

Integrated Periodic Timetabling and Vehicle Circulation Scheduling

Rolf N. van Lieshout¹

¹Econometric Institute and ECOPT, Erasmus University Rotterdam, 3000 DR Rotterdam, The Netherlands. Email address: vanlieshout@ese.eur.nl

Econometric Institute Report Series - EI2019-27

Abstract

Periodic timetabling is one of the most well researched problems in the public transport optimization literature. However, the impact timetabling has on the number of required vehicles, which directly translates to operator costs, is rarely considered. Therefore, in this paper, we consider the problem of jointly optimizing the timetable and the vehicle circulation schedule, which specifies the cyclic sequences of trips vehicles perform. In order to be able to solve realistic instances, we improve an earlier proposed formulation by contraction techniques, valid inequalities and symmetry-breaking constraints. Ultimately, this allows us to explore the trade-off between the number of vehicles and the attractiveness of the timetable from the passengers' perspective. An extensive computational study demonstrates the effectiveness of the improved formulation. Moreover, using this approach we are able to find timetables requiring substantially fewer vehicles at the cost of minimal increases of the average travel time of passengers.

1 Introduction

The timetable is the foundation of any public transportation system. It determines the travel times of passengers and the paths they take to travel from their origins to their destinations. Furthermore, the extent to which disturbances such as delays propagate through the system also highly depends on the timetable. All in all, it is widely accepted that the timetable determines the attractiveness of the public transportation system from the passengers' point of view. Because of the importance of a good timetable, and the fact that finding such a timetable is a non-trivial exercise, the timetabling problem has triggered many researchers in the field of public transport optimization. However, the impact the timetable has on the costs of the operator is rarely considered in timetabling models.

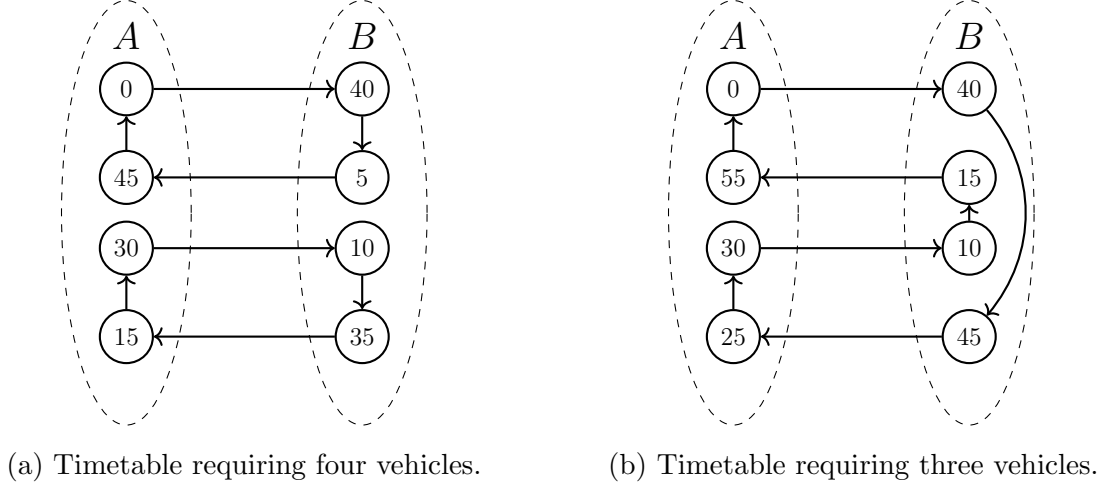


Figure 1: Two timetables for a line (A, B) with a frequency of 2 per hour and a driving time of 40 minutes.

In this paper, we consider *periodic timetabling*. In a periodic timetable, the schedule of a single period (e.g. one hour) is repeated throughout the day. Periodic timetables are favoured by passengers as they can easily memorize the departure times of services (e.g. a traveller can always take the train at $xx:06$). Furthermore, operators that use a periodic timetable only need to plan services for one period, after which they can repeat the timetable during the entire day.

Most periodic timetabling models are based on the Periodic Event Scheduling Problem (PESP). Serafini and Ukovich (1989) introduce the PESP and prove it is NP-complete. In its traditional form, the PESP is a feasibility problem. Later, researchers have started to consider optimization variants of the PESP, where the objective typically is to minimize passenger travel time or to maximize the robustness of the timetable. The PESP can be solved using constraint programming or SAT solvers (Gattermann, Großmann, Nachtigall, & Schöbel, 2016; Großmann et al., 2012; Kroon et al., 2009; Schrijver & Steenbeek, 1993), the modulo-simplex heuristic (Borndörfer, Hoppmann, Karbstein, & Löbel, 2016; Goerigk & Schöbel, 2013; Nachtigall & Opitz, 2008) or mixed-integer programming based approaches (Goerigk & Liebchen, 2017; Liebchen, 2008; Maróti, 2017; Polinder, Breugem, Dollevoet, & Maróti, 2019).

The goal of this paper is to find timetables that are both attractive for the passengers and for the operator. Specifically, we analyze the trade-off between the travel time of passengers and *the number of vehicles required to operate the timetable*. To illustrate the impact a timetable has on the required number of vehicles, consider the two timetables in Figure 1. The timetable in Figure 1a induces two vehicle circulations, cyclic sequences of trips performed by the same vehicles, both requiring two vehicles. In contrast, the timetable in Figure 1b induces one vehicle circulation requiring three vehicles. Furthermore, both timetables yield the same travel times for passengers. Therefore, as vehicles constitute a large proportion of the operator costs, the second timetable is clearly preferred over the first one.

Traditionally, timetabling and vehicle scheduling are performed sequentially; first the timetable is optimized to minimize travel times of passengers and subsequently the vehicles are scheduled to minimize operator costs. However, this is a greedy approach that leads to globally sub-optimal solutions. More recently, integrated public transport planning, where the goal is to find a good overall line plan, timetable and vehicle schedule, has seen increasing attention (Burggraeve, Bull, Vansteenwegen, & Lusby, 2017; Kaspi & Raviv, 2013; Pätzold, Schiewe, Schiewe, & Schöbel, 2017; Schöbel, 2017; Van Lieshout & Bouman, 2019). The potential of jointly optimizing the timetable and vehicle schedule has already been demonstrated for non-periodic timetables by Desfontaines and Desaulniers (2018); Fonseca, Van der Hurk, Roberti, and Larsen (2018); Ibarra-Rojas, Giesen, and Rios-Solis (2014) and Schmid and Ehmke (2015). Even when only limited modifications to an initial timetable are allowed, large savings in operational costs can be obtained without dramatic increases in travel times.

In contrast, only a few contributions consider the integration of periodic timetabling and vehicle scheduling. Peeters (2003) notes that if the operated vehicle circulations are fixed a priori, the number of required vehicles can easily be counted within the PESP. Furthermore, the author shows that choosing the circulation can be incorporated in the PESP (without introducing auxiliary variables) for the case that at most two services terminate at each station. However, this only works under the *sufficient spread* assumption, which states that there is never more than one vehicle idle at a station. Under the same assumption, Kroon, Peeters, Wagenaar, and Zuidwijk (2013) are able to choose vehicle turnarounds within the PESP between an arbitrary number of services. Nührenberg (2016) proposes an explicit matching approach that is also valid when the sufficient spread assumption does not hold, which requires extending the PESP with matching variables and constraints. In this model, both the timetable and the vehicle circulations are entirely determined. We refer to this completely integrated problem as the Vehicle Circulation Periodic Event Scheduling Problem (VC-PESP). However, the proposed formulation contains many big- M constraints, causing it to not perform well, even on medium-sized instances. Next to that, the author purely focuses on minimizing the number of required vehicles, without regard to the passengers' perspective.

In this paper, we improve the formulation of Nührenberg (2016), in order to be able to tackle realistic instances of the VC-PESP and gain insight in the trade-off between the average travel time a timetable offers and the number of vehicles it requires. To do so, we first formalize and analyze the vehicle circulation scheduling problem (given a timetable), which, in contrast to traditional non-periodic vehicle scheduling (see e.g. Bunte and Kliwer (2009)), has not received much attention in the literature. Next, we use the properties of the vehicle circulation scheduling problem to develop a stronger formulation of the VC-PESP. In particular, to reduce the number of big- M constraints, we develop a modeling technique that computes the minimum turnaround time at stations in a contracted graph. Moreover, we present new valid inequalities that bound the minimum turnaround time of vehicles, strengthening the linear programming relaxation. Finally, we proof that a greedy

algorithm is optimal for the vehicle circulation scheduling problem and use this observation to develop symmetry-breaking inequalities based.

Besides strengthening the formulation of the VC-PESP, we also extend the problem by considering restrictions on the number of lines that may occur in a single vehicle circulation. In practice, public transport operators prefer short circulations where vehicles perform only a single line or alternate between two lines, as such circulations are more easy to operate and avoid dependencies between different parts of the network. We show how such practical requirements can be incorporated in the VC-PESP.

Computational experiments with different instances based on the Dutch railway network demonstrate the effectiveness of the proposed improved formulation. When minimizing the number of required vehicles, the improved formulation finds timetables requiring fewer vehicles compared to the original formulation, reducing the average optimality gap from 5.0 to 2.7 percent. An analysis of the trade-off between the number of vehicles and the perceived travel time of timetables illustrates the benefit of integrating circulation scheduling within the timetabling problem: minimizing average travel time without taking vehicles into account results in relatively costly timetables. In more than half of the considered instances, our approach finds timetables that require fewer vehicles without any increase in the travel time. Moreover, if one permits a travel time increase of 0.1 percent, we are typically able to reduce the number of vehicles with about 10 percent. Furthermore, the experiments show that in most instances vehicle circulations consisting of at most two lines suffice for realizing very efficient timetables and only small decreases in travel time can be attained if longer vehicle circulations are allowed.

Summarizing, the main contributions of this paper are (1) the analysis of the vehicle circulation scheduling problem, (2) the improvement of a mathematical formulation to efficiently solve periodic timetabling problems while controlling the number of required vehicles, (3) showing that neglecting vehicle circulations in the periodic timetabling problem leads to very inefficient timetables; considerable savings in the number of vehicles can generally be achieved without compromising the passengers' perspective and (4), demonstrating that a large proportion of these savings can already be obtained with vehicle circulations consisting of at most two lines, which are attractive for operators.

