

Case studies in planning crew members

J. P. Martins & E. Morgado

SISCOG, Sistemas Cognitivos SA, Portugal

Abstract

CREWS is a software product for planning and managing the work of rail-staff, which provides solutions to one of the core problems that railway companies face today – effective management of resources.

This paper reports some results of a long-term development work in the application of both Artificial Intelligence and Operations Research techniques to the planning and managing of staff (drivers, guards, and station personnel). The paper presents case studies emerging from the application of CREWS, both to railway and to subway companies. CREWS has grown to be a strong and mature product that provides decision-support in the task of planning the daily work of more than 20,000 staff members across Europe. CREWS-based systems are in routine use in the Dutch Railways, the Norwegian State Railways, the Finnish Railways, the Danish State Railways, the Suburban trains of Copenhagen, and the London Underground.

Keywords: crew scheduling, crew planning.

1 Problem description

Railways are going through a golden expansion period. The pressure for environmentally friendly transportation associated with the growing needs for mass transport and the possibility of competing with airlines in medium-range routes are generating a rail expansion that had not been seen for almost a century. On the other hand, productivity concerns are leading railway companies to introduce new management styles to improve the results of the business. In particular, new tools are being searched to improve the use of resources. A railway company needs to manage three main types of resources: track, rolling stock, and crew. In this paper, we just address the management of crew.



The planning and management of crew members is constrained by the timetable, the rolling stock roster, the labor rules, the number of available workers and their skills, and a large number of other operational constraints. From a global point of view, this problem is addressed in a sequence of phases:

1. *Long-term planning*. Produces a plan of the work without taking into account the names of workers (but rather the available skills) nor the dates where the work is being planned (but its frequency). Long-term planning deals with the regularity involved in the operation. It is done in two sequential steps:
 - (a) *Duty planning* arranges tasks into duties (sequences of tasks to be done by one person in one day). Constraints in this phase include the maximum duration of a duty, space continuity between tasks, transfer times between tasks with different equipment, places and times for meal breaks, compatibility of route and equipment skills.
 - (b) *Roster planning* arranges the duties produced in (a) in sequences of duties, rest time, and days off. Constraints in this phase include the maximum week working time, rest time between duties, weekly days off, and available skills.
2. *Staff allocation*. Associates crew members to the weeks of the rosters produced in 1 (b), giving rise to an allocated roster. The instantiation of the allocated roster for a certain period of days produces a *daily plan*, a plan for the work of each staff member for each individual calendar day in the period.
3. *Short-term planning*. Deals with predicted irregularity. It amends the daily plan for certain calendar days, either because some of the tasks have changed (e.g., a change in the timetable due to track work) or because a person is not available to work. Short-term planning works with particular days and with particular persons. The constraints handled in this phase include all the constraints in Phases 1 and 2, plus constraints associated with particular persons, such as maximum working hours per year and holidays.
4. *Dispatching*. Dispatching is performed on the day of operation. It is similar to short-term planning, but works in real-time. The changes introduced at this level may be due to delays, breakdowns, or unexpected absences of crew.
5. *Control*. Control is done after the work has been completed. It checks the work that was performed, comparing it with the work that was planned in Phases 2 and 3, updates personnel records, and feeds the payroll system.

2 Approach

When CREWS started being developed, in the second half of the 1980s, rail-staff planning had been approached by Operational Research (OR) techniques [2], but the results obtained with automatic optimization algorithms, based on a “black-box” approach, have proven to be unsatisfactory, mainly because when faced with



a full size problem, these solutions tend to need computational resources that far exceed what is available.

Since human planners could build acceptable plans where algorithmic solutions failed, SISCOG took the challenge of using Artificial Intelligence (AI) techniques to approach this problem. One of our initial goals was to produce a “white-box” system, in the sense that the planner could perceive what was going on, could interact with the system, proposing alternatives or querying decisions, and could adapt the behavior of the system to changing circumstances. The resulting system should play the role of a “digital colleague” interacting with planners to build plans in a co-operative way. This approach was much influenced by the early Expert Systems [5]. Furthermore, SISCOG took the additional challenge of building a product that contains the basic knowledge for crew planning, remains constant across companies, and only needs to be extended with the particularities of each one (domain, labor rules, planning strategies, and objectives).

