

# A Study of Examination Timetabling with Multiobjective Evolutionary Algorithms

*Luís F. Paquete\**

*Carlos M. Fonseca†*

\* Faculdade de Economia, Universidade do Algarve  
Campus de Gambelas, 8000 Faro, Portugal  
Email: lpaquete@ualg.pt

† ADEEC & ISR, Faculdade de Ciências e Tecnologia, Universidade do Algarve  
Campus de Gambelas, 8000 Faro, Portugal  
Email: cmfonsec@ualg.pt

## 1 Introduction

Despite the success of multiobjective evolutionary algorithms (EAs) in many distinct application areas, only recently have they been applied to an Examination Timetabling Problem (ETP) [1] and to a Course Timetabling Problem [2]. In this study, a real ETP is considered, involving both “within-group” and “between-group” edge constraints. A multiobjective evolutionary algorithm [7] based on a direct encoding of the mapping between exams and time slots is used to minimise the number of violations of each type of constraints as separate objectives.

Since the number of time slots to which different exams can be assigned varies wildly in this problem (from 2, in some cases, up to a maximum of 30, in the general case), the mutation operator is designed to take this information into account. In addition, and exploiting the ability to map constraint violations to specific events, the mutation operator is biased further towards such events, in a violation-directed strategy [4].

An experimental study of the performance of several variants of this algorithm, based on attainment function methodology [9] and closed hypothesis testing procedures [12], provides insight into the effects of the multiple objective handling method, mutational bias, and type of mutation operator used. The trade-off between solution quality and run time, inherent to the very notion of optimiser performance, is accommodated by considering time (number of generations) as an additional objective in the attainment function framework, as suggested in [9].

## 2 ETP Formulation

An examination timetable of a typical educational institution can be represented as a mapping  $t : E \rightarrow T$ , where  $E$  is a set of events and  $T$  is a set of time slots. An ETP  $(E, T, f)$  is the problem of finding an optimal mapping  $t_{opt} \in H$ , where  $H$  denotes the set of all possible mappings from  $E$  to  $T$ , such that, given a cost function  $f : H \rightarrow R^+$ ,  $f(t_{opt}) \leq f(t)$ ,  $\forall t \in H$ .

Most EA approaches to timetabling formulate  $f$  as a weighted sum of numbers of constraint violations, where the weights are chosen according to the importance of the corresponding constraints, as described in [4]. In a true multiobjective formulation, objective functions are treated separately, and, thus, no

weight-setting is required. In addition, if a perfect solution does not exist, a number of alternative, compromise solutions is produced.

### 3 A Multiobjective Evolutionary Algorithm for ETP

In a multiobjective EA, as used in this study, individual evaluation is performed separately for each objective. Fitness can then be assigned based on a Pareto-ranking of the population [7]. The choice of representation requires greater care. Although indirect representations, which code for how a timetable is to be generated by a heuristic scheduler, are known to be successful [14], and have grown in popularity, they may not be a suitable choice for a multiobjective EA. The use of heuristic search to deliver timetables which minimise certain types of constraint violations implicitly biases the search away from solutions which establish a compromise between the violation of those and of other constraints.

On the other hand, in a multiobjective EA with a direct representation, where each position in the chromosome corresponds to one examination and can take values according to the number of alternative time slots available for that examination, search bias should primarily arise from the definition of the objectives themselves.

As for the mutation operator, independent mutation of each position of the chromosome is considered. The setting of the mutation probability should be made while taking into account the selective pressure,  $\mu$ , imposed by the selection process [8]. Specifically, the probability of an individual surviving mutation should not be less than  $1/\mu$ . For a binary chromosome of length  $L$ ,  $P_s = (1 - P_m)^L$ , which implies that  $P_m \leq 1 - \mu^{-1/L}$ . This probability of mutation is slightly below the error threshold for infinite populations defined in [13].

These considerations can be extended to suit the general mapping case arising in ETP. Assuming that there is a number  $a_i$  of alternative time slots for each examination  $i$ , setting the probability of mutation  $P_{m_i} = 1 - p^{(a_i-1)}$ , where  $p = P_s^{1/\sum_{i=1}^L (a_i-1)}$ , ensures the given probability of survival while assigning greater probability of mutation to those examinations for which there are more alternative time slots.

Since it is possible to pin-point which examinations are involved in constraint violations, one may set  $P_{m_i} = 1 - q^{(a_i-1)(\beta c_i+1)}$ , where  $c_i$  denotes the number of constraints which examination  $i$  violates,  $\beta$  is a bias constant, and  $q = P_s^{1/\sum_{i=1}^L (a_i-1)(\beta c_i+1)}$ . By increasing  $\beta$ , mutation is directed towards those examinations which need to be rescheduled, in order to accelerate the optimisation process as in [5]. In the limit, as  $\beta \rightarrow \infty$ , only those exams which violate any constraints at all may be mutated. The best value of  $\beta$  is expected to depend on problem difficulty.

