Multiobjective Memetic Algorithms applied to University Timetabling Problems

Nuno Leite[1,2], Fernando Melício[1,2] and Agostinho Rosa[2,3]
[1]*ISEL - Lisbon Polytechnic Institute, R. Conselheiro Emídio Navarro, 1, 1959-007 Lisboa, Portugal*
[2]*LaSEEB-System and Robotics Institute, Av. Rovisco Pais 1, TN 6.21, 1049-001 Lisboa, Portugal*
[3]*Department of Bioengineering/IST, TU-Lisbon, Av. Rovisco Pais, 1, 1049-001 Lisboa, Portugal*

1 STAGE OF THE RESEARCH

1.1 Introduction

The present document describes the Ph.D. thesis proposal of Nuno Leite, doctoral candidate in Electrical and Computer Engineering at the Technical University of Lisbon (IST/TU-Lisbon). The thesis is titled "Multiobjective Memetic Algorithms applied to University Timetabling Problems" and is supervised by Drs. Agostinho Rosa (IST/TU-Lisbon) and Fernando Melício (ISEL/Polytechnic Institute of Lisbon). The research work is developed in the LaSEEB (Evolutionary Systems and Biomedical Engineering Lab) at the IST/TU-Lisbon. The five year doctoral program enrolled by Nuno Leite started in December 2010 and ends in December 2015. The curricular part of the doctoral program include the realization of four Ph.D. courses.

1.2 Scope of Research

Nuno's thesis concerns the study and implementation of population based metaheuristics for solving university timetabling problems, e.g. the examination timetabling and post-enrolment course timetabling problems. The study will be focused in approaching these problems as single and multiobjective problems using memetic algorithms.

1.3 Developed Work

Nuno has now completed three of the four Ph.D. courses. These courses are: *Nonlinear Optimization* (completed), *Intelligent Optimization* (completed), *Multiobjective Optimization with Evolutionary Algorithms* (completed), and *Statistical Learning* (not completed yet). Related to the thesis work, Nuno has published one conference paper (Leite et al., 2012). This paper was further selected so that a revised and extended version of this paper will be published by Springer-Verlag in a SCI Series book (Leite et al., 2013b). In May 2013, the authors submitted a position paper to the IJCCI/ECTA 2013 (Leite et al., 2013a), which was accepted for publication as a short paper. In the paper (Leite et al., 2012), the authors present a bi-objective memetic algorithm for solving a real instance of the uncapacitated examination timetabling problem from the ISEL/Polytechnic Institute of Lisbon. The automatic algorithm implemented was able to produce feasible timetables with lower number of conflicts in students's exams compared with the manual solution, and in a shorter time. In the paper (Leite et al., 2013b), the previous work was further extended to solve the capacitated case. Also, rooms are allocated automatically by the algorithm. In the novel work (Leite et al., 2013a), the application of a recently proposed single-objective memetic algorithm (known as SFLA – Shuffled Frog-Leaping Algorithm) to the examination timetabling problem is studied. The proposed adaptation was evaluated on the standard Toronto benchmark datasets with comparable results to the state-of-the-art algorithms. We now intend to study forms of improving the proposed algorithm which include: (i) implement a diversity management procedure in order to control the population diversity, (ii) study the problem landscape characteristics and propose efficient neighbourhood operators to be incorporated in the local search step of the algorithm, and (iii) propose an extension of the algorithm by managing multiple objectives and experiment with the Toronto and the 2nd International Timetabling Competition (ITC2007) benchmark data.

Leite N., Melício F. and Rosa A. (2013).
Multiobjective Memetic Algorithms applied to University Timetabling Problems.
In *Doctoral Consortium*, pages 29-37

2 OUTLINE OF OBJECTIVES

2.1 General Objectives

The research work undertaken has the following objectives:

1. Solve hard multiobjective combinatorial optimization problems by approaching them with Multiobjective Memetic Algorithms. The study will focus on University Timetabling Problems such as the Examination timetabling problem or the Post-enrolment course timetabling problem.
2. Propose a novel multiobjective problem formulation and benchmark data for examination timetabling problem instances found in Portuguese universities.
3. Implement the proposed algorithms in a parallel way using existing frameworks and libraries.
4. Implement an application with a graphical user interface that helps the timetable planner in the multiobjective (many-objective) optimization process.

2.2 Description

The following items describe in more detail the objectives of the research work.

