Investigation of fairness measures for nurse rostering

Pieter Smet · Simon Martin · Djamila Ouelhadj · Ender Özcan · Greet Vanden Berghe

Received: date / Accepted: date

Abstract The nurse rostering problem presents a combinatorial optimisation problem in which shifts must be assigned to nurses who are subject to a large number of workforce related constraints. In the literature on nurse rostering, the fairness of the constructed rosters has often been neglected. Solutions are typically evaluated by means of a weighted sum objective function which does not explicitly account for the fair distribution of individual high quality rosters. The present contribution aims at incorporating fairness measures in existing solution methods for the nurse rostering problem. Preliminary experimental results show that fairer solutions are obtained when applying the new fairness measures.

Keywords Nurse rostering · Fairness · Evaluation

1 Introduction

Several studies have shown that high quality work rosters contribute to the job satisfaction of nurses. This is an important result, because by increasing the job satisfaction of nurses, their retention rate is also likely to increase. Factors that influence the quality of a roster include the working hours and whether

P. Smet and G. Vanden Berghe
CODeS, KAHO St.-Lieven,
Gebr. De Smetstraat 1, 9000 Gent, Belgium
E-mail: {Pieter.Smet, Greet.VandenBerghe}@kahosl.be

S. Martin and D. Ouelhadj
Logistics and Mathematics Management Group,
University of Portsmouth, Department of Mathematics, UK
E-mail: {Simon.Martin, Djamila.Ouelhadj}@port.ac.uk

E. Özcan
Automated Scheduling, Optimisation and Planning Research Group,
University of Nottingham, Department of Computer Science, UK
E-mail: Ender.Ozcan@nottingham.ac.uk

or not nurses can work a minimum number of consecutive days. Furthermore, we assume that a fair distribution of work also attributes to a higher rate of satisfaction with regard to the roster. We present a novel approach that tries to guarantee a fair distribution of individual rosters among nurses in automatically generated solutions.

In the literature, several automated approaches have been proposed as decision support systems for nurse rostering. [2] describe an auction based self rostering system in which nurses have a number of points they can use to bid for shifts or days off. After the bidding phase, an algorithm determines the winner of each auction and tries to find a feasible schedule which includes the winning bids. However, fairness is not guaranteed since experienced nurses can easily misuse the auction system such that they always make the winning bids. [5] present an automated preference scheduling approach in which nurses can request particular shifts or days off. Linked to each request is a grade indicating its importance. All requests are passed on to an algorithm that produces a feasible schedule, while respecting the preferences as much as possible. The system explicitly attends to fairness by means of an extra variable in the objective function. This variable is used to maximise the number of requests of the least favoured nurse, thus taking into account the worst individual roster. Other automated approaches often model fairness as balance constraints on working time [1].

Nurse rostering problems can be represented as constraint optimisation problems using 5-tuples $\langle N, D, S, K, C\rangle$ with N the set of nurses, D the set of days in the planning period and all relevant days in the previous and next planning periods, S the set of shift types, K the set of skill types and C the set of constraints. Solutions are typically evaluated using a weighted sum of soft constraint violations (Equation 1). This objective function has the advantage that it is both easy to understand and to implement. However, algorithms optimising WO can generate unfair solutions in which bad individual rosters are compensated by other high quality rosters. These solutions will have a high overall quality but also an unfair distribution of individual rosters. We propose an alternative for WO which better guarantees fair solutions and can be easily added to models for nurse rostering.

$$WO = \sum_{\forall n \in N} \sum_{\forall c \in C} \#\text{violations}_{n,c} \times \text{weight}_c \tag{1}$$

2 Fairness objective

Inspired by the approach of [5], we propose the objective function FO (Equation 2), whereby the quality of the worst individual roster determines the overall solution quality. In doing so, nurses' rosters will not be improved at the expense of the worst individual roster.

$$FO = max_{\forall n \in N} \sum_{\forall c \in C} \#\text{violations}_{n,c} \times \text{weight}_c \tag{2}$$

When optimising FO, the nurse rostering problem becomes a min-max problem. For algorithms that performed well when optimising WO, it will be more difficult to find good solutions. Algorithms must deal with a trade-off between a fair distribution of individual rosters and the overall solution quality. In the next section we present preliminary experimental results showing the effect of FO on the fairness and overall quality of solutions.