CREWS uses an architecture that is based on a central server linked to a database and to modules that can either be located centrally or locally in the depots, together with internet access. The modules of CREWS (shown in Figure 1) follow very closely the phases presented in Section 1: there are two modules for the long-term planning phase, the *Duty Scheduler* and the *Roster Scheduler*; one module for the allocation phase, the *Staff Allocator*; one module for the short-term planning phase, the *Short-term Scheduler*, that is complemented by *communication functions* that allow staff to consult their plans via internet and to introduce requests regarding changes in work; one module for the dispatching phase, the *Real-time Dispatcher*; and two modules for the control phase, the *Work Recorder* and the *Web-based Work recorder* (enabling the workers to introduce requests for the correction of the work that has been done).

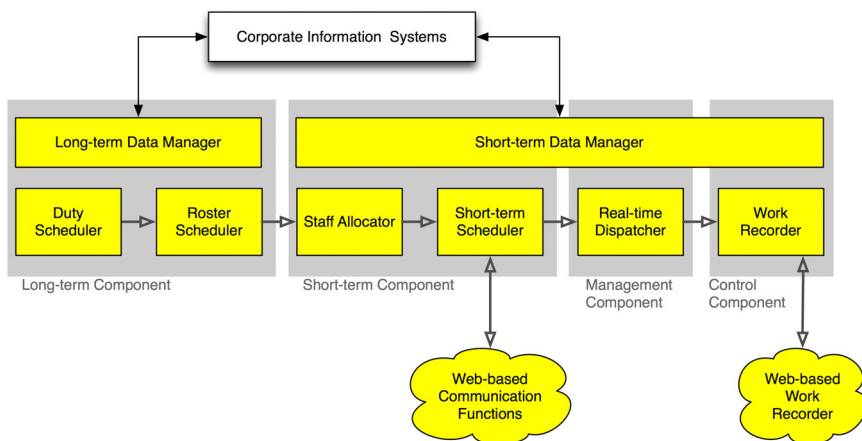


Figure 1: Modules of CREWS.

Another module, the *Data Manager*, acts as the interface with corporate information systems and supports the preparation of input data, handles change in data, enables simulation of hypothetical data situations, and maintains the consistency and completeness of data (both before and during the planning process).

The “white-box” approach used by CREWS, provides mechanisms for supporting manual, semi-automatic and automatic modes of operation in an integrated way. Each operation mode offers a different level of user-support:

1. In *manual mode*, the user operates the system, using drag-and-drop, as if he/she was using pen and paper. The information is displayed in a screen that contains two windows, the top window shows the work that has yet to be planned and the bottom window shows the plans that have been constructed (Figure 2 illustrates this concept with a screen from the Duty Scheduler). Whenever an operation is performed, the system verifies all constraints imposed upon the plan and tells the user the constraints that are violated by the operation. If the user chooses to violate a constraint, the plan is shown with a violation indication;
2. In *semi-automatic mode*, the system gives hints about how to construct the plan, by computing a set of alternatives following an optimizing strategy. The role of the user is to select the proposal that he/she thinks is best;

