Convergence Characteristics of Keep-Best Reproduction

Kay Wiese and Scott D. Goodwin
Department of Computer Science, University of Regina
Regina, Saskatchewan, S4S 0A2, Canada
Phone: (306)-585-5210, Fax: (306)-585-4745
email: wiese@cs.uregina.ca, goodwin@cs.uregina.ca

Key Words: Selection Schemes, Family Competition, Convergence Models, Selection Intensities

Abstract

This paper presents theoretical convergence characteristics of *Keep-Best Reproduction* (KBR), a selection strategy for genetic algorithms (GAs). KBR was previously introduced and encouraging results were reported in the traveling salesman domain [16, 18] where KBR was compared with the standard replacement strategy of replacing the two parents by their two children. Here we demonstrate that in a non-operator environment as well as in the ONEMAX domain KBR has the same convergence characteristics as 2-tournament selection and elitist recombination (ELR) [13]. We also show how a modification of ELR suggested in [15] can be utilized to tune the selection pressure of KBR. These analytical models are fairly simplistic and cannot accurately model the convergence characteristics in more complex domains where building blocks are correlated, such as the TSP domain. We will give some empirical results of a comparison of KBR and ELR in this domain.

1 Introduction

Selection plays a vital role in every evolutionary algorithm. Without selection the search process becomes random and promising regions of the search space would not be favoured over non-promising regions. In order to have an efficient and effective search there must be a search criteria (the fitness function) and a selection process that gives individuals with higher fitness a higher chance of being selected for reproduction, mutation and survival.

Depending on the selection strategy that is used, there will always be a tradeoff between *exploration* and *exploitation* of the search space. Informally, exploration is the part of the search that "discovers" new and hopefully promising regions in the search space, while exploitation is the part of the search that stays in one region of the search space and tries to improve individuals in that locality. Both exploration and exploitation are important for a successful search and do usually compete with each other in the sense that too much exploration means too little exploitation and vice versa.

The selection strategy that is employed determines how much exploration and how much exploitation is performed. One measure to quantify selection strategies is the *selection pressure*. Informally, a high selection pressure means that highly fit individuals are selected in disproportionate numbers of samples, while individuals of lower fitness are often not selected at all and are lost during the search process.

A higher selection pressure leads to less exploration of the search space, and more exploitation of so-called "super-individuals", which often leads to a loss of diversity in the population and ultimately to premature convergence. On the other hand, a selection strategy with low selection pressure does not differentiate as much (or at all) between good and bad individuals. This ultimately leads to less exploitation of highly fit individuals and can slow down the convergence speed of the search. Too low a selection pressure could prevent the search from converging at all if the selection strategy forces the search to "jump" from one part of the search space to another all the time.

Section 2 reviews related work on selection schemes. Section 3 reviews KBR, a selection scheme based on family competition. In Sect. 4 we derive the selection intensity of KBR and a convergence model of KBR. Section 5 compares KBR with elitist recombination in the TSP domain. Section 6 contains the conclusions while some ideas for further research are discussed in Sect. 7.

2 Selection Schemes in Evolutionary Computation

There is a variety of different selection schemes for GAs, genetic programming and evolution strategies. Many of them were designed for one category of evolutionary algorithm, but in most cases they can be applied in other categories. Below we give a brief summary of selection schemes and classify them into one of three categories: 1) parent selection schemes, 2) global competition and replacement schemes, and 3) local competition (or family competition) schemes.

2.1 Parent Selection

The parent selection scheme decides who is allowed to pass on their genes into the next generation either through cloning, recombination with another parent, or mutation of the parent. In general we distinguish between fitness proportional selection and ranking methods. Two popular fitness proportional schemes are roulette wheel selection and stochastic universal sampling. Linear and exponential ranking as well as tournament selection are examples of ranking methods. For an overview of parent selection schemes we refer to [2].