A simplification of this mutation operator consists of deciding first whether the chromosome should be mutated with probability  $P_s$ , and then assigning each chromosome position a probability of mutation proportional to  $(a_i-1)(\beta c_i+1)$  and selecting a single position to be mutated by means of roulette-wheel selection. No recombination operator was used in this study.

### 4 The Problem

The problem considered is the assignment of 249 examinations to 30 time slots in the former Unit of Exact and Human Sciences (UCEH) of the University of Algarve. Examinations are grouped according to the curricula of 13 degree programmes, in a total of 53 groups. The typical group is composed of 8 to 10 examinations, and many examinations are common to more than one group. In addition to staff availability restrictions, which are domain constraints, and are incorporated at the encoding level, two types of constraints were formulated:

**Within-group** Examinations in the same group must occur on different days. This was the case for all 53 groups.

**Between-group** Examinations in a group must not occur simultaneously with any examination in certain other groups. This was the case for 41 pairs of groups.

The total number of violations of each type of constraints was considered an objective to be minimised.

## 5 Performance Assessment of Multiobjective Evolutionary Algorithms

Since the EAs are stochastic optimisers, different runs tend to produce different results. Therefore, multiple runs of the same algorithm on a given problem are needed to statistically describe their performance on that problem. Optimiser performance may be understood either as the time taken to produce a solution of a given quality (run-time view [11]) or as the quality of solutions produced within a given time (solution quality view [6]). If the optimiser is stochastic, both of these quantities are random, and studying optimiser performance is reduced to studying the corresponding distributions.

Adopting the solution-quality view of performance, the outcome of a multiobjective evolutionary algorithm run consists of the set of non-dominated objective vectors evaluated during the run. One suitable measure of performance is the probability that an arbitrary goal (vector in objective space) is attained during a single run of the algorithm. This leads to the definition of *attainment function* (see [9]). The attainment function can be estimated directly from data collected from multiple optimisation runs by counting the number of runs which attained each and every goal in objective space, and normalising the results by the total number of runs.

The comparison of the performance of different multiobjective optimisers can be made by testing for the equality of the corresponding attainment functions, in a way analogous to the Kolmogorov-Smirnov test [3], and using permutation arguments [10] to determine the distribution of the test statistic under the null hypothesis [15].

## 6 Experimental Methodology

Three different aspects of the proposed multiobjective EA are considered of interest in this study:

**Multiple objective handling technique** Pareto-ranking is compared against linear-ranking of the sum of the objectives.

**Mutation operator** Independent mutation is compared against single-position mutation.

**Mutation bias** Five levels of the mutation-bias parameter  $\beta$  are considered: 0, 1,  $L$ ,  $2L$  and  $\infty$ , where  $L$  denotes chromosome length.

Each algorithm obtained by each and every combination of the above aspects was run 10 times for 5000 generations, in a complete block design. Fitness was assigned linearly based on ranking with selective pressure  $\mu = 2$ , and generational replacement of the 200 individual population was used. The mutation survival probability was set to  $P_s = 1/\mu = 0.5$ . For each run, four sets of data were collected:

1. The minimum sum of the objective values found until the end of the run

2. The minimum sum of the objective values found up to and including each generation, and the corresponding generation number
3. The set of non-dominated objective vectors accumulated until the end of the run
4. The set of non-dominated objective vectors, augmented by the corresponding generation number, accumulated up to and including each generation

In order to assess the overall effect of each of the above aspects of the algorithms on optimisation performance, several statistical tests were performed. In the first two cases, these were formulated as two-sided, two-sample tests by dividing the blocks into two groups, according to the aspect of interest, and computing the maximum absolute difference between the corresponding attainment functions. The distribution of the test statistic was estimated by randomisation restricted to matching blocks [10].

In the third case, the statistical tests were formulated as two-sided, five-sample tests. Where the null hypothesis was rejected, a closed testing method was used in order to attempt to detect which groups of samples actually differed [12]. Tests analogous to the Birnbaum-Hall test [3] were used in the more-than-two sample case.

## 7 Summary of Results

Regarding the quality of the final solutions, significant differences were detected between the algorithms based on Pareto-ranking and those based on the sum of the objectives, with the sign of the differences observed indicating better performance of the Pareto-ranking approach. On the other hand, no difference was detected between the two mutation operators. As for the values of  $\beta$ , significant performance differences were observed between groups where  $\beta = 0$ ,  $\beta = 1$ , and  $\beta \in \{L, 2L, \infty\}$ .