The remainder of this paper is structured as follows. In Section 2, we discuss the periodic timetabling problem and in Section 3, the vehicle circulation scheduling problem. In Section 4, we describe the mathematical formulation for the VC-PESP that is proposed in Nührenberg (2016). In Section 5, we analyze the computational complexity of the integrated problem. In Section 6, we present the improved formulation for the VC-PESP. In Section 7, we describe how restrictions on the number of lines in a circulation can be included in the formulation. In Section 8, we discuss the results of the conducted computational experiments. Finally, in Section 9, we conclude the paper and describe future research directions.

2 Periodic Timetabling

In periodic timetabling, the goal is to determine departure and arrival times for a given set of services, such that all operational and service requirements are met. As the schedule repeats every period (e.g. one hour), it suffices to determine departure times and arrival times for a single period.

2.1 The Periodic Event Scheduling Problem

The most widely applied method to formulate periodic timetabling problems is through the Periodic Event Scheduling Problem (PESP), introduced in Serafini and Ukovich (1989). The input to the PESP is a directed network $N = (E, A)$ referred to as *event-activity network*, a period or cycle time T and lower and upper bounds l_a and u_a for all $a = (e, f) \in A$. Every node $e \in E$ represents a repeating event that needs to be scheduled at some time $\pi_e \in [0, T)$ and every arc $a \in A$ represents an activity for which the duration should be between the specified bounds:

Definition 1 (PESP). Given a period T , an event-activity network $N = (E, A)$ and lower and upper bounds l_a and u_a for all $a = (e, f) \in A$, the Periodic Event Scheduling Problem is to determine event times $\pi_e \in [0, T)$ for all $e \in E$, satisfying

$$(\pi_f - \pi_e) \text{ modulo } T \in [l_a, u_a] \text{ for all } a = (e, f) \in A$$

or to conclude that no such schedule exists.¹

Serafini and Ukovich (1989) prove the NP-completeness of the PESP. Odijk (1997) shows the problem is NP-complete even for fixed $T \geq 3$ by a reduction from graph coloring.

The PESP can be used to formulate a wide range of timetabling constraints. The events that need to be scheduled are the arrivals and departures of services at all stations. Activities are used to model various operational and service requirements of the timetable, covering e.g. driving and dwell times, minimum and maximum transfer times and minimum separation times. For a more comprehensive overview, we refer to Peeters (2003) and Liebchen and Möhring (2007).

2.2 Mixed Integer Programming Formulations

By introducing an integer variable p_a for all $a \in A$, the PESP can be formulated as the following mixed integer program:

$$\text{find } (\boldsymbol{\pi}, \mathbf{p}) \tag{1}$$

$$\text{s.t. } l_a \leq \pi_f - \pi_e + Tp_a \leq u_a \quad \forall a = (e, f) \in A, \tag{2}$$

$$0 \leq \pi_e \leq T \quad \forall e \in E, \tag{3}$$

$$p_a \in \mathbb{Z} \quad \forall a \in A. \tag{4}$$

¹For this definition, we assume w.l.o.g. that $0 \leq l_a \leq u_a \leq T$ (Peeters, 2003).

In this formulation, the p_a variables take the role of the modulo operator. For example, consider an activity $a = (e, f)$ with $l_a = 5$ and $u_a = 15$ and let $T = 30$. If $\pi_f = 5$ and $\pi_e = 25$, p_a can take the value 1, such that $\pi_f - \pi_e + Tp_a = 5 - 25 + 30 = 10 \in [5, 15]$.

The standard formulation of the PESP (1)-(4) has a very weak linear programming (LP) relaxation and does not perform well on reasonably sized instances. A stronger formulation can be obtained by formulating the PESP in terms of the activity durations or *tensions* $x_a = \pi_f - \pi_e + Tp_a$ and exploiting the *cycle periodicity* property discovered by Odijk (1996):

Lemma 1 (Cycle Periodicity (Odijk, 1996)). *Consider a feasible PESP solution $(\boldsymbol{\pi}, \mathbf{p})$ and let $x_a = \pi_f - \pi_e + Tp_a$ for $a = (e, f) \in A$. Let C denote a cycle in the event-activity network, with forward arcs C^+ and backward arcs C^- . It holds that*

$$\sum_{a \in C^+} x_a - \sum_{a \in C^-} x_a = q_C T \text{ for some integer } q_C.$$

Moreover, $a_C \leq q_C \leq b_C$, with

$$a_C = \left\lceil \frac{\sum_{a \in C^+} l_a - \sum_{a \in C^-} u_a}{T} \right\rceil \text{ and } b_C = \left\lfloor \frac{\sum_{a \in C^+} u_a - \sum_{a \in C^-} l_a}{T} \right\rfloor$$

Lemma 1 shows that the cycle periodicity property is a necessary condition for a feasible solution to the PESP. Odijk (1996) proves that it is in fact both necessary and sufficient, so if the property holds for all cycles, the solution is feasible. The final ingredient for the cycle periodicity formulation is asserted by Nachtigall (1994), who shows that if the periodicity property holds for the cycles in a carefully chosen subset of the set of all cycles, it actually holds for all cycles in the event-activity network. Specifically, if the cycle periodicity property holds for all cycles in an *integral cycle basis* of the event-activity network, it holds for any cycle (Liebchen & Peeters, 2009).

From these results, it follows that the PESP can be formulated in terms of tension variables x_a for all arcs and periodicity variables q_C for all cycles in an integral cycle basis \mathcal{B} . The entire cycle periodicity formulation reads as follows:

$$\text{find } (\mathbf{x}, \mathbf{q}) \tag{5}$$

$$\text{s.t. } l_a \leq x_a \leq u_a \quad \forall a \in A, \tag{6}$$

$$\sum_{a \in C^+} x_a - \sum_{a \in C^-} x_a = q_C T \quad \forall C \in \mathcal{B}, \tag{7}$$

$$a_C \leq q_C \leq b_C \quad \forall C \in \mathcal{B}, \tag{8}$$

$$x_a \in \mathbb{R} \quad \forall a \in A, \tag{9}$$

$$q_C \in \mathbb{Z} \quad \forall C \in \mathcal{B}. \tag{10}$$

Constraints (6) impose the lower and upper bounds for the tensions. Constraints (7) model the periodicity of the timetable through the cycle periodicity property. Constraints (8) are not required for the correctness of the formulation but are included to strengthen the LP relaxation.

The cycle periodicity formulation contains only $|\mathcal{B}| = |A| - |E| + 1$ integer variables, compared to $|A|$ in the standard formulation. Furthermore, the equality constraints (7) lead to a stronger propagation of branching decisions. Combined, this makes the cycle periodicity formulation the preferred model for solving the PESP.

3 Vehicle Circulation Scheduling

Once the timetable is known, operators need to schedule the vehicles in order to cover all trips. This includes determining sequences of trips performed by vehicles, and in e.g. railway contexts also coupling and decoupling decisions in order to better meet demand. In this paper, we limit ourselves to *vehicle circulation scheduling*, which involves determining which trips are operated consecutively *within the periodic pattern*, under the assumption that a single vehicle suffices to cover a trip. As the timetable is periodic, every trip should have exactly one successor and one predecessor, which results in cycles of trips, referred to as circulations, that are performed by one or multiple vehicles. For example, if T is 60 minutes and a certain circulation takes 180 minutes, three vehicles are needed to operate this circulation. The vehicle circulation scheduling problem is to find the set of circulations that cover all trips and minimize the number of required vehicles.

To model vehicle transitions between trips, we let $E_{\text{end}}(s) \in E$ and $E_{\text{start}}(s) \in E$ denote the ends and starts of trips at station s , and extend the event-activity network with *turnaround activities* $A_{\text{turn}}(s)$ linking end events $e \in E_{\text{end}}(s)$ with start events $f \in E_{\text{start}}(s)$. An example is depicted in Figure 2. In general, these activities could also involve a vehicle driving from one station to another without passengers (a deadhead trip). However, in this paper we assume that there are only turnaround activities between ends and starts of trips at the same station. Furthermore, we assume that all transitions between trip ends and trip starts at the same station are possible, such that the *turnaround graph* $G_{\text{turn}}(s) = (E_{\text{end}}(s), E_{\text{start}}(s), A_{\text{turn}}(s))$ is a complete bipartite graph. We let $E_{\text{end}}, E_{\text{start}}$ and A_{turn} denote the union of the sets $E_{\text{end}}(s), E_{\text{start}}(s)$ and $A_{\text{turn}}(s)$ over all stations s , respectively. Finally, without loss of generality $l_a = 0$ and $u_a = T$ for all $a \in A_{\text{turn}}$, as the minimum time between trips is incorporated in the final dwell activity before the end event.




Figure 2: Event-activity network with turnaround activities indicated by the dashed arcs.

Let A_{veh} denote all drive and dwell activities in the event-activity network. Then, a vehicle

circulation schedule is a set of directed cycles \mathcal{C} in the event-activity network, consisting of vehicle activities and turnaround activities, such that all vehicle activities are covered exactly once. The cycles represent the sequences of trips performed by vehicles in the periodic pattern. As trips are connected by turnaround activities, there is a direct correspondence between the vehicle circulation schedule and the selected turnaround activities, denoted as $\mathcal{T}(\mathcal{C})$. Given a feasible timetable, the sum of the tensions of a vehicle circulation are an integer multiple of T by Lemma 1. Hence, for a given timetable \mathbf{x} , the number of vehicles required to operate the timetable using vehicle circulation schedule \mathcal{C} , denoted as $\nu(\mathbf{x}, \mathcal{C})$, equals

$$\nu(\mathbf{x}, \mathcal{C}) = \frac{1}{T} \sum_{c \in \mathcal{C}} \sum_{a \in c} x_a = \frac{1}{T} \left(\sum_{a \in A_{\text{veh}}} x_a + \sum_{a \in \mathcal{T}(\mathcal{C})} x_a \right).$$

For a given timetable, only the selected turnaround activities still need to be determined. To model this problem, let $\delta^+(e)$ denote the set of turnaround activities leaving e for $e \in E_{\text{end}}$ and let $\delta^-(e)$ denote the set of turnaround activities entering e for $e \in E_{\text{start}}$. Next, we introduce the binary decision variable y_a for $a \in A_{\text{turn}}$ which is equal to 1 if a is selected and 0 otherwise. Then, for a given timetable \mathbf{x} we can find a vehicle circulation schedule that minimizes the number of required vehicles by solving the following problem:

$$\min \quad \sum_{a \in A_{\text{turn}}} x_a y_a \quad (11)$$

$$\text{s.t.} \quad \sum_{a \in \delta^+(e)} y_a = 1 \quad \forall e \in E_{\text{end}}, \quad (12)$$

$$\sum_{a \in \delta^-(e)} y_a = 1 \quad \forall e \in E_{\text{start}}, \quad (13)$$

$$y_a \in \{0, 1\} \quad \forall a \in A_{\text{turn}}. \quad (14)$$

The objective (11) is to minimize the total turnaround time. Constraints (12) and (13) ensure that every end event has a successor and every start event has a predecessor. As already observed by Orlin (1982) and Serafini and Ukovich (1989), it follows that finding a vehicle circulation schedule requiring the fewest number of vehicles, comes down to solving a weighted perfect matching problem in a bipartite graph. Furthermore, as we assume that deadheading is not allowed, the bipartite graph decomposes into complete bipartite graphs per station, such that a global optimal solution can be found by solving a weighted perfect matching problem for each terminal station independently. Borndörfer, Karbstein, Liebchen, and Lindner (2018) prove that solving this formulation is actually equivalent with solving a traditional non-periodic vehicle scheduling problem (see e.g. Bunte and Kliwer (2009)) on a roll out of the periodic timetable over the day.

