Combined Working Time Model Generation and Personnel Scheduling

Maik Günther and Volker Nissen

Ilmenau University of Technology, Information Systems in Services, D-98684 Ilmenau
maik.guenther@gmx.de, volker.nissen@tu-ilmenau.de

Abstract. Workforce management is comprised of several phases, such as working time model generation and personnel scheduling. The combination of these phases has significant potential, especially for volatile personnel demand. This article shows that the concepts for the automatic generation of working time models already used in retail can be transferred to personnel scheduling in the logistics industry. Through this, the assignment of personnel can be accurately adapted to personnel demand. The results suggest the use of heuristics, especially meta-heuristics such as the evolution strategy or constructive methods which are adapted to the problem at hand.

Keywords: working time model, workforce scheduling, workforce management, evolution strategy, constructive heuristic.

1 Introduction

Despite many advantages in automation, logistics is still very labour intensive. At the same time logisticians are under considerable pressure, among other things due to legal regulations, stronger customer orientation and increasingly tough international competition. Above all, though, the current economic situation forces logisticians to take measures in order to remain competitive. An important parameter for cost-cutting is demand-oriented workforce management.

Employees spend up to 36% of their working time unproductively, depending on the branch [8]. Major reasons for this include a lack of planning and controlling. The problem can be dealt with using demand-oriented workforce management. Key planning goals are increased productivity, reduction of personnel costs, prevention of overtime and better motivation of employees [13].

In order to achieve these goals, accurate personnel assignment adapted to demand is required. A multi-level approach to workforce management in seperate steps can be very inefficient, especially for volatile demand. In retail, where personnel demand generally depends on customer frequency, which is very volatile, other methods must be used. For several years working time models have been created automatically during personnel scheduling, meaning these models are not preset. Rather, there are rules which are applied directly during personnel assignment planning (e.g., the minimum and maximum length of the models). The success of this concept in retail suggests its use in other fields in which

W. Dangelmaier et al. (Eds.): IHNS 2010, LNBIP 46, pp. 210–221, 2010.

volatile personnel demand also occurs. Therefore, the present work attempts to investigate the automatic generation of working time models as a direct part of personnel scheduling using practical examples.

First, the four phases of workforce management for volatile demand are explained. Possibilites for combining individual phases are shown at the end of section 2, in particular the combination of working time model generation and personnel scheduling. Section 3 discusses the necessity of accurate working time models using an example from logistics. For this, the limitations of the current practice (two-phase approach) are highlighted. The remedy is the automatic generation of suitably adapted working time models directly in personnel scheduling, as has already been done in retail. This concept is presented in section 4. Also, corresponding practical problems from retail, which have already been worked on, are discussed as well as solution methods used to integrate model generation and personnel scheduling. The transfer of these approaches from retail to logistics is done in section 5 using a practical application. Finally, the paper concludes with a summary and some indications for future work.

2 The Four Phases of Demand-Driven Workforce Management

Workforce management is a central component of business action and is intricately linked to many other processes in the organisation. Several software manufacturers follow an idealised approach, which is shown in figure 1 [12]. This method is quite practical and useful; however, the combination of phases can bring about further potential for usage. The individual phases are explained below and the possibilities for combining them are then discussed.

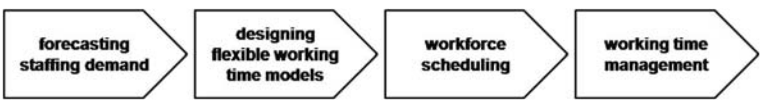

Fig. 1. Four phases of demand-driven workforce management

2.1 Staffing Demand Forecasts (Phase 1)

It is necessary for efficient personnel assignment to determine personnel demand as exactly as possible. According to a study by Miebach Consulting [7], errors in this phase alone can cause up to 15% higher personnel costs. In practice, various methods are employed depending on the planning horizon, with simple approximation dominating for short-range estimation of demand.

