FUTURE DEMAND UNCERTAINTY IN PERSONNEL SCHEDULING: INVESTIGATING DETERMINISTIC LOOKAHEAD POLICIES USING OPTIMIZATION AND SIMULATION

Michael Römer
Institute of Information Systems and Operations Research
Martin Luther University Halle-Wittenberg, Germany
Email: michael.roemer@wiwi.uni-halle.de

Taïeb Mellouli
Institute of Information Systems and Operations Research
Martin Luther University Halle-Wittenberg, Germany
Email: mellouli@wiwi.uni-halle.de

KEYWORDS

Personnel Scheduling; Uncertainty, Optimization, Simulation; Computational Stochastic Optimization; Lookahead Policies

ABSTRACT

One of the main characteristics of personnel scheduling problems is the multitude of rules governing schedule feasibility and quality. This paper deals with an issue in personnel scheduling which is both relevant in practice and often neglected in academic research: When evaluating a schedule for a given planning period, the scheduling history preceding this period has to be taken into account. On the one hand, the history restricts the space of possible schedules, in particular at the beginning of the planning period and with respect to rules a scope transcending the planning period. On the other hand, the schedule for the planning period under consideration affects the solution space of future planning periods. In particular if the demand in future planning periods is subject to uncertainty, an interesting question is how to account for these effects when optimizing the schedule for a given planning period. The resulting planning problem can be considered as a multistage stochastic optimization problem which can be tackled by different modeling and solution approaches. In this paper, we compare different deterministic lookahead policies in which a one-week scheduling period is extended by an artificial lookahead period. In particular, we vary both the length and the way of creating demand forecasts for this lookahead period. The evaluation is carried out using a stochastic simulation in which weekly demands are sampled and the scheduling problems are solved exactly using mixed integer linear programming techniques. Our computational experiments based on data sets from the Second International Nurse Rostering Competition show that the length of the lookahead period is crucial to find good-quality solutions in the considered setting.

INTRODUCTION

Personnel scheduling problems have been widely studied in the Operations Research literature. The interest in this class of problems stems from the fact that, compared to scheduling problems dealing with more "simple" resources, personnel-related problems are considerably more complex: In personnel scheduling problems, a multitude of rules and objectives need to be considered. An important source of these rules is the human need for daily, weekly and annual recreation. Furthermore, quality-of-life aspects often play an important role in the objective function of personnel scheduling problems: People value having a full weekend off; furthermore, they have individual preferences for days and shifts off and other schedule attributes. On the one hand, these individual preferences lead to to the fact that there are often different contracts such as full-time and part-time contracts and sometimes even different payment schemes. On the other hand, personalized preferences expressed e.g. via requests for days off lead to the fact that personnel scheduling problems often cannot simply be regarded on a resource-aggregated level but every staff member needs to be considered individually.

While the challenging characteristics of personnel scheduling problems sketched so far form a good source of challenges for the scientific community in the field of Operations Research – for a recent overview of the research dealing with personnel, see e.g. (Van den Bergh et al., 2013) – many of those studies deal with simplified and stylized problem instances. This paper deals with two practically relevant issues to be considered in real-world personnel scheduling often neglected in academic papers: The multistage nature of personnel scheduling problems and the uncertainty of the demand to be covered in future planning periods. In practice, when evaluating a schedule for a given planning period, the scheduling history preceding this period has to be taken into account. On the one hand, this restricts the space of possible schedules, in particular at the beginning of the planning period and with respect to rules a scope transcending the planning period. On the other hand, the schedule for the planning period under consideration affects the solution space of future planning periods. In particular if the demand in future planning periods is subject to uncertainty, an interesting question is how to account for these effects when optimizing the schedule for a given planning period.

In this paper, we show that the planning problem resulting from considering the described multistage characteristic along with uncertain future demands can be considered as a multistage stochastic optimization problem. Furthermore, we propose deterministic lookahead policies for this problem in which at each stage, a mixed-integer linear programming problem is solved for the planning period under consideration augmented with a lookahead period. Finally, we evaluate these policies using publicly available problem instances from the Second International Nurse Rostering Competition (INRC-II).

