

Towards a reference model for timetabling and rostering

Patrick De Causmaecker · Greet Vanden Berghe

Published online: 2 March 2010
© Springer Science+Business Media, LLC 2010

Abstract Timetabling and rostering research often starts from particular real world problems. The last two decennia have seen a large number of papers discussing cases, models and approaches. This large body of publications does not presently constitute a structured domain that provides guidelines for addressing particular problem instances, nor does it allow identifying gaps where new research is needed.

In this paper, we instigate a structured model for timetabling and rostering. We present a basic structure for the integrated staff planning and rostering problem. The model can only result in a valid and efficient tool through the input of multiple disciplines. We give a first formal description for short term nurse rostering. Specific problems are positioned in this model.

Keywords Reference model · Timetabling · Rostering

1 Introduction

Scientists from different domains have studied a large number of instances of timetabling and rostering problems. Among the many application domains figuring as suppliers of case studies, were school and university timetabling, sports timetabling as well as employee timetabling, agenda building and more. Authors often presented a detailed description of a specific case. They developed data models, identified the relevant hard and soft constraints and defined goals for the algorithms which they subsequently developed and tested. Alternatively, they studied historical information in order to find a case base for guiding the decision making. Not surprisingly, given the model complexity and the required level of flexibility,

P. De Causmaecker (✉)

CODES, Computer Science KUL@K and Information Technology, K.U.Leuven, Campus Kortrijk,
Etienne Sabbelaan 53, 8500 Kortrijk, Belgium
e-mail: Patrick.DeCausmaecker@kuleuven-kortrijk.be

G. Vanden Berghe

CODES, Information Technology, KaHo Sint-Lieven, Gebroeders Desmetstraat 1, 9000 Gent, Belgium
e-mail: Greet.VandenBerghe@kahosl.be

the algorithms were almost always heuristics designed according to some metaheuristic paradigm.

In doing so, timetabling researchers have achieved progress along at least three important lines:

- They discovered and studied a large collection of situations for which they created sophisticated models;
- They built an interesting set of large and difficult instances for metaheuristic algorithms, studied algorithmic behaviour and proposed new techniques and devices;
- They exhibited the need for—and partly established—systematic testing and evaluation procedures based on well balanced benchmarks, often drawn from real world situations.

The presently available corpus on timetabling is sufficiently mature to find its way to industrial applications. Examples of such applications have emerged (Bilgin et al. 2008; Ho et al. 2001; Kingston 2005). Two important questions arise:

- Which specific input is needed and in which form should it be presented to allow industry to make optimal use of the present and future scientific state of the art.
- How can we facilitate future developments and position them in the corpus of knowledge.

We present a **multidisciplinary** effort addressing the gap between theory and practice in university course timetabling, as diagnosed by McCollum (2007). This **reference model** allows one to discover opportunities for new developments and supporting communication, interchange and collaboration. We want to improve the input of real world information and distinguish the different levels of abstraction.

There are two main aspects.

1. Input

The domain models and constraint sets that have been described in the nurse rostering literature (Bard and Purnomo 2004; Bourdais et al. 2003; Burke et al. 2001, 2006) have been collected by individual researchers in individual cases. The necessary information was commonly obtained through interviews with people in the field, for example, hospital managers, head nurses and regular nurses. This is true for the hard constraints and regulations, although these could often be inferred from the problem at hand, from available legislation or from contract agreements (Meyer auf'm Hofe 2001; Petrovic et al. 2003). It is even more true for the soft constraints expressing preferences and working comfort expectations of the personnel (nurses, doctors, administrative staff, cleaners) or of the patients. It is hard to bring structure to this material gathered in an ad hoc fashion without a scientific guideline. This guideline can come from other scientific disciplines within the operational research community as well as in humanities. We should investigate what law, social, economical and management sciences may have to say (De Causmaecker et al. 2007).

