




BERENIKE MASING¹, NIELS LINDNER², CHRISTIAN
LIEBCHEN³

Periodic Timetabling with Integrated Track Choice for Railway Construction Sites

¹  0000-0001-7201-2412

²  0000-0002-8337-4387

³  0000-0002-4311-2024

Zuse Institute Berlin
Takustr. 7
14195 Berlin
Germany

Telephone: +49 30 84185-0
Telefax: +49 30 84185-125

E-mail: bibliothek@zib.de
URL: <http://www.zib.de>

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Periodic Timetabling with Integrated Track Choice for Railway Construction Sites

Berenike Masing ^{a,1}, Niels Lindner ^{a,b}, Christian Liebchen ^c

^a Department Network Optimization, Zuse Institute Berlin
Takustr. 7, 14195 Berlin, Germany

¹ E-mail: masing@zib.de, Phone: +49 (0) 30 84185223

^b Institut für Mathematik, Freie Universität Berlin
Arnimallee 3, 14195 Berlin, Germany

^c Technical University of Applied Sciences Wildau
Hochschulring 1, 15745 Wildau, Germany

Abstract

We propose a mixed-integer linear programming model to generate and optimize periodic timetables with integrated track choice in the context of railway construction sites. When a section of a railway network becomes unavailable, the nearby areas are typically operated close to their capacity limits, and hence carefully modeling headways and allowing flexible routings becomes vital. We therefore discuss first how to integrate headway constraints into the Periodic Event Scheduling Problem (PESP) that do not only prevent overtaking, but also guarantee conflict-free timetables in general and particularly inside stations. Secondly, we introduce a turn-sensitive event-activity network, which is able to integrate routing alternatives for turnarounds at stations, e.g., turning at a platform vs. at a pocket track for metro-like systems. We propose several model formulations to include track choice, and finally evaluate them on six real construction site scenarios on the S-Bahn Berlin network.

Keywords

Railway Timetabling, Periodic Timetabling, Periodic Event Scheduling, Train Routing, Turnarounds

1 Introduction

In almost all industrial and developing countries it is a political goal to perform a modal shift from private car use to public transport, among others for environmental reasons. In agglomerations, local railways and metro systems are of particular importance, because of their large capacities. These systems are likely to be operated at a high frequency, e.g., every 5 minutes during peak hours, and often no more than 30 minutes off-peak.

In order to keep such systems in a safe condition, not only the trains, but also the tracks are subject to regular maintenance. Small maintenance action is typically planned during periods with no regular service, e.g., during nights that do not touch any weekend. However, there is also maintenance work that exceeds such a short night break. Then, the regular service is affected and needs to be rescheduled on the infrastructure that remains available.

The relevance of automated timetabling in particular for construction sites becomes even more clear, when considering the planning process. Throughout an entire year, compared to

the annual regular timetable, the planning staff might spend more time on planning timetables around all construction sites – but these timetables typically have relatively short validity periods (e.g., one weekend vs. 364 days for the annual timetable).

In the present paper we consider the problem of designing a periodic schedule around railway construction sites. Our approach is based on the standard graph model for periodic timetabling, the Periodic Event Scheduling Problem (PESP), introduced by Serafini and Ukovich (1989). Compared to the classical periodic timetabling problem (see e.g., Liebchen, 2008; Kroon et al., 2009), there are certain key differences: At first, the actual planning area (close to the construction site) is typically only a small part of the network, and the timetable to be designed must fit to the regular (annual) timetable for the rest of the network. Secondly, the volume of operations near the blocked construction site is likely close to the tracks’ capacity, in particular regarding station capacity for turnarounds. Thirdly, timetable regularity (even spaced headways) becomes even more important.

For example, when a standard metro station with one platform and two tracks is the last one that can be served, there can already be multiple paths for the trains to change their direction, see Figure 7. Moreover, if this station’s capacity is then highly occupied, then the routing of the trains through this temporal terminus station becomes key.

In the area of railway optimization, the Dutch railways serve as a kind of reference (Kroon et al., 2009). But even in their line of development, for complexity reasons it had been a very early fundamental decision *not* to consider timetabling on the tracks between the stations and the routing of the trains within a station in an integrated way (Kroon et al., 1997).

In more recent years, there had been proposed a way to model the integration of these two planning steps. In Wüst et al. (2018), extra variables by which the optimization model may turn on (off) PESP constraints that connect events which are related in (not) selected routings, are added. Yet, these are not fully able to meet the practical requirements of periodic timetabling for railway construction sites, in particular within stations and long dwell times of the trains there, e.g., in the context of turnarounds. In Section 3, we propose new ways to integrate track choice into periodic timetabling.

Enforcing conflict-free periodic timetables has received much attention in the literature. PESP-based methods can prevent overtaking of trains by ensuring minimum headway times. Yet, during long dwell times at platforms or sidings another type of conflict can arise, namely *overlapping track occupation*. To the best of our knowledge, this has never been examined, and we describe four methods to overcome this issue.

Section 2 introduces our problem, recapitulates PESP, extensively discusses conflict-freeness, and the aspect of regularity. We present a *turn-sensitive* event-activity network, which allows to flexibly model routing choices and to formulate our full mixed-integer programming model in Section 3. Section 4 compares the performance of our model and its configurations, including several alternative headway formulations, on real-world railway construction sites for the S-Bahn Berlin network. We conclude the paper in Section 5.

2 Periodic Timetabling Model

2.1 General Setting

We briefly describe the basic setting for our timetabling problem and refer to Sections 2 and 3 for details. We are given a graph \mathcal{G} modeling the stations of a railway network and the links between the stations on a macroscopic level. On \mathcal{G} , we are given a line plan, i.e., a set

\mathcal{L} of simple paths together with frequencies $f : \mathcal{L} \rightarrow \mathbb{N}$. Moreover, there is an infrastructure graph \mathcal{I} refining \mathcal{G} and modeling the railway network on at least a mesoscopic level, e.g., with platform tracks as vertices. The task is now to find an optimal conflict-free mesoscopic periodic timetable for the macroscopic lines in \mathcal{L} respecting the frequency assignment f . In particular, it is part of the timetabling task to determine a mesoscopic routing in \mathcal{I} for each line in \mathcal{L} . The objective is measured in terms of operational and user costs.

As our approach is targeted towards railway construction sites, we introduce further prerequisites and constraints: We denote by \mathcal{ZI} a subgraph of \mathcal{I} that signifies currently unavailable infrastructure, and let \mathcal{ZG} be the corresponding subgraph of \mathcal{G} . We assume that all lines in \mathcal{L} are contained in $\mathcal{G} \setminus \mathcal{ZG}$. Moreover, we suppose that an original timetable is given on \mathcal{I} for an original set of lines in \mathcal{G} . When only a small part of infrastructure is affected, it is not reasonable to re-optimize the timetable of the full network. We therefore consider a *planning subgraph* $\mathcal{PI} \subseteq \mathcal{I} \setminus \mathcal{ZI}$ and a corresponding planning subgraph $\mathcal{PG} \subseteq \mathcal{G} \setminus \mathcal{ZG}$, and demand that line routing and timetable are fixed to the original line routing and timetable outside \mathcal{PG} resp. \mathcal{PI} . In particular, for the purpose of timetabling, we can discard lines that do not touch the planning subgraph. Finally, we want to compute periodic timetables for periodic vehicle circulations (Borndörfer et al., 2018; Van Lieshout, 2021), i.e., we want each line $l \in \mathcal{L}$ in $\mathcal{G} \setminus \mathcal{ZG}$ to be routed as some closed walk in $\mathcal{I} \setminus \mathcal{ZI}$.