Demand determination based solely on experience is often suboptimal. However, the planner frequently lacks high-performance support tools. Modern solutions are able to predict the expected work volume using past data. Another form of demand determination is event-oriented. In logistics, for instance, if it is known ahead of time when a lorry (or van, aeroplane or train) of a specific type will arrive or depart again, the personnel demand can be ascertained.

2.2 Designing Flexible Working Time Models (Phase 2)

A working time model defines work beginning and ending time on a particular day, among other things. Starting with personnel demand, working time models are generated which cover that demand as well as possible. Legal and contractual regulations must be taken into account. The amount of effort required to generate models is usually quite large. Sometimes, they must also be authorised by the works council. A set of working time models, onces generated, are usually rarely changed and are used for later planning. An extreme negative example for an ill-suited working time model would be a plan with only one model (work begins at 7:00, ends at 15:00, variable break of 30 minutes, no flexitime). If all employees are scheduled using this model, sub-daily, weekly and even seasonal variations in personnel demand cannot be compensated for. Unnecessary costs due to idle time and overtime as well as bad service and low employee motivation would be just some of the effects.

2.3 Workforce Scheduling (Phase 3)

This phase involves the actual plan generation. It is determined which employee is assigned to which workstation at what time or rather which task is to be completed. The size of the planning horizon can be variable, as is the point in time at which the staff can view a new plan. Various constraints are considered during the planning, e.g., qualifications, desires of the employees, absences, balance of flexitime hours, legal and contractual regulations, aspects of fairness and so on.

2.4 Working Time Management (Phase 4)

Time management involves more than just recording the entering and exiting times of employees. In this phase, the flexitime balances of each employee are calculated based on their working times. These calculations and the regulations on which they are based can be quite complex. Absences are also planned and tracked. Additionally, as part of time management all data are gathered which are necessary for input into the payroll system.

2.5 Combination of Phases

Generally the four phases are carried out in sequence but there are sensible exceptions. The integration of personnel demand determination (phase 1) into personnel scheduling (phase 3) can be found, for example, in call centres, where personnel demand cannot be separated from scheduling. If the service level sinks sufficiently due to a lack of manpower, many customers will phone again in subsequent periods, affecting the personnel demand of that period as well.

The combination of working time model generation (phase 2) and personnel scheduling (phase 3) is also of interest. These approaches are sometimes used in retail when working time models are dynamically created during each planning session (phases 2 and 3 in one step). This combination significantly increases the flexibility of assignment planning and makes accurate scheduling adapted to demand possible. Sections 4 and 5 discuss the combination of these phases in more detail.

Even more optimisation potential lies in a combination of all three phases 1 to 3, as one could not only fit personnel supply to demand, but also shift demand in a favourable way for planning. However, this is only possible when personnel demand is fairly deterministic and the precise moment of occurence is not completely externally determined. Such flexibility is available, for instance, when customer orders can be individually planned ahead and to some extent be moved between different time slots.

Various approaches can be found with respect to phase 4; it can be arranged before or parallel to other phases. In this way, personnel scheduling can influence time management. On the other hand, if the time management data flow into the next planning session, flexitime balances and absences can be taken into account. Planning for days off alone would require a parallel approach because providing days off reduces personnel capacity, possibly requires adjusted working time models and even influences personnel demand (refusal/postpone of customer orders due to reduced personnel capacity).

3 Demand-Oriented Working Time Models

This section discusses the effects from the generation of demand-oriented working time models. The limits of this approach are also examined. The differences between fixed and demand-oriented working time models are illustrated using a case from logistics.

Figure 2 shows real data from a logistics service provider, who mainly carries out loading and unloading activities. The personnel demand fluctuates throughout the five days. However, it is not possible with only three available working time models to cover the demand well. The result indicates large phases of over- and understaffing due to the volatile demand with unadapted personnel assignments. This situation is critical for the logistics company because a high service level is contractually obligated to the customers.

