

A categorisation of nurse rostering problems

Patrick De Causmaecker · Greet Vanden Berghe

Published online: 16 December 2010
© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract Personnel rostering has received ample attention in recent years. Due to its social and economic relevance and due to its intrinsic complexity, it has become a major subject for scheduling and timetabling researchers. Among the personnel rostering problems, nurse rostering turned out to be particularly complex and difficult. In this paper, we propose a notation for nurse rostering problems along the lines of the $\alpha|\beta|\gamma$ notation for scheduling. The system allows extension as well as refinement. It is the aim of the notation to facilitate problem description, classification and systematic study. It enables authors to position the problems in the vaster body of research on the subject. By developing this notation for nurse rostering, we hope that an extension of it will be applicable to a broader domain of personnel rostering.

Keywords Nurse rostering · Personnel scheduling · Classification

1 Introduction

In this paper, we build on the work of the last decades to produce a classification system for nurse rostering problems. This system is based on a notation for which we borrow the main lines from the production scheduling domain. The

classification will considerably simplify problem representation. It allows one to classify problems according to a set of categories.

Specifying the characteristics of a problem unambiguously positions it in the classification. Hence, it is possible to compare problem definitions and pinpoint the differences. A problem may be next to other problems in the classification, differing in characteristics that are not expressed in the categorisation. An author may study the differentiation that originates from and now is clearly linked to the extra characteristics. Eventually, the observation may motivate refinement by the newly discovered problem characteristics. These refinements will be managed at <http://ingenieur.kahosl.be/vakgroep/it/nurse/archive.htm>. We follow a similar notation as the one that was used to classify production scheduling problems (Graham et al. 1979). As in scheduling, we may succeed in differentiating between hard and easy cases in specific categories. Developers will eventually have a better structured catalog of solution methods acting as a roadmap towards more efficient applications. This paper proposes such a categorisation and presents the analysis of a limited set of instances in a specific category as an example. The initial version of this notation was presented in a plenary talk at the PATAT 2008 conference (De Causmaecker 2008). It is accepted for publication in the post conference selected papers (De Causmaecker and Vanden Berghe 2011).

Although our study is mainly focussed on the nurse rostering literature, we believe that an extension of it can be applied to a broader field of personnel rostering. Many decades of research into automated methods to solve nurse rostering problems have brought to us a large collection of expressive models and powerful methods. Several authors (Siferd and Benton 1992; Burke et al. 2004) have scrutinised real world problems and published on their findings in great de-

P. De Causmaecker (✉)
Department of Computer Science, K.U. Leuven Campus Kortrijk,
E. Sabbelaan 53, 8500 Kortrijk, Belgium
e-mail: patrick.decausmaecker@kuleuven-kortrijk.be

G. Vanden Berghe
Information Technology, KaHo St.-Lieven, Gebr. Desmetstraat 1,
9000 Gent, Belgium
e-mail: greet.vandenbergh@kahosl.be

tail. The nurse rostering problem has, due to the required expressiveness and flexibility of the models and due to the apparent hardness of the resulting optimisation problems, become an attractive problem for many methods in optimisation, such as exact methods, heuristic procedures and metaheuristics. We will not review the nurse rostering literature here. A thorough review was published in Burke et al. (2004). At the end of this paper, we will demonstrate our notation by classifying papers that mainly appeared after that review was published. Here we solely mention a limited number of examples to illustrate the diversity that we want to cover. Some authors discuss exactly solvable models, either constructed (e.g. Wong and Chun 2004) or real world examples (e.g. Azaiez and Al-Sharif 2005; Bard and Purnomo 2005a). When we turn to more complex real world problems for which no exact solution is found, the situations become more diverse. Even after this many years of nurse rostering research, a clear and crisp definition of the problem is still to emerge. This inhibits a systematic comparative study of models and solution approaches. There is a need for a clear definition of the different types of nurse rostering problem together with a larger number of shared datasets. The nurse rostering literature encompasses problems with very different characteristics. Certain authors concentrate on a cost per time unit/person (Aickelin and Li 2007; Maenhout and Vanhoucke 2008), while others deal with labor ergonomics considerations (Bourdais et al. 2003; Lodree et al. 2009) that concern the quality of the entire personal schedules (Gutjahr and Rauner 2007; Osogami and Imai 2000; Wong and Chun 2004). The latter are in general more characteristic for nurse and personnel scheduling problems encountered in the literature. Recently, many publications address real-life nurse rostering problems, for which data have been retrieved from hospitals (Beddoe and Petrovic 2007; Bilgin et al. 2011; Bourdais et al. 2003; Burke et al. 2008; Moz and Pato 2007; Parr and Thompson 2007). It is not the case that all of them provide public datasets. The problems addressed in Bard and Purnomo (2006) (the nurse addition problem) and in Moz and Pato (2007) (the nurse rerostering problem) differ slightly from the majority of problems discussed in the nurse rostering literature in that (1) the number of (extra) nurses needs to be minimised or (2) unexpected variations in workload need to be addressed while minimising the number of shift changes in the roster. Obviously, the applicability or the success of optimisation approaches strongly depends on the problem under consideration and, consequently, relating observations and problem characteristics or complexity could be useful.

Unlike some other well known optimisation problems (e.g. the scheduling problem (Brucker 2004; Herroelen et al. 1999; Pinedo 2002)), the hardness of the nurse rostering problem has only been analysed in exceptional cases (Osogami and Imai 2000; Vanhoucke and Maenhout 2009).

The effect of certain objectives or constraints (e.g. addressing coverage, consecutive assignments, fairness, etc.) on the complexity of the problem is still unclear.

The constraints and models encountered in nurse rostering papers are extensive and complex. It is easy to see how the shift structures and requirements for workers in a production unit can be mapped into a simple nurse rostering model. Moreover, the last decades have seen a multitude of publications discussing specific cases and describing algorithms to tackle them. The availability of this vast body of literature together with the apparent complexity motivates the development of a systematic notation and classification system. Such a typology was developed and recently improved for cutting and packing problems (Dyckhoff 1990; Wäscher et al. 2007). A classification for the resource constrained project scheduling problems was undertaken in Brucker et al. (1999). The earliest effort in scheduling led researchers to develop a specific notation. In Graham et al. (1979) the $\alpha|\beta|\gamma$ notation was introduced. They created the framework to embody over 30 years of research. We will follow the structure of this notation and rephrase its categories to match the requirements of the personnel rostering field. We hope this will support empirical as well as theoretical study of approaches for personnel rostering.