2.2 Global Competition and Replacement Schemes

In a standard GA the parents are chosen by one of the parent selection schemes discussed in the previous subsection, then the genetic operators are applied and the two offspring are inserted into the next generation, while the parents usually die. While operators can be disruptive, it might be worthwhile checking the fitness of the parents and offspring first and then deciding who will go to the next generation and who will die. One of the earliest competitive schemes were the (μ, λ) and the $(\mu + \lambda)$ evolution strategy schemes developed by [10]. These strategies work as follows: μ parents produce λ offspring. In the $(\mu + \lambda)$ strategy the μ best individuals out of the union of the μ parents and the λ offspring form the next parent generation, while in the (μ, λ) strategy only the μ best out of the λ offspring ($\mu < \lambda$) form the next parent population. Note that in these cases there is no explicit parent selection from a pool of individuals; the μ best individuals simply become the next parent population. Eshelman's CHC algorithm [5] employs a similar idea. The children population and the parent population are merged and ranked according to fitness. The top s individuals of this merged set are selected to be the next generation. In addition, highly disruptive operators are used to have a more explorative search. All of the selection schemes discussed so far are generational. Contrary to those, Syswerda's steady state GA (SSSGA) [12] and Whitley's GENITOR [19] use a steady state replacement scheme, where only one recombination takes place during each generation. The single offspring usually replaces the worst individual in the population.

2.3 Local Competition and Replacement Schemes

In contrast to global competition, local competition takes place within a family of usually two parents and two offspring. Mahfoud has developed an algorithm called Deterministic Crowding where a parent competes with its genotypical (or phenotypical) most similar child [8]. His study has shown encouraging results for multimodal function optimization. Culberson has developed the Genetic Invariance Genetic Algorithm (GIGA), where the pair of children competes against the pair of parents for replacement [4]. Altenberg has developed Upward Mobility Selection in the genetic programming domain. The basic idea is that offspring is only inserted into the new population if it is fitter than its parents [1]. Elitist Recombination (ELR) was proposed by Goldberg and Thierens [13]. Parents are randomly chosen and mated. The two best out of the union of the two parents and the two offspring are inserted into the next generation. Essentially, Upward Mobility Selection and ELR are basically equivalent implementations of the same family competition scheme.

3 Keep-Best Reproduction: A Family Competition Scheme

What properties do we want a selection strategy to have? First, we want to preserve previous good genetic material, so it can be exploited further. Second, we want the search to make progress in the form of highly fit individuals, so that new promising regions of the search space can be explored. Third, we want to have fast convergence (by increasing the selection pressure) but avoid premature convergence (by maintaining diversity).

In order to preserve good genetic information as well as to introduce new, good genetic information into the population, we had previously proposed the following intermediate selection strategy: Keep only the best of the 2 offspring chromosomes and replace the other by the best parent. Since this ensures that both the best offspring and parent chromosome are kept, we call this technique *Keep-Best Reproduction* or short KBR. It is intuitively clear that KBR has a higher selection pressure than the standard replacement technique of replacing both parents by their two offspring. We refer to the latter selection strategy as *Standard Selection* or short STDS. Both use the same parent selection strategy but KBR employs an additional selection step on the parents and children in order to decide who will survive into the next generation. By keeping the best child we seek to achieve fast convergence. Through controlling the selection pressure by keeping the best parent we seek to prevent premature convergence.

KBR should not be confused with tournament selection. Tournament selection is a parent selection method that randomly chooses s individuals from the population and the best of those s individuals becomes a parent. Here s is the size of the tournament. Increasing s increases the selection pressure. KBR works locally only on the set of the parents and the set of the children. It does not have the same random component as tournament selection has. Also tournament selection only decides who is chosen for reproduction, KBR decides who will live into the next generation.

Also, it is obvious how KBR differs from ELR. However, in Sect. 4 we will argue that 2-tournament selection, ELR and KBR have the same convergence characteristics in non-operator models, as well as in the ONEMAX domain with crossover.

4 Selection Intensity and Convergence Models

The selection pressure of a selection scheme is usually quantified by its selection intensity I:

$$I(t) = \frac{S(t)}{\sigma(t)} = \frac{\overline{f^s(t)} - \overline{f(t)}}{\sigma(t)} \tag{1}$$

Here the selection differential $S(t)$ is the difference between the average fitness of the parent population at generation t, $\overline{f^s(t)}$, and the population mean fitness at generation t, $\overline{f(t)}$. $\sigma(t)$ is the standard deviation from the mean fitness $\overline{f(t)}$ at generation t.