2.2 Periodic Event Scheduling

Periodic timetabling problems are usually modelled with the *Periodic Event Scheduling Problem* (PESP), originally introduced by Serafini and Ukovich (1989). For an event-activity network (simple digraph) \mathcal{N} , a period time $T \in \mathbb{N}$, lower and upper bounds $\ell_{ij}, u_{ij} \in \mathbb{N}$, and weights $w_{ij} \in \mathbb{N}$ for the activities $(i, j) \in \mathcal{A}(\mathcal{N})$, PESP can be formulated as follows:

$$\text{Minimize} \quad \sum_{(i,j) \in \mathcal{A}(\mathcal{N})} w_{ij} y_{ij} \quad (1)$$

$$\text{s.t.} \quad \pi_j - \pi_i + T p_{ij} = y_{ij} + \ell_{ij} \quad (i, j) \in \mathcal{A}(\mathcal{N}) \quad (2)$$

$$0 \leq y_{ij} \leq u_{ij} - \ell_{ij} \quad (i, j) \in \mathcal{A}(\mathcal{N}) \quad (3)$$

$$0 \leq \pi_i \leq T - 1 \quad i \in \mathcal{V}(\mathcal{N}) \quad (4)$$

$$p_{ij} \in \mathbb{Z} \quad (i, j) \in \mathcal{A}(\mathcal{N}) \quad (5)$$

The goal is to assign a timestamp π_i to each event $i \in \mathcal{V}(\mathcal{N})$ within the period time T such that for each pair $(i, j) \in \mathcal{A}(\mathcal{N})$ of activities, there is a *periodic tension* value x_{ij} within the activity bounds $\ell_{ij} \leq x_{ij} \leq u_{ij}$ satisfying $x_{ij} \equiv \pi_j - \pi_i \pmod{T}$. These tension values essentially model activity durations between the events' points in time, which are in turn given by the *periodic timetable* π . In the context of railway timetabling, events are arrivals and departures, and typical activity types comprise, e.g., driving between two stations, waiting at a platform, or turnarounds. With the help of periodic offset variables $p_{ij} \in \mathbb{Z}$ for $(i, j) \in \mathcal{A}(\mathcal{N})$, the modulo constraints can be linearized as $x_{ij} = \pi_j - \pi_i + T p_{ij}$. Considering the periodic slack $y_{ij} = x_{ij} - \ell_{ij}$ instead of x_{ij} for all arcs $(i, j) \in \mathcal{A}(\mathcal{N})$, we arrive at the above MIP model (1)-(5).


If $0 \leq u_{ij} - \ell_{ij} \leq T$ – what we always assume w.l.o.g. – then (5) can be replaced by $p_{ij} \in \{0, 1, 2\}$. If even $u_{ij} < T$ holds, p_{ij} can be further restricted to $\{0, 1\}$. We refer to (Liebchen, 2006; Liebchen and Möhring, 2007; Nachtigall, 1998; Odijk, 1994; Lindner and

Reisch, 2022) for an intensive treatment of MIP formulations, modeling capabilities, and the hardness of PESP.

2.3 Modeling Conflict-Free Timetables

To model additional operational requirements, in particular that multiple trains are not scheduled on the same tracks at the same time, one can add *headway activities*, e.g., to ensure that a minimal distance between two successive trains is kept. These can be added into \mathcal{N} as additional arcs with lower and upper bounds and treated as any normal activity. However, standard headway activities have a weakness, as they separate the occurrences of two events only, not of their activities. This can lead to unwanted results, such as illegal overtaking, which can occur, e.g., when two trains with different speeds use the same track:

Example 2.1 (Illegal overtaking). Consider Figure 1a for which we assume a period time T of, say, $T = 10$, and therein the two black activities (i_1, j_1) and (i_2, j_2) with bounds $[2, 8]$ each. Suppose that these activities use the same infrastructure, e.g., driving of trains on the same physical track. If the minimum time distance between i_1 and i_2 resp. j_1 and j_2 is supposed to be at least three, then this can be modeled with the two orange headway activities with bounds $[3, T - 3]$ each. For any feasible timetable, this guarantees $|\pi_{i_1} - \pi_{i_2}| \geq 3$ and $|\pi_{j_1} - \pi_{j_2}| \geq 3$. However, the timetable π indicated in blue is feasible, but the two trains overtake each other on the same track. Note that $p_{i_1 j_1} + p_{j_1 j_2} = 0 + 1 = 1$ and $p_{i_1 i_2} + p_{i_2 j_2} = 0 + 0 = 0$.




(a) Illegal overtaking despite respecting minimum headway times. (b) Preventing illegal overtaking by subdivision: The blue timetable from Figure 1a is no longer feasible, as the tension between v_1 and v_2 can impossibly lie within $[3, T - 3]$.

Figure 1: Overtaking problem with solution approach

Preventing overtaking trains has received much attention in the literature (Lindner, 2000; Peeters, 2003; Schrijver and Steenbeek, 1993): One approach is to require an additional constraint on the periodic offsets in a situation as in Example 2.1, namely

$$p_{i_1 j_1} + p_{j_1 j_2} = p_{i_1 i_2} + p_{i_2 j_2}. \quad (6)$$

Another method is to subdivide the arcs (i_1, j_1) and (i_2, j_2) as in Figure 1b until the span becomes smaller than twice the minimum headway time (Liebchen and Möhring, 2007). Observe that the former approach is formally leaving the graph model of PESP, while the latter is staying immediately within PESP.



(a) Illegal occupation of tracks despite respecting minimum headway times and no overtaking. (b) The subdivision approach for overtaking: The blue timetable from Figure 2a remains feasible, no further subdivision will resolve this issue.

Figure 2: Track Occupation Problem

However, none of the solutions to the overtaking problem could finally address all issues related to periodic railway scheduling, namely the *track occupation problem*.

Example 2.2 (Track occupation problem). Again, assume a period time $T = 10$ and consider the feasible timetable displayed in blue in Figure 2a. If the black arcs $(i_1, j_1), (i_2, j_2)$ are driving activities of trains t_1 and t_2 on the same infrastructure, this might be a reasonable schedule. In this case, it could be possible that t_1 is reasonably far away from the station for t_2 to enter the tracks at time 3. In contrast, if the arcs are waiting or turnaround activities on the same infrastructure, this schedule poses a problem: Train t_2 is scheduled to enter the tracks while t_1 still occupies the tracks. Moreover, no amount of subdividing the two problematic activities will resolve this issue as becomes evident from Figure 2b. The constraint (6) is satisfied, too, as $p_{i_1 j_1} + p_{j_1 j_2} = 0 + 1 = 1$ and $p_{i_1 i_2} + p_{i_2 j_2} = 0 + 1 = 1$.

To the best of our knowledge, this issue has not yet been addressed in the literature. In particular, recall from Kroon et al. (1997, 2009) that even in one of the most prominent applications of mathematical optimization to railway planning a very clear separation is practiced: At Dutch railways, since the very early beginning, on the one hand, the network-wide timetable on the tracks is computed by the software module CADANS, while on the other hand, conflict-free routings of the trains within a station are computed by STATIONS.

The track occupation problem mostly occurs when the time-interval of an activity blocking a part of the infrastructure is fairly large. Turnarounds take quite a bit of time, as the train sometimes needs to be emptied and the driver usually needs to move from the front all the way to the back. There can be further safety constraints, as the train might then operate on the left-hand side of the tracks with fewer signals, etc. In conjunction with the limited station capacities in construction site scenarios, turnarounds become hence a very difficult aspect to schedule.

In practice, it is not enough to ensure that no two trains occupy the tracks at the same time, but there are usually further safety or technical requirements: Each track needs to be unoccupied for a certain amount $\varepsilon > 0$ of time, before the next train may arrive. Moreover, it is often desirable to separate two subsequent arrivals by at least $h > 0$ units of time.

To address the track occupation problem including the safety requirements as given by the parameters h and ε , we will first define occupation intervals, which represent the physical occupation of the tracks by trains including safety buffers.