We are aware that this does apply to more than one set of problems, especially but not exclusively in timetabling. By studying the nurse rostering problem, we hope to gain experience in how to tackle personnel rostering in particular and eventually the other domains.

In Sect. 2, we look at different initiatives for benchmarking or classifying nurse rostering problems. We identify the essential properties and present a novel notation for nurse rostering problems in Sect. 3. Each of the categories of that notation is detailed in the subsequent sections: the personnel environment α in Sect. 4; the work characteristics β in Sect. 5 and the optimisation objective γ in Sect. 6. In Sect. 7, we give examples of problems tackled in some of the recent nurse rostering research papers (most of which appeared after 2004) and categorise them according to the new notation. Section 8 presents conclusions and lines for future work.

2 Classifying and benchmarking nurse rostering problems

The nurse rostering problem still is a challenging research subject that is widely considered to be a hard optimisation problem to solve. Osogami and Imai (2000) prove that the nurse rostering problem is NP-hard. In fact, they prove that the timetabling problem, which is NP-complete, can be transformed into a decision version of the nurse rostering problem with only a subset of the real world constraints that apply to it. It is worth mentioning that some of the constraints characterising the nurse rostering problem, have not

been taken into account in this analysis. Examples of such constraints include restrictions with respect to consecutive assignments, constraints dealing with the simultaneous assignment of different people, etc. These constraints consume more evaluation time than the ones that involve counting assignments (Burke et al. 2001b). In addition, experiments show that problems become much harder to solve when including consecutiveness (Messelis et al. 2009). Unlike for the scheduling domain, where the $\alpha|\beta|\gamma$ notation (Lawler et al. 1993) is commonly accepted, the nurse rostering research lacks a general format to describe problems. For a long time, benchmark instances were missing. A few recent initiatives enable systematic study and offer a framework for the classification and formalisation of nurse rostering problems and general personnel scheduling problems. Two main efforts are in classification and benchmarking.

2.1 Classification

The problem description in Burke et al. (2004) puts forward general terminology and classifies a large number of papers according to the characteristics of the problems that they describe. These might assist researchers to position their particular problem or instance among other problems that have been reported on in the literature, and as such, build upon previous knowledge to improve the state-of-the-art approaches. The classification is not suitable for distinguishing, for example, complex problems from simple ones. However, it does provide some criteria for grouping particular problem types. Examples of such criteria are:

- *General problem characteristics: the amount of detail in which shift types, skills, coverage, etc. are defined, the flexibility of setting or defining problem parameters, whether the rosters are cyclic or not.*

Coverage, for example, is treated differently in many research papers. It is important to know whether under- and/or overstaffing is tolerated, whether coverage is feasible within an interval, etc. Problems addressing the same constraints but treating them differently (hard/soft) are thought of as quite different. It determines the adaptability of the model to address various types of nurse rostering problem in different health care environments, regions and countries. Flexible models require flexible approaches. Cyclical or semi-cyclical rosters are subject to additional constraints. The approaches for dealing with cyclical schedules may differ strongly from algorithms addressing non-cyclical problems.

- *Objectives: optimisation problems or decision problems, optimisation objectives such as minimising violations of constraints, minimising the number of staff, etc. and all kinds of combinations of criteria.*

Usually, a personnel rostering problem deals with multiple objectives. The nature of the objectives influences the requirements to the solution approach.

- *Constraints: the number of constraints, the type of constraints, whether they are hard or soft and whether they are set or configurable.*

We are not aware of a straightforward relation between the number of constraints and the complexity of the problem. Real world (and simplified) problem instances often deal with conflicting and/or redundant constraints. As a consequence, the number of constraints is not always an issue. Rather the kind of constraints and the presence of constraint combinations determines which approaches are most appropriate.

- *Dimensions of the problem: the planning horizon, the number of nurses, the number of shift types.*

Positioning new problems with respect to these characteristics should help to compare particular problems and to compare the performance of algorithms. However, a comparison of approaches probably makes sense only when applied to problems that have the same characteristics with respect to each of the criteria. Future papers will be more effective if researchers manage to clearly situate their work among the existing literature. The classification criteria in Burke et al. (2004) offer guidance.

A different approach is outlined in Maenhout and Vanhoucke (2008), where it is suggested to perform calculations of some instance parameters prior to the search, in order to improve the selection of optimisation approaches.

2.2 Datasets and benchmarks

A collection of personnel scheduling—mainly nurse rostering—problems and datasets is provided at <http://www.asap.cs.nott.ac.uk/projects/nmhpr/data/> (Burke et al. 2007). The data, most of which have been derived from real world problems, are available in XML format. It allows researchers to download a variety of problems for comparing the performance of their algorithms. The author of the web page also offers a graphical user interface, a parser and a set of solvers to tackle the problem. He keeps track of the best results for each problem instance. We want to support this initiative and hope that researchers in the field of nurse rostering will make an effort to describe their particular problems according to the same format and thus contribute to what should become a benchmark for this research domain.

The web site <http://allserv.kahosl.be/~burak/project.html> also presents a set of complex real problem instances and different scenarios associated with each set (Bilgin et al. 2008). The scenarios represent issues that the real life wards, from which the data have been taken, often deal with.

NSPLib (<http://www.projectmanagement.ugent.be/nsp.php>) (Vanhoucke and Maenhout 2007) is a complementary initiative in that it is built around a problem generator for nurse rostering. In contrast to the other benchmark site, the problems in NSPLib were not derived from real life ones but

they were constructed with that problem generator. A number of complexity indicators were applied to create instances that are as diverse as possible. Another set of instances is called the ‘realistic’ set mainly because the problems have a planning horizon of 28 days. The advantage of NSPLib is that, due to its size, it allows for proper statistical analysis of different approaches.

These recently published datasets are of high importance for researchers who want to compare their nurse rostering approach to the best performing approaches on particular datasets. Moreover, they facilitate data collection for (starting) researchers who, until recently, had to generate test sets themselves or collect them in collaboration with hospitals. In Sect. 7, we show how the instances from these datasets fit in the model that we propose.

An archive of problem instances keeps track of the online problem repositories among which are the two above. It is maintained at <http://ingenieur.kahosl.be/vakgroep/it/nurse/archive.htm>.