1. Study university timetabling problems in its different dimensions: as scalar (single-objective) and as multiobjective optimization problems. Typically, multiobjective approaches solve intermediate scalar versions of the problem at hands (e.g. multiobjective evolutionary algorithms that use decomposition, MOEA/D (Zhang and Li, 2007)). In particular, in examination timetabling, we will consider the following problems:
   - single-objective problems of Toronto [1] and ITC2007 [2] benchmark datasets;
   - bi-objective problems of Toronto benchmark datasets;
   - many-objective problems of ITC2007 benchmark datasets.
2. Propose scalar and multiobjective memetic algorithms that can efficiently solve the studied problems. Compare the proposed algorithms with current state-of-the-art algorithms.
3. Propose a model and benchmark data for a new multiobjective examination timetabling problem found typically in Portuguese universities. This model extends the ITC2007 model and cover a more realistic model comprising:
   - different interrelated examination sessions that include intermediate evaluation tests (usually two or three tests or a global test in some cases) during the academic semester or year. These are followed by a final examination period, which is formed, normally, by two examination sessions;
   - consider more hard and soft constraints. Some characteristics of the problem are the following. The intermediate evaluation tests during the semester should take place in the days of one of the classes. In this scenario the classes's timetables have to be considered in order to know the days of each class for schedule the evaluation tests. For examinations with few students (e.g. less than 100), the examinations should take place in the rooms of one of the classes. Otherwise, if the examination has a large number of students, then it should be scheduled into examination special rooms. The date of the last intermediary examination test must be far apart from the corresponding first final examination date. Also, all examinations should be spread out for each student, number of used periods minimized, among other constraints.

   This new problem will be also solved by the proposed multiobjective algorithm.
4. Study and design a multiobjective memetic algorithm parallel implementation.
5. Study forms of interacting with the decision maker in many-objective optimization. Design an interactive multiobjective memetic algorithm and an application with a graphical user interface that expose the algorithm functionality.

In the following section, we describe the research problem in detail.

3 RESEARCH PROBLEM

In this thesis we address the class of university timetabling problems, which include problems such as course scheduling and examination timetabling. The problem will be approached in a generic way but we intend to conduct practical experiments on specific problems such as the examination timetabling

[1] University of Toronto benchmark data proposed in (Carter et al., 1996).

[2] 2nd International Timetabling Competition benchmark data (McCollum et al., 2012).

problem (ETTP) found in universities. The university timetabling problems belong to the general class of NP-complete Combinatorial optimization problems (Lewis, 2008). This implies that deterministic exact algorithms, which guarantee finding of optimal solutions, can only be applied to small sized problem instances. Due to this, there's a growing interest in using metaheuristics (Glover and Kochenberger, 2003), (Gendreau and Potvin, 2010) to solve real instances of these hard combinatorial problems.

The university timetabling is a subset of the general family of timetabling and scheduling problems. The aim of these problems is to assign a set of *entities* (e.g. tasks, events, people, vehicles) to limited number of *resources* over time, respecting a set of predefined schedule requirements (Lewis, 2008). In this sense, the university timetabling problem is generically defined as the task of assigning a number of *events*, such as lectures, exams, or meetings, to a limited set of timeslots and rooms, satisfying a set of *constraints*. A categorization of typical constraints can be found in (Lewis, 2008). The constraints are further classified in *hard* and *soft* constraints. The first ones are related to the feasibility of the timetable (e.g. a teacher cannot be scheduled to give two different lectures at the same timeslot in course timetabling, or a student cannot take two different exams at the the same timeslot, in examination timetabling), and cannot be violated. The soft constraints should be satisfied, if possible, as they contribute to the improvement of the timetable quality for the different stakeholders (students, institution, professors/invigilators, etc.), but they could be violated. An example of a soft constraint is the *Period spread* (McCollum et al., 2012) in examination timetabling, with the aim of spreading the exams for all students as much as possible.

Real instances of university timetabling problems encountered in practice, such as the ETTP, have several hard and soft constraints. As in many other problems in scheduling and other areas, these constraints can be seen as multiple objectives, conflicting or not, which means that university timetabling problems are multiobjective optimization problems (MOPs) in nature. Due to the associated extra complexity, in practice these problems have been approached mainly as single-objective problems and very few studies consider multiobjective approaches.

The approaches found in literature to solve the ETTP were first categorised by Carter (Carter, 1986) into four types: sequential methods, cluster methods, constraint-based methods and generalised search. Later, Petrovic and Burke (Petrovic and Burke, 2004) added more categories: hybrid evolutionary algorithms, metaheuristics, multi-criteria approaches, case based reasoning techniques, hyperheuristics and adaptive approaches.

