

Greedy Seeding Procedure for GAs Solving a Strip Packing Problem

Carolina Salto
Universidad Nacional de La Pampa
General Pico, Argentina
saltoc@ing.unlpam.edu.ar

and

Enrique Alba - Juan M. Molina
E.T.S.I. Informática, Universidad de Málaga
Málaga, España
{eat,jmmb}@lcc.uma.es

and

Guillermo Leguizamón
Universidad Nacional de San Luis
San Luis, Argentina
legui@unsl.edu.ar

Abstract

In this paper, the two-dimensional strip packing problem with 3-stage level patterns is tackled using genetic algorithms (GAs). We evaluate the usefulness of a greedy seeding procedure for creating the initial population, incorporating problem knowledge. This is motivated by the expectation that the seeding will speed up the GA by starting the search in promising regions of the search space. An analysis of the impact of the seeded initial population is offered, together with a complete study of the influence of these modifications on the genetic search. The results show that the use of an appropriate seeding of the initial population outperforms existing GA approaches on all the used problem instances, for all the metrics used, and in fact it represents the new state of the art for this problem.

Keywords: Genetic Algorithms, Strip Packing, Seeding

1 INTRODUCTION

The two-dimensional strip packing problem (2SPP) arise in many real-world applications such as in the paper or textil industries. Typically, the 2SPP consists of a set of M rectangular pieces, each defined by a *width* and a *height*, which have to be packed in a larger rectangle with a fixed width W and unlimited length, designated as the *strip*. The search is for a layout of all the pieces in the strip that minimizes the required strip length with the following restrictions: all pieces have to be packed with their sides parallel to the sides of the strip, without overlapping,

and rotations are not allowed. This problem is similar to the one of cutting the pieces out of the strip by means of orthogonal cuts, minimizing the consumed strip.

Additionally, another constraint is included in our problem: we consider only 3-stage level (guillotine) packing patterns. In these patterns, pieces are packed by horizontal levels (parallel to the bottom of the strip). Inside each level pieces are packed bottom left justified and, when there is enough room in the level, pieces with the same width are stacked one above the other. Many real application of 2D cutting and packing in the glass, wood, and paper industries consider 3-stage level patterns, for that the importance of incorporating this restriction in the problem formulation.

The 2SPP is NP-hard [9]. A few exact approaches for this problem are known [6, 14], although using metaheuristics is the usual way to solve it. Regarding the existing surveys of metaheuristics in the literature, Hopper and Turton [8, 9] review the approaches developed to solve 2D packing problems using GAs; Simulated Annealing, Tabu Search, and artificial Neural Networks are also considered. They conclude that Evolutionary Algorithms (EAs) [4, 15] are the most widely investigated meta-heuristics in the area of cutting and packing. Lodi et. al [13] consider several methods for the 2SPP in their survey and discussed also mathematical models; specially the case where the items have to be packed into rows forming levels are discussed in detail. Other representative works about GAs applied to solve level packing problems are [5], [10], [11], and [16] and 3-stage guillotine cuts [17, 18, 21, 22, 23]. The majority of the approaches imposing this last restriction deal with bin packing problems, for that the difficulty to find works in strip packing to compare with.

In this article we use a GA as the general driving force to locate the region in which a solution of minimum length is located. GAs deal with a population of tentative solutions, on which genetic operators are applied in an iterative manner to progressively compute new solutions of higher quality. We study a hybrid approach where a GA is combined with a heuristic placement routine. The GA is used to determine the sequence in which the pieces are to be packed, and the placement routine determines the layout of the pieces onto the strip in order to generate a 3-stage guillotine pattern.

The efficiency of a GA could be improved by increasing the quality of the initial population. This can be done by adopting simple rules. We here investigate the advantages of seeding the initial population using a set of greedy rules including information of the problem (such as the pieces width, the pieces area, etc.), resulting in a more specialized initial population. The main goal of this paper is to build an improved GA to solve larger problems than the ones found in the literature at present, and to quantify the effects of including a seeding procedure into the algorithms, looking for the best trade-off between exploration and exploitation in the search process, what is considered the key point to perform accurately and efficiently on complex applications.

The organization of the paper is as follows. The components of the used algorithm are described in Section 2. Section 3 describes the greedy generation of the initial population. In Section 4, we explain the parameter settings of the algorithms used in the experimentation. Section 5 reports on the algorithm performances. Finally, in Section 6, we give some conclusions and analyze future search directions.

