Evolutionary Operator" $\mathcal{G} = \mathcal{M} \circ \mathcal{F}$ acting on a vector which in some way represents the current population (or probability distributions thereon). \mathcal{M} represents the combined action of mutation and recombination, and \mathcal{F} represents the action of selection, in most cases explicitly incorporating information about the fitness function. The approaches consider the fixed points of the operator \mathcal{G} . Unfortunately at present there exists no general method for finding these when \mathcal{G} is nonlinear, for example when \mathcal{M} includes crossover as well as mutation or when \mathcal{F} is complex. For this reason analytical approaches have tended to use a generational mutation-selection model with fitness-proportionate selection (for which van Nimwegen (1999) presents a linearisation). Although Wright (1999) extended the dynamical systems model to include truncation selection in non-generational models, considering a wider range of replacement strategies in SSEAs presents the same problems of finding fixed points.

Thus while these approaches have much to offer, for various reasons none of them currently provide the ability to conduct a widespread comparison between different replacement methods. Other descriptors of behaviour are more amenable to analysis, but tend to only consider one facet of the EA in isolation. Some of these are rooted in a pragmatic dimensional analysis approach to understanding EAs, whereby the key facets affecting behaviour are identified and their effects examined in isolation. Goldberg (2002) describes such a philosophy of using combinations of indicators when designing EAs to tackle specific problems. Thus different indicators describing selection might be considered in conjunction with indicators describing other factors—such as the "mixing time" (Thierens and Goldberg, 1993). As will be seen, a variety of methods have been proposed to characterise the effect of selection in isolation, and permit comparison between different operators. Part of the purpose of this paper is to present a preliminary study of whether these have any predictive value for the observed behaviour of EAs in which all other details are left unchanged.

One common indicator used to examine the affect of different selection operators in this kind of analysis is the takeover time (or probability), which reflects the expectation of the time taken (or likelihood) for a single fitter individual to take over a finite population under selection and replacement alone (Goldberg and Deb, 1991). Chakraborty (1995; 1997) used a Markov Chain analysis to model takeover times for a number of selection mechanisms in both GEAs and SSEAs. Using a numerical approach with fixed population sizes, results were presented for SSEAs using Replace–Worst, Replace–Random and deletion by exponential ranking. Smith and Vavak (1999a) used this model to examine the behaviour of SSEAs using several other replacement schemes, again for a single fixed population size. More recently Rudolph (2000; 2001) generalised this model to derive some analytical results for takeover times and probabilities as a function of population size.

Rudolph has also argued that takeover time alone is not a sufficient indicator of the selection pressure exerted. An indicator η representing the mean fraction of non-best individuals available as parents prior to takeover was proposed in (Rudolph, 2000). In a similar vein, following Blickle and Thiele (1997), Motoki (2002) proposed the “Expected Loss of Diversity” indicator to quantify the change in the diversity between a given population and the (same-sized) set of parents selected from it. However this was not extended to cover SSEAs, for reasons which will be discussed later.

1.2 Goals and Contributions

This paper has two principal aims. The first is to bring together models for the behaviour of seven commonly used replacement mechanisms and derive expressions for the indicators mentioned above, i.e., Takeover Time, η , and the Expected Loss of Diversity. Since first moments (means) on their own are often insufficient to discriminate between different distributions, analytic expressions for the variance of takeover times are also provided. These theoretical indicators are used to compare the behaviour of the different strategies. The second aim of this paper is to examine their practical relevance as predictors for EA performance in terms of optimisation time and reliability. It must be emphasised that it is not the intention of this paper to come up with a “one size fits all” recommendation for choice of replacement strategy. As will be seen, although some strategies may have little to recommend them, the relative ranking of others will largely depend on the intended use of the EA to be implemented, as reflected in the choice of performance metrics.

Rudolph (2000; 2001) has already produced analytic expressions for the expected takeover times for the following replacement strategies: k –ary Select-Replace Tournaments, Replace-Worst (GENITOR), Elitist-Random-Replacement, and Elitist Kill-Tournaments, together with bounds on the takeover probabilities for the non-elitist versions of the latter two strategies. In addition he has produced expressions for the value of η for k –ary Select-Replace Tournaments and Replace-Worst (GENITOR). These results are included for the sake of completeness. Note that “elitist” in this sense means that the final member of the fitter class is never permitted to be deleted.

All other analytic results, namely the expected takeover times and η for the various Conservative and the Elitist-FIFO strategies, η for Elitist versions of Replace Random and Kill tournaments, the bounds on the takeover probabilities for the strict FIFO strategy, all the expressions describing the variance in takeover time, and all the expressions describing the Expected Loss of Diversity are new to this paper, as is the statistical analysis of the predictive power of the various indicators.

1.3 Organisation of This Paper

The rest of this paper proceeds as follows. Section 2 briefly reviews the principal mathematical techniques used within this paper and formally defines the quantities of interest: takeover time, diversity and expected loss of diversity. Section 3 contains descriptions and models of the different algorithms considered. This section is divided into two parts. Section 3.1 discusses the replacement strategies which guarantee takeover of the fittest individual in a selection-replacement only regime. Section 3.2 then considers those algorithms where the fittest member may be lost due to the stochastic nature of the replacement strategy. Section 4 provides a summary and comparison of the values of various indicators for the replacement methods. Following this, Section 5 details results from a number of experiments designed to test the relative merit of the different indicators as predictors of EA behaviour. Finally in Section 6 these results are discussed and some conclusions drawn about the merits of different strategies and indicators.

2 Modelling Selection with Markov Chains

2.1 Background

The use of Markov Chains to model different aspects of EA behaviour is now well established in various contexts. The sheer size of the full models of Nix and Vose (1992) renders them generally intractable for non-trivial systems. However various papers have used Markov chains to model different aspects of EA behaviour in isolation, such as genetic drift in the absence of selection, e.g. (Asoh and Mühlenbein, 1994; Goldberg and Segrest, 1987). A frequently studied indicator is the “Takeover Time” (Goldberg and Deb, 1991), namely the number of selection events needed for a single member of an fitter class to take over a population. Rudolph (2000) proposed an indicator that characterises the fraction of non-optimal solutions evaluated before takeover, as an alternative tool for predicting the merit of an operator in an optimisation setting.

These approaches divide the population of an EA into two equivalence classes of different fitness. They model the evolution of the proportion belonging to the fitter class as a function of time and population size N . In the context of SSEAs each discrete time-step consists of a parent being chosen by some parental selection mechanism, and a copy of it being made to put in to the population according to the replacement strategy used. This analysis demonstrates the ability of the selection and replacement operators to exploit fitter individuals, and also the degree to which they are prone to stochastic effects, namely the loss of (possibly all) members of the fitter class.

Each selection-replacement event can be considered to occupy a single discrete time-step, so the evolution of the population can be modelled as a random sequence $X(t) : 0 \leq X(t) \leq N$, where $X(t)$ denotes the number of members of the population belonging to the fitter class at time t . Since only one member of the population is changed at a time, $X(t) - X(t - 1) \in \{-1, 0, 1\}$ and the system is effectively a random walk. For most selection or replacement operators (all but FIFO) the probability that $X(t + 1)$ takes some value j depends solely on $X(t)$, and so the random sequence $X(t), X(t + 1), \dots$, is a time-homogeneous Markov Chain with $N + 1$ discrete states represented by the different values of $X(t)$. The corresponding transition matrix P is either bi- or tri-diagonal with non-zero elements defined by the selection and replacement strategies. Given P it is possible to calculate the n -step transition matrix P^n which defines the transition probabilities after n steps.

As there is no introduction of variety due to reproduction $P_{N,N-1} = P_{0,1} = 0$. In terms of random walks, the state $X(t) = N$ is termed an absorbing barrier. For any

non-zero value $X(t) = i$, there is a positive probability that $X(t + N - i) = N$ by successive increases, hence all states $X(t) \in \{1, \dots, N - 1\}$ are transient. In the general case $X(t) = 0$ is an absorbing barrier. In the special case of algorithms which can never lose the final copy of the fittest solution, $P_{1,0} = 0$, and so $X(t) = 0$ is not accessible from any other state, and is effectively a reflecting barrier.

2.2 Mean and Variance of Takeover Times

For any given instance of the random walks that are considered here, the takeover time T is defined as $T = \min(t : X(t) = N \mid X(0) = 1)$, with corresponding expectation $E(T)$ and variance $V(T)$.

Several of the methods described do not permit the replacement of a member of the fitter class with a less fit one, so $X(t)$ is non-decreasing, and each of the transitions from state i to state $i + 1$ are independent events. If $T_{i,i+1}$ denotes the number of time steps taken between transitions, then this variable is geometrically distributed with mean $1/P_{i,i+1}$ and variance $(1 - P_{i,i+1})/P_{i,i+1}^2$. Thus starting from one copy, the takeover time T has mean and variance:

$$E(T) = \sum_{i=1}^{N-1} \frac{1}{P_{i,i+1}}, \quad V(T) = \sum_{i=1}^{N-1} \frac{1 - P_{i,i+1}}{P_{i,i+1}^2} = \sum_{i=1}^{N-1} \left(\frac{1}{P_{i,i+1}} \right)^2 - E(T). \quad (1)$$

In the general case when $X(t)$ can decrease, the method above is not valid and more complex methods are necessary. Considering a random walk with two absorbing barriers, the first and last rows and columns of the corresponding transition matrix P concern the absorbing states $X(t) \in \{0, N\}$. If these are removed from P , the remaining matrix S gives the transition probabilities for movements between the transient states. The fundamental matrix is defined as $Q = (I - S)^{-1}$, with I an identity matrix of appropriate dimension. This fundamental matrix has the property that $Q_{i,j}$ is the expected number of occurrences of the transient state j prior to absorption if the system starts in the transient state i . Noting that by convention $X(0) = 1$, the expected number of time-steps prior to absorption in the general case is given by the sum of the number of occurrences of each of the transient states, and hence for elitist algorithms with $P_{1,0} = 0$:

$$E(T) = \sum_{j=1}^{N-1} Q_{1,j}. \quad (2)$$

Helpfully, expressions for these terms for a random walk with absorbing barriers, which can be used to derive $E(T)$ are given in (Rudolph, 1999). Derivation of the variance $V(T)$ in the general cases is more complex. It is possible to derive from first principles expressions for the variance of the number of occurrences of each state, based on the probability that a state j is ever reached from some state i . These probabilities can in turn be derived from the entries in the fundamental matrix. However, in the general case with more transitions between states it is not reasonable to consider the number of occurrences of each state as independent events. Therefore in this case calculating $V(T)$ simply as the sum of these single-state variances, without taking into account the covariance between them, will lead to overestimates. Given these qualms, and the long-winded nature of the calculations, derivation of $V(T)$ in the general case remains for future work.