Evolutionary approaches are well suited to tackle multiobjective problems because a population of solutions is already maintained, therefore a set of non-dominated solutions (trade-off solutions) could be determined at the end of one generation of the algorithm. The Multiobjective Evolutionary Algorithms (MOEAs) have been extensively studied in the last 20 years (Van Veldhuizen and Lamont, 2000) (Zhou et al., 2011). It is well known that combining a local search procedure into an evolutionary algorithm greatly improves the algorithm's performance. These algorithms were named Memetic Algorithms (MAs) (Moscato and Norman, 1992) (Neri and Cotta, 2012). Recently, researchers began combining the ideas of MOEAs with memetic evolution, and propose several Multiobjective Memetic Algorithms (MOMAs). These were used recently for solving hard Combinatorial optimization problems with good success (Jaszkiewicz et al., 2012). Hybridizations of evolutionary algorithms with other techniques other than local search metaheuristics, for instance with exact algorithms, seem also to be very promising (Raidl, 2006). In (Ehrgott and Gandibleux, 2008) the authors survey Multiobjective (meta)heuristics for solving multiobjective combinatorial optimization problems, focusing on recent approaches, where metaheuristics are hybridized with exact methods.

In Multiobjective Optimization (MOO), the problems are usually classified in two categories. If the number of objectives is below three or four, the optimization problem is termed simply as MOP; otherwise, if for a number of objectives greater than three or four, the problem is termed Many-objective optimization problem (Zhou et al., 2011).

The early algorithms proposed for Evolutionary Multiobjective Optimization (EMO), such as NSGA-II (Deb et al., 2002), don't work well for the many-objective case (Ishibuchi et al., 2008) (Zhou et al., 2011). As such, many-objective problems pose new challenges for algorithm design, visualisation and implementation.

Multiobjective optimization involves three phases (Coello et al., 2006), (Branke et al., 2008): model building, optimization, and decision making. So, converting a MOP into a single-objective problem puts the decision making task before optimization, and before the existing alternatives are known. The task of a multiobjective algorithm, as shown in Figure 1, is to find a good approximation of the true Pareto front to the decision maker (DM). Three aspects of the solution set quality are usually considered:

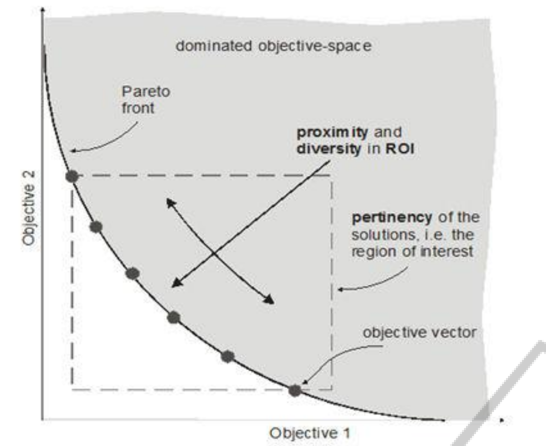

Figure 1: Requirements of a MOEA. (Illustration from Prof. Peter J. Fleming site, http://www.shef.ac.uk/acse/staff/peter_fleming/manyoo).

1. *Proximity*, the produced set should contain solutions close to the true Pareto set.
2. *Diversity*, the solutions in the approximated set should be well distributed.
3. *Pertinency*, the produced set should contain solutions in the DM's region of interest (ROI). The definition of a ROI is especially relevant as the number of objectives increase, because the number of solutions of the approximated set also increases, thus increasing the difficulty of the DM to analyse all the alternatives.

4 STATE OF THE ART

4.1 Single-objective Approaches to the ETTP

The development of systems to automate the construction of university examination timetables has begun in the 1960 decade. The paper (Qu et al., 2009) constitute a survey of the recent (from 1995 to 2008) techniques and algorithmic approaches used to solve this problem. Lewis (Lewis, 2008) surveys metaheuristic-based techniques applied to University Timetabling problems. These two surveys describe techniques applied to Toronto, Nottingham and Melbourne benchmarks. Müller (Müller, 2009) and Gogos (Gogos et al., 2012) propose two efficient scalar optimization techniques to the ITC2007 benchmark data (these authors were the winner and second classified in the ITC2007 competition, respectively). Müller's algorithm is based on a general Constraint solver developed by the author which was adapted to the ETTP. Gogos apply a multi-stage algorithm comprising several metaheuristics and also a step where a Branch & Bound based method was applied.