2 A HYBRID GA FOR THE 2SPP

In Algorithm 1 we can see the structure of the basic steady-state GA ($(\mu + 1)$ -GA) we use for solving the 2SPP. This algorithm creates an initial population $P(0)$ of μ solutions in a random




Figure 1: cutting pattern for the permutation 2 4 1 5 6 3 8 7.

(uniform) way, and then evaluates these solutions. The evaluation uses a placement algorithm to arrange the pieces in the strip to construct a feasible packing pattern. After that, the population goes into a cycle where it undertakes evolution. This cycle involves the selection of two parents by binary tournament and the application of recombination and mutation operators to create a new solution ($P'(t)$). The new generated individual replaces the worst individual in the population only if it is fitter. In this study, the stopping criterion for the cycle is to reach a maximum number of evaluations (*max_evaluations*). The best solution is identified as the best individual ever found which minimizes the required strip length.

Algorithm 1 Genetic algorithm

```

GA
 $t = 0$ ; {current generation}
initialize( $P(t)$ );
evaluate( $P(t)$ );
while ( $t < \text{max\_evaluations}$ ) do
   $P'(t) = \text{evolve}(P(t))$ ; {recombination and mutation}
  evaluate ( $P'(t)$ );
   $P(t+1) = \text{select}(P'(t), P(t))$ ;
   $t = t + 1$ ;
end while

```

In most cases, applying a GA means devising a customized representation for encoding a candidate solution. We encode a packing pattern into a chromosome as a sequence of pieces that defines the input for the layout algorithm. Therefore, a chromosome will be a permutation $\pi = (\pi_1, \pi_2, \dots, \pi_M)$ of M natural numbers (piece identifiers). In order to generate a 3-stage level pattern, a modified *next-fit* heuristic (NF) is used here—in the following referred as *modified next-fit*, or *MNF*—which was proved to be very efficient in [18, 20]. This heuristic gets a sequence of pieces as its input and constructs the packing pattern by placing pieces into stacks, and then stacks into levels in a greedy way, i.e., once a new stack or a new level is started, previous ones are never reconsidered. See Figure 1 for an illustrative example. Deeper explanation of the MNF procedure can be found in [23].

GAs are guided by the values computed by an objective function for each tentative solution until the optimum or an acceptable solution is found. In our problem, the objective is to minimize the strip length needed to build the layout corresponding to a given solution π . An important consideration is that two packing patterns could have the same length—so their fitness will be equal—although, from the point of view of reusing the trim loss, one of them

can be actually better because the trim loss in the last level (which still connects with the remainder of the strip) is greater than the one present in the last level of the other layout. Therefore we are using the following fitness function:

$$F(\pi) = \text{strip.length} - \frac{l.waste}{l.length * W} \quad (1)$$

where *strip.length* is the length of the packing pattern corresponding to the permutation π , *l.waste* is the area of reusable trim loss in the last level l of the packing pattern and *l.length* is the length of l . Hence, $F(\pi)$ is both simple and accurate.

Regarding the genetic operators, the recombination operator used here, *Best Inherited Level Recombination (BILX)* (introduced in [23] as BIL), incorporates some problem-specific knowledge into their mechanism in order to improve the trim loss. The BILX operator transmits the best levels of the parent to the child, i.e. those with the highest filling rate (*fr*) or, equivalently, with the least trim loss. This rate is calculated as follows, for a given level l :

$$fr(l) = \sum_{i=1}^n \frac{width(\pi_i) \times length(\pi_i)}{W \times l.length} \quad (2)$$

where π_1, \dots, π_n are the pieces in l , *width*(π_i) and *length*(π_i) are the piece dimensions.

Actually, BILX works as follows. Let be nl the number of levels in one parent $parent_1$. In the first step the filling rates of all nl levels from $parent_1$ are calculated. After that, a probability of selection, proportional to its filling rate, is assigned to each level and a number ($nl/2$) of levels are selected from $parent_1$. The pieces π_i belonging to the inherited levels are placed in the first positions of the child. Meanwhile, the remaining positions are filled with the pieces which do not belong to that levels, in the order they appear in the other parent $parent_2$.