Assuming a standardized normal distribution of the initial population's fitnesses, i.e. $N(\overline{f}, \sigma) = N(0,1)$, the selection intensity I simply becomes the expected average fitness of the population after applying the selection scheme to the original population. Thus we can write the selection intensity I independent of the generation t as:

$$I = \overline{f^s} \tag{2}$$

For the remainder of this paper we will use this notion of I being independent of the generation number t. This is exactly the model that Blickle and Thiele used to compute selection intensities [3]. They have derived the selection intensity for tournament selection of size s to be

$$I(s) = \int_{-\infty}^{\infty} s\, x \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \left(\int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy \right)^{s-1} dx \tag{3}$$

Note that I is dependent on s, the size of the tournament, but not dependent on the generation number t. According to [3] these integral equations can be solved analytically for the cases $s = 1, \ldots, 5$. For example for a tournament of size 1 the selection intensity is $I(1) = 0$ and for a tournament of size 2 it is $I(2) = \frac{1}{\sqrt{\pi}}$. For tournaments of size $s > 5$, the integral equation has to be solved numerically. Alternatively for tournament sizes of $s > 5$ Blickle and Thiele derived an approximation formula with a relative error of less than 1%:

$$I(s) \approx \sqrt{2(\ln(s) - \ln(\sqrt{4.14 \ln(s)}))} \tag{4}$$

This approximation formula can also be used for $s \in [2,5]$ with a relative error of less than 2.4%. Table 1 shows the selection intensities of tournament selection for various tournament sizes.

Table 1: Selection intensities $I(s)$ for tournaments of size s

s	1	2	3	4	5
I(s)	0.00	0.56	0.85	1.03	1.16

For a tournament of size 2, Thierens and Goldberg [14] derive the same selection intensity (in the form of the population average fitness increase from one generation to the next in the ONEMAX domain) but in a completely different manner as Blickle and Thiele. Thierens and Goldberg's formulation can not be extended to other tournament sizes. However, they also derive a convergence model for 2-tournament in the ONEMAX domain:

$$p(t) = \frac{1}{2}(1 + \sin(\frac{t}{\sqrt{\pi l}})) \qquad (5)$$

$p(t)$ is the proportion of 1-bits in the total population at generation t, while l is the bitlength of the chromosomes. For a randomly initialized population $p(0) = 0.5$ can be assumed. To compute g_{conv}, the total number of generations the population needs to fully converge, we set $p(g_{conv}) = 1$, and solve for g_{conv}. We find

$$g_{conv} = \frac{\pi}{2}\sqrt{\pi l} \qquad (6)$$

In the same paper, Thierens and Goldberg showed that their ELR algorithm has the exact same selection intensity and convergence characteristics as 2-tournament selection. In [13] they showed that when optimizing the ONEMAX function, the best parent will go to the next generation and the worst parent will be replaced by the best child. This is easy to understand if we consider that in the ONEMAX domain the total number of 1-bits before and after crossover remains the same.

This is exactly what KBR explicitly does. Originally we proposed to use roulette wheel selection as a parent selection for KBR. If instead we use random parent selection, the selection intensity for KBR is the same as for 2-tournament selection and for elitist recombination, namely $I = \frac{1}{\sqrt{\pi}}$. The convergence model of KBR with random parent selection is the same as for 2-tournament selection and ELR in the ONEMAX domain:

$$p(t) = \frac{1}{2}(1 + \sin(\frac{t}{\sqrt{\pi l}})) \qquad (7)$$

and

$$g_{conv} = \frac{\pi}{2}\sqrt{\pi l} \qquad (8)$$

While the selection intensity of tournament selection can be tuned by changing the size of the tournament, both ELR and KBR have fixed selection intensities. Thierens [15] proposed a modified elitist recombination that allows to tune the selection pressure of ELR, much in the same way as this can be done for tournament selection by modifying the size of the tournament. He has proposed to select one parent via a tournament of size s and to select the other parent randomly from the population. Then the local ELR family competition is applied. Thierens' model assumes a heritability of 1, which he achieves by not applying genetic operators. Since Thierens' model does not take into account genetic operators, children are simply copies of their parents. We have already argued that the two fittest individuals from the two parents and children are the fittest parent and its child clone, which are then inserted into the next generation by ELR. Again, this is exactly what KBR does explicitly. We can use Thierens' modified parent selection to tune selection intensities for KBR in the same way it is done for ELR.

The selection intensity of this modified KBR can be computed as the $(s+1)^{th}$ order statistics of a random sample of size $s+1$, which is also the population mean fitness increase since the standard deviation of the starting population is 1. Using the notation of Blickle and Thiele [3] we can write the selection intensity of the modified KBR as:

$$I(s) = \int_{-\infty}^{\infty} (s+1)\, x \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \left(\int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy \right)^{s} dx \qquad (9)$$

This integral equation is analytically solvable for $s \leq 4$. For $s > 4$ it can be solved numerically or by using a modification of Blickle and Thiele's formula:

$$I(s) \approx \sqrt{2(\ln(s+1) - \ln(\sqrt{4.14 \ln(s+1)}))} \qquad (10)$$

The relative error of this approximation is less then 2.4% for $s \in [2,4]$ and less than 1% for $s > 4$. Table 2 shows the selection intensities of this modified KBR. We can conclude that the modified KBR with tournament size s has the same selection intensity as regular tournament selection with tournament size $s+1$.