This problem can be countered with the introduction of further working time models. Figure 3 clearly shows how the personnel assignment is oriented toward demand. By better utilising the normal working times of all employees, over- and understaffing can be greatly reduced while maintaining the same number of employees and the same level of effort. The effects which result from this are obvious: cost reduction through improved utilisation of employee time and fewer temporary workers, increase of turnover through a higher level of service and a rise in employee motivation.

Through an increase in the number of working time models, even short-term fluctuations in personnel demand can be covered economically. This, of course,

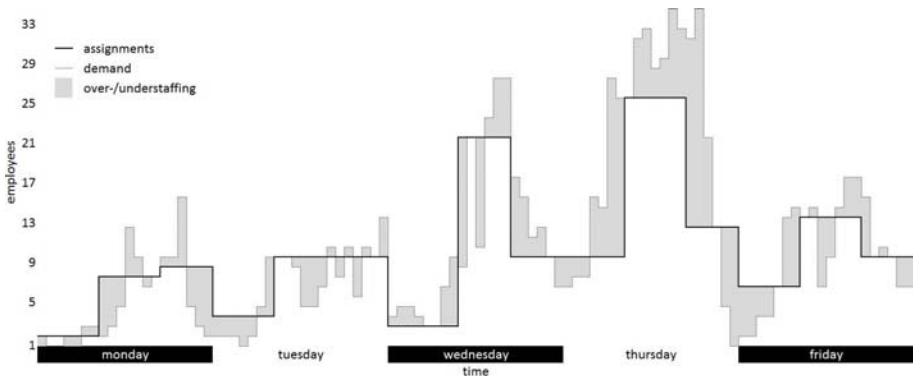

Fig. 2. Personnel demand and personnel assignment with three working time models

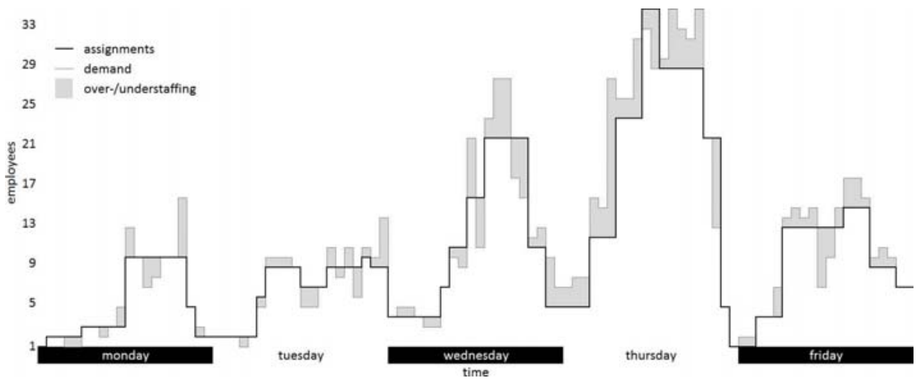

Fig. 3. Personnel demand and personnel assignment using more working time models

means a rise in complexity and effort needed for the generation and maintenance of the working time models. It is impractical to do this manually or by using spreadsheets, but a powerful software solution is needed.

On the basis of the working time models generated in this phase, employees are subsequently assigned working time models and workstations or orders for a defined time period in the assignment planning phase (phase 3). During this phase, qualifications, flexitime balances, teams, sub-daily workstation switches, planned absences and so on are taken into account. Using an approach in which phases 2 and 3 are not integrated, the latter phase must necessarily leave some constraints and some data unaccounted for. An otherwise optimal solution for phase 2 and subsequently for phase 3 could be determined to be sub-optimal when the result is regarded as a whole. Therefore, a combination of both phases is desirable.

Such as combination, however, requires that personnel demand can be forecasted accurately, using, for instance, past demand data, fixed customer events or information from returns data. However, this already applies, when the phases

2 and 3 are not integrated, but only the number of working time models is increased. Any attempt for demand-oriented personnel planning must ultimatively be based on reliable demand figures. Moreover, stochastic influence (dynamic demand) should be fairly low. Otherwise, there are more effective measures to capture stochastic demand peaks, for instance employing temporary workers, who can be ordered to work or sent home at short notice (but at extra cost).