On the other side, the idea behind the mutation operator used in this work is to change the location of some pieces so that the final cost is reduced; it was successfully tested in [23]. Named as *Best and Worst Stripe Exchange (BW_SE)*, this mutation changes the location of the best and the worst level. The pieces of the best level (the one with highest filling rate) are allocated in the first positions of the new packing pattern while the pieces of the worst level are assigned to the last positions. The middle positions are filled with the remaining pieces in the order they appeared in the original packing pattern. In BW_SE, the movements can help to the involved levels or their neighbors to accommodate pieces from neighboring levels, thus improving their trim loss.

3 INITIAL SEEDING

The performance of a GA is often related to the quality of its initial population. This quality depends on two important issues: the average fitness of individuals in the population and the diversity in the population. By having an initial population with better fitness values, better final individuals can be found faster [1, 2, 19]. Besides, high diversity in the population inhibits early convergence to a locally optimal solution.

There are many ways to arrange this initial diversity. The idea in this work is to start with a seeded population created by following some building rules, hopefully allowing to reach good solutions in early stages of the search. The rules will include some characteristics from the problem such as piece sizes, and also incorporate ideas from the *best fit* (BF) and *first fit* (FF) heuristics [12].

Table 1: Rules to generate the initial population

#	Rule Description	#	Rule Description
1	sorts pieces by decreasing width.	2	sorts pieces by increasing width.
3	sorts pieces by decreasing length.	4	sorts pieces by increasing length.
5	sorts pieces by decreasing area.	6	sorts pieces by increasing area.
#	Rule Description		
7	sorts pieces by alternating between decreasing width and length.		
8	sorts pieces by alternating between decreasing width and increasing length.		
9	sorts pieces by alternating between increasing width and length.		
10	sorts pieces by alternating between increasing width and decreasing length.		
11	the pieces are reorganized following the BF heuristic.		
12	the pieces are reorganized following the FF heuristic.		

Individuals are generated in two steps. In the first step, the packing patterns are randomly sampled from the search space with a uniform distribution. After that, each of them are modified by one rule, randomly selected, with the purpose of improving the piece location inside the random packing pattern. Each application of a rule yields a (possibly) different solution because of the randomization used in the first step.

The rules for the initial seeding are listed in the Table 1. These rules are proposed in order to produce individuals with improved fitness values and also for introducing diversity in the initial population. Hence, sorting the pieces by their width will hopefully increase the probability of stacking pieces, and then produce more dense levels. On the other hand, sorting by length will generate levels with smaller wasted space above the pieces, especially when the pieces lengths are very similar. BF and FF relocate the pieces with the goal of reducing the trim loss inside a level. The possible new layout obtained in this way has to be transmitted to the chromosome in such a way that we can obtain the same layout by applying MNF to the chromosome. Finally, rules 7 to 10 have been introduced for increasing the initial diversity. As we will see, these rules are not only useful for initial seeding: several of them can be used as a simple greedy algorithm for local search to help during the optimization process.

4 IMPLEMENTATION

The specific $(\mu + 1)$ -GA we have implemented is analyzed here using three different methods of seeding the initial population: (i) by means of a random generation (GA), (ii) by applying one determined rule from the Table 1 (GA_i where i stand for a rule number) and (iii) by applying a rule from Table 1, but a randomly selected one for each individual of the population (GA_{Rseed}). We also consider the random generation of individuals as other applicable rule in this case.

The population size was set to 512 individuals. The maximum number of evaluations was fixed to 2^{16} . The recombination operator was applied with a probability of 0.8, while the mutation probability was set to 0.1. Parameters (population size, stop criterion, probabilities, etc) were not chosen at random, but rather by an examination of values previously used with success (see [21]).

These algorithms were run in MALLBA [3], a C++ software library fostering rapid prototyping of hybrid and parallel algorithms, and the platform was an Intel Pentium 4 at 2.4 GHz and 512 MB RAM, linked by Fast Ethernet, under SuSE Linux with 2.4.19-4GB kernel version.

We have considered five randomly generated problem instances with M equal to 100, 150, 200, 250 and 300 pieces and a known global optimum equal to 200 (the minimal length of the strip). These instances belong to the subtype of the guillotine patterns only, i.e. the optimum value does not correspond to the 3-stage guillotine pattern constraint. They were obtained by

an own implementation of a data set generator, following the ideas proposed in [25] with the length-to-width ratio of all M rectangles in the range $1/3 \leq l/w \leq 3$. These instances are publicly available at <http://mdk.ing.unlpam.edu.ar/~lisi/2spp.htm>.