2. Levels

Although rostering is typically concerned with the most detailed level of decision making, it is not independent of decisions taken at more abstract levels like staffing and planning. The goal for the timetabling activity—what determines a good working roster—cannot be defined independently at the different levels of decision making. After all, these goals are mostly set at the more strategic levels. We will thus not be able to study and define general reference models without taking all these levels into some account. As a consequence, the reference model will ultimately not only describe how timetabling components for the most detailed, short term decision making should be conceived, but

also how decision support tools for the mid and long term should behave and interact. We want to pinpoint the best level for each decision act.

We introduce a **formal model** for nurse rostering allowing to position specific contributions. This work is complementary with the inspiring contribution of Schaerf and Di Gaspero (2007) on measurability and reproducibility in university timetabling research. The **categorisation** system will allow researchers to evaluate and compare approaches. A study of complexity requires accurate description of problems as provided by this categorisation.

Only a few authors have tried to define relevant complexity measures for nurse rostering, e.g. Glass and Knight (2009), Messelis et al. (2009) and Vanhoucke and Maenhout (2009). If the reference model has to provide reliable information about the applicability of a specific technique for a specific problem, such measures are indispensable.

The structure of the paper is as follows. The interactions between the three levels in personnel management (long, mid and short term planning) require a multidisciplinary approach with input from different fields in humanities. This is discussed in Sect. 2. A systematic study of the behaviour of algorithms for personnel rostering problems requires a systematic classification. In Sect. 3, we propose a notation that can support such a classification for the field of nurse rostering. It is our belief that the specific notations presented towards the end of the paper can serve as a starting point for more generic approaches.

2 Reference model

Many papers in the timetabling and rostering literature discuss instances inspired by real world problems (Qu et al. 2009; Burke et al. 2004). As Glass and Knight (2009) remark, the difference between the models that are actually handled and the original problems is often substantial. Essential ingredients are often left out for the sake of the theoretical scope. Kellogg and Walczak (2007) identify differences between daily practice in a hospital and the academic literature explaining the lack of interaction.

The short term level is influenced by decisions taken for the mid and long term (Tien and Kamiyama 1982; Warner 1976). Also for these higher levels, there is a body of literature that investigates the application of automated decision support for real world cases (de Vries 1987; Easton et al. 1992). Part of the problem with the interaction between theory and practice can be explained by the lack of interaction between activities at the different levels. As Glass and Knight (2009) remark, legal and ergonomic criteria play a role in rostering practice. All these observations illustrate the need for input from other than the operational research domain. Such a multidisciplinary effort requires an adequately structured reference model.

Short term planning and scheduling of people's efforts in healthcare, for example, typically requires to collect data concerning the available means and their associated constraints. Part of this information is output of higher level decision making.

This situation is comparable to production planning and scheduling where one traditionally distinguishes the following levels.

- At the management level: product mix, product production trees, machine park layout.
- At the business level: production planning, lotsizing.
- At the lower level these production plans are decomposed in jobs and operations for specific machines.

The situation is further complicated when multiple companies are involved in such a production system. The *supply chain* defines a planning problem that crosses company borders.

The detailed scheduling problem has been studied in great depth (see e.g. Halldorsson et al. 2007). Business systems and planning models at this level have emerged leading to a number of standards. Supply chain management is a discipline of its own. It is clear that this decomposition leads to suboptimal solutions. Only recently, people have started addressing the global optimisation problem. It turns out to be hard to translate the detailed terminology of the short term scheduling to the mid term planning level. Stated differently, it has for a long time been unknown which quantities can be exchanged between the two levels. This situation has led to a specific niche, the domain of *manufacturing execution systems*, which plays the role of an intelligent interface between the business systems and the production line. Systems at this interface typically provide sophisticated logging, tracking and reporting together with some level of detailed scheduling.

In the field of timetabling and more specifically nurse rostering, the situation is comparable but further complicated by the fact that humans are involved. One typically distinguishes decisions to be taken for the long term (years), mid term (months) and short term (weeks).