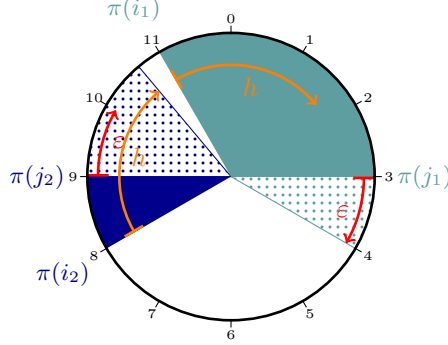


Figure 3: (ε, h) -occupation intervals of two trains for $T = 12$: The solid areas show the physical blockage of a track by two trains, the dotted indicate the time span that the tracks need to stay unoccupied by the safety parameters ε and h .

Definition 2.3 (Periodic Intervals). For $a, b \in \mathbb{R}$ with $a < b$ we define the periodic interval

$$[a, b]_T := \{x \bmod T \mid x \in [a, b]\} \subseteq [0, T[,$$

i.e., the image of the half-open interval $[a, b[$ w.r.t. the modulo T map with values in $[0, T[$.

For example, $[2, 8]_{10} = [2, 8[, [8, 12]_{10} = [0, 2[\cup [8, 10[$, and $[2, 18]_{10} = [0, 10[$.

Definition 2.4 ((ε, h) -Occupation Intervals). Let π be a periodic timetable w.r.t. an event-activity network \mathcal{N} , and let x be a corresponding periodic tension. For $h, \varepsilon > 0$, the (ε, h) -occupation interval of an activity $(i, j) \in \mathcal{A}(\mathcal{N})$ is the periodic interval

$$[\pi_i, \max(\pi_i + h, \pi_i + x_{ij} + \varepsilon)]_T.$$

Let $\mathcal{H} \subseteq \mathcal{A}(\mathcal{N}) \times \mathcal{A}(\mathcal{N})$ denote the set of pairs of arcs in \mathcal{N} whose activities make use of a common piece of infrastructure. We will assume that if $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$, then $((i_2, j_2), (i_1, j_1)) \notin \mathcal{H}$.

Definition 2.5 ((ε, h) -Conflict-Free Timetable). A periodic timetable of an event-activity network \mathcal{N} is (ε, h) -conflict-free if the (ε, h) -occupation intervals of (i_1, j_1) and (i_2, j_2) are disjoint for any pair $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$.

For a better understanding of (ε, h) -occupation intervals and conflict-freeness, consider Figure 3, where the occupation of infrastructure by two trains is displayed on an abstract clock for a pair $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$. For $t \in \{1, 2\}$, $\pi(i_t)$ marks the arrival of train t and $\pi(j_t)$ its departure. The solid areas show how long the track is blocked by each train physically, and the dotted areas show how much time is blocked by the safety parameters h and ε . The entire shaded regions show the two (ε, h) -occupation intervals of the two trains, respectively. If these two areas do not intersect, those trains are scheduled conflict-free.

Remark 2.6. To be (ε, h) -conflict-free, it is necessary for a periodic timetable π with tension x that both $x_{i_1 j_1} < T - \varepsilon$ and $x_{i_2 j_2} < T - \varepsilon$ hold for all pairs $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$. Moreover, h needs to be larger than 0 and smaller than T .

In the sequel, we present four approaches to model (ε, h) -conflict-free timetables. Due to Remark 2.6, we assume w.l.o.g. $u_{i_1 j_1}, u_{i_2 j_2} < T - \varepsilon$ for all $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$, and furthermore $0 < h < T$. We will need a technical lemma, whose proof will be omitted:

Lemma 2.7. *Let $a, b, c, d \in \mathbb{R}$. Then $[a, b[_T \cap [c, d[_T \neq \emptyset$ if and only if $(c - a) \bmod T \in [0, b - a[_T$ or $(a - c) \bmod T \in [0, d - c[_T$.*

Q3 – Triangular Headway Formulation

Our first approach introduces four activities and two constraints on the periodic offset variables per pair $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$.




Figure 4: Q3 headway activities with lower and upper bounds

Proposition 2.8. *For each pair $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ add four activities (i_1, i_2) , (i_2, i_1) , (j_1, i_2) , (j_2, i_1) with bounds as in Figure 4. Then a periodic timetable is (ε, h) -conflict-free if and only if for the corresponding periodic offset vector p holds*

$$p_{i_1 i_2} = p_{i_1 j_1} + p_{j_1 i_2} \quad \text{and} \quad p_{i_2 i_1} = p_{i_2 j_2} + p_{j_2 i_1} \quad (7)$$

for all $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$.

Proof. Suppose that π is not (ε, h) -conflict-free. Then we find $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ with

$$[\pi_{i_1}, \max(\pi_{i_1} + h, \pi_{i_1} + x_{i_1 j_1} + \varepsilon)]_T \cap [\pi_{i_2}, \max(\pi_{i_2} + h, \pi_{i_2} + x_{i_2 j_2} + \varepsilon)]_T \neq \emptyset,$$

and this implies by Lemma 2.7 that $(\pi_{i_2} - \pi_{i_1}) \bmod T \in [0, \max(h, x_{i_1 j_1} + \varepsilon)]_T$ or $(\pi_{i_1} - \pi_{i_2}) \bmod T \in [0, \max(h, x_{i_2 j_2} + \varepsilon)]_T$. By Remark 2.6, we have $h < T$, $x_{i_1 j_1} < T - \varepsilon$, and $x_{i_2 j_2} < T - \varepsilon$. Thus, the above periodic intervals are identical with the standard half-open intervals.

Further, as $x_{i_1 i_2}, x_{i_2 i_1} \in [h, T - h] \subseteq [0, T[$, we have $h \leq x_{i_1 i_2} = (\pi_{i_2} - \pi_{i_1}) \bmod T$ and $h \leq x_{i_2 i_1} = (\pi_{i_1} - \pi_{i_2}) \bmod T$. We infer $x_{i_1 i_2} < x_{i_1 j_1} + \varepsilon$ or $x_{i_2 i_1} < x_{i_2 j_2} + \varepsilon$, which implies by the lower bounds on $x_{j_1 i_2}$ and $x_{j_2 i_1}$ that $x_{i_1 i_2} < x_{i_1 j_1} + x_{j_1 i_2}$ or $x_{i_2 i_1} < x_{i_2 j_2} + x_{j_2 i_1}$. Inserting the constraint (3) and observing that the π -variables cancel out, we obtain $p_{i_1 i_2} < p_{i_1 j_1} + p_{j_1 i_2}$ or $p_{i_2 i_1} < p_{i_2 j_2} + p_{j_2 i_1}$, which violates (7).

Assume now that π is (ε, h) -conflict-free. Using Lemma 2.7, we conclude that $(\pi_{i_2} - \pi_{i_1}) \bmod T \notin [0, \max(h, x_{i_1 j_1} + \varepsilon)]_T$ and $(\pi_{i_1} - \pi_{i_2}) \bmod T \notin [0, \max(h, x_{i_2 j_2} + \varepsilon)]_T$ for all $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$. If $\pi_{i_2} > \pi_{i_1}$, then $p_{i_1 i_2} = 0$, $p_{i_2 i_1} = 1$, and

$$p_{i_1 j_1} + p_{j_1 i_2} = \frac{x_{i_1 j_1} + x_{j_1 i_2} + \pi_{i_1} - \pi_{i_2}}{T} \leq \frac{x_{j_1 i_2} - \varepsilon}{T} \leq \frac{T - 2\varepsilon}{T} < 1,$$

$$p_{i_2 j_2} + p_{j_2 i_1} = \frac{x_{i_2 j_2} + x_{j_2 i_1} + \pi_{i_2} - \pi_{i_1}}{T} > \frac{x_{j_2 i_1}}{T} \geq \frac{\varepsilon}{T} > 0,$$

and we conclude that (7) is satisfied due to the integrality of the periodic offset p . The case $\pi_{i_1} > \pi_{i_2}$ is analogous, and the case $\pi_{i_1} = \pi_{i_2}$ is impossible, because of $\varepsilon > 0$. \square