3 Experimental results

We define the individual roster quality of each nurse n q_n, the average individual roster quality μ and the standard deviation σ. The relative standard deviation RSD is defined as the ratio of σ on μ, and allows for comparison of fairness for different scenarios. Low values for RSD indicate fair solutions, in which little variance exists between individual roster quality. The relative quality gap between the best and the worst individual roster is defined as $diff$. A relatively small difference between the best and worst individual roster is indicated by low values for $diff$. The overall solution quality calculated with WO is defined as $q^{WO}_{solution}$. This metric allows us to compare the overall quality of solutions obtained with different objective functions.

Experiments are performed on real world data collected from four different wards in two Belgian hospitals. For each ward, two scenarios are considered: one in which each nurse has the same contract and one in which each nurse has a personalised contract. The hyper-heuristic approach presented in [6] is used to find solutions. Each run is repeated ten times, with computation time limited to ten minutes.

Table 1 shows that fairer solutions are obtained when optimising FO. For all instances, solutions obtained with FO have lower values for RSD than those obtained with WO. This means that the quality of individual rosters varies less when using FO, thus producing fairer solutions. The reported values of $diff$ support this result. The difference between the worst and best individual roster is smaller when optimising FO. From the results in Table 1 it can also be concluded that, for both objectives, it is easier to find fairer solutions when all nurses have identical contracts. Looking at the overall solution quality $q^{WO}_{solution}$, there exists no clearly identifiable trend. In some cases it is clear that the new objective makes it harder for the hyper-heuristics to find good solutions. However, the opposite is also true for some instances. These inconsistent results require further investigation into the structure of each problem instance (e.g. which constraints are present).

4 Discussion

Preliminary experimental results show that fairer solutions are obtained when optimising FO than when optimising WO. Looking at the overall solution quality, a trade-off exists between the fairness of a solution and its overall

Instance	RSD		$diff$		$q^{WO}_{solution}$		
	WO	FO	WO	FO	WO	FO	gap
Emergency-i	19,24%	7,40%	54,00%	24,11%	192569	169096	-12,19%
Emergency-d	19,10%	8,35%	62,02%	30,91%	215297	167200	-22,34%
Geriatrics-i	32,23%	16,07%	65,08%	45,93%	35930	46986	30,77%
Geriatrics-d	52,53%	27,67%	79,37%	66,69%	53718	69733	29,81%
Psychiatry-i	15,97%	10,75%	47,18%	42,56%	147960	145491	-1,67%
Psychiatry-d	35,44%	29,15%	66,91%	65,20%	129340	125936	-2,63%
Reception-i	55,20%	35,94%	79,18%	59,42%	80909	103832	28,33%
Reception-d	60,40%	38,70%	95,41%	89,24%	57749	61478	6,46%

Table 1 Results for the centralised approach. Instance-i refers to cases with identical contracts, instance-d refers to cases with different contracts.

solution quality. However, this result is not consistent for all instances under study and deserves further investigation. One obvious future step incorporates optimising FO while at the same time attempting to improve WO, without decreasing the quality of the worst individual roster. Furthermore, decentralised approaches, e.g. agent-based frameworks [3], present other interesting possibilities for defining new fairness measures in which individual nurses' objectives can be optimised at the same time [4].

Acknowledgements This research was jointly carried out within the IWT project (IWT 110257) and the research sabbatical sponsored by the University of Portsmouth, Department of Mathematics.

References

1. E. K. Burke, P. De Causmaecker, G. Vanden Berghe and H. Van Landeghem, The state of the art of nurse rostering, Journal of Scheduling, 7(6):441-499 (2004)
2. M. L. De Grano, J. D. Medeiros, and D. Eitel, Accommodating individual preferences in nurse scheduling via auctions and optimization, Health Care Management Science, 12(3):228-242 (2009)
3. S. Martin, D. Ouelhadj, P. Beullens and E. Özcan, A generic agent-based framework for cooperative search using pattern matching and reinforcement learning, Technical report 5861, University of Portsmouth (2012)
4. D. Ouelhadj, S. Martin, P. Smet, E. Özcan and G. Vanden Berghe, Fairness in nurse rostering, Technical report, KAHO Sint-Lieven (2012)
5. E. Rönnberg and T. Larsson, Automating the self-scheduling process of nurses in Swedish healthcare: a pilot study, Health Care Management Science, 13:35-53 (2010)
6. P. Smet, B. Bilgin, P. De Causmaecker and G. Vanden Berghe, Modelling and evaluation issues in nurse rostering. Annals of Operations Research, 1-24 (2012)