The paper is structured as follows: In the next section, we describe the problem setting of the INRC-II forming the source of the data sets used in the computational experiments. Then, we provide a short review of related work followed

Proceedings 30th European Conference on Modelling and Simulation ©ECMS Thorsten Claus, Frank Herrmann, Michael Manitz, Oliver Rose (Editors)
ISBN: 978-0-9932440-2-5 / ISBN: 978-0-9932440-3-2 (CD)

by a characterization of the personnel scheduling problem considered in this paper as a multistage stochastic optimization problem. Next, we present the deterministic lookahead policies to be investigated in this paper followed by the description of the experimental design and the computational results.

PROBLEM SETTING

The problem setting along with the data sets considered in this paper stems from the Second International Nurse Rostering Competition (INRC-II). For a detailed description of this competition and its problem setting, see (Ceschia et al., 2015). In this section, we will briefly sketch its main characteristics relevant to our investigation.

The INRC-II problem consists in finding a cost-minimal schedule (a sequence of shift assignments and days off) for a given set of nurses covering the given shift-wise demand and respecting all hard roster legality rules for a given planning horizon. The cost function consists of a linear combination of penalties for violating soft rules of the problem. Furthermore, each nurse has a given set of skills and certain contract (e.g. full time, half time and part time) governing certain rule-related parameters such as the maximum number of working days in the planning horizon. Using the numbering and notation from (Ceschia et al., 2015) in which **H** stands for a hard and **S** for a soft constraint, the constraints used for evaluating a solution are as follows:

H1 A nurse can be assigned at most one shift per day.

H2 For each (day/shift/skill) combination, the assigned number of nurses must cover the minimum requirement.

H3 Two consecutive shifts of one nurse must form a legal shift type succession (e.g., early must not follow night shift)

S1 Respect the the optimal requirement for each (day/shift/skill) combination

S2 Respect the minimum and maximum number of consecutive work days (in general and for each shift type)

S3 Respect the minimum and maximum number of consecutive days off

S4 Respect the shift off requests for each nurse

S5 At a weekend, a nurse should either work both days or no day at all

S6 Respect the minimum and maximum number of total work days in the planning horizon

S7 Respect the maximum number of total working weekends.

Since they affect the legality of blocks of days on and days off, in the following discussion, we will refer to the rules H3, S2, and S3 as "block-related" rules. Similarly, since they affect the full planning horizon, we refer to the rules S6 and S7 as "full-horizon" rules.

One of the main features of INRC-II competition is the fact that the problem to be solved is a multistage problem under uncertainty: While each instance consists of a four- or eight-week scheduling horizon, demand and request information only becomes available for a single week in each stage. Thus, in each stage, a single-week scheduling problem needs to be solved – under consideration of the history from the previous week(s) affecting the evaluation of the full-horizon rules and of the block-related rules at the beginning of the week. Given the last statement, it becomes clear that a main challenge of the described problem lies in both finding a good schedule for the week under consideration and leaving a history allowing finding good schedules for the subsequent week(s) – for which both demand and request information is unknown. Following Powell (2014), this type of problem can be considered as a multistage stochastic resource allocation problem.

Each publicly available data set for the INRC-II consists of multiple files: A scenario file containing nurse-, contract- and rule-related data, multiple history files which can serve as a history for the first week and 10 week data files containing demand and preference information. Note that in order to reduce the computational burden for for the extensive experiments performed for the present paper, for each of the considered INRC-II instances, we only consider the skill "trainee". On the one hand, this reduces the number of nurses to be regarded per instance. On the other hand, since in all instances, the trainee nurses only have a single skill, the originally multi-skill setting of the INRC-II problem is turned into a single-skill setting.

RELATED WORK

For a recent and comprehensive survey of the Operations Research literature dealing with variants of personnel scheduling problems, see (Van den Bergh et al., 2013). This section intends to provide a short overview of work closely related to the problem sketched in the previous section and to our approaches used in the following sections.

When it comes to solving (deterministic) personnel scheduling problems, as discussed by Van den Bergh et al. (2013), one can distinguish exact and heuristic methods. Many of the most successful exact approaches are based on Mathematical Programming, making use of state-of-the art solvers complemented by problem-specific valid inequalities (see e.g. Santos et al., 2014) and/or advanced techniques such as branch-and-price (see e.g. Burke and Curtois, 2014). Note that while we use a Mixed-Integer Linear Programming approach to solve the scheduling problem in each stage, the focus of the present paper is on evaluating policies for handling the multistage stochastic nature of the problem. Consequently, besides the fact that we use an exact approach for solving these problems in order to avoid issues with regard to solution quality introduced by heuristic approaches, the choice of the modeling and solution approach is not of primary importance for the results of the present paper.