3 An $\alpha|\beta|\gamma$ notation

The main aim of this paper is the introduction of a framework for categorising nurse rostering problems. If we want to systematically study performance of algorithms as well as hardness of problems, we need a notation that enables distinguishing between nurse rostering problems. Such a notation should eventually allow situating the problem at hand and offer a tool for comparison of models, optimisation results and algorithm performance. It is clear that an algorithm that was developed for solving a particular class of problems does not necessarily behave well when applied to another class. In the next paragraphs, such a notation is proposed. Future refinements to the notation may be necessary.

It is important that the notation bridges between problem definitions and solution methodology. We start from the following observations. Any nurse rostering problem requires scheduling personnel in order to perform an amount of work over a certain period of time. It demands to do so while optimising one or more criteria. This situation is similar to the situation in scheduling problems, which require scheduling a number of machines in order to perform a number of jobs while optimising one or more criteria. Similar motivations as the ones developed in the current paper, led back in the seventies to the adoption of the $\alpha|\beta|\gamma$ classification system for scheduling problems (Graham et al. 1979). We found it useful to apply a similar type of categorisation to rostering problems. The distinction between machine environment α , work characteristics β and optimisation objectives γ is present in a nurse rostering setting when replacing machines by personnel. When specifying work characteristics and personnel environment, our main concern will be

describing problems and distinguishing between them. One confusing point may be that soft constraints actually are optimisation objectives. We will not move them there because we think that they inherently describe personnel or work. We will leave room for indicating the presence of such soft constraints in the description of the optimisation objectives.

In the rest of this section, we develop this categorisation and give examples of how it may be used. An overview is given in Table 1 and the following subsections discuss the various components in depth.

As can be seen in the table, the description of a nurse rostering problem is decomposed in three main parts or categories. These categories often refer to constraints, but also describe structural properties and optimisation goals.

Category α refers to the description of the personnel environment with information about the number of staff, their skills and their availabilities. Various constraints on the individual workload and the organisation of work over periods of time, like nights and weekends, can be defined. We decided to put the skill characteristic in the α category although one could argue that it also fits in the β category. In the detailed definition of the problem class, dependencies between skills may be given. For example, some problems allow members of personnel to be replaced by higher qualified people while other instances define strictly distinct skill groups. The number of personnel is a size parameter and is not explicitly present in the categorisation, although of course it will be an important attribute of a specific instance. Tables 2 to 5 include examples for each of the α characteristics.

Category β refers to work characteristics describing the actual services to be delivered and the time structure. The latter structure is an essential ingredient for rostering problems. It impacts both the way the demand for services can be specified and the working conditions for the personnel. Again detailed definitions can be given for a specific class of problem instances, and, as in category α , we leave out size parameters such as the planning horizon. Examples of each β characteristic are presented in the Tables 6 and 7.

Category γ refers to optimisation objectives and allows the distinction between various modes of decision support. Rostering problems usually display a significant number of soft constraints of which a minimal violation should be pursued. Although the constraints are actually defined in α , the objective of minimising their violations belongs to the γ category.

We must stress that the model in Table 1 is not meant to be a fixed and rigid model. The categories have been defined after careful analysis of the nurse rostering literature. They can be detailed further if new problem properties come up. As we will demonstrate in Sect. 7, the categorisation allows for an unambiguous definition of a specific problem

Table 1 Classification of nurse rostering problems

	Personnel constraints		Skill interactions	
α Personnel environment	A	Availability	2, 3, ...	Fixed number
	S	Sequences	N	Variable number
	B	Balance	I	Individual skill definitions
	C	Chaperoning		
	Coverage constraints		Shift type	
β Work characteristics	R	Range	2, 3, ...	Fixed number
	T	Time Intervals	N	Variable number
	V	Fluctuating	O	Overlapping
	Objective		Mode	
γ Optimisation objective	P	Personnel constraints	M	Multi-objective
	L	Coverage constraints		
	X	Number of personnel		
	R	Robustness		
	G	General		

and hence can be used for comparing problems. When comparing or recommending approaches, the most reliable results will be obtained when looking for the algorithms corresponding to problems in the same category.

In the coming sections we present detailed discussions of each of the categories α , β and γ .

4 Constraints applying to personnel—category α

The constraint categories *A* and *S* in Table 1 refer to single nurses or, more generally, to single members of personnel. Constraint categories $\alpha : B$ and $\alpha : C$ imply interactions between members of personnel. We discuss these two cases separately.

4.1 Categories $\alpha : A$ and $\alpha : S$ constraints applying to individual members of personnel

Constraints on the schedule of a single nurse originate from legal, local and personal regulations. Such elements constrain the shift pattern a specific nurse can be expected to perform. Apart from legal requirements, these constraints are typically negotiated and often influenced by group culture and personal preferences. They result in contracts or work regulations that define a list of *A* and *S* constraints. We call $\alpha : A$ the category of constraints related to minimum and maximum availabilities of the nurses. The wide variety of constraints in this category are evaluated by one evaluation procedure, as was shown in Burke et al. (2001b). Category $\alpha : S$ groups the constraints handling sequences, like minimum or maximum work stretches. They require an evaluation procedure different from the one applied for $\alpha : A$

constraints. We refer to Burke et al. (2001b) for a more formal definition of the different constraint classes and their evaluation.

In practice, ward managers will try complying as much as possible and generate schedules that do not violate these constraints strongly. Although a solution that does not contain any constraint violation often turns out not to be achievable, support from a tool that tries approximating the ideal situation is appreciated.

The notation can be expanded in many ways. This expansion is open and researchers may be guided by their specific modelling needs. As an example, Table 2 and 3 contain a list of constraints that we take into account in the expansion of $\alpha : A$ and $\alpha : S$, respectively. While not being exhaustive, the constraints in this list and their generalisations are, to our experience, sufficient for a broad range of cases. We expand on the short descriptions in Table 2 and 3.