Table 2: Selection intensities $I(s)$ for the modified KBR. One parent is selected by a tournament of size s, while the other parent is selected at random

s	1	2	3	4	5
I(s)	0.56	0.85	1.03	1.16	1.27

Thierens and Goldberg [13] have performed experiments in the ONEMAX domain as well as on bounded fully deceptive functions and empirically compared tournament selection of size 2 and standard replacement with ELR. Their conclusions were that in the ONEMAX domain the two selection schemes show very little difference. In the deceptive function domain ELR performed slightly better than tournament selection. This became even more evident when the populations were undersized.

We have not performed empirical studies in the ONEMAX domain that compare KBR with ELR and tournament selection, since in this domain KBR and ELR are literally the same algorithms and one cannot expect that our results for KBR would differ from the ones found in [13]. The above convergence analysis of KBR is limited to either the ONEMAX domain or a model with heritability one, meaning that the offspring are no different from the parents. We believe that the difference of KBR and ELR can only be shown in a domain where longer building blocks need to be processed and operators can have a disruptive effect on these building blocks. We choose to compare KBR and ELR in the TSP domain which is also a problem domain of practical interest.

5 Comparison of KBR and ELR in the TSP domain

We have compared KBR and ELR on a 100 city asymmetric travelling salesman problem. The cost between two cities was a random integer number between 0 and *maxcost*, where *maxcost* was set to 100 times the number of cities. The parent selection was done via roulette wheel selection. The mutation operator was a simple swap operation that picks two random locations in the tour, and exchanges the two cities in those locations. The crossover operator we used was the partially mapped crossover (PMX). A detailed description of PMX can be found in [6]. The fitness function we used was $f_i = c_{max} - c_i$, where c_i is the actual tour cost of individual i and c_{max} is the maximum cost in the population. All results were averaged over 10 independent runs with different random seeds. For each selection strategy 100 different combinations of crossover probability P_c and mutation probability P_m were used with P_c ranging from 0.1 to 1.0 with increments of 0.1 for STDS, KBR, and ELR and P_m ranging from 0.01 to 0.1 with increments of 0.01 for STDS and P_m ranging from 0.1 to 1.0 with increments of 0.1 for KBR and ELR. All results shown are with finetuned operator probabilities for optimal results.

5.1 Recombination Alone

Figure 1 shows that both KBR and ELR converge prematurely. ELR more so than KBR. For comparison we have also depicted the convergence of STDS in Fig. 1. Note that STDS, although applying selection pressure through roulette wheel selection, does not suffer from the same premature convergence. The premature convergence of KBR and ELR is due to a loss of diversity. In order to reintroduce diversity we use higher mutation rates in combination with crossover.

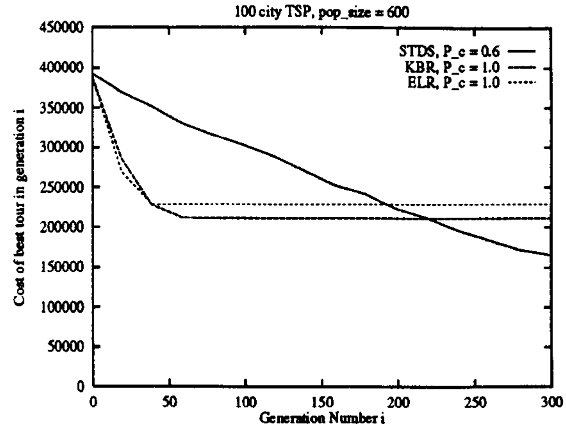

Figure 1: The x-axis shows the generation number i, while the y-axis shows the cost of the best tour after i generations for a 100 city random TSP. No mutation was used. The crossover probability was set to $P_c = 0.6$ for STDS and to $P_c = 1.0$ for KBR and ELR. The population size was 600

5.2 Recombination and Mutation

With KBR we were able to speed up the convergence of the GA by using higher mutation rates. This should not come as a surprise. While with STDS, mutation is performed and the mutated chromosomes are inserted into the next generation, KBR only keeps the best child. In case mutation lowers the fitness of the offspring, there is always the good genetic material of the best parent that is kept. So higher mutation rates are not as disruptive as with STDS. On the other hand, without mutation, KBR very rapidly converges to local optima of low quality (see Subsect. 5.1). The higher mutation rates reintroduce diversity and help steer the GA away from these inferior local optima. A similar argument can be made for ELR. According to Fig. 2, ELR converges even more rapidly, but fails to find better solutions than KBR. The cost of the cheapest tour found for a population size of 600 was 86,719 with ELR and 85,541 with KBR after 600 generations. In fact the population for ELR was fully converged after 220 generations, while for KBR after 600 generations there was still diversity in the population and room for exploration. Similar findings were made with other population sizes and problem sizes.

