

The Modeling Power of the Periodic Event Scheduling Problem: Railway Timetables — and Beyond*

Christian Liebchen and Rolf H. Möhring

TU Berlin, Institut für Mathematik, Straße des 17. Juni 136, D-10623 Berlin
{liebchen,moehring}@math.tu-berlin.de

Abstract. In the planning process of railway companies, we propose to integrate important decisions of network planning, line planning, and vehicle scheduling into the task of periodic timetabling. From such an integration, we expect to achieve an additional potential for optimization.

Models for periodic timetabling are commonly based on the Periodic Event Scheduling Problem (PESP). We show that, for our purpose of this integration, the PESP has to be extended by only two features, namely a linear objective function and a symmetry requirement. These extensions of the PESP do not really impose new types of constraints. Indeed, practitioners have already required them even when only planning timetables autonomously without interaction with other planning steps. Even more important, we only suggest extensions that can be formulated by mixed integer linear programs.

Moreover, in a selfcontained presentation we summarize the traditional PESP modeling capabilities for railway timetabling. For the first time, also special practical requirements are considered that we prove not being expressible in terms of the PESP.

1 Introduction

Traditionally, the planning process of railway companies is subdivided into several tasks. From the strategic level down to the operational level, the most prominent subtasks are network planning, line planning, timetable generation, vehicle scheduling, crew scheduling, and crew rostering, see Figure 1.

For a detailed description of these planning steps, as well as for an overview of solution approaches, we refer to Bussieck, Winter, and Zimmermann [4]. Notice that network planning and line planning are of course part of the strategic planning process of public transportation companies. In contrast, vehicle scheduling and crew scheduling are of operational nature. In between, timetabling forms the linkage between service and operation. An important reason for the division into at least five subtasks is the high complexity of the overall planning process ([4], [7]).

* Supported by the DFG Research Center “Mathematics for key technologies” in Berlin.

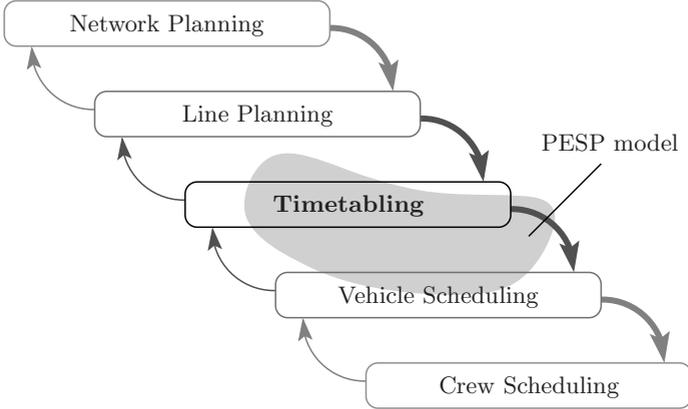


Fig. 1. Planning phases covered by the PESP beforehand

During the last years, a trend towards the integration of several planning steps has emerged. For example, vehicle and crew scheduling were successfully combined by Borndörfer, Löbel, and Weider [3] and by Haase, Desaulniers, and Desrosiers [8]. Similarly, a combination of line planning and network planning is the objective of Borndörfer, Grötschel, and Pfetsch [2]. Periodic timetabling has also served as a starting point for such attempts. Kolonko and Engelhardt-Funke [6] consider investments into infrastructure by using multi-criteria optimization. Nachtigall [20] computes timetables that require only few rolling stock for a specific vehicle schedule. Lindner [16] integrates the choice of rolling stock types in a non-linear model. Liebchen and Peeters [14] provide a linear model that serves as a good approximation for minimizing rolling stock while optimizing periodic timetables.

In this paper, we demonstrate how periodic timetable construction can be combined with other planning steps. Further, we incorporate other practical conditions on timetables such as timetable symmetry, line planning, and even infrastructure decisions. We show that this can in fact be achieved with only slight variations of the commonly used model for periodic timetable construction, the PESP model introduced by Serafini and Ukovich [30] in 1989. The variations keep much of the properties of the PESP model and are again mixed integer programs over a feasibility domain with essentially the same structure as the original PESP. In particular, all of the valid inequalities for the PESP stay valid, and some of the new formulations even speed up the solution time with standard MIP solvers. But there have also been proposed other solution techniques for PESP instances: constraint programming ([29]) and genetic algorithms ([21]). Hence, in this paper we will restrain ourselves to the pure modeling capabilities of the general PESP model—with only two small exceptions. But these exceptions have already been asked explicitly by practitioners for their own sake.

In the discussion of these modeling features, we will also lay out large parts of the map of the borderline between what still fits into the traditional PESP model,

and what requires new features, and at which cost. To this end, we also review the traditional PESP modeling issues, thus altogether providing a self-contained presentation of the PESP modeling capabilities and its extensions to symmetry, line planning, and network planning. Any of our suggestions for integrating these features can be formulated as a MIP, in particular not involving any quadratic terms.

The paper is organized as follows. Section 2 introduces the PESP. It presents its main formulations as a graph theoretic potential problem and as a mixed integer program, and reports on its complexity and a useful characterization of periodic timetables.

Section 3 discusses requirements for cyclic timetables that can be met by the PESP. These include simple requirements such as collision-free traffic on single tracks and headway between successive trains, but also more sophisticated ones such as bundling of lines, train coupling and sharing, fixed events in connection with hierarchical planning, and also disjunctive constraints and soft constraints.

Section 4 is devoted to timetable requirements that are beyond the scope of the traditional PESP, such as balanced reduction of service and symmetry of timetables. We show that the PESP or its MIP model only needs to be extended slightly in order to accommodate symmetry requirements.

Finally, in Section 5, we consider the integration of aspects of other planning steps into periodic timetable construction, in particular vehicle scheduling (minimization of rolling stock), line planning (simultaneous construction of line plan and timetable), and network planning (making infrastructure decisions). This integration makes essential use of the flexibility of the PESP, in particular disjunctive constraints, uses symmetry and, as a new technique, integrates aspects of graph techniques into the PESP in order to handle line planning.

All model features are illustrated by examples from our practical experience with timetable construction at Deutsche Bahn AG, S-Bahn Berlin GmbH, and BVG (Berlin Underground).

2 The Periodic Event Scheduling Problem (PESP)

In 1989, Serafini and Ukovich [30] introduced the Periodic Event Scheduling Problem (PESP), by which periodic timetabling instances may be formulated in a very compact way. Since then, this model has been widely used ([29,18,24,16,26]). In the Periodic Event Scheduling Problem (PESP), we are given a period time T and a set V of events, where an event models either the arrival or the departure of a directed traffic line at a certain station. Furthermore, we are given a set of constraints A . Every constraint $a = (i, j)$ relates a pair of events i, j by a lower bound ℓ_a and an upper bound u_a .

A solution of a PESP instance is a node assignment $\pi : V \mapsto [0, T)$ that satisfies

$$(\pi_j - \pi_i - \ell_a) \bmod T \leq u_a - \ell_a, \quad \forall a = (i, j) \in A, \quad (1)$$

or $\pi_j - \pi_i \in [\ell_a, u_a]_T$ for short. We call a feasible node potential π a feasible *timetable*. Notice that we can scale an instance such that $0 \leq \ell_a < T$, and for the *span* $d_a := u_a - \ell_a$ of a *feasible interval* $[\ell_a, u_a]_T$ we may assume w.l.o.g. $d_a < T$.

Furthermore, for every fixed event i_0 , every fixed point of time $t_0 \in [0, T)$, and every feasible timetable π there exists an equivalent timetable π' with $\pi'_{i_0} = t_0$. This is achieved by performing the simple shift $\pi'_i := (\pi_i - (\pi_{i_0} - t_0)) \bmod T$. Let us denote by $D = (V, A, \ell, u)$ the *constraint graph* modeling a PESP instance.

There are several practical aspects of periodic timetabling which profit from the presence of a linear objective function of the form

$$\sum_{a=(i,j) \in A} w_a \cdot (\pi_j - \pi_i - \ell_a) \bmod T,$$

with weights w_a . In our opinion, the most striking one is the integration of central aspects of vehicle scheduling, cf. section 5.1.

Another perspective on periodic scheduling can be obtained by considering tensions instead of potentials. In a straightforward way, define for a given node potential π its *tension*

$$\hat{x}_a := \pi_j - \pi_i, \forall a = (i, j) \in A.$$

We call a set of edges $C \subseteq A$ an *oriented cycle* if re-orienting a subset of its edges yields a directed circuit. The *incidence vector* γ_C of an oriented cycle C is a vector in $\{-1, 0, 1\}^A$, where the entry minus one indicates a backward arc of the oriented cycle. The cycle space \mathcal{C} of a directed graph D is defined as

$$\mathcal{C} := \text{span}\{\gamma_C \mid C \text{ oriented cycle in } D\}.$$

Recall that a vector \hat{x} is a tension (or potential difference), if and only if for some cycle basis B of \mathcal{C} , and each of its oriented cycles $C \in B$ with incidence vectors γ_C it holds that $\gamma_C \hat{x} = 0$ (e.g. [1]). This yields the following MIP formulation

$$\left. \begin{array}{l} \min c^t(\hat{x} + pT) \\ \text{s.t. } \Gamma \hat{x} = 0 \\ \ell \leq \hat{x} + pT \leq u \\ p \in \mathbb{Z}^A, \end{array} \right\} \text{ or } \left. \begin{array}{l} \min c^t x \\ \text{s.t. } \Gamma(x - pT) = 0 \\ \ell \leq x \leq u \\ p \in \mathbb{Z}^A, \end{array} \right\} \quad (2)$$

where $\Gamma \in \{-1, 0, 1\}^{(|A|-|V|+1) \times |A|}$ denotes the cycle-arc incidence matrix (*cycle matrix*) of some cycle basis of the directed graph D . The x variables are in fact a *periodic tension*, which we formally define for a given node potential π to be

$$x_{ij} := (\pi_j - \pi_i - \ell_{ij}) \bmod T + \ell_{ij}.$$

Sometimes, it is useful to define *slack variables* $\tilde{x}_a := x_a - \ell_a$.

Recall that cycle matrices are totally unimodular ([28]). This is the main observation to prove the following lemma.

Lemma 1 ([23]). *Let \mathcal{I} denote an instance of PESP with integral vectors ℓ and u and an integer period time T . If \mathcal{I} admits some feasible timetable $\pi \in [0, T)^V$, then it also admits an integral feasible timetable $\pi' \in \{0, \dots, T-1\}^V$.*

Already Serafini and Ukovich made the following simple but useful observation.

Lemma 2 (Serafini and Ukovich [30]). *If we relax the requirement $\pi \in [0, T]^V$ to $\pi \in \mathbb{Q}^V$, then for every spanning tree H and every feasible timetable π there exists an equivalent feasible timetable π' which induces $p_a = 0$ for $a \in H$.*

Notice that we may interpret the remaining non-zero integer variables as the representants of the elements of a (strictly) fundamental cycle basis. A generalization to integral cycle bases yields many variants of problem formulation 2, some of which are easier to solve for MIP solvers ([12]).

Periodic tensions can be characterized similarly to classic aperiodic tensions.

Lemma 3 (Cycle Periodicity Property). *A vector $x \in \mathbb{Q}^A$ is a periodic tension, if and only if for every cycle C with incidence vector $\gamma_C \in \{-1, 0, 1\}^A$, there exists some $z_C \in \mathbb{Z}$, such that*

$$\gamma_C x = z_C T. \quad (3)$$

The PESP is \mathcal{NP} -complete, since it generalizes Vertex Coloring ([23]). To see this, orient the edges of a Coloring instance arbitrarily and assign feasible periodic intervals $[1, T - 1]_T$ to each of them. Solution methods for the PESP include Constraint Programming ([29]), Genetic Algorithms ([21]), and of course integer programming techniques. For a computational study in that these substantially different approaches are compared to each other, we refer to [15]. For the MIP approach, a very important ingredient is

Theorem 1 (Odijk [24]). *An integer vector p allows a feasible solution for the MIP (2), if and only if for every oriented cycle C of the constraint graph, the following cycle inequalities hold*

$$\underline{p}_C := \left\lfloor \frac{1}{T} \left(\sum_{a \in C^+} \ell_a - \sum_{a \in C^-} u_a \right) \right\rfloor \leq \sum_{a \in C^+} p_a - \sum_{a \in C^-} p_a \leq \left\lceil \frac{1}{T} \left(\sum_{a \in C^+} u_a - \sum_{a \in C^-} \ell_a \right) \right\rceil =: \bar{p}_C, \quad (4)$$

where C^+ and C^- denote the forward and the backward arcs of the cycle C .

We close this section by listing other totally different practical applications which can be modeled via the PESP ([30]). The most prominent ones are the scheduling of systems of traffic lights, and periodic job shop scheduling.