3.1 Properties

We now present properties of the vehicle circulation scheduling problem. Later, we use these properties to develop a strong formulation for the VC-PESP.

Our analysis is based on the observation that the arc weights x_a of the weighted perfect matching problem are generated from a special metric. Specifically, there exist end times $\pi_e \in [0, T)$ for all $e \in E_{\text{end}}$ and start times $\pi_f \in [0, T)$ for all $f \in E_{\text{start}}$ such that $x_a = \pi_f - \pi_e$ modulo T for all $a = (e, f)$. Hence, the end events and start events can be represented by points on the circle with circumference T , such that we can view the problem as *clockwise bipartite perfect matching on a circle*. To illustrate, Figure 3a depicts an example of such a matching problem and a possible solution. The following lemma provides a first result of having this structure:

Lemma 2. *Let y^1 and y^2 denote feasible solutions to (11)-(14). For a feasible timetable \mathbf{x} , it holds that $\sum_{a \in A_{\text{turn}}} x_a y_a^1 - \sum_{a \in A_{\text{turn}}} x_a y_a^2 = zT$ for some integer z .*

Proof. Intuitively, this statement must be true since the number of vehicles required to operate a timetable must always be integer. For a formal proof, let $M^1 = \{a \in A_{\text{turn}} : y_a^1 = 1\}$ and $M^2 = \{a \in A_{\text{turn}} : y_a^2 = 1\}$. From matching theory, it is known that the symmetric difference $M^1 \oplus M^2 = (M^1 \setminus M^2) \cup (M^2 \setminus M^1)$ is the union S of vertex-disjoint cycles. Within these cycles, arcs from M^1 and M^2 alternate. It follows that

$$\begin{aligned} \sum_{a \in A_{\text{turn}}} x_a y_a^1 - \sum_{a \in A_{\text{turn}}} x_a y_a^2 &= \sum_{C \in S} \left(\sum_{a \in C: a \in M^1} x_a - \sum_{a \in C: a \in M^2} x_a \right) \\ &= \sum_{C \in S} q_C T && \text{(Lemma 1)} \\ &= \left(\sum_{C \in S} q_C \right) T \end{aligned}$$

for integers q_C . □

To further analyze the problem, we define the *inventory function* of a matching $I_s(t, M)$, which denotes the number of idle vehicles at station s under matching M at time t . The inventory function of the matching visualized in Figure 3a is presented in Figure 3b. Using this concept, the following lemma characterizes optimal solutions to the matching problem.

Lemma 3. *Consider a station s with turnaround arcs $A_{\text{turn}}(s)$ and activity durations x_a for all $a \in A_{\text{turn}}(s)$. Let $M \subseteq A_{\text{turn}}(s)$ and $M' \subseteq A_{\text{turn}}(s)$ denote perfect matchings from the associated end events to start events. The following statements are true:*

- (i) $\sum_{a \in M} x_a = \int_0^T I_s(t, M) dt$.
- (ii) $I_s(t, M) - I_s(t, M') = z$ for all $t \in [0, T)$ for some constant integer z .
- (iii) M is a minimum cost perfect matching if and only if $I_s(t, M) = 0$ for some $t \in [0, T)$.

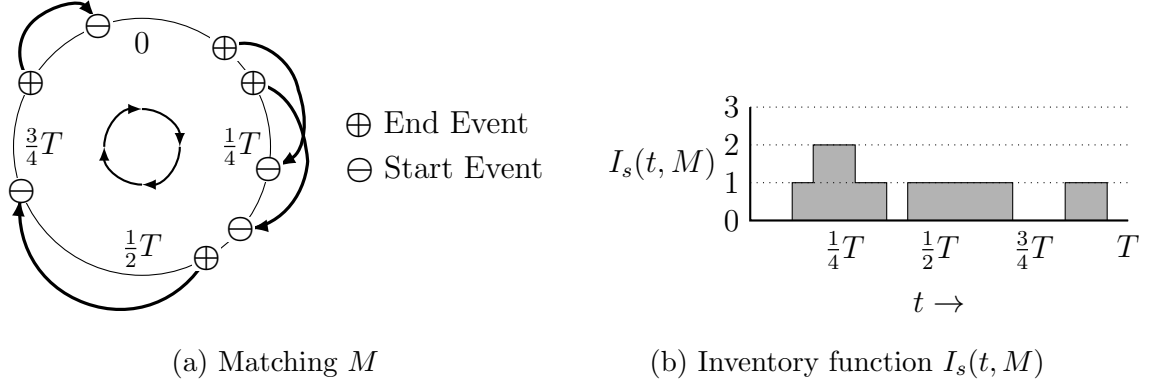


Figure 3: Example of a matching at a station and the associated inventory function.

Proof. (i) Consider an arc $a = (e, f)$. Let $I_s(t, a)$ be equal to 1 if a "covers" time instant t and 0 otherwise. Formally, if $\pi_e \leq \pi_f$, a covers t if $t \in [\pi_e, \pi_f]$, otherwise a covers t if $t \in [0, \pi_f] \cup [\pi_e, T]$. It holds that $I_s(t, M) = \sum_{a \in M} I_s(t, a)$. Hence, $\int_0^T I_s(t, M) dt = \int_0^T \sum_{a \in M} I_s(t, a) dt = \sum_{a \in M} \int_0^T I_s(t, a) dt = \sum_{a \in M} x_a$.

(ii) Let $\text{end}(a, b)$ and $\text{start}(a, b)$ denote the end and start events in the interval $[a, b]$, respectively. It holds that $I_s(t, M) = I_s(0, M) + \text{end}(0, t) - \text{start}(0, t)$. Therefore, $I_s(t, M) - I_s(t, M') = I_s(0, M) - I_s(0, M')$, which is integer.

(iii) First assume that $I_s(t', M) = 0$ for some $t' \in [0, T)$. Let M^{opt} denote the optimal matching. Clearly, $I_s(t', M^{\text{opt}}) \geq 0$, such that $I_s(t', M) - I_s(t', M^{\text{opt}}) \leq 0$. From (ii), it follows that $I_s(t, M) - I_s(t, M^{\text{opt}}) \leq 0$ for all t . From (i), we then find that

$$\sum_{a \in M} x_a - \sum_{a \in M^{\text{opt}}} x_a = \int_0^T I_s(t, M) - I_s(t, M^{\text{opt}}) dt \leq 0,$$

hence M is optimal. To prove the reverse direction, assume that M is optimal but $I_s(t, M) > 0$ for all $t \in [0, T)$. Then, there exist arcs $(e_1, f_1), \dots, (e_m, f_m)$ that cover all $t \in [0, T)$. Without loss of generality, we can assume that

$$\pi_{e_1} < \pi_{f_m} < \pi_{e_2} < \pi_{f_1} < \dots < \pi_{e_m} < \pi_{f_{m-1}}$$

Clearly, the cost of the matching can be decreased by selecting the arcs $(e_1, f_m), (e_2, f_1), \dots, (e_m, f_{m-1})$, contradicting the optimality of M . \square

Next, we present a greedy algorithm for vehicle circulation scheduling in Algorithm 1. The algorithm simply iteratively matches unmatched end events with the unmatched start event that gives the shortest turnaround time.

Theorem 4. *The matching M^{greedy} returned by Algorithm 1 is optimal for the vehicle circulation scheduling problem.*

Algorithm 1 Greedy algorithm for vehicle circulation scheduling

Input: Turnaround graph $G_{\text{turn}}(s) = (E_{\text{end}}(s), E_{\text{start}}(s), A_{\text{turn}}(s))$ and
turnaround times $x_{e,f} \forall (e, f) \in A_{\text{turn}}(s)$

Output: Minimum cost perfect matching M in $G_{\text{turn}}(s)$

- 1: $M = \emptyset, U_{\text{end}} = E_{\text{end}}(s), U_{\text{start}} = E_{\text{start}}(s)$
 - 2: **while** M is not a perfect matching **do**
 - 3: Pick any event $e \in U_{\text{end}}$
 - 4: Find $f = \operatorname{argmin}_{g \in U_{\text{start}}} x_{e,g}$
 - 5: $M \leftarrow M \cup (e, f)$
 - 6: $U_{\text{end}} \leftarrow U_{\text{end}} \setminus e, U_{\text{start}} \leftarrow U_{\text{start}} \setminus f$
 - 7: **Return** M
-

Proof. Assume M^{greedy} is not optimal. Then, by part (iii) of Theorem 3, $I_s(t, M^{\text{greedy}}) > 0$ for all $t \in [0, T)$. This implies that there exist arcs $(e_1, f_1), \dots, (e_m, f_m)$ that cover all $t \in [0, T)$. Without loss of generality, we can assume that

$$\pi_{e_1} < \pi_{f_m} < \pi_{e_2} < \pi_{f_1} < \dots < \pi_{e_m} < \pi_{f_{m-1}}$$

We can further assume that e_1 is the end event that is matched first in the greedy algorithm (of the events e_1, \dots, e_m). This event is matched with f_1 . However, it holds that $x_{e_1, f_m} < x_{e_1, f_1}$. Therefore, the greedy algorithm would select the arc (e_1, f_m) instead of (e_1, f_1) , resulting in a contradiction. It follows that M^{greedy} is optimal. \square

As a corollary, we find that the vehicle circulation scheduling problem can be solved in time $O(n^2 \log n)$, with $n = |E_{\text{end}}| = |E_{\text{start}}|$

Corollary 5. *The vehicle circulation scheduling problem can be solved in $O(n^2 \log n)$.*

Proof. Algorithm 1 is correct for any order in which the end events are matched. Hence, a valid approach is to first sort all turnaround activities in the entire network according to their process time, and then choose the first n (starting with the shortest process time) non-conflicting activities in this list as the selected turnarounds. Since there are at most n^2 activities in the graph, this results in a complexity of $O(n^2 \log n^2) = O(n^2 \log n)$. \square

As a final result, we provide a lower bound on the minimum turnaround time at a station based on a weighted sum of the durations of all turnaround activities.