5.3 Varying Operator Probabilities

Both KBR and ELR also make the underlying GA more robust in the sense that small changes in genetic operator probabilities do not lead to a large change

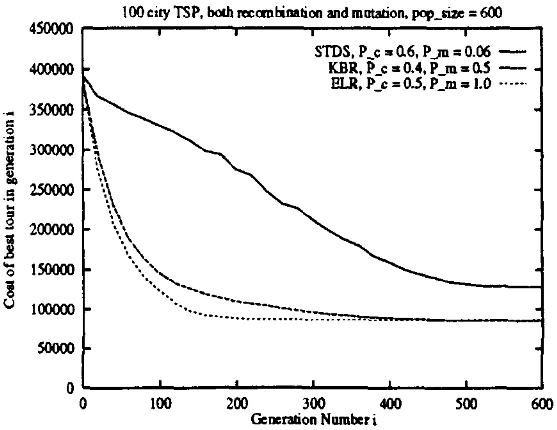

Figure 2: The x-axis shows the generation number i, while the y-axis shows the cost of the best tour after i generations for a 100 city random TSP. The population size was 600

Table 3: Performance difference after small change of operator probabilities. While the mutation probability P_m was not changed, the crossover probability P_c was changed from 0.6 to 0.7

STDS	KBR	ELR
best tour P_c, P_m	best tour P_c, P_m	best tour P_c, P_m
129,020 0.6, 0.05	89,200 0.6, 0.5	90,718 0.6, 0.5
225,868 0.7, 0.05	89,345 0.7, 0.5	92,030 0.7, 0.5

in algorithm performance. KBR is not affected as much by small changes to operator probabilities as is STDS. Table 3 shows the effect of a small change in the crossover probability from 0.6 to 0.7. While after the change in crossover probability, the tour-cost increases by 0.16% for KBR and by 1.45% for ELR, it increases by 75.1% if STDS is used.

6 Conclusion

We have derived the selection intensity of KBR in a simplistic non-operator environment as well as for the ONEMAX domain to be $I = \frac{1}{\sqrt{(\pi)}}$. Also we have derived a convergence model for KBR in those two domains. Both selection intensity and convergence model of KBR are identical to the selection intensity and convergence models for 2-tournament selection and ELR. We have demonstrated how an idea introduced by Thierens [15] can be used to tune the selection intensity of KBR in the same way the selection intensity can be tuned for tournament selection by modifying the tournament size. We believe the difference between KBR and ELR can only be shown in a domain where building blocks are correlated, such as the TSP domain. In this domain both KBR and ELR show similar advantages when compared with the standard selection strategy of replacing both parents by their offspring, such as a more efficient and more effective search. Also both KBR and ELR work well with smaller population sizes when compared to STDS. One of the differences between KBR and ELR in the TSP domain is that ELR converges more rapidly, but usually towards solutions of lesser quality than the ones found by KBR. For the tests that we have performed the tours found by ELR were on average about 3% more expensive than the tours found by KBR.

7 Further Research and Discussion

The analysis of KBR in Sect. 4 assumed that the parent population is chosen randomly without replacement. For the empirical evaluation, however, we have used roulette wheel selection. This could give an advantage to KBR, since KBR always keeps the better parent. It remains to be seen how well KBR and ELR compare to each other when the parents are selected randomly. We believe that KBR will benefit from a fitness biased parent selection while ELR probably works best with random parent selection. Currently we are working on developing a convergence model for both KBR and ELR if they are combined with fitness biased parent selection.

Since KBR as a selection strategy works on fitness values only, it should work well in other problem domains. This however remains to be shown empirically in every individual case. KBR can easily be added to any existing generational GA, which makes it easy for others to try KBR with their particular application.

It would be interesting to see how other GA techniques such as other parent selection strategies (tournament, random, ...) and – for the TSP domain – other genetic operators (OX, PBX, CX, ...) affect the performance of KBR and ELR.