3 Timetabling Requirements Covered by the PESP

This section gives a broad overview of the timetable modeling capabilities of the PESP. Contrary to the following sections, practical requirements to be modeled are limited to those arising in periodic timetabling. Nevertheless, there are many facts we have to discuss in order to give a self-contained overview.

However, let us start by naming two facts which are definitely beyond the scope of the PESP: routing of trains through stations or even alternative tracks, and routing of the passenger flow. Hence, throughout this paper we assume fixed routes for both trains and passengers. A short motivation for these assumptions will be given at the beginning of Section 4.

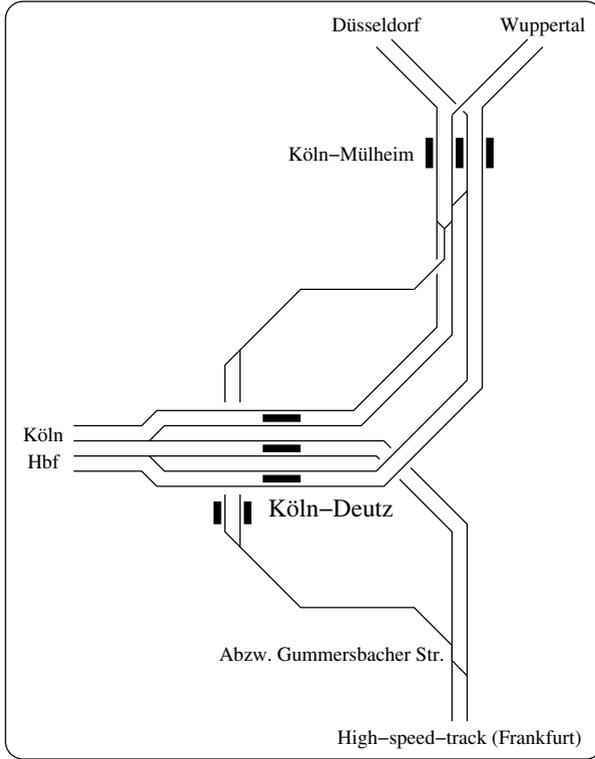


Fig. 2. Track map of Köln-Deutz (Cologne) — based on [11]

For the vast majority of practical requirements to be modeled, we provide examples which are close to practice. However, in particular time and track information might not always reflect practice exactly. Depending on the fact to be modeled, we provide a track map, a line plan, a visualization¹ of the timetable of a given track by means of a time-space diagram, and last but not least the resulting PESP subgraph. For readers not familiar with the first three types of charts, we refer to any textbook on railway engineering.

Most of our real-world examples are taken from the surroundings of the station Köln-Deutz (Cologne), which is part of the German ICE/IC-network. Figure 2 displays the general track map of Köln-Deutz. Unless stated otherwise, we assume a period time of $T = 60$ minutes.

3.1 Elementary Requirements

Both, for sake of completeness and in order to introduce the notation used in the following figures, we start by modeling the three most elementary actions within public transportation networks: trips, stops, and changeovers.

¹ In German: “Bildfahrplan.”

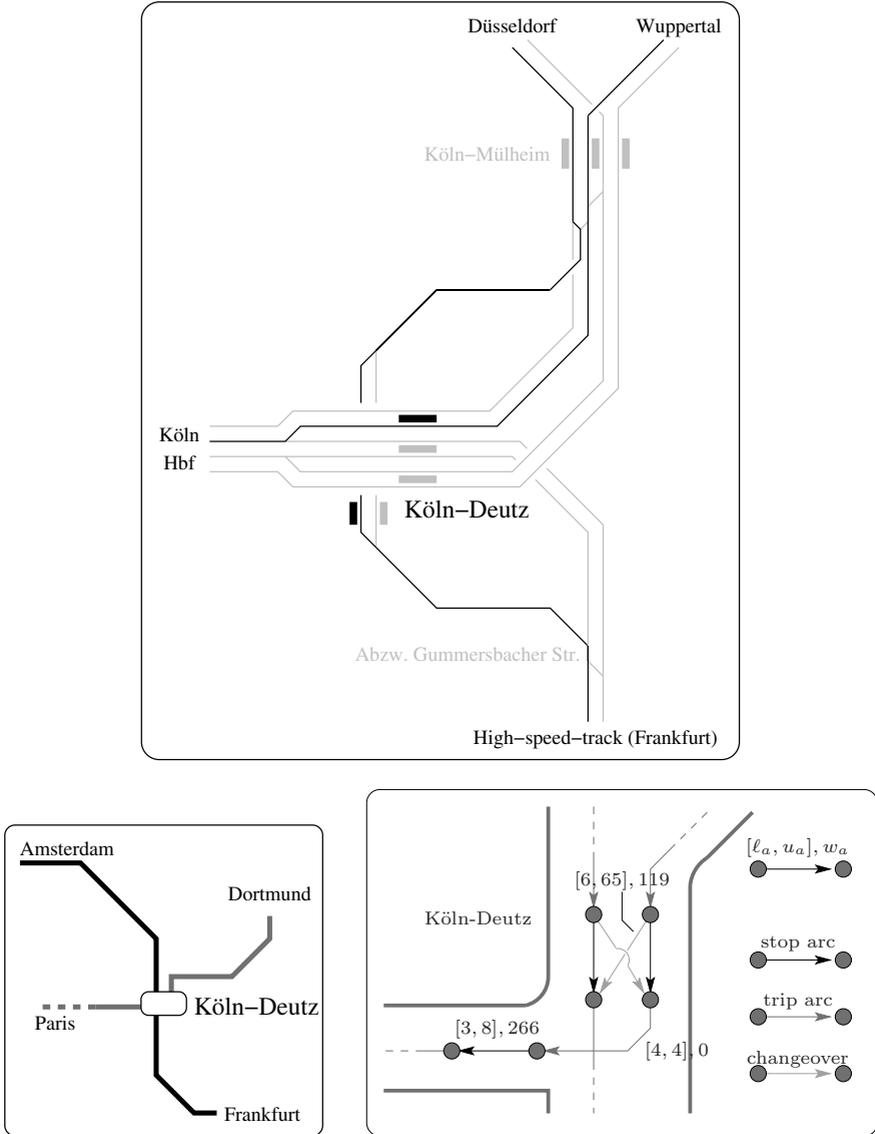


Fig. 3. Modeling elementary requirements: (a) two disjoint routes of lines serving Köln-Deutz; (b) the corresponding line plan; (c) PESP constraints modeling running activities, stopping activities, and changeover activities

In Figure 3 (a), we highlight the tracks used by two lines which cross at Köln-Deutz. The lines themselves are given in Figure 3 (b). Finally, we provide the constraint graph which models running, stopping, and changeover activities of these lines at Köln-Deutz in Figure 3 (c) as PESP constraints. For instance, the

trip arc with the constraint $[4, 4]_{60}$ ensures a trip time of precisely four minutes from Köln-Deutz to Köln Hbf. Within Köln Hbf, the minimum stopping time is set to three minutes such that passengers can board and alight the train. Finally, the increase of travel time for passengers that stay within the train is bounded by additional five minutes, providing an upper bound of $3 + 5 = 8$.

Notice that we ensure changeover quality by linearly penalizing changeover times which exceed a certain minimal changeover time required for changing platforms. In our example, a minimal changeover time of six minutes is assumed when connecting from Dortmund to Frankfurt. Using this approach, changeover arcs typically have a wide span.

An alternative way of modeling changeovers is to require some important ones not to exceed a maximal amount of effective waiting time. Then, we end up with rather small spans for changeover arcs. Schrijver and Steenbeek [29] follow this approach, which seems to be very suitable for constraint programming solvers.

Stopping arcs typically have very small span. In rather unimportant stations, in general it is a good choice to fix the span to zero, in particular if there is neither a junction of tracks, nor a single track, nor any changeovers.

Just as trip arcs, stopping arcs with span zero constitute redundancies which can be eliminated very efficiently in a preprocessing step. For example, one can contract any *fixed arc*, i.e. having zero span, together with its target node. Doing so, the arcs which were incident with the contracted target node only have to be redirected to the source node of the contracted arc, after having shifted their feasible intervals appropriately. Moreover, an arc being (anti-) parallel to another one can be eliminated, if its feasible interval is a superset of the other arc. In addition to nodes with degree at most two, Lindner [16] gives further situations in which the graph can be simplified.

If there are several lines using the same track into the same direction, sometimes a balanced service might be required. For n lines, this can easily be achieved by introducing arcs with feasible interval $[\frac{T}{n}, T - \frac{T}{n}]_T$ between any unordered pair of events that represent the departure at the first station of the common track. Certainly, strict balancedness may be relaxed by increasing the feasible interval.

Safety Requirements. If, in contrast to the previous discussion, there is no need for a balanced service, then at least a minimal headway h between any two of them has to be ensured. In the easiest case, the lines are operated with the same type of trains, and their running time is fixed. Then, we can sufficiently separate any two lines by introducing constraints similar to the above ones, having feasible interval $[h, T - h]_T$. These can be inserted either at the beginning or at the end of their common track. The more sophisticated constellation of trains involving different speeds will be discussed in Section 3.2.

But two trains may also use the same track in opposite directions. This is mainly the case for single tracks, see Figure 4 (a). Obviously, a train may not enter the single track until the train of the opposite direction has left it. In Figure 4 (b), we give a timetable visualization that is extremely useful in particular for single tracks. We assume a fixed local signaling, and the grey boxes visualize

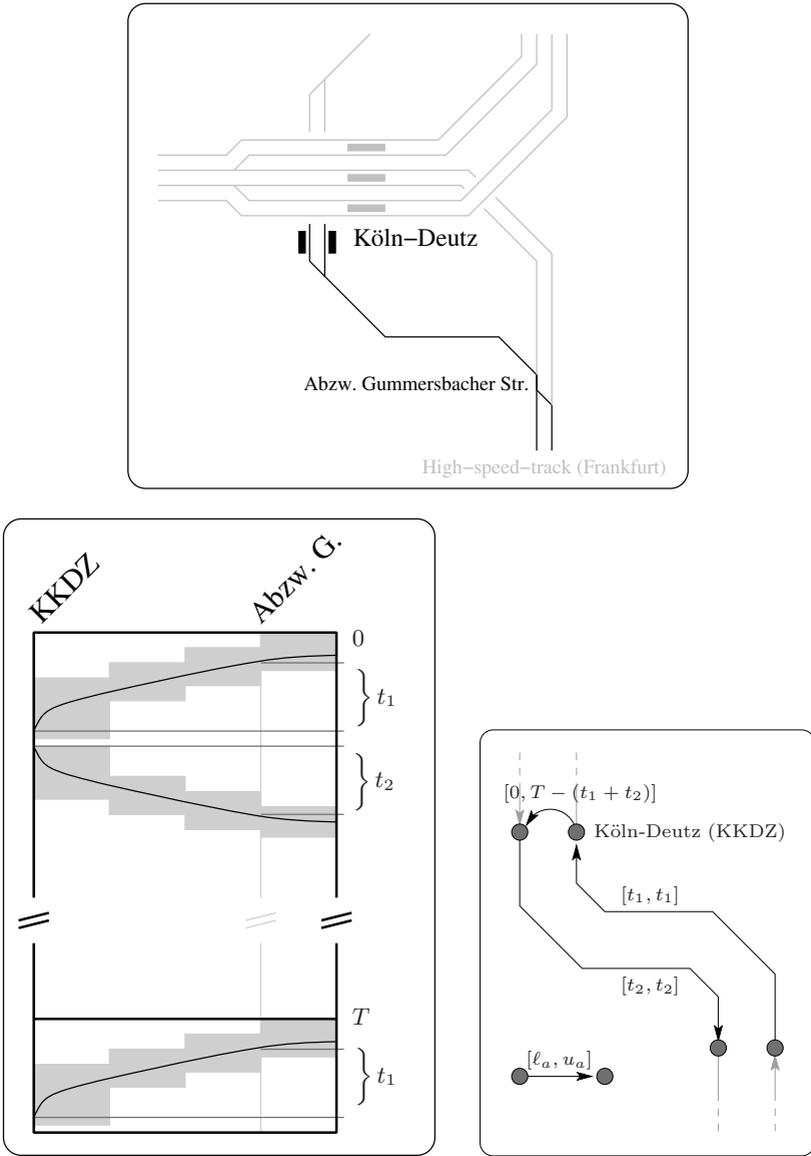


Fig. 4. Modeling single tracks: (a) a single track south of Köln-Deutz; (b) visualization of a feasible timetable for that single track; (c) PESP constraints ensuring safety distance for single track

the time a train blocks a certain part of the track. Surprisingly, there is only one single constraint needed to prevent two trains of opposite directions from colliding within the single track, as can be seen in Figure 4 (c). To that end, consider the western entry point to the single track. A train may only enter the

single track after a train of the opposite direction has left ($\ell_a = 0$). But it also must have left the single track before the next train of the opposite direction may enter the single track ($u_a = T - (t_1 + t_2)$).

