European Journal of Operational Research 202 (2010) 379–389

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Discrete Optimization

The nurse rostering problem: A critical appraisal of the problem structure

Celia A. Glass *, Roger A. Knight

Cass Business School, City University London, 106 Bunhill Row, London EC1Y 8TZ, United Kingdom

ARTICLE INFO

Article history:
Received 5 September 2008
Accepted 29 May 2009
Available online 24 June 2009

Keywords:
Integer programming
Nurse rostering

ABSTRACT

This paper is concerned with the problem of nurse rostering within hospitals. We analyse a class of four benchmark instances from the nurse rostering literature to provide insight into the nature of the problem. By highlighting the structure of the problem we are able to reduce the relevant solution space. A mixed integer linear programme is then able to find optimal solutions to all four instances of this class of benchmark problems, each within half an hour. Our second contribution is to extend current mathematical approaches to nurse rostering to take better account of the practical considerations. We provide a methodology for handling rostering constraints and preferences arising from the continuity from one scheduling period to the next.

1. Introduction

The nurse rostering problem is a complex combinatorial optimisation problem and is known to be NP-hard. The problem is to decide which members of a team of nurses should be on duty at any time, during a rostering period of, typically, one month. This involves determining the days on and off duty for each nurse, as well as their shift start and finish times on each working day. The underlying aim is to ensure that staffing levels are sufficient to cover the nursing requirement, while individual nurses duties are within the terms of their employment contract. These competing goals are encapsulated in two sets of constraints: *staffing* constraints ensure that sufficient nurses of each type are on duty at any particular time, while *schedule* constraints relate to the sequences and combinations of shifts to be worked by each nurse. It is not always possible to satisfy both sets of constraints simultaneously. The modelling approach therefore involves reducing selected constraints to soft constraints with measurements of their violation. The objective is then to minimise some trade-off of the measures of violation of these soft constraints. An attractive aspect of this approach is the introduction of measures of roster quality beyond contractual obligations.

The quality of a roster can be judged in terms of compliance with legal requirements (for example, the European Working Time Regulations[1]), employment contracts, and the collective and individual preferences of the nurses themselves. There are a wide range of considerations which constitute rostering "good practice". The UK Government's Health and Safety Executive (HSE) published a number of guidelines in HSE (2006), and a set of 10 ergonomic principles was presented by Kundi (2003). Poorly designed rosters can lead to inadequate periods of rest and recovery time, and to socially unacceptable patterns of working. Such outcomes are stressful to the nurses, and the negative impacts upon their productivity and levels of engagement are well documented in the literature on the ergonomics of shift working (see, for example, Totterdell (2005)). The complexity of the task has led to the use of automated rostering systems in some European hospitals, although in the UK the task is usually carried out manually. For these reasons, there is a need to develop nurse rostering methodologies which have the potential for practical application.

There is a large body of research literature relating to nurse rostering. The survey papers by Ernst et al. (2004a,b) and Burke et al. (2004) give a detailed overview of the literature up to 2005. However, the latter concludes that only a few papers are based on real world data, or address the development of rostering systems for implementation in hospitals. A recent paper, Burke et al. (2008), was jointly authored by researchers at Nottingham University and at ORTEC,[2] a major supplier of rostering software. This paper presents a practical, heuristic-based methodology and compares results from variable neighbourhood local search heuristics with those obtained using ORTEC's Harmony software on a range of problem instances. Since the context is the same of the recent article by Burke et al. in this journal, we omit a detailed literature review in order to avoid unnecessary repetition.

* Corresponding author. Tel.: +44 20 7040 8959.
E-mail address: C.A.Glass@city.ac.uk (C.A. Glass).

[1] European Working Time Directive (93/104/EC).

[2] ORTEC, Groningenweg 6-33, 2803 PV Gouda, Holland, The Netherlands.