5 COMPUTATIONAL ANALYSIS

In this section we summarize the results of applying the proposed algorithm with its seeding variants to all the problem instances. Our aim is to offer meaningful results and check them from a statistical point of view.

For each algorithm we have performed 30 independent runs per instance using the parameter values described in the previous section. Also, the evaluation considers two important issues for any search process: the capacity for generating new promising solutions and the pace rate of the progress in the surroundings of the best found solution (fine tuning or intensification). For a meaningful analysis we consider the average fitness values and the entropy measure (as proposed in [7]), which is computed as follows:

$$entropy = \frac{\sum_{i=1}^M \sum_{j=1}^M \left(\frac{n_{ij}}{\mu}\right) \ln\left(\frac{n_{ij}}{\mu}\right)}{M \ln M} \quad (3)$$

where n_{ij} represents the number of times the piece i is set into the position j in the population of size μ . This function takes values in $[0..1]$ and a value of 0 indicates that all the individuals in the population are identical.

Table 2 shows the results obtained for the different methods for generating the initial population for all instances (GA , GA_i and GA_{Rseed}). The most relevant metrics used in this comparison are: the average objective value of the initial population (column *avg_{ini}*), the best found feasible solution (column *best*) and the average value of the best found feasible solution, in the 30 independent runs, along with its standard deviation (column *avg_{±σ}*). The optimal *best* values are printed in bold.

From that results we can point out that any seeded GA starts the search process from better fitness values than in the case of randomly generated initial populations (GA). Best initial populations, in average, are obtained using GA_4 but having poor genetic diversity (see Figure 2 which shows the mean entropy values in the initial population for each algorithm and different M values). Rule 4 arranges the pieces by their height so the generated levels have very similar piece heights and consequently the produced free space inside a level is small. On the other side, a poor performance is obtained with both GA_7 and GA_9 , which present in mean the worst initial population and also with poor genetic diversity. The random selection of rules to generate the initial solution (GA_{Rseed}) works quite well (being in the medium positions of the rule ranking) with an acceptable initial population diversity (entropy value near to 0.8).

Furthermore, a neat conclusion of this study is that the GA_{Rseed} (random seeded initial populations) outperforms significantly the GA with a traditional initialization in all the metrics (the p -value is close to 0) and also the rest of the seeded GAs. This suggests that the efficiency of a GA could be improved simply by increasing the quality of the initial population.

Figure 2 shows that the majority of the seeded GAs present poor initial genetic diversity (less than 0.5), except GA_{12} , GA_{11} and GA_{Rseed} with an initial diversity (near to 1) close to one of the no seeded GA . The two first algorithms apply NF and BF heuristics (respectively) to a random generated solution, i.e. the possibly changes in the piece positions generated by the heuristics does not take into account piece dimensions hence the original piece positions inside the chromosome are more probably maintained. On the other hand, GA_{Rseed} combine

Table 2: Experimental results for the GA regarding different seeding methods.

Alg	M = 100			M = 150			M = 200			M = 250			M = 300		
	avg _i	best	avg±σ	avg _i	best	avg±σ	avg _i	best	avg±σ	avg _i	best	avg±σ	avg _i	best	avg±σ
GA	417.34	235.75	248.13	504.76	241.78	254.60	485.28	249.64	258.50	489.08	243.81	255.59	531.22	253.75	263.77
GA ₁	353.24	240.68	253.25	408.29	239.73	247.79	382.10	243.70	250.72	354.45	236.86	244.25	409.46	285.50	295.50
GA ₂	298.24	231.67	241.05	369.14	241.73	248.58	352.55	232.83	243.65	331.02	231.77	242.58	345.53	251.68	261.55
GA ₃	326.18	264.20	266.20	274.18	229.71	234.71	295.05	230.64	233.64	284.19	232.81	234.39	273.28	245.71	245.96
GA ₄	282.59	230.20	231.12	240.31	228.71	230.70	244.43	231.58	231.96	239.65	229.76	230.45	239.13	234.59	234.70
GA ₅	381.31	342.41	346.56	427.88	333.56	338.02	415.83	295.63	329.08	377.02	271.57	278.35	401.03	326.50	329.00
GA ₆	371.87	294.29	299.42	377.88	297.21	308.33	359.04	265.59	291.18	357.10	313.27	318.02	375.12	321.53	324.97
GA ₇	397.19	256.58	265.98	486.73	259.46	267.51	452.62	252.74	272.19	411.87	252.81	263.83	454.54	281.71	287.63
GA ₈	298.95	240.28	242.03	274.66	249.71	250.70	275.21	240.62	241.50	262.77	234.64	238.38	259.60	246.43	246.49
GA ₉	384.78	255.63	261.95	499.52	264.74	269.33	460.18	248.54	261.45	450.60	256.55	268.84	486.68	280.60	294.07
GA ₁₀	354.03	262.47	279.10	353.80	262.69	290.02	368.99	258.55	271.16	321.36	256.62	270.76	336.15	275.73	287.45
GA ₁₁	277.87	229.78	236.66	290.06	239.72	243.44	278.44	226.80	234.69	273.94	222.70	233.27	283.25	229.58	237.57
GA ₁₂	281.23	233.77	240.05	294.41	242.71	248.36	282.96	231.69	237.84	276.63	229.71	234.10	287.98	227.67	241.60
GA _{Rseed}	340.57	228.59	232.30	369.50	229.04	231.66	358.89	220.71	228.38	340.64	221.61	227.36	360.68	225.49	233.47





