A Direct MILP Approach Based on State-Expanded Network Flows and Anticipation for Multi-Stage Nurse Rostering under Uncertainty

Michael Römer · Taïeb Mellouli

1 Introduction

This paper sketches the approach used in our solver submitted to the Second International Nurse Rostering Competition (INRC-II) – for details regarding the problem setting and the rules of the competitions, see (Ceschia et al, 2015). The INRC-II considers a multi-stage setting under uncertainty: Each stage consists of a single planning week for which request and demand data is known while future data is uncertain. Thus, when solving a planning problem for a given week, the solver needs to take the (uncertain) future weeks into account. As a result, every solver needs to deal with two main challenges: First, it needs to be able find high-quality schedules for a single week quickly (strict time limits are imposed in the competition). Second, it needs to find a way to deal with the multi-stage setting under uncertainty, that is, to find rosters not only forming good solutions for the week under consideration but also leaving favorable conditions for future scheduling periods. In the following sections, we discuss our approaches for addressing these challenges.

2 Modeling and Solution Approach

We tackled the problem by formulating it as a network flow-based MILP directly solved by the open source solver Coin CBC without resorting to decomposition approaches such as Branch-and-Price. Our formulation can be charac-

M. Römer
Institute of Business Informations Systems and Operations Research
Martin Luther University Halle-Wittenberg, Germany
Tel.: +49 345 55 23404
E-mail: michael.roemer@wiwi.uni-halle.de

T. Mellouli
Institute of Business Informations Systems and Operations Research
Martin Luther University Halle-Wittenberg, Germany

Proceedings of the 11th International Confenference on Practice and Theory of Automated Timetabling (PATAT-2016) – Udine, Italy, August 23–26, 2016

terized as a multi-commodity flow MILP formulation based on state-expanded networks; for similar formulations used in the context of airline crew scheduling, see e.g. (Mellouli, 2001) and (Römer and Mellouli, 2011). Each nurse is associated with a directed acyclic network layer in which arcs correspond to shift and days off assignments; a flow from the source to the sink in that layer can be interpreted as a roster. The network is expanded by states in a way that each node is associated not only with a time and shift index but also with state information needed to evaluate a subset of the rules (mostly rules considering blocks of work and days off). As a consequence, each shift assignment is associated with multiple arcs emanating from different state nodes; the state-dependent soft rule penalties are associated with the flow variables. The rules not embedded in the network (that is, rules concerning the total number of assignments and work weekends) are enforced by supplementary constraints. The model considers the assignment of shifts and days off individually for each nurse. In order to reduce model symmetry, the assignment of the skill in which a nurse works is handled by an aggregated model component. The detailed skill assignment is then performed in a postprocessing step.

While the sketched formulation leads to comparatively big model instances in terms of variables and constraints, the formulation is strong, that is, the instances exhibit a small duality gap. In addition, the formulation allowed us to heuristically reduce the size of the model instances by removing nodes and arcs associated with certain states we considered as unlikely to occur in a good solution (e.g. states involving big violations of block-related soft rules). Finally, our solver recognizes separable subproblems and solves them independently: If the set of skills in an instance involves disjunct subsets of skills connected by the skillsets of the nurses, the rostering problems for these disjunct skillsets can be solved separately which reduces the computational burden.

We performed an extensive testing with different parameter settings of the CBC solver we employed for solving our MILP model instances in order to achieve feasible high-quality solutions in a short amount of time. For example, our solver uses the barrier solver for the LP relaxation, preprocessing and strong branching are turned off and the diving heuristics are turned on.

3 Dealing with Multi-Stage Characteristic and Uncertainty

In order to handle the multi-stage characteristic combined with the uncertainty regarding future demands and requests in the INRC-II setting, our solver employs a deterministic lookahead approach: Using artificially generated demand data, the planning period is extended by an anticipation period. The generation of demand data makes use of known demand information from the given planning week and from past weeks.

In order to reduce the effect of the anticipation component on the size and computational complexity of the model instances, we employ three main ideas: First, in the anticipation, we do not consider each nurse separately but we aggregate the nurses according to contracts and skills. Second, we more aggressively apply the heuristic model reductions described above in the anticipation component. Third, for all variables in the anticipation component, we remove the integrality requirements – given the small duality gap of our formulation, the impact of this relaxation on the quality of the anticipation can be neglected.

4 Competition Results

The competition involved a qualification phase based on known instances followed by a final in which the solvers of seven qualified teams were evaluated with a set of hidden instances. The ranks were computed by taking the average ranking on all runs (for each instance, multiple runs with different seeds were performed). Our solver ranked first on both the known and hidden instances in the competition with an average rank of 1.76 on the hidden instances. It is interesting to observe that both our solver and the column generation-based solver by the team from Montréal ranking second (avg. rank 1.86) employ "exact" methods based on mathematical programming while the other finalist approaches using heuristics performed significantly worse.

When looking at the details of the results, one can observe that our solver found the best solution in 394 out of 600 runs. In 36 runs, however, our solver did not find a feasible solution within the given time limit. This behavior is due to the fact that the open source MILP solver was not able to find an integer solution within the given time limits. Note that this problem may be addressed by choosing more conservative values for the parameters heuristically reducing the size of the instances. In addition, given the fact that the majority of the rules in the competition are soft, an integer solution can quickly be obtained from an LP solution by a simple problem-specific rounding heuristic.

References

Ceschia S, Dang NTT, De Causmaecker P, Haspeslagh S, Schaerf A (2015) Second International Nurse Rostering Competition (INRC-II) - Problem Description and Rules. Tech. Rep. arXiv:1501.04177

Mellouli T (2001) A Network Flow Approach to Crew Scheduling based on an Analogy to a Train/Aircraft Maintenance Routing Problem. In: Voss S, Daduna J (eds) Computer-Aided Scheduling of Public Transport, LNEMS, vol 505, Springer, Berlin, pp 91–120

Römer M, Mellouli T (2011) Handling Rest Requirements and Preassigned Activities in Airline Crew Pairing Optimization. In: Pahl J, Reiners T, Voß S (eds) Network Optimization, vol 6701, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 643–656