2.3 Takeover Probabilities

For those replacement strategies in which there is a positive probability for even the best individuals to disappear entirely, it is necessary to talk about takeover probabilities rather than takeover times. For a starting configuration of the EA with $X(0) = i$, the probability that $X(t) = N$, (i.e., that the class has taken over the population at time t) is $P_{i,N}^t$. By convention $X(0) = 1$ and so the takeover probability at time t is defined to be $P_{1,N}^t$. Rather than calculating this explicitly, note that the state N can only be reached directly from state $N - 1$, hence the probability of takeover is $P_{N-1,N} \cdot Q_{1,N-1}$. As noted above, expressions for the elements of Q are given in (Rudolph, 1999), and for the takeover probabilities of three strategies in (Rudolph, 2000).

2.4 The Diversity Indicator η

Rudolph (2000) defined his diversity indicator η as “the mean fraction of non-best individuals that were available for the generation of candidate solutions prior to absorption.” The general form can be stated using the expected number of occurrences of each state as:

$$\eta = \frac{\sum_{j=1}^{N-1} (N-j) \cdot Q_{1,j}}{N \cdot \sum_{j=1}^{N-1} Q_{1,j}} = 1 - \frac{\sum_{j=1}^{N-1} j \cdot Q_{1,j}}{N \cdot E(T)} \quad (3)$$

This was formulated in terms of a variable B_T , denoting the total number of members of the fitter class prior to absorption, which has an expected value equivalent to the numerator of the second term in Eq. 3. It is shown that matters can be simplified if the symmetry property holds i.e., $P_{i,i+1} = P_{n-i,n-i+1}$, in which case $\eta = 0.5$. For the general non-decreasing case the system may be thought of as a series of independent transitions between states $i, i+1, i \in \{1, \dots, N-1\}$, so each of these has expected duration $Q_{1,i} = 1/P_{i,i+1}$.

2.5 An Alternative Indicator: Expected Loss of Diversity

Following Blickle and Thiele (1997), Motoki (2002) defined D_x , the Expected Loss of Diversity, as the expectation of the difference in the frequency of occurrence of individuals between the population at a given time-step, and the pool of parents resulting from the application of a selection operator x . In that paper various rank-based selection methods for generational EAs are compared under the assumptions that each individual parent has a different fitness, parents are selected one at a time with replacement, and there is no overlap between generations. Under these circumstances $D_x(N) = \sum_{k=1}^N (1 - p_{s,x}(k))/N$, where $p_{s,x}(k)$ is the probability of selecting an individual of rank k under any given scheme x .

In a steady-state model, the offspring will generally be inserted into the population. This will change the rank of some of the surviving parents, and in a selection-only model will violate the assumption of distinct fitnesses. Therefore it is not possible to consider the effects of N selection-replacement events, which would permit comparison between steady-state and generational schemes. However it is relatively straightforward to consider the expected effect of a single event, and the single-step expected loss of diversity will be defined as:

$$D_{x,y}(N) = \frac{1}{N} \cdot \sum_{k=1}^N (1 - p_{s,x}(k)) \cdot p_{d,y}(k) \quad (4)$$

where $p_{d,y}(k)$ denotes the probability of deleting an individual of rank k under replacement scheme y . For the purposes of comparison the same parent selection method x is

used for each replacement strategy. Appendix D contains the derivations of $D_{x,y}(N)$ for various schemes y , and the results for the elitist schemes are shown in Table 3.

3 Description and Models of Algorithms

This section defines the different replacement strategies under consideration, and provides details of the Transition Matrix P for the analysis of the Markov chain. Where applicable this is done both for the general case (using a term p_s to denote the probability that a member of the fitter class is chosen for copying) and then for the case of deterministic binary tournaments (DBT) for parent selection. In the latter this probability will be denoted as $p_{s(i)}$ for a population containing i members of the fitter class, and is given by: $p_{s(i)} = i \cdot (2N - i) / N^2$. In all cases $P_{N,N} = P_{0,0} = 1$ and $P_{N,N-1} = P_{0,1} = 0$.

3.1 Non-Decreasing Methods

The first class of methods are *Non-Decreasing*, i.e., $X(t + 1) \geq X(t)$ and the matrix P is bidiagonal with $P_{i+1,i} = 0$, $P_{ii} = 1 - P_{i,i+1} \quad \forall i \in \{0, \dots, N - 1\}$. As noted above, in these cases $E(T)$, $V(T)$ and η can be fairly simply derived in terms of the elements $P_{i,i+1}$ of the transition matrix.

3.1.1 k -ary Select-Replace Tournaments

This algorithm was suggested by Rudolph (2000). Here k individuals are chosen (with replacement) from the population, and a copy of the fittest one replaces the least fit. In other words the parent-selection and replacement strategies are merged. In this case the number of copies of the fitter class will increase unless all k members of the tournament belong to the same class, so the relevant entries of the transition matrix are:

$$\begin{aligned} P_{i,i+1} &= 1 - (1 - i/N)^k - (i/N)^k & i \in \{1, \dots, N - 1\}, & \quad k \in \{1, \dots, N\} \\ &= 2i(N - i)/N^2 & i \in \{1, \dots, N - 1\}, & \quad k = 2. \end{aligned} \quad (5)$$

3.1.2 Replace Worst

This is the replacement strategy used in Whitley's original Steady-State Genetic Algorithm GENITOR (Whitley and Kauth, 1988) and often referred to by that name. In this case the current least fit member of the population is replaced. Thus a member of the less fit class will always be replaced until takeover has occurred, and so the defining elements of P are given by:

$$\begin{aligned} P_{i,i+1} &= p_s & i \in \{1, \dots, N - 1\}, & \quad \text{general case} \\ &= i(2N - i)/N^2 & i \in \{1, \dots, N - 1\}, & \quad \text{DBT.} \end{aligned} \quad (6)$$

3.1.3 Conservative-FIFO Selection

This operator was presented in Vavak et al. (1997) for use in noisy and/or non-stationary environments. It combines a First-In-First-Out (FIFO, also known as "Delete-Oldest") replacement strategy with a modified deterministic binary tournament selection operator. In this "Conservative" selection, the choice of member to be replaced is made *before* the parents are chosen. This member takes part in all of the tournaments to become parents at that iteration, the other candidates being chosen at random. The number of these tournaments is one for the select-replace model, and is otherwise determined by the arity of the recombination operator. Since the tournaments are deterministic, elitism is implicitly assured under selection alone, and also in

the presence of “respectful” recombination (Radcliffe, 1991), although not in the presence of mutation or other variation operators.

With a FIFO strategy the probability of deleting one of the members of the fittest class depends on when they were created, so the transition probabilities are not time-independent. Therefore the progression of $X(t)$ for this replacement strategy cannot be exactly modelled as a Markov Chain with $N + 1$ states, rather, for any given value $X(t) = i$, the system can be in any one of $N!/i!(N - i)!$ states, depending on the age-distribution of the fitter members. The process does form a Markov chain and the transition probabilities between them can be derived, but this remains for future work. However, as shown by Smith and Vavak (1999a), it is possible to get a prediction which is a close match to the empirical behaviour using a Markov model based on the following approximation. Let p_{old} denote the probability that the oldest member is in the fitter class, and note that in this case it will win the tournament, so $X(t)$ is non-decreasing, and can only increase if the oldest is not a member of the fitter class and the other (randomly chosen) candidate in the tournament is. In Appendix A.1 a general approximation for p_{old} is derived under the assumption that the value decreases linearly with the age of the member, so that the Markov condition holds. Equation A-2 in Appendix A.2 gives the specific form for this modified binary tournament selection, so the defining elements of P are given by:

$$\begin{aligned}
 P_{i,i+1} &= (i/N) \cdot (1 - p_{old}) & i \in \{1, \dots, N - 1\} & \text{general case} \\
 &= 2i(N - i)/(2N^2 - iN) & i \in \{1, \dots, N - 1\} & \text{linear approximation.} \quad (7)
 \end{aligned}$$

3.1.4 Conservative–Random

Smith and Vavak (1999a) tested the conservative selection operator using a random rather than age-based choice of the member to be replaced. Again, $X(t)$ will not change unless a less fit member is chosen to be replaced, and the other participant in the tournament is a member of the fitter class, yielding:

$$\begin{aligned}
 P_{i,i+1} &= p_s \cdot (1 - p_s) & \text{general case} \\
 &= i(N - i)/N^2 & \text{DBT} \quad (8)
 \end{aligned}$$

since both participants in this binary tournament are chosen at random. Note that this transition probability is exactly half that of binary selection-replacement tournaments shown in Eq. 5, since in this method the first randomly drawn member is replaced, rather than the less fit of the two. There is another difference that occurs with recombinative algorithms, since in one case the algorithm makes independent random drawings for each member of each tournament, whereas in the conservative tournaments the member selected to be replaced (either according to age or at random) takes part in each tournament.

3.2 Decreasing Methods

Into this class of strategies fall those algorithms where the value of $X(t)$ may decrease as well as increase, i.e. $X(t + 1) - X(t) \in \{-1, 0, 1\}$, and the corresponding transition matrices are tridiagonal. In general these methods have the potential to lose the sole copy of the fittest class, so it is no longer meaningful to talk about takeover times as the system can be absorbed into either state $X(t) = 0$ or $X(t) = N$. Instead, it is more appropriate to consider the probability that takeover actually occurs.

As with generational EAs, the use of elitism is frequently incorporated—that is the algorithms are modified so that the sequence $X(t) = 1, \dots, 1, 0$ never occurs. In terms

of the modified transition matrix, $P_{1,0} = 0$ and since $P_{1,2}$ is the same, $P_{1,1} = 1 - P_{1,0} - P_{1,2}$ is larger, but all other elements of P are unchanged. Rudolph (2001) derived the elements $Q_{1,j}$ of the fundamental matrix for elitist versions of Replace-Random and Kill Tournaments, which using Eq. 2 yield $E(T)$ for these strategies. Details of these calculations, plus new derivations of η are given in Appendix B.

3.2.1 Replace Random

This strategy relies on a random choice of member of the population to be replaced at each time-step. Since the population size is finite the mean life expectancy of any given individual is N time steps, as it is for the FIFO strategy. However there can be severe variance, with consequent effects for the behaviour of the algorithm as noted by DeJong and Sarma (1992). Although it has been much studied theoretically, primarily due to the simplicity with which it can be modelled, empirically there seems to be little evidence in its favour—see for example the poor results reported with this algorithm compared to other strategies (Smith and Vavak, 1999a).

Since the probability of picking a member for replacement is simply i/N for the fitter class, and $(N - i)/N$ for the others, the defining elements of P are given by:

$$\begin{aligned}
 P_{i,i+1} &= p_s \cdot (N - i)/N & i \in \{1, \dots, N - 1\} & \text{general case} \\
 &= i(2N - i)(N - i)/N^3 & i \in \{1, \dots, N - 1\} & \text{DBT} \\
 P_{i,i-1} &= (1 - p_s) \cdot i/N & i \in \{1, \dots, N - 1\} & \text{general case} \\
 &= i(N - i)^2/N^3, & i \in \{1, \dots, N - 1\} & \text{DBT.} \quad (9)
 \end{aligned}$$

Rudolph (2000) has shown via the fundamental matrix route that for DBT that the takeover probability is $0.5 + 1/2N$.