On the long term, decisions are closely linked to selected strategic options. At this level the existence of a unit or the discontinuation of an activity are decided. Expected costs and revenues are weighted. Information that is needed is at what cost a specific unit can operate, what its relative impact may be on the other activities and whether or not profitable synergies can be expected. The market evolution, expected demand and prospects for the future are taken into account.

Mid term decisions are tactical and influence the operational capacities of the units as they exist. Tuning allocation of staff to expected performance and productivity is the main subject at this level of decision making.

Short term decision making starts in general from a given manpower and demand to produce feasible rosters and working schemes that meet both the demand and the constraints on the individual working schedules.

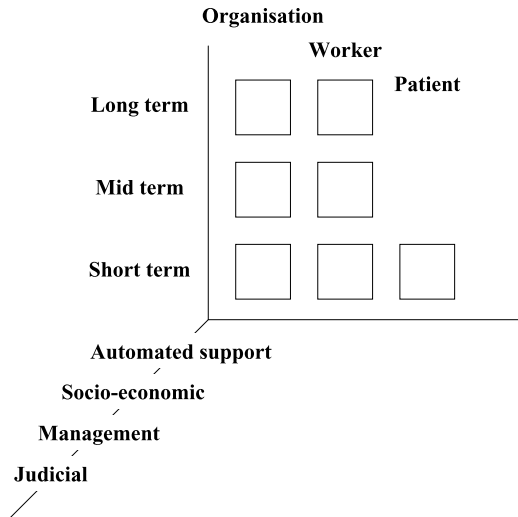
These three levels clearly do not operate independently. Among the subjects to be studied, we mention:

- We need instruments to feed the higher level decision making from below. Experience in building rosters at the short time level should be aggregated to serve as input for the mid and long.
- Decisions at the staffing level should take specific constraints into account. Feasibility of a certain construction should be checked. We need standard measures and simulation tools (e.g. Gemmel et al. 2008).
- Rules for ‘good’ rosters should include, along with legal constraints and local habits or conventions, results from the theory of labour ergonomics (see e.g. Baaijens et al. 2005; Bourdais et al. 2003).

A possible scheme for health care timetabling is presented in Fig. 1. Each rectangle in this scheme is a target for a line of research and/or for software developers. The interactions between these rectangles could become subject of an interdisciplinary study (De Causmaecker et al. 2007). Some of the disciplines are shown on the third axis in Fig. 1.

3 Categorisation

Interesting modelling, solution techniques and real world issues have come up in the recent operational research literature (De Cesco et al. 2008). Some of this work indicates the need for standardisation and classification and for interchangeable formats (Glass and Knight 2009).

Fig. 1 Parties, planning horizons and disciplines

Adoption of such a standardisation by the community could lead to a platform for comparative study of problems and algorithms. Until today however, such adoption has been very limited or non existing. Researchers in the field seem to feel that the domain has not been sufficiently explored for setting a final definition.

3.1 Example models

3.1.1 Languages

A number of authors have presented language constructs to describe timetabling problems. We mention Kingston's STTL (Burke et al. 1997; Kingston 2001) and Özcan's TTML scheme (Özcan 2003). Together with our co-authors, we proposed to devise an ontology for timetabling and used it in a framework (De Causmaecker et al. 2005). The competitions have defined their own formats. These efforts are still far from being commonly used. They were built from the experience of the respective authors and they do offer a systematic way to describe specific classes of timetabling problems.

3.1.2 Modelling devices

Recently, Ranson and Ahmadi (2007) proposed an extensible modelling framework for timetabling. They stress the importance of being language independent and in this sense extend on the previous contributions mentioned above.