Researchers recognized that next to the algorithm used to optimize this problem, the way the neighbourhoods are constructed and sampled is important (Demeester et al., 2012) (Abdullah et al., 2010). The Kempe chains, introduced by Morgenstern (Morgenstern, 1989) for solving graph colouring problems, is one of the most successful neighbourhoods. The Kempe chain-based neighbourhood was first applied on the ETTP by Thompson and Dowsland (Thompson and Dowsland, 1998), which used the Simulated Annealing metaheuristic. The Kempe chain heuristic maintains feasibility by swapping infeasible exams between time slots. The simple neighbourhood introduced in Burke and Bykov (Burke and Bykov, 2008) is also frequently used. This heuristic swaps all exams of two randomly chosen time slots, and so it maintains feasibility.

4.2 Multiobjective Approaches to the ETTP

More recently, the ETTP has been approached like a Multiobjective/Multi-criteria Optimization problem, recognizing the true dimensions of real world problems, that typically have many facets to consider (proximity costs between student exams, timetable lengths, room assignment, invigilator availability, etc.). Silva et al. (Silva et al., 2004) present an overview on multiobjective timetabling metaheuristics. Multi-criteria techniques were proposed in (Burke et al., 2001) and (Petrovic and Bykov, 2003). Pais and Amaral (Pais and Amaral, 2009) also approach ETTP (Toronto benchmark) as a multiobjective problem, and consider the four objectives studied in (Petrovic and Bykov, 2003). Other recent works (Côté et al., 2004), (Wong et al., 2004), (Cheong et al., 2009) and (Mumford, 2010), applied MOEAs to solve the ETTP. The current authors propose in (Leite et al., 2012) a MOEA based on NSGA-II and applied it to a real problem instance of a Portuguese university. The crossover and mutation operators were adapted from the ones in (Cheong et al., 2009). A computational analysis involving MOEAs for the ETTP is given in (Paquete and Fonseca, 2001). Burke et al. (Burke et al., 2008) and later (McCollum et al., 2012) present the ITC2007 problem as an MOP. They consider four objectives, each representing a compromise between the various interested parties or stakeholders. These parties are: (1) student interests for a good individual timetable, (2) interests of exam invigilators (and students), (3) the front load (representing the desire of the exam markers to receive the largest exams as soon as possible so as to have more time for marking), and (4) the estate management interests in avoiding the use of some rooms and periods. Due to the complexity of this four objective problem, in Burke et al. (Burke et al., 2008) work the objectives are further grouped in two objectives concerning the students interests and the administration interests. The authors then apply a solver (CPLEX solver) to solve smaller versions of the real datasets. It is mentioned that due to the unoptimized formulation used, the full instances could not be solved.

4.3 Parallel Metaheuristics

Alba et al. (Alba et al., 2013) present recent advances and new trends on Parallel metaheuristics. Talbi (Talbi and Hasle, 2013) present an editorial review of the application of metaheuristics on GPUs. Coello et al. (Coello et al., 2006) also dedicate a chapter to MOEA parallelization. In (Nebro and Durillo, 2010), a MOEA/D parallel implementation was proposed.

4.4 Many-objective Optimization Algorithms

In addition to the previous referred research on multi and many-objective optimization, we include some recent developments on the NSGA-II and the MOEA/D algorithms. In (Deb and Jain, 2012), the authors propose a many-objective NSGA-II (MO-NSGA-II). Recently (Jain and Deb, 2013), the same authors present an improved MO-NSGA-II. In (yan Tan et al., 2013), it is proposed an improved MOEA/D for many-objective optimization.

5 METHODOLOGY

I'm developing a portable library containing single-objective and multi-objective algorithms. The language used is the C++ programming language. In order to reuse some code and to implement the parallel MOEA we are using the ParadisEO library (Cahon et al., 2004).

5.1 Neighbourhood and Variation Operators

I intend to implement some efficient neighbourhoods such as Kempe chains and time slot swaping, mentioned in Subsection 4.1. We will extend the recently developed Shuffled Frog-Leaping Algorithm to include these neighbourhoods. We will analyse and implement other neighbourhoods and variation operators such as the ones proposed in (Abdullah et al., 2010) and (Cheong et al., 2009). The performance of these operators will be studied first on the smaller dataset (Leite et al., 2012), in order to visualize its operation on the timetables produced, and later we will test with the bigger datasets of Toronto and ITC2007 benchmarks.