3.2.2 Kill Tournament

The strategy of selecting a member for replacement by a “kill tournament” has many of the advantages of a Linear Ranking strategy, without the need for continually resorting the population according to fitness, and a number of possibilities arise. As for the parent selection tournament, it is possible to pick a number of parents at random and replace the worst, or to pick two random parents and replace the worse with some probability $0.5 \leq d \leq 1.0$. Thus a range of selection pressures can be achieved. Other authors have suggested tournaments between the parents and their offspring (see Thierens (1997) for a good discussion).

Assuming that the participants in the kill tournament are selected randomly with replacement, the probability p_{kill} of deleting a member of the fitter class will again be a function of $X(t)$ and is defined by analogous expressions to those for parent selection. For a k -ary tournament with the worst always replaced, all of the participants in the tournament must be of the fitter class if one is to be replaced, i.e., $p_{kill} = (i/N)^k$. For the general case with binary tournaments this probability is given by:

$$\begin{aligned}
 p_{kill} &= i^2/N^2 + 2i(N - i)(1 - d)/N^2 \\
 &= (2d - 1) \cdot i^2/N^2 + (2 - 2d) \cdot i/N. \quad (10)
 \end{aligned}$$

The defining elements of P are therefore:

$$\begin{aligned}
 P_{i,i+1} &= p_s \cdot (1 - p_{kill}) & i \in \{1, \dots, N-1\} & \text{general case} \\
 &= i(2N - i)(N^2 - i^2)/N^4 & i \in \{1, \dots, N-1\} & \text{DBT.} \\
 P_{i,i-1} &= p_{kill} \cdot (1 - p_s) & i \in \{1, \dots, N-1\} & \text{general case} \\
 &= i^2(N - i)^2/N^4 & i \in \{1, \dots, N-1\} & \text{DBT.} \quad (11)
 \end{aligned}$$

Using the fundamental matrix does not yield an analytical expression in this case, but the bounds $(1 - 1/N) \leq p_{takeover} \leq (1 - 1/5N)$ are derived in (Rudolph, 2000).

3.2.3 First-In-First-Out (FIFO)

One simple method of avoiding some of the variability induced by a replace-random strategy is to ensure an identical life-expectancy for each member of the population. This can be simply achieved via an age-based replacement strategy, as is used in the canonical Simple Genetic Algorithm. Smith and Vavak (1999b) reported that this sort of strategy has advantages on dynamical optimisation problems since it provides a simple means of systematically reducing the retention of outdated fitness information.

As for the Conservative-FIFO strategy, the age-dependency means that the variable $X(t)$ does not satisfy the Markov condition, but it is possible to derive a reasonable model by assuming a form for p_{old} that is linear in $X(t)$ in which case the defining elements of P in the generic form are:

$$\begin{aligned}
 P_{i,i+1} &= p_s \cdot (1 - p_{old}) \\
 P_{i,i-1} &= p_{old} \cdot (1 - p_s) \quad (12)
 \end{aligned}$$

In Appendix A.1 an approximation $p_{old} = 2i/N - p_s$ is derived based on this assumption. If a stochastic binary selection tournament is used with a FIFO strategy, the expression derived in Eq. A-1 reduces to $p_{old} = 2i(1 - s)/N + (2s - 1)i^2/N^2$, where s is the probability of accepting the fitter member. Comparing this with the equivalent probability for a stochastic binary kill tournament, (p_{kill} in Eq.10), shows that the two are the same provided $s = d$. This shows that *provided the linear model is accurate*, when using a binary selection tournament with a probability $0.5 \leq s \leq 1.0$ of selecting the fitter, a FIFO policy is equivalent (at least in terms of takeover) to a binary kill tournament where the worst is deleted with the same probability.

Smith and Vavak (1999a) explored this approximation for a single fixed population size and found it to be accurate apart from minor deviations in the behaviour during the initial phases (low values of $X(t)$). In Appendix C bounds are derived for the takeover probability which do not make the assumption of linearity. These have been confirmed experimentally, and although initially high (for low N), they decrease with N , rapidly converging to $1 - 2e^{-2} + e^{-4}$ for the lower bound and $1 - e^{-2} - 2e^{-4}$ for the upper. This rapid decrease with N to fixed limits is also seen with Replace-Random, but not with kill tournaments, where the limits only converge asymptotically to 1.

It is straightforward to modify the FIFO strategy so that the last copy of the fitter class is not replaced by a less-fit individual, and Smith and Vavak (1999a) demonstrated that under these conditions, the linear approximation, and hence for binary tournament selection the equivalence to Kill Tournaments was accurate. Hence the takeover times and η in this case are the same as for elitist Kill Tournaments.

Table 1: Comparison of Indicators for different strategies based on takeover analysis. Columns from left are strategy, expectation and variance of takeover time. “-” indicates value not available, H_N is the N^{th} harmonic number, $\zeta(N)$ is Riemann’s zeta function.

Method	$E(T)$	$V(T)$
Worst	$\frac{1}{2} \cdot N \cdot H_{2N-1} - \frac{1}{2}$	$\frac{1}{4} \cdot N^2 \cdot \zeta_{2N-1}^2 - \frac{N}{4} H_{2N-1}$
Binary SR	$1 \cdot N \cdot H_{N-1}$	$\frac{1}{2} \cdot N^2 \cdot \zeta_{N-1}^2$
E-Tourn.	$1 \cdot N \cdot H_{N-1} + \frac{1}{2} \cdot H_{N-1}$	-
E-FIFO	$1 \cdot N \cdot H_{N-1} + \frac{1}{2} \cdot H_{N-1}$	-
C-FIFO	$\frac{3}{2} \cdot N \cdot H_{N-1} - \frac{N}{2N-2}$	$\frac{5}{4} \cdot N^2 \cdot \zeta_{N-1}^2 + \frac{N}{2} H_{N-1} - \frac{4N^3 - 5N^2 + 2N}{4(N-1)^2}$
E-Rand	$2 \cdot N \cdot H_{N-1} - \frac{2N}{N+1} \cdot H_{N-1}$	-
C-Rand	$2 \cdot N \cdot H_{N-1}$	$2 \cdot N^2 \cdot \zeta_{N-1}^2$

4 Comparison of Indicators

4.1 Mean and Variance of Takeover Time

Details of the derivation of $E(T)$, $V(T)$, and η for those methods which guarantee absorption to $X(t) = N$ are given in Appendix B. For the sake of brevity, from this point on the prefix C will be used to denote the Conservative selection tournaments, and E to denote the elitist variant of the algorithms. “Kill Tournaments” will be abbreviated to Tourn, Replace-Worst to Worst, and Replace-Random to Rand.

Table 1 shows the expectation and variance of takeover time for these strategies ranked by increasing $E(T)$. The expression $H_N = \sum_{i=1}^N 1/i$, which is termed the N^{th} Harmonic number, is bounded by $\log(N) < H_N < \log(N) + 1$. $\zeta_N^2 = \sum_{i=1}^N 1/i^2$ denotes Riemann’s Zeta function with exponent 2 truncated to its first N terms. Since all terms are positive this converges asymptotically, and the limit is $\pi^2/6$.

In order to verify these predictions 1000 runs were made with each algorithm for each population size in the range $\{2, \dots, 1000\}$. These results confirmed the predictions when exact models were available. Figure 1 shows a comparison of the predictions with experimental observations for those strategies with a FIFO element. Specifically it contrasts $E(T)$ with the observed mean takeover time for C-FIFO and E-FIFO, and the positive square root of $V(T)$ with the observed standard deviation of the takeover time for C-FIFO. The observed standard deviation of takeover time for E-FIFO is also shown for completeness.

As can be seen the observed variance in takeover times confirms the predicted $V(T)$ for C-FIFO, and the observed values for the two strategies are very similar. In both cases the predicted $E(T)$ over-estimates the observed mean takeover time, but the discrepancy in each case is less than one standard deviation throughout the spread of values of N tested. Under these circumstances a hypothesis test at any given value of N would be highly unlikely to measure these differences as significant. Therefore




Figure 1: Comparisons of predictions of Mean, $(E(T))$, and Variance $(V(T))$, of Takeover Times as a function of Population Size, N for the C-FIFO and -FIFO strategy according to the linear model and experimental results.

Figure 1 would appear to lend weight to the use of the linear approximation for p_{old} , especially for predicting the variance in takeover time.

The results for $E(T)$ in Table 1 show that although they are all $\Theta(N \log(N))$, the actual takeover times vary by as much as a factor of 4. The fastest to converge are those methods that use fitness as a criterion for selecting members to be replaced, followed by those using a FIFO strategy, then those using a random strategy. For both the FIFO and Rand strategies, the Elitist versions have faster takeover than the Conservative variants. The same pattern holds for the variance where available, although the difference in magnitude between Worst and C-Rand is now eightfold.

By way of comparison with these values for $E(T)$, Goldberg and Deb (1991) derive the expression $t^* = (\log(N) + \log(\log(N)))/\log(2)$ for the takeover time (in generations) of a generational EA with the same deterministic binary tournaments for parental selection. Assuming N selection and replacement events per generation, noting that $\log(2) \approx 2/3$, and using the upper bound of H_{N-1} , gives $E(T) \approx 3N/2 \cdot H_{N-1} + 3N \cdot (\log(\log(n)) - 1)/2$. In terms of the ranking by $E(T)$, this is close to halfway between C-FIFO and E-Rand (this has been verified experimentally).

In a "canonical" GEA with age-based replacement, the population is completely replaced by its offspring in each generation. Thus each member of the population potentially participates in exactly N selection tournaments before it is replaced. The same is true for a SSEA with the FIFO strategy, and on average for the Rand strategy. Therefore one might intuitively expect the takeover times to be the same in these three cases. In practise the stochastic nature of Random replacement means that the takeover time is 1/3 larger for this type of SSEA than for the corresponding GEA, since as $X(t)$ increases the probability of randomly selecting a less-fit member to be replaced diminishes, unlike the systematic replacement in age based schemes such as GEAs or SSEAs with FIFO replacement.

It has long been hypothesised that SSEAs might yield performance advantages

over GEAs because they are able to utilise good solutions as soon as they are made, rather than waiting for the next generation for them to become parents. Since E-FIFO has no explicit fitness component in the choice of member to be replaced, the fact that its expected takeover time $E(T)$ is only $2/3$ that for the corresponding GEA is a good illustration of the ability of the SSEA to make copies of high-fitness solutions sooner. It could be argued that this faster takeover occurs purely as a result of genetic drift, rather than any fitness-based effect. However, the close match between observed and predicted behaviour, and the fact that the linear approximation for p_{old} suggests that E-FIFO and E-Tourn (which definitely is fitness-biased) are equivalent belies this suggestion. The reason for this is that the linear approximation implicitly assumes a constantly increasing mean fitness *within* the equivalent of a generation of a GEA.

Finally, note that the takeover time of E-Tourn is half that of E-Rand. This suggests a direct relationship between $E(T)$ and the rate of increase in mean fitness which is cited by Rogers and Prügel-Bennet (1999) as being twice as high for fitness-based replacement as for Random under their particular selection conditions (see Section 1.1).