The analogue of the linear or integer program in commonly used metaheuristics is a model of the solution space and a neighbourhood structure. Timetabling has a number of general characteristics and much effort has been spent on catching these characteristics in metaheuristic approaches. Looking at the literature, authors often refer to each other's neighbourhood structures (moving timeslots of exams, moving single exams) (Qu et al. 2009). Some authors have tried to develop generic neighbourhoods (e.g. Kempe chains Thompson and Dowsland 1996). An example of transfer of ideas happened with the tiles concept of Kingston (2005) that formed the basis of the pillar neighbourhood in De Causmaecker et al. (2009). Lectures first are grouped into clusters called tiles, after which they are organised

within these tiles. It is intuitive for the users and thus gives them some understanding of the solution method. By allowing users to define the tiles, the algorithm can be given hints as to where to look for good solutions. Although Kingston introduced tiles in the context of high school course timetabling, they are generic as was demonstrated in De Causmaecker et al. (2009) and Ryckbosch et al. (2008) where we modified them for use, respectively in a university course and sports timetabling problem.

The need to apply experience and knowledge from the user is explicitly addressed in case based reasoning approaches (Petrovic et al. 2003) and is an important rationale for the introduction of hyperheuristics (Burke et al. 2003). Acquiring data about the efficiency of individual steps in an algorithm may be used in the systematic study that we envision (Bilgin et al. 2007). An attempt to compare metaheuristic and case based reasoning based on properties from the application environment was presented by Petrovic and Vanden Berghe (2008).

3.2 Study of the hardness of rostering problems

One tool that could help at the level of staffing is a predictor for timetabling problems. Such a predictor could e.g. tell to what extent the constraints can be expected to be satisfied. Given a possible staffing decision, the decision makers could use the tool to judge its feasibility given the current regulations and expected workload. Predictors have been established in other domains where they could often be used in a preprocessing phase to decide which algorithm is the most appropriate for a certain instance (Leyton-Brown et al. 2002; Xu et al. 2007). Haspeslagh et al. (2009), Messelis et al. (2009) recently obtained promising results indicating that such a predictor is feasible.

If we want to obtain generic results along these and similar lines, we need a categorisation system of problems allowing to compare results from different application environments. A classification in nurse rostering was developed by Burke et al. (2004). Here we elaborate on this classification and make it more formal.

3.3 A categorisation of nurse rostering problems

Rostering problems are defined by three elements:

- The available personnel, the people that can be assigned a certain task at certain points in time. We will call this the *personnel environment*.
- The work that has to be performed, the set of tasks and the associated time structure. We will call this the *work characteristics*.
- The conditions to be optimised, the objective that is to be met by the schedule. We call this the *optimisation objective*.

These three elements bear close resemblance to the $\alpha|\beta|\gamma$ categorisation in production scheduling (e.g. Brucker 2004; Graham et al. 1979; Pinedo 2006). We adopt this structure and notation for nurse rostering problems, see Table 1.

- **Personnel environment** (α) The personnel environment distinguishes the categories *A* (availabilities), *S* (sequences), *B* (balance) and *C* (chaperoning). *A* describes the availability of each member of personnel and allows expression of preferences. *S* is used for constraints on the consecutiveness of assignments, e.g. the allowed number of consecutive night shifts. *B* allows the expression of balance constraints on the schedule, e.g. for enforcing some level of fairness. *C* can be used to express the need for combined tasks

Table 1 Classification of nurse rostering problems

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

where one member of personnel must necessarily work together with another one. Independently from these categories, personnel may have individual skills. In the default situation, skills do not apply. A number indicates that the number of possible skills in the problem is fixed, while N denotes a variable number. When each member of personnel appears to have an individual set of qualifications, the label is I .

- **Work characteristics** (β) The work characterised by the demand to be covered (coverage) and the shift structure. We distinguish R (range), V (fluctuating) and T (time interval) types of coverage. The default situation occurs when the coverage is expressed as a number of people required per shift type—skill set. R denotes that it is formulated as a range of feasible coverage, e.g. minimum and maximum coverage. When the coverage varies over time, we call it fluctuating: V . In case the coverage is expressed as a number of people required per time interval, we label it as T . These time intervals do not necessarily correspond to the shift types.