4.2 Categories $\alpha : B$ and $\alpha : C$ constraints applying to groups of people

Apart from social relations, feelings of fairness in the group and contractual obligations, individual working schedules strongly influence personal well-being and even long term health conditions. Hence the importance of these constraints as has been recognised by many authors, such as Burke et al. (2001a), Bellanti et al. (2004), Chiarandini et al. (2000). The balance constraints in Table 1 restrict the variation of the workload per nurse, including the degree of violation of his/her preferences and work regulations as far as these are legal. They can be expressed as a maximum value for the difference of the quantities measuring this workload, or as a limitation on the statistical properties of the distribution as

Table 2 Category $\alpha : A$

$\alpha : A$	Availabilities
a	<p>General availabilities</p> <p>This constraint expresses the availability of a nurse at certain points in time. Nurses may not be available at specific dates or for specific shifts. These availability constraints are not part of the contract.</p> <ul style="list-style-type: none"> – short term absence or unavailability, – long term unavailability, e.g. maternity leave, – holiday period, – etc.
b	<p>Number of assignments</p> <p>The total number of assignments to a nurse's roster must be between certain margins, e.g. assignments during reference periods such as days, particular days of the week, weekends, weekdays, weeks, bank holidays, the entire planning horizon. It is important to note that the reference period need not be connected (e.g. weekends). Minima, maxima or both may be given. These availability constraints are always related to the person's contract.</p> <ul style="list-style-type: none"> – full time contracts, e.g. minimum 18 and maximum 22 assignments per month, – part time contracts, e.g. minimum 12 and maximum 15 assignments per month, – weekend contracts, e.g. for people who only perform weekend work, – weekday contracts, for people who never carry out weekend work, – etc.
c	<p>Number of shift types out of a set</p> <p>This constraint bears close resemblance to the previous one and can be expressed similarly. It controls the number of assignments to particular shift types over the reference periods defined in $\alpha : A b$.</p> <ul style="list-style-type: none"> – night contracts, for people who only perform night shifts, – day contracts, for people who never perform night shifts, – Wednesday afternoons off, in case of a contractual agreement about such unavailability, – Minimum/maximum number of night shifts, – etc.
d	<p>Number of hours</p> <p>The number of hours worked over a reference period should be within given limits.</p> <ul style="list-style-type: none"> – Minimum/maximum number of hours per day/week/month, etc., – Maximum number of overtime hours per day/week/month/year, – Minimum/maximum number of hours on bank holidays. – etc.
e	<p>Preferences</p> <p>The constraint expresses individual preferences to perform a (particular) shift on a particular date or to be off work on a particular shift/day/period. The constraints are not part of the contract.</p> <ul style="list-style-type: none"> – request for a day off at a particular date, – request for one or more shifts off at a particular date, – request to carry out a particular shift type at a particular date, – request to be on duty at a particular date (no matter for which shift type), – etc.

such. The chaperoning constraints limit the rosters of two people simultaneously. Table 4 and 5 present examples of the balance and chaperoning constraints.

4.3 Categories $\alpha : 2, 3, \dots, N$ and I definition of skills

Table 1 refers to different classes of skill definition.

A situation in which all the members of personnel are qualified to perform the same duties, is the default case (single skill). When a problem considers a set of completely independent skills, for which the workload is expressed separately, then the default case also holds. Indeed, such a problem boils down to a set of smaller independent subproblems,

Table 3 Category $\alpha : S$

$\alpha : S$	Sequences
a Maximum or minimum number of consecutive assignments.	<p>The aim is to prevent nurses from working continuously during too long periods, or to prevent that one single assignment disturbs a healthy pattern. The reference period need not be connected. It can be equal to a particular set of shift types, days, weekends, etc.</p> <ul style="list-style-type: none"> – Minimum/maximum number of consecutive working days. Often, work stretches that are shorter than 2 days or longer than 6 days are not acceptable, – Minimum/maximum number of consecutive days with particular shifts assigned. Sometimes a general constraint over all the shift types is defined but in many real world situations, more detail is required. The acceptable work stretch for consecutive morning shifts may be between 2 and 6, while for the night shifts it may be between 2 and 4, for example, – Minimum/maximum number consecutive working weekends, – etc.
b Maximum or minimum breaks between working periods	<p>The constraint prevents breaks that are too short to be healthy as well as very long breaks. Night shift sequences e.g. may require a longer break to recover.</p> <ul style="list-style-type: none"> – Minimum number of hours between two consecutive assignments. It should be at least a number of hours long. This time often depends on the shift types that are involved (e.g. minimum 11 hours rest required between day shift types; more rest is required when one of the considered shift types is a night shift), – Minimum/maximum number of consecutive free days. The minimum is often two. Particular contracts with long work stretches of night shifts require longer breaks. The maximum number is used to prevent people from spending all the holidays that their contract allows within one single stretch, – etc.
c Compatibility constraints on shift types	<p>Legislation does not allow certain shift types to follow directly after some other.</p> <ul style="list-style-type: none"> – No night shift after a free day (because the necessary information flow between personnel needs to be secured), – Forward rotation, e.g. no early shift after a late or night shift, – etc.
d Other	<p>We refer to particular sequence constraints such as cyclical patterns, weekend work, etc.</p> <ul style="list-style-type: none"> – Cyclical patterns. That very particular constraint defines a pattern that generates constraints by mapping it sequentially onto the planning period. It determines on which days a (not necessarily specified) shift type should be assigned or not, – Minimum number of assignments during weekends, if there are assignments at all (for example, assignments on both Saturday and Sunday). It avoids that many people's weekends are occupied with only one shift assignment, – No night shift before a free weekend, – etc.

each dealing with one single skill. Such problems essentially do not differ from the default single skill problems.

When interactions are allowed between skills, the situation is different. Sometimes hierarchical substitution between different skill categories is tolerated. A very common example is the following:

- People belong to one of the three skill categories: carers, regular nurses and head nurses. Regular nurses are entitled to replace carers and the head nurses could replace people from both other classes, whenever the availability in these classes would be too low.

In such cases, the number of skill classes is 2, 3, ... for a fixed number of classes or N for a variable number. The

Table 4 Category $\alpha : B$ $\alpha : B$ Balance

The balance constraints explicitly affect the rosters of multiple people at the same time. It becomes more complex when the constraint affects people with different contracts. In that case, the difference should be relative to the expected workload for each of the contracts. The constraint can be defined over different reference periods (e.g. a week, month, planning horizon, semester, year, etc.).

- Balance the number of hours worked over all the personnel, such that the difference between the workload for different members of personnel is lower than, e.g. 8 hours, 5%, etc.,
- Balance the number of assignments, such that the difference over a reference period is, for example, less than two,
- Balance the number of assignments on bank holidays,
- Balance the number of night shifts,
- etc.

Table 5 Category $\alpha : C$ $\alpha : C$ Chaperoning

The chaperoning constraint requires a specific nurse always, or at certain time units, to be accompanied by another nurse for reasons related to communication, specific task issues, car pooling, etc. Another usage of the chaperoning category allows one to express that people should not work at the same time. The constraint takes the rosters of multiple nurses into account.