Figure 2: Average entropy of initial population. Figure 3: Mean number of evaluations to reach the best value for each instance.

all the previous considerations with random generation of the piece positions, so high genetic diversity was expected.

As supposed, the GA_{Rseed} significantly reduces the number of evaluations required to find good solutions (see Figure 3): they are near two times smaller in mean than GA . To confirm these observations, we used the t -test, which indicates that the difference among the algorithms are significant under this metric (p -values near to 0). GA_7 , GA_9 and GA present similar effort in order to locate their best solutions: all of them have quite similar initial population averages (see Table 2) but GA presents best final solution qualities.

With respect to mean execution times there are a little difference against any seeded GA due to the initialization phase, but this difference is negligible.

After the 30 runs of GA_{Rseed} for all M values, we analyzed the contribution of the solutions

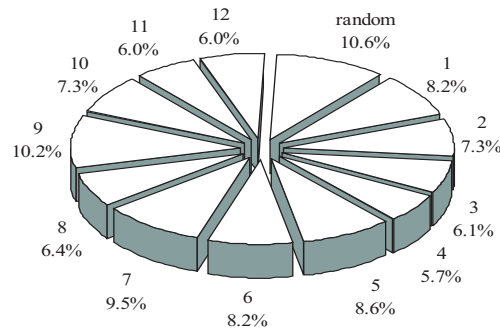
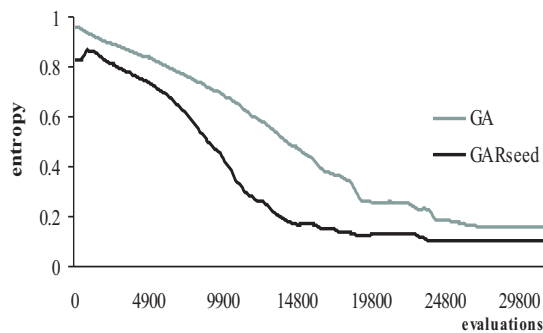



Figure 4: Rule Rank.

Figure 5: Population entropy for $M=200$.Figure 6: Average population fitness for $M=200$.

generated by each rule to the average fitness in the initial population. Regarding that measurement, the rule rank, from best to worst, was as follows: 13, 9, 7, 5, 1, 6, 10, 2, 8, 3, 12, 11 y 4 (see Figure 4), although the differences in percentages are quite small. The rule preserving the original piece position (random generation) are first in the rankings followed by the ones sorting pieces by alternating between increasing/decreasing width and height. By surprise, rules incorporating the BF and FF heuristics appear in last positions, when the expectation was that they could create good packing patterns due to the improvement in the layout they achieved. Also a remarkable point here is the rank position of rule 4 which appeared in the last position, although its use for the initialization of the whole population brings to a better performance than the rest.

5.1 Additional Discussion

If we seek among all the proposed algorithms, we can state that the GA_{Rseed} , which initializes the population considering the whole set of rules including information of the problem, converges to a better final solution, allowing a highly reduced number of evaluations comparing to that of GA using a random sampled initial population. So, for the next study we compare the performance of GA_{Rseed} with the performance of the plain GA .