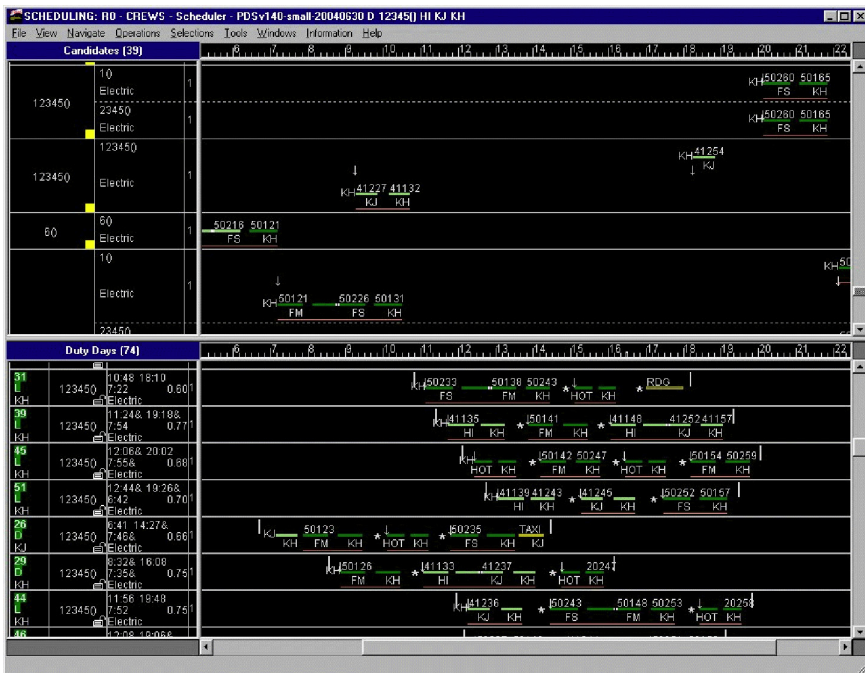


Figure 2: Duty Scheduler screen.

3. In *automatic mode*, the system generates the plan by itself, resorting to an optimizer. CREWS's optimizer uses a combination of Artificial Intelligence and Operational Research techniques;
4. In *mixed mode*, the user constructs the plan by resorting to an arbitrary combination of the other modes of operation. It provides a full cooperation between the user and the system, showing what is going on, providing explanations about the decisions taken by the system, enabling the interaction of the user on the work being done, and taking the bulk of work from the user, when he selects to do so.

All modes of operation use the same set of labor rules, which are defined independently from the system and could be changed by the customer to adapt the system to a new reality. The labor rules are validated during the construction of the plans. In order to comply with these rules, the system was capable of creating new tasks to be combined with the tasks to be planned, such as positioning trips, meal breaks, sign-on and sign-off tasks.

3 Case studies

3.1 Dutch railways (NS)

NS (Nederlandse Spoorwegen) is the main passenger railway operator in The Netherlands. Every day, over a million persons travel by train in The Netherlands over the busiest railway transportation network in the world. NS operates 4,700 trains per day on a network with 1,500 miles, with a rail staff of 5,200 (2,500 drivers and 2,700 guards).

The development of a CREWS-based system started in 1993. Its scope was limited to planning long-term duties [4]. Although duties are distributed over 29 depots, planning is done centrally. The project development work was quite long, spanning from 1993 to 1997, due to the fact that (1) CREWS was not mature and (2) to changes in the system's initial architecture.

The system started live production in 1998. In the version originally deployed, only the manual mode was being used because the solutions provided by the optimizer were not satisfactory. Even with just the manual mode, NS started to obtain benefits from the system: all the tedious work of checking which trains were covered was handled by the system; all the labor rules were in the system and were automatically checked whenever duties were constructed or when there was a change in the timetable; furthermore, since all labor rules were in the system, detailed knowledge of these rules was not required of planners. Differences in rules and regulations for drivers and guards started to be seen as insignificant. As a consequence, planners could perform planning for both personnel groups and the planning units could be integrated. The initial use of the system originated 10% decrease in the number of planners.

Although the system was appreciated by planners and the management of NS appreciated the flexibility in the change of labor rules, the optimizer remained as



a problem because of its very local view of the planning process, producing sub-optimal solutions. NS started to look for an external optimizer that could be linked into the system, ending up with TURNI [1] that was used for a few years.

In the early 2000s, SISCOG launched a new optimizer based on OR. Upon the presentation of the results that were being achieved by the optimizer, NS decided to set up 10 scenarios and to benchmark the results achieved by the CREWS-based optimizer and TURNI. In 8 of the scenarios CREWS-based optimizer produced gains between 0.5% and 1.5%; 2 scenarios were solved by the CREWS-based optimizer and not by TURNI. As a result, NS switched back to a full fledged system based on CREWS.