- Two particular nurses should always be on duty at the same time,
- A nurse should not be on duty, unless another one is working as well (for example a trainee and a supervisor),
- Two particular nurses should never be on duty at the same time,
- etc.

personnel from a higher skill class can cover for the lower skills in case of personnel shortage. Note that we do not explicitly label hierarchical substitution because it is a more common procedure than other skill interactions. The presence of a number or N is sufficient.

Sometimes the interactions between different skills are more complex than hierarchical replacements. We label these problems I , which stands for individual skill definitions. It occurs when a person's characterisation is based on a complex mix of qualifications. In that case, substitution among skill classes purely depends on the individual nurses. The latter category often compares better to real world situations where, e.g. the eldest, most experienced or best trained 'lower' category nurse will be best suited to replace a superior one.

5 Constraints concerning the workload—category β

Constraints in these categories describe the workload for the team. Categories $\beta : R, T$ and V in Table 1 describe coverage constraints. These typically correspond to the demand for personnel support at each point in time. They express numbers of personnel needed and specify the required skill as in category α . Categories $\beta : 1, 2, 3, N$ and O describe the shift structure. This category is very diverse. Virtually every

organisation has implemented its own structure to meet the needs of its clients and fit the expectations of its personnel.

5.1 Categories $\beta : R, T$ and V constraints describing the coverage

The required coverage is the number of personnel of a specific skill needed at each time interval or for each shift. In the review paper (Burke et al. 2004), problems surged with time intervals and fluctuating coverage specifications. The review mentions one paper and one working paper with time interval coverage. Since then, several authors have been addressing this issue. We include their papers in our review in Table 8. We consequently distinguish the constraints that are presented in Table 6.

These constraints are present in the nurse rostering problems described in the literature. In each of the above situations coverage can be modelled as a hard or a soft constraint (Sect. 2.1).

The label $\beta : T$ denotes that the coverage constraints are expressed over time intervals. Assuming that the time intervals do not correspond to the durations of the shifts, a conversion is required in order to solve the problem with the defined shifts. Category $\beta : R$ covers the problems in which

Table 6 Category $\beta : R, T$ and V

$\beta : R, T$ and V	Workload
<i>R</i> Range	The required coverage is bounded from above and/or from below at each time unit. This may be a hard or a soft constraint. – minimum 3 and maximum 4 early shifts, – minimum 2 and maximum 4 late shifts, – minimum 1 and maximum 2 night shifts, – etc.
<i>T</i> Time intervals	By default, coverage constraints are expressed per shift type. This may not be appropriate, in which case time intervals can be used. The time intervals can be as small as a number of minutes but they usually are related to the start and end time of the shift types. – 7 am–8 am: at least 2 regular nurses, at least 3 carers, – 8 am–10 am: at least 4 regular nurses, at least 2 carers, – 10 am–12 am: at least 3 regular nurses, at least 2 carers, at least one head nurse, – etc.
<i>V</i> Fluctuating	By default, coverage constraints are constant over time. If this is not the case, the label V and a detailed description are given. In hospitals, fluctuating demands are common, due to the activities that are assigned to particular days of the week (e.g. surgery on a set day has an impact on the level of activity in the wards on subsequent days). – Sunday 31 December: minimum 2 and maximum 2 early shifts Sunday 31 December: minimum 2 and maximum 2 late shifts Sunday 31 December: minimum 1 and maximum 1 night shifts Monday 1 January: minimum 2 and maximum 2 early shifts Monday 1 January: minimum 2 and maximum 2 late shifts Monday 1 January: minimum 1 and maximum 1 night shifts Tuesday 2 January: minimum 4 and maximum 5 early shifts Tuesday 2 January: minimum 3 and maximum 4 late shifts Tuesday 2 January: minimum 2 and maximum 3 night shifts Wednesday 3 January: minimum 2 and maximum 3 early shifts Wednesday 3 January: minimum 2 and maximum 3 late shifts Wednesday 3 January: minimum 2 and maximum 3 night shifts etc.

coverage constraints are satisfied when the assignments are within a range of minimum and maximum acceptable coverage. Absence of the label $\beta : V$ implies that the coverage constraints are constant over time.

The impact of these constraints on the difficulty and the attainable degree of optimisation has not been systematically investigated. It is clear that when the demands are raised, the same team of personnel may not be sufficient. At the other end of the spectrum, too low demands put the team in a situation of work shortage. That could be a problem in case many of the $\alpha : A$ constraints are expressed as minimum requirements. The impact of a fluctuating demand is another issue as it may make it hard to meet the sequence constraints.

5.2 Categories $\beta : 2, 3, \dots, N$ and O constraints describing the shift structure

In the literature, various shift structures can be found (Burke et al. 2004). Authors typically study either a set of problems with a fixed and usually small number of shifts or a set over which this number varies and is usually larger.

6 Optimisation objective—category γ

The nurse rostering literature presents numerous objectives and optimisation modes. We identify the objectives $\gamma : P, L, X, R$ and G that determine the goal function, and the category $\gamma : M$ when the problem is set as a multi-objective decision support problem.

Table 7 Category $\beta : 2, 3, \dots, N$ and O

$\beta : 2, 3, \dots, N$ and O		Shift structure
2, 3, ...	Fixed number of shifts	The default case (with only one shift type) applies to organisations that express requirements in terms of daily attendance of people, e.g. those that provide only nine to five services. A fixed number of different shift types (notation: 2, 3 or a higher number) occur frequently. Three shifts are sufficient for those organisations that use early, late or early, late and night shifts. Note that a ‘dummy shift type’ is quite common in papers. It refers to ‘no assignment’ on the particular day. We do not categorise it as an extra shift type. The required coverage is bounded from above and/or from below at each time unit. This may be a hard or a soft constraint.
N	Variable number of shifts	The category N encompasses all cases in which the number of shifts is variable within the set of problem instances.
O	Overlap	O can be added as a label to $\beta : 2, 3, \dots$, and N to indicate that shift types are not fully separated in time. This is important if the organisation needs communication at the time of the team switch. It did, in our applications (Burke et al. 2006), have an impact on the algorithm performance in case of time interval requirements ($\beta : T$).

6.1 Categories $\gamma : P, L, X, R$ and G

$\gamma : P$ and $\gamma : L$ denote that at least some violations of personnel constraints (category α) and of coverage constraints (category β) should be minimised. In case the notation of a nurse rostering problem contains no $\gamma : P$, all the constraints of the α category, if there are any, should be fully satisfied. Similarly, if the notation contains no $\gamma : L$, all the constraints of the β category should be treated as hard constraints. Note that many of these problems still allow for considerable flexibility if the personnel constraints or the coverage constraints are defined in terms of intervals.