5.2 Hybrid EMO Methodology

As mentioned in the objectives outline, we intend to tackle the university timetabling problems using EMO algorithms combined with other (meta)heuristics, forming the so called memetic algorithms.

We intend to extend the recently proposed SFLA algorithm and implement a multiobjective SFLA. Some efforts have been done in this path (Rahimi-Vahed and Mirzaei, 2007).

For the local search we intend to explore several of the most successful techniques in the literature such as Hill Climbing, Simulated Annealing, Tabu Search, among others. The acceptance criteria implemented is very important as shown, for example, by the method implemented by (Burke and Bykov, 2008), which achieves the best results on some Toronto datasets.

Other idea will be explored in the final project of the fourth and last doctoral course, yet to be completed. In this project I intend to implement a learning algorithm, called STAGE (Boyan and Moore, 2001), and use it to solve the ETTP. The STAGE algorithm works by learning an evaluation function that predicts the outcome of a local search algorithm, such as hillclimbing, from features of states visited during search. The learned evaluation function is then used to bias future search trajectories toward better optima on the same problem. In (Boyan and Moore, 2001) the authors use the STAGE algorithm to solve several known large-scale optimization domains, such as bin-packing, channel routing, cartogram design or Boolean satisfiability problem.

We intend to study the application of the STAGE algorithm on the ETTP and test it with the benchmarks available. We also want to study other regression strategies such as SVM (*Support Vector Machines*) regression.

5.3 Variable Timetable Length Solutions

One of the typical objectives of timetabling optimization is to find solutions with the lower number of time slots. This was explored in some single-objective (Carter et al., 1996) (Caramia et al., 2008) and multiobjective methods (Côté et al., 2004) (Cheong et al., 2009) for the Toronto benchmark. Manipulating variable timetable length solutions, by creating new time slots needed to accommodate infeasible exams generated by the variation operators, and then trying to repair those solutions, seems another tool for conveniently exploring the search space. Associated to this are the methods proposed in the graph colouring literature, specifically for solving the *Minimum Vertex Colouring* problem. As analysed in the timetabling literature, the ETTP considering just the student clash hard constraint with k time slots is equal to the $k-$vertex colouring problem. So, graph colouring heuristics are also of interest to our analysis. One example of an effective memetic algorithm for tackling the graph colouring problem is the work (Lü and Hao, 2010).

So, we intend to create operators that vary the time slots as needed in order to rearrange the exams. Repairing schemes and specific constraint-handling techniques to handle infeasible solutions will be devised.

5.4 Algorithm Comparison and Performance Measures

For a fair algorithm comparison of the proposed algorithms on benchmark data for the Toronto dataset, we intend to use similar computers to those tested in the literature. For the ITC2007 data, the organizers of the competition provided an application for estimating the CPU time available for running the optimization algorithm (for single-objective/non-parallel algorithm comparison purposes). This application only applies to single CPU machines.

Desirable characteristics of a MOEA include generation of nondominated set close as much as possible to the true Pareto front, and, secondly, that solutions in the obtained nondominated set should be distributed in a diverse and uniform manner. For evaluating these two aspects we will consider the use of performance metrics suggested in the literature (Zitzler et al., 2003) and that are used by researchers in this field to compare the several MOEAs.

5.5 Activity Plan until December 2015

In the next months, I intend to explore thoroughly the concepts mentioned earlier about many-objective optimization. At the same time, we will improve (by exploring the described points mentioned in the earlier sections) and extend the single-objective proposed in (Leite et al., 2013a) to the ITC2007. Next, we will study forms of extending it to a multiobjective framework and implement also a parallel version of the algorithm. The last objective mentioned of creating a graphical interface for the MOEA is optional in this phase of the work.

5.5.1 Goals for Publication

We intend to prepare in the mean time communications and articles to submit to related conferences and journals. We are considering the following international conferences: Evolutionary Computation, Theory and Applications (IJCCI-ECTA), Practice and Theory of Automated Timetabling (PATAT), Multidisciplinary International Scheduling Conference (MISTA). For the journals we are considering the following ones: Elsevier's European Journal of Operational Research (EJOR), Springer's Journal of Scheduling (JSched), Springer's Annals of Operations Research (AOR) and IEEE Transactions on Evolutionary Computation (TEVC).