In 2008, NS and SISCOG joined forces on a research project for a further improved version of the optimizer. The goal of the new version was to be able to run the entire planning problem for NS in one single run. This means simultaneously generating duties for all depots in the country, about 1,100 duties for a single day, resulting in a much larger planning instance than can be typically found in the OR literature. At the end of this project, not only the system was able to perform the duty generation for the entire country, but also achieved gains of 3% over what was used.

In 2009, a benchmark was performed using 1999 data and rules to measure the overall improvement that the several versions of the optimizers had introduced. The results of the tests have shown that the new optimizer gives about 6% efficiency improvement with respect to the manually created plan. In the mean time, the several changes to rules and constraints that have been introduced (and were not considered in the benchmark) give additional efficiency. It is also important to say that the generation of duties subject to the new rules and constraints would have been impossible without the use of an IT optimizer. The number of planners was reduced by 60% from the initial count in 1999.

3.2 Norwegian state railways (NSB)

Norges Statsbaner, commonly known as NSB, is the largest passenger railway company in Norway. NSB operates 1,300 trains per day during the week and 770 in weekends, over a network with 2,540 miles, with a staff of 1,800 (1,000 drivers and 800 guards). Staff is divided into 41 depots for drivers and 38 depots for guards, spread over the country.

The project for the development of a CREWS-based system, named TPO, started in 1998. The goal of the system was to have an integrated environment for producing duties, rosters, staff allocation, short-term changes, and recording and reporting [3]. NSB produced a detailed set of requirements based on the analysis of working methods. The job was well done, but as the project progressed, NSB realized that it was impossible to completely specify the application in advance. The development was done in several phases, each one addressing one of the modules of the system.

One of the big difficulties of this project were the labor rules. NSB has a very large number of quite complex labor rules. Rules had different interpretations for



planners across the country, almost each depot in charge of rostering and short-term planning had a unique interpretation for several of them. One of the jobs in the project was the clarification and the agreement on the meaning of the labor rules and their introduction on the system.

Many changes took place in NSB during the period the system was being implemented. Production was subdivided into business areas for guards. Freight was split from production and is now handled by a separate company. Numerous organizational changes affected the way tasks are subdivided among planners. For these reasons, it was not possible to create a basis for comparative analysis. The results of the use of the system show that NSB has more effective production of plans. Planners use less time to produce a complete plan with ready printouts and all calculations of statistics. Management has been able to check consequences of new rules while negotiations with the unions were taking place. This had never been done before. Gained time is primarily used to make better plans and to create more alternatives for the next and later timetable periods. This also had never been done before. More reliable and complete statistics have contributed to improvement of cost control and pricing of services. Hidden costs became visible and the use of TPO has contributed to a reorganization of the planning process and a clarification of many issues.

TPO was deployed in several stages. Conductors started long-term production in 2000; drivers, who have a much complicated set of labor rules, had the long-term planning fully deployed in 2003; short-term planning, together with recording and reporting, was progressively placed in production from 2003 to 2006. Taking a new system into production is a long-term investment and the key to successful implementation is commitment from management. The users do not experience advantages from day one. There is a threshold to overcome before the advantages become apparent. Endurance is required in the start up phase.

Since the end of the project, the TPO system has been extended with new modules of CREWS, the Real-time Dispatcher, Communication Functions, and the Web-based Work recorder.

3.3 Finnish railways (VR)

VR (VR-Yhtymä Oy) is a state-owned railway in Finland. Since the density of population in most parts of Finland is low, the majority of railway traffic is composed of long-distance trains. Commuter services are rare outside Helsinki area, but there are express train connections between most of the cities. VR provides both passenger and freight traffic, running more than 1,000 trains a day, over a rail network with 5,250 miles, with a staff of 2,800 (2,000 drivers and 800 guards).