4 Integrating the Phases 2 and 3: Lessons from Retail

Through the automatic generation of working time models directly within workforce scheduling it possible to generate a plan without fixed working time models. More specifically, automatic working time model generation refers to separating oneself from the idea of using only a few rigid working time models, because they do not always provide optimum coverage of the actual personnel demand. Instead, the planner only needs to provide some rules for the generation of working time models suitably adapted to demand. These include the minimum and maximum duration of the models as well as the limits for their beginning and end. Individual working time models are then generated for each employee for each day. Naturally, aspects as employee availability and planned weekly working hours are accounted for.

This apporach has been successfully implemented at many retailers and several software manufacturers provide specialized methods for this task. This stresses that the combination of phases 2 and 3 in workforce management is both possible and indeed successful. However, only a few scientific sources have dealt with this topic, yet. Below, we discuss this related work.

4.1 Work of Others

In 2007 Sauer and Schumann [11] introduced a constructive method for a personnel planning problem with automatically generated working time models in retail. The procedure was implemented in the workforce management system of a software manufacturer. Demand is viewed as deterministic and forecasted based on historical data. No demand dynamics are considered at the time of planning. Due to the fact that the scheduling algorithm was designed for interactive scheduling, greedy heuristics were chosen for two reasons. Expert knowledge can easily be integrated into the algorithm to generate acceptable plans for the human planner. And greedy heuristics generate a solution quickly. Unfortunately, this solution method cannot take into account more than one workstation or sub-daily workstation switches. Additionally, the planning horizon is limited to a maximum of one week.

Prüm [9] also addresses different variants of a personnel planning problem from retail. He uses 20 test problems with 4 to 168 timeslots as well as 2 to 576 employees. Prüm creates working time models parallel to assignment planning with varying demand in one-hour intervals. However, only one workstation is present in the scenarios of Prüm, and sub-daily workstation rotations are not included. He experiments with simplex, branch & bound as well as a hybrid method. His results indicate that problems of realistic size with constraints can in general not be successfully solved with exact methods.

4.2 Own Work

A practical retail problem with two workstations and also including planned sub-daily job rotations is considered in our own, previously unpublished work. Here, working time models are generated directly as part of personnel assignment planning. Fifteen employees work in the department for ladies' wear at a department store. The store is open Monday to Saturday from 10:00 to 20:00 and closed on Sunday and holidays. Six employment contracts exist, differing in the amount of planned weekly working time from between 10 and 40 hours. Employees are assigned to two different workstations (till and sales), with all employees trained for both stations. Qualifications are therefore not required to be taken into account. Many other factors influence planning, such as regulations, employee availability, promotions and time sheets. The personnel demand can be forecasted fairly accurately. It is given in one-hour intervals and centrally determined based on past data.

In order to find a solution to the problem described, automatic working time model generation is performed. Because of high fluctuations in demand, sub-daily workstation changes are allowed. Several hard and soft constraints must be considered during the planning process. They include rules supplied by the personnel planner of the company to ensure reasonable working time models. One year is planned ahead based on available historical data, resulting in a very complex search space with 131,400 dimensions (decision variables). If the assigment plan violates soft constraints, this is punished with error points that reflect the companies requirements as inquired through interviews. The planning problem can now be modeled as an optimisation problem where the sum of error points must be minimised. For real data sets and benchmarks see [14].

Staff scheduling is a hard optimisation problem. Garey and Johnson [3] demonstrate that even simple versions of staff scheduling problems are NP-complete. Kragelund and Kabel [6] show the NP-hardness of the general employee timetabling problem. Thus, heuristic approaches appear justified for the combined tasks of automatically generating working time models and scheduling the workforce in a company. In particular, we compare heuristics based on the evolution strategy (ES) [1], particle swarm optimisation (PSO) [5], multi-start local search (MLS) and a commercially successful constructive heuristic.