The first important aspect considered in this paper is the multistage characteristic of the problem described in the previous section. Note that the multitude and the complexity of the schedule legality rules makes this issue particularly relevant in personnel scheduling problems: On the one hand, there are full-horizon rules transcending the planning period of each stage, on the other hand, there are local or block-related rules affecting the start of the planning period. While this issue is often ignored in the personnel scheduling literature, recently, Salassa and Vanden Berghe (2012) address this issue in a deterministic setting and show that neglecting the multistage characteristic leads to inferior results in practice. Furthermore, it can be expected that the INRC-II competition will draw a certain amount of research interest towards this issue (which actually happened to the authors of this paper).

The second important aspect of the problem sketched in the previous section is the fact that the demand and request information in future stages is subject to uncertainty. As shown in (Van den Bergh et al., 2013), there are some works dealing with uncertainty in personnel scheduling, mostly using stochastic programming and robust optimization approaches. Most papers however, deal with a two-stage setting in which the decision stages deal with different types of decisions: For example, in (Kim and Mehrotra, 2015), the first stage involves staffing decisions and the second stage involves the selection of weekly scheduling patterns under uncertain demand. Other papers such as (Campbell, 2011) deal with the problem of creating a schedule under demand uncertainty which then has to be adjusted in an operational setting forming the second-stage subproblem. This contrasts with the problem setting described in the previous section in which the type of decisions taken at each stage have the same type: Constructing a schedule for a full week for which demand is known while demand for the subsequent weeks is subject to uncertainty. Following Powell (2014), this type of problem can be characterized as a sequential or multistage stochastic optimization problem; it consists in finding an optimal policy, that is a function mapping states to decisions. While in general, this optimization problem is intractable, approaches from the field of *approximate dynamic programming* have been successfully applied to multistage stochastic resource allocation problems, see various case studies e.g. on fleet and driver management in the less-than-truckload industry described in the monograph (Powell, 2011).

To the best of our knowledge, however, there is no study applying similar techniques to personnel scheduling problems as described in the last section. In the following sections, we demonstrate that the problem can in fact be interpreted and modeled as a multistage stochastic optimization problem. Furthermore, we propose variants of deterministic lookahead policies and evaluate these policies using stochastic simulation.

MODELING AS A MULTISTAGE STOCHASTIC OPTIMIZATION PROBLEM

The problem addressed in this paper can be considered as as a (stochastic) dynamic resource allocation problem. While there are different modeling and solution frameworks for this class of problems, among which is (multistage) stochastic programming, in this paper, we will model the problem using the more general framework for multistage stochastic optimization problems discussed in (Powell, 2014). In the personnel scheduling problem under consideration, each stage t consists of a weekly scheduling problem for which the demand is assumed to be known. The future demand from stage $t+1$ to the final stage T (in the instances considered in this paper, T is 4 or 8), however, is uncertain. Note that an interesting feature of the personnel planning problem under consideration is that each stage itself consists of a dynamic resource allocation problem (which in addition has to account for the impact of the resulting schedule on future planning weeks with uncertain demand).

According to Powell (2014), a multistage stochastic optimization can be be modeled using a framework encompassing five elements, each of which we will shortly explain and apply to the personnel scheduling problem addressed in this paper.

The first element is the characterization of the **state** S_t of the system at the time t before a decision is made. S_t is the so-called state variable, which, following Powell (2014), can be defined as "the minimally dimensioned function of history that is necessary and sufficient to compute the decision function, the transition function and the contribution function". In the case of the personnel planning problem considered here, the elements of the state variable encompass the resource state at the beginning of the planning week and the demand and request information (in the INRC-II problem, this information is supplied in form of the history and week data files). Note that in the case study under consideration, the probability distribution of demand and request data does not depend on the week index t. The resource state represents all rule-relevant information such as number of days worked so far, number consecutive work days up to the border of the planning period etc.

The next element are the **decisions** x_t to be determined by the chosen policy in stage t. In the personnel scheduling problem under consideration, the vector x_t encompasses all assignment decisions, that is, there is a binary decision variable for each combination of nurse, day d in the week t under consideration and shift type. Note that the ensuring the feasibility of the decision vector x_t is one of the things to consider when designing a policy π.