When designing an EA to be utilised in real-world application, factors other than the average speed of finding good solutions may be important, and it is worth considering the other information these analytic results contain. One type of application for EAs is in situations where they are used repetitively on similar problems—for example doing daily scheduling, vehicle routing, production planning or other management tasks. In such cases it may be more important that the EA reliably finds a good solution in a fixed amount of time, rather than occasionally finding very good solutions, but at the cost of unpredictable run-times. Under these circumstances it is clear that $V(T)$, the variance in takeover time, may be an important factor in choosing between replacement strategies, since the difference between strategies is up to eightfold.

4.2 Mean Fraction and Number of Less-Fit Points Selected Prior to Takeover

Table 2 shows the predicted fraction and number of non-best individuals for different strategies, ranked by increasing number. Rudolph (2000) argues that η may be a better predictor of EA behaviour than $E(T)$, since the latter does not take into account the diversity of solutions produced. However, possibly more relevant to the discovery of global optima on complex landscapes is not the fraction, but the *total* number of different solutions evaluated prior to takeover. As can be seen from Table 2, many of the strategies have $\eta = 0.5$, that is the fraction of non-best individuals selected may be equal, but the total number of such individuals can vary by as much as a factor of two between different algorithms.

The rank order of the different strategies is largely the same as that by $E(T)$ with the following two exceptions.

- The Worst strategy, despite having the fastest takeover time, can select more non-best points than the Binary SR, E-Tourn and E-FIFO strategies.
- C-FIFO is now very close to the maximal values observed for the Rand strategies, rather than lying midway between them and the E-Tourn/E-FIFO ones.

4.3 Expected Loss of Diversity

Details of the calculations of Expected Loss of Diversity (D) per select-replace event are given in Appendix D and the results are shown in Table 3, with the strategies ranked by decreasing D ; that is by increasing preservation of diversity. These are dominated by the co-efficient of the term in $1/N$. Thus the various strategies again fall into two

Table 2: Comparison of Diversity Indicators for different strategies, ranked by $\eta \cdot E(T)$. Columns from left are strategy, mean fraction (η) and number ($\eta \cdot E(T)$) of non-best individuals selected prior to takeover .

Method	η	$E(T) \cdot \eta$
Binary SR	$\frac{1}{2}$	$\frac{1}{2} \cdot N \cdot H_{N-1}$
E-Tourn.	$\frac{1}{2}$	$\frac{1}{2} \cdot N \cdot H_{N-1} + \frac{1}{4} \cdot H_{N-1}$
E-FIFO	$\frac{1}{2}$	$\frac{1}{2} \cdot N \cdot H_{N-1} + \frac{1}{4} \cdot H_{N-1}$
Worst	$\approx 1 - \frac{2 \log(2)}{\log(2N)}$	$\approx \frac{1}{2} \cdot N \cdot H_{2N-1}$
C-FIFO	$\frac{2}{3} - \frac{6N^2 - 4N - 3}{18N^2(N-1)H_{N-1} - 6N^2}$	$\approx 1 \cdot N \cdot H_{N-1} - \frac{N}{3(N-1)}$
E-Rand	$\frac{1}{2}$	$\frac{N}{N+1} \cdot N \cdot H_{N-1}$
C-Rand	$\frac{1}{2}$	$1 \cdot N \cdot H_{N-1}$

camps, the two conservative strategies (with C-FIFO this time higher than C-Rand) having lower D values than the rest. Notably, the loss of diversity is greater for Worst than all the other strategies, in contrast to the values for η and $\eta \cdot E(T)$. It is worth noting that although the table above shows the values for the E-Rand and E-Tourn, the non-elitist versions are interesting. Rand has $D = (N - 1)/N^2$ —the same as Binary SR, and Tourn has $D = (N - 1)/N^2 + 1/N^4$, which for all but small N is also very similar.

5 Evaluating the Indicators for Predictive Value

5.1 Background

The indicators derived above do not consider the representation of a candidate solution, the variation operators that act on it, or the particular fitness function. Therefore it would be surprising if they were able to provide accurate predictions for the *absolute* values observed for any performance metric used to compare algorithms on any given problem. However, if all other aspects of the algorithm are kept identical, it is possible that they may provide indicators of, or insight into, the *relative* performance observed with different replacement strategies. This is the hypothesis which will be examined in the next section.

5.2 Experimental Method

In order to evaluate the predictive value of these indicators a series of experiments were conducted on three problems with the same binary representation but different properties. These were the simple One-Max problem, concatenated copies of Deb’s four-bit deceptive problem (Deb and Goldberg, 1992)—hereafter named Trap, and a version of the type R1 Royal Road problems (Mitchell et al., 1992) with 4-bit blocks—hereafter named Royal. Formulations of these are given in Appendix E. For all three problems, the number of decision variables $L \in \{64, 128\}$ were used.

Table 3: Comparison of Expected Loss of Diversity ($D_{x,y}$) per selection-replacement event with binary tournament selection (x) and different replacement strategies (y), ranked by $D_{x,y}$ decreasing down the table.

Method (y)	$D_{x,y}$, $x = \text{Deterministic Binary Tournaments}$
Worst	$D_{x,y} = 1/N - 1/N^3$
E-FIFO	$1/N - 1/N^2 + 2/N^3 \leq D_{x,y} \leq 1/N - 1/N^3$
E-Tourn.	$D_{x,y} = 1/N - 2/3N^2 + 1/N^3 - 5/3N^4 + 1/N^5$
Binary SR	$D_{x,y} = 1/N - 1/N^2$
E-Rand	$D_{x,y} = 1/N - 1/N^2 - 2/N^3$
C-FIFO	$1/N^2 \leq D_{x,y} \leq 1/2N + 1/2N^2$
C-Rand	$D_{x,y} = 1/N - 1/2N - 1/2N^2$

OneMax is a simple separable unimodal problem. The other two present contrasting types of barriers that the EA must surmount in order to solve each sub-problem, although both have the same size sub-problems and so the same likelihood of an optimal sub-solution to any one of them being present in the initial population. For the trap problems, the barrier is fitness-based: generating the global optimum requires identifying and mixing the sub-solutions before a high fitness solution with some sub-optimal blocks can take over the population. This will aid in evaluating the effects of fitness-based components on the loss of diversity, and how this affects search. By contrast the Royal problem presents an entropic barrier to the EA. The fitness problem for each sub-problem is a “needle in the haystack” type. This requires the maintenance of diversity (entropy) to search each plateau, followed by a dramatic reduction as the population moves to the next higher fitness level. This tests the effect of the non-fitness based factors such as genetic drift on the search process.

With the exception of Binary SR (which makes little sense if recombination is used) all of the guaranteed takeover strategies listed in Table 1 were implemented within a Steady State Genetic Algorithm. Both parents were selected by deterministic binary tournaments (the conservative variant being used where appropriate) and a single offspring generated by two point crossover then subjected to mutation with bit-wise probability $1/L$. Population sizes $N \in \{50, 100, 250, 500, 1000\}$ were used.

For each combination of strategy, population size, problem and problem size, 25 runs of the EA were made, each being continued until either the optimum was found, or 500000 evaluations were made, whichever was sooner. Each solution generated in a run was stored, so that the total number of *distinct* solutions could be compared. The various indicators derived above were compared for their predictive value according to two commonly used quantifiers of EA behaviour, namely the Success Rate (number of runs in which the global optimum was located) and the Average Evaluations to Success (AES). The latter was measured using both the total and distinct number of solutions

Table 4: Relative ordering of methods according to Success Rates on Trap problems, and according to different indicators. The use of $<$ for the empirical results (top line) denotes that the results were significantly different with 95% confidence.

	Ordering
Success Rates	Worst $<$ E-Tourn $<$ E-Rand $<$ E-FIFO $<$ C-FIFO = C-Rand
$1-D$	Worst $<$ E-FIFO $<$ E-Tourn $<$ E-Rand $<$ C-FIFO $<$ C-Rand
$E(T)$	Worst $<$ E-Tourn = E-FIFO $<$ C-FIFO $<$ C-Rand = E-Rand
$E(T) \cdot \eta$	E-Tourn = E-FIFO $<$ Worst $<$ C-FIFO $<$ E-Rand $<$ C-Rand

evaluated. Also considered was the variance in the number of evaluations to success and the diversity, which was measured as the ratio of the number of distinct points evaluated to the total number of points evaluated during each run.

5.3 Predicting Relative Orderings of Success Rates

For the OneMax and Royal problems, all runs located the global optimum. This suggests that entropic barriers of this size may not present a major factor, i.e., that the rate of convergence due to genetic drift was slower than the mean time to discover a fitter solution.

Analysis of variance (ANOVA) of the number of runs which located the global optimum for each combination of population size and replacement strategy for the two different sizes of Trap problem, and post-hoc testing using Tamhane's (1979) T2 test, reveals that the results fall into significantly different groups with 95% confidence. Table 4 shows the relative rankings of the different strategies according to Success Rates on the Trap problems, and also according to various indicators. Inspection of this table shows that none of the indicators correctly predict the relative rankings of the strategies. However, it also shows that in every case the position of E-FIFO is different to expected, and if this is ignored, then decreasing D (Expected Loss of Diversity) correctly predicts the rankings.

5.4 Average Evaluations to Solution

Table 5 shows the relative ordering of the takeover times $E(T)$, and also of the number of evaluations before the global optimum is *first* located on successful runs, i.e., not including runs which do not find the optimum, and not including the time taken for the whole population to converge to the optimum. Within the table use of \leq denotes that observed difference between a pair of results is not statistically significant at the 5% level according to Tamhane's T2 test, which does not assume equal variances.

The results according to the distinct number of points sampled were in almost every case identical to those presented here, apart from occasional differences in the probabilities that two groups of results were significantly different. The sole point of interest occurring from the comparison of distinct with total evaluations came when the values for the unsuccessful runs were included in the Trap row. In this case the Worst and E-Tourn strategies move to the other end of the ordering, because of the high cut-off (half a million evaluations). However the Trap-distinct ordering remains unchanged, indicating premature convergence on runs not locating the global optimum.

It can be seen that the relative ordering of both total and distinct evaluations is predominantly as predicted by $E(T)$. There are some cases where a predicted difference is observed, but is not statistically significant with 95% confidence, but this might be

Table 5: Relative ordering of methods according to $E(T)$ (top line) and Average Evaluations to Success. Use of bold type indicates deviation from predicted relative order. Use of \leq indicates observed difference was not significant at the 95% confidence level.