Q4 – Butterfly¹ Headway Formulation

A drawback of the Q3 approach is the large amount of additional constraints: For each pair in \mathcal{H} two additional constraints on the integer variables p are needed. As an equivalent alternative, one can also combine them so that only one constraint per pair is needed. The following is a consequence of Proposition 2.8 and Odijk’s cycle inequalities:

Lemma 2.9. *For each pair $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ add four activities (i_1, i_2) , (i_2, i_1) , (j_1, i_2) , (j_2, i_1) with bounds as in Figure 4. Then a periodic timetable is (ε, h) -conflict-free if and only if for the corresponding periodic offset p and all $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ holds $p_{i_1 j_1} + p_{j_1 i_2} + p_{i_2 j_2} + p_{j_2 i_1} = 1$.*

QT – Traditional Headway Constraints

As shown before, Liebchen and Möhring (2007) solve the overtaking problem by refining the graph without any additional constraints. In this section we show an approach similar in spirit, in which we subdivide each of the problematic activities and add additional headway arcs. For such a refinement approach, we first need to answer the question of how high of a resolution is needed, however. To this end, we first consider a special case with particularly small upper bounds in which the condition on the p -variables is automatically fulfilled, as can be seen by an application of Odijk’s cycle inequalities:

Lemma 2.10. *For each pair $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ add four activities (i_1, i_2) , (i_2, i_1) , (j_1, i_2) , (j_2, i_1) with bounds as in Figure 4. If $u_{i_1 j_1} < h + \varepsilon$ and $u_{i_2 j_2} < h + \varepsilon$ hold for all $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$, then any periodic timetable is (ε, h) -conflict-free.*

Lemma 2.10 can be used to prove the validity of the following refinement technique:

Lemma 2.11. *Consider a refinement of a pair of activities $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ in the following way: Subdivide both activities into k sub-activities (v_r^{l-1}, v_r^l) with $v_r^0 = i_r$ and $v_r^k = j_r$ with integral bounds such that*

$$\begin{aligned} 0 \leq \ell_{v_r^{l-1}, v_r^l} &\leq u_{v_r^{l-1}, v_r^l} < \min(h + \varepsilon, 2\varepsilon), & l \in \{1, \dots, k\}, r \in \{1, 2\}, \\ \sum_{l=1}^k \ell_{v_r^{l-1}, v_r^l} &= \ell_{i_r, j_r}, & r \in \{1, 2\}, \\ \sum_{l=1}^k u_{v_r^{l-1}, v_r^l} &= u_{i_r, j_r}, & r \in \{1, 2\}, \end{aligned}$$

cf. Figure 5. Then any feasible timetable on the refined instance restricts to a (ε, h) -conflict-free timetable on the original instance.

Q0 – Hyperedge Headway Formulation

It is also possible to model the track occupation problem directly, introducing a different type of constraint which takes into account data from three vertices and two arcs (therefore

¹We derive this name from the four used arcs in the new constraint forming a butterfly shape in Figure 4.




Figure 5: A refinement of arc (i_1, j_1) to prevent overlapping track occupation, assuming for simplicity that ℓ and u are both divisible by k to maintain the integrality of the bounds.

“hyperedge”). A huge drawback of all the headway formulations above is the large amount of additional integer variables: For the Q3 and Q4 formulations, we add four activities and hence four integral variables to the model for each pair in \mathcal{H} . This issue is exacerbated in the QT-formulation due to the many headway arcs arising from the refinement. In contrast, we propose to introduce only two new activities with corresponding binary offset variables.

Lemma 2.12. *For each pair $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ add two activities (i_1, i_2) , (i_2, i_1) with bounds $[h, T - h]$ each. Then a periodic timetable π with corresponding tension x and periodic offset p is (ε, h) -conflict-free if and only if for all $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ holds*

$$\pi_{i_2} + Tp_{i_1 i_2} \geq \pi_{i_1} + x_{i_1 j_1} + \varepsilon \quad \text{and} \quad \pi_{i_1} + Tp_{i_2 i_1} \geq \pi_{i_2} + x_{i_2 j_2} + \varepsilon. \quad (8)$$

2.4 Objective and Regularity Recommendation

The main goal of construction scheduling is to provide a timetable which is viable from an operational perspective. However, once feasibility is guaranteed, the goal is to provide a good timetable. The operator’s objective are generally low operational costs, while the passengers’ objective is frequent – preferably regular – and in particular fast service.

We model the operational costs as the sum of total cycle times per line, because these essentially correspond to the number of trains that are required for operating that line, see e.g., Liebchen (2008). For the user costs, we suggest to consider only activities involving passengers, such as driving and waiting within non-terminal stations. If information on travel chains or transfers is available, these may also be represented in the user costs. However, such data could not be available, hence we suggest an alternative approach: regularity.

If, e.g., within a 20-min interval service is provided with frequency two, then an arrival of trains every ten minutes is preferred over two trains arriving in short succession and then a long waiting time before the next train departs again at the beginning of the next period: The expected waiting time for passengers arriving uniformly at random is then shorter.

Also from the operator’s point of view, regularity is of value, as this means that the trains are scheduled time-wise as far apart as possible, meaning that short delays can be absorbed better and result in less knock-on delays for succeeding trains. While it is possible to enforce regular timetables with the help of headway arcs a with fixed bounds $\ell_a = u_a$, this often risks to lead to infeasibility or undesired artificial waiting times. Soft constraints relaxing these headway constraints have been introduced by Nachtigall (1996), but we propose to include regularity in a more flexible way, hereby formally leaving the PESP graph model.

To that end, for all pairs $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ of activities on the same infrastructure, we introduce two new continuous variables $\xi_{i_1 i_2}$ and $\xi_{i_2 i_1}$. These *regularity variables*




Figure 6: Regularity values for two timetables of a line with frequency 2 and $T = 20$.

measure every time unit by which the scheduled headway from i_k to i_l exceeds the ideally balanced headway $r_{i_l i_k}$ for $k, l \in \{1, 2\}$ with $k \neq l$. Formally, we bound $\xi_{i_k i_l}$ by

$$\xi_{i_k i_l} \geq \pi_{i_l} - \pi_{i_k} + T p_{i_k i_l} - r_{i_l i_k}, \quad \text{and} \quad \xi_{i_k i_l} \geq 0. \quad (9)$$

The value for $r_{i_l i_k}$ should be set to $r_{i_l i_k} = T/\bar{f}$ for \bar{f} being the aggregated frequency of lines using the infrastructure of (i_1, j_1) (and thus (i_2, j_2)). Note that in all four headway formulations, the arc (i_k, i_l) and hence $p_{i_k i_l}$ exist. Since $\pi_{i_l} - \pi_{i_k} + T p_{i_k i_l} \equiv (\pi_{i_l} - \pi_{i_k}) \bmod T$, the constraints model the positive deviation from the optimal timewise distance $r_{i_l i_k}$.

Example 2.13. To illustrate the behaviour of ξ , we refer to Figure 6. Both subfigures show exemplary timetables for frequency 2, such that the optimal time between two arrival events would be $r_{i_1 i_2} = 10$ with period $T = 20$. A schedule with optimal regularity can be seen on the right-hand side, where i_2 is scheduled at time 10, while i_1 is scheduled at 0 with aggregated regularity values $\xi_{i_1 i_2} + \xi_{i_2 i_1} = 0$. On the left however, we have a deviation by 5 from the optimum, the regularity value is $\xi_{i_1 i_2} + \xi_{i_2 i_1} = 5 > 0$. With respect to the aggregated regularity, the right schedule is therefore preferred.

Remark 2.14. Note that $(\pi_{i_2} - \pi_{i_1} + T p_{i_1 i_2}) + (\pi_{i_1} - \pi_{i_2} + T p_{i_2 i_1}) = T$ implies that the non-negativity condition (9) on the regularity variable is essential. In fact, we have $\xi_{i_1 i_2} + \xi_{i_2 i_1} \geq T - 2r_{i_1 i_2}$, and equality holds if and only if the schedule is entirely regular.