Theorem 6. *Consider a station s with end events $E_{\text{end}}(s)$, turnaround arcs $A_{\text{turn}}(s)$ and activity durations x_a for all $a \in A_{\text{turn}}(s)$ and let $M(s) \subseteq A_{\text{turn}}$ denote any feasible perfect matching. It*

holds that

$$\sum_{a \in M(s)} x_a \geq \frac{1}{|E_{\text{end}}(s)|} \sum_{a \in A_{\text{turn}}(s)} x_a - T_s \frac{|E_{\text{end}}(s)| - 1}{2}. \quad (15)$$

Proof. We first consider the case where end events and start events alternate in the periodic pattern. Let $k = |E_{\text{end}}(s)|$. Then, we can partition the arcs $A_{\text{turn}}(s) = A_1(s) \cup A_2(s) \cup \dots \cup A_k(s)$, where $A_1(s)$ contains all activities from end events to the next start, $A_2(s)$ all the activities from end events to start events that 'skip' one start event and in general $A_i(s)$ contains the activities from end events to start events that 'skip' i departures (see Figure 4 for an illustration). It can be observed that $\sum_{a \in A_i(s)} x_a = \sum_{a \in A_1(s)} x_a + (i - 1)T_s$. Therefore, we have that

$$\sum_{a \in A_{\text{turn}}(s)} x_a = \sum_{i=1}^k \sum_{a \in A_i(s)} x_a = k \sum_{a \in A_1(s)} x_a + T_s \frac{k(k-1)}{2},$$

or equivalently,

$$\sum_{a \in A_1(s)} x_a = \frac{1}{k} \sum_{a \in A_{\text{turn}}(s)} x_a - T_s \frac{k-1}{2}.$$

The theorem follows from the observation that $A_1(s)$ minimizes the total turnaround time.

For the case where end and start events do not alternate, we first relate the theorem to the difference between the minimum turnaround time and the average turnaround time. Let $\mathcal{M}(s)$ denote the set of all perfect matchings at station s and let $x(M) = \sum_{a \in M} x_a$ denote the turnaround time of matching M . Since the turnaround graph is a complete bipartite graph, there are $|E_{\text{end}}(s)|!$ distinct perfect matchings and every arc is contained in $|E_{\text{end}}(s) - 1|!$ matchings, such that

$$\begin{aligned} \frac{1}{|\mathcal{M}(s)|} \sum_{M \in \mathcal{M}(s)} x(M) &= \frac{1}{|E_{\text{end}}(s)|!} \sum_{a \in A_{\text{turn}}(s)} \sum_{M \in \mathcal{M}(s): a \in M} x_a \\ &= \frac{1}{|E_{\text{end}}(s)|!} \sum_{a \in A_{\text{turn}}(s)} |E_{\text{end}}(s) - 1|! x_a \\ &= \frac{1}{|E_{\text{end}}(s)|} \sum_{a \in A_{\text{turn}}(s)} x_a. \end{aligned}$$

In other words, the theorem gives a lower bound on the turnaround of any matching based on the average turnaround time over all possible matchings. Hence, what remains to prove is that the difference between the minimum turnaround time and the average turnaround time is maximized when end events and start events alternate.

Let M^{opt} denote the matching attaining the minimum turnaround time and let M' denote any other matching. Let $\pi_e \in [0, T_s)$ denote the event time of event e . As ends and starts do not alternate, there are end events e_1 and e_2 and start events f_1 and f_2 such that $(e_1, f_1) \in M^{\text{opt}}$, $(e_2, f_2) \in M^{\text{opt}}$ and

$$0 = \pi_{e_1} < \pi_{e_2} < \pi_{f_1} < \pi_{f_2}.$$

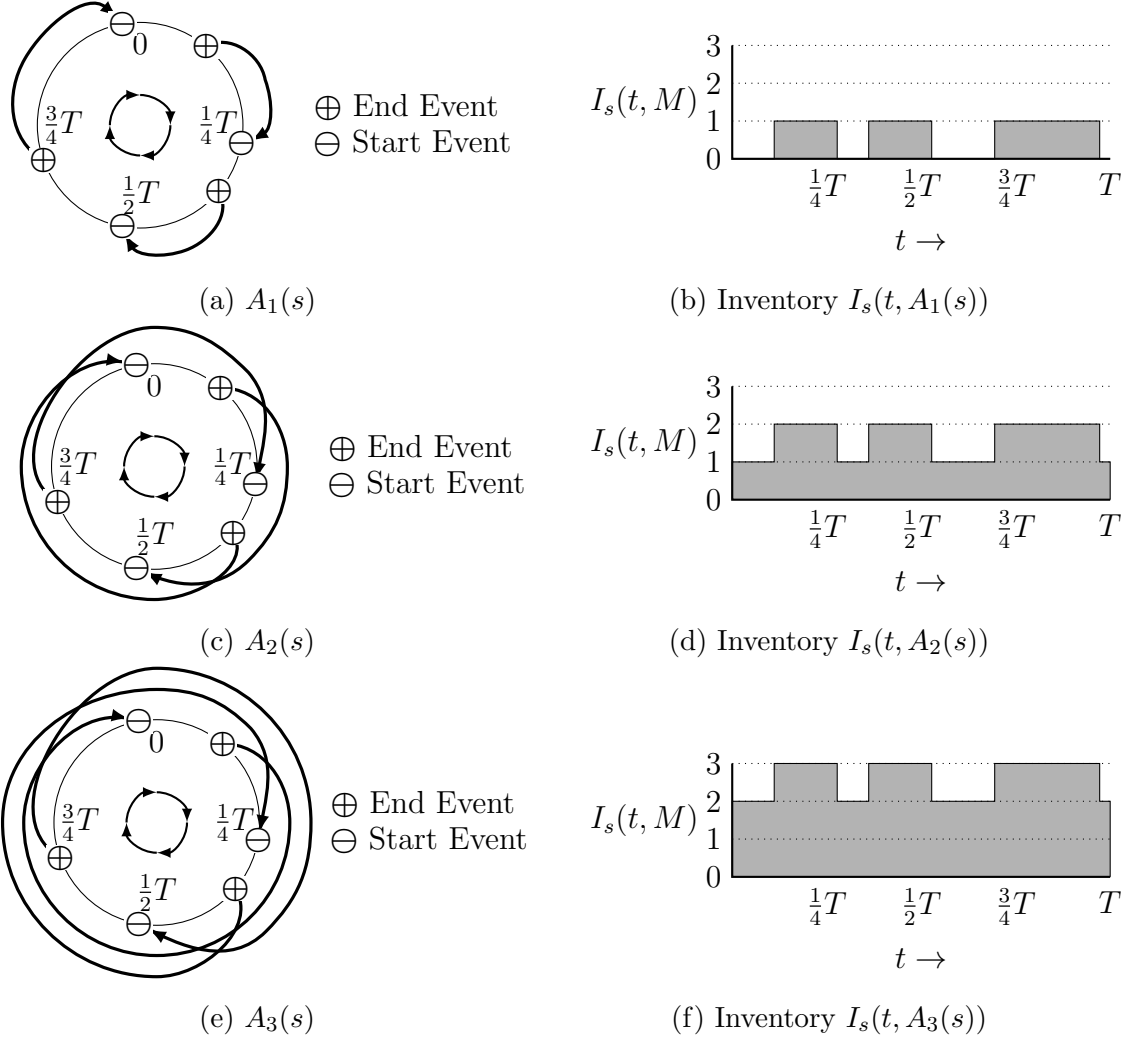


Figure 4: Illustration of the decomposition of all turnaround arcs into $A_1(s)$, $A_2(s)$ and $A_3(s)$, along with the corresponding inventory functions.

Now, consider shifting f_1 with δ such that $\pi'_{f_1} = \pi_{f_1} - \delta$ and $\pi_{e_1} < \pi'_{f_1} < \pi_{e_2}$, i.e. after the shift the end events e_1 , e_2 and start events f_1 , f_2 do alternate. Let $x'(M)$ denote the turnaround time of matching M after shifting f_1 . Clearly, $x'(M^{\text{opt}}) = x(M^{\text{opt}}) - \delta$. Let e_{alt} denote the end event that turns on f_1 in matching M' . Then,

$$x'(M') = \begin{cases} x(M') + T - \delta & \text{if } \pi_{f_1} - \delta < \pi_{e_{\text{alt}}} < \pi_{f_1} \\ x(M) - \delta & \text{else.} \end{cases}$$

It follows that $x'(M') - x'(M^{\text{opt}}) \geq x(M') - x(M^{\text{opt}})$. Therefore,

$$\frac{1}{|\mathcal{M}(s)|} \sum_{M' \in \mathcal{M}(s)} x'(M') - x'(M^{\text{opt}}) \geq \frac{1}{|\mathcal{M}(s)|} \sum_{M \in \mathcal{M}(s)} x(M') - x(M^{\text{opt}}).$$

This shows that by disentangling non-alternating end events and start events, the difference between

the minimum turnaround time and the average turnaround time cannot decrease. Therefore, the difference is maximized when all ends and starts alternate. \square

Note that for stations with only a single terminating line, Equation (15) directly gives the minimum turnaround time, since the end events and start events of one line always alternate. Hence, for such stations, the minimum turnaround time can be computed without explicitly solving a matching problem.