Model	Ordering
$E(T)$	Worst < E-Tourn = E-FIFO < C-FIFO < C-Rand = E-Rand
1Max-64	Worst \leq E-Tourn \leq E-FIFO < C-FIFO \leq C-Rand \leq E-Rand
1Max-128	Worst \leq E-Tourn < E-FIFO < C-FIFO \leq C-Rand \leq E-Rand
Royal-64	Worst \leq E-Tourn < E-FIFO < C-FIFO \leq C-Rand < E-Rand
Royal-128	Worst \leq E-Tourn < C-FIFO \leq E-FIFO \leq C-Rand < E-Rand
Trap-64	Worst \leq E-Tourn < C-FIFO \leq C-Rand < E-FIFO < E-Rand
Trap-128	Worst \leq E-Tourn < E-Rand \leq C-FIFO \leq C-Rand \leq E-FIFO

resolved by increasing the number of runs performed. The exceptions are the relative positions of E-FIFO on the Trap and 128-bit Royal problems, and the position of E-Rand in the 128 bit Trap problem. The former is worthy of separate discussion, but the latter is an artifact of the low number of successful runs, and the observed differences in mean values are rarely statistically significant.

In order to further investigate the behaviour on the Trap problem, the number of evaluations taken to discover the fittest point found on each run was noted. In this case the ordering on the Trap problems becomes: Worst < E-Tourn < C-FIFO \leq C-Rand \leq E-FIFO \leq E-Rand, which except for E-FIFO is as predicted by $E(T)$.

The predictive value of $E(T)$ for the *absolute* value of AES was also analysed. To this end the Pearson Correlation co-efficients between this value and $E(T)$ were calculated for each combination of population size and problem, averaged over all replacement strategies except E-FIFO (of which more later). Although this analysis revealed a high correlation in each case, the underlying models derived were different.

Overall, $E(T)$ appears to be a good predictor of the relative time taken for the algorithm to find the best solution it is going to within the (large) number of evaluations allowed. However, for some problems, exemplified by our choice of the Trap problem, these points may well be locally, but not globally, optimal. In other words the Worst strategy is fastest to find the best solution it is going to reach, but this may be at the expense of getting trapped on local optima - which ties in with the Success Rate results. Nevertheless, the use of $E(T)$ to predict *absolute* values of AES can be discounted on the basis of these results. Although the correlation co-efficients may be high (> 0.8), the underlying statistical models are different. Thus when a linear regression is built using data from all three problems, it is only able to account for about 25% of the observed variations in the number of evaluations taken to reach the best solution in any given run.

5.5 Behaviour of the E-FIFO strategy

The primary deviations from predicted behaviour according to both metrics above concern the E-FIFO strategy, and it is worth considering the reasons for this. These predictions were based heavily on the use of the linear approximation for p_{old} which implicitly assumes that the mean fitness is continually stochastically increasing. However detailed analysis of the run-time results show a well known phenomenon. For two of the problems the evolving populations frequently exhibit periods during which the mean fitness does not increase, interspersed by the occasional discovery of new higher fitness

solutions and subsequent rises in mean fitness. In practice fitter members will rarely completely take over the population because of the diversifying effects of crossover and mutation. During these “transition” periods the behaviour of E-FIFO is characterised by the linear approximation and the behaviour is like that of E-Tourn. The differences arise from the periods when there are no new high fitness solutions to take over the population. If the mean fitness is static when averaged over a period of more than N time steps, then clearly the distribution of fitter members of the population according to age has become more uniform. In the extreme case the probability that the oldest member of the population is one of i members of the fitter class is simply i/N , in which case the behaviour of E-FIFO will become that of E-Rand.

The OneMax problem has sub-problems of length 1, and hence no barriers to improvement. Examining Table 5 in closer detail, it is apparent that for this problem the ordering of methods is as predicted by the linear approximation for E-FIFO: the AES values are smaller than for E-Rand, and there is often no significant difference between its behaviour and that of E-Tourn. Empirically the mean fitness was seen to increase steadily until the optimum was found.

For the other two problems the sub-problems have a larger size, and the “epochs” or periods when the mean fitness is not increasing, correspond to the cases where the population has converged onto the optimal sub-string for some sub-problems, but not others. For the Royal Road problem, all other sub-strings have equal fitness, so assuming random drift, an “average” population member which has not solved a given sub-problem will require mutation to simultaneously flip two bits from 0 to 1 in order to “solve” that sub-problem. By contrast, for the Trap problem the all-zeroes substring is also an attractor for the system, so on average discovery of the optimal substring will require more bits to be flipped by mutation, and take correspondingly longer. Because of this, it might be hypothesised that the periods of waiting to discover new higher fitness solutions, during which E-FIFO behaves more like E-Rand, will be longer for the Trap problem than for the Royal problem, and so differences between the observed behaviours of E-FIFO and E-Tourn will be more marked in this case. This is exactly what can be seen in the empirical results in Table 5. The rank ordering is the same for Royal as for the predictions, there is now a statistically significant difference between the results for E-FIFO and E-Tourn. For the trap problem, E-FIFO displays behaviour which is most like that of E-Rand, and is significantly slower than the other methods.

5.6 Variance in Evaluations to Success

The observed variance in the total number of evaluations to success for the three strategies Worst, C-Rand and C-FIFO was subjected to ANOVA. This confirmed that the problem type and size, population size and the strategy all had significant effects with 95% confidence. A post-hoc analysis using Tamhane’s test of the effect of the strategy revealed that the observed variances were, as predicted, in the order Worst < C-FIFO < C-Rand with greater than 99% confidence in the significance of the differences. Also as predicted, the ratio of the variance values was approximately 5:8 for the two conservative strategies. The value for the Worst strategy could not reasonably be compared there were far fewer successful runs on the Trap problem.

However, as was the case for Average Evaluations to Success, the predominant factor in determining the variance was the choice of problem. Thus although the values of $V(T)$ are a good predictor for the relative rankings of different strategies, they should not be used to predict the absolute values of the variances in the number of evaluations before the global optimum is located.

5.7 Diversity

Table 6 shows the relative ranking of the strategies according to the indicators η and D , and also according to observations of the ratio of distinct to total points evaluated during each successful run; in total, and also grouped by problem. Use of \leq denotes that observed difference between this pair is not statistically significant at the 5% level according to Tamhane’s test. Only successful runs were considered to avoid introducing spurious effects as a result of upper level of 500,000 evaluations. Note that because of this, the values for Worst and E-Tourn strategies on the Trap are based on far fewer runs, especially with the lower population sizes where diversity is more of an issue. When all runs (successful or otherwise) are considered, the order is the same as that for Royal, and all differences are significant except for those between E-FIFO and E-Rand.

It is clear that neither of the indicators is predictive of the relative order. For example, according to η , C-Rand, E-FIFO, E-Rand and E-Tourn will all select the same fraction of parents that are not the current best. A full ANOVA taking all experimental factors into account shows a significant difference between these different strategies with 95% confidence. Equally, C-FIFO (according to both indicators) and C-Rand (according to D) should exhibit more diversity than their counterparts using standard tournament selection, but in practise this is never the case.

These deviations arise because η and D are predictors of diversity in the choice of parents, whereas the observed variable relates to new offspring. These are created by the processes of crossover and mutation, both of which may increase diversity. Furthermore, the “best” solution is changing in practise as the problems are solved, whereas the models assume it is initially present.

Table 6: Relative ranking of methods according to different indicators of the increasing diversity of search. *Opt* denotes runs which found the global optimum. Use of \leq denotes that observed difference between this pair is not statistically significant at the 5% level according to Tamhane’s test.

Indicator	Relative Ordering of Diversity
η	C-Rand = E-FIFO = E-Tourn = E-Rand < C-FIFO < Worst
D	Worst < E-FIFO < E-Tourn < E-Rand < C-FIFO < C-Rand
Distinct:Total	
Royal	Worst \leq E-Tourn < C-Rand < C-FIFO < E-Rand \leq E-FIFO
OneMax	C-Rand < C-FIFO < Worst \leq E-Tourn < E-FIFO \leq E-Rand
Trap- <i>Opt</i>	C-Rand \leq C-FIFO < E-Rand \leq E-FIFO \leq E-Tourn \leq Worst
All- <i>Opt</i>	C-Rand \leq Worst \leq C-FIFO \leq E-Tourn < E-Rand \leq E-FIFO

6 Conclusions

For a variety of reasons, the study of the effects of selection within steady state EAs has been relatively unexplored compared to their generational counterparts, despite their widespread use. This is unfortunate since most practical, and in particular theoretical work on GEAs assumes an age-based “replace all parents with offspring” policy, and relatively few theoretical results are available for guidance on this topic. This paper has gone some way to addressing this problem by considering a range of indicators that relate to selection/replacement alone.

While theoretical indicators for comparing algorithms are of interest in themselves, the assumptions made to make such complex systems tractable to analytical results often open them to doubts about their practical relevance, and there is currently a wide discrepancy between the analytical indicators available and the type of performance metrics typically of interest to practitioners. This paper has not only been concerned with developing analytical indicators, but also with examining the value of these indicators for predicting the relative and absolute behaviour of SSEAs incorporating those strategies on combinatorial optimisation problems. To this end, a series of experiments were run using different population sizes and problem types and sizes. Depending on the perspective of the user, a variety of performance metrics may be used. Four were examined, namely the Success Rate (number of runs in which the global optimum was located), the Average Evaluations to Success (AES), the variance in the number of evaluations to success, and the ratio of distinct to total solutions evaluated per run, an indicator of the diversity of the search.

It should be reiterated at this stage that some of these indicators were not intended to be used on their own—for example Takeover Time analysis, which does not consider the problem encoding or variation operators, was originally intended to be used in tandem with a mixing analysis, which does. Thus as expected, the values predicted by the indicators are not correlated with observed absolute values of Success Rate, the mean (AES) or Variance of the time taken to solution when averaged over a range of problems. It should also be noted that any results can only be taken as tentative, since the test suite is so small. However, as seen, even this limited test suite presents strong evidence that some of these indicators do not have any predictive value for the behaviour of “real” EAs.

The Success Rate results showed that although none of the indicators based on the Markov models acted as a good predictor, decreasing Expected Loss of Diversity ($1 - D$) correctly predicted the relative ordering of the strategies, excluding E-FIFO.

Considering the AES indicator, which only considers successful runs, $E(T)$ was a good predictor for the relative rankings of the strategies, with deviations again mostly arising from the approximations used to satisfy the Markov condition for the Elitist-FIFO strategy. Furthermore, analysis showed that when the number of evaluations taken to find the best (i.e., not necessarily globally optimal) solution seen per run is considered, $E(T)$ correctly predicts the relative order on all three different problem types (again apart from the E-FIFO strategy).

A similar analysis reveals that the variance in takeover time $V(T)$ is also a good predictor for the relative variance in run-times. Statistical analysis yielded high correlation co-efficients between AES- $E(T)$ and Variance- $V(T)$, even on Trap problems, although the underlying models, in particular the regression co-efficients, differed between problems. The strong correlations observed between these strictly fitness-based indicators, and the empirical performance metrics, even in the presence of multiple fitness plateaus, suggests that differential rates of genetic drift (which are not considered by $E(T)$ or $V(T)$) may be less important than fitness components when comparing between strategies.

Concerning the diversity of the search, two indicators were available: η , which groups all non-fittest solutions into a single equivalence class, and D , which assumes that all members of the population have a distinct fitness, in other words that there are no duplicates. It was found that the neither had any value as a predictor of the relative rankings of the different strategies according to the ratio of distinct to total number of solutions evaluated. This is perhaps more of an artifact of the empirical indicator, since

it was noted above that the Expected Loss of Diversity per select-replace event correctly predicted the relative success rates of different strategies. However the η indicator appears to have little predictive value.