Note that so far we did not care about any buffer times and blocking times when setting the feasible interval to $[0, T - (t_1 + t_2)]_T$. Assuming a *minimal crossing time* b at both endpoints of the single track, i.e. the time that has to pass from a train leaving the single track until a train in opposite direction may enter, we obtain the following feasible interval

$$[b, T - (t_1 + t_2 + b)]_T.$$

Again, if there are several lines that have to be scheduled on a single track, one constraint for every unordered pair of opposite directions is needed.

Some authors ([9]) consider situations at crossings, where trains are shortly using the track of the opposite direction (cf. Figure 5), as another modeling feature. But this is just a special case of single tracks, if the network is modeled at an appropriate granularity. Abzw. Gummersbacher Straße has to be split into a northern station and a southern station which are linked by an eastern and a western track, where the western track can be traversed in both directions.

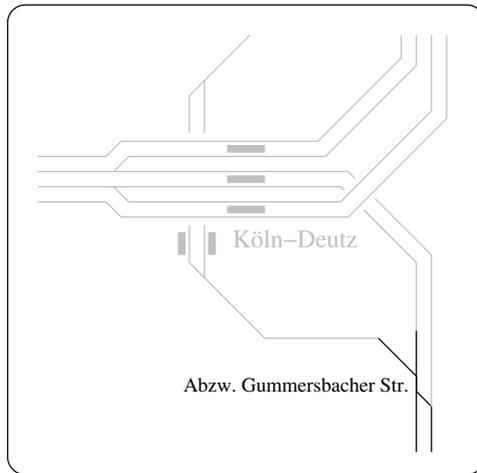


Fig. 5. Crossing of track of the opposite direction south of Köln-Deutz

3.2 More Sophisticated Requirements

Whereas the practical requirements discussed in the previous section might arise in almost every railway network, the following aspects are of a more specialized nature.

Fixed Events. When planning a timetable hierarchically, e.g. from international trains down to local trains, one has to consider the fixed settings of previous

hierarchies without replanning their times. Hence, the capability to fix an event to a certain point of time is another important modeling feature.

Fortunately, due to the periodic nature of the PESP, we may shift every feasible timetable such that a fixed event i_0 is fixed to a desired point in time $t_0 \in [0, T)$, i.e. $\pi_{i_0} = t_0$, and the objective value remains unchanged. By defining one of the events to be fixed as a kind of “anchor” event, we can easily relate the other events i_j to be fixed to certain points of time t_j by introducing arcs $a_j = (i_0, i_j)$ with $\ell_{a_j} = u_{a_j} = t_j - t_0$.

Bundling of Lines. Hierarchical planning gives rise to a further challenging aspect of timetabling. Notice that if a track is used by trains of different speeds, the capacity of that track significantly depends on the ordering of the trains. The first two parts of Figure 6 visualize this effect. In the first scenario, slow and fast trains alternate, which implies that only two hourly lines of each of the two train types can be scheduled. However, if lines are bundled with respect to their speeds, three lines of the same two types of trains can be scheduled without having to invest into infrastructure, cf. Figure 6 (b).

On the one hand, when only planning the high-speed lines in the first step of a hierarchical approach, it may happen that decisions on a higher level result in infeasibility on a lower level. On the other hand, hierarchical decomposition might have been chosen because an overall planning was considered to be too complex.

In order to keep the advantage of decomposition but limit the risk of infeasibility on lower levels, we propose to only bundle the lines of the current level of hierarchy. Figure 6 (c) gives the complete set of lines which should be operated on the track in question. In Figure 6 (d), we provide the PESP graph for the ICE/IC network. To bundle the three active lines, we introduce an artificial event and require each of the departure events to be sufficiently close to that artificial event. Hereby, the departure events will be close to each other as well.

In particular, we must not choose one of the existing events as “anchor”, because this would predict the corresponding line to be the head of the sequence of bundled lines. This must definitively be avoided, because — contrary to assumptions made by Krista [9] — the ordering of lines is indeed a major result of timetabling. Finally, based on profound estimates on passengers’ behaviour the management has to decide whether it is more important to operate as many trains as possible—and hereby bundle the trains of the same type—or whether a balanced service within the different types of trains should be preferred.

Train Coupling/Train Sharing. During the last decade, in railway passenger traffic a trend emerged towards train units which can easily be coupled and shared. Doing so, more direct connections can be offered without increasing the capacity of some bottleneck tracks.

In Figure 7 (a), we display a line which is operated by two coupled train units between Berlin and Hamm. They split in Hamm to serve the two major routes of the Ruhr area, hereby offering direct connections from Berlin to the most important cities of that region. Still, this line occupies for example the high-speed track between Berlin and Hannover only once per hour.

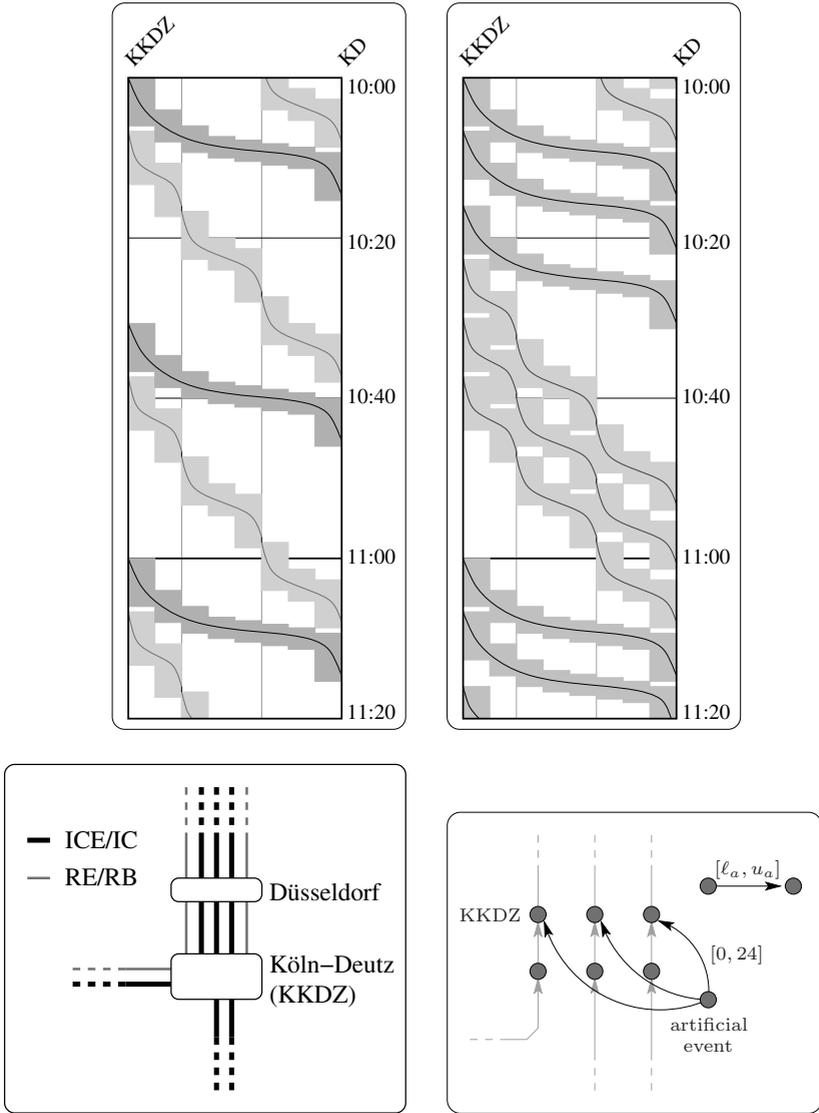


Fig. 6. Bundling of lines: (a) poor capacity if slow and fast trains are alternating; (b) capacity increase by bundling trains of the same type; (c) complete line plan for all the types of lines; (d) PESP constraints ensuring enough capacity for RE/RB lines already when planning only ICE/IC lines within the first step of a hierarchical planning

In Figure 7 (b), we provide PESP constraints which ensure the time for splitting the two train units in Hamm to be at least five minutes. Furthermore, for the two departing trains, a safety distance of four minutes is guaranteed.

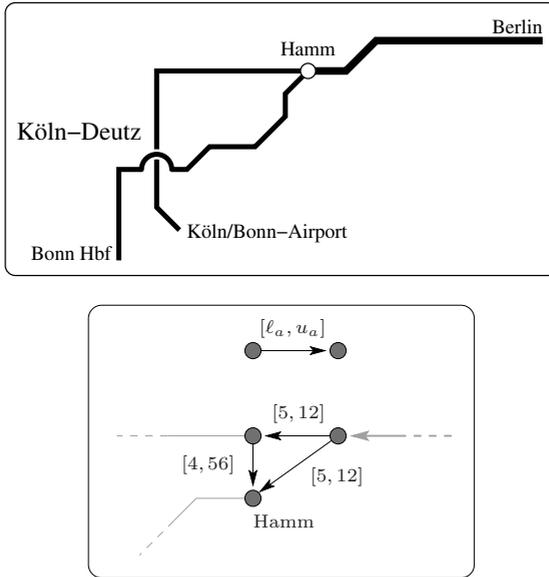


Fig. 7. Modeling train sharing: (a) line plan for the line Berlin-Hamm- $\{\text{Bonn Hbf} \mid \text{Köln/Bonn-Airport}\}$; (b) PESP constraints ensuring safety distance and time to split train units, but not specifying the ordering of departures

Notice that we do not need to specify which train should leave Hamm first. This decision will be made implicitly, and in an optimized way, by the PESP solver.

Variable Trip Times. As long as trip times are fixed, a usual safety constraint prevents two identical trains from overtaking each other. With h being the minimal headway for the track, we put an arc with feasible interval $[h, T - h]_T$ between the two events of entering the common track. If the line at the tail of the constraints is by f time units faster than the line at the tail of the constraints, overtaking can be prevented by modifying the constraint to $[h + f, T - h]_T$. This can be understood easily by having again a look at the corresponding situation in Figure 6 (a).

But this is no longer guaranteed if the model includes variable trip times. Even ensuring the minimal headway at the end of the track, too, does no longer prevent overtaking (even of trains of the same type) if the span in the trip times is at least twice the safety distance h , i.e. $u_a - \ell_a \geq 2h$. Schrijver and Steenbeek [29], Lindner [16], and Kroon and Peeters [10] tackle this phenomenon by adding extra constraints on the integer variables of the MIP formulations. Hereby, they leave the PESP model. In addition, Kroon and Peeters [10] provide some sufficient conditions on trip times, safety distance, and on the degree of flexibility of the trip times that prevent trains from overtaking.

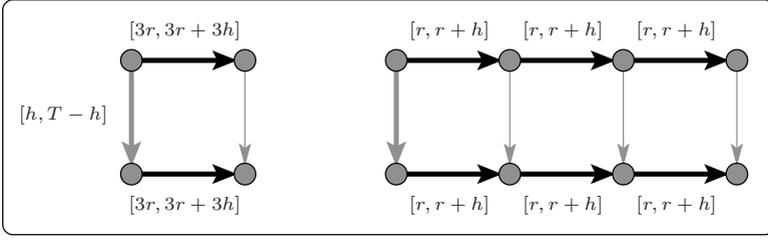


Fig. 8. Overtaking and variable trip times: (a) standard granularity does not prevent overtaking; (b) finer granularity prevents overtaking

In order to stay within the PESP model, we propose to subdivide² an initial trip arc into new smaller ones such that $u_a - \ell_a < 2h$ for every new trip arc. For an example, we refer to Figure 8, where **bold** arcs represent arcs of the spanning tree for which we set $p_a = 0$, cf. Lemma 2, and $3r$ is the minimum running time for the track.

Although this might seem to expand the model, the approach behaves rather well. More precisely, in every feasible timetable, the integer variables which we have to introduce for our additional arcs are in fact fixed to zero. This can simply be seen by applying the cycle inequalities (4) to any of the three squares in Figure 8 (b),

$$\underline{p} = \left\lfloor \frac{1}{T}(r+h - (T-h) - (r+h)) \right\rfloor = \left\lfloor \frac{h-T}{T} \right\rfloor = 0,$$

$$\bar{p} = \left\lceil \frac{1}{T}((r+h) + (T-h) - h - r) \right\rceil = \left\lceil \frac{T-h}{T} \right\rceil = 0.$$

Notice that the corresponding bounds for the initial formulation are only -1 and 1. But this is very natural, because there are three different types of timetables possible, of which we have to cut off two. The value one, for instance, models the fact that the second (lower) train is overtaking the first (upper) train.

Although we showed that the inconveniences caused by flexible running times can be overcome, we will assume fixed running times throughout the remainder of this paper.