Our solution approaches based on adapted versions of ES, PSO and MLS create working time models automatically in the following way: The problem is represented as a two-dimensional matrix of employees and time periods, where the cells are filled with workstation assignments. A dummy workstation is used to mark times where an employee is not available. The heuristics change these workstation assignments, implicitly generating working-time models that are evaluated through the penalties in the objective function. Faulty models (wrong length, gaps etc.) are additionally corrected by a repair heuristic within ES and

PSO. For each heuristic, different strategy parameter settings were tested. 30 runs were performed for each setting using 400,000 evaluated solutions as the uniform termination criterion.

We compare these results to a commercial workforce management software package that delivers an adequate constructive method, capable of solving the problem at hand. All restrictions are supported and the associated error points can be entered into the software. Unfortunately, no code is available and no information is given in the documentation as to how the working time models are generated, so it must be considered a black box. However, our application of the constructive heuristic was supported by the software manufacturer, so errors in software handling can be excluded.

Table 1 presents results for the different solution approaches to the retail problem. For ES and PSO only the best strategy parameter sets are given, i.e. a swarm size of 20 particles in a gBest neighbourhood structure for PSO, and a (1, 5)-selection strategy for the evolution strategy. The constructive heuristic can be regarded as a benchmark, since it is part of a commercially successful software package that is actually used at around 300 companies for personnel planning.

ES with classical Gaussian mutation (Type G) yields the best mean error points. ES with a mutation adapted to combinatorial search spaces using the concept of maximum entropy (Type E) [10] finds the best of all solutions. If one inspects the assignment plans generated with the ES, one sees that they can hardly be improved upon, even with highly complex manual changes. For this reason, and because of the superior performance over the other solution methods, these plans can be regarded as very usable.

ES and PSO significantly outperform the commercial constructive heuristic. The constructive method delivers unsatisfactory results, because it is unable to cope with the various employee contracts. However, one run requires only some 10 minutes as opposed to 6 hours for a single run with the ES or PSO. Multi-start local search, as some form of 'brute-force' solution approach, is the worst heuristic and not able to generate reasonable solutions at all.

5 Transfer to an Application from Logistics

5.1 Description of the Problem

It can be seen in table 1 that several solution methods were able to almost eliminate over- and understaffing using automatically generated working time models. While this form of personnel planning has already gained significance in retail, practical applications in logistics are unknown to us. Very small personnel budgets, a tendency for understaffing and intense competition have promoted automatic generation of working time models in retailing.

However, logistics in many areas is also characterized by tough competition, cost pressure and volatile demand. Therefore, a transfer of workforce management approaches originally developed in a retail context appear useful. At this stage, the applicability of individual solution methods is not our focus. Rather,

Table 1. Comparison (minimisation of error points) of the different approaches, based on 30 independent runs each. Best results are bold and underlined.

heuristic	error		number of job-changes	under-staffing in (h/yr)	over-staffing in (h/yr)	too much weekly working time in (h/yr)	more than one working time model per employee per day
	mean	min					
constructive heuristic	84,690.0	84,690	0.0	25.0	0.0	1,386.5	0.0
multistart local search	4,265,931.9	3,807,375	2,434.5	209.5	1,420.8	955.8	410.8
PSO (20)	37,117.9	14,385	389.9	13.9	0.3	597.9	0.0
ES (1,5) Type G	**8,267.1**	5,924	214.3	13.9	0.1	120.2	0.0
ES (1,5) Type E	8,464.8	**4,911**	248.0	14.0	0.1	122.8	0.0

the objective is to check whether this type of integrated work force planning can be reasonably transferred from retail to a given logistics service provider scenario, which is typical for many companies in that sector. Again, the idea is to automatically generate working time models as part of workforce scheduling. For this reason, three variations with diverse preset working time models will be compared to the combination of model generation and assignment planning. The fixed working time models were directly taken from the human resources department of the respective company.