These algorithms were compared with respect to the phenotypic entropy (Figure 5) as well as the evolution of the average population fitness (Figure 6). These comparisons are illustrated over the instance with $M=200$ (similar results are obtained with the rest of the instances). From these figures, we observe that the GA with seeding has a fast phenotypic entropy decrease to the

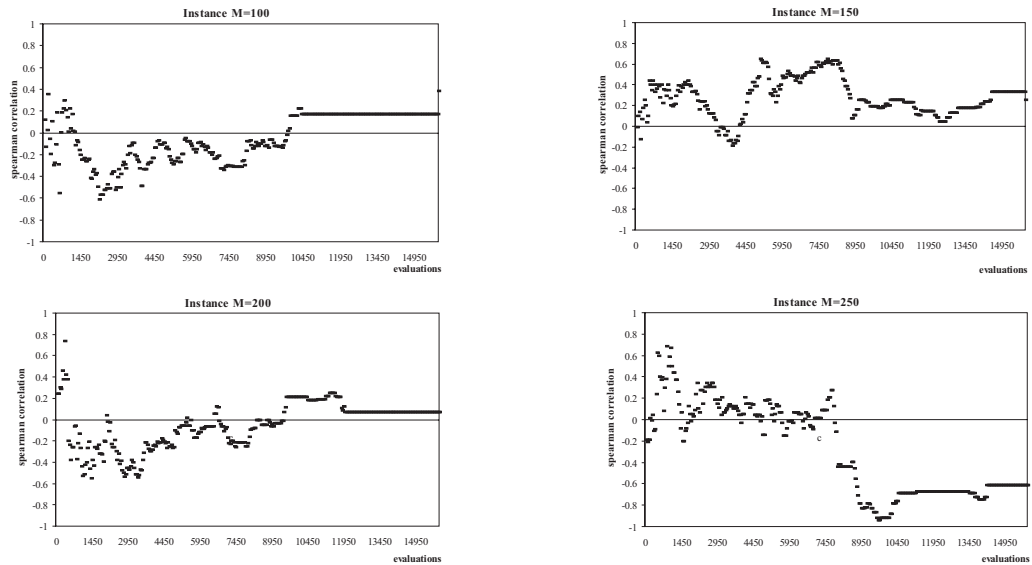


Figure 7: Evolving populations correlation between best fitness in each population and entropy measure. Each point represents the correlation between 15 populations from a 15 runs.

best final solutions (resulting in faster convergence), but the genetic diversity decreases slower than phenotypic diversity thus. In fact, we can see that the seeded initial populations present a high entropy value (near to 0.85 value), i.e., good diversity, but with fitter individuals.

In order to provide a deeper explanation of what is happening during the search of the algorithms we analyze the relationship between the fitness and the population entropy (genotypic diversity), which will have a strong effect on search difficulty. A way to examine this is by the use of the Spearman correlation measure (a nonparametric correlation test), which ranks two sets of variables and tests for a linear relationship between the variables rank. The Spearman correlation coefficient is computed (from [24]) as follows:

$$1 - \frac{6 \sum_{i=1}^N d_i^2}{N^2 - N} \quad (4)$$

where N is the number of items, and d_i is the difference between each pair of fitness and entropy ranks. The correlation values range from +1 (perfect positive correlation), through 0 (no correlation), to -1.0 (negative correlation). For our study, if we see ideal low fitness values, which will be ranked in ascending order (1=best, ..., 15=worst) and high diversity, ranked in ascending order (1=lowest diversity and 15=highest diversity), then the correlation coefficient should be strongly negative. Alternatively a positive correlation indicates that either bad fitness accompanies high diversity or good fitness accompanies low diversity.

Figure 7 shows the correlation between entropy and best fitness for each generation. Note that each point represents the correlation between 15 populations, sampled from 15 runs where there is a dependency of later evaluations on preceding ones. Experiments with instance $M=250$ show a period of fluctuation until evaluation 8000, then the correlation coefficient became strongly negative, indicating that a good fitness is present with good diversity. Experiments with instances with $M=100$ and $M=200$ contain a period of fluctuation until evaluations 11000 approximately, after which entropy lost correlation with best fitness (coefficient correlations near to zero). For these two instances, the relationship between fitness and diversity becomes less important, probably due to other critical relationships. Finally, after a period of early