3.3 General Modeling Capabilities

There are also important non-timetabling features which can be modeled by the PESP in a very elegant way. The types of such constraints are disjunctive constraints and soft constraints. Although they were originally introduced for their own sake, they turn out to be very useful for even more specialized requirements, which practitioners required to be modeled.

² This approach has also been discussed with Peeters [25,26] several years ago.

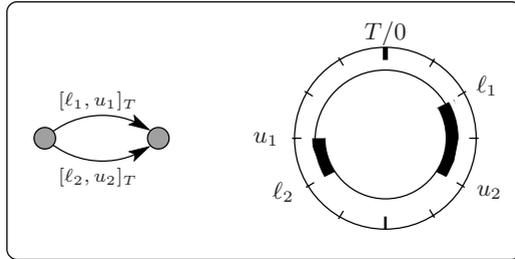


Fig. 9. Disjunctive constraints

Disjunctive Constraints. The feasible region of MIPs are commonly given as the intersection of finitely many half-spaces, plus some integrality conditions. If disjunctive constraints have to be modeled, usually artificial integer variables are introduced. However, the PESP offers a much more elegant way.

When introducing the PESP, Serafini and Ukovich [30] already made the important observation that the intersection of two PESP constraints is not always again a single PESP constraint. Rather, the feasible interval for a tension variable can become the *union* of two PESP constraints, e.g.

$$\pi_j - \pi_i \in [\ell_1, u_1]_T \cap [\ell_2, u_2]_T \Leftrightarrow \pi_j - \pi_i \in [\ell_1, u_2]_T \cup [\ell_2, u_1]_T.$$

We illustrate their observation in Figure 9. Nachtigall [20] observed that any union of k PESP constraints can be formulated as the intersection of at most k PESP constraints.

As an immediate practical application of disjunctive constraints, we consider optional operational stops. Long single tracks with no stop may cause the timetable of a line to be fixed within only small tolerances. In such a situation, Deutsche Bahn AG considers the option of letting the ICE/IC trains of one direction stop somewhere, although there is no ICE/IC station. In the current timetable, this takes place on the line between Stuttgart and Zurich, at Epfendorf.

If we want periodic timetable optimization to be competitive, we should enable the PESP to introduce an additional stop as well. We do so by introducing a pair of disjunctive constraints. The first constraint is a usual stop arc a_1 . We set the lower bound ℓ_{a_1} to zero, which models the option of not introducing an additional stop. The upper bound u_{a_1} is set to the sum of the minimal increase b of travel time occurring from braking and accelerating, plus the maximal amount of stopping time s at the station. For the effected increase \tilde{x}_a of travel time, this translates to

$$\tilde{x}_a \in \{0\}_T \cup [b, b + s]_T,$$

which is a disjunctive constraint. Notice that additional waiting time should be penalized in this situation similarly to an extension of a regular service stop. Moreover, if there are other lines operating on the same track, we have to take precautions that were discussed in the paragraph on variable trip times. However,

optional operational stops make most sense within long single tracks. But there, in many cases there are not several lines using that large bottleneck.

Obviously, the introduction of an additional stop can also be due to the construction of a new station. Since such decisions are a part of network planning, we postpone this discussion until Section 5.3.

Soft Constraints. Nachtigall [19] investigated the combination of two antiparallel arcs $a_1 = (i, j)$ and $a_2 = (j, i)$. If they have an identical coefficient in the objective function and if neither of them can become infeasible for any vector π , or x resp., then they model a *soft constraint*.

Classically, if a certain tension value x_a does not satisfy a given PESP constraint $[\ell_a, u_a]_T$, one would declare the complete timetable as infeasible. But sometimes, it can be an alternative only to produce a significant penalty in the objective function, if a constraint is not satisfied.

To that end, we relax the upper bound of the original constraint to $\ell + T - 1$ —we may assume the instance being scaled such that the precondition of Lemma 1 is satisfied. Further, we introduce a new antiparallel arc with feasible interval according to Figure 10. Then, these two constraints yield a piecewise constant

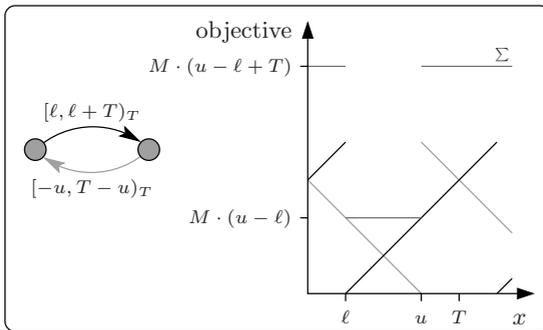


Fig. 10. Soft constraints

behavior of the objective function, which serves as an indicator for the violation of the original constraint, but without guaranteeing feasibility. For an initial constraint $x_a \in [\ell_a, u_a]$ consider the corresponding pair of artificial constraints a_1 and a_2 —each of these having having cost coefficient M . They contribute to the objective function

$$M \cdot (x_{a_1} + x_{a_2}) = \begin{cases} M \cdot (u - \ell) & \text{if } x_{a_1} \in [\ell_a, u_a]_T, \text{ and} \\ M \cdot (u - \ell + T) & \text{otherwise,} \end{cases}$$

hereby indicating whether the original constraint a is satisfied for the tension vector x .

In our cooperation with Berlin Underground, we were asked to construct a timetable that, among the top 50 most important connections, maximizes the

number of connections having a waiting time of at most five minutes. In fact, soft constraints are well-suited for letting MIP solvers produce a timetable being optimal subject to this kind of objective function.

4 Timetabling Requirements *Not* Covered by the PESP

Although the most important practical requirements for a periodic timetable can be modeled within the PESP, we are still aware of some special features for which the PESP fails. To the best of our knowledge this is the first time that practical requirements of timetabling are proven to be beyond the scope of the PESP.

First, one may think of situations in which it is not fixed which trains are operated on which track, for example within stations. Consider a station having two tracks in the same direction and three lines serving that direction. Then, we cannot decide a priori which pair of lines shall be within the station at the same time, hence omitting the sequencing constraint between these two lines. This observation is the motivation for the DONS system to be subdivided into CADANS, covering the timetabling step, and STATIONS, covering the routing aspect ([31]).

Notice that this requirement even affects the strategies for parking trains at terminus stations. Consider the following example, which has been inspired by the situation at Warschauer Straße of the Berlin fast train network, although there are further alternatives within that station. Within 20 minutes, two lines end at that station, both sharing the same track for arrival and departure, cf. Figure 11. For instance, arcs a_1 and a_2 measure the time that the two trains stay within this terminus station. We assume a turnover to take at least four

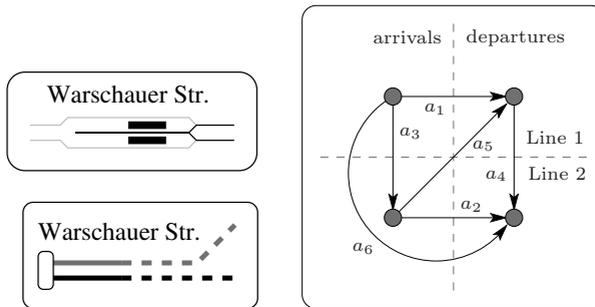


Fig. 11. Routing within terminus stations

minutes at the platform, or at least eight minutes when visiting the parking facility. Further, any arriving and departing trains block the platform for 59 seconds before and after their arrival and departure, respectively. To ensure that all passengers alighted before driving to the parking facility, the platform position is blocked for one more minute.

Proposition 1. *For every set of PESP constraints either timetables which are operable are classified as infeasible, or timetables which are not operable are classified as feasible.*

Proof. We start by analyzing the two major strategies individually: both lines turn at the platform, or line 2 turns in the parking facility, w.l.o.g. Table 1 provides tight lower and upper bounds for the six arcs in Figure 11 (c) with respect to these two scenarios. More precisely, with a strategy specified, we have that for every arc $a = (i, j)$ and every value $t_a \in [u_a, \ell_a]$ there exists an operable

Table 1. Tight interval bounds for different turning strategies at Warschauer Straße

Arc	Both at platform		Line 2 to parking		not specified	
	ℓ_a	u_a	ℓ_a	u_a	ℓ_a	u_a
a_1	4	12	4	13	4	13
a_2	4	12	8	18	4	18
a_3	6	14	6	17	6	17
a_4	6	14	2	14	2	14
a_5	10	18	7	18	7	18
a_6	10	18	6	18	6	18

timetable π such that

$$(\pi_j - \pi_i - \ell_a) \bmod T = t_a - \ell_a.$$

Further, by simple case inspection one can verify that every operable timetable which implements that specific strategy respects each of the given bounds. Hence, in order to provide general PESP constraints which characterize the operable timetables without having specified any parking strategy a priori, the feasible intervals must include the feasible intervals of both scenarios.

However, there exists a vector π which respects the six PESP constraints thus obtained (see the last two columns of Table 1), but which does *not* encode an operable timetable, because the two trains would be at the platform at the same time: line 1 arrives at minute 00 and departs only at minute 13, although line 2 already arrives at minute 06 and departs at minute 16. But for each of the $\binom{4}{2}$ potential differences between these four events there also exist operable timetables that attain the very same tension value. \square

Hence we cannot establish a set of PESP constraints that precisely identifies practically operable timetables as feasible solutions.

Apart from the rather important routing requirement, which unfortunately is simply out of scope for the PESP, we will analyze a very special situation in more detail, namely the balanced reduction of service. Finally, we will introduce the important notion of *symmetry*. On the one hand, symmetry slightly exceeds the original PESP, but on the other hand, when added explicitly, gives rise to a mechanism to include important aspects of line planning into the very same planning step as periodic timetabling and vehicle scheduling.

4.1 Balanced Reduction of Service

The Berlin fast train company (S-Bahn Berlin GmbH) aims at operating only one timetable for one whole day. The late evening service differs from the rush hour only in that some trains are omitted. Hence, the timetable must respect the available capacity during the rush hour, and it has to offer a balanced service in the late evening as well.

From a pure operations point of view, it could seem strange to sidestep an intraday change of the timetable structure. It is for sure that the information technology available in the 21st century could cope with this. But it is still the policy of the company. It is given as a motivation that customers really expect to have only one single timetable to be kept in mind for their station.

Consider the approximately 10 km long track from Zoo station to Berlin East station. On it, a minimal headway of 2.5 minutes has to be respected. The period time is 20 minutes and eight³ lines (having identical train types) per period and direction have to be scheduled. In the late evening service, there are four trains every 20 minutes, two of them being fixed to a 10 minutes time lag. We call these two lines *core-lines*.

Of course it would be ideal to have a five minutes time lag between two consecutive trains in the evening. But this is impossible because one of the evening trains is required to serve Potsdam every 10 minutes together with a rush hour train. Hence, one should ensure that the maximal time lag between two consecutive trains does not exceed 7.5 minutes.

But this simple requirement cannot be covered by the PESP. Consider the two types of timetables given in Table 2. Timetables of type 1 satisfy our requirement

Table 2. Possible timetables for the late evening service from Zoo station to Berlin East station. This table only shows the core-lines that are actually running in the evenings. Each of the – entries is a joker for a rush-hour train.

Timetable	Departure times ($T = 20$ minutes)								
Type 1	0.0	–	–	7.5	10.0	12.5	–	–	(20.0)
Type 2	0.0	2.5	–	7.5	10.0	–	–	–	(20.0)

by bounding the maximum distance between two consecutive trains to 7.5 minutes, but type 2 does not because there we have a gap of 10 minutes.

Proposition 2. *For every set of PESP constraints either timetables of both types are feasible, or timetables of both types are infeasible.*

Proof. There are two types of constraints to be analyzed:

- i. one constraint between the two non-core lines,
- ii. four constraints between one of the two core lines and one of the two non-core lines.

³ One of them only serves as a free slot for occasional non-passenger trips.

Since we must not specify the sequence of the lines in advance, only symmetric constraints $[\ell, T - \ell]_T$ make sense. Moreover, all constraints of type (ii) have to be identical for the same reason.

To guarantee feasibility of type 1 timetables, we deduce $\ell \leq 5$ for the constraint of type (i) and $\ell \leq 2.5$ for the constraints of type (ii). But then, timetables of type 2 stay feasible as well. Hence, in order to cut off timetables of type 2, we have to increment one of the given bounds. But since they are tight, this would immediately cut off timetables of type 1 as well. \square

Notice, however, that other railway companies implement other strategies to attain a balanced reduction of service. We will present an approach which turns out to be easier for timetabling, but slightly more complex for operation and customers.