The problem shown below exhibits many similarites to the logistics problem discussed in [4]. Personnel demand varies greatly at sub-daily intervals and for different workstations. Sub-daily workstation rotations are to be planned as well. Actual data are available for an entire year. The logistics provider operates between 6:00 and 19:00 with 480 employees. Their contracts are the same in most cases, so that the contractual weekly amount of work varies only slightly. In order to keep personnel assignments flexible, flexitime balances are kept and workers may maintain the balance within the range of -30 to +80 hours. The work volume is subject to economic and seasonal fluctuations, but can be regarded as fairly deterministic. Sub-daily jumps in personnel demand occur as well. In order to automatically generate working time models, the regulations found below are in effect and violations are assigned an amount of error points that reflects the companies requirements as inquired through interviews:

- compliance with the given allocations (no over- or understaffing)
- employees should not be assigned to workstations with no demand
- compliance with the contractually arranged maximum number of hours
- working time models should not be shorter than 3 hours
- working time models should not be longer than 9 hours
- working time models must not be split up during the course of a working day

- working time models must not begin or end outside operating hours
- working time models begin and end on the hour
- work must not be done on Sundays and holidays
- absent employees (holiday etc.) may not be planned for the corresponding day
- sub-daily workstation changes are only allowed on the hour
- no unnecessary workstation changes (no improvement of assignment plan)

5.2 Results and Discussion

Table 2 shows the effects of increasing the number of working time models and adding automatic model generation. The calculations were done for an entire calendar year, with the focus of the optimisation in each case on reducing over- and understaffing. All of the working time models operate on the period between 6:00 and 19:00. As a rule, working time models were generated whenever possible with a length of 7.5 hours. Stronger consideration of employee desires or unplanned absences would worsen the results. However, the main tendency seen in table 2 remains even then.

Personnel demand is viewed as deterministic and planning is based on historical data. In practice, dynamic changes (late truck arrival etc.) occur, but only rarely. Nevertheless, this points to the limits of personnel planning. Personnel demand must be known in advance and stochastic influence should be low to allow for demand-oriented, integrated planning as suggested here.

It is clear from the table that over- and understaffing in personnel assignment are reduced through the corresponding increase in the number of models. The fluctuating demand can be covered significantly better with more models available. In principle, a maximum of 62 different models could be generated given the existing restrictions. The generation, administration and possible later adjustment of these models, however, is significantly time consuming and error-prone. These problems can be avoided through the use of automatic model generation in parallel to workforce scheduling. Moreover, the problem of over- and understaffing is completely eliminated with this method, as can be seen in table 2. This demonstrates that an integrated approach to automatically generate working time models and schedule personnel to work at given workstations can indeed be very beneficial in logistics. It must be pointed out that the resulting working time models respect the constraints (rules) on model generation as provided by the planner. Moreover, the resulting plan can be made available to employees long before the relevant working periods actually begin, so that also contractual regulations regarding plan stability are covered.

The results in table 2 were generated with the constructive heuristic also shown in table 1. While this approach performed poorly for the retail problem, it proved to be very useful for the logistics problem. The constructive method is more suited to staff whose contractual amount of hours are similar, as is the case in the logistics problem. The retail problem, by contrast, presented highly variable contracts (full-time, student helpers, so-called 400€ workers etc.). With this diversity in retail, the constructive approach had significant difficulties and, consequently, was unable to produce adequate sub-daily staffing changes.

Table 2. Effects of increasing the number of (preset) working time models and introducing automatic working time model generation on over- and understaffing in a representative logistics personnel planning application

working time models	overstaffing (in h/yr)	understaffing (in h/yr)	total (in h/yr)
5 different working time models	26,094	528	26,622
6 different working time models	16,488	4,734	21,222
15 different working time models	3,426	438	3,864
automatically generated working time models	0	0	0

As a rule, constructive heuristics are suitable for automatic generation of working time models when the method is specifically designed to solve the particular problem at hand. Small changes of the application area can be enough to severely limit the quality of the results, though. Because the constructive heuristic fits well to the logistics problem and produces very satisfactory results in a short time, the implementation of further and potentially more time-consuming approaches, such as metaheuristics, was dispensed with for now.