The project for the development of a CREWS-based system, named VIP (VR's Integrated Planner) started in 2001. The goal was to develop a long- and short-term planning system for the drivers of both freight and passenger trains. The start of the project was not easy. On the one hand, VR had already tried the development of a similar system with another supplier and the development went



wrong, this generated suspicions towards a new supplier that was located in a country considered to be in a lower IT-development state than Finland. On the other hand, VIP was the first system that was developed by SISCOG based on a long and very precise list of technical and delivery-time requirements, involving quite a large team from the customer side. Any delay of a delivery or of an acceptance phase was considered by VR as a big problem. VR placed quite a stress on project management based on monthly reports. SISCOG had to learn how to deal with these aspects.

The planning problem itself was not easy. Although VR did not have the complexity of the labor rules existing in Norway, the possibility of combining passenger and freight trains in a single duty, together with the existence of several different types of duties and a large number of rules for transfer times, introduced quite a complexity in the problem. VR's rostering process, based on 3-week rosters, was quite different from what was used in other railways. On the other hand, an endless number of rules for defining the compensations to be paid to workers further hindered the problem.

A major project milestone was the formal definition of all duty types and their relationship with the different applicable rules. This was a job that VR never thought to be possible. Although the VIP system was developed with a large team from VR's side, expert planners had a limited time to devote to the project and were called upon their planning job whenever necessary. This fact introduced much longer acceptance periods for the software deliveries than was foreseen. The system was placed into production in two phases, the long-term component went into production in 2003 and the short-term component was deployed in 2004.

In 2005, SISCOG started another project with VR, extending the VIP system to plan the work of 800 guards and the work of 2,000 staff that perform railway-yard shunting tasks. Due to the mutual trust that, in the meantime, was built between the members of the project team, this second project had a very smooth development.

As a result of the two projects with the VIP system, there was a reduction of the number of staff. The exact reduction was not quantified, but nowadays duties are about 98% efficient, and there was a reduction in planning time to about one third. Staff is very satisfied with their duties considering that they provide a balance between work and free time.

In the period between 2008 and 2009, VIP was, again, extended to other classes of personnel, including ticket selling, contact center, and traffic control personnel. The work of some of these persons is not based on train timetables but rather on workloads that define peaks of work for certain periods of the day. The depth of the system was also extended to include real-time dispatching, work recording and web-based interfaces to crew members.

3.4 London Underground (LUL)

The London Underground, LUL, is a subway system serving a large part of Greater London and neighboring areas in the UK. The London Underground transports



more than 3,000,000 passengers per day over a network of 252 miles, organized as 11 lines, serving 275 stations and using more than 3,000 drivers.

The development of a CREWS-based system, named TSS, started in 2005. The goal of TSS was to produce duty and roster plans for all the drivers of LUL. Up to the start of TSS plans had always been done by hand. For a single major line, this job took over 13 weeks. LUL had a team of 15 planners, including trainees, with each planner focusing on a line. Training time for a competent planner was 2 to 3 years and 3 to 5 years for being able to plan a major line. LUL had three major goals for TSS: (1) to produce more robust and consistent plans; (2) to enable more efficient use of drivers; and (3) to allow for a larger number of planning options to be created and considered in a shorter period of time.

The work involved for SISCOG was much larger than initially foreseen. The plans of the 11 lines are internally grouped into 9 train-staff plans. Although the underlying labor rules for drivers are common, there are several characteristics that make the planning process different between lines. Each plan has its own specific set allowances. The details within each station regarding walking and transfer-times are huge and depend on the time of the day and on the day of the week.

LUL did not compromise for anything but perfection. Any plan produced had to be better than what was produced by the best planners. Different lines had different goals. For example, in the Northern Line, the goal was to increase the robustness of the plans, while keeping the same number of drivers; in the Piccadilly Line, the goal was to reduce the number of drivers, while keeping the robustness of the plan. During the project, the planners were asked to investigate the possibility of opening a new depot at Brixton (on the Victoria Line). The project team felt that this was a good opportunity to test the use of TSS in addressing “what-if” scenarios. Many of these aspects lead to further changes in the CREWS-based optimizer. Nevertheless, all challenges placed to CREWS were solved, with significant benefits to LUL. Due to contract confidentiality clauses, we are not allowed to disclose the exact gains that were achieved. However, the system paid itself after the first year of operation.