Consider the track Niederh ochststadt-Langen (Hessen) via Frankfurt Hbf of S-Bahn Frankfurt. Compare the regular service hourly pattern to the weak-traffic service hourly pattern, which are given in Table 3. For the weak-traffic service,

Table 3. Timetables for regular service and weak-traffic service between Niederh ochststadt and Langen (Hessen)[17]

Line	regular service				weak traffic	
	S4	S3	S4	S3	S4	S3
Bad Soden	–	20	–	50	–	50
Kronberg	09	–	39	–	24	–
Niederh�ochststadt	14	29	44	59	29	59
Langen (Hessen)	56	11	26	41	11	41
Darmstadt Hbf	–	25	–	55	–	55

every second train is omitted. To prevent a 45 minutes gap every hour, one of the two lines is shifted by 15 minutes and uses the slot of the train of the other line, which has just been skipped.

If we assume none of the lines to share a track with other lines outside their common part, then we can easily deduce a feasible timetable for the weak-traffic service from a periodic timetable, which is feasible for the regular service. In case of single tracks along the peripheral segments, the only thing to be ensured is that the shift of 15 minutes appears simultaneously for the two directions. Hereby, every meeting point for the weak-traffic service is already a meeting point for the regular service — hence, single tracks stay respected. Trivially, along the common track no conflicts will appear either.

4.2 Symmetry of a Periodic Timetable

Throughout our discussion of symmetry, we assume that for every directed line there exists another directed line serving the same stations just in opposite order. Moreover, the concept of symmetry makes only sense, if for every traffic line, the running and stopping times of its two opposite directions are the same. Also

for the minimum headways and other operational constraints we require them to be identical in both directions. Furthermore, the passenger flow is assumed to be symmetric.

First, observe that in every periodic timetable with period time T , every train meets some train of the opposite direction of its line twice within the period time. In general, every line can have different times for these meetings.

A periodic railway timetable is called *symmetric with (global) axis s* , if at time s every train in the network meets a train of the opposite direction of its line. From the above considerations we deduce that we may assume w.l.o.g. $s \in [0, \frac{T}{2})$.

For the arrival or departure event of a directed line at a certain station, we denote by its *complementary event* the departure or arrival, resp., of the opposite line at the same station. In the sequel, we provide two characterizations of symmetric timetables.

Lemma 4. *A timetable is symmetric with axis s , if and only if for every pair i and \bar{i} of complementary events there holds*

$$\frac{(\pi_i + \pi_{\bar{i}}) \bmod T}{2} = s. \quad (5)$$

Proof. Let i and \bar{i} be any two complementary events. By definition, they are part of the two opposite directions of the same line. Moreover, they are located in the same station S .

In a symmetric timetable, the trains of the two opposite directions meet at times s and $s + \frac{T}{2}$. Consider two virtual events j and \bar{j} of passing the meeting point M . As the trains meet there, we have $\pi_j = \pi_{\bar{j}} \in \{s, s + \frac{T}{2}\}$.

We assumed the travel times of two opposite trains to be identical and denote the travel time between S and M by t . Hence, w.l.o.g.

$$(\pi_i + \pi_{\bar{i}}) \bmod T = ((\pi_j + t) + (\pi_{\bar{j}} - t)) \bmod T = (2 \cdot \pi_j) \bmod T. \quad \square$$

To define a counterpart of condition (5) for the tension formulations (2), we define two arcs $a = (i, j)$ and $\bar{a} = (\bar{j}, \bar{i})$ to be *complementary*, if $\{i, \bar{i}\}$ and $\{j, \bar{j}\}$ are complementary, and we have $\ell_a = \ell_{\bar{a}}$ and $u_a = u_{\bar{a}}$. With these definitions at hand, we are able to define a symmetric instance of PESP: A constraint graph is called *symmetric*, if every arc connects either two complementary events, or if for every arc $a \in A$ there exists some complementary arc $\bar{a} \in A \setminus \{a\}$.

Lemma 5. *Consider an instance of PESP that is modeled by a connected symmetric constraint graph. Let π be a feasible timetable with corresponding periodic tension x . There exists some $s \in [0, \frac{T}{2})$ such that Condition (5) holds for every pair of symmetric events, if and only if every pair of complementary arcs a and \bar{a} fulfills*

$$\tilde{x}_a = \tilde{x}_{\bar{a}}. \quad (6)$$

Proof. “ \Rightarrow ”: Let $a = (i, j)$ and $\bar{a} = (\bar{j}, \bar{i})$ denote two complementary arcs of the constraint graph. Then, we have

$$\begin{aligned} \tilde{x}_a &= x_a - \ell_a \stackrel{(2)}{=} (\pi_j - \pi_i - \ell_a) \bmod T \\ &\stackrel{(5)}{=} (2s - \pi_{\bar{j}} - (2s - \pi_{\bar{i}}) - \ell_a) \bmod T \\ &= (\pi_{\bar{i}} - \pi_{\bar{j}} - \ell_{\bar{a}}) \bmod T = x_{\bar{a}} - \ell_{\bar{a}} = \tilde{x}_{\bar{a}}. \end{aligned}$$

“ \Leftarrow ”: Let x be the periodic tension of some feasible timetable π . We show that there exists one global symmetry axis s such that Condition (5) is satisfied for π .

We compute s from an arbitrary fixed event, say i ,

$$s := \frac{(\pi_i + \pi_{\bar{i}}) \bmod T}{2}.$$

Now, we consider an arbitrary pair of complementary events j and \bar{j} . Since D is connected and symmetric, there exists a path P from i to j or \bar{j} that only contains arcs a such that $\bar{a} \in A \setminus \{a\}$. We assume w.l.o.g. that P starts at i and ends at j . By setting

$$x_P := \sum_{a \in P^+} x_a - \sum_{a \in P^-} x_a,$$

we obtain $\pi_j = (\pi_i + x_P) \bmod T$. As for every $a \in P$ there exists its complementary arc $\bar{a} \in A \setminus \{a\}$, the complementary path \bar{P} of P from \bar{j} to \bar{i} is well-defined. Equation (6) ensures $x_{\bar{P}} = x_P$.

In total, we obtain

$$\frac{(\pi_j + \pi_{\bar{j}}) \bmod T}{2} = \frac{(\pi_i + x_P + \pi_{\bar{i}} - x_{\bar{P}}) \bmod T}{2} = \frac{(\pi_i + \pi_{\bar{i}}) \bmod T}{2} = s. \quad \square$$

Remark 1. If the line plan of a traffic network is connected and the constraint graph is symmetric, we are able to give an even more compact characterization of symmetry. Then, a feasible tension encodes a symmetric timetable, if and only if Condition (6) is satisfied for changeover arcs and stopping arcs. In fact, in the proof of Lemma 5 we can then find a path that only uses such arcs, plus trip arcs, which we assume to have zero span.

Surely, one can introduce a certain tolerance Δ on the symmetry requirement. But notice that in this case, condition (6) has to be blown up by a new integer variable.

Example 1 (Deutsche Bahn AG). Figure 12 shows two real-world timetable queries for opposite directions. These are representative for large parts of central European countries, such as Germany and Switzerland, which are operated with symmetry axis zero within only minor tolerances. Hence, if not stated otherwise we assume $s = 0$ throughout this paper for ease of notation.

We check the three characterizations of symmetry. Most striking, the changeover waiting time is almost the same in both directions, cf. Remark 1 and

Station/Stop	Date	Time	Platform	Products	Comments	
Berlin Zoologischer Garten	05.06.03	dep 09:54	4	ICE 952	InterCityExpress BordRestaurant	
Wolfsburg		dep 10:54				
Hannover Hbf		dep 11:31				
Bielefeld Hbf		dep 12:24				
Hamm(Westf)		dep 12:54				
Hagen Hbf		dep 13:25				
Wuppertal Hbf		dep 13:42				
Köln-Deutz		dep 14:11				
Köln Hbf		arr 14:14				6
Köln Hbf		05.06.03				dep 15:13
Aachen Hbf	dep 15:52					
Aachen Süd(Gr)	05.06.03	arr 17:46				
Liege-Guillemins						
Bruxelles-Midi						

Duration: 7:52; runs daily

All information is issued without liability. Software/Data: HAFAS 5.00.DB.4.5 - 20.05.03 [5.00.DB.4.5/v4.05.p0.13_data:59e79704]

Station/Stop	Date	Time	Platform	Products	Comments	
Bruxelles-Midi	05.06.03	dep 12:16		ICE 15	InterCityExpress Onboard meeting place	
Liege-Guillemins		dep 13:28				
Aachen Süd(Gr)	05.06.03	dep 14:10	3			
Aachen Hbf		arr 14:46				
Köln Hbf		arr 14:46				
Köln Hbf	05.06.03	dep 15:47	2	ICE 953	InterCityExpress BordRestaurant	
Köln-Deutz		dep 15:51				
Wuppertal Hbf		dep 16:17				
Hagen Hbf		dep 16:35				
Hamm(Westf)		dep 17:10				
Bielefeld Hbf		dep 17:37				
Hannover Hbf		dep 18:31				
Wolfsburg		dep 19:05				
Berlin Zoologischer Garten		arr 20:02				1

Duration: 7:46; runs Mo - Fr, not 29. May, 9. Jun, 21. Jul, 15. Aug, 11. Nov

Hint: Prolonged stop

All information is issued without liability. Software/Data: HAFAS 5.00.DB.4.5 - 20.05.03 [5.00.DB.4.5/v4.05.p0.13_data:59e79704]

Fig. 12. Symmetric timetables in practice

Equation (6). To check Condition (5), we consider the arrival of ICE 952 in Köln Hbf and the complementary departure of ICE 953. The two events sum up to $(14 + 47) \bmod 60 \approx 0$, and the same can be observed for the Brussels trains. Finally, notice that the Berlin line has one of its meeting points between Köln-Deutz and Wuppertal Hbf, at minute zero, of course. To that end, we have to know that the trains from Berlin arrive at Köln-Deutz at minute 09, which is two minutes before its departure at minute 11.

Some practitioners consider the changeover condition in Remark 1 to be an important advantage of symmetric timetables. Even though this might depend on personal preferences, we do *not* consider this really to be a striking argument for symmetry. Actually, there are examples which prove that symmetric timetables are only suboptimal, even if the input data is symmetric ([13]).

Apparently there are not yet many discussions of symmetric timetables available. But among further motivations for symmetry, as they can be found in [13], the most convincing one seems to be that symmetry halves the complexity of an instance. This can in particular be useful if there are complex interfaces to

international trains or to regional traffic, and when planning is performed manually. However, this argument should become less important in the future, as we think that PESP solvers achieve some more progress in performance, and hence find their way into practice.

With the following theorem, we are able to prove a conjecture that has been stated in [13].

Theorem 2. *Symmetry of periodic timetables cannot be guaranteed by only using PESP constraints (1).*

Proof. Consider the PESP instance in Figure 13 (b). The PESP constraints that relate the two opposite directions of the line to be considered model the two single tracks of the track map that is shown in Figure 13 (a). The minimum

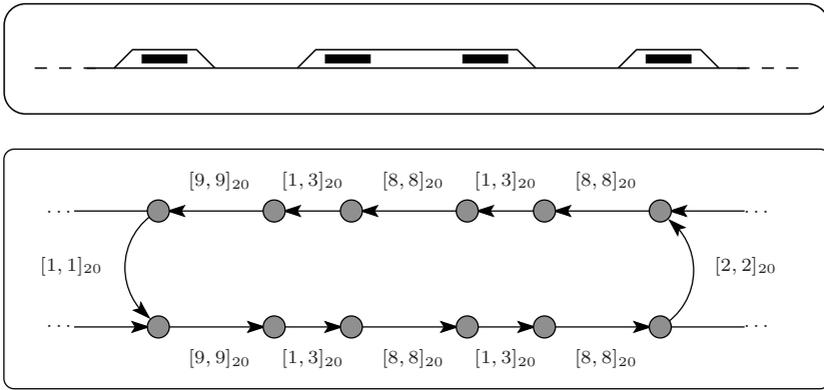


Fig. 13. A track map (a) on which an instance of PESP (b) does not admit any integral symmetric solution

crossing times (cf. Section 3.1) that apply to a certain single track depend on the infrastructure and the signaling system. For the western single track, we assume minimum crossing times of one time unit at both of its endpoints, for the eastern single track we assume two time units at both of its endpoints. Hence, given a period time of $T = 20$ time units and the indicated one-way running times of eight and nine time units, the single track constraints become tight.

Summing up the lower bounds of the constraints of the directed cycle yields 57, summing up the upper bounds provides 63. Hence, there exist feasible timetables. Moreover, due to the cycle periodicity property (Lemma 3), we know that in each of the feasible solutions the tension values sum up to 60. Hence, a slack of three time units has to be distributed on the four arcs with positive span.

In every symmetric feasible timetable, both of the directions obtain 1.5 time units of slack, hereby implying non-integral tension values. In contrast, by Lemma 1 every feasible system of PESP constraints (1) admits a feasible integral timetable. \square

Hence, we will have to add non-PESP constraints to the MIP formulations of a PESP instance in order to ensure symmetry. This is really required in practice, because in particular with national railway companies, we gained the experience that the symmetry requirement is really a knockout criterion.