0377-2217/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2009.05.046

Our main research interest is in employee rostering for call centres. In this domain, the application of Mixed Integer Programming (MIP) techniques is more common than the use of heuristics. The traditional approach to call centre rostering is to decompose the problem into two stages. First, a set of weekly *tours* are scheduled, where each tour consists of a set of attendance details; days on and off duty, daily shift start and finish times, and even lunch and coffee break timings. At a second stage, each tour is assigned to an individual employee, taking account of availability and preferences. The problem is complicated by the large number of feasible shifts, which may start or finish in any reasonable quarter-hour interval, unlike in nurse rostering where generally only a handful of shifts are considered. The most comprehensive tour scheduling model in the call centre literature is that of Brusco and Jacobs (2000). This model allows the specification of a *bandwidth* parameter which limits the variation in shift start-times within a working week. A major problem is that this approach does not distinguish between forward and backward rotation. For example, a typical restriction in practice is that a day shift cannot immediately follow a night shift, as indicated by constraint H5 in Appendix A. If this condition is imposed through a bandwidth restriction, then the same constraint prevents a night shift from appearing *anywhere* in the same tour as a day shift. Hence, bandwidth has little relevance in the context of nurse rostering, where specific rules relating to night shift working (for example, constraints H3, H5, H10, S10) are typically imposed. In the rostering literature relating to healthcare, Isken (2004) developed a tour scheduling model which treated start time variations in the same way as Brusco and Jacobs (2000), albeit with a different formulation. The Isken model therefore suffers from the same drawback.

We have developed an alternative tour scheduling model, presented in our paper Glass and Knight (2008). We recognise that while it is preferable to schedule the same shift start time (or later) on consecutive working days, there is usually no objection to any shift change following a period of one or two days off-duty.

An important issue in both call centre and nurse rostering is the continuity from one rostering period to the next. The nurse rostering benchmark instances are designed only to produce rosters for an isolated period, applying penalties in accordance with the convention that all potential violations are counted at the beginning of the period, and ignored at the end. We recognise that the benchmark instances are intended as a basis for comparison between alternative rostering methodologies, and that the consideration of an isolated rostering period serves this purpose. However, in a practical environment, information relating to one rostering period is carried forward to the next, creating additional considerations of "continuity".

Our contribution to the mathematical nurse rostering literature is as follows:

(1) analysis of a benchmark problem which identifies a lower bound on the solution, and which brings to light aspects of the problem structure which we are later able to exploit in order to radically reduce the size of solution space;
(2) improved results for four related benchmark problem instances, by providing an optimal solution to each instance, within a practical execution time, using an MIP approach;
(3) improvements in handling shift changes within an MIP, extended from the context of call centre rostering to nurse rostering; and
(4) a methodology for handling continuity between rostering periods.

The paper is organised in two parts. In the first part we analyse a class of four benchmark problem instances initially specified by ORTEC, each of which involves a single rostering period. We give a structural analysis of the interrelationship between constraints leading to a lower bound on penalty costs for one of the problems. We also provide optimal solutions for the four problems and compare our results to those obtained previously.

In the second part of this paper we propose a more flexible approach, handling those requirements which relate to the continuity between rostering periods. We are aware that these issues are well understood and taken account of in practice. Indeed, the benchmark problems originate from practice. Our contribution is to formalise the continuity goals. We illustrate how our approach achieves these objectives by producing a solution which continues from one month to the next.

2. Benchmark problems

In this section, we address a class of four benchmark problems originally provided by ORTEC. The first two problems are referred to as "GPost" and "GPost-B", and the second two as "ORTEC01" and "ORTEC02". These problems are among a number of nurse rostering benchmark problems which are maintained and made available by Nottingham University.[3]

"GPost" is a small introductory problem for eight nurses across a rostering period of exactly 4 weeks. There are only two nurse contracts, full and part time, and two shift types, day and night. However, this problem contains many of the structural elements of the larger, more complex nurse rostering examples. "GPost-B" is a relaxation of this problem with some period-to-period continuity constraints removed.

"ORTEC01" is a larger problem for 16 nurses, for the 31 days of January 2003. There are four shifts: Early, Day, Late and Night (E, D, L and N). "ORTEC02" is a variant on "ORTEC01", with an added constraint relating to a period of annual leave. The problem is fully described in Burke et al. (2008), except for a couple of individual rostering constraints which are not explicitly listed in the academic papers but are embedded in the solutions presented on the Nottingham University website. The full set of rostering rules for both problems are listed in Appendix A for completeness.

Our general approach to these problems is to initially treat those soft constraints with high penalties as being hard constraints. Should this lead to infeasibility, we relax the constraints incrementally in order of their cost, until a feasible solution is found. This approach is fairly standard and aligns that of Burke et al. (2008), where a rule relating to the number of night shifts to be worked in a 5 week period is listed both as a hard constraint, and also as a soft constraint with a high penalty. All our experiments were performed using a desktop PC with a P4 2.67 GHz processor.

We observe that, in all four instances, the staffing requirement is fairly consistent on a daily basis, and throughout the rostering period. This imposes a high degree of structure on the problem which we can exploit. We start by analysing the problem structure before addressing each of the instances individually. The key to the constraint labels is given in Appendix A.

[3] http://www.cs.nott.ac.uk/~tec/NRP/.