Nurse rostering problems that try to minimise the number of personnel belong to the category $\gamma : X$. Often, the objective appears to be minimising the number of additional members of staff that are seconded internally from another ward or hired through an external agency.

The objective $\gamma : R$ indicates that robustness criteria apply in the evaluation of a roster and/or that rescheduling operations try to minimise the number of disturbances in an existing roster. As concluded in Burke et al. (2004), robustness is an important issue in nurse rostering. Since then, a few papers explicitly addressed rescheduling. Examples are:

- minimise the difference between the new and the original roster,
- enable minimal modifications to a roster,
- generate rosters that are stable against minimal disruptions.

We categorise all the other objectives in $\gamma : G$. It is clear that there is a large variety of optimisation objectives and we do not want to introduce labels for every single objective that occurs in one paper. Some of the objectives in $\gamma : G$

may evolve into separate labels if they are recognised to be generally important by the community. Some examples are:

- minimise the personnel cost (e.g. by minimising expensive overtime, minimising assignments to tasks that should be carried out by other people, etc.),
- minimise the difference between the actual assignments and the personnel capacity, which allows for maximising the usage of available personnel,
- etc.

In $\gamma : G$, we capture objectives that apply to particular schedules and that are not covered by any of the other labels. The possibility exists to extend the notation with new objectives when they have relevance for a class of nurse rostering problems.

6.2 Category $\gamma : M$

Although most nurse rostering problems are multi-objective problems by nature, they are often modelled by a weighted cost function that deals with all the objectives at once. In the latter case, the quality of a solution is expressed by a single number. In some problem settings, this is not sufficient. The user may want to see the effect on other objectives of an increased quality in one specific direction. In this case, the problem itself is a multi-objective one and it will be labelled $\gamma : M$.

7 Examples

In order to illustrate the $\alpha|\beta|\gamma$ notation introduced in Sect. 3, we present some problems that were described in the recent literature. The problems fit into the introduced categories,

Table 8 Classification of recent research papers

Notation	References	Approach
AN 2(3) P	Aickelin et al. (2007)	estimation of distribution algorithm (EDA), ant-mining
AN 2(3) P	Aickelin and Li (2007)	EDA, Bayesian network
ASBN V2 PL	Azaiez and Al-Sharif (2005)	linear programming
AS(N) N P(L).XR	Bard and Purnomo (2005c, 2006)	integer programming (IP), branch and price
ASBN T3 PX	Bard and Purnomo (2005a, 2005b)	column generation (IP and heuristics)
AS1 PX	Bard and Purnomo (2007)	integer programming + heuristic
ASN R3 P	Beddoe et al. (2009), Beddoe and Petrovic (2006, 2007)	case based reasoning
AS RN P	Beste et al. (2007)	tabu search
ASB RVNO PLR	Bilgin et al. (2008)	local search
ASB RTNO PL	Bourdais et al. (2003)	constraint programming
ASBC RVNO PR	De Causmaecker and Vanden Berghe (2003)	metaheuristics
ASBC RVNTO PR	Burke et al. (2006)	metaheuristics
ASBCN NO P	Brucker et al. (2010), Burke et al. (2010)	decomposition and local search, scatter search
AS N P	Burke et al. (2008)	heuristic ordering and variable neighbourhood search
ASN R P	Chun et al. (2000)	constraint propagation and heuristics
S RV2(3) P	Dias et al. (2003)	genetic algorithm, tabu search
AS R3 PLM	Duenas et al. (2003)	genetic algorithm
A2(3) 3(N) P	Parr and Thompson (2007)	simulated annealing with SAWing and noising
AS VNO PLG	Gutjahr and Rauner (2007)	ant colony optimisation
AS 3 PL	Maenhout and Vanhoucke (2008)	genetic algorithm
AS N PR	Moz and Pato (2007)	genetic algorithms
AS 3 PL	Osogami and Imai (2000)	local search
S 2 P	Ozcan (2005)	memetic algorithms
ASB R3V M	Topaloglu and Selim (2007)	fuzzy multi-objective programming
AS V3 PL	Wong and Chun (2004)	constraint programming

which does not mean that they cover the complete category. Each notation in Table 8 is based on the problem description in the corresponding paper. It is worth noting that the majority of papers does not refer to or provide public datasets. The text is the only reference. In addition, we briefly report on the general approach applied in that paper.

We pay special attention to the problem instances that have been made publicly available. The nurse rostering datasets provided in Burke et al. (2008) and Brucker et al. (2010) belong to the category $AS|N|P$. Those that are made available in Bilgin et al. (2008) are categorised as $ASBI|RVNO|PL$. The NSPLib instances (Vanhoucke and Maenhout 2007) can be classified as $AS|3|PL$. Maenhout and Vanhoucke (2008) did very valuable research but the number of constraints that they deal with in e.g. the S category is rather small. They did not undertake any systematic study across the parameter domains in order to reveal their real impact on the hardness.

The comparison tables in Burke et al. (2004) lead straightforwardly to the new notation. We decided to only present the more recent papers. The paper by Bourdais et al. (2003) presents an alternative view on the categories in the present paper. Their ‘rule categories’ (that are based on related literature and on real data) can be linked to our notation as follows:

- DEM: minimum coverage constraints (β),
- AVA: availability constraints including holiday requests (α),
- DIS: fair distribution of work among people (α),
- ERG: a heterogeneous category of constraints related to patterns of work, shift sequences, etc. (α).

For each of these rules, Bourdais et al. (2003) give further refinements that are described in words. Demand and availability constraints are always considered hard but some of the distribution and ergonomic constraints may be soft. When describing the problems in Aickelin et al. (2007), Aickelin and Li (2007), Parr and Thompson (2007), we put 2(3) for the shift types because they do not deal with the night shifts and the 2 day type shifts at the same time. The notation for the problems in Bard and Purnomo (2005c, 2006) contains an N between brackets because both the default problems with one single shift, and the problems with a variable number of hierarchical skills are modelled.

As can be expected, papers from the same authors elaborate on the same or similar problems. Table 8 is far from covering all possible combinations of labels. We do indeed not expect all combinations to occur in practice and it would in fact be interesting to position the set of studied problems within the space of all the possibilities. More interestingly, the table highlights similarity of problems in independent studies, for example (Osogami and Imai 2000) and (Maenhout and Vanhoucke 2008); (Beddoe et al. 2009; Beddoe and Petrovic 2006, 2007) and (Chun et al. 2000).

