# Operational Railway Crew Scheduling with Equity Concerns

B.T.C. van Rossum[a,*] T. Dollevoet[a], and D. Huisman[a,b]

[a] Econometric Institute and Erasmus Center for Optimization in Public Transport, Erasmus University Rotterdam, Rotterdam, The Netherlands
vanrossum@ese.eur.nl, dollevoet@ese.eur.nl, huisman@ese.eur.nl
[b] Process quality and Innovation, Netherlands Railways, Utrecht, The Netherlands
[*] Corresponding author

*Extended abstract submitted for presentation at the 11th Triennial Symposium on Transportation Analysis conference (TRISTAN XI)*
*June 19-25, 2022, Mauritius Island*

January 14, 2022

Keywords: Railway optimisation, Crew scheduling, Equity, Column generation

## 1 INTRODUCTION

Netherlands Railways (NS) is the major Dutch railway operator and employs over 3,000 drivers and 3,000 guards divided over 28 crew bases. To ensure a fair allocation of work over the different crew bases, so-called Sharing-Sweet-and-Sour rules were introduced (Abbink *et al.*, 2005). These rules specify that no crew base has to perform a disproportionate amount of attractive (sweet) or less attractive (sour) work. For example, sweet work includes work on modern Intercity trains, whereas sour work includes work on lines with a higher risk of passenger aggression.

In the current crew planning process at NS, the Sharing-Sweet-and-Sour rules are incorporated in the tactical crew scheduling problem, i.e., when creating a crew schedule based on the generic, annual timetable. In the operational scheduling phase, where the crew schedule is updated to account for weekly deviations from the generic timetable, these rules are not explicitly enforced, however. As such, it is unclear to what extent the realised crew schedules satisfy the Sharing-Sweet-and-Sour rules. Moreover, since the Sharing-Sweet-and-Sour rules are formulated on the crew base level, they do not necessarily guarantee a fair division of work between individual crew members.

To tackle these two issues, we propose to (i) formulate Sharing-Sweet-and-Sour rules on the individual level and (ii) consider them explicitly in the operational planning phase. For these rules to be meaningful, we must consider a planning horizon of reasonable length, i.e., at least several weeks. The goal is then to construct a sequence of crew schedules, e.g. for each week in the period, such that the individual Sharing-Sweet-and-Sour rules are satisfied over the full planning period. To this end, we propose a sequential solution approach, where large deviations from the target levels are penalised. In particular, we propose to steer the crew schedules through the objective functions of the weekly crew scheduling problems in the planning period.

The contributions of this work are twofold. First, we formulate the operational crew scheduling problem with Sharing-Sweet-and-Sour rules on the individual level and propose a sequential solution approach to solve the problem. Second, we apply our sequential solution approach to a real-life instance from NS, and find that our approach is able to significantly raise the percentage of rosters that satisfies individual Sharing-Sweet-and-Sour rules.

## 2 PROBLEM DESCRIPTION

We consider the operational crew planning phase over a planning period of $T_{PLAN}$ days. In this phase, we aim to construct duties that match the rosters constructed in the tactical phase, and that cover all tasks in the weekly timetables. We assume that, for each crew member, an individual template-based roster is known. In other words, for each day of the planning period a template is assigned to each crew member, and this template in turn specifies the time window in which the crew member is allowed to work. Furthermore, we assume that the weekly timetables are constructed and released $T_{RELEASE}$ days in advance, where $T_{RELEASE} < T_{PLAN}$. For each week of the planning period, we require that a feasible crew schedule is constructed at least $T_{MIN}$ days in advance. A crew schedule is feasible when the duties in the schedule match the templates in the roster, and when the duties jointly cover all tasks in the weekly timetable. Realistic values for $T_{RELEASE}$ and $T_{MIN}$ might be 4 and 2 weeks, respectively.

Each roster consists of an assignment of duties to crew members. A duty is a sequence of tasks that constitutes a day of work, and should satisfy various constraints. First, each duty should start and end at the same crew base. Second, no duty should exceed a maximum workload, which is determined by the start time of the duty. Third, each duty should contain a meal break at a dedicated station with a canteen, and the workload before and after the break may not exceed a given maximum. Note that a duty can only be assigned to a crew member when the duty lies within the time window specified by the templates in the roster.

Finally, we consider a set of individual Sharing-Sweet-and-Sour rules that must be satisfied over the planning period. Each rule corresponds to a duty attribute, and requires that the average score of all duties performed by a crew member, with respect to this attribute, lies in a certain target interval. In this work we consider the following three duty attributes, which jointly aim to capture the (un)attractiveness of work:

1. Duty length (in hours): the length of a duty, defined as the difference in start and end time of the duty.

2. Percentage of type-A work: type-A work is considered more desirable, and mainly includes work on Intercity trains.

3. Percentage of aggression work: aggression work includes trips with a high probability of encountering passenger aggression, and is considered less desirable. Note that this attribute applies to guards only, and not to drivers.