As was stated at the outset, it was never the intention of this paper to compare a number of replacement strategies and declare one the “winner”. There are now good reasons to believe such a quest to be fruitless, and this partly motivated the deliberately limited set of test problems chosen. Moreover, as was shown, the different strategies show different relative rankings according to the choice of performance metric. However, it is possible to draw some fairly broad conclusions.

First, there appears to be little merit in a Replace-Random scheme. Regardless of the choice of analytical indicator or performance metric, there are other algorithms which perform as well. In particular the equally computationally trivial FIFO mechanism exhibits the same success rate, faster location of optima and greater diversity of search. When combined with the conservative method of coupling selection and replacement, the C-Rand strategy performs reasonably well, but again not as well as C-FIFO according to any indicator or performance metric tested here.

Second, the two fitness-based methods (Worst and E-Tourn) perform poorly in terms of Success Rates with low population sizes in the presence of fitness barriers to the identification of the optimum. However, Worst does converge the fastest of all the methods. In particular it does so much more reliably, in the sense of having a smaller variance in takeover time and time to solution than all the other methods. The indicators suggest that a replace Worst strategy will find the best solution it is going to on average three times as fast as the C-FIFO strategy, and with only a fifth of the variance in run-times. This is borne out by the results, but these also show that this comes at the expense of success in locating the global optimum. For repetitive problems, as long as the solutions found are of reasonable quality, this reliability will often be of benefit.

It is hoped that this paper will go some way towards providing a more informed basis for the choice of strategy according to intended use. Future work intends to expand the analysis of the variance in takeover time, and also to extend the analysis to explicitly take into account the action of the variation operators.

Appendix

A Modelling FIFO replacement

A.1 A Linear Approximation for p_{old}

In order to model p_{old} in an evolving population it is necessary to consider the probability $p_{class}(x)$, that an individual of age x ($0 \leq x \leq N - 1$) belongs to the fitter class, where x denotes the number of individuals inserted into the population after that member. If $p_{class}(x)$ is independent of x the model is identical to that for Replace-Random seen in Eq. 9, but simulation reveals that this is not the case. This corresponds to the intuition that if i (the current value of $X(t)$) is increasing with time, which means that selection and replacement are increasing the mean fitness of the population, then younger members of the population are more likely to belong to the fitter class than older ones.

Under the assumption that $p_{class}(x)$ decreases linearly with age x across the population, it can be modelled as $p_{class}(x) = a + bx$, where a and b are constants. The value of a is $p_{class}(0)$, and so by definition $a = p_s(t - 1) \approx p_s(t)$, where p_s is a function of i , the current value of $X(t)$. Since there is only one individual of each age, the expected number of members of the fitter class for each age is equal to the probability

$p_{class}(x)$. Summing over all ages between 0 and $N - 1$, the expected number of members in the population is the sum of the expected number for each age, i.e., the sum of p_{class} . Recognising that this sum must equal i gives:

$$i = \sum_{x=0}^{N-1} p_{class}(x) = Na + b \sum_{x=0}^{N-1} x \approx Np_{s(i)} + bN(N-1)/2$$

therefore $b = \frac{2(i - Np_{s(i)})}{N(N-1)}$.

Substituting these values for a and b in the linear model gives the probability that the oldest member of the population belongs to the fitter class as:

$$p_{old} = p_{class}(N-1) = 2i/N - p_{s(i)}. \tag{A-1}$$

A.2 Selection and Replacement Probabilities with Conservative-FIFO Tournaments

Using the same notation as above, if a C-FIFO strategy is used, the probability of selecting a member of the fitter class to be parent is:

$$p_s = p_{old} + p_{random} \cdot (1 - p_{old}) = p_{old} + \frac{i}{N}(1 - p_{old}).$$

Substituting this expression into Eq. A-1 yields:

$$p_{old} = i/(2N - i) \quad \text{and} \quad p_{s(i)} = \frac{i(3N - 2i)}{N(2N - i)}. \tag{A-2}$$

B Calculation of Takeover Indicators for Strategies Which Guarantee Takeover

B.1 k -ary SR Tournaments

The generalised form of this result is given in (Rudolph, 2000). Here the binary version is shown for completeness, and for comparison with the other strategies where binary tournaments are used for parent selection. From Eq. 5 setting $k = 2$ gives $p_{i,i+1} = 2i(1 - i/N)/N$, which when inserted into Eq. 1 gives:

$$E(T) = \frac{N}{2} \sum_{i=1}^{N-1} \frac{N}{i(N-i)} = N \sum_{i=1}^{N-1} \frac{1}{i} = N \cdot H_{N-1}$$

giving the bounds: $N \log(N-1) < E(T) < N(\log(N-1) + 1)$.

Similarly for the variance:

$$\begin{aligned} \sum_{i=1}^{N-1} \frac{1}{p_{i,i+1}^2} &= \frac{N^4}{4} \sum_{i=1}^{N-1} \frac{1}{i^2(N-i)^2} = N \cdot \sum_{i=1}^{N-1} \frac{1}{i} + \frac{N^2}{2} \cdot \sum_{i=1}^{N-1} \frac{1}{i^2} \\ &= N \cdot H_{N-1} + \frac{N^2}{2} \zeta_{N-1}^2 \\ \text{therefore } V(T) &= \frac{N^2}{2} \zeta_{N-1}^2 \rightarrow \frac{N^2 \pi^2}{12}, \end{aligned}$$

where ζ_{N-1}^2 denotes Riemann's Zeta function with exponent 2 truncated to its first $N-1$ terms, which since all terms are positive converges asymptotically, and the limit is $\pi^2/6$.

Since from Eq. 5 $p_{i,i+1} = 1 - ((N-i)/N)^k - (i/N)^k$, which is symmetric with respect to N and $N-i$, $\eta = 0.5$.

B.2 Replace–Worst

Inserting the form of Eq. 6 for deterministic binary tournaments into Eq. 1 gives:

$$\begin{aligned} E(T) &= N^2 \cdot \sum_{i=1}^{N-1} \frac{1}{i(2N-i)} = \frac{N}{2} \cdot \sum_{i=1}^{N-1} \left(\frac{1}{i} + \frac{1}{2N-i} \right) \\ &= \frac{N}{2} (H_{N-1} + H_{2N-1} - H_N) = \frac{N}{2} \left(H_{2N-1} - \frac{1}{N} \right), \end{aligned}$$

which yields the bounds: $\frac{N}{2} (\log(2N-1) - \frac{1}{N}) < E(T) < \frac{N}{2} (\log(2N-1) + 1)$.

For the variance:

$$\begin{aligned} \sum_{i=1}^{N-1} \frac{1}{p_{i,i+1}^2} &= \sum_{i=1}^{N-1} \frac{N^4}{i^2(2N-i)^2} = \frac{N}{4} \sum_{i=1}^{N-1} \left(\frac{i+N}{i^2} + \frac{3N-i}{(2N-i)^2} \right) \\ &= \frac{N}{4} \sum_{i=1}^{N-1} \left(\frac{1}{i} + \frac{1}{2N-i} + \frac{N}{(2N-i)^2} + \frac{N}{i^2} \right) \\ &= \frac{N}{4} (H_{N-1} - H_N + H_{2N-1} + N \cdot \zeta_{N-1}^2 - N \cdot \zeta_N^2 + N \cdot \zeta_{2N-1}^2) \\ &= \frac{N^2}{4} \cdot \zeta_{2N-1}^2 + \frac{N}{4} \cdot H_{2N-1} - \frac{1}{2} \end{aligned}$$

$$\text{therefore } V(T) = \frac{N^2}{4} \cdot \zeta_{2N-1}^2 - \frac{N}{4} \cdot H_{2N-1}.$$

An expression for η is given in (Rudolph, 2000) as $\eta \approx 1 - \frac{2 \log(2)}{\log(2N)} \rightarrow 1$.

B.3 Conservative–Random

Inserting binary tournament selection probability $p_{s(i)}$ into Eq. 8 yields $p_{i,i+1} = \frac{N-i}{N} \cdot \frac{i}{N}$. Inspection shows this probability to be half that for binary SR tournaments with $k=2$, which leads immediately to the results:

$$2N \log(N-1) < E(T) < 2N(\log(N-1) + 1) \quad (\text{A-3})$$

and

$$V(T) = 2N^2 \cdot \zeta_{N-1}^2. \quad (\text{A-4})$$

Since the transition probability is symmetric with respect to i and $N-i$, $\eta = 0.5$.

B.4 Conservative–FIFO

The linear approximation $p_{i,i+1} = \frac{2i(N-i)}{N(2N-i)}$ was derived in Appendix A.2. In (Smith and Vavak, 1999b) this model was tested against simulations for a single fixed population size and was shown to be reasonably accurate, but that the assumption introduced errors into the prediction in the initial stages. In order to make corrections for these errors the takeover time is split into two phases. The second of these deals with the transitions $i \in \{2, \dots, N-1\}$ and uses the equation above. The first deals explicitly with the time taken to first duplicate the fitter solution, which will be denoted E_{12} . Ignoring the minor adjustments required by the fact that the number of copies stays the same every N^{th} cycle, the transition to having two copies occurs if the randomly

selected member of the tournament belongs to the fitter class, which happens with probability $1/N$, hence $E_{12} = N$. Putting these together:

$$\begin{aligned} E(T) &= N + \sum_{i=2}^{N-1} \frac{N(2N-i)}{2i(N-i)} = N + \sum_{i=1}^{N-1} \frac{N(2N-i)}{2i(N-i)} - \frac{N(2N-1)}{2(N-1)} \\ &= \sum_{i=1}^{N-1} \frac{N}{2} \cdot \left(\frac{2}{i} + \frac{1}{N-i} \right) - \frac{N}{2(N-1)} = \frac{N}{2} \left(3H_{N-1} - \frac{1}{N-1} \right) \end{aligned}$$

giving the bounds:

$$\frac{N}{2} \left(3 \log(N-1) - \frac{1}{(N-1)} \right) < E(T) < \frac{N}{2} \left(3 \log(N-1) + 3 - \frac{1}{(N-1)} \right). \quad (\text{A-5})$$

Combining the transition probabilities for the two different phases gives:

$$\begin{aligned} \sum_{i=1}^{N-1} \frac{1}{p_{i,i+1}^2} &= N^2 + \frac{N^2}{4} \cdot \sum_{i=2}^{N-1} \frac{(2N-i)^2}{i^2(N-i)^2} \\ &= N^2 \cdot \left(1 - \left(\frac{2N-1}{2N-2} \right)^2 \right) + \frac{N^2}{4} \sum_{i=1}^{N-1} \left(\frac{1}{N-i} + \frac{2}{i} \right)^2 \\ &= N^2 \cdot \left(1 - \left(\frac{2N-1}{2N-2} \right)^2 \right) + \frac{5N^2}{4} \zeta_{N-1}^2 + 2N \cdot H_{N-1} \\ &= \frac{5N^2}{4} \zeta_{N-1}^2 + 2N \cdot H_{N-1} - \frac{N^2}{4(N-1)^2} \cdot (4N-3) \end{aligned}$$

therefore $V(T) = \frac{5N^2}{4} \zeta_{N-1}^2 + \frac{1}{2} \cdot N \cdot H_{N-1} - \frac{N(4N^2 - 5N + 2)}{4(N-1)^2}$.