4 Mathematical Formulation for the VC-PESP

We are now ready to present a mathematical formulation for the integrated timetabling and vehicle circulation scheduling problem. To cover the timetabling part, we introduce the same decision variables as in the cycle periodicity formulation of the PESP: continuous tension variables x_a for all $a \in A$ and integer valued periodicity variables q_C for all $C \in \mathcal{B}$, where \mathcal{B} is an integral cycle basis. For the vehicle circulation scheduling part, we introduce binary matching variables y_a which are equal to 1 if turnaround activity $a \in A_{\text{turn}}$ is selected in the matching and 0 otherwise. Furthermore, we introduce an integer variable n representing the total number of vehicles required to operate all circulations. The VC-PESP can then be formulated as the following bi-objective problem:

$$\min \sum_{a \in A} \lambda_a x_a \quad (16)$$

$$\min n \quad (17)$$

$$\text{s.t. } l_a \leq x_a \leq u_a \quad \forall a \in A, \quad (18)$$

$$\sum_{a \in C^+} x_a - \sum_{a \in C^-} x_a = q_C T \quad \forall C \in \mathcal{B}, \quad (19)$$

$$a_C \leq q_C \leq b_C \quad \forall C \in \mathcal{B}, \quad (20)$$

$$\sum_{a \in \delta^+(e)} y_a = 1 \quad \forall e \in E_{\text{end}}, \quad (21)$$

$$\sum_{a \in \delta^-(e)} y_a = 1 \quad \forall e \in E_{\text{start}}, \quad (22)$$

$$n \geq \frac{1}{T} \left(\sum_{a \in A_{\text{veh}}} x_a + \sum_{a \in A_{\text{turn}}} x_a y_a \right) \quad (23)$$

$$x_a \in \mathbb{R} \quad \forall a \in A, \quad (24)$$

$$q_C \in \mathbb{Z} \quad \forall C \in \mathcal{B}, \quad (25)$$

$$y_a \in \{0, 1\} \quad \forall a \in A_{\text{turn}}, \quad (26)$$

$$n \in \mathbb{Z}. \quad (27)$$

The objectives are to minimize the average (perceived) travel times of all passengers (which can be achieved by setting appropriate weights λ_a) and the operator costs (captured by the number of required vehicles to operate the timetable). Constraints (18)-(20) cover the timetabling part of the formulation. Constraints (21) and (22) handle the matching part of the formulation, making sure that every trip has exactly one successor and one predecessor. Constraint (23) provides the link between the process time variables, the matching variables and the vehicle variable. The right hand side represents the total driving, dwelling and turnaround time of the chosen timetable in combination with the selected turnaround, divided by the period length. In an integer solution, the sum of tensions of each cycle will necessarily be an integer multiple of the period, such that the division by the period gives an integer result. The remainder of the constraints state the domains of the variables.

Formulation (16)-(27) is nonlinear, as constraint (23) contains the product of binary and continuous variables. However, it can be linearized by defining auxiliary continuous variables w_e for all $e \in E_{\text{end}}$, representing the turnaround time between end event e and the start event it connects to, and replacing constraint (23) by

$$n \geq \frac{1}{T} \left(\sum_{a \in A_{\text{veh}}} x_a + \sum_{e \in E_{\text{end}}} w_e \right). \quad (28)$$

To ensure that the w -variables attain the correct values, we add the following constraints:

$$w_e \geq x_a - M_1(1 - y_a) \quad \forall e \in E_{\text{end}} \quad \forall a \in \delta^+(e), \quad (29)$$

$$w_e \leq x_a + M_1(1 - y_a) \quad \forall e \in E_{\text{end}} \quad \forall a \in \delta^+(e), \quad (30)$$

$$w_e \geq 0 \quad \forall e \in E_{\text{end}}. \quad (31)$$

Note that it suffices to set $M_1 = T$, since x_a is smaller than T for turnaround arcs. The resulting formulation is the one studied by Nührenberg (2016), except that we include both the average travel time and the number of vehicles as objectives, instead of only the number of vehicles.

5 Complexity

As the PESP is NP-complete, it follows that the VC-PESP is NP-complete as well. However, it turns out that even if all timetabling complexity is removed from the problem (such that any timetable is feasible), it is still a hard problem to decide whether there exists a timetable that can be operated with a certain number of vehicles.

Specifically, let us define the SIMPLE-VC-PESP as follows. Consider a line plan \mathcal{L} , where each $l = (s_1, s_2) \in \mathcal{L}$ has a minimum unidirectional trip time t_l and a frequency f_l . From this line plan, we derive an event-activity network with events for the departures and arrivals of the lines

at the terminals, trip activities linking the departure of a line at one terminal to the arrival of the line at the other terminal, synchronizing activities linking different events of the same line and turnaround activities linking arrivals with departures at the same terminal. The synchronizing activities are added for lines with frequencies larger than 1 and impose that the departures of a line with frequency f are separated by exactly T/f time units. In this problem, there are no safety or transfer activities and all drive and dwell activities are contracted in a single trip activity with a fixed duration t_l . This implies that the arrival times of lines directly follow from the departure times. Therefore, the SIMPLE-VC-PESP contains only two timetabling decision variables $\pi_l^{s_1}$ and $\pi_l^{s_2}$ per line $l = (s_1, s_2)$, representing the departure times at the terminals.

Definition 2 (SIMPLE-VC-PESP). Given a line plan \mathcal{L} specifying the lines and frequencies and given a maximum number of vehicles n , the Simple Vehicle Circulation Periodic Event Scheduling Problem is to determine departure times at the terminals $\pi_l^{s_1}$ and $\pi_l^{s_2}$ for all lines $l = (s_1, s_2) \in \mathcal{L}$ such that at most n vehicles are needed to operate the periodic schedule, or to conclude that no such schedule exists.

Note that in the SIMPLE-VC-PESP, all timetables $\boldsymbol{\pi}$ are feasible. Hence, the only remaining complexity lies in making sure that all arrivals have short turnarounds to departures, such that only a small number of vehicles is required.

Theorem 7. *The SIMPLE-VC-PESP is NP-complete.*

Proof. We show this by reduction from the strongly NP-hard 3-PARTITION problem (Garey & Johnson, 1979). Let $S = \{s_1, s_2, \dots, s_m\}$ be a set of integers such that $\sum_{i=1}^m s_i = \frac{m}{3}B$ and $\forall i$ it holds that $\frac{B}{4} < s_i < \frac{B}{2}$. A 3-partition instance is a YES-instance if it is possible to partition S into $k = \frac{m}{3}$ triplets S_1, S_2, \dots, S_k such that each triplet sums to B .

We reduce the 3-PARTITION instance to a star network where we have a central hub station $v_0 \in V$ and $m + 1$ outer stations $v_1, \dots, v_{m+1} \in V$. We put $T = k + \sum_{i=1}^m s_i$ time units. For every element $s_i \in S$, we create a line $l_i \in \mathcal{L} = \{v_0, v_i\}$ with frequency $f_i = 1$ and a trip time $t_i = \frac{1}{2}s_i$. Furthermore, we create a line $l_{m+1} = \{v_0, v_{m+1}\}$ with frequency $f_{m+1} = k$ and a trip time $t_{m+1} = \frac{1}{2k}$.

Note that the total driving time of the created SIMPLE-VC-PESP instance equals $\sum_{i=1}^{m+1} 2f_i t_i = k + \sum_{i=1}^m s_i = T$. Therefore, at least 1 vehicle is required to operate the line plan. We put $n = 1$. Furthermore, w.l.o.g. we can assume that the vehicles have a turnaround time of 0 time units at the outer stations (i.e. $\pi_{l_i}^{v_i} = \pi_{l_i}^{v_0} + t_i \bmod T$). Therefore, only the departure times at the central station need to be determined. We claim that the 3-PARTITION instance is a YES-instance if and only if the SIMPLE-VC-PESP instance is a YES-instance.

If the 3-PARTITION instance is a YES-instance, we can create a solution using only a single vehicle. Let S_1, \dots, S_k denote the triplets. We set $\pi_{l_{m+1}} = 0$, so line $m + 1$ has departures at

$0, \frac{T}{k}, \dots, \frac{(k-1)T}{k}$. For triplet $S_j = (s_p, s_q, s_r)$, we set $\pi_{l_p} = \frac{(j-1)T}{k} + 1$, $\pi_{l_q} = \pi_{l_p} + 2t_p$ and $\pi_{l_r} = \pi_{l_q} + 2t_q$. Then, it is possible to create a circulation that performs the first service of line $m + 1$, then the lines corresponding to S_1 , then the second service of line $m + 1$, then all lines corresponding to S_2 et cetera. As this circulation takes exactly one period, only a single vehicle is needed to operate the timetable, such that the SIMPLE-VC-PESP instance is a YES-instance.

If the SIMPLE-VC-PESP instance is a YES-instance, it must be that there is a single circulation covering all lines with a duration of T . W.l.o.g. we can set $\pi_{l_{m+1}} = 0$. We can therefore write this circulation as

$$l_{m+1}^1, \underbrace{l_1', \dots, l_p'}_{\mathcal{L}_1}, l_{m+1}^2, \underbrace{l_1'', \dots, l_q''}_{\mathcal{L}_2}, l_{m+1}^3, \dots, l_{m+1}^k, \underbrace{l_1'', \dots, l_r''}_{\mathcal{L}_k}$$

As the services of line $m + 1$ should be spread exactly $\frac{1}{k}T$ time units and the total circulation time equals T , it holds that

$$\sum_{l \in \mathcal{L}_1} t_l = \sum_{l \in \mathcal{L}_2} t_l = \dots = \sum_{l \in \mathcal{L}_k} t_l = B.$$

The corresponding elements in S partition S into k subsets with a sum of B . Moreover, since $\frac{B}{4} < s_i < \frac{B}{2}$, all these subsets must contain exactly three elements. We conclude that the 3-PARTITION instance also is a YES-instance. \square

6 A Stronger Formulation

The mathematical formulation of the VC-PESP given in Section 4 is investigated by Nührenberg (2016). However, the author finds that the formulation is weak and results in large optimality gaps even for medium-sized instances. This is likely caused by the big- M constraints (29-30) that are necessary to linearize the formulation. Furthermore, the formulation also allows for many symmetric solutions, as there can be multiple matchings that attain the same turnaround time.