8 Acknowledgements

We would like to thank the anonymous reviewers for their insightful comments that have helped to improve the technical quality of this paper.

References

[1] Altenberg, L., "The Evolution of Evolvability in Genetic Programming", *Advances in Genetic Programming*, pp. 47–74, 1994.

[2] Baeck, T., *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*, Oxford University Press, 1996.

[3] Blickle, T., and Thiele, L., *A Comparison of Selection Schemes Used in Genetic Algorithms*, TIK Report Nr.11, 2nd Edition, Swiss Federal Institute of Technology, 1995.

[4] Culberson J., "Crossover versus Mutation: Fueling the Debate: TGA versus GIGA", *Proceedings of the 5th International Conference on Genetic Algorithms ICGA-93*, pp. 632, 1993.

[5] Eshelman, L.J., "The CHC Adaptive Search Algorithm: How to Have Safe Search When Engaging in Nontraditional Genetic Recombination", *Foundations of Genetic Algorithms, FOGA-I*, pp. 265–283, 1991.

[6] Goldberg, D.E. and Lingle, R., Jr., "Alleles, Loci, and the Traveling Salesman Problem", *Proceedings of the 1st International Conference on Genetic Algorithms ICGA-85*, pp. 154–159, 1985.

[7] Holland, J.H., *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI, University of Michigan Press, 1975.

[8] Mahfoud, S.W., "Crowding and Preselection Revisited", *Proceedings of Parallel Problem Solving from Nature PPSN-II*, pp. 27–36, 1992.

[9] Oliver, I.M., Smith, D.J., and Holland, J.R.C.,, "A Study of Permutation Crossover Operators on the Traveling Salesman Problem", *Proceedings of the 2nd International Conference on Genetic Algorithms ICGA-87*, pp. 224-230, 1987.

[10] Schwefel, H.P., "Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie", *Interdisciplinary Systems Research*, Volume 26, Birkhaeuser, Basel, 1977.

[11] Starkweather, T., McDaniel, S., et. al, "A Comparison of Genetic Sequencing Operators", *Proceedings of the 4th International Conference on Genetic Algorithms ICGA-91*, pp. 69–76, 1991.

[12] Syswerda, S., "Uniform Crossover in Genetic Algorithms", *Proceedings of the 3rd International Conference on Genetic Algorithms, ICGA-89*, pp. 2–9, 1989.

[13] Thierens, D. and Goldberg, D.E., "Elitist Recombination: An Integrated Selection Recombination GA", *Proceedings of the 1st IEEE World Congress on Computational Intelligence*, pp. 508–512, 1994.

[14] Thierens, D. and Goldberg, D.E., "Convergence Models of Genetic Algorithm Selection Schemes", *Parallel Problem Solving from Nature PPSN-III*, Lecture Notes in Computer Science Vol. 866, Springer Verlag, pp. 119-129, 1994.

[15] Thierens, D., "Selection Schemes, Elitist Recombination and Selection Intensity", *Proceedings of the 7th International Conference on Genetic Algorithms ICGA-97*, pp. 152–159, 1997.

[16] Wiese, K. and Goodwin, S.D., "Keep-Best Reproduction: A Selection Strategy for Genetic Algorithms", *Proceedings of the 1998 ACM Symposium on Applied Computing SAC'98*, pp. 343–348, 1998.

[17] Wiese K. and Goodwin, S.D., "Parallel Genetic Algorithms for Constrained Ordering Problems", *Proceedings of the 11th International Florida Artificial Intelligence Research Symposium, FLAIRS'98*, pp. 101–105, 1998.

[18] Wiese, K. and Goodwin, S.D., "The Effect of Genetic Operator Probabilities and Selection Strategies on the Performance of a Genetic Algorithm", *Advances in Artificial Intelligence*, Lecture Notes in Artificial Intelligence Vol. 1418, Springer Verlag, Germany, pp. 141–153, 1998.

[19] Whitley, D., "The GENITOR Algorithm and Selection Pressure: Why Rank-Based Allocation of Reproductive Trials is Best", *Proceedings of the 3rd International Conference on Genetic Algorithms ICGA-89*, pp. 116-121, 1989.

[20] Whitley, D., Starkweather, T., and D'Ann Fuquay, "Scheduling Problems and Traveling Salesman: The Genetic Edge Recombination Operator", *Proceedings of the 3rd International Conference on Genetic Algorithms ICGA-89*, pp. 133–140, 1989.