To summarize, besides a linear objective function, symmetry is the second important requirement arising in the practice of periodic railway timetabling, by which the initial PESP model should be extended. Fortunately, in computations on real-world data sets it has been observed that MIP solvers may profit from the addition of symmetry constraints, in particular in formulation (6) ([13]). Such a generalized MIP model even inherits large parts of the structure of a pure PESP model. Most important, the cycle inequalities (4) remain valid.

5 Further Planning Steps Covered by the PESP

In the following, we will demonstrate that the modeling capabilities of the PESP are not limited only to periodic timetabling. Rather, central aspects of both preceding and succeeding planning steps in the sense of Figure 1 can be integrated.

We start this discussion with the well-established technique of minimizing the number of vehicles required to operate a periodic timetable by penalizing waiting times of vehicles. Hereafter, we provide first ideas for the integration of important decisions of line planning. We close this section by proposing a way to model some specialized decisions arising in network planning.

5.1 Aspects of Vehicle Scheduling

Almost all companies in public transportation have in common that they want to minimize the amount of rolling stock required to serve their networks. Notice that the quality of the vehicle schedule for a fully periodic timetable, i.e. with no peak trips included, is largely determined by the timetable.

Consider for example the hourly line displayed in Figure 14 (a). Assume the minimal travel times between the two endpoints to be 235 minutes for each direction. Given strict minimal turnover times of 45 and 60 minutes, respectively, the minimal number of vehicles required to operate this line is precisely

$$N := \left\lceil \frac{1}{60}(235 + 235 + 45 + 60) \right\rceil = 10.$$

A timetable which lets the trains leave at the full hour from Frankfurt and Amsterdam can indeed be operated with only 10 trains, at least if the stopping times are extended only moderately. On the contrary, a timetable in which only the trains starting at Frankfurt depart at minute 00, but the trains from Amsterdam leave at minute 30 requires at least 11 vehicles. Hence, the amount of vehicles depends on the timetable.

We will analyze in which special cases pure PESP constraints are able to control the number of trains required. After that, we show that a linear objective function covers many more of the practical cases.

Proposition 3 (Nachtigall [20]). *Consider a fixed traffic line with period time T . If we assume trains always to serve only this line, and if we do not allow to insert additional stopping time, then there exist upper bounds u for the turnover activities, such that the only feasible timetables are those which can be operated with the minimal amount of trains.*

Proof. We present a proof of this simple fact, both in order to provide the notation used in the following paragraphs, and because it avoids modulo-notation.

Denote the endpoints of the line by A and B . Let ℓ_{AB} denote the minimal travel time from A to B , i.e. the sum of the minimal stopping and running times of the activities of this directed traffic line. Moreover, denote by ℓ_B the minimal amount of time a train has to stay in endpoint B between two consecutive trips.

The minimal number N of trains required to operate this line is precisely

$$N = \left\lceil \frac{\ell_{AB} + \ell_B + \ell_{BA} + \ell_A}{T} \right\rceil.$$

From the cycle periodicity property (3) we know that every feasible timetable x fulfills

$$x_{AB} + x_B + x_{BA} + x_A = zT, \quad (7)$$

for some $z \in \mathbb{Z}$. Hence, we must ensure $z = N$. To that end, consider the slack

$$\sigma := NT - (\ell_{AB} + \ell_B + \ell_{BA} + \ell_A) \quad (8)$$

of this traffic line, implying $(x_A - \ell_A) + (x_B - \ell_B) = \sigma$. But since $\sigma < T$, by setting

$$u_A := \ell_A + \sigma \quad (9)$$

we even ensure $x_{AB} + x_B + x_{BA} + x_A < (N + 1)T$. \square

Let us now analyze the case in which additional stopping times may be inserted, i.e. $u_{AB} > \ell_{AB}$. We will show that together with the constraints (9), some timetables which require an additional train may become feasible.

On the one hand, consider a timetable for which we have $x \equiv \ell$ for all activities, except for the turnover time in one endpoint. This timetable can still be operated with the minimal number of trains, showing that decreasing the value (9) for u_A would cut off timetables we seek for.

On the other hand, assume $x_{AB} = u_{AB}$ and $x_{BA} = u_{BA}$. If

$$(u_{AB} - \ell_{AB}) + (u_{BA} - \ell_{BA}) + \sigma \geq T, \quad (10)$$

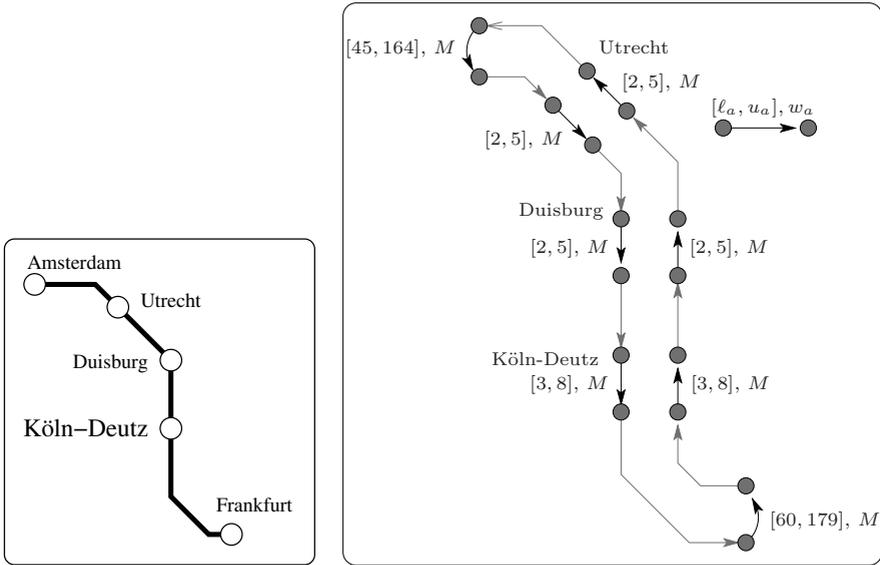


Fig. 14. Modeling aspects of vehicle scheduling: (a) line plan; (b) PESP constraints measuring the number of trains required to operate the line

then we can extend x to a timetable that still respects (9), but which requires at least one additional train. For instance, if inequality (10) is tight, then for $x \equiv u$ we have

$$\begin{aligned}
 x_{AB} + x_B + x_{BA} + x_A &= u_{AB} + u_B + u_{BA} + u_A \\
 &\stackrel{(9)}{=} (u_{AB} - \ell_{AB}) + (\ell_B + \sigma) + (u_{BA} - \ell_{BA}) + \\
 &\quad + (\ell_A + \sigma) + \ell_{AB} + \ell_{BA} \\
 &\stackrel{(10)}{=} T + \sigma + \ell_{AB} + \ell_B + \ell_{BA} + \ell_A \\
 &\stackrel{(8)}{=} (N + 1)T.
 \end{aligned}$$

The above dilemma is our main motivation for the need of a linear objective function. Such a function takes advantage of equation (7): By assigning a value M to the arcs modeling a traffic line, every additional train pays $M \cdot T$ to the objective function value. Of course, it suffices to consider arcs with positive span, cf. Figure 14 (b). If the value for M is chosen relatively large compared to the passenger weights, the objective function essentially models the piecewise constant behavior of the cost of the rolling stock for operating the railway network.

From a more local perspective, we just penalize idle time of trains. But this can even be done without knowing a priori the circulation plan of the trains. Although a straight-forward exact model involves a quadratic objective function, Liebchen and Peeters [14] report that a simple linear relaxation in terms of the PESP yields results of high quality.

Very recently, Nyhave, Hove, and Clausen [22] proposed an integer linear model to precisely count the number of trains required to operate a timetable, even if trains are allowed to switch lines in their endpoints. This approach does not depend on additional assumptions as synchronization constraints or pre-defined time-windows for turnaround times, as they were used by Peeters [26]. In the sequel, we translate their ideas into the PESP plus some additional variables and constraints.

Consider a station S that is a terminus for the two lines 1 and 2. Denote by a_i and d_i the arrival and departure events in station S of line i . We introduce the following arcs

$$\begin{aligned} a_{11} &= (a_1, d_1) \quad \text{and} \quad a_{22} = (a_2, d_2), \\ a_{12} &= (a_1, d_2) \quad \text{and} \quad a_{21} = (a_2, d_1). \end{aligned}$$

The effective waiting times for the trains in S are $\tilde{x}_{11} + \tilde{x}_{22}$ if trains stay on their lines, or $\tilde{x}_{12} + \tilde{x}_{21}$ if trains switch lines. Notice that $(a_{11}, a_{21}, a_{22}, a_{12})$ is an oriented cycle. In particular, there exists some $z \in \mathbb{Z}$ such that

$$x_{11} + x_{22} = x_{12} + x_{21} + zT.$$

In most cases, we have $\ell_{a_{11}} = \ell_{a_{12}}$ and $\ell_{a_{21}} = \ell_{a_{22}}$. Then, we even know that there exists some $r \in [0, T)$ and $b_1, b_2 \in \{0, 1\}$ such that

$$r = \tilde{x}_{11} + \tilde{x}_{22} - b_1 \cdot T = \tilde{x}_{12} + \tilde{x}_{21} - b_2 \cdot T.$$

Hence, in an optimal vehicle schedule the total effective waiting time in station S amounts to $r + \min\{b_1, b_2\} \cdot T$. To obtain a MIP-formulation, we introduce a new (rational) variable w and require

$$w \geq b_1 + b_2 - 1 \quad \text{and} \quad w \geq 0.$$

Finally, station S contributes

$$M \cdot (r + w \cdot T)$$

to the objective function, where M again denotes the cost factor for vehicle waiting time.

5.2 Aspects of Line Planning

Our main idea for letting PESP solvers even take decisions of line planning is to combine — or match — pre-defined line-segments. To that end, we will make intensive use of disjunctive constraints. Unfortunately, we will only be able to ensure symmetric line plans if we require symmetry also within the stations where lines are matched.

We are aware of only one other approach for integrating the planning phases of line planning, timetabling and vehicle scheduling ([32]). Whereas that approach is based on the assumption that the line plan contains no cycles, our ideas do not

require any restrictive assumptions on the topology of the network. Rather, we are able to keep even very important technical restrictions such as single tracks.

Notice that bad decisions at the level of line planning may cause very bad results also for vehicle scheduling. Consider the four line segments displayed in Figure 15. We assume a period time of $T = 60$ minutes and a minimal turnover time of 30 minutes at each of the four terminus stations. The time for a one-way trip from the matching station to one of the endpoints is indicated at the corresponding edge.

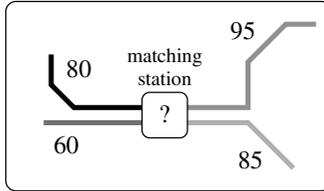


Fig. 15. Line segments where only one matching provides good vehicle schedules

In fact, the vehicle schedule is fixed due to the distinct endpoints. Combining the south-west segment with the north-east segment causes this line to require at least

$$\left\lceil \frac{1}{60}(60 + 95 + 30 + 95 + 60 + 30) \right\rceil = \left\lceil \frac{370}{60} \right\rceil = 7 \text{ trains.}$$

The other line of the same matching requires seven trains, too.

In contrast, the other matching implies seven trains only for the northern line consisting of the two top line segments. But the other line can be operated with only six trains. Hence, already the line plan has a major impact on the cost of operation. Claessens et al. [5] consider this phenomenon in their approach for constructing cost-optimal line plans.

However, they omit the important intermediate linking step of computing a timetable. Therefore, their approach must also consider possible constellations in which there is no feasible timetable using only six trains for the southern line. This would be the case, if there was a single track with travel time 25 minutes for every direction just at the end of the south-east segment. The same holds if it is required that the two lines together form an exact half-hourly service along the backbone of the network.

We consider a track that has to be served in the same direction by n directed lines which are operated by trains of identical type. We denote the matching station by S which resides between the two endpoints of the common track. We consider n line segments L_1^a, \dots, L_n^a which have station S as their common endpoint, and n line segments L_1^d, \dots, L_n^d having station S as their common starting point. Any (bipartite) perfect matching between the arriving and the departing line segments induces a line plan.

But from the perspective of timetabling, there are only n arrival events a_1, \dots, a_n as well as n departure events d_1, \dots, d_n visible. Hence, we must deduce only from their arrival times π_{a_i} and their departure times π_{d_j} which arriving line segment L_i^a should be matched with which departing line segment L_j^d .

This can be done in a canonical way, if we choose the matching station S such that it has only one track in the direction of the line segments we consider. If necessary, we add an artificial station in the middle of some track. Then, at most one train can be in S at the same time. Timetables respecting this constraint can be characterized very easily as follows.