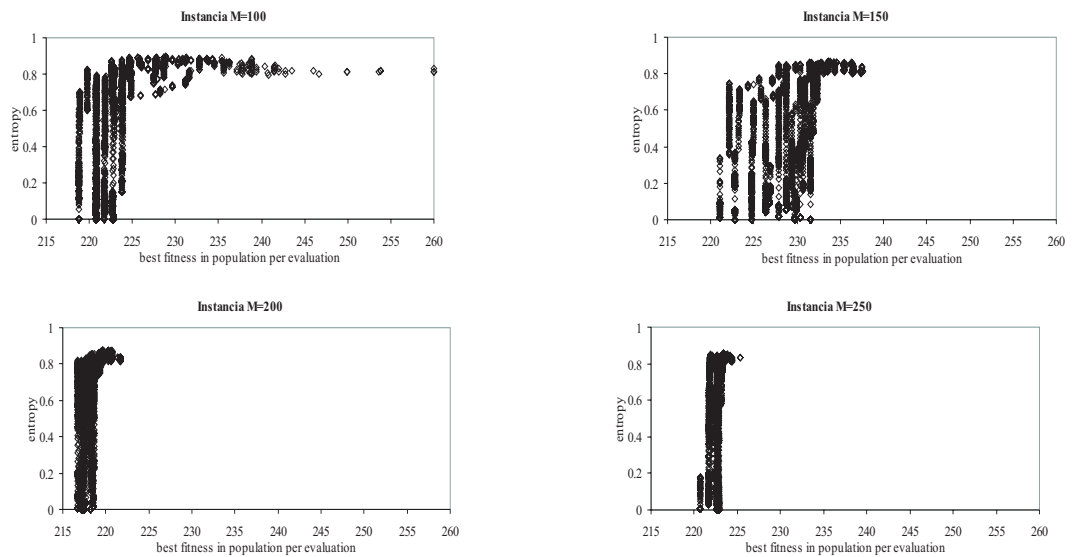


Figure 8: Best fitness plotted against that population's diversity.

fluctuations, the correlation for instance $M=150$ is positive, this can be owing to the fact that, as the best fitness is achieved early in a run, the population is made up of an increasingly larger number of copies of the best fit individual.

Figure 8 plots the best fitness found in the population along the x axis and population's diversity on the y axis. From there we can see some populations with bad fitness (low values are better) occurring with higher entropy. Once the algorithm find a good fitness, the entropy values begin to decrease (the populations are less diverse), i.e. the population is made up of an increasingly larger number of copies of the best fit individual.

6 CONCLUSIONS

In this paper we have investigated different greedy methods of generating the initial population in a traditional GA to solve the 3-stage 2SPP. The study, validated from a statistical point of view, analyzes the capacity of the seeding policy to generate new potentially promising individuals and the ability to keep a diversified population. The problem-aware seeding includes pieces dimensions and known level heuristics.

Our results show that an improvement in the GA performance was observed by using a problem-aware seeding, regarding both efficiency (effort) and quality of the solutions found. Moreover, the random selection of rules to build the initial population works properly, providing good genetic diversity of initial solutions. The performance of a GA is sensitive to the quality of its initial population. By having an initial population with better fitness values, we typically get better final individuals.

As future works we plan to go into the analysis of correlation between mean population fitness and genetic diversity in any depth. Also we propose to investigate non-permutation representations and a direct mapping to the final disposition of the pieces, as well as to construct parallel versions of the algorithms studied in this work.

ACKNOWLEDGEMENTS

This work has been partially funded by the Spanish Ministry of Education and the European FEDER under contract TIN2005-08818-C04-01 (the OPLINK project, <http://oplink.lcc.uma.es>). We also acknowledge the Universidad Nacional de La Pampa, and the ANPCYT in Argentina from which we received continuous support.

