Review Report on the PhD Thesis of Ing. Roman Václavík – "Algorithms for Personnel Scheduling Enhanced by Machine Learning Techniques" by Broos Maenhout (01-03-2019)

1. General assessment

In the operations research community, research is characterised by the drive for more efficient and effective solution algorithms for a wide variety of problems and problem variants. This is exemplified in many domains by the use of a benchmark dataset for a specific problem aiming researchers to push their boundaries and to come up with the best performing algorithms. In that perspective, the student has been able to perform some ground-breaking research as significant acceleration mechanisms were developed leading to (approximately) the same high-quality solutions or even outperforming the state-of-the-art literature. Despite the research efforts of the OR community, academic state-of-the-art solution methodologies have rarely combined machine learning techniques in algorithms for combinatorial optimisation. This feature is truly accomplished in the underlying PhD research. The student is one of the first to combine combinatorial optimisation in the domain of employee timetabling with machine learning techniques to increase the efficiency of the already known and proven solution techniques exploiting intermediate data.

The proposed methodologies are adequate for researchers in the domain of employee timetabling since the optimisation techniques used are widespread (local-search based meta-heuristics, branch-and-price, single-pass scheduling heuristics) and common to solve personnel rostering problems. The machine learning techniques are quite recent and already established in machine learning. Both these items lead to the fact that researchers can easily learn from the insights provided by this PhD thesis. The methodologies have been applied to a variety of problems and problem variants, denoting the robust performance, and significantly reduced the computational effort to come up with high-quality solutions, which confirms that the PhD student has met the objectives of the dissertation. In addition, the proposed techniques have been proven useful not only for artificial instances, generated for academic purposes, but also for real-world problem instances. Real-world problems are often too complex in terms of problem size and/or problem characteristics, which may lead to unreasonable long running times. The thesis shows how running times can be reduced significantly and in this way the conducted research helps to close the gap between academic research and providing solutions for complex real-life problems in a suitable time span.

2. Detailed comments related to the different chapters

Chapter 1

The introductory chapter provides a sufficient introduction to the topic positioning the conducted research well in the academic literature. However, my comments/questions are as follows:

- The motivation of the research to devise, i.e. small CPU time and robust performance, could be better embedded in the literature. Not a single reference has been used to state the importance of both the main motivations of the thesis. Are these items also recognized by other researchers?
- The nurse rostering problem could be better introduced in this introductory chapter. The main context of the nurse rostering problem is missing, i.e.
    - Which decisions are assumed to be taken (tour scheduling problem, staffing problem, etc)? Do you tackle multiple personnel scheduling problems in the PhD thesis?
    - Only De Causmaecker and Vanden Berghe (2011) is indicated as a reference to nurse rostering. Why? Other references are missing, e.g. the literature overview of Van den Berghe et al. (2013).
    - The motivation to select the nurse rostering problem as the principal problem should be better documented. Many standard combinatorial optimization problems involve highly complex benchmark datasets. Comment on the characteristics of the different instances of the dataset to validate your approach is indeed well designed leading to a robust performance.
- The model-based methodology and the data-driven approach can be better explained. Describe their main perspective. Note that also the data-driven approach constructs a model based on the obtained data. Identify the synergies between both types of approaches based upon the literature.

Chapter 2

Chapter 2 of the thesis provides a contribution to the academic community as the use of classifiers based on neural networks have been proven to reduce the computational effort significantly for the nurse rostering problem at the expense of a little decrease of the solution quality. In heuristic procedures, most of the computational effort is devoted to evaluating nurse rosters. For complex real-world problems, in particular, this abundant use of CPU time hinders the application of sophisticated solution approaches. The research in this chapter comes towards this lacuna in academic research by integrating two well-known methodologies, i.e. meta-heuristic optimisation and machine learning techniques for the well-known nurse rostering problem. However, I have to admit that I have some doubts related to the generality of the proposed methodology as a result of the employed meta-heuristic, the visited neighbourhoods and the employed classifiers, i.e.

- A change to a nurse roster is performed by simple local neighborhoods such as the swapping of assignments or the single-shift neighborhood. However, it is not clear what the impact is of these neighborhoods on the efficiency of the proposed evaluation methods. In the literature, there are much larger neighborhoods defined, e.g. the day-based local search of Maenhout and Vanhoucke (2007), or even approaches have been developed using mathematical programming, which may change the structure of

a solution completely (e.g. local branching). Other approaches (e.g. population-based meta-heuristics) are far more disruptive and may generate completely new solution points. In your approach, the used meta-heuristics are all local-search based heuristics, which start from a current solution point and only adapt this solution point in a very small extent. Given these insights, I want you to elaborate on the usefulness of the proposed evaluation methods linked to the optimization method used.

- Adding to the previous question, you should elaborate on the generality of the applied classifiers. The model (p671) does not allow the calculation of performance measures related to a set of nurses (e.g. in many studies in the literature, the coverage requirements are set soft). In addition, is a distinction possible between different types of classifiers based on the type of constraints (e.g. sequence versus counter constraints)?
- The conclusion claims that the proposed methodology can also be applied to other timetabling problems. However, this is not clear as no computational proof has been given.

Other comments concerning this chapter:

- Tabu search: There is claimed that tabu search is one of the most successful meta-heuristics for the nurse rostering problem but no adequate reference is given (p673). In In addition, the pseudocode could be better explained
    - How is employee b selected?
    - Discuss lines 30-36.
- How are thresholds set to distinguish between a very good and a good solution (or a very bad and a bad solution)?
- Why is it better to use relative values compared to absolute values? Is this inherent to the optimization technique (i.e. a local-search based meta-heuristic)?
- Neural networks is a machine learning technique that will come up with a model to evaluate the solution points based upon the generated data. Can you discuss the relationship of the constructed model with the objective function value?
- Comparison is mainly done based on CPU time but not on solution quality, which is worse given the stop criterion used. Given a specific stop criterion in terms of CPU time, the proposed methodology allows to evaluate more schedules. Does the evaluation of a higher number of schedules lead to a better solution quality, which is the eventual goal of optimization?