We can measure the overall regularity simply by aggregation, i.e., by including the term $\sum_{((i_1, j_1), (i_2, j_2)) \in \mathcal{H}} \xi_{i_1 i_2} + \xi_{i_2 i_1}$ in our objective.

3 Integrating Track Choice

3.1 Event-Activity Network

At the heart of the scheduling problem is the *event-activity network* \mathcal{N} . Its vertices correspond to departure or arrival events of trains, while the arcs model the activities, such as driving, waiting, turning events and is used to model time and location of the vehicles simultaneously. We extend the standard notion to encode vehicle directions when occupying infrastructure to model the turnarounds realistically. Within PESP, safety and operational requirements such as a headway between vehicles are also modelled with headway arcs as part of the set of activities. We will treat headways separately however, and do not consider them as part of \mathcal{N} . Moreover, we will ignore transfer activities and any other activities that connect events belonging to different lines for the ease of exposition. As a byproduct, \mathcal{N} decomposes into several connected components, one component $\mathcal{N}(l)$ for each line $l \in \mathcal{L}$.

We demand that $\mathcal{N}(l)$ should reflect any route that a train can feasibly drive. As our lines are to be routed as closed walks, this includes routes for turnarounds in stations. However, we assume that when not in the process of turnarounds, trains use only right-hand tracks, unless declared otherwise on small sections of the network.

In the following, we assume that the infrastructure graph \mathcal{I} contains *platform vertices*, which single out tracks at platform in stations, and *pocket track vertices*, which model tracks that are not located at passengers' platforms, but are suitable for changing the direction of a vehicle. The infrastructure network is allowed to be much finer, but we restrict to those two classes of vertices to extract the essential information for the network \mathcal{N} .

Definition 3.1 (Turn-Sensitive Event-Activity Network). *The turn-sensitive event-activity network $\mathcal{N}(l)$ of a single line $l \in \mathcal{L}$ can be constructed as a digraph with the following arcs:*

- *For every platform p and direction of travel $\omega \in \{\triangleleft, \triangleright\}$, add a waiting activity $((arr, p, \omega), (dep, p, \omega))$ if there is a routing of the line l in the infrastructure graph such that platform p is used in direction ω without turning.*
- *For two successive stations with platforms p and p' in line l in direction ω , add a driving activity $((dep, p, \omega), (arr, p', \omega))$ if there is a routing of line l such that the train can depart platform p and arrive at p' in direction ω without travelling over other platforms or pocket tracks.*
- *For every pocket track s at a turning station of line l add a turnaround activity $((arr, p, \triangleright), (dep, p, \triangleleft))$ if there is a routing of l via s .*
- *For every platform p at a turning station and $\omega \in \{\triangleleft, \triangleright\}$, add a turnaround activity $((arr, p, \omega), (dep, p, -\omega))$ if there exists a routing of l such that p can be used to change direction from ω to $-\omega$.*
- *For every platform - pocket track pair (p, s) at a turning station, add a driving activity $((dep, p, \omega), (arr, s, \triangleright))$ if there is a routing of the line l going over platform p in direction ω directly to the pocket track s .*
- *For every platform - pocket track pair (p, s) at a turning station, add a driving activity $((dep, s, \triangleleft), (arr, p, \omega))$ if there is a routing of the line l going from pocket track s directly to platform p in direction ω .*

The entire \mathcal{N} is the union of the disjoint event-activity networks w.r.t. the line plan, i.e., $f(l)$ distinct copies for each event-activity network $\mathcal{N}(l)^{(k)}$, $k \in \{1, \dots, f(l)\}$, $l \in \mathcal{L}$.

Example 3.2. To illustrate the construction of \mathcal{N} , consider the section of an exemplary track layout displayed in Figure 7. Suppose line $l \in \mathcal{L}$ goes from Station 1, turns at Station 2 and goes back to Station 1. We assume that l does not turn in Station 1, thus, it would have to get from platform P1 to P2. To do so, it can use either platforms P3, P4 or the pocket track S5 to turn, which results in the six routings displayed in Figure 9.

Figure 8 then shows the resulting \mathcal{N} . S5 is only used for turning, thus there is one turning activity connecting an arrival and departure event. On platform P3, there are three options: A train either turns, or waits before continuing in or against standard direction of travel (lowest two blue arcs), similarly for P4. Each path from the leftmost white node to the nearby yellow node corresponding to arrival at P1 and departure at P2, respectively, corresponds to one of the turning maneuvers described above. It is important to distinguish




Figure 7: Track layout




Figure 8: Event-activity network for a line operating on the track layout of Figure 7 and turning at Station 2. Departure events are marked in yellow, and arrivals in white. Driving activities are black, waiting activities are blue, and turnarounds are red. The triangles inside the nodes indicate ω , the direction of travel of the train.




Figure 9: Possible vehicle routes for a turn in Station 2

the different directions of the events, as there might be different activity bounds. More importantly, this is needed for completeness, as otherwise, not all turns could be expressed by paths.

Remark 3.3. As our focus is on construction scenarios, we include only platform and pocket track vertices of infrastructure within the planning subgraph \mathcal{PI} in $\mathcal{N}(l)$. Outside of the planning area changes are not allowed, such that for any line leaving \mathcal{PI} , all events are fixed to the original timetable and can be aggregated into a single node.

3.2 Routing Alternatives

The standard PESP model assumes that the routing of lines is predetermined. However, as is evident from the small example in Figure 8, the exact track choice of a line is not trivial to decide, and the routing choice of one line can make the routing of another impossible. Particularly in construction scenarios, the available space needs to be administered and assigned efficiently to offer as much service as possible. It is thus essential to include the routing of the lines as part of the optimization process.

A feasible routing of a line l corresponds to a simple cycle in $\mathcal{N}(l)$. An enumeration of all simple cycles however, does not seem efficient. We therefore propose a partial enumer-

ation only. The main idea is to identify activities or events that are part of every routing of l and then enumerate all paths between two inevitable events.

The set of *inevitable arcs* $\mathcal{A}_{\text{inev}}$, i.e., arcs that are part of every routing choice, is a feedback arc set. To determine this feedback arc set, we proceed as follows: For each arc $a \in \mathcal{N}(l)$, we consider the subgraph $\mathcal{N}(l)^a := \mathcal{N}(l) \setminus \{a\}$ and check if there still is a cycle in $\mathcal{N}(l)^a$ that contains a driving activity for each pair of subsequent stations in l . This can be done by solving a flow-type problem on $\mathcal{N}(l)^a$.

Let $S(l)$ denote the subgraph of $\mathcal{N}(l)$ that arises from removing all arcs in $\mathcal{A}_{\text{inev}}(l)$. Each connected component of $S(l)$ is acyclic, has a unique source s with in-degree 0, and a unique sink t with out-degree 0, so that we can label each component by $C(l, s, t)$.

Definition 3.4 (Sheaf of Alternatives). *For a line l and a component $C(l, s, t)$ of $S(l)$, we call the set S of all s - t -paths of $C(l, s, t)$ a sheaf of alternatives for l .*

We define $S_{\text{inev}} := \{\bigcup_{l \in \mathcal{L}} \mathcal{A}_{\text{inev}}(l)\}$ and denote by \mathcal{S} the set comprised of S_{inev} and all sheaves of alternatives for all lines. Each element of \mathcal{S} is a set of subsets of arcs of \mathcal{N} , and each arc in \mathcal{N} belongs to exactly one $S \in \mathcal{S}$, although it might be part of several sets in \mathcal{S} .

Example 3.5. Returning to the excerpt of a line network $\mathcal{N}(l)$ depicted in Figure 8, one can envision the concept of sheaves of alternatives nicely: As platform P1 needs to be used when driving towards and P2 when coming from Station 2, the two left most waiting arcs are inevitable ones. Each of the six described paths uses them both. The sheaf of alternatives in this case are the sub-paths of the previously described ones, starting at the departure node (yellow, \triangleright) at P1, ending the arrival node (white, \triangleleft) at P2.