REFERENCES

- [1] R.K. Ahuja and J.B. Orlin. Developing fitter genetic algorithms. *INFORMS Journal on Computing*, 9(3):251–253, 1997.
- [2] R.K. Ahuja, J.B. Orlin, and A. Tiwari. A greedy genetic algorithm for the quadratic assignment problem. *Computers & Operations Research*, 27(3):917–934, 2000.
- [3] E. Alba, J. Luna, L.M. Moreno, C. Pablos, J. Petit, A. Rojas, F. Xhafa, F. Almeida, M.J. Blesa, J. Cabeza, C. Cotta, M. Díaz, I. Dorta, J. Gabarró, and C. León. *MALLBA: A Library of Skeletons for Combinatorial Optimisation*, volume 2400 of *LNCS*, pages 927–932. Springer, 2002.
- [4] T. Bäck, D. Fogel, and Z. Michalewicz. *Handbook of evolutionary computation*. Oxford University Press, New York, 1997.
- [5] A. Bortfeldt. A genetic algorithm for the two-dimensional strip packing problem with rectangular pieces. *European Journal of Operational Research (article in press)*, 2005.
- [6] S.P. Fekete and J. Schepers. On more-dimensional packing III: Exact algorithm. Technical Report ZPR97-290, Mathematisches Institut, Universität zu Köln, available from the first author at Department of Mathematics, 1997.
- [7] J.J. Grefenstette. *Genetic Algorithms and Simulated Annealing*, chapter Incorporating problem specific knowledge into genetic algorithms, pages 42–60. Morgan Kaufmann Publishers, 1987.
- [8] E. Hopper. *Two-dimensional packing utilising evolutionary algorithms and other meta-heuristic methods*. PhD thesis, University of Wales, Cardiff, U.K., 2000.
- [9] E. Hopper and B. Turton. A review of the application of meta-heuristic algorithms to 2d strip packing problems. *Artificial Intelligence Review*, 16:257–300, 2001.
- [10] S. Hwang, C. Kao, and J. Horng. On solving rectangle bin packing problems using genetic algorithms. *IEEE International Conference on Systems, Man, and Cybernetics - Humans, Information and Technology*, 2:1583–1590, 1994.
- [11] B. Kroger. Guillotineable bin-packing: a genetic approach. *European Journal of Operational Research*, 84:645–661, 1995.
- [12] A. Lodi, S. Martello, and M. Monaci. Recent advances on two-dimensional bin packing problems. *Discrete Applied Mathematics Journal of Operation Research*, 141:241–252, 2002.
- [13] A. Lodi, S. Martello, and M. Monaci. Two-dimensional packing problems: a survey. *European Journal of Operational Research*, 141:241–252, 2002.
- [14] S. Martello, S. Monaci, and D. Vigo. An exact approach to the strip-packing problem. *Informes Journal on Computing*, 15:310–319, 2003.
- [15] M. Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer, third revised edition, 1996.
- [16] C.L. Mumford-Valenzuela, J. Vick, and P.Y. Wang. *Metaheuristics: Computer Decision-Making*, chapter Heuristics for large strip packing problems with guillotine patterns: An empirical study, pages 501–522. Kluwer Academic Publishers BV, 2003.
- [17] J. Puchinger and G.R. Raidl. Models and algorithms for three-stage two-dimensional bin packing. Technical report, Technische Universität Wien, Institut für Computergraphik und Algorithmen, 2004.
- [18] J. Puchinger, G.R. Raidl, and G. Koller. *Solving a Real-World Glass Cutting Problem*, volume 3004 of *LNCS*, pages 162–173. Springer, 2004.
- [19] C.R. Reeves. A genetic algorithm for flowshop sequencing. *Computers & Operations Research*, 22(1):5–13, 1995.
- [20] C. Salto, J.M. Molina, and E. Alba. Sequential versus distributed evolutionary approaches for the two-dimensional guillotine cutting problem. *Proceedings of International Conference on Industrial Logistics (ICIL 2005)*, pages 291–300, 2005.
- [21] C. Salto, J.M. Molina, and E. Alba. Analysis of distributed genetic algorithms for solving cutting problems. *International Transactions in Operational Research*, 13(5):403–423, 2006.

- [22] C. Salto, J.M. Molina, and E. Alba. A comparison of different recombination operators for the 2-dimensional strip packing problem. *Proceedings of the XII Congreso Argentino de Ciencias de la Computación (CACIC'06)*, 2006.
- [23] C. Salto, J.M. Molina, and E. Alba. Evolutionary algorithms for the level strip packing problem. *Proceedings of the Workshop on Nature Inspired Cooperative Strategies for Optimization NICSO 2006*, pages 137–148, 2006.
- [24] S. Siegel. *Nonparametric Statistics for the Behavioral Sciences*. New York: McGraw-Hill, 1956.
- [25] P.Y. Wang and C.L. Valenzuela. Data set generation for rectangular placement problems. *EJOR*, 134:378–391, 2001.