The third element of the framework is the vector of **exogenous information** W_t becoming available in period t: In case of the personnel scheduling problem, the exogenous information involves the demand for each combination of day and shift as well as the shift off requests.

The next element is the **transition function** $S^M(S_t, x_t, W_{t+1})$ describing the transition from state S_t to S_{t+1} given the decisions x_t taken in t and the exogenous information becoming available in $t+1$. In the personnel scheduling problem considered in this paper, the transition to the resource state R_{t+1} is performed by computing the new schedule history information for each nurse based on the information R_t and on the assignment information contained in the vector x_t. Since the demand and request information is not history-dependent, the transition with regard to this part of the state variable is performed by replacing the information from period t with the newly arrived exogenous information contained in W_{t+1}.

The last element of the modeling framework is the **objective function** stating the overall objective of minimizing the expected costs over the planning horizon. Note that in the case of the planning problem under consideration, the planning horizon is finite and consists of T periods. The objective function (which is linear and does not involve a discount factor) can be formulated as follows:

$$\min_{\pi \in \Pi} \mathbb{E}^{\pi} \sum_{t=1}^{T} C(S_t, X_t^{\pi}(S_t)) \tag{1}$$

Note that the optimization problem (1) consists in finding the cost-minimal policy π from the set Π of all policies. Since π is a function, the problem forms a search in a function space.

DETERMINISTIC LOOKAHEAD POLICIES

At the time of this writing, no computationally tractable method for solving problem (1) exactly is known. Nonetheless, there exist methods for approximately tackling this type of problem. According to Powell (2014), there are four fundamental classes of policies (and hybrids of these classes) typically employed to address multistage stochastic optimization problems:

Policy function approximations (PFA) represent an analytic function mapping a state to an action or a set of decisions. For example, a PFA may be a simple lookup table, a decision rule or a linear or polynomial function. Note that applying a PFA does not involve solving an optimization problem.

(Myopic) Cost function approximations (CFA) are formed by modifying the objective function and/or the constraints of the decision problem to be solved in stage t in a way that the resulting policy does not only involve the single-stage cost function but also accounts for the impact of the decisions in stage t on the future stages.

Value function approximations (VFA) involve constructing and calibrating a function approximating the future value of the state resulting from taking a decision. For example, in a resource allocation problem, value function approximation can be designed around the post-decision state, that is, the state after having taken a decision, of the resources to be allocated.

Lookahead policies involve solving a multistage decision problem at each stage by explicitly considering the future decision stages for a lookahead horizon to be specified. Lookahead policies come in two main flavors: In a deterministic lookahead policy, the uncertain parameters in the exogenous information process are replaced by point estimates resulting in a deterministic multistage optimization problem to be solved in each stage. In a stochastic lookahead policy, the uncertain parameters are modeled in a stochastic scenario tree; the resulting problem to be solved then forms a (multistage) stochastic programming problem (which may be approximated by a two-stage stochastic programming problem in order to reduce the computational complexity of the problem).

For the personnel scheduling problem considered in this paper, we first tried to design myopic cost function approximations and simple value function approximations – however, we were not satisfied with the results: In fact, it was difficult to even get an intuitive understanding of what makes a "good" post-decision state (that is, the state after having taken the scheduling decisions for the week under consideration) and thus, at least to our impression, approximating the future value of a given state is difficult for this type of problem.

As a result, we started experimenting with deterministic lookahead policies which yielded much better results. This was in line with our intuition that a lookahead model would be a good approach for dealing with the impact of the block-related rules at the end of the scheduling week: Instead of explicitly stating what makes a good state at the end of a week, adding a lookahead period H may implicitly yield good end-of-the week states since the impact of the schedule for the planning week under consideration on the subsequent week is accounted for by considering the lookahead period.

The resulting deterministic lookahead policy π, parameterized by the forecasting strategy *fs* and the length H of the lookahead horizon can be written as follows:

$$X_t^{\pi(fs,H)}(S_t) = \underset{x_t \in \mathcal{X}_t}{\arg\min} \left(c_{tt} x_{tt} + \sum_{t'=t+1}^{t+H} c_{tt'} x_{tt'} \right) \quad (2)$$

Using this policy, in each stage t, the optimization problem does not involve the decisions variables (and possibly needed supplementary variables) affecting t (represented by the vector x_{tt}), but also the decisions $x_{tt'}$ affecting stages contained in the lookahead period ranging from $t+1$ to $t+H$. Note that the full vector of decision variables from all these stages is denoted with x_t; the set $\mathcal{X}_t$ depends on the state S_t and denotes the set of all feasible vectors x_t.