In this section, we present several ways to enhance the formulation of the VC-PESP by exploiting the special matching structure of the vehicle circulation scheduling problem. First, we discuss how the number of binary matching variables can be reduced. Next, we give multiple valid inequalities that strengthen the linear relaxation. Finally, we propose symmetry-breaking constraints, that make sure only a single matching that attains the minimum turnaround time is feasible.

6.1 Computing the Matching in a Contracted Network

Even though the period of the entire timetable might be equal to T , often the period is smaller from a local perspective. We can use this property to reduce the number of matching variables necessary to compute the total turnaround time at a station. In particular, if the events at station s repeat every T_s time units, we say that the *local period* at station s is T_s . In case the local

period is strictly smaller than the global period, it turns out to be possible to reduce the number of matching variables at s with a factor of $\left(\frac{T}{T_s}\right)^2$.

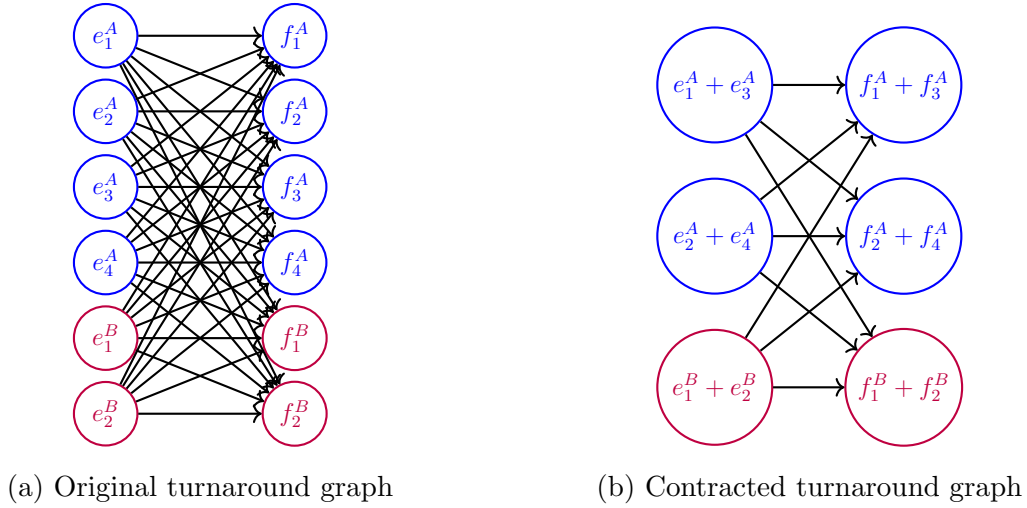


Figure 5: Contraction of the turnaround graph at a station with lines A and B with frequencies 4 and 2, respectively.

For instance, consider a station that serves as a terminal for two lines, line A that runs every 15 minutes and line B that runs every 30 minutes. Regardless of the global period, the local period T_s at this station is 30 minutes, since all events repeat at least every 30 minutes. This implies that the inventory function $I_s(t, M)$ also repeats every 30 minutes, for any matching. Hence, we can simply optimize the matching for all events occurring in 30 minutes and repeat this pattern in the subsequent periods. Figure 5 visualizes the contraction for this station for the case that the global period is 60 minutes. The events e_1^A and e_3^A are end events of line A that are separated with exactly 30 minutes, so they are combined in a single node in the contracted graph. The same holds for other pairs of events in the original graph. This results in a reduction of the number of turnaround arcs from 36 to 9.

To formally model the total turnaround time in the contracted graphs, we define for each line l that terminates at s and has global frequency f_l , the local line frequency $f_l^s := \frac{f_l T_s}{T}$. Next, we create a contracted graph in which every line is represented with f_l^s contracted end events and f_l^s contracted start events. Each of such contracted events corresponds to $\frac{T}{T_s}$ events in the regular network, in such a way that these events are all separated by a multiple of T_s minutes. Let $E_{\text{end}}^{\text{con}}(s)$ and $E_{\text{start}}^{\text{con}}(s)$ denote all end events and start events in the contracted graph at s . The set $A_{\text{turn}}^{\text{con}}(s)$ denotes the set of all turnaround arcs in the contracted networks, and contains an arc for each pair of contracted end and start events, such that the contracted network is complete. It follows that each contracted arc a corresponds to $\left(\frac{T}{T_s}\right)^2$ arcs in the regular network. We let $A_{\text{turn}}(a)$ denote the set of turnaround arcs in the regular network that correspond to the contracted arc a . Finally,

the set of all terminal stations is denoted as S_{term} .

To perform the matching in the contracted network, we introduce a binary matching variable γ_a and a continuous tension variable χ_a for all $a \in A_{\text{turn}}^{\text{con}}(s)$ and a turnaround time variable ω_e for all $e \in E_{\text{end}}^{\text{con}}(s)$. The following constraints then model the matching in this contracted networks:

$$\sum_{a \in \delta^+(e)} \gamma_a = 1 \quad \forall e \in E_{\text{end}}^{\text{con}}, \quad (32)$$

$$\sum_{a \in \delta^-(e)} \gamma_a = 1 \quad \forall e \in E_{\text{start}}^{\text{con}}, \quad (33)$$

$$\chi_a \geq x_{a'} - T + T_s \quad \forall s \in S_{\text{term}}, \quad \forall a \in A_{\text{turn}}^{\text{con}}(s), \quad \forall a' \in A_{\text{turn}}(a), \quad (34)$$

$$\chi_a \leq x_{a'} \quad \forall s \in S_{\text{term}}, \quad \forall a \in A_{\text{turn}}^{\text{con}}(s), \quad \forall a' \in A_{\text{turn}}(a), \quad (35)$$

$$\omega_e \geq \chi_a - M_2(1 - \gamma_a) \quad \forall e \in E_{\text{end}}^{\text{con}}, \quad \forall a \in \delta^+(e), \quad (36)$$

$$\omega_e \leq \chi_a + M_2(1 - \gamma_a) \quad \forall e \in E_{\text{end}}^{\text{con}}, \quad \forall a \in \delta^+(e), \quad (37)$$

$$n \geq \frac{1}{T} \left(\sum_{a \in A_{\text{veh}}} x_a + \sum_{s \in S_{\text{term}}} \frac{T}{T_s} \sum_{e \in E_{\text{end}}^{\text{con}}(s)} \omega_e \right) \quad (38)$$

$$\omega_e \geq 0 \quad \forall e \in E_{\text{end}}^{\text{con}}, \quad (39)$$

$$\chi_a \in \mathbb{R} \quad \forall a \in A_{\text{turn}}^{\text{con}}, \quad (40)$$

$$\gamma_a \in \{0, 1\} \quad \forall a \in A_{\text{turn}}^{\text{con}}. \quad (41)$$

Constraints (32)-(33) are the matching constraints for the contracted network. Constraints (34)-(35) link the tensions of the activities of the contracted network with those of the original network. Here, we use that the tensions of the arcs in the original network associated to an arc in the contracted network are $x, x + T_s, x + 2T_s, \dots, x + T - T_s$ for some $x \in [0, T_s)$. For example, if the local period is 15 minutes and the contracted arc a corresponds to tensions of 12, 27, 42 and 57 minutes in the original network, these constraints enforce that $\chi_a = \min\{12, 27, 42, 57\} = 12$.

Constraints (36)-(37) make sure that the turnaround time of every contracted event is equal to the tension of the arc that is selected to be in the matching. Since it holds that $\chi_a \in [0, T_s)$, it now suffices to set $M_2 = T_s$ instead of T , which also strengthens the formulation. Constraint (38) counts the number of vehicles, and now takes into account that turnaround times of contracted events should be counted according to their multiplicity.

For the sake of notation, in the remainder of this section, we omit the superscript "con" to denote the contracted turnaround graph and write x_a instead of χ_a , y_a instead of γ_a and w_e instead of ω_e .




Figure 6: Example illustrating the weak linear programming relaxation of the formulation.

6.2 Strengthening the LP-relaxation

A potential weakness of the original formulation is that its linear programming (LP) relaxation is very weak due to the linearization of the quadratic terms $x_a y_a$. Even if the periodicity variables q_C are integer for all $C \in \mathcal{B}$ (so the timetable is feasible) and only the integrality of the matching variables y_a is relaxed, the turnaround time variables w_e often still equal 0 for all end events. To illustrate this, consider the timetable displayed in Figure 6, with two vehicles turning at a certain station. We assume the event times are fixed and that the end events take place at time instants $-\varepsilon$ and ε , and the start events at $\frac{1}{2}T_s - \varepsilon$ and $\frac{1}{2}T_s + \varepsilon$, for some small ε . Clearly, the minimum turnaround time, which is actually obtained by both possible matchings, is T_s . However, if the matching variables y_a are equal to $\frac{1}{2}$, the turnaround time variables w_e can take the value ε , such that the total turnaround time is smaller than 2ε . In order to decrease the gap between the actual minimum turnaround time and the turnaround time in the relaxation, we develop three classes of valid inequalities.

6.2.1 Min-Mean Cuts

The first class of valid inequalities is based on the lower bound on the minimum turnaround time at a station given by Theorem 6. This theorem directly implies that the following inequalities are valid for the VC-PESP:

$$\sum_{e \in E_{\text{end}}(s)} w_e \geq \frac{1}{|E_{\text{end}}(s)|} \sum_{a \in A_{\text{turn}}(s)} x_a - T_s \frac{|E_{\text{end}}(s)| - 1}{2} \quad \forall s \in S_{\text{term}}. \quad (42)$$

We refer to these inequalities as the *min-mean* cuts, as Theorem 6 relates the matching with the minimum turnaround time with the mean turnaround time over all matchings. Recall that for stations with only a single terminating line, these cuts directly give the minimum turnaround time, since the end events and start events of one line always alternate, so for these stations it is no longer necessary to add any binary variables.

For the example in Figure 6, the min-mean cut states that $\sum_{e \in E_{\text{end}}(s)} w_e \geq \frac{1}{2}T_s$. Therefore, the original LP solution where both turnaround time variables have value ε is cut off.