6 EXPECTED OUTCOME

With the realization of this Ph.D. thesis I expect to give novel contributions to the study of the University Timetabling field. The expected outcome of the thesis is:

- Extend the study of university timetabling problems, exploring them as multiobjective problems, which they are in nature. The majority of the existing studies consider single-objective problems and very few consider two-objective problems. We will extend this framework by considering formulations with three and more objectives.
- Contribute to the understanding of the search space of the analysed timetabling problems.
- Provide efficient neighbourhood operators for the selected problems.
- Provide efficient multiobjective memetic algorithms that improve upon the results of the state-of-the-art algorithms for the multiobjective and single-objective problem instances of the analysed problems.

- The proposed algorithms should be optimized in terms of time performance due to the their parallel implementation. This provides for efficient usage of the multicore machines available in nowadays mainstream market, and, more importantly, provide for new modifications to a given timetable to be done in a much short period of time. The parallelism is also important because multiobjective optimization algorithms are more complex and manipulate more data compared to single-objective optimization algorithms.
- New multiobjective examination timetabling problem formulation and benchmark data from practical problem found in Portuguese universities.
- Publish research results in related refereed conferences and journals.
- User application providing a graphical user interface that helps the decision maker to search interactively given a selected set of objectives, guiding the search process towards a preferred Pareto subset.

ACKNOWLEDGEMENTS

Nuno Leite's work is partially supported by the FCT SFRH/PROTEC/67953/2010 grant. The authors work is also supported by the PEst-OE/EEI/LA0009/2011 grant.

REFERENCES

Abdullah, S., Turabieh, H., McCollum, B., and McMullan, P. (2010). A tabu-based memetic approach for examination timetabling problems. In Yu, J., Greco, S., Lingras, P., Wang, G., and Skowron, A., editors, *RSKT*, volume 6401 of *Lecture Notes in Computer Science*, pages 574–581. Springer.

Alba, E., Luque, G., and Nesmachnow, S. (2013). Parallel metaheuristics: recent advances and new trends. *International Transactions in Operational Research*, 20(1):1–48.

Boyan, J. and Moore, A. W. (2001). Learning evaluation functions to improve optimization by local search. *J. Mach. Learn. Res.*, 1:77–112.

Branke, J., Deb, K., Miettinen, K., and Słowiński, R., editors (2008). *Multiobjective Optimization: Interactive and Evolutionary Approaches*. Lecture Notes in Computer Science. Springer Berlin Heidelberg, Berlin, Heidelberg.

Burke, E., Bykov, Y., and Petrovic, S. (2001). A Multicriteria Approach to Examination Timetabling. In Burke, E. and Erben, W., editors, *Practice and Theory of Automated Timetabling III*, volume 2079 of *Lecture Notes in Computer Science*, pages 118–131. Springer Berlin / Heidelberg.

Burke, E. K. and Bykov, Y. (2008). A Late Acceptance Strategy in Hill-Climbing for Exam Timetabling Problems. In *PATAT '08 Proceedings of the 7th International Conference on the Practice and Theory of Automated Timetabling*.

Burke, E. K., McCollum, B., McMullan, P., and Parkes, A. J. (2008). Multi-objective aspects of the examination timetabling competition track. *Proceedings of PATAT 2008*.

Cahon, S., Melab, N., and Talbi, E.-G. (2004). ParadisEO: A Framework for the Reusable Design of Parallel and Distributed Metaheuristics. *J. Heuristics*, 10(3):357–380.

Caramia, M., Dell'Olmo, P., and Italiano, G. F. (2008). Novel local-search-based approaches to university examination timetabling. *INFORMS Journal on Computing*, 20(1):86–99.

Carter, M., Laporte, G., and Lee, S. Y. (1996). Examination Timetabling: Algorithmic Strategies and Applications. *Journal of the Operational Research Society*, 47(3):373–383.

Carter, M. W. (1986). A survey of practical applications of examination timetabling algorithms. *Oper. Res.*, 34(2):193–202.

Cheong, C., Tan, K., and Veeravalli, B. (2009). A Multi-objective Evolutionary Algorithm for Examination Timetabling. *Journal of Scheduling*, 12:121–146.

Coello, C. A. C., Lamont, G. B., and Veldhuizen, D. A. V. (2006). *Evolutionary Algorithms for Solving Multi-Objective Problems (Genetic and Evolutionary Computation)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.