Clearly, (2) forms a very high-level statement of a deterministic lookahead policy. For our experiments, we formulated (2) as a mixed-integer linear program – that is, $\mathcal{X}_t$ is formulated as a set of linear inequalities and integer constraints – we solve using a standard solver. Note that instead of using a compact formulation as discussed in Santos et al. (2014) or an extreme-point formulation as proposed by Burke and Curtois (2014), we use a multi-commodity flow formulation in the spirit of (Mellouli, 2001) and (Römer and Mellouli, 2011) dealing with airline crew scheduling.

In order to capture the soft full-horizon constraints (in our case, concerning the limitations regarding number of days of work and of working weekends) in the model, the upper bounds belonging to the full planning horizon were adjusted according to the relative amount of time passed after the end of the lookahead horizon $t + H$.

When designing a deterministic lookahead model it is necessary to make point forecasts for the exogenous information $W_{t'}$ for the lookahead period from $t + 1$ to $t + H$ for which demand is uncertain at stage t. In this paper, we consider two strategies *fs* for obtaining these forecasts. The first approach, referred to by $fs := R$, uses a simple resampling strategy: For each of the periods t', one of the ten week data files in the INRC-II dataset under consideration is randomly drawn and the demand information from this week data file is used as a forecast.

The second approach is based on (conditional) averaging and referred to by $fs := A$ in the experimental results section. Based on an analysis of the weekly demand data of all INRC-II data sets, we figured out that in most cases, the demand of weekdays is significantly different from the demand on weekends. As a result, for each data set, we compute the average demand for weekdays and weekend days across all week data files and use these average values as point forecasts for the respective types of days (the weekday average is thus use from Monday to Friday and the weekend average for Saturday and Sunday).

In addition to varying the forecasting strategy, we experiment with the length of the lookahead horizon. Note that we do not only use full week increments, but also consider a lookahead horizon of three days. In this case, the H takes the fractional value $3/7$ and the lookahead model is constructed accordingly.

EXPERIMENTAL DESIGN AND RESULTS

The problem instances for evaluating the policies described above are publicly available and stem from the Second International Nurse Rostering Competition (INRC-II) described in (Ceschia et al., 2015). As discussed in the problem description above, in order to manage the computational burden of the experiments, we artificially reduced these instances by only considering trainee nurses, resulting in a reduction with respect to the number of nurses to approximately a quarter of the original number (the tables below display the number of trainee nurses per instance).

For each of the resulting instances, we evaluate the performance of the policies described in the previous section by a simulation-based approximation of the objective function (1): We randomly sample N paths for the full scheduling horizon T (that is, we randomly select a history file for the first week and then select a random sequence of week data files for each of the data sets). The estimated performance $\hat{F}^{\pi(fs,H)}$ of a policy $\pi(fs, H)$ is then estimated by applying the policy to each sample path and averaging the resulting costs:

$$\hat{F}^{\pi(fs,H)} = \frac{1}{N}\sum_{n=1}^{N}\left(\sum_{t=1}^{T} C(S_t, X_t^{\pi(fs,H)}(S_t))\right) \qquad (3)$$

As discussed in the previous section, in this paper, we consider deterministic lookahead policies varying with respect to two dimensions: The forecasting strategy *fs* and the length H of the lookahead horizon. Concerning the forecasting strategy, we use an averaging strategy (signified by A) and a resampling strategy (signified by R). When it comes to the lookahead horizon H, we consider multiple values: First of all, we consider a myopic policy with $H = 0$ in which there is no lookahead at all in order to figure out whether a lookahead is useful at all. Then, we use a lookahead period of three days ($H = 3/7$) in order to consider a short lookahead horizon – the main motivation is to evaluate whether such a short lookahead with a comparably small computational burden is able to address the border effects resulting from block-related constraints such as shift sequence rules. Furthermore, for both the four- and the eight-week instances, lookahead horizons of 1, 2 and 3 weeks are considered (note that for a four-week instance, $H = 3$ means that in the first week, the full planning horizon of $T = 4$ is considered in the lookahead model). Moreover, for the eight-week instances, $H = 5$ and $H = 7$ are considered.