Remark 3.6. \mathcal{N} and \mathcal{S} can easily be adapted to model routing choices beyond turnarounds by expanding the graph, e.g., if several platform tracks should be considered at a larger station, or if driving on the left track for a certain driving activity should be an option.

3.3 Model Formulation

We expand the model introduced in Section 2 to include track choice. Our approach generalizes the model introduced by Wüst et al. (2018), but can cover more practical situations: To our understanding, Wüst et al. (2018) make a track choice based on events grouped in so-called operation points, and headways (and only headways) are activated whenever two events are assigned to the same track. However, their model does not cover routing intricate paths within an operation point, and turnarounds do not play any role. For example, the decision if a train should be turned at a platform or at a siding, cannot be modeled.

Our main idea is to solve a “standard” PESP model on the entire network \mathcal{N} and to introduce two new kinds of variables for each arc $(i, j) \in \mathcal{A}(\mathcal{N})$, namely (implicit) binary variables $h_{i,j}$ indicating whether the arc (i, j) is part of a chosen routing – and the corresponding PESP constraints being thus active –, as well as $\bar{y}_{i,j}$ reflecting the periodic slack on an arc if it is chosen, and zero otherwise. For each alternative path F per sheaf of alternatives $S \in \mathcal{S}$, a binary variable $b_F \in \{0, 1\}$ indicates whether F is chosen or not.

The track choice can be tied to the timetabling constraints in different ways, some of which we will discuss in detail. Our entire model reads as follows, where the constraints modeling the SLACK ASSIGNMENT (11), ARC ACTIVATION (16) and HEADWAY CONDITIONS (17) can be expressed in multiple ways, which we discuss below. We also refer to the end of this section for details about the OBJECTIVE (10).

$$\begin{aligned}
\min \quad & \text{OBJECTIVE} & (10) \\
\text{s.t.} \quad & \text{SLACK ASSIGNMENT (IB/AB)} & (11) \\
& y_{ij} \leq u_{ij} - \ell_{ij} + (T - 1 - u_{ij} + \ell_{ij})(1 - h_{ij}) & (i, j) \in \mathcal{A}(\mathcal{N}) \quad (12) \\
& \bar{y}_{ij} \leq (u_{ij} - \ell_{ij})h_{ij} & (i, j) \in \mathcal{A}(\mathcal{N}) \quad (13) \\
& \bar{y}_{ij} \geq y_{ij} - (T - 1)(1 - h_{ij}) & (i, j) \in \mathcal{A}(\mathcal{N}) \quad (14) \\
& \sum_{F \in S} b_F = 1 & S \in \mathcal{S} \quad (15) \\
& \text{ARC ACTIVATION (B/H)} & (16) \\
& \text{HEADWAY CONDITIONS (Q0/Q3/Q4/QT)} & (17) \\
& \xi_{i_1 i_2} \geq \pi_{i_2} - \pi_{i_1} - T p_{i_1 i_2} + r_{i_1 i_2} + M(h_{i_1 j_1} + h_{i_2 j_2} - 2) & ((i_1, j_1), (i_2, j_2)) \in \bar{\mathcal{H}} \quad (18) \\
& \xi_{i_1 i_2} \leq M h_{i_1 j_1} & ((i_1, j_1), (i_2, j_2)) \in \bar{\mathcal{H}} \quad (19) \\
& \xi_{i_1 i_2} \leq M h_{i_2 j_2} & ((i_1, j_1), (i_2, j_2)) \in \bar{\mathcal{H}} \quad (20) \\
& 0 \leq \pi_i \leq T - 1 & i \in \mathcal{V}(\mathcal{N}) \quad (21) \\
& 0 \leq h_{ij} \leq 1 & (i, j) \in \mathcal{A}(\mathcal{N}) \quad (22) \\
& y_{ij} \geq 0 & (i, j) \in \mathcal{A}(\mathcal{N}) \quad (23) \\
& \bar{y}_{ij} \geq 0 & (i, j) \in \mathcal{A}(\mathcal{N}) \quad (24) \\
& p_{ij} \in \{0, 1, 2\} & (i, j) \in \mathcal{A}(\mathcal{N}) \quad (25) \\
& b_F \in \{0, 1\} & F \in \mathcal{S}, S \in \mathcal{S} \quad (26) \\
& \xi_{i_1 i_2} \geq 0 & ((i_1, j_1), (i_2, j_2)) \in \bar{\mathcal{H}} \quad (27)
\end{aligned}$$

Constraint (15) ensures that exactly one of the alternative paths F per sheaf S is chosen. The arc activation constraint (16) ensures that if an arc (i, j) is part of a chosen alternative, then h_{ij} will be set to 1, cf. Section 3.3.

The variables y_{ij} and p_{ij} model the periodic slack and offsets on the entire network, even on unactivated arcs. Similarly, the timetable values π_i are set even for unused events $i \in \mathcal{V}(\mathcal{N})$. If $h_{ij} = 1$, then (12) corresponds to the standard PESP slack constraint (3). For non-chosen arcs, i.e., those with $h_{ij} = 0$, the right-hand side in (12) results in $T - 1$, meaning that for any choice of π_i and π_j , valid y_{ij} and p_{ij} can be chosen.

The constraints (13) and (14) set \bar{y}_{ij} to zero if $(i, j) \in \mathcal{A}(\mathcal{N})$ is not part of the routing and to $\bar{y}_{ij} \geq y_{ij}$ otherwise, meaning that \bar{y}_{ij} models the periodic slack of active arcs.

The constraints (18), (19) and (20) incorporate track choice into our modeling approach of regularity as introduced in Section 2.4. Regular arrivals should obviously only be included along chosen routes. By introducing big- M constraints, in particular, for $M \geq 2T - 1$, the constraints (19) and (20) set $\xi_{i_1 i_2}$ to zero if, for a pair of arcs on the same infrastructure $((i_1, j_1), (i_2, j_2)) \in \bar{\mathcal{H}}$, not both (i_1, j_1) and (i_2, j_2) are part of an active route, while (18) results in the original regularity constraints (9) if both arcs are active.

Finally, (21)-(27) describe the variable bounds, cf. (3)-(5).

Since we only want to re-schedule on the infrastructure of the planning subgraph \mathcal{PT} , events in \mathcal{N} which use infrastructure outside of \mathcal{PT} are fixed to the regular annual timetable.

It should be pointed out that neither the regularity variables, nor the active slack variables \bar{y}_{ij} are necessary for feasibility. However, they are important for the objective func-

tion, as the goal is to evaluate the timetable depending on the route choice.

Slack Assignment

We have evaluated two different approaches to model the slack assignment. Firstly, one can stick to the standard PESP approach, namely by setting

$$(IB) \quad y_{ij} + \ell_{ij} = \pi_j - \pi_i + Tp_{ij}.$$

To keep the constraint structure for non-activated arcs, free upper bounds $u_{ij} = \ell_{ij} + T - 1$ must be used, leaving $p_{ij} \in \{0, 1, 2\}$ for most arcs. But for non-activated arcs, $\ell_{ij} = 0$ and $u_{ij} = T - 1 < T$ suffice. Thus, we propose the following bound-activation variant on ℓ_{ij}

$$(AB) \quad y_{ij} + \ell_{ij}h_{ij} = \pi_j - \pi_i + Tp_{ij}.$$

Together with (12) and in function of h_{ij} , this allows us to refine (25) to $p_{ij} \in \{0, 1\}$, only.