Côté, P., Wong, T., and Sabourin, R. (2004). Application of a Hybrid Multi-Objective Evolutionary Algorithm to the Uncapacitated Exam Proximity Problem. In *Proceedings of the 5th International Conference on Practice and Theory of Automated Timetabling (PATAT 2004)*, pages 151–167.

Deb, K. and Jain, H. (2012). Handling many-objective problems using an improved NSGA-II procedure. In *Evolutionary Computation (CEC), 2012 IEEE Congress on*, pages 1–8.

Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A Fast and Elitist Multiobjective Genetic Algorithm : NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197.

Demeester, P., Bilgin, B., Causmaecker, P. D., and Berghe, G. V. (2012). A hyperheuristic approach to examination timetabling problems: benchmarks and a new problem from practice. *J. Scheduling*, 15(1):83–103.

Ehrgott, M. and Gandibleux, X. (2008). Hybrid Metaheuristics for Multi-objective Combinatorial Optimization. In Blum, C., Aguilera, M., Roli, A., and Sampels, M., editors, *Hybrid Metaheuristics*, volume 114 of *Studies in Computational Intelligence*, pages 221–259. Springer Berlin / Heidelberg.

Gendreau, M. and Potvin, J.-Y. (2010). *Handbook of Metaheuristics*. Springer Publishing Company, Incorporated, 2nd edition.

Glover, F. W. and Kochenberger, G. A. (2003). *Handbook of Metaheuristics*. Springer.

Gogos, C., Alefragis, P., and Housos, E. (2012). An improved multi-staged algorithmic process for the solution of the examination timetabling problem. *Annals OR*, 194(1):203–221.

Ishibuchi, H., Tsukamoto, N., and Nojima, Y. (2008). Evolutionary many-objective optimization: A short review. In *IEEE Congress on Evolutionary Computation*, pages 2419–2426. IEEE.

Jain, H. and Deb, K. (2013). An improved adaptive approach for elitist nondominated sorting genetic algorithm for many-objective optimization. In Purshouse, R. C., Fleming, P. J., Fonseca, C. M., Greco, S., and Shaw, J., editors, *EMO*, volume 7811 of *Lecture Notes in Computer Science*, pages 307–321. Springer.

Jaszkiewicz, A., Ishibuchi, H., and Zhang, Q. (2012). Multiobjective memetic algorithms. In Neri, F., Cotta, C., and Moscato, P., editors, *Handbook of Memetic Algorithms*, volume 379 of *Studies in Computational Intelligence*, pages 201–217. Springer Berlin Heidelberg.

Leite, N., Melício, F., and Rosa, A. C. (2013a). Solving the Examination Timetabling Problem with the Shuffled Frog-Leaping Algorithm. Accepted on the IJCCI-ECTA 2013 conference as a position paper.

Leite, N., Neves, R. F., Horta, N., Melício, F., and Rosa, A. C. (2012). Solving an Uncapacitated Exam Timetabling Problem Instance using a Hybrid NSGA-II. In Rosa, A. C., Correia, A. D., Madani, K., Filipe, J., and Kacprzyk, J., editors, *IJCCI*, pages 106–115. SciTePress.

Leite, N., Neves, R. F., Horta, N., Melício, F., and Rosa, A. C. (2013b). Solving a Capacitated Exam Timetabling Problem Instance using a Bi-objective NSGA-II. In Madani, K., Correia, A. D., Rosa, A. C., and Filipe, J., editors, *IJCCI (Selected Papers)*, Studies in Computational Intelligence. Springer (in press).

Lewis, R. (2008). A survey of metaheuristic-based techniques for university timetabling problems. *OR Spectrum*, 30(1):167–190.

Lü, Z. and Hao, J.-K. (2010). A memetic algorithm for graph coloring. *European Journal of Operational Research*, 203(1):241 – 250.

McCollum, B., McMullan, P., Parkes, A. J., Burke, E. K., and Qu, R. (2012). A New Model for Automated Examination Timetabling. *Annals of Operations Research*, 194:291–315.

Morgenstern, C. (1989). *Algorithms for General Graph Coloring*. PhD thesis, Department of Computer Science, University of New Mexico, Albuquerque, New Mexico.

Moscato, P. and Norman, M. (1992). A "Memetic" Approach for the Traveling Salesman Problem Implementation of a Computational Ecology for Combinatorial Optimization on Message-Passing Systems. In *Proceedings of the International Conference on Parallel Computing and Transputer Applications*, pages 177–186. IOS Press.

Müller, T. (2009). ITC2007 solver description: a hybrid approach. *Annals of Operations Research*, 172(1):429–446.