The deployment of TSS was done in phases, each phase being associated with a line or a set of lines. In 2008, the first plan produced by TSS went live in production. Nowadays all long-term plans for all lines are in full production with TSS. In early 2010, LUL contracted an extension to TSS to handle a smooth interaction between long-term plans and short-term plans resulting from work on the track and special days of operation.

4 Benefits

The main goal of the development of CREWS was to increase the efficiency of railway companies through the optimized use of staff resources. This goal has largely been achieved; CREWS-based systems are able to obtain a 3-5% cost reduction in the use of staff, when compared with the manual construction of plans. In companies with several thousand workers, this quickly adds up to savings of



several million dollars a year. However, many other benefits, that were not fully apparent when the development started, have surfaced along the years.

CREWS-based systems eliminate the burden of repetitive and tedious work that was placed on the work of planners, enabling them to concentrate on their planning work with a user-friendly help provided by the system.

Another benefit corresponds to the possibility of generating alternative solutions. When a non-IT approach is used, the company is usually satisfied when a plan is produced to cover the work on all trains. The generation of a staff plan, when done manually, is a task that involves several full-time planners over a period of several months and the delivery of the plans is usually just-in-time. With CREWS, a more than 10-fold speed increase in the production of plans has been achieved, together with the possibility of having several planners, working in parallel, producing alternative plans. Companies that use CREWS could, for the first time, start comparing other aspects of the solution, such as social benefits given to workers (other than the aspects that are prescribed in the rules), which enable the balance of productivity costs, robustness issues, and social aspects. Thus, gained time is primarily used to make better plans and to create more alternatives for the next and later timetable periods.

CREWS provides a “what-if” simulation mode, where new rules and conditions can be tested by simple changes in parameters. In this way, the company can evaluate the consequences of new rules being discussed with trade unions and the effect of providing new skills to staff located in different personnel bases.

A benefit that is not easily quantified corresponds to the preservation of the planning knowledge within the company. In fact, all the knowledge pertaining the task at hand is stored within the system. Differences in rules and regulations for different classes of personnel, seen by most companies almost as dealing with different realities, become insignificant when CREWS is used. As a consequence, the same person can plan different personnel groups and planning units can be integrated.

Acknowledgements

SISCOG would like to acknowledge the important role that our customers have played in the development of CREWS. SISCOG also thanks all its staff, current and past, for their work and dedication and the long night and weekend hours spent to make CREWS and its applications a success.

References

- [1] Caprara A., Fischetti M., Toth P., Vigo D., and Guida P. L., “Algorithms for railway crew management”, *Mathematical Programming* 79, pp. 125–141, 1997.
- [2] Ernst A. T., Jiang H., Krishnamoorthy M., Owens B., and Sier D., “An Annotated Bibliography of Personnel Scheduling and Rostering”, *Annals of Operations Research* 127, pp. 21–144, 2004.



- [3] Martins J. P., Morgado E. M., and Haugen R., “TPO: A System for Scheduling and Managing Train Crew in Norway”, *Proc. of the Annual Conference on Innovative Applications of Artificial Intelligence*, Riedl and Hill (eds.), pp. 25–32, Menlo Park, CA: AAAI Press, 2003.
- [4] Morgado E. M. and Martins J. P., “CREWS_NS: Scheduling Train Crew in The Netherlands”, *Proc. of the Ninth Annual Conference on Innovative Applications of Artificial Intelligence*, pp. 893–902, Menlo Park, CA: AAAI Press, 1997.
- [5] Nilsson N. J., *The Quest for Artificial Intelligence*, Cambridge, UK: Cambridge University Press, 2010.