The problem we aim to solve then consists of constructing feasible crew schedules throughout the planning period, at each point in time only using the released timetable information, such that the individual Sharing-Sweet-and-Sour rules of all crew members are satisfied at the end of the period.

## 3 METHODOLOGY

We propose to tackle this problem by solving a sequence of crew scheduling problems of fixed size. In particular, at time $t$ we solve the crew scheduling problem corresponding to period $[t + T_{MIN}, t + T_{MIN} + T_{SOLVE}]$. Here $T_{SOLVE}$ governs the size of the crew scheduling problems we tackle, and naturally we require that $T_{SOLVE} \leq T_{RELEASE} - T_{MIN}$. Based on the solution to this problem, we fix the crew schedules for the period $[t + T_{MIN}, t + T_{MIN} + T_{FIX}]$, where $T_{FIX} \leq T_{SOLVE}$. When fixing the schedules, we update the intermediate Sharing-Sweet-and-Sour scores of all crew members. We use the intermediate scores to penalise or reward the allocation of duties with certain attributes to specific crew members, thereby steering the crew schedules towards ones that satisfy the Sharing-Sweet-and-Sour rules. We then set $t = t + T_{FIX}$ and repeat this procedure until feasible crew schedules have been obtained for the complete planning period, thereby solving a total of $T_{PLAN}/T_{FIX}$ scheduling problems. The procedure is illustrated in Figure 1.

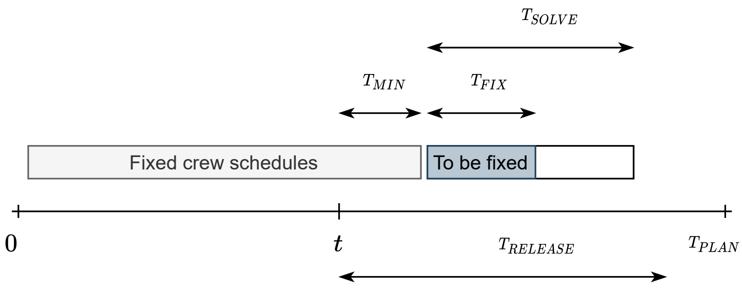


Figure 1 – *Illustration of sequential solution approach.*

Every single crew scheduling problem to be solved consists of a regular crew scheduling problem with side constraints. The side constraints govern the assignment of duties to templates, and potentially model any rostering constraints. For example, it might be necessary to incorporate a maximum workload constraint over the period $T_{SOLVE}$, to include forward rotation within a week, or to model potential overtime pay. The objective function of each problem is used to penalise or reward the allocation of certain duty attributes to crew members. In addition, a penalty cost might be used for any reserve crew members that are necessary to cover all tasks. All in all, the structure of the scheduling problems is highly similar to that of the integrated crew re-planning problem considered in Breugem *et al.* (2022), the main difference being the fact that our objective function contains penalties on duty attributes. As in their work, we use a column generation based diving heuristic to solve the crew scheduling problems.

Let us now illustrate how the objective function can be used to steer solutions towards satisfying the Sharing-Sweet-and-Sour rules. Suppose, without loss of generality, that each Sharing-Sweet-and-Sour rule specifies a lower bound $LB_j$ for a certain attribute $j$. Moreover, let $x_{ij}$ be the average score of the past duties of crew member $i$ with respect to attribute $j$ at some point in the moving horizon algorithm. We can then define a simple scoring function

$$f_j(x) = \max\left\{0, \alpha_j \cdot \frac{LB_j - x}{LB_j}\right\}, \tag{1}$$

where $\alpha_j$ is an attribute-specific weighting factor, and assign the value of $f_j(x_{ij})$ to all tasks in the duties of crew member $i$ that contribute towards attribute $j$ in the next crew scheduling problem. This will steer optimal solutions of the next problem towards ones where crew member $i$ performs duties that score high on attribute $j$, and will increase the likelihood that the Sharing-Sweet-and-Sour rules of this crew member are satisfied. Note that this objective-based approach is easily amendable to the pricing problems of the crew scheduling problem.