All experiments were conducted on a personal computer with 8 GB Ram, with an Intel Core i7 4-core CPU with 3.4 GHz. The simulation and the model generation are implemented in C++. For solving the mixed integer linear programs (MILP) within the deterministic lookahead approach, we used IBM ILOG CPLEX 12.6. For each instance and for each policy, we carried out $N = 1000$ replications in order to obtain a good estimate of the true value of each policy. Note that the number of (non-trivial) MILP to solve within one replication corresponds to the number of weeks in the instance. As a consequence, since in total we had to solve 912 000 MILP, we set the mipgap parameter to 0.5 % and imposed a time limit of 4 minutes per MILP. Note, however, that this time limit was rarely ever hit.

TABLE I. AVERAGE PERFORMANCE FOR 4-WEEK INSTANCES ($N = 1000$)

		lookahead horizon H (weeks)				
# nurses	*fs*	0	3/7	1	2	3
5	A	575.3	492.8	444.5	441.1	438.6
	R	575.3	494.4	447.7	443.2	439.9
8	A	614.1	613.0	559.5	544.8	532.7
	R	614.1	625.6	551.4	539.3	539.1
10	A	637.2	564.0	441.8	424.0	422.6
	R	637.2	578.8	445.6	432.8	428.9
20	A	920.3	797.3	566.5	509.9	508.6
	R	920.3	829.3	561.8	519.1	512.6
25	A	829.2	692.0	509.7	441.4	433.0
	R	829.2	680.1	514.6	449.1	443.3
30	A	1694.5	999.3	688.0	619.4	590.2
	R	1694.5	974.6	683.3	610.7	585.4

The results for the four-week instances are presented in Table I. The most evident observation concerns the effect of the lookahead horizon: For all instances (and for both forecasting strategies), increasing the lookahead horizon leads to a significant cost reduction. For all but one instance, even a three-day lookahead yields a major improvement (more than 10 % on average) compared to a myopic policy without lookahead. Then, for every additional week, the solution quality is further improved.

A very interesting result is that the positive effect of a long lookahead period grows with the size of the instances in terms of the number of nurses: The relative impact of using a long lookahead horizon is much bigger for bigger instances (e.g., the best solution is 25 % better than the myopic solution in the 5-nurse instance while it is 75 % better for the 30-nurse instance). Furthermore, the marginal positive impact of additional weeks to the horizon is higher for big instances.

Regarding the choice of the lookahead strategy, the result is not as clear: In some cases, the resampling-based strategy (R) works better, in other cases the average-based strategy (A) is more effective. Even for a single instance, one strategy does not necessarily dominate the other when considering all tested lookahead horizons. On average, however, the strategy A turns out to perform slightly better. In general, given the described results and given the fact that both strategies are fairly simple it may be beneficial to develop more sophisticated forecasting strategies in order to achieve better overall results.

Most of the statements made regarding the results for the four-week instances also hold for the results for the eight-week instances displayed in Table II. In general, a longer forecasting horizon leads to better results and in most (but not all) cases the averaging-based forecasting strategy performs better than the resampling strategy. For some instances, interestingly, even if there is some benefit of using a very long lookahead horizon, the largest part of the improvement compared to a myopic policy can be obtained with a 1- or 2-week lookahead period.

TABLE II. AVERAGE PERFORMANCE FOR 8-WEEK INSTANCES ($N = 1000$)

# n.	*fs*	lookahead horizon H (weeks) 0	3/7	1	2	3	5	7
5	A	807.9	722.8	652.7	637.4	630.2	633.8	633.6
	R	807.9	746.8	650.5	633.7	620.6	612.8	613.2
8	A	1318.3	1572.1	1134.1	1112.6	1085.2	1065.5	1069.9
	R	1318.3	1552.2	1134.1	1106.2	1086.5	1073.0	1070.1
10	A	1374.0	1316.5	1001.8	973.2	942.6	910.3	904.2
	R	1374.0	1351.8	1012.5	973.9	944.0	913.3	909.8
20	A	1645.7	1826.7	1289.0	1194.6	1167.8	1154.3	1154.0
	R	1645.7	1832.9	1278.1	1201.5	1166.4	1153.1	1147.6
25	A	1821.6	1697.2	1204.0	1149.6	1123.8	1057.0	1040.3
	R	1821.6	1700.5	1200.2	1154.0	1120.1	1065.3	1050.7
30	A	3416.7	2082.0	1047.4	834.4	761.7	671.4	657.2
	R	3416.7	2006.0	1017.0	825.1	767.5	682.4	670.2

Again, the relative effects of the lookahead policy grow with the instance size.