For the calculation of η begin by considering the number of copies of the best individual prior to takeover, again using the two-phase approximation.

$$\begin{aligned} E(B_T) &= \sum_{i=1}^{N-1} \frac{i}{p_{i,i+1}} = N + \sum_{i=2}^{N-1} i \cdot \frac{N(2N-i)}{2i(N-i)} \\ &= \frac{N}{2} \cdot \left(2 - \frac{2N-1}{2(N-1)} + \sum_{i=1}^{N-1} \frac{(2N-i)}{(N-i)} \right) \\ &= \frac{N}{2} \cdot \left(N + NH_{N-1} - \frac{1}{2(N-1)} \right). \end{aligned} \quad (\text{A-6})$$

Using $\eta = 1 - E(B_T)/(N \cdot E(T))$ and inserting the quantities from Eq. A-5 and Eq. A-6 gives:

$$\eta = 1 - \frac{H_{N-1} + 1 - 1/2N(N-1)}{3H_{N-1} - 1/(N-1)} = \frac{2}{3} - \frac{6N^2 - 4N - 3}{6N^2(3(N-1)H_{N-1} - 1)},$$

which for large N gives $\eta \rightarrow 2/3$ from below.

B.5 Elitist Replace-Random

Rudolph (2001) contains the following results:

$$Q_{i,j} = \frac{N^2}{N+1} \cdot \left(\frac{1}{j} + \frac{1}{N-j} \right) \quad \therefore E(T) = \frac{2N^2}{N+1} H_{N-1}.$$

Inserting these results into the second expression for η in Eq. 3, and noting that $j \cdot \left(\frac{1}{j} + \frac{1}{N-j} \right) = \frac{N}{N-j}$ yields:

$$\eta = 1 - \frac{\frac{N^2}{N+1} \sum_{j=1}^{N-1} j \cdot \left(\frac{1}{j} + \frac{1}{N-j} \right)}{\frac{2N^3}{N+1} \cdot H_{N-1}} = 1 - \frac{\frac{N^2}{N+1} \sum_{j=1}^{N-1} \frac{N}{N-j}}{\frac{2N^3}{N+1} \cdot H_{N-1}} = 1/2.$$

B.6 Elitist Kill-Tournaments

Rudolph (2001) contains the following results:

$$Q_{i,j} \leq \left(\frac{N}{2} + \frac{1}{4} \right) \cdot \left(\frac{1}{j} + \frac{1}{N-j} \right) \quad \therefore E(T) \leq \left(N + \frac{1}{2} \right) H_{N-1}.$$

Inserting these results into the first form for η in Eq. 3 yields:

$$\begin{aligned} \eta &\geq 1 - \frac{\sum_{j=1}^{N-1} j \cdot \left(\frac{N}{2} + \frac{1}{4} \right) \cdot \left(\frac{1}{j} + \frac{1}{N-j} \right)}{N \cdot \sum_{j=1}^{N-1} \left(\frac{N}{2} + \frac{1}{4} \right) \cdot \left(\frac{1}{j} + \frac{1}{N-j} \right)} \\ &= 1 - \frac{N \cdot \sum_{j=1}^{N-1} \frac{1}{N-j}}{N \cdot \sum_{j=1}^{N-1} \left(\frac{1}{j} + \frac{1}{N-j} \right)} = 1/2. \end{aligned}$$

C Bounding the Takeover Probability for FIFO Replacement

As noted in (Smith and Vavak, 1999b), the linear approximation for p_{old} in non-elitist FIFO tends to be inaccurate for low i . In this section upper and lower bounds will be derived for the probability of absorption to the state $X(t) = N$, without using this approximation, and thus without recourse to standard techniques for Markov Chains.

It is possible to derive an upper bound for the probability that absorption to state N happens by considering that this always happens unless either the original member of fitter class is deleted before it is selected to be copied, or that a single copy is made at some time $c : 1 \leq c \leq N$, but that no further copies are made, and that both of these are deleted after their lifespan of N events. In other words an upper bound is given by assuming that absorption to $X(t) = N$ occurs unless the progression of $X(t)$ is $1 \rightarrow 0$ or $1 \rightarrow 2 \rightarrow 1 \rightarrow 0$. Denoting the probabilities of these two events as p_{10} and p_{1210} respectively, $p_{takeover} < 1 - p_{10} - p_{1210}$. The takeover probability will be strictly less than this upper bound, since other sequences of events terminating in state 0 are possible, even from state $N-1$. A tighter bound could be derived by considering progressions such as $1 \rightarrow 2 \rightarrow 3 \rightarrow 2 \rightarrow 1 \rightarrow 0$, or $1 \rightarrow 2 \rightarrow 1 \rightarrow 2 \rightarrow 1 \rightarrow 0$ etc., but these corrections become increasingly small since at any time step the probability of selecting one of i member of the fitter class to be copied rises quadratically with i .

To derive the necessary probabilities is a relatively straightforward exercise in combinatorics, making use of the identity $\lim_{N \rightarrow \infty} (1 - a/N)^N = e^{-a}$. Thus straight away

$p_{10} = \left(\frac{N-1}{N}\right)^{2N} \rightarrow e^{-2}$. To calculate p_{1210} it is necessary to account for the sequence of events given creation of the second copy of the fitter member at some time $1 \leq c \leq N$

$$\begin{aligned} P_{1210} &= \sum_{c=1}^N \left(\frac{N-1}{N}\right)^{2(c-1)} \cdot \frac{(2N-1)}{N^2} \cdot \left(\frac{N-2}{N}\right)^{2(N-c)} \cdot \left(\frac{N-1}{N}\right)^{2c} \\ &= \frac{(2N-1)}{(N-1)^2} \cdot \left(\frac{N-2}{N}\right)^{2N} \cdot \sum_{c=1}^N \left(\frac{(N-1)^2}{N(N-2)}\right)^{2c} \\ &= \frac{(2N-1)}{(N-1)^2} \cdot \left(\frac{N-2}{N}\right)^{2N} \cdot \sum_{c=1}^N (1 + O(N^{-2}))^{2c} \rightarrow 2e^{-4}. \end{aligned}$$

Thus for the upper bound, ignoring terms in N^{-4} and smaller:

$$p_{takeover} < 1 - \left(\frac{N-1}{N}\right)^{2N} - N \cdot \frac{(2N-1)}{(N-1)^2} \cdot \left(\frac{N-2}{N}\right)^{2N}$$

noting that the expression $(1 - a/N)^N$ converges to e^{-a} from above:

$$\lim_{N \rightarrow \infty} p_{takeover} < 1 - e^{-2} - 2e^{-4}.$$

Calculation of a lower bound for the takeover probability is more complex, since the probability that $X(t+N) = 0 \mid X(t) = N-1$ is positive non-zero. However, as noted in Section 3.1.3, the system as a whole can be described as a Markov chain, and by similar arguments to those for the simpler systems, it is clear that only the states $X(t) = 0$ and $X(t) = N$ are absorbing. Thus it is possible to proceed by noting that the lower bound for probability of reaching the (single) state characterised by $X(t) = N$ will be 1—the upper bound for the probability of convergence to the (single) state characterised by $X(t) = 0$.

If it is considered that loss of the fitter solution happens unless there are at least two copies of the fitter member in the population after N time-steps (i.e. $X(N) > 1$), then loss happens with probability $p_{10} + p_{121}$, where p_{121} denotes that the variable $X(c) = 2$ for some $1 \leq c < N$ and $X(N) = 1$.

By similar arguments to those above:

$$\begin{aligned} p_{121} &= \sum_{c=1}^N \left(\frac{N-1}{N}\right)^{2(c-1)} \cdot \frac{(2N-1)}{N^2} \cdot \left(\frac{N-2}{N}\right)^{2(N-c)} \\ &= \frac{2N-1}{2N-3} \cdot \left(\frac{N-2}{N-1}\right)^2 \cdot \left(\left(\frac{N-1}{N}\right)^{2N} - \left(\frac{N-2}{N}\right)^{2N} \right) \rightarrow e^{-2} - e^{-4} \end{aligned}$$

which gives;

$$p_{loss} \leq 1 - \left(1 - \frac{2}{N}\right)^{2N} - \frac{2N-1}{2N-3} \cdot \left(\frac{N-2}{N-1}\right)^2 \cdot \left(\left(1 - \frac{2}{N}\right)^{2N} - \left(1 - \frac{4}{2N}\right)^{2N} \right)$$

$$\therefore \lim_{N \rightarrow \infty} p_{takeover} > 1 - 2 \cdot e^{-2} + e^{-4}.$$

D Expected Loss of Diversity

For steady state algorithms the expected loss of diversity at each time step is defined as $D_{x,y}(t, N) = \sum_{k=1}^N (1 - p_{s,x}(k)) \cdot p_{d,y}(k)/N$, where the population is assumed to have different fitness values, and $k \in \{1, \dots, N\}$ denotes the rank of each member, N being the fittest. The exact form of $p_{s,x}(k)$ will depend on the parent selection method x . Here binary deterministic tournaments are assumed for consistency, so using the abbreviated notation $p_s(k) = (2k - 1)/N^2$ and $1 - p_s(k) = (N^2 - 2k + 1)/N^2$. For clarity the subscript denoting the parent selection strategy is omitted, and the calculation of D for the various strategies is as follows:

- Delete-Worst: By definition $p_{d,\text{worst}}(k) = 0 \ \forall \ k > 1$ hence

$$D_{\text{worst}}(N) = \frac{(1 - p_s(1))}{N} = \frac{1}{N} - \frac{1}{N^3}.$$

- Binary Select-Replace Tournaments: For any k , it gets replaced without copying if it is one of a randomly chosen pair and the other is fitter hence

$$D_{\text{k-SR}}(N) = \frac{1}{N} \sum_{k=1}^N 2 \cdot \frac{1}{N} \cdot \frac{N - k}{N} = \frac{1}{N} - \frac{1}{N^2}.$$

- Replace-Random: By definition each member has an equal chance of being replaced, i.e. $p_{d,\text{Rand}}(k) = 1/N \ \forall \ k$ hence

$$D_{\text{Rand}}(N) = \frac{1}{N} \sum_{k=1}^N \frac{1 - p_s(k)}{N} = \frac{1}{N} - \frac{1}{N^2}.$$

- Elitist Replace-Random: In this case $p_{d,\text{E-Rand}}(N) = 0$, i.e., the summation limit changes from N to $N - 1$ thus:

$$D_{\text{E-Rand}}(N) = \frac{1}{N} \sum_{k=1}^{N-1} \frac{1 - p_s(k)}{N} = \frac{1}{N} - \frac{1}{N^2} + \frac{2}{N^3}.$$