Mumford, C. (2010). A Multiobjective Framework for Heavily Constrained Examination Timetabling Problems. *Annals of Operations Research*, 180:3–31.

Nebro, A. J. and Durillo, J. J. (2010). A Study of the Parallelization of the Multi-Objective Metaheuristic MOEA/D. In Blum, C. and Battiti, R., editors, *LION*, volume 6073 of *Lecture Notes in Computer Science*, pages 303–317. Springer.

Neri, F. and Cotta, C. (2012). A primer on memetic algorithms. In Neri, F., Cotta, C., and Moscato, P., editors, *Handbook of Memetic Algorithms*, volume 379 of *Studies in Computational Intelligence*, pages 43–52. Springer Berlin Heidelberg.

Pais, T. C. and Amaral, P. A. (2009). Weight aggregation in a multiobjective approach for exams timetabling problems. Centro de Matemática e Aplicações. Pre-print CMA 4-2009, available on: http://www.cma.fct.unl.pt/sites/www.cma.fct.unl.pt/files/documentos/publicacoes/pdf_2009/CMA%204-2009.pdf.

Paquete, L. F. and Fonseca, C. M. (2001). A study of examination timetabling with multiobjective evolutionary algorithms. *Proceedings of the 4th Metaheuristics International Conference (MIC 2001)*, pages 149–154.

Petrovic, S. and Burke, E. (2004). University timetabling. In *Handbook of Scheduling: Algorithms, Models, and Performance Analysis, chapter 45*. Chapman Hall/CRC Press.

Petrovic, S. and Bykov, Y. (2003). A Multiobjective Optimisation Technique for Exam Timetabling Based on Trajectories. In Burke, E. and Causmaecker, P., editors, *Practice and Theory of Automated Timetabling IV*, volume 2740 of *Lecture Notes in Computer Science*, pages 181–194. Springer Berlin Heidelberg.

Qu, R., Burke, E., McCollum, B., Merlot, L. T. G., and Lee, S. Y. (2009). A Survey of Search Methodologies and Automated System Development for Examination Timetabling. *Journal of Scheduling*, 12:55–89.

Rahimi-Vahed, A. and Mirzaei, A. H. (2007). A hybrid multi-objective shuffled frog-leaping algorithm for a mixed-model assembly line sequencing problem. *Computers & Industrial Engineering*, 53(4):642 – 666.

Raidl, G. (2006). A Unified View on Hybrid Metaheuristics. In *Hybrid Metaheuristics*, pages 1–12.

Silva, J. D. L., Burke, E. K., and Petrovic, S. (2004). An introduction to multiobjective metaheuristics for scheduling and timetabling. In Gandibleux, X., Sevaux, M., Sörensen, K., and T'kindt, V., editors, *Metaheuristics for Multiobjective Optimisation*, volume 535 of *Lecture Notes in Economics and Mathematical Systems*, pages 91–129. Springer Berlin Heidelberg.

Talbi, E.-G. and Hasle, G. (2013). Metaheuristics on gpus. *J. Parallel Distrib. Comput.*, 73(1):1–3.

Thompson, J. M. and Dowsland, K. A. (1998). A robust simulated annealing based examination timetabling system. *Computers & Operations Research*, 25(78):637 – 648.

Van Veldhuizen, D. A. and Lamont, G. B. (2000). Multiobjective evolutionary algorithms: Analyzing the state-of-the-art. *Evol. Comput.*, 8(2):125–147.

Wong, T., Côté, P., and Sabourin, R. (2004). A Hybrid MOEA for the Capacitated Exam Proximity Problem. In *Congress on Evolutionary Computation*, volume 2, pages 1495 – 1501.

yan Tan, Y., chang Jiao, Y., Li, H., and kuan Wang, X. (2013). MOEA/D + uniform design: A new version of MOEA/D for optimization problems with many objectives. *Computers & Operations Research*, 40(6):1648 – 1660.

Zhang, Q. and Li, H. (2007). MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition. *IEEE Trans. Evolutionary Computation*, 11(6):712–731.

Zhou, A., Qu, B.-Y., Li, H., Zhao, S.-Z., Suganthan, P. N., and Zhang, Q. (2011). Multiobjective evolutionary algorithms: A survey of the state of the art. *Swarm and Evolutionary Computation*, 1(1):32 – 49.

Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C. M., and da Fonseca, V. G. (2003). Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Trans. Evolutionary Computation*, 7(2):117–132.