When considering time as an additional objective, the same differences were observed for the  $\beta$  factor. However, differences were observed between algorithms depending on mutation operator, with the sign of the differences observed indicating better performance of independent mutation over one-position mutation. Finally, performance differences were also observed depending on objective handling technique, with each method performing better in its own sense: Pareto-ranking provided a better covering of the objective space, and linear aggregation of the objectives was more effective in minimising the total number of constraint violations across the runs.

Versions of the evolutionary algorithms described here have been in use at UCEH (and now FCT) since 1999. This study raises interesting issues regarding the use of direct representation in timetabling and other resource allocation problems, as well as regarding the relative merits of different objective handling techniques. It also demonstrates how the performance of multiobjective EAs may be assessed based on attainment function methodology, and highlights the fact that both run time and solution quality may be reconciled in a single view of multiobjective optimiser performance.

## Acknowledgement

This work has been partly funded by the Fundação para a Ciência e a Tecnologia under the PRAXIS XXI programme (Project PRAXIS-P-MAT-10135-1998), with the support of the European Social Fund, FEDER, and the Portuguese State.

## References

- [1] E. Burke, Y. Bykov, S. Petrovic. *Multicriteria Approach in British Universities*. In E. Burke, W. Erben, editors, *Proceedings of the 3rd International Conference on the Practice and Theory of Automated Timetabling (PATAT '00)*, page 180, 2000.
- [2] M. Carrasco, M. Pato. *Automated Timetabling Using a Multiobjective Genetic Algorithm*. In E. Burke, W. Erben, editors, *Proceedings of the 3rd International Conference on the Practice and Theory of Automated Timetabling (PATAT '00)*, page 127, 2000.
- [3] J. Conover. *Practical Nonparametric Statistics*, John Wiley & Sons, New York, 1980.
- [4] D. Corne, P., Ross, H.-L. Fang. *Fast Practical Evolutionary Timetabling*. In T. Fogarty, editor, *Evolutionary Computing*, pages 250–263, 1994.
- [5] C. M. Fonseca, E. M. Mendes, P. J. Fleming, S. A. Billings. *Non-linear Model Term Selection with Genetic Algorithms*. In IEE/IEEE Workshop on Natural Algorithms in Signal Processing, vol. 2, pages 27/1–27/8, 1993.
- [6] C. M. Fonseca, P. J. Fleming. *On the Performance Assessment and Comparison of Stochastic Multiobjective Optimizers*. In H.-M. Voigt, W. Ebeling, I. Rechenberg, H.-P. Schwefel, editors, *Proceedings of the 4th Conference on Parallel Problem Solving from Nature*, Springer Verlag, pages 584–593, 1996.
- [7] C. M. Fonseca, P. J. Fleming. *Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms – Part I: A Unified Formulation*. In *IEEE Transactions on Systems, Man and Cybernetics*, 28 (1), pages 26–37, 1998.
- [8] C. M. Fonseca, P. J. Fleming. *Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms – Part II: Application Example*. In *IEEE Transactions on Systems, Man and Cybernetics*, 28 (1), pages 38–47, 1998.
- [9] V. Grunert da Fonseca, C. M. Fonseca, A. O. Hall. *Inferential Performance Assessment of Stochastic Optimisers and the Attainment Function*. In *Evolutionary Multi-Criterion Optimization, First International Conference*, vol. 1993 of *Lecture Notes in Computer Science*, pages 213–225, Springer-Verlag, 2001.
- [10] P. Good. *Permutation Tests: A Practical Guide to Resampling Methods for Testing Hypotheses*, Springer-Verlag, New York, 2000.
- [11] H. Hoos, T. Stützle. *Evaluating Las Vegas Algorithms – Pitfalls and Remedies*. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, pages 238–245, Morgan Kaufmann, 1998.
- [12] J. Hsu. *Multiple Comparisons - Theory and Methods*, Chapman & Hall/CRC, 1996.
- [13] G. Ochoa, I. Harvey, H. Buxton. *Error Thresholds and their Relation to Optimal Mutation Rates*. In D. Floreano, J.-D. Nicoud, F. Mondada, editors, *Proceedings of the 5th European Conference on Artificial Life (ECAL'99)*, Springer-Verlag, 1999.
- [14] B. Paechter, A. Cumming, H. Luchian, M. Petriuc. *Two Solutions to the General Timetable Problem using Evolutionary Methods*. In *Proceedings of the First IEEE Conference on Evolutionary Computation*, pages 300–305, 1994.
- [15] K. J. Shaw, A. L. Nortcliffe, M. Thompson, C. M. Fonseca, J. Love, P. J. Fleming. *Assessing the Performance of Multiobjective Genetic Algorithms for Optimization of a Batch Process Scheduling Problem*. In *Proceedings of the Congress on Evolutionary Computation (CEC99)*, vol. 1, pages 37–45, 1999.