Finally, it should be noted that an issue which may be considered in more detail are the characteristics of the problem instances under consideration and their effect on the choice of the policy. While the tables only show the number of nurses for describing a problem instance, the instances may vary with respect to several other factors such as the relative frequency of nurses with part-time contracts, the values of rule-related parameters such as the maximum number of consecutive work days, the demand level in relation to available nurses and the variability of demand. It can be suspected that these characteristics have a certain impact on the effectiveness of policies and may help to explain the differences in the results between the instances which cannot simply be explained by the number of available nurses.

CONCLUSIONS AND FUTURE RESEARCH OPPORTUNITIES

In this paper, we consider a class of personnel scheduling problems under future demand uncertainty in which scheduling is carried out and fixed for a planning period (e.g. one week) and demand for the future planning periods is unknown. We argue that this type of problem forms a multistage (also called sequential) stochastic optimization problem consisting in finding an optimal policy, that is an optimal decision function for each stage. Among the classes of policies which can be used to tackle this type of problem (see Powell, 2014), in this paper, we investigate variants of deterministic lookahead policies with different forecasting strategies and lookahead horizons.

The computational results impressively show that even fairly simple deterministic lookahead policies can lead to huge improvements compared to a naive myopic approach. The improvements tend to grow with the instance size with respect to the number of nurses and to the length of the scheduling horizon T as well as with the length of the lookahead horizon chosen for the lookahead policy.

The results presented in this paper form a first step to investigate personnel scheduling problems in the framework of computational stochastic optimization as discussed in (Powell, 2014). Besides trying to improve the deterministic lookahead policies investigated in this paper, a natural next step would be to consider a stochastic lookahead policy by formulating a two- or multistage stochastic programming model. Furthermore, after thoroughly examining the problem structure and the different solutions from these policies, it might be tried to find a good value function approximation allowing to solve larger problem instances. In addition, it appears promising to develop hybrid policies, e.g. by combining a lookahead policy with a cost function approximation and/or a value function approximation.

Finally, in order allow addressing bigger problem instances, e.g. the full INRC-II instances involving all skills and multiple overlapping skill sets with up to 120 nurses, it is important to improve the modeling and solution approaches used to solve the lookahead problems.

REFERENCES

Burke, E. K. and Curtois, T. (2014), New approaches to nurse rostering benchmark instances, *European Journal of Operational Research* 237(1), 71–81.

Campbell, G. M. (2011), A two-stage stochastic program for scheduling and allocating cross-trained workers, *Journal of the Operational Research Society* 62(6), 1038–1047.

Ceschia, S., Dang, N. T. T., De Causmaecker, P., Haspeslagh, S. and Schaerf, A. (2015), Second International Nurse Rostering Competition (INRC-II) - Problem Description and Rules, Technical Report arXiv:1501.04177 [cs.AI], arXiv:1501.04177 [cs.AI].

Kim, K. and Mehrotra, S. (2015), A Two-Stage Stochastic Integer Programming Approach to Integrated Staffing and Scheduling with Application to Nurse Management, *Operations Research* 63(6), 1431–1451.

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

Powell, W. B. (2011), *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, John Wiley & Sons.

Powell, W. B. (2014), Clearing the Jungle of Stochastic Optimization, Bridging Data and Decisions, INFORMS Tutorials in Operations Research, INFORMS, pp. 109–137.

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

Salassa, F. and Vanden Berghe, G. (2012), A stepping horizon view on nurse rostering, D. Kjenstad, E. K. Burke and B. McCollum, eds, Proceedings of the 9th International Conference on the Practice and Theory of Automated Timetabling., PATAT, Son, Norway, 28-31 August 2012, pp. 161–173.

Santos, H. G., Toffolo, T. A. M., Gomes, R. A. M. and Ribas, S. (2014), Integer programming techniques for the nurse rostering problem, *Annals of Operations Research* pp. 1–27.

Van den Bergh, J., Beliën, J., De Bruecker, P., Demeulemeester, E. and De Boeck, L. (2013), Personnel scheduling: A literature review, *European Journal of Operational Research* 226(3), 367–385.