Possible shift structures are strongly connected to work characteristics. We distinguish problems with a fixed number of shift types defined: 2, 3, ... and problems in which the number is variable: N . When the shift structures are overlapping, the label O is used.

- **Optimisation objective** (γ) Possible objectives are P (personnel constraint violation), L (coverage), X (number of personnel) and G (general). These may be multi-objective (M).

3.4 Example

A detailed description of the elements in Table 1, together with the example application of the notation to published problems, is presented in De Causmaecker and Vanden Berghe (2008). In this paragraph we give a simple example. Suppose a ward is working with nurses in three skill categories, say x, y and z , where a nurse from skill type y can work as a nurse in type z (label = 3). It has regulations constraining the number of consecutive assignments, assignments in weekends, minimum length of the breaks and so on (S). Nurses have to work

a minimal number of hours and can ask for some days off (A). The workload for the nurses should be balanced (B). This activates labels A, S, B and skill type 3 in category α .

Furthermore, the organisation uses five shift types (label = 5), which partly overlap in time (O). The coverage is given as a range per shift (R), no time intervals are used in this specification, but the coverage varies per shift (V).

As an objective, the strategy is to always meet the coverage constraints, possibly violating the other constraints at some given cost which must be approximately minimised (P).

The resulting problem can thus be classified as $ASB3|RV5O|P$. Our notational system should at any time be open for extension to accommodate new problem types coming up, and the level of detail should also be set according to the needs of the research at hand. We did keep the size of the problem out of the notation as we think that this should be a parameter in any systematic investigation of a set of rostering instances. What we do propose is to investigate the resulting categories systematically and exhaustively in order to determine the hardness of the instances included. We went through part of the nurse rostering literature since 2004 and were able to classify the problems discussed (De Causmaecker and Vanden Berghe 2008). We furthermore performed a systematic analysis of a specific category. These results are presented in Messelis et al. (2009).

4 Conclusion

In this contribution we joined in and extended on ideas presented by McCollum (2007), Schaerf and Di Gaspero (2007). We presented a reference model for timetabling and rostering.

An essential step in the further development of this model is a multi disciplinary effort in which all levels of personnel planning systems can be integrated.

The paper introduces a classification of problems to be used for study and comparison. We propose a notation for rostering to classify most problems in the recent literature. Once a category has been defined, a systematic study on the influence of the parameters in the category on the hardness of the problem can be undertaken. This allows to draw charts to support the development of automated rostering systems.

Future work includes an extensive study of the different categories, the extension of the notation to other fields of timetabling and the interdisciplinary effort to set up an integrated reference model.

We believe that these ideas can be taken forward to arrive at a comprehensive methodology for the field of timetabling and expect ever more generic approaches to emerge in the literature.

Acknowledgements Kristof Salomez of Vrije Universiteit Brussel, Paul Gemmel of Ghent University, Christine Baaijens and Luc Berlinger of Déhora were excellent partners from respectively law, management and socio-economical sciences in the interdisciplinary discussions. We further thank PhD students Stefaan Haspeslagh, Burak Bilgin and Tommy Messelis for their assistance in checking the model.

References

- Baaijens, C., van Doorne-Huiskes, J., & Schippers, J. (2005). Do dutch employees want to work more or less hours than they actually do? In B. Peper, J. van Doorne-Huiskes, & L. den Dulk (Eds.), *Flexible working and organisational change: the integration of work and personal life*. Cheltenham: Edgar Elgar.
- Bard, J. F., & Purnomo, H. W. (2004). Real-time scheduling for nurses in response to demand fluctuations and personnel shortages. In E. K. Burke & M. Trick (Eds.), *Proceedings of the 5th international conference on the practice and theory of automated timetabling, PATAT* (pp. 67–87). Pittsburgh, August 2004.