6.2.2 Reference Matching Cuts

To improve the effectiveness of the min-mean cuts, we also include the following constraints. Let M_s denote a reference matching for station s and introduce the integer variable z_s for all $s \in S_{\text{term}}$. By Lemma 2, the difference in turnaround time between two matchings must be an integer multiple of T_s , hence the following equality is valid:

$$\sum_{e \in E_{\text{end}}(s)} w_e = \sum_{a \in M_s} x_a - T_s z_s \quad \forall s \in S_{\text{term}}, \quad (43)$$

$$z_s \in \mathbb{Z} \quad \forall s \in S_{\text{term}}. \quad (44)$$

To illustrate the usefulness of these cuts, consider again the example in Figure 6. In case the reference matching variable z_s is integer, this implies that $\sum_{e \in E_{\text{end}}(s)} w_e$ is either 0 or T_s . Next, the min-mean cut states that $\sum_{e \in E_{\text{end}}(s)} w_e \geq \frac{1}{2}T_s$. Hence, provided that z_s is integer, $\sum_{e \in E_{\text{end}}(s)} w_e$ must be equal to T . The z_s variables are therefore useful variables to branch on in a branch-and-bound context.

6.2.3 Headway Cuts

Headway cuts can be used if all departures at a station should be separated by a headway time $h_s > 0$. In this case, if a start event f is scheduled at time π_f and an end event e does *not* connect to f , it follows that e connects to a departure outside the interval $[\pi_f - h_s, \pi_f + h_s]_{T_s}$. This implies that the following inequalities are valid:

$$w_e \geq x_a + h_s - T_s \quad \forall s \in S, \quad \forall e \in E_{\text{end}}(s), \quad \forall a \in \delta^+(e), \quad (45)$$

In case the end event e connects to f , the inequality clearly holds as $h_s < T_s$. Otherwise, e connects to a different start event, which must be separated at least h_s time units from f . Taking periodicity into account, the right hand side of the inequality represents a lower bound on the turnaround time. For example, if $T_s = 60$ minutes, $h_s = 3$ minutes, $x_{e,f} = 59$ and e does not connect to f , it should hold that $w_e \geq 2$.

These inequalities cut off a relatively small portion of the feasible region, since they only have an effect if $x_{e,f} + h_s > T_s$ and h_s usually is relatively small. By preprocessing the event-activity network, the minimum separation between departures can usually be strengthened. Furthermore, these constraints can easily be incorporated by setting $M_2 = T_s - h_s$ in constraints (36)-(37).

6.3 Symmetry Breaking

In the VC-PESP, there can be a very large number of matchings that lead to the same number of required vehicles. This even holds for a fixed timetable. As an example, consider the timetable




Figure 7: Example illustrating the existence of many symmetric solutions.

at a station depicted in Figure 7. All 24 possible matchings at this station attain the same (and therefore minimum) turnaround time. To understand why this symmetry occurs, recall that the greedy algorithm described in Section 3 minimizes the total turnaround time, regardless of the order in which the events are matched. In other words, any order leads to an optimal matching with the same total turnaround time, but can lead to a different solution. Furthermore, different orders do lead to different solutions when the end events and start events do not alternate. As a result, branching on the matching variables y_a leads to many similar nodes being created, increasing the computational burden.

The symmetries can be broken if we specify an order in which the end event nodes are matched in the greedy algorithm a priori, and incorporate this order in the mathematical formulation. Let $E_{\text{end}}^{<e}$ ($E_{\text{end}}^{>e}$) denote the set of end events that come earlier (later) in the ordering than end event e and consider a turnaround activity (e, f) at some station s with tension $x_{e,f}$. Recall that h_s denotes the minimum separation between departures at station s . In an integer solution, the following cases can occur:

$$w_e \leq \begin{cases} x_{e,f} & \text{if } e \text{ connects to } f \text{ (so } y_{e,f} = 1), \\ x_{e,f} + M_3 & \text{if } e' \text{ connects to } f \text{ with } e' \in E_{\text{end}}^{<e}, \\ x_{e,f} - h_s & \text{if } e' \text{ connects to } f \text{ with } e' \in E_{\text{end}}^{>e}. \end{cases} \quad (46)$$

Here M_3 is a sufficiently large constant. Intuitively, if e connects to f , the turnaround time will be exactly $x_{e,f}$ by constraints (29). In the second case, some e' that comes earlier in the ordering than e turns on f , such that we cannot restrict w_e . In the final case, some e' that comes later in the ordering than e connects to f , which implies that the turnaround time w_e should be at most $x_{e,f} - h_s$, as e has a higher priority and will hence connect to a start event that departs earlier than f . The above stated relations can be captured in the following constraints:

$$w_e \leq x_{e,f} + M_3 \sum_{e' \in E_{\text{end}}^{<e}} y_{e',f} - h_s \sum_{e' \in E_{\text{end}}^{>e}} y_{e',f} \quad \forall s \in S, \quad \forall e \in E_{\text{end}}(s), \quad \forall f \in E_{\text{start}}(s). \quad (47)$$

These constraints ensure that for every possible timetable, only one matching that attains the minimum turnaround time is feasible, namely the one induced by the specified order. A sufficiently large value for M_3 is $T_s - h_s$.

7 Short Circulations

In practice, operators typically prefer shorter vehicle circulations over longer ones, since short circulations are more robust and easier to manage. If a disturbance occurs in a short circulation, it affects only a small part of the network. In contrast, if a disturbance occurs in a long circulation, the disturbance may propagate through all lines in the circulation, and therefore may have a large impact on the operations.

In this paper, we consider three types of vehicle circulations. *Fixed* circulations consist of services corresponding to a single line. Hence, if the entire network is operated using fixed circulations, there are dedicated vehicles for all lines. *Combined* circulations consist of services associated with two lines, so vehicles can alternate between lines. *Flexible* circulations consist of services belonging to three or more lines.

Fixed circulations are easily imposed within the VC-PESP by removing all turnaround activities between different lines. For the case where both fixed and combined circulations are allowed, we consider the slightly more restrictive case that every line can be combined with at most one other line. Then, let \mathcal{L} denote the set of lines, let $A_{\text{turn}}^{\text{comb}}$ denote the set of all turnaround activities that connect different lines and let $\text{lines}(a)$ denote the pair of lines a combined turnaround connects. If we introduce binary decision variables v_{lm} for all pairs of lines $l, m \in \mathcal{L}$ that are equal to one if the lines are combined and zero else, the following constraints model combined circulations:

$$\sum_{m \in \mathcal{L}, l \neq m} v_{lm} \leq 1 \quad \forall l \in \mathcal{L}, \quad (48)$$

$$y_a \leq v_{\text{lines}(a)} \quad \forall a \in A_{\text{turn}}^{\text{comb}}, \quad (49)$$

$$v_{lm} \in \{0, 1\} \quad \forall l \in \mathcal{L}, \quad \forall m \in \mathcal{L}, \quad l \neq m. \quad (50)$$

Constraints (48) ensure that every line is combined with at most one other line. Constraints (49) link the matching variables with the line combination variables.

Finally, note that when restrictions are imposed on the number of lines in vehicle circulations, it is not necessary to include all machinery developed in the previous section. If all circulations are fixed, the min-mean cuts (42) suffice for directly modeling the turnaround times, so none of the other techniques are needed. In case combined circulations are imposed, the matching problems at the different stations are linked, implying that they cannot be solved for each station independently. As such, the symmetry-breaking constraints (47) are no longer necessary (and potentially might even cut off the optimal solution).

8 Computational Experiments

In order to demonstrate the effectiveness of the proposed formulation, we present the results of a series of experiments. First, we focus on minimizing the number of vehicles to analyze the impact of the contraction techniques, valid inequalities and symmetry-breaking constraints. Thereafter, we consider the trade-off between the number of required vehicles and the quality of the timetable from the passengers' perspective (measured in terms of average perceived travel time).

8.1 Instances

The instances used for the experiments are derived from a part of the Dutch railway system. The considered network is visualized in Figure 8. In this network, two types of lines are operated; *regional* lines that have stops at all intermediate stations and *intercity* lines that only have stops at intercity stations. As different types of rolling stock is used to operate different line types, the mathematical model is (trivially) extended to account for vehicle types.

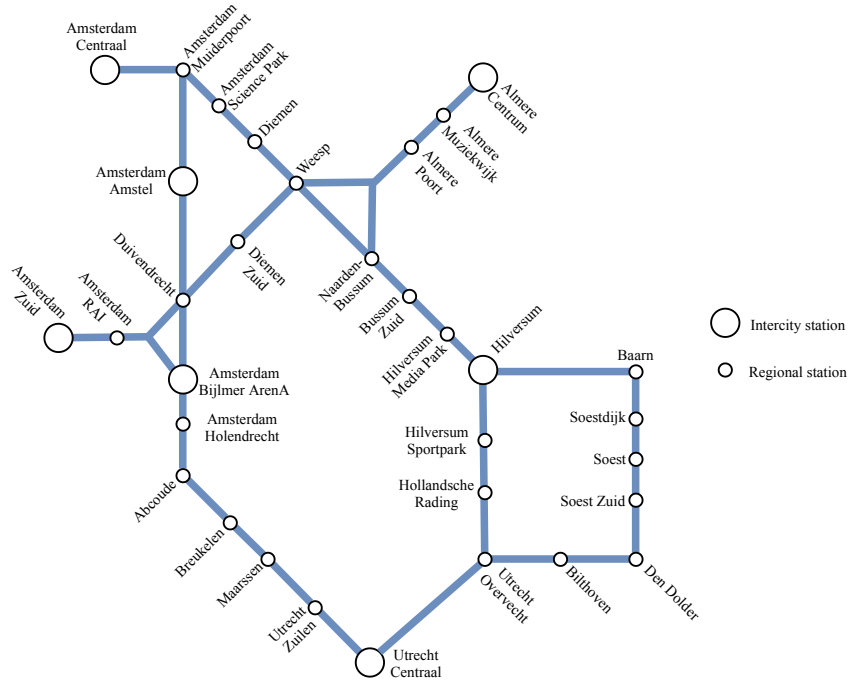


Figure 8: Part of the Dutch railway system used in the computational study.

To construct the instances, we create ten different line plans for the network from which we derive the event-activity networks. Safety constraints are generated with a headway of 3 minutes based on the railway infrastructure. Table 1 gives for all line plans the number of lines $|\mathcal{L}|$, the period in minutes T , the number of events $|E|$, the number of activities $|A|$, the number of turnaround activities $|A_{\text{turn}}|$ and a lower bound on the number of required vehicles \underline{n} (computed using minimum

driving and dwell times). In order to optimize the perceived travel time, we employ an origin destination matrix with passenger counts. The origin destination pairs are routed over the network based on the line plan, giving the number of passengers per activity. In this paper, we then define the perceived travel time of a route as the sum of expected waiting time for the first service, the in-vehicle time and an additional transfer penalty of 10 minutes plus twice the transfer time. The activity weights λ_a are set accordingly.