Definition 1 (Alternating timetable). *For a fixed station S and a fixed direction, a periodic timetable π with n pairwise different arrival times $0 \leq \pi_{a_1} < \dots < \pi_{a_n} < T$ and n pairwise different departure times $0 \leq \pi_{d_1} < \dots < \pi_{d_n} < T$ is called alternating at S , if either $\pi_{a_i} \leq \pi_{d_i} < \pi_{a_{i+1}}$ for every $i = 1, \dots, n$, or $\pi_{d_i} < \pi_{a_i} \leq \pi_{d_{i+1}}$ for every $i = 1, \dots, n$, where we define $\pi_{n+1} := \pi_1 + T$.*

Lemma 6. *A timetable π ensures that there is always at most one train at station S if and only if it is alternating at S .*

Hence, for an alternating periodic timetable, we combine the arriving line segment L_i^a with the departing line segment L_j^d , if and only if the latter marks the unique first possible departure. In the sequel, we will give PESP constraints ensuring every feasible timetable to be alternating at S . Thus, every feasible timetable will encode some unique matching and the associated line plan.

The first two sets of constraints ensure the minimal headway d in front of and behind the matching station S :

$$\forall i, j \in \{1, \dots, n\} : \pi_{a_j} - \pi_{a_i} \in [d, T - d]_T, \quad (11)$$

$$\forall i, j \in \{1, \dots, n\} : \pi_{d_j} - \pi_{d_i} \in [d, T - d]_T. \quad (12)$$

Notice that (11) and (12) can only be fulfilled if $0 \leq d \leq \frac{T}{n}$. Moreover, we relate arrival events to departure events by the following disjunctive constraints

$$\forall i, j \in \{1, \dots, n\} : \pi_{d_j} - \pi_{a_i} \in [0, T - d + h]_T, \quad (13)$$

$$\forall i, j \in \{1, \dots, n\} : \pi_{d_j} - \pi_{a_i} \in [d, T + h]_T, \quad (14)$$

where we denote by h the maximal stopping time for a train at station S . Together, these constraints (13) and (14) yield

$$(\pi_{d_j} - \pi_{a_i}) \bmod T \in [0, h] \dot{\cup} [d, T - d + h]. \quad (15)$$

Trivially, $0 \leq h < d$ is necessary for every feasible timetable π to be alternating at S .

Theorem 3. *Let π be a timetable respecting constraints (11) to (14). Then for every departure event d_j , there exists a unique arrival event a_i satisfying*

$$\pi_{d_j} - \pi_{a_i} \in [0, h]_T, \quad (16)$$

if and only if $h < (n + 1)d - T$.

Since $0 \leq h$, from $h < (n + 1)d - T$ we conclude $\frac{T}{n+1} < d$.

Proof. “ \Rightarrow ”: We assume $h \geq (n + 1)d - T$. Since $d = \frac{T}{n}$ would imply $h \geq d$, we must only investigate the case that $d < \frac{T}{n}$. We will construct a timetable which respects the constraints (11) to (14), but which contradicts (16).

Define $\pi_{a_i} := (i - 1)d$, for all $i = 1, \dots, n$, and $\pi_{d_j} := j \cdot d$, for all $j = 1, \dots, n$. By construction, all the constraints are satisfied. However, since $\pi_{a_n} + h < n \cdot d = \pi_{d_n}$, for departure π_{d_n} none of the arrival events fulfills (16), q.e.d.

“ \Leftarrow ”: We assume there exists a timetable π having one departure event d_0 such that

$$\forall i = 1, \dots, n : (\pi_{d_0} - \pi_{a_i}) \bmod T > h,$$

but which respects the constraints (11) to (14). We may assume w.l.o.g. that for the cyclic predecessor arrival a_1 of d_0 we have $\pi_{a_1} = 0$. As π is feasible, it satisfies (15). From our assumption, we conclude $d \leq \pi_{d_0}$ and $\pi_{d_0} + (d - h) \leq \pi_{a_2}$, and hence $\pi_{a_2} - \pi_{a_1} \geq 2d - h$. Event a_1 also takes place at time T . For notational convenience, we define $\pi_{a_{n+1}} := T$. With this notation, we have $\pi_{a_{i+1}} - \pi_{a_i} \geq d$, for all $i = 2, \dots, n$. By the definition of $\pi_{a_{n+1}}$, we know that

$$\sum_{i=1}^n (\pi_{a_{i+1}} - \pi_{a_i}) = \pi_{a_{n+1}} - \pi_{a_1} = T.$$

Summing up the lower bounds yields $T \geq (n + 1)d - h$, which contradicts the hypothesis of Theorem 3. \square

Corollary 1. *If $h < (n+1)d - T$, then every timetable which respects constraints (11) to (14) is an alternating timetable.*

In Figure 16, we provide an example for the easiest case, namely matching two lines. As usual, we assume the period time to be 60 minutes.

Remark 2. There are of course alternating periodic timetables in the case $d \leq \frac{T}{n+1}$. PESP solvers are able to detect even those, if we were able to pre-define sufficiently many empty slots. By an “empty slot” we understand an artificial line which we have to schedule in the same way as the original lines, hereby separating the lines before and after the empty slot.

In more detail, let us assume that $\frac{T}{n^*+1} < d \leq \frac{T}{n^*}$ for some $n^* > n$, and that h satisfies the assumptions of Theorem 3 for n^* . We then introduce $n^* - n$ artificial dummy arrival and departure events a_i and d_i , $i = n + 1, \dots, n^*$. To prevent the original line segments from being matched with an artificial event, we require $\pi_{d_i} - \pi_{a_i} \in [0, h]$ for all $i = n + 1, \dots, n^*$.

By construction, the only feasible timetables let the original arrivals and departures alternate. However, perfectly balanced timetables, i.e. $\pi_{a_i} := (i - 1)\frac{T}{n}$, are infeasible under these settings if $n^* < 2n$, since they do not provide $n^* - n$ empty slots.

Recall that so far we have considered only one direction. Hence, there is no mechanism yet to bind the matching of one direction to that of the opposite

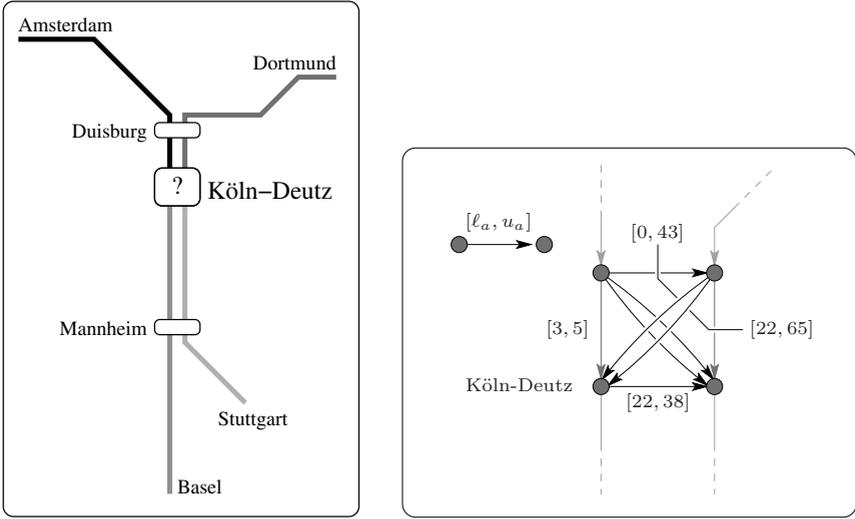


Fig. 16. Modeling aspects of line planning: (a) line segments; (b) PESP constraints ensuring the segments to be matched

direction. But the matchings of opposite directions must fulfill the symmetry assumption that we gave at the beginning of Section 4.2. Otherwise, the trains from direction A could pass the matching station S in order to continue towards B , but the trains from B pass S before continuing in direction C . Thus, it would not be possible to communicate the line plan in the way customers are used to, because it may no more be visualized by an undirected graph. However, limited asymmetries in operation are accepted in practice.

Example 2 (S-Bahn Berlin GmbH). We consider the line S2 serving the route Blankenfelde-Lichtenrade-Buch-Bernau. Between Lichtenrade and Buch, a ten minutes frequency must be offered, for the remaining parts 20 minutes suffice.

In the current timetable ([27]), this line is served in an asymmetric way. In order to cope with the single tracks (which are present at both endpoints) to limit the total amount of stopping time, and to ensure an efficient employment of the rolling stock, an asymmetric service is offered, and we present it in table 4.

In order to ensure symmetric line plans, we have to guarantee the following condition. If we combine the arrival event a_i with the departure event d_j in one direction, then in the opposite direction the complementary arrival event a'_j must be combined with the departure event d'_i . More precisely, when considering the corresponding tension variables $x_{a_i d_j}$ and $x_{a'_j d'_i}$, they must fulfill

$$x_{a_i d_j} \in [0, h] \Leftrightarrow x_{a'_j d'_i} \in [0, h]. \tag{17}$$

In fact, this condition is quite similar to the symmetry constraints (6). What makes things more complicated is the fact that we must not predict in advance

Table 4. Asymmetric service of line S2 (Berlin)

Blankenfelde dep		10:09			arr o	11:14	
Lichtenrade dep ↓		10:15	10:25		arr o	11:05	11:15
Buch arr o		11:06	11:16		dep ↑	10:14	10:24
Bernau arr o		11:21			dep		10:10

for which pairs (i, j) requirement (17) has to hold, and for which pairs it may be violated. Hence, we propose to guarantee property (17) for the matched pairs by imposing symmetry requirements on *every* pair of complementary junctions. But it is clear that this approach cuts off feasible timetables for symmetric line plans just because such timetables need not to be symmetric, see e.g. example 3.

Example 3 (S-Bahn Berlin GmbH). Consider the current timetable ([27]) of the ring subnetwork of S-Bahn Berlin GmbH, of which we provide an excerpt in table 5. Obviously, the line plan is symmetric. But the timetable is not symmetric.

Table 5. Symmetric line plan but asymmetric timetable

Direction A							
Line		S45	S46	S8	S9	S47	S8
Origin		BFHS	BKW	BGA	BFHS	BSPF	BZN
Schöneweide	dep ↓	xx:01	xx:06	xx:10	xx:13	xx:15	xx:18
Baumschulenweg	arr o	xx:03	xx:09	xx:13	xx:16	xx:17	xx:21
Destination		BHMS	BGS	BPKR	BZOO	BWES	BPKR
Direction B							
Line		S8	S46	S9	S47	S8	S45
Origin		BPKR	BGS	BZOO	BWES	BPKR	BHMS
Baumschulenweg	dep ↓	xx:02	xx:06	xx:08	xx:13	xx:14	xx:19
Schöneweide	arr o	xx:05	xx:08	xx:10	xx:15	xx:17	xx:21
Destination		BGA	BKW	BFHS	BSPF	BZN	BFHS

This can be seen by calculating the symmetry axes of lines S47 and S9 at station Schöneweide. Departure and arrival of line S47 sum up to 30, hence the trains of this line meet at times 5 and 15. For line S9 the sum yields 23, providing a symmetry axis of 1.5. An easier argument for asymmetry is that the sequence of the trains in direction B is not the inverse of the one in direction A.

There are two main objectives for the matching approach. First, we want to offer direct trips for as many passengers as possible. Second, the timetable should require only few trains for operation.

For the second criterion, in the case $h = 0$, no additional weight on arcs within the matching node is required in order to minimize the amount of rolling stock required to operate the timetable. In the case $h > 0$, one could put the vehicle weight on the arcs with feasible interval $[0, T - d + h]$. But this would

no longer yield the desired exact piecewise-constant behavior of the objective, because some double counting can appear.

For maximizing the number of direct travelers, we consider the number of passengers w_{ij} starting their trip before the common track on a train covering line segment L_i^a , and finishing their trip after the common endpoint on a train covering line segment L_j^d . The value w_{ij} is added to the weight of the arc $a = (a_i, d_j)$ with $\ell_a = 0$ and $u_a = [0, T - d + h]$. The resulting cost coefficients in the objective function make even sense for pairs of line segments which are not matched, because long changeover times of many passengers are penalized.

Notice that the values w_{ij} are only well-defined if the two line segments do not serve a second matching station. This shows that the decisions to be taken within a matching station are of a rather local nature.

Summarizing, there are important scenarios in which the PESP can integrate relevant aspects of line planning into a model suited for timetabling and key issues of vehicle scheduling. This is in particular the case if symmetric timetables and balanced sequences along the common tracks, i.e. $d > \frac{T}{n+1}$, are requested for their own sake. Moreover, we observed that the larger the distance between two matching stations, the more reliable the passenger weight that we propose.

We think that fast train networks of European agglomerations, such as Frankfurt, Munich, or Paris (RER), are well-suited candidates for this approach. There, many passengers might have their origin or destination somewhere on the backbone route, and balanced sequences must be ensured due to the large number of lines per period.