8 Conclusion

We propose a notation for the classification of a set of personnel rostering problems. It builds on earlier attempts to characterise nurse rostering problems (Burke et al. 2004) and adopts the $\alpha|\beta|\gamma$ notation for scheduling. We developed a number of categories originating from the problems and approaches that we found in the literature as well as in our own nurse rostering research. The categories are slots for *problem* features. The categories exhibit a model of the problem domain, where we tried to catch features with relevance for the hardness of the problem. Problems sharing these features will end up in the same category and will be candidates for closer investigation of their differences. Where these differences turn out to have an important impact on the solution approaches and can be attributed to clearly identifiable problem related features, they will become candidates for refinement of the categorisation. The features allow classifying all the problems from the papers in Burke et al. (2004). Furthermore, we presented the classification of 30 papers that appeared over the last years.

We decided to concentrate on problem related features and not to address solution approaches, procedures or abstract models. An alternative focus could be to concentrate on these more generic features. The present classification could be extended with more quantitative measures of complexity, which would probably require a higher level of abstraction. This is an important subject for further investigation of a more generic and theoretical nature. The current framework concentrates on the real world context in which this domain evolves. It enables classifying problems based on real world observable properties. Our hope is that it will serve as a useful tool to bring order to the myriad of problem definitions and applications and as a starting point for further systematic research.

With the publication of this notation, we hope to start a discussion resulting in an agreed upon progressively refined classification system in which researchers and developers can place their contributions and locate related approaches and results. With this goal in mind we opened up a website at <http://ingenieur.kahosl.be/vakgroep/it/nurse/archive.htm>.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Aickelin, U., & Li, J. (2007). An estimation of distribution algorithm for nurse scheduling. *Annals of Operations Research*, 155(1), 289–309.
- Aickelin, U., Burke, E. K., & Li, J. (2007). An estimation of distribution algorithm with intelligent local search for rule-based nurse rostering. *Journal of the Operational Research Society*, 58(12), 1574–1585.

- Azaiez, M. N., & Al-Sharif, S. S. (2005). A 0-1 goal programming model for nurse scheduling. *Computers and Operations Research*, 32(3), 507–491.
- Bard, J. F., & Purnomo, H. W. (2005a). A column generation-based approach to solve the preference scheduling problem for nurses with downgrading. *Socio-Economic Planning Sciences*, 39, 139–213.
- Bard, J. F., & Purnomo, H. W. (2005b). Preference scheduling for nurses using column generation. *European Journal of Operational Research*, 164, 510–534.
- Bard, J. F., & Purnomo, H. W. (2005c). Hospital-wide reactive scheduling of nurses with preference considerations. *IIE Transactions*, 37, 589–608.
- Bard, J. F., & Purnomo, H. W. (2006). Incremental changes in the workforce to accommodate changes in demand. *Health Care Management Science*, 9, 71–85.
- Bard, J. F., & Purnomo, H. W. (2007). Cyclic preference scheduling of nurses using a Lagrangian-based heuristic. *Journal of Scheduling*, 10(1), 5–23.
- Beddoe, G. R., & Petrovic, S. (2006). Selecting and weighting features using a genetic algorithm in a case-based reasoning approach to personnel rostering. *European Journal of Operational Research*, 10(1), 649–671.
- Beddoe, G. R., & Petrovic, S. (2007). Enhancing case-based reasoning for personnel rostering with selected tabu search concepts. *Journal of the Operational Research Society*, 58(12), 1586–1598.
- Beddoe, G., Petrovic, S., & Li, J. (2009). A hybrid metaheuristic case-based reasoning system for nurse rostering. *Journal of Scheduling*, 12(2), 99–119.
- Bellanti, F., Carello, G., Della Croce, F., & Tadei, R. (2004). A greedy-based neighbourhood search approach to a nurse rostering problem. *European Journal of Operational Research*, 153(1), 28–40.
- Bester, M. J., Nieuwoudt, I., & Van Vuuren, J. H. (2007). Finding good nurse duty schedules: a case study. *Journal of Scheduling*, 10, 387–405.
- Bilgin, B., De Causmaecker, P., Rossie, B., & Vanden Berghe, G. (2011). Local search neighbourhoods for dealing with a novel nurse rostering model. *Annals of Operations Research*, accepted.
- Bilgin, B., De Causmaecker, P., Rossie, B., & Vanden Berghe, G. (2008). Local search neighbourhoods to deal with a novel nurse rostering model. In *Proceedings of the 7th international conference on practice and theory of automated timetabling*, Montreal, August 2008.
- Bourdais, S., Galinier, Ph., & Pesant, G. (2003). Hibiscus: A constraint programming application to staff scheduling in health care. In *Lecture notes in computer science. CP 2003* (Vol. 2833, pp. 153–167).
- Brucker, P. (2004). *Scheduling algorithms* (4th ed.). Berlin: Springer.
- Brucker, P., Drexl, A., Möhring, R., Neumann, K., & Pesch, E. (1999). Resource-constrained project scheduling: Notation, classification, models, and methods. *European Journal of Operational Research* 112, 3–41.
- Brucker, P., Burke, E. K., Curtois, T., Qu, R., & Vanden Berghe, G. (2010). Adaptive construction of nurse schedules: A shift sequence based approach. *Journal of Heuristics* 16(4), 559–573.
- Burke, E. K., Cowling, P., De Causmaecker, P., & Vanden Berghe, G. (2001a). A memetic approach to the nurse rostering problem. *Applied Intelligence*, 15, 199–214. Special issue on simulated evolution and learning.
- Burke, E. K., De Causmaecker, P., Petrovic, S., & Vanden Berghe, G. (2001b). Fitness evaluation for nurse scheduling problems. In *Proceedings of the congress on evolutionary computation (CEC2001)* (pp. 1139–1146), Seoul, Korea, May 27–30, 2001. New York: IEEE Press.
- Burke, E. K., De Causmaecker, P., Vanden Berghe, G., & Van Landeghem, H. (2004). The state of the art of nurse rostering. *Journal of Scheduling*, 7(6), 441–499.
- Burke, E. K., De Causmaecker, P., Petrovic, S., & Vanden Berghe, G. (2006). Metaheuristics for handling time interval coverage constraints in nurse scheduling. *Applied Artificial Intelligence*, 20(9), 743–766.
- Burke, E. K., Curtois, T., Qu, R., & Vanden Berghe, G. (2007). *A time pre-defined variable depth search for nurse rostering* (Technical report). University of Nottingham.
- Burke, E. K., Curtois, T., Post, G., Qu, R., & Veltman, B. (2008). A hybrid heuristic ordering and variable neighbourhood search for the nurse rostering problem. *European Journal of Operational Research*, 188, 330–341.
- Burke, E. K., Curtois, T., Qu, R., & Vanden Berghe, G. (2010). A scatter search approach for the nurse rostering problem. *Journal of the Operational Research Society*, 61, 1667–1679.
- Chiarandini, M., Schaerf, A., & Tiozzo, F. (2000). Solving employee timetabling problems with flexible workload using tabu search. In E. K. Burke & W. Erben (Eds.), *Proceedings of the 3th international conference on the practice and theory of automated timetabling*, PATAT (pp. 298–302). Konstanz, Germany, August 2000.
- Chun, A. H. W., Chan, S. H. C., Lam, G. P. S., Tsang, F. M. P., Wong, J., & Yeung, D. W. M. (2000). Nurse rostering at the hospital authority of Hong Kong. In *Proceedings of the 17th national conference on AAAI and 12th conference on IAAI* (pp. 951–956).
- De Causmaecker, P. (2008). Towards a reference model for timetabling and rostering. In E. K. Burke & M. Gendreau (Eds.), *The practice and theory of automated timetabling, proceedings of the 5th international conference*, Montreal, August 2008.
- De Causmaecker, P., & Vanden Berghe, G. (2003). Relaxation of coverage constraints in hospital personnel rostering. In E. K. Burke & P. De Causmaecker (Eds.), *Lecture notes in computer science: Vol. 2740. Practice and theory of automated timetabling* (pp. 129–147). Berlin: Springer.
- De Causmaecker P., & Vanden Berghe G. (2011). Towards a reference model for timetabling and rostering. *Annals of Operations Research*, accepted.
- Dias, T. M., Ferber, D. F., de Sousa, C. C., & Moura, A. V. (2003). Constructing nurse schedules at large hospitals. *International Transactions in Operational Research* 10, 245–265.
- Duenas, A., Mort, N., Reeves, C., & Petrovic, D. (2003). Handling preferences using genetic algorithms for the nurse scheduling problem. In G. Kendall, E.K. Burke, & S. Petrovic (Eds.), *Proceedings of the 1st multidisciplinary international conference on scheduling: Theory and applications (MISTA)* Nottingham, UK (pp. 180–196).
- Dyckhoff, H. (1990). A typology of cutting and packing problems. *European Journal of Operational Research*, 44, 145–159.
- Graham, R. L., Lawler, E. L., Lenstra, J. K., & Rinnooy Kan, A. H. G. (1979). Optimization and approximation in deterministic sequencing and scheduling: a survey. *Annals of Discrete Mathematics*, 4, 287–326.
- Gutjahr, W. J., & Rauner, M. S. (2007). An ACO algorithm for a dynamic regional nurse-scheduling problem in Austria. *Computers & Operations Research*, 34(3), 642–666.
- Herroelen, W., Demeulemeester, E., & De Reyck, B. (1999). A classification scheme for project scheduling. *International series in operations research and management science: Vol. 14* (pp. 77–106). Project scheduling—recent models, algorithms and applications. Kluwer Academic Publishers.
- Lawler, E. L., Lenstra, J. K., Rinnooy Kan, A. H. G., & Shmoys, D. B. (1993). *Handbook in operations research and management science*. Vol. 4. Amsterdam: North Holland.
- Lodree, E. J. Jr., Geiger, C. D., & Jiang, X. (2009). Taxonomy for integrating scheduling theory and human factors: Review and research opportunities. *International Journal of Industrial Ergonomics*, 39(1), 39–51.