Chapter 3

In my opinion, the paper in Chapter 3 is the strongest research in this PhD, making a significant contribution to the literature. One of the disadvantages of exact approaches such as branch-and-price is the required computational time to solve the pricing problem each time to optimality. In response, many authors have proposed acceleration mechanisms to speed up the performance of branch-and-price methods such as heuristics, reduced networks for the pricing problem, etc. However, the design of a learning method based on intermediate data has not been proposed in the literature so far. Using this approach, the developed algorithm reduces the CPU time enormously, whereas the optimality of the solution is still guaranteed. The results have proven to outperform the state-of-the-art literature and the computational results show very clearly the contribution of each design choice by the authors. Given the

increased power of mathematical programming, branch-and-price techniques will only gain in popularity and this paper will certainly inspire many other authors. My questions, related to this chapter are as follows:

- Regression analysis comes up with a model, that approximately explains the behavior of the system. This model is formalized using mathematical expressions. Can you elaborate on the link between the model constructed by the regression and the pricing problem? Can you explain/give insight why the learning model is very performant when the dual prices are aggregated week by week? What if the problem structure was not defined based on weekly patterns? It is a little bit surprising why so little information is used from the pricing problem structure in the regression model and the data can simply be aggregated on (high) level of a week. Discuss this question in relationship to the structure of the pricing problems of the tackled nurse rostering instances.
- In my opinion, for nicely structured pricing problems with clear features and a clear calculation of the objective (e.g. day/shift preference cost) the technique may be very performant. However, in the literature there is also referred to irregular structures of the pricing problem, which refer to nurse schedules for which the objective is not simply the sum of the cost of different features (e.g. shift assignments) but the cost depends on the overall structure of the line-of-work (e.g. Aickelin and Dowsland, 2000). Hence, the objective is the result of a specific sequence of duties or a line-of-work and learning will become very difficult. Hence, I want you to elaborate on the applicability of your approach or at least indicate in the conclusions the limitations of your research. In other words, discuss the generality of your approach? How structured should the pricing problem be? You indicated that if the pricing problem is nicely structured, learning may not be necessary. However, if the problem is not structured at all, learning may not be possible. What are the boundaries of your approach?
- Some tables only present results for a limited number of instances (Table 5, 8, 9). Motivate the choice of selection here based on the problem characteristics of the instances. Are the results similar for other instances too?

Chapter 4

From a practical point-of-view, Chapter 4 assesses most probably the largest benefits of combining machine learning techniques with decision-making techniques in the perspective of online scheduling, for which immediate decisions are required. Learning methods improve the solution quality significantly compared to simple heuristic rules. The chapter provides a detailed comparison of different optimisation and learning strategies and the conclusions are well-grounded. In addition, the trade-offs between the two different objective function criteria is mapped in an accurate and scientific way. Plaudits are in place for the very well worked out test design and structure of the computational experiments in this chapter.

- Which characteristics are included in the feature vector $x^{(i)}$? Are there also instance-related characteristics included (such as number of constraints, type of constraints, etc.), which may be the key to develop a well-performing hyper-heuristic, which may be more suitable to tackle different kinds of real-life rostering problems (cf. Smith-Miles and Lopes (2012)).
- What is the impact of the number of segments on the results? Is it really necessary to consider 10 different segments? For your information: Multi-mode research (still under review) has indicated that a limited well-chosen subset of available modes is sufficient for finding an appropriate solution quality.

- The proposed results are very depending upon the used procedure to solve the scheduling problem and the objective function structure of the problem under study. Elaborate how this impacts your results. For most optimisation procedures, we can observe a quick convergence to a solution of satisfactory quality (presumably a local optimum) such that the number of different segments may be too high. For finding better solutions, significantly longer CPU times are required, which may be prohibitive for an online scheduling tool.
- How would Figure 15 look like if you would construct a stochastic Pareto front, i.e. a Pareto front based on confidence intervals and not on the average performance?

3. Summary

It is clear that the student has done a large amount of high-quality work. The PhD thesis comes forward to different gaps in the research domain of employee timetabling. The candidate has demonstrated creativity and originality to master the challenging problems by combining optimisation methods and machine learning techniques. He has proven to come up with innovative ideas to improve the computational performance of established algorithms. The ideas have been proven to be suitable for different problems and problem variants, showing the robustness of the developed methodology. Overall the candidate has made a number of significant contributions to the domain of employee timetabling and these will facilitate the further development of new research directions in the domain of employee timetabling.

The author of the thesis proved to have an ability to perform research and to achieve scientific results. I do recommend the thesis for presentation with the aim of receiving the Degree of PhD.

Yours sincerely,

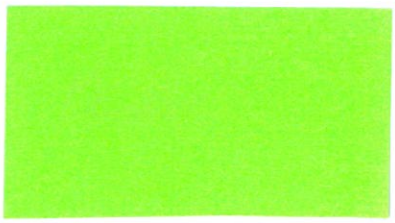

Prof dr Broos Maenhout
Associate Professor at the Ghent University
Department of Business Informatics and Operations Management
Tweekerkenstraat 2, 9000 Gent, Belgium
Tel: +32 9 264 98 32
Email: Broos.Maenhout@Ugent.be