Arc Activation

We propose to handle arc activation in two ways. The first is the intuitive approach of explicitly setting h_{ij} as binary variables which should be set to 1 if and only if an alternative using said arc (i, j) is active. As exactly one alternative F per sheaf S is chosen, we can use an aggregated form:

$$(B) \quad h_{ij} = \sum_{S \in \mathcal{S}} \sum_{F \in S: (i,j) \in F} b_F \quad \text{and} \quad h_{ij} \in \{0, 1\}.$$

In order to reduce the number of integral variables, one can also consider a relaxed version, such that the variables h_{ij} are only implicitly binary:

$$(H) \quad h_{ij} \geq \sum_{S \in \mathcal{S}} \sum_{F \in S: (i,j) \in F} b_F \quad \text{and} \quad h_{ij} \leq 1.$$

This constraint ensures that h_{ij} is 1 if any of the chosen alternatives uses (i, j) . As for all arcs $(i, j) \in \mathcal{A}(\mathcal{N})$ all h_{ij} are part of the objective (cf. Section 3.3), h_{ij} is set to 0 in an optimal solution if (i, j) is not used.

Headway Conditions

For each of the headway formulations to model (ε, h) -conflict-free timetables, track activation needs to be handled in a slightly different manner. With the exception of the QT-formulation, which we want to keep as close to the standard PESP framework as possible, we want to introduce as few new variables as possible. Recall that we do not consider headways to be part of \mathcal{N} , except for QT. For the formulations of Q0, Q3, and Q4, we only introduce the corresponding offset variables, but the slacks are modelled only implicitly, and headway activation is handled via big- M constraints. As a discussion of all four model formulation would include many technical details, we will show how to include track choice for Q3 only. The other formulations can be modeled in a similar fashion.

Q3 The headway constraints of Q3 should be active for a pair of arcs on the same infrastructure $((i_1, j_1), (i_2, j_2)) \in \mathcal{H}$ only if both arcs (i_1, j_1) and (i_2, j_2) are part of a chosen

alternative. If one of them is not used, we turn this constraint insignificant as follows:

$$(Q3) \quad h(h_{i_1 j_1} + h_{i_2 j_2} - 1) \leq \pi_{i_2} - \pi_{i_1} + T p_{i_1 i_2} \leq (T - h)(3 - h_{i_1 j_1} - h_{i_2 j_2}) \quad (28)$$

$$\varepsilon(h_{i_1 j_1} + h_{i_2 j_2} - 1) \leq \pi_{j_2} - \pi_{i_1} + T p_{i_1 j_2} \leq (T - \varepsilon)(3 - h_{i_1 j_1} - h_{i_2 j_2}) \quad (29)$$

$$-(2 - h_{i_1 j_1} - h_{i_2 j_2}) \leq p_{i_1 j_1} + p_{j_1 i_2} - p_{i_1 i_2} \leq 2(2 - h_{i_1 j_1} - h_{i_2 j_2}) \quad (30)$$

$$p_{i_1 i_2}, p_{j_1 i_2} \in \{0, 1\} \quad (31)$$

Here (28) and (29) bound the (implicit) slack on the arcs (i_1, i_2) and (j_1, i_2) respectively if both arcs of the pair $((i_1, j_1), (i_2, j_2))$ are active. Otherwise, this constraint can be fulfilled trivially as for any choice of $\pi_{i_1}, \pi_{i_2} \in [0, T - 1]$. The inequality (30) turns exactly into

$$p_{i_1 j_1} + p_{j_1 i_2} - p_{i_1 i_2} = 0, \quad (32)$$

thus ensuring that the timetable is conflict-free by Proposition 2.8 if both infrastructure arcs (i_1, j_1) and (i_2, j_2) are active. The left term of (32) is between -1 and 2 , so that condition (30) is fulfilled by default if at least one of the arcs is not active, i.e., $h_{i_1 j_1} + h_{i_2 j_2} \leq 1$.

Objective

We model the objective as a combination of user and operator costs, and include the regularity values. Like introduced in Section 2.4 in the basic model, we consider operator costs to be the sum of cycle times per line and the user costs as the aggregated user-related activity durations for all chosen arcs. More precisely, we model operator and user costs by

$$c_{\text{operator}} = \sum_{(i,j) \in \mathcal{A}(\mathcal{N})} \bar{y}_{ij} + \ell_{ij} h_{ij}, \quad c_{\text{user}} = \sum_{(i,j) \in U} \bar{y}_{ij} + \ell_{ij} h_{ij},$$

respectively, where $U \subseteq \mathcal{A}(\mathcal{N})$ is the set of driving and waiting arcs, which are not within a terminal station of a line. Note that we minimize the periodic tension and not the periodic slack in the user and operator costs. In contrast to the standard PESP model, slack and tension cannot be used equivalently: As arc-choice is part of our optimization, it might be more cost efficient to use an alternative with smaller lower bounds and larger slack than an alternative with small slack if the lower bounds are then significantly larger.

As we are interested in regular timetables, our full objective then minimizes

$$\sum_{(i,j) \in \mathcal{A}(\mathcal{N})} \bar{y}_{ij} + \ell_{ij} h_{ij} + \sum_{(i,j) \in U} \bar{y}_{ij} + \ell_{ij} h_{ij} + \sum_{((i_1, j_1), (i_2, j_2)) \in \overline{\mathcal{H}}} \xi_{i_1 i_2}.$$

4 Computational Experiments

4.1 Overview

We implemented our model for generating and optimizing periodic timetables with integrated track choice for railway construction choice, including the various modeling variations described in Section 2 and Section 3. DB Netz AG, Germany's largest railway infrastructure manager, provided us with real-world data from the S-Bahn Berlin network. This suburban railway network comprises 168 stations in the Berlin metropolitan area and is operated with 16 lines, where each line is running at an interval of 5, 10 or 20 minutes. Already on this network, there are several hundred construction sites per year.




Figure 10: Construction site scenarios on the S-Bahn Berlin network. The red parts are unavailable (\mathcal{ZG}), the cyan parts are the planning subgraph (\mathcal{PG}).

Scenario	\mathcal{ZG}		$\mathcal{PG} \setminus \mathcal{ZG}$		\mathcal{N}		$ \mathcal{H} $	$ \mathcal{S} $	$\sum_{S \in \mathcal{S}} S $	$\sum_{l \in \mathcal{L}} f_l$
	$ V $	$ E $	$ V $	$ E $	$ \mathcal{V} $	$ \mathcal{A} $				
1	4	4	1	0	16	18	3	3	4	2
2	4	4	7	6	121	145	82	9	24	5
3	8	10	22	20	262	317	155	19	99	14
4	1	2	13	11	361	467	512	34	443	10
5	3	6	45	43	835	950	1079	46	166	22
6	4	5	41	42	996	1137	1152	67	240	16

Table 1: Size metrics of our scenarios: Vertices and edges of the unavailable subgraph \mathcal{ZG} , the planning subgraph \mathcal{PG} , and the resulting turn-sensitive event-activity network \mathcal{N} ; activity pairs using the same infrastructure \mathcal{H} , sheaves of alternatives \mathcal{S} , total number of alternatives and sum of scheduled line-frequencies.

We investigate six real construction scenarios of various sizes in 2021 and 2022, see Table 1 and Figure 10. For each of these scenarios, we consider the line plan that was actually put into operation, and try to compute and optimize a conflict-free periodic timetable with turnaround routings according to our model. The period time is 20 minutes, but as DB Netz AG is planning timetables with a resolution of 0.1 minutes, we scale by a factor of 10 to maintain integrality, so that $T = 200$.

We evaluate the performance and compare our various modeling features on our choice of six scenarios in the subsequent subsection. Let us mention that a discussion of the characteristics of the computed timetables is beyond the scope of this submission, both, due to space restrictions and ongoing work in progress with the practitioners.

4.2 Computational Results

We compared each of the 16 model configurations described in Section 3.3 on the six scenarios. We solved each scenario-model pair with the MIP solver Gurobi 9.1.2 (Gurobi Optimization, LLC, 2022) on an Intel i7-9700K CPU machine with 64 GB RAM with a wall time limit of 90 minutes. We were able to find a solution for all of our scenarios – in all but the last one even within the first 30 seconds. The three smallest scenarios could be solved to optimality. The other three instances terminated with optimality gaps below 7%.