5.3 Aspects of Network Planning

We propose to also model two questions which arise in network planning within the PESP: the extension of existing tracks, and thus lines, beyond their current endpoints, and the construction of faster tracks as substitution for existing ones. Taking into account that, in these questions, we have to select one option out of a small number of disjoint options, it is evident that we will make intensive use of disjunctive constraints, cf. Section 3.3. Recall that there, we already discussed the introduction of optimal additional stops. With appropriate weights that reflect amortisation—see below—these may also cover the construction of new stations along an existing track.

We only discuss the construction of faster tracks in detail. But the reader will have no difficulty to adapt our suggestions to the very similar task of the extension of tracks.

In Figure 17, we provide a constraint graph which offers the option of a new track between Aachen and Köln, being then part of the European high-speed line PBK (Paris-Brussels-Köln). We provide the status quo, with one intermediate stop, only for illustration purposes. In the future, we have the option to either use the current tracks, thus keeping a trip time of 38 minutes, or to establish the new high-speed track, hereby reducing the trip time down to 26 minutes.

To define appropriate weights for the arcs, we have to take into account three different types of objectives: The number of customers c who profit from a new

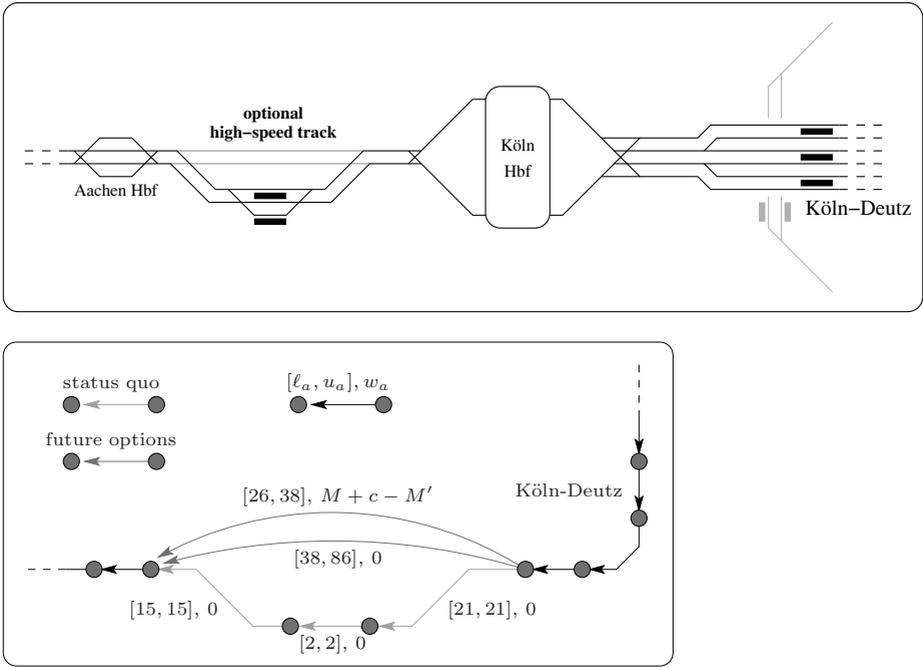


Fig. 17. Modeling aspects of network planning: (a) infrastructure including optional high-speed track; (b) PESP constraints taking into account the two infrastructural alternatives

track by shorter travel times, the trip times of the trains which may allow to reduce the number of trains requires $(M, c.f. \text{Section } 5.1)$, and the cost M' of the investment. One can imagine that it is an absolutely non-trivial management decision to derive an hourly weight M' from the total cost of the investment.

Similarly to line planning, investments into infrastructure will only make sense if they are effected for both directions at the same time. Again, we ensure symmetric investments by requiring the timetable to be symmetric.

Let us now analyze the situation in which several lines have the option of using the same new, faster track. Of course, we want to ensure that infrastructure is only paid once in terms of the objective function. Hence, we have to partition the total cost onto all of the concerned lines. But what if in a solution of a PESP instance only one line is routed over the new track?

But a reasonable allocation of the total costs is only possible, if we know in advance how many lines will have to use the new track. Unfortunately, we are only able to ensure this with constraints of the types already introduced, if *all* the lines must use the same track. This would, e.g., be the case when analyzing two mutually exclusive variants of constructing a new track.

We can guarantee that all the lines use the same track simply by enforcing the same running time for each line. This is achieved by introducing constraints

of type (6). But notice that in this case we cheat a bit, because those constraints no longer relate only pairs of complementary arcs to each other... Anyway, the MIP formulation of this even slightly more extended model incorporates many of the computational aspects of the pure PESP model.

6 Conclusion

Our discussion of the PESP model shows that it has a great modeling power and extendability. We have demonstrated that many non-standard requirements for periodic timetables and also important aspects of other – traditionally separate – planning phases can be integrated into the PESP. Figure 18 displays the gain by this modeling power over the traditional use of the PESP displayed in Figure 1.

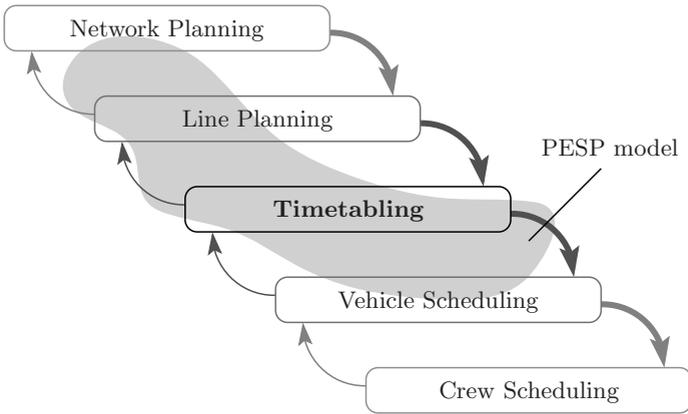


Fig. 18. Planning phases covered by the PESP with our contribution

Interestingly, this integration into the PESP has been possible without seemingly complicating it too much. In all cases, we obtained mixed integer programs that still have the characteristics of a PESP. Hence we believe that these extended models stay computationally tractable also for networks of relevant sizes. So far, our belief is confirmed by a confidential study for S-Bahn Berlin GmbH for two of its three major subnetworks.

We therefore hope that these models, through their integrative approach to vehicle scheduling, timetabling, line planning, and infrastructure planning, will eventually lead to better decision making in practice.

Acknowledgments

We want to thank the staff of Deutsche Bahn AG, S-Bahn Berlin GmbH, and Berliner Verkehrsbetriebe (BVG) for providing us with both real-world data and very detailed requirements of their specific periodic timetabling problems. Moreover, we thank the referees for their *very* detailed suggestions.

References

1. Bollobás, B.: *Modern Graph Theory*. Graduate Texts in Mathematics, vol. 184 (2nd printing) Springer, Heidelberg (2002)
2. Borndörfer, R., Grötschel, M., Pfetsch, M.E.: *A path-based model for line planning in public transport*. Technical Report 05-18, Zuse Institute, Berlin (2005)
3. Borndörfer, R., Löbel, A., Weider, S.: *Integrierte Umlauf- und Dienstplanung im öffentlichen Nahverkehr*. In: HEUREKA '02: Optimierung in Transport und Verkehr, Tagungsbericht, number 002/72. FGSV Verlag (in German) (2002)
4. Bussieck, M.R., Winter, T., Zimmermann, U.: *Discrete optimization in public rail transport*. *Mathematical Programming B* 79, 415–444 (1997)
5. Claessens, M.T., van Dijk, N.M., Zwaneveld, P.J.: *Cost optimal allocation of rail passenger lines*. *European Journal of Operational Research* 110(3), 474–489 (1998)
6. Engelhardt-Funke, O., Kolonko, M.: *Analysing stability and investments in railway networks using advanced evolutionary algorithms*. *International Transactions in Operational Research* 11, 381–394 (2004)
7. Grötschel, M., Löbel, A., Völker, M.: *Optimierung des Fahrzeugumlaufs im öffentlichen Nahverkehr*. In: Hoffmann, K.H., Jäger, W., Lohmann, T., Schunck, H. (eds.) *Mathematik - Schlüsseltechnologie für die Zukunft*, Springer, Heidelberg (in German) (1997)
8. Haase, K., Desaulniers, G., Desrosiers, J.: *Simultaneous vehicle and crew scheduling in urban mass transit systems*. *Transportation Science* 35(3), 286–303 (2001)
9. Krista, M.: *Verfahren zur Fahrplanoptimierung am Beispiel der Synchronzeiten*. Ph.D. thesis, Technische Universität Braunschweig (in German) (1997)
10. Kroon, L.G., Peeters, L.W.P.: *A variable trip time model for cyclic railway timetabling*. *Transportation Science* 37, 198–212 (2003)
11. Leuschel, I.: *Der Fernverkehrsfahrplan 2003 der Deutschen Bahn AG*. *Eisenbahntechnische Rundschau (In German)* 51(7–8), 452–464 (2002)
12. Liebchen, C.: *Finding short integral cycle bases for cyclic timetabling*. In: Di Battista, G., Zwick, U. (eds.) *ESA 2003*. LNCS, vol. 2832, pp. 715–726. Springer, Heidelberg (2003)
13. Liebchen, C.: *Symmetry for periodic railway timetables*. *Electronic Notes in Theoretical Computer Science* 92, 34–51 (2004)
14. Liebchen, C., Peeters, L.: *Some practical aspects of periodic timetabling*. In: Charni, P., Leisten, R., Martin, A., Minnemann, J., Stadtler, H. (eds.) *Operations Research 2001*, Springer, Heidelberg (2002)
15. Liebchen, C., Proksch, M., Wagner, F.H.: *Performance of algorithms for periodic timetable optimization*. In: Hickman, M. (ed.) *Computer-Aided Transit Scheduling— Proceedings of the Ninth International Workshop on Computer-Aided Scheduling of Public Transport*. *Lecture Notes in Economics and Mathematical Systems*, Springer, Heidelberg (to appear, 2005)
16. Lindner, T.: *Train Schedule Optimization in Public Rail Transport*. Ph.D. thesis, Technische Universität Braunschweig (2000)
17. Rhein-Main Verkehrsverbund: *Fahrplanbuch Gesamtausgabe (gültig ab 14. Dezember 2003)* (2003)
18. Nachtigall, K.: *A branch and cut approach for periodic network programming*. *Hildesheimer Informatik-Berichte* 29, Universität Hildesheim (1994)
19. Nachtigall, K.: *Cutting planes for a polyhedron associated with a periodic network*. *Institutsbericht IB 112-96/17*, Deutsche Forschungsanstalt für Luft- und Raumfahrt e.V (July 1996)

20. Nachtigall, K.: Periodic Network Optimization and Fixed Interval Timetables. Habilitation thesis, Universität Hildesheim (1998)
21. Nachtigall, K., Voget, S.: A genetic algorithm approach to periodic railway synchronization. *Computers and Operations Research* 23(5), 453–463 (1996)
22. Nielsen, M.N., Hove, B., Clausen, J.: Constructing periodic timetables using MIP—a case study from DSB S-train. *International Journal of Operations Research*, 1 (2005)
23. Odijk, M.A.: Construction of periodic timetables, Part 1: A cutting plane algorithm. Technical Report 94-61, TU Delft (1994)
24. Odijk, M.A.: A constraint generation algorithm for the construction of periodic railway timetables. *Transportation Research B* 30(6), 455–464 (1996)
25. Peeters, L.W.P.: Personal Communication (2000)
26. Peeters, L.W.P.: Cyclic Railway Timetable Optimization. Ph.D. thesis, Erasmus Universiteit Rotterdam (2003)
27. S-Bahn Berlin GmbH: S-Bahn-Fahrplan (gültig ab 16. Juni 2003) (2003)
28. Schrijver, A.: *Theory of Linear and Integer Programming*, 2nd edn. Wiley, Chichester (1998)
29. Schrijver, A., Steenbeek, A.G.: Dienstregelingontwikkeling voor Nederlandse Spoorwegen N.S. Rapport Fase 1, Centrum voor Wiskunde en Informatica (Oktober 1993)
30. Serafini, P., Ukovich, W.: A mathematical model for periodic scheduling problems. *SIAM Journal on Discrete Mathematics* 2(4), 550–581 (1989)
31. van den Berg, J.H.A., Odijk, M.A.: DONS: Computer aided design of regular service timetables. In: Murthy, T.K.S., Mellitt, B., Brebbia, C.A., Sciotto, G., Sone, S. (eds.) *Computers in Railways IV (COMPRAIL)*—vol. 2: Railway Operations. WIT Press (1994)
32. Völker, M.: Ein multikriterieller Algorithmus zur automatisierten Busliniennetzplanung. Lecture on the OR Workshop Optimierung im öffentlichen Nahverkehr (2003)