6 Conclusion and Future Work

It was shown that the demand-oriented generation of working time models can contribute significantly to accurate workforce scheduling adapted to personnel demand. This enables the minimisation of overtime and idle time, among other things, which positively affects personnel costs, workload on individual employees and service level. The limits of a two-phase approach for working time model generation and subsequent personnel assignment planning were shown. The remedy is the combination of both phases into one planning step. This approach has already gained significance in retail, where personnel demand is extremely volatile. The concept was successfully transferred from the retail domain to a logistics personnel planning application that is representative for many similar problems in the logistics sector. This opens up the possibility to use a practically tested instrument for the flexibilisation of workforce scheduling under volatile personnel demand in logistics. The corresponding heuristic solution methods should be integrated in todays workforce management software to fully leverage the benefits.

Future work will involve the implementation and comparison of further solution methods for the mentioned logistics application as well as other scheduling applications from the logistics and retail domains.

References

1. Beyer, H.G., Schwefel, H.P.: Evolution strategies: a comprehensive introduction. Natural Computing 1, 3–52 (2002)
2. Ernst, A.T., Jiang, H., Krishnamoorthy, M., Owens, B., Sier, D.: An Annotated Bibliography of Personnel Scheduling and Rostering. Annals of OR 127, 21–144 (2002)

3. Garey, M.R., Johnson, D.S.: Computers and Intractability. A Guide to the Theory of NP-Completeness. Freeman, New York (1979)
4. Günther, M., Nissen, V.: Sub-Daily Staff Scheduling for a Logistics Service Provider. KI Künstliche Intelligenz 25 (2010) (accepted)
5. Kennedy, J., Eberhart, R.C., Shi, Y.: Swarm Intelligence. Kaufmann, San Francisco (2001)
6. Kragelund, L., Kabel, T.: Employee Timetabling. An Empirical Study. Master's Thesis, Dep. of Computer Science, Univ. of Aarhus (1998)
7. Miebach Consulting GmbH (ed.): Personalbedarfsplanung. Personalplanungs- und Flexibilisierungsmethoden in der logistischen Praxis. Frankfurt am Main (2008)
8. Proudfoot Consulting (ed.): Produktivitätsbericht 2007. Company Report (2007)
9. Prüm, H.: Entwicklung von Algorithmen zur Personaleinsatzplanung mittels ganzzahliger linearer Optimierung. Untersuchung eines bedarfsorientierten Modells mit flexibler Schichtkonfiguration. Master's Thesis, FH Trier, FB Informatik (2006)
10. Rudolph, G.: An Evolutionary Algorithm for Integer Programming. In: Davidor, Y., Männer, R., Schwefel, H.-P. (eds.) PPSN 1994. LNCS, vol. 866, pp. 139–148. Springer, Heidelberg (1994)
11. Sauer, J., Schumann, R.: Modelling and Solving Workforce Scheduling Problems. In: Sauer, J., Edelkamp, S., Schattenberg, B. (eds.) Proc. of PUK 2007, pp. 93–101 (2007)
12. Scherf, B.: Ein Vorgehensmodell zur Personaleinsatzplanung. In: Fank, M., Scherf, B. (eds.) Handbuch Personaleinsatzplanung, Datakontext, pp. 15–53 (2005)
13. Scherf, B.: Wirtschaftliche Nutzenaspekte der Personaleinsatzplanung. In: Fank, M., Scherf, B. (eds.) Handbuch Personaleinsatzplanung, Datakontext, pp. 55–83 (2005)
14. Sub-Daily Staff Scheduling Data Sets and Benchmarks, `http://www.tu-ilmenau.de/fakww/2608+M54099f70862.0.html`