- Bilgin, B., Özcan, E., & Korkmaz, E. E. (2007). An experimental study on hyper-heuristics and exam timetabling. In E. K. Burke, & H. Rudová (Eds.), *Lecture notes in computer science: Vol. 3867. Practice and theory of automated timetabling VI* (pp. 394–412). Berlin: Springer.
- Bilgin, B., De Causmaecker, P., Rossie, B., & Vanden Berghe, G. (2008). Local search neighbourhoods to deal with a novel nurse rostering mode. 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: Vol. 2833. CP 2003* (pp. 153–167). Berlin: Springer.
- Brucker, P. (2004). *Scheduling algorithms* (4th ed.). Berlin: Springer.
- Burke, E. K., Pepper, P. A., & Kingston, J. H. (1997). A standard data format for timetabling instances. In *Lecture notes in computer science* (Vol. 1408, pp. 213–222). Berlin: Springer.
- Burke, E. K., Cowling, P., De Causmaecker, P., & Vanden Berghe, G. (2001). A memetic approach to the nurse rostering problem. *Applied Intelligence*, 15, 199–214. Special issue on *Simulated Evolution and Learning*.
- Burke, E. K., Kendall, G., & Soubeiga, E. (2003). A tabu-search hyper-heuristic for timetabling and rostering. *Journal of Heuristics*, 9(6), 451–470.
- 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.
- De Causmaecker, P., & Vanden Berghe, G. (2008). *Categorisation of personnel rostering problems* (Technical report). K.U.Leuven.
- De Causmaecker, P., Custers, N., Demeester, P., & Vanden Berghe, G. (2005). Semantic components for timetabling. In *Lecture notes in computer science: Vol. 3616. Practice and theory of automated timetabling V* (pp. 17–33). Berlin: Springer.
- De Causmaecker, P., Vanden Berghe, G., Baaajens, C., Gemmel, P., & Salomez, K. (2007). *Workforce planning and scheduling: an integrated approach for health care* (Technical report). K.U.Leuven.
- De Causmaecker, P., Demeester, P., & Vanden Berghe, G. (2009). A decomposed meta-heuristic approach for a real-world university timetabling problem. *European Journal of Operational Research*, 195(1), 307–318.
- De Cesco, F., Di Gaspero, L., & Schaerf, A. (2008). Benchmarking curriculum-based course timetabling: Formulations, data formats, instances, validation, and results. In *Proceedings of the 7th international conference on practice and theory of automated timetabling*. Montreal.
- de Vries, G. (1987). Nursing workload measurement as management information. *European Journal of Operational Research*, 29, 199–208.
- Easton, F., Rossin, D., & Borders, W. (1992). Analysis of alternative scheduling policies for hospital nurses. *Production and Operations Management*, 1(2), 159–174.
- Gemmel, P., Vandaele, D., & Tambreur, W. (2008). Hospital process orientation (hpo): the development of a measurement tool. *Total Quality Management and Business Excellence*, 19(12), 1207–1217.
- Glass, C. A., & Knight, R. A. (2009). The nurse rostering problem: A critical appraisal of the problem structure. *European Journal of Operational Research*.
- 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.
- Halldorsson, A., Kotzab, H., Mikkola, J. H., & Skjoett-Larsen, T. (2007). Complementary theories to supply chain management. *Supply Chain Management: An International Journal*, 12(4), 284–296.
- Haspelslagh, S., Messelis, T., Bilgin, B., De Causmaecker, P., & Vanden Berghe, G. (2009). *Hardness studies on rostering problems* (Technical report). K.U.Leuven and KaHo Sint-Lieven.
- Ho, W. K., Lim, A., & Oon, W. C. (2001). Maximizing paper spread in examination timetabling using a vehicle routing method. In *Proceedings of 13th IEEE international conference on tools with artificial intelligence* (pp. 359–366).
- Kellogg, D. L., & Walczak, S. (2007). Nurse scheduling: From academia to implementation or not? *Interfaces*, 37(4), 355–369.
- Kingston, J. H. (2001). Modeling timetabling problems with sttl. In E. K. Burke, & W. Erben (Eds.), *Lecture notes in computer science: Vol. 2079. Practice and theory of automated timetabling II* (pp. 309–321). Berlin: Springer.
- Kingston, J. H. (2005). A tiling algorithm for high school timetabling. In M. Trick, & E. K. Burke (Eds.), *Lecture notes in computer science: Vol. 3616. Practice and theory of automated timetabling V* (pp. 208–225). Berlin: Springer.
- Leyton-Brown, K., Nudelman, E., & Shoham, Y. (2002). Learning the empirical hardness of optimization problems: The case of combinatorial auctions. In *Constraint programming (CP)* (pp. 556–572).