- Conservative-Random: The item to be deleted is chosen at random, so $p_{d,\text{C-Rand}}(k) = 1/N \ \forall \ k$. The member to be deleted is not selected to be a parent if the other randomly selected item is fitter, hence by comparison with the form for binary SR above:

$$D_{\text{C-Rand}}(N) = 1/2N - 1/2N^2.$$

- Kill Tournaments: For a tournament of size t , the deletion probability is $p_{d,\text{Tourn}}(k) = ((N - k + 2)^t - (N - k)^t)/N^t$ which for $t = 2$ reduces to $p_{d,\text{Tourn}}(k) = (2N - 2K + 1)/N^2$. Hence

$$D_{\text{Tourn}}(N) = \frac{1}{N} \sum_{k=1}^N \frac{(N^2 + 1) - 2k}{N^2} \cdot \frac{(2N + 1) - 2k}{N^2} = \frac{1}{N} - \frac{1}{N^2} + \frac{1}{N^4}.$$

- **Elitist Kill Tournaments:** Again the elitist variant simply requires a change in the summation limit, which for $t = 2$ yields:

$$D_{\text{E-Tourn}}(N) = \frac{1}{N} \sum_{k=1}^{N-1} \frac{(N^2 + 1) - 2k}{N^2} \cdot \frac{(2N + 1) - 2k}{N^2} = \frac{1}{N} - \frac{2}{3N^2} - \frac{1}{N^3} + \frac{5}{3N^4} - \frac{1}{N^5}.$$

- **FIFO:** In this case it is not possible to derive an exact result formula since the dependence of rank upon age is ill-defined. However it is possible to derive bounds by considering the two limiting cases. In the first case, when there is no relation between the age of an individual and its rank (i.e., the population mean fitness is not improving) then FIFO is equivalent to Replace-Random. If on the contrary it is assumed that the fitness of the offspring produced is constantly increasing, so that the dependence of rank upon age is linear, then the oldest member will always be the least fit, so the situation is the same as for Replace-Worst. Taking these two extremes together yields

$$\frac{1}{N} - \frac{1}{N^2} \leq D_{\text{FIFO}}(N) \leq \frac{1}{N} - \frac{1}{N^3}.$$

- **Elitist-FIFO:** This is as above, with the lower bound corrected as per the Elitist Replace-Random:

$$\frac{1}{N} - \frac{1}{N^2} + \frac{2}{N^3} \leq D_{\text{E-FIFO}}(N) \leq \frac{1}{N} - \frac{1}{N^3}.$$

- **Conservative-FIFO:** By similar arguments, the lower bound on the loss of diversity is the same as for Conservative-Random. For the upper bound, when the population is steadily increasing in fitness, the oldest member will always be the least fit, i.e. $p_{\text{DC-FIFO}}(k) = 0 \quad \forall k > 1$. That member will only be copied if the randomly chosen member of the tournament is the same one, which has probability $1/N$ yielding:

$$\frac{1}{2N} - \frac{1}{2N^2} \leq D_{\text{C-FIFO}}(N) \leq \frac{1}{N} - \frac{1}{N^2}.$$

E Test Problems Used

The following problems were used in the experiments for candidate solutions of the form $x \in \{0, 1\}^L$ with $L \in \{64, 128\}$. All are normalised to give values between 0 and 100.

OneMax: This is a simple unimodal problem.

$$f(x) = \frac{100}{L} \cdot \sum_{i=1}^L x(i).$$

Royal: This is a type R1 Royal Road (Forrest and Mitchell, 1992) where the fitness is given by the number of blocks “aligned” to the target string in a problem with L blocks, each of length K :

$$f(x) = \frac{400}{L} \cdot \sum_{i=0}^{(L-1)/K} \prod_{k=1}^K x(4i + k).$$

For the experiments reported here the epistatic blocks had size $K = 4$.

Trap: This problem is composed of $L/4$ four-bit deceptive problems as given by Deb in (Bäck et al., 1997) concatenated together:

$$f(x) = \frac{400}{L} \cdot \sum_{i=0}^{(L-1)/4} \max \left(\prod_{k=1}^4 x(4i+k), 0.6 - 0.2 \cdot \sum_{k=1}^4 x(4i+k) \right).$$

References

- Asoh, H. and Mühlenbein, H. (1994). On the mean convergence time of evolutionary algorithms without selection and mutation. In Davidor, Y., Schwefel, H.-P., and Männer, R., editors, *Proceedings of the 3rd Conference on Parallel Problem Solving from Nature*, number 866 in Lecture Notes in Computer Science, pages 88–97. Springer, Berlin.
- Bäck, T. (1995). Generalised convergence models for tournament and (μ, λ) selection. In Eshelman, L., editor, *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 2–8. Morgan Kaufmann, San Francisco.
- Bäck, T., Fogel, D., and Michalewicz, Z., editors (1997). *Handbook of Evolutionary Computation*. Institute of Physics Publishing, Bristol, and Oxford University Press, New York.
- Banzhaf, W. and Reeves, C., editors (1999). *Foundations of Genetic Algorithms 5*. Morgan Kaufmann, San Francisco.
- Blickle, T. and Thiele, L. (1997). A comparison of selection schemes used in genetic algorithms. *Evolutionary Computation*, 4(4):261–294.
- Chakraborty, U. (1995). An analysis of selection in generational and steady state genetic algorithms. In *Proceedings of the National Conference on Molecular Electronics*. NERIST (A.P.) India.
- Chakraborty, U., Deb, K., and Chakraborty, M. (1997). Analysis of selection algorithms: A Markov Chain approach. *Evolutionary Computation*, 4(2):133–167.
- De Jong, K. and Sarma, J. (1992). Generation gaps revisited. In (Whitley, 1993), pages 19–28.
- Deb, K. and Goldberg, D. (1992). Analyzing deception in trap functions. In (Whitley, 1993), pages 93–108.
- Forrest, S. and Mitchell, M. (1992). Relative building block fitness and the building block hypothesis. In (Whitley, 1993), pages 109–126.
- Goldberg, D. (2002). *The Design of Innovation: Lessons from and for Competent Genetic Algorithms*. Kluwer Academic Publishers, Boston.
- Goldberg, D. and Deb, K. (1991). A comparative analysis of selection schemes used in genetic algorithms. In Rawlins, G., editor, *Foundations of Genetic Algorithms*, pages 69–93. Morgan Kaufmann, San Francisco.
- Goldberg, D. and Segrest, P. (1987). Finite Markov chain analysis of genetic algorithms. In Grefenstette, J., editor, *Proceedings of the 2nd International Conference on Genetic Algorithms and Their Applications*, pages 1–8. Lawrence Erlbaum, Hillsdale, New Jersey.

- Mitchell, M., Forrest, S., and Holland, J. (1992). The royal road for genetic algorithms: Fitness landscapes and ga performance. In Varela, F. and Bourgine, P., editors, *Toward a Practice of Autonomous Systems: Proceedings of the 1st European Conference on Artificial Life*, pages 245–254. MIT Press, Cambridge, MA.
- Motoki, T. (2002). Calculating the expected loss of diversity of selection schemes. *Evolutionary Computation*, 10(4):397–422.
- Nix, A. and Vose, M. (1992). Modelling genetic algorithms with Markov chains. *Annals of Mathematics and Artificial Intelligence*, pages 79–88.
- Prügel-Bennet, A. and Shapiro, J. (1994). An analysis of genetic algorithms using statistical mechanics. *Physical Review Letters*, 72(9):1305–1309.
- Radcliffe, N. (1991). Forma analysis and random respectful recombination. In Belew, R. and Booker, L., editors, *Proceedings of the 4th International Conference on Genetic Algorithms*, pages 222–229. Morgan Kaufmann, San Francisco.
- Rogers, A. and Prügel-Bennett, A. (1999a). Genetic drift in genetic algorithm selection schemes. *IEEE Transactions on Evolutionary Computation*, 3(4):298–303.
- Rogers, A. and Prügel-Bennett, A. (1999b). Modelling the dynamics of a steady-state genetic algorithm. In (Banzhaf and Reeves, 1999), pages 57–68.
- Rudolph, G. (1999). The fundamental matrix of the general random walk with absorbing boundaries. Technical report, Collaborative Research Centre Computational Intelligence CI-75, University of Dortmund.
- Rudolph, G. (2000). Takeover times and probabilities of non-generational selection rules. In Whitley, D., Goldberg, D., Cantu-Paz, E., Spector, L., Parmee, I., and Beyer, H.-G., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)*, pages 903–910. Morgan Kaufmann, San Francisco.
- Rudolph, G. (2001). Takeover times of noisy non-generational selection rules that undo extinction. In Kurková, V., Steele, N. C., Neruda, R., and Kárný, M., editors, *Proceedings of International Conference on Neural Networks and Genetic Algorithms, ICANN-NGA 2001*, pages 268–271. Springer, Berlin, Heidelberg, New York.
- Smith, J. and Vavak, F. (1999a). Replacement strategies in steady state genetic algorithms: dynamic environments. *Journal of Computing and Information Technology*, 7(1):49–60.
- Smith, J. and Vavak, F. (1999b). Replacement strategies in steady state genetic algorithms: static environments. In (Banzhaf and Reeves, 1999), pages 219–234.
- Tamhane, A. (1979). A comparison of procedures for multiple comparisons of means with unequal variances. *Journal of the American Statistical Association*, 74:471–480.
- Thierens, D. (1997). Selection schemes, elitist recombination and selection intensity. In Bäck, T., editor, *Proceedings of the 7th International Conference on Genetic Algorithms*, pages 152–159. Morgan Kaufmann, San Francisco.
- Thierens, D. and Goldberg, D. (1993). Mixing in genetic algorithms. In Forrest, S., editor, *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 38–45. Morgan Kaufmann, San Francisco.

- van Nimwegen, E. and Crutchfield, J. (2001). Optimizing epochal evolutionary search: Population-size dependent theory. *Machine Learning*, 45(1):77–114.
- van Nimwegen, E., Crutchfield, J., and Mitchell, M. (1999). Statistical dynamics of the Royal Road genetic algorithm. *Theoretical Computer Science*, 229:41–102.
- Vavak, F. (1997). *Genetic Algorithm Based Self-Adaptive Techniques for Direct Load Balancing in Nonstationary Environments*. PhD thesis, University of the West of England.
- Vose, M. (1999). *The Simple Genetic Algorithm*. MIT Press, Cambridge, MA.
- Whitley, L., editor (1993). *Foundations of Genetic Algorithms - 2*. Morgan Kaufmann, San Francisco.
- Whitley, L. and Kauth, J. (1988). Genitor: A different genetic algorithm. In *Proceedings of the 3rd Annual Rocky Mountain Conference on AI*, pages 118–130. Rocky Mountain Society for Artificial Intelligence.
- Wright, A. and Zhou, Y. (1999). Markov models of genetic algorithms. In Banzhaf, W., Daida, J., Eiben, A., Garzon, M., Honavar, V., Jakiela, M., and Smith, R., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-1999)*, pages 734–741. Morgan Kaufmann, San Francisco.