# 4 CASE STUDY AT NETHERLANDS RAILWAYS

## 4.1 Instance

We consider the guards of crew base Utrecht, the largest crew base of NS, during weeks 40 to 46 of 2021, a period in which the timetable was relatively unaffected by anti-covid measures. For each of the 219 guards of this crew base, we construct an individual template-based roster by splitting up the cyclic rosters for this period and extending it to seven weeks if necessary. Each duty in the original, cyclic roster yields a template with time window identical to that of the duty, but expanded with half an hour at the start and end. For each of the $T_{PLAN} = 49$ days of the planning period, we require that all tasks, which were originally covered by guards of crew base Utrecht, are covered. This corresponds to approximately 1,000 tasks per day of the planning period.

The Sharing-Sweet-and-Sour rules we consider are as follows. We require that the average duty length is at most 7.5 hours, that at least 30% of the working time is on Type-A trains, and that at most 18% of the time worked involves a higher risk of passenger aggression. The sample averages for these attributes are 7.2 hours, 32%, and 16%, respectively. Finally, we set $T_{SOLVE} = T_{FIX} = 1$ day. Using the column generation heuristic of Breugem *et al.* (2022), we can obtain crew schedules for the complete planning period in less than 1.5 hours of computation time.

### 4.2 Results

We now analyse the percentage of rosters that satisfies the individual Sharing-Sweet-and-Sour rules when we do not actively steer versus when we apply our feedback mechanism. In particular, we compare the results of applying our algorithm with $\alpha = 0$ (no steering) and $\alpha = 3,000$ (active steering). Note that we apply the same weighting factor to all three attributes. Without steering, the percentage of rosters that satisfies the duty length, Type-A, and aggression rule equals 72%, 69%, and 65%, respectively. With steering, these percentages increase to 97%, 90%, and 76%. Figure 2 compares the distribution of the roster type-A scores with and without steering. We find that steering reduces the score discrepancies between rosters, and steers the large majority of rosters above the Sharing-Sweet-and-Sour threshold.

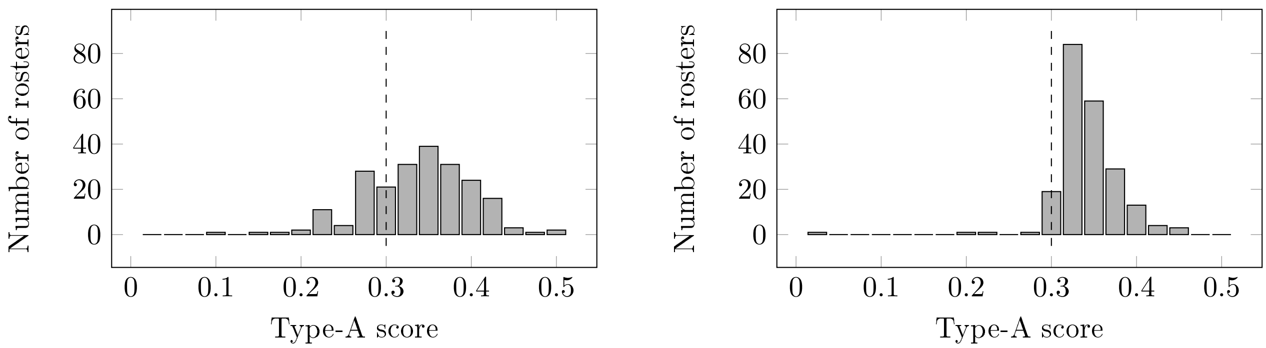


(a) $\alpha = 0$

(b) $\alpha = 3,000$

Figure 2 – *Distribution of roster type-A scores with (*$\alpha = 0$*) and without (*$\alpha = 3,000$*) steering mechanism. The dashed line indicates the lower bound of 0.3*

## 5 CONCLUSION

In this work, we formulated the operational crew scheduling problem with individual Sharing-Sweet-and-Sour rules. We proposed a sequential solution approach to solve the problem, and evaluated this on a real-life instance from NS. We find that our approach is able to significantly raise the percentage of rosters that satisfies individual Sharing-Sweet-and-Sour rules. Several interesting directions for future research remain. First, it would be interesting to employ a more sophisticated scoring function, that avoids assigning unattractive work to rosters that are currently just above an attribute bound. Second, one could attempt to include more (complex) roster constraints in order to make the problem more realistic. Third, look-ahead information from the individual rosters, possibly combined with predictions based on historical timetable data, could be incorporated in the scoring function to ensure timelier steering towards target attribute scores.

## References

Abbink, E, Fischetti, M, Kroon, L, Timmer, G, & Vromans, M. 2005. Reinventing crew scheduling at Netherlands Railways. *Interfaces*, **35**(5), 393–401.

Breugem, T, van Rossum, BTC, Dollevoet, T, & Huisman, D. 2022. A column generation approach for the integrated crew re-planning problem. *Omega*, **107**, 102555.