- Maenhout, B., & Vanhoucke, M. (2008). Comparison and hybridization of crossover operators for the nurse scheduling problem. *Annals of Operations Research*, 159(1), 333–353.
- Messelis, T., Haspeslagh, S., Bilgin, B., De Causmaecker, P., & Vanden Berghe, G. (2009). Towards prediction of algorithm performance in real world optimisation problems. In *Proceedings of the 21st Benelux conference on artificial intelligence* (Vol. 21, pp. 177–183), BNAIC. Eindhoven, October 2009.
- Moz, M., & Pato, M. V. (2007). A genetic algorithm approach to a nurse rostering problem. *Computers & Operations Research*, 34(3), 667–691.
- Osogami, T., & Imai, H. *Classification of various neighborhood operations for the nurse scheduling problem* (Technical Report 135). The Institute of Statistical Mathematics.
- Ozcan, E. (2005). Memetic algorithms for nurse rostering. In *Lecture notes in computer science: Vol. 3733. Computer and information sciences—Proceedings ISCIS* (pp. 482–492).
- Parr, D., & Thompson, J. M. (2007). Solving the multi-objective nurse scheduling problem with a weighted cost function. *Annals of Operations Research*, 155(1), 279–288. Special issue on personnel scheduling and planning.
- Pinedo, M. (2002). *Scheduling, theory, algorithms, and systems* (2nd ed.). New York: Prentice Hall.
- Siferd, S. P. & Benton, W.C. (1992). Workforce staffing and scheduling: Hospital nursing specific models. *European Journal of Operational Research*, 60, 233–246.
- Topaloglu, S. & Selim, H. (2007). Nurse scheduling using fuzzy multiple objective programming. In H. G. Okuno & M. Ali (Eds.), *Lecture notes in artificial intelligence: Vol. 4570. IEA/AIE 2007* (pp. 54–63). Berlin: Springer.
- Vanhoucke, M., & Maenhout, B. (2007). NSPLib—a nurse scheduling problem library: a tool to evaluate (meta-)heuristic procedures. In S. Brailsford & P. Harper (Eds.), *Operational research for health policy: making better decisions, proceedings of the 31st annual meeting of the working group on operations research applied to health services* (pp. 151–165).
- Vanhoucke, M., & Maenhout, B. (2009). On the characterisation and generation of nurse scheduling problem instances. *European Journal of Operational Research*, 196(2), 457–467.
- Wäscher, G., Haussner, H., & Schumann, H. (2007). An improved typology of cutting and packing problems. *European Journal of Operational Research*, 183, 1109–1130.
- Wong, G. Y. C. & Chun A.H.W. (2004). Constraint-based rostering using meta-level reasoning and probability-based ordering. *Engineering Applications of Artificial Intelligence*, 17(6), 599–610.