- McCollum, B. (2007). A perceptive on bridging the gap between theory and practice in university timetabling. In E. K. Burke, & H. Rudová (Eds.), *Lecture notes in computer science: Vol. 3867. Practice and theory of automated timetabling VI* (pp. 3–23). Berlin: Springer.
- 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, number 21 in BNAIC*, Eindhoven, October 2009.
- Meyer auf'm Hofe, H. (2001). Nurse rostering as constraint satisfaction with fuzzy constraints and inferred control strategies. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 57, 67–99.
- Özcan, E. (2003). Towards an xml based standard for timetabling problems: Ttml. In G. Kendall, E. K. Burke, & S. Petrovic (Eds.), *Proceedings of the 1st multidisciplinary international conference on scheduling: theory and applications, MISTA2003* (pp. 566–569).
- Petrovic, S., & Vanden Berghe, G. (2008). Comparison of algorithms for nurse rostering problems. In *Proceedings of the 7th international conference on practice and theory of automated timetabling*, Montreal.
- Petrovic, S., Beddoe, G., & Vanden Berghe, G. (2003). Storing and adapting repair experiences in personnel rostering. In E. K. Burke, & P. De Causmaecker (Eds.), *Lecture notes in computer science: Vol. 2740. Practice and theory of automated timetabling IV* (pp. 148–165). Berlin: Springer.
- Pinedo, M. L. (2006). *Springer series in operations research. Planning and scheduling in manufacturing and services*. Berlin: Springer.
- Qu, R., Burke, E. K., McCollum, B., Merlot, L. T., & Lee, S. Y. (2009). A survey of search methodologies and automated system development for examination timetabling. *Journal of Scheduling*, 12(1), 55–89.
- Ranson, D., & Ahmadi, S. (2007). An extensible modelling framework for timetabling problems. In E. K. Burke, & H. Rudová (Eds.), *Practice and theory of automated timetabling VI* (Vol. 3867, pp. 383–393). Berlin: Springer.
- Ryckbosch, F., Vanden Berghe, G., & Kendall, G. (2008). A heuristic approach for the travelling tournament problem using optimal travelling salesman tours. In *Proceedings of the 7th international conference on practice and theory of automated timetabling*. Montreal.
- Schaerf, A., & Di Gaspero, L. (2007). Measurability and reproducibility in university timetabling research. In E. K. Burke, & H. Rudová (Eds.), *Lecture notes in computer science: Vol. 3867. Practice and theory of automated timetabling VI* (pp. 40–49). Berlin: Springer.
- Thompson, J., & Dowsland, K. (1996). Variants of simulated annealing for the examination timetabling problem. *Annals of Operational Research*, 63, 105–128.
- Tien, J., & Kamiyama, A. (1982). On manpower scheduling algorithms. *Society for Industrial and Applied Mathematics*, 24(3), 275–287.
- Vanhoucke, M., & Maenhout, B. (2009). On the characterisation and generation of nurse scheduling problem instances. *European Journal of Operational Research*, 196(2), 457–467.
- Warner, M. (1976). Scheduling nursing personnel according to nursing preference: A mathematical programming approach. *Operations Research*, 24.
- Xu, L., Hoos, H. H., & Leyton-Brown, K. (2007). Hierarchical hardness models for sat. In *Lecture notes in computer science: Vol. 4741. Principles and practice of constraint programming (CP-07)* (pp. 696–711). Berlin: Springer.