Available online at www.sciencedirect.com
ScienceDirect
European Journal of Operational Research 188 (2008) 330–341
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/ejor

Discrete Optimization

# A hybrid heuristic ordering and variable neighbourhood search for the nurse rostering problem

Edmund K. Burke [a,*], Timothy Curtois [a], Gerhard Post [b], Rong Qu [a], Bart Veltman [b]

[a] School of Computer Science and Information Technology, University of Nottingham, Jubilee Campus, Wollaton Road, Nottingham NG8 1BB, UK
[b] ORTEC, Groningenweg 6-33, 2803 PV Gouda, Holland, The Netherlands

Received 24 January 2006; accepted 24 April 2007
Available online 29 April 2007

**Abstract**

This paper is concerned with the development of intelligent decision support methodologies for nurse rostering problems in large modern hospital environments. We present an approach which hybridises heuristic ordering with variable neighbourhood search. We show that the search can be extended and the solution quality can be significantly improved by the careful combination and repeated use of heuristic ordering, variable neighbourhood search and back-tracking. The amount of computational time that is allowed plays a significant role and we analyse and discuss this. The algorithms are evaluated against a commercial Genetic Algorithm on commercial data. We demonstrate that this methodology can significantly outperform the commercial algorithm. This paper is one of the few in the scientific nurse rostering literature which deal with commercial data and which compare against a commercially implemented algorithm.


*Keywords:* Variable neighbourhood search; Heuristics and metaheuristics; Nurse rostering; Hybrid methods

## 1. Introduction

It is clear that the efficient rostering of healthcare personnel can lead to the more effective utilisation of valuable resources. Healthcare institutions around the world are becoming increasingly interested in the deployment of decision support technology to aid the personnel scheduling process. A very general form of the nurse rostering problem could be described as follows: Given a set of shifts and a set of nurses over a certain time period, assign each shift to a nurse subject to a set of constraints. The constraints are usually defined by regulations, working practices and the preferences of the nurses.

The problem of nurse rostering is relatively easily described but like most real world search problems it is far from easy to automatically generate very high quality solutions. Indeed, there have been many papers over the years from across operational research and artificial intelligence that have tackled the problem in one form or another. A wide range of approaches and techniques have been investigated and used. Ernst et al. [19,20] identified 28 different

* Corresponding author.
*E-mail address:* ekb@cs.nott.ac.uk (E.K. Burke).

0377-2217/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2007.04.030

categories of methods that have been used on personnel scheduling problems. These include constraint logic programming, constructive heuristics, expert systems, genetic algorithms, integer programming, set partitioning, simple local search and simulated annealing. A recent review of automated nurse rostering approaches found that, although there has been a lot of research in the area, surprisingly few of the methods were tested on real world data [17]. The paper went on to conclude that even fewer have actually been implemented in real hospital wards.

Of those techniques that have been applied on real-world problems metaheuristic methods seem to dominate. One approach which has been applied in multiple real world hospitals is a hybrid tabu search [15]. The tabu search is integrated with techniques which are usually observed in manual scheduling approaches. The algorithm has been incorporated into software that has been used to create nurse rosters in over 40 Belgian hospitals and copes with many shift types, work regulations and skill categories. This work was hybridised with an evolutionary approach [12] to produce a methodology which could generate higher quality solutions but at the cost of increased computational time. Variable Neighbourhood Search [21,28] has also been applied and tested on highly constrained real world nurse rostering data [14]. The authors found that VNS could be effectively used to escape from the local optima found using single neighbourhood heuristics. They also commented ''After reaching a local optimum, we recommend the exploration of wider environments''. Evaluation methods for challenging real world problems are presented and discussed in [13]. A methodology which can handle a more flexible approach to real world nurse rostering than the traditional fixed period based approach is presented in [11]. An overview of the work carried out by these authors in Belgian hospitals is presented in [16].

Another investigation on real data explored a genetic algorithm approach [1], which successfully exploits problem specific knowledge in tackling the problem. Although the method is tailored for that particular problem instance, the underlying concepts could be applied to other nurse rostering problems. In 1998, Dowsland was also able to match the quality of schedules produced by an expert human scheduler using a highly developed tabu search [18]. The algorithm 'oscillates' between trying to improve the cover and improving the preference costs. As well as using tabu lists, candidate lists and diversification strategies, the search also uses a large neighbourhood created by looking for chains of overall improving swaps. Aickelin and Li [2,3] have since experimented with the application of bayesian optimisation and classifier systems to similar nurse rostering problems. The results are close to those produced by an optimal integer programming method and the authors concluded that with further effort and experimentation the algorithms could well improve even more. Bellanti et al. [9] tackled a problem with hard constraints and objectives (or soft constraints) using various local search techniques. The authors presented good results for a tabu search and iterated local search which use neighbourhoods defined by changing the assignment of night shifts.

Another methodology that has been tested on complex real-world data from a UK hospital is case-based reasoning [30]. This approach avoids the use of evaluation functions but instead aims to imitate how an expert human scheduler would produce a good schedule. This is done by storing observed methods for repairing violations in schedules and retrieving, adapting and performing these repairs or moves whenever a similar violation is encountered again. As an extension to their work, the authors also suggest methods in which it could be combined with a meta-heuristic approach [7,8]. Another relatively recent methodology is a combination of constraint networks and knowledge-based rules [24]. The approach was implemented in a commercial software package and has been successfully used in a number of hospitals.

Berrada et al. [10] developed a multi-objective mathematical programming model to represent a real world problem containing both hard and soft constraints in a Canadian hospital. The schedules produced met the standards required by the head nurses. The authors also experimented with a tabu search and found that although it required greater computational time it was useful in some circumstances. Valouxis and Housos [32] approach a nurse rostering problem using an approximate integer linear programming model to produce initial solutions. The initial solutions are then further optimised using a local search with a '2-opt' neighbourhood and a tabu search. Their method compared very favourably with a constrained programming approach. In 2004, Bard and Purnomo [5] employed a combination of heuristic and integer programming methods to solve nurse preference scheduling problems with up to one hundred nurses and approximately 13 hard and soft constraints. Individ-