For a qualitative comparison of the different model configurations, we examined when a first solution is found, as quickly finding any solution can be of interest, particularly to answer the question if a given line concept can be operated at all. Figure 11 displays when the first solution is found, with marks at ∞ if none was found. The model configurations are denoted by their SLACK ASSIGNMENT, HEADWAY CONDITIONS, and ARC ACTIVATION.




Figure 11: Time to first found solution in seconds

In the tiny first scenario, all models find the optimal solution directly, while with increasing complexity of the instance, differences emerge: E.g., in Scenario 5, all of the models find a solution within the first 90 seconds, except for QT-models, for which the values range from 242 up to 1359 seconds. In fact, we observe that QT performs badly in comparison with the other formulations as well: In Scenario 4, a solution is found significantly later for both QT with B- formulations; for QT in combination with the H-form, no solution is found at all. In the largest scenario, Scenario 6, none of the QT-models produced a solution.

When it comes to comparing the bound activation, there do not seem to be any clear winners with respect to time to the first solution, as both IB and AB perform similarly for the same HEADWAY CONDITIONS and ARC ACTIVATION. A preferred choice for ARC ACTIVATION is not evident either: In Scenario 5, for example, the B- seems to behave better than the H-formulation in combination with IB, while the opposite is true for AB.

Interesting is also the behaviour in Scenario 6: Here, only five of the configurations are solvable at all. In this instance, IB_Q4_B finds a primal solution significantly earlier than the other four, yet none of the other three Q4-formulations provide a solution in the given time. In contrast, three Q3-formulations find a solution.

These results support our assumption that the headway constraints are one of the main difficulties when solving track-choice PESP: The more activity pairs \mathcal{H} on the same infrastructure, the later a primal solution is found. However, there also seems to be a correlation to the number of alternatives: The two scenarios with the largest amount of alternatives in total, namely Scenarios 4 and 6, were the only ones in which some model formulations did not find any feasible solution within 90 minutes.

Apart from feasible solutions, we are interested in optimality. Figure 12 displays the relative best dual bounds of each model after 90 minutes, i.e., $\text{dual}(m)/\text{dual}(m_{\text{best}})$, where $\text{dual}(m)$ is the lower bound of model m and $\text{dual}(m_{\text{best}})$ is the maximum lower bound of all models of that scenario. Most prominent is the weak performance of QT: As before, it behaves significantly worse than the other formulations. When comparing IB with AB, our results are inconclusive again. However, for the arc-activation formulation – with a few singular exceptions – the B-form seems to provide better results than the H-form.

When comparing the quality of the primal solutions after the run time of 90 min, all model formulations, except for QT are comparable: In Scenarios 1-3, which we were able




Figure 12: Relative best dual values $dual(m)/dual(m_{best})$

to solve to optimality, all formulation found the optimal solution. In Scenarios 4 and 5, the worst solution was less than 0.01% worse than the best. The models that found a feasible solution in Scenario 6 differed by less than 3% from the best found solution, which was provided by model AB_Q0_H.

We conclude that including track choice in periodic timetabling for our construction site scenarios is well feasible: For small scenarios all of our models performed similarly well. For larger scenarios, however, we do not recommend QT – the large increase of integral variables caused by the refinement seems to have a large negative impact. In particular, for our application case this underlines the benefit of MIP models that go beyond pure implementations of PESP. In contrast, while Q0 with the least amount of integral variables performed well, it provided qualitatively similar results as other headway formulations. If a certificate of optimality is of importance, we suggest one of the B formulations, as they consistently provided good dual bounds.

5 Outlook

We demonstrated that the automatic generation of conflict-free periodic timetables for railway constructions sites is feasible in practice. However, two main issues remain: The choice of a line plan is essential and not at all clear. We therefore aim to integrate the timetabling model into a larger framework that also includes line planning focusing on both passengers’ and operators’ perspectives. Moreover, given that a large quantity of heuristics is known for PESP (see, e.g., Borndörfer et al., 2020), it is a natural question whether and how these methods extend to our model. We believe that using primal heuristics should be able to speed up the solution process significantly.

Acknowledgements

Special thanks go to the team of Jens Brandenburger at DB Netz AG for our continuous, fruitful exchange, sharing their practical experience and for the provision of real-world data.

References

- Borndörfer, R., Karbstein, M., Liebchen, C., and Lindner, N. (2018). A simple way to compute the number of vehicles that are required to operate a periodic timetable. In Borndörfer, R. and Storandt, S., editors, *Proceedings of ATMOS 2018*, volume 65 of *OpenAccess Series in Informatics (OASICs)*, pages 16:1–16:15, Dagstuhl, Germany. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Borndörfer, R., Lindner, N., and Roth, S. (2020). A concurrent approach to the periodic event scheduling problem. *Journal of Rail Transport Planning & Management*, 15:100175. Best Papers of RailNorrköping 2019.
- Gurobi Optimization, LLC (2022). Gurobi Optimizer Reference Manual.
- Kroon, L. G., Huisman, D., Abbink, E., Fioole, P.-J., Fischetti, M., Maroti, G., Schrijver, A., Steenbeek, A., and Ybema, R. (2009). The new Dutch timetable: The OR revolution. *INTERFACES*, 39(1):6–17.
- Kroon, L. G., Romeijn, H. E., and Zwaneveld, P. J. (1997). Routing trains through railway stations: complexity issues. *Eur. J. Oper. Res.*, 98:485–498.
- Liebchen, C. (2006). *Periodic timetable optimization in public transport*. PhD thesis, Dissertation.de, Berlin.
- Liebchen, C. (2008). The first optimized railway timetable in practice. *Transp. Sci.*, 42(4):420–435.
- Liebchen, C. and Möhring, R. H. (2007). The modeling power of the periodic event scheduling problem: Railway timetables — and beyond. In Geraets, F., Kroon, L., Schoebel, A., Wagner, D., and Zaroliagis, C. D., editors, *Algorithmic Methods for Railway Optimization*, pages 3–40, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Lindner, N. and Reisch, J. (2022). An analysis of the parameterized complexity of periodic timetabling. *Journal of Scheduling*, 25(2):157–176.
- Lindner, T. (2000). *Train Scheduling in Public Rail Transport*. PhD thesis, Technische Universität Braunschweig.
- Nachtigall, K. (1996). Cutting planes for a polyhedron associated with a periodic network. Institutsbericht IB 112-96/17, Deutsche Forschungsanstalt für Luft- und Raumfahrt e.V.
- Nachtigall, K. (1998). *Periodic Network Optimization and Fixed Interval Timetables*. Habilitation thesis, Universität Hildesheim.
- Odiijk, M. A. (1994). Construction of periodic timetables, part 1: A cutting plane algorithm. Technical Report 94-61, TU Delft.
- Peeters, L. (2003). *Cyclic Railway Timetable Optimization*. PhD thesis, Erasmus Universiteit Rotterdam.
- Schrijver, A. and Steenbeek, A. (1993). Spoorwegdienstregelontwikkeling. Technical report, Centrum voor Wiskunde en Informatica.

- Serafini, P. and Ukovich, W. (1989). A Mathematical Model for Periodic Scheduling Problems. *SIAM Journal on Discrete Mathematics*, 2(4):550–581.
- Van Lieshout, R. N. (2021). Integrated Periodic Timetabling and Vehicle Circulation Scheduling. *Transportation Science*, 55(3):768–790. Publisher: INFORMS.
- Wüst, R., Bütikofer, S., Ess, S., Gomez, C., Steiner, A., Laumanns, M., and Szabó, J. (2018). Periodic timetabling with 'track choice'-pesp based on given line concepts and mesoscopic infrastructure. In Fortz, B. and Labbé, M., editors, *OR Proceedings 2018*, pages 571–578. Springer.