The experiments are performed on a machine with an Intel Xeon E5-2650 v2 2.60Ghz processor and 64 GB RAM. CPLEX 12.8.0 with default settings is used to solve the mixed-integer programs.

Table 1: Characteristics of the instances used in the computational study.

Instance	$ \mathcal{L} $	T	$ E $	$ A $	$ A_{\text{turn}} $	\underline{n}
ns1	10	60	522	732	90	15
ns2	15	60	684	1000	122	20
ns3	14	60	752	1105	138	21
ns4	9	60	908	1405	200	26
ns5	10	60	1030	1643	290	28
ns6	10	30	522	734	90	29
ns7	12	60	1122	1899	318	31
ns8	15	60	1158	1932	380	33
ns9	14	30	714	1028	124	40
ns10	11	20	550	747	78	47

8.2 Performance of the Formulations

To examine the effectiveness of the proposed techniques, we define the following formulations:

- *Orig*: the original formulation presented in Section 4.
- *Con*: the formulation with the contraction technique described in Section 6.1.
- *Val*: the formulation with the valid inequalities described in Section 6.2.
- *Sym*: the formulation with the symmetry-breaking constraints described in Section 6.3.

Next to these formulations, we also test combining the proposed techniques, which we indicate using plus signs, so e.g. *Con+Val* stands for the formulation with both the contraction technique and the valid inequalities.

As the improvements to the original formulations are developed to better model the number of vehicles, the performance of the formulations is tested by solving the VC-PESP with the goal

of minimizing the number of required vehicles, with flexible circulations. In the first experiment, we solve the VC-PESP using the different formulations. To reduce the impact of the seed used by CPLEX, we solve each instance with five different seeds and consider the average performance. For this experiment, we use a time limit of 15 minutes.

In Table 2, we present the average number of vehicles and the optimality gap obtained on the instances with the formulations *Orig*, *Con*, *Con+Val*, *Con+Sym* and *Con+Val+Sym*. The original formulation is clearly the worst performing formulation. In all instances except ns2 and ns9, the original formulation results in the largest number of vehicles and optimality gap. Introducing the contraction technique considerably improves the performance, most notably on ns5, ns7 and ns8. Adding the valid inequalities on top of the contraction does not yield further improvements. On the other hand, the symmetry-breaking constraints do reduce the number of vehicles and the average gap. However, it turns out that the valid inequalities are effective when added in combination with the symmetry-breaking constraints, as the *Con+Val+Sym* formulation attains the lowest average number of vehicles and the lowest average gap. Even more, on all instances this formulation finds the same or a lower number of vehicles as the other formulations. Compared to the original formulation, *Con+Val+Sym* reduces the average number of required vehicles by 0.8 (2.6 percent) and almost halves the average optimality gap. Therefore, we conclude that the *Con+Val+Sym* formulation is the most effective formulation for solving the VC-PESP, yielding considerable improvements over the original formulation.

Table 2: The average number of required vehicles and the average optimality gap obtained with the different formulations.

Instance	Orig		Con		Con+Val		Con+Sym		Con+Val+Sym	
	Vehicles	Gap (%)	Vehicles	Gap (%)	Vehicles	Gap (%)	Vehicles	Gap (%)	Vehicles	Gap (%)
ns1	16	1.2	16	0.0	16	0.0	16	0.0	16	1.2
ns2	20	0.0	20	0.0	20.2	1.0	20.2	1.0	20	0.0
ns3	22.8	7.8	22.4	6.1	21.6	2.6	21.6	2.6	21.4	1.8
ns4	26.4	1.5	26	0.0	26	0.0	26	0.0	26	0.0
ns5	31.2	10.2	30.2	7.3	30.2	7.3	30	6.7	30	6.7
ns6	30	0.0	30	0.0	30	0.0	30	0.0	30	0.0
ns7	35.2	11.8	33.4	7.2	33.2	6.6	33.2	6.6	33	6.1
ns8	38.4	14.0	37	10.8	37.2	11.3	37.2	11.3	36.8	10.3
ns9	41.4	3.4	40.8	1.9	42	4.7	40.6	1.4	40.2	0.5
ns10	48	0.0	48	0.0	48	0.0	48	0.0	48	0.0
Average	30.9	5.0	30.4	3.3	30.4	3.3	30.3	3.0	30.1	2.7

To gain more insight on the effect the contraction technique and the valid inequalities have on the tightness of the formulation, we perform a second experiment where we consider the *partial linear programming relaxation*, where all integer variables, except the periodicity variables, are relaxed. As feasible solutions to this partial relaxation correspond to feasible timetables, this

allows us to measure the discrepancy between the number of required vehicles as indicated by the relaxed model and the actual number of vehicles required to operate the timetable. Formally, we let n^{PLP} denote the number of vehicles in the optimal solution to the partial relaxation (which in general will be fractional) and let $n(\mathbf{x}^{PLP})$ the actual number of vehicles the found timetable requires. In a tight formulation, the gap between n^{PLP} and $n(\mathbf{x}^{PLP})$ should be small. In this experiment, if the relaxation is not solved to optimality within 2 minutes, the values of the best solution found so far are reported.

In Table 3, the results of solving the partial relaxation are presented for the formulations *Orig*, *Con*, *Val* and *Con+Val*. For instance ns1, the original formulation finds a solution with 14.34 vehicles. However, evaluating this solution gives that 21 vehicles are required to operate the corresponding timetable, such that the gap is quite large at 31.7 percent. The formulations with the improvements also find solutions with 14.34 vehicles, but these solutions require only 19, 18 and 18 vehicles, for *Con*, *Val* and *Con+Val* respectively. For all instances, formulations *Orig* and *Con* find the same minimum number of vehicles in the partial relaxation. With the contraction technique however, the gap between n^{PLP} and $n(\mathbf{x}^{PLP})$ and is 17.5 percent on average, compared to 19.7 percent with the original formulation. When the valid inequalities are included in the formulation, the number of vehicles in the partial relaxation increases in seven of the ten instances, on average from 28.11 to 28.47 for formulation *Val* and 28.50 for formulation *Con+Val*. Furthermore, the gap drastically decreases, to 12.9 and 12.4 percent, respectively. This shows that the contraction technique and especially the valid inequalities lead to a tighter formulation of the VC-PESP.

Table 3: Results of minimizing the number of vehicles with all integer variables relaxed, except the periodicity variables. The reported gap is the relative gap between n^{PLP} , the number of vehicles in the partial relaxation and $n(\mathbf{x}^{PLP})$, the actual number of vehicles required by the solution of the partial relaxation.

Instance	Orig			Con			Val			Con+Val		
	n^{PLP}	$n(\mathbf{x}^{PLP})$	Gap (%)	n^{PLP}	$n(\mathbf{x}^{PLP})$	Gap (%)	n^{PLP}	$n(\mathbf{x}^{PLP})$	Gap (%)	n^{PLP}	$n(\mathbf{x}^{PLP})$	Gap (%)
ns1	14.34	21	31.7	14.34	19	24.5	14.34	18	20.3	14.34	18	20.3
ns2	18.96	26	27.1	18.96	26	27.1	18.96	23	17.6	18.96	23	17.6
ns3	20.17	27	25.3	20.17	27	25.3	20.17	26	22.4	20.17	25	19.3
ns4	24.35	31	21.4	24.35	29	16.0	24.93	28	11.0	24.93	28	11.0
ns5	27.50	36	23.6	27.50	33	16.7	28.01	32	12.5	27.86	32	12.9
ns6	28.68	32	10.4	28.68	32	10.4	28.70	31	7.4	28.70	31	7.4
ns7	29.95	37	19.0	29.95	35	14.4	30.52	33	7.5	31.03	35	11.3
ns8	32.66	41	20.3	32.66	41	20.3	33.74	39	13.5	33.26	37	10.1
ns9	38.51	43	10.4	38.51	44	12.5	39.05	43	9.2	39.41	42	6.2
ns10	45.96	50	8.1	45.96	50	8.1	46.29	50	7.4	46.29	50	7.4
Average	28.11	34.4	19.7	28.11	33.6	17.5	28.47	32.3	12.9	28.50	32.1	12.4

8.3 Trade-off between the number of vehicles and the average travel time

To analyze the trade-off between the number of vehicles and the average travel time, we use the concept of a Pareto-efficient timetable, where it is impossible to reduce the number of vehicles without increasing the average travel time, and vice-versa. To compute the set of Pareto-efficient timetables, we first separately minimize the average travel time and the number of vehicles, resulting in an upper and lower bound on the minimum number of required vehicles. These first two problems are solved with a time limit of 30 minutes. Next, for every integer m between the lower bound and the upper bound, the VC-PESP is solved with the average travel time as the objective and the constraint that the number of vehicles is at most m . For these problems, the time limit is set to 15 minutes. The best performing formulation found in the previous section, *Con+Val+Sym*, is used in these experiments.

We compute the Pareto-efficient solutions for the cases with fixed circulations, combined circulations and flexible circulations. To speed up the computations, we use the obtained solutions with fixed circulations as starting solutions for the combined circulations, and the obtained solutions with combined circulations as starting solutions for the flexible circulations.

The obtained efficient solutions are visualized in Figures 9a-9j. The figures also depict the sequential solutions obtained by first optimizing the timetable to minimize the travel time and subsequently optimizing the vehicle circulations for the found timetable. It can be observed that sequential optimization leads to very inefficient timetables. In many cases, the number of vehicles of these solutions can be reduced without any increases in travel time. Furthermore, if an increase in travel time is required to reduce the number of vehicles, this increase is typically very limited. In general, it can be observed that reductions in the number of vehicles only require small increases in average travel time, but the required increases progressively become larger when the number of vehicles approaches its minimum.

A second observation from Figures 9 is that, as expected, the flexible solutions dominate the combined solutions, which in turn dominate the fixed solutions. Especially the difference between the combined solutions and the fixed solutions is considerable. The benefit of flexible circulations over combined circulations is less consistent over the instances, but still clearly noticeable for ns5, ns9 and ns10. For seven of the instances, flexible circulations do allow finding solutions with fewer vehicles than with combined circulations, but these timetables on the far end of the Pareto-curve are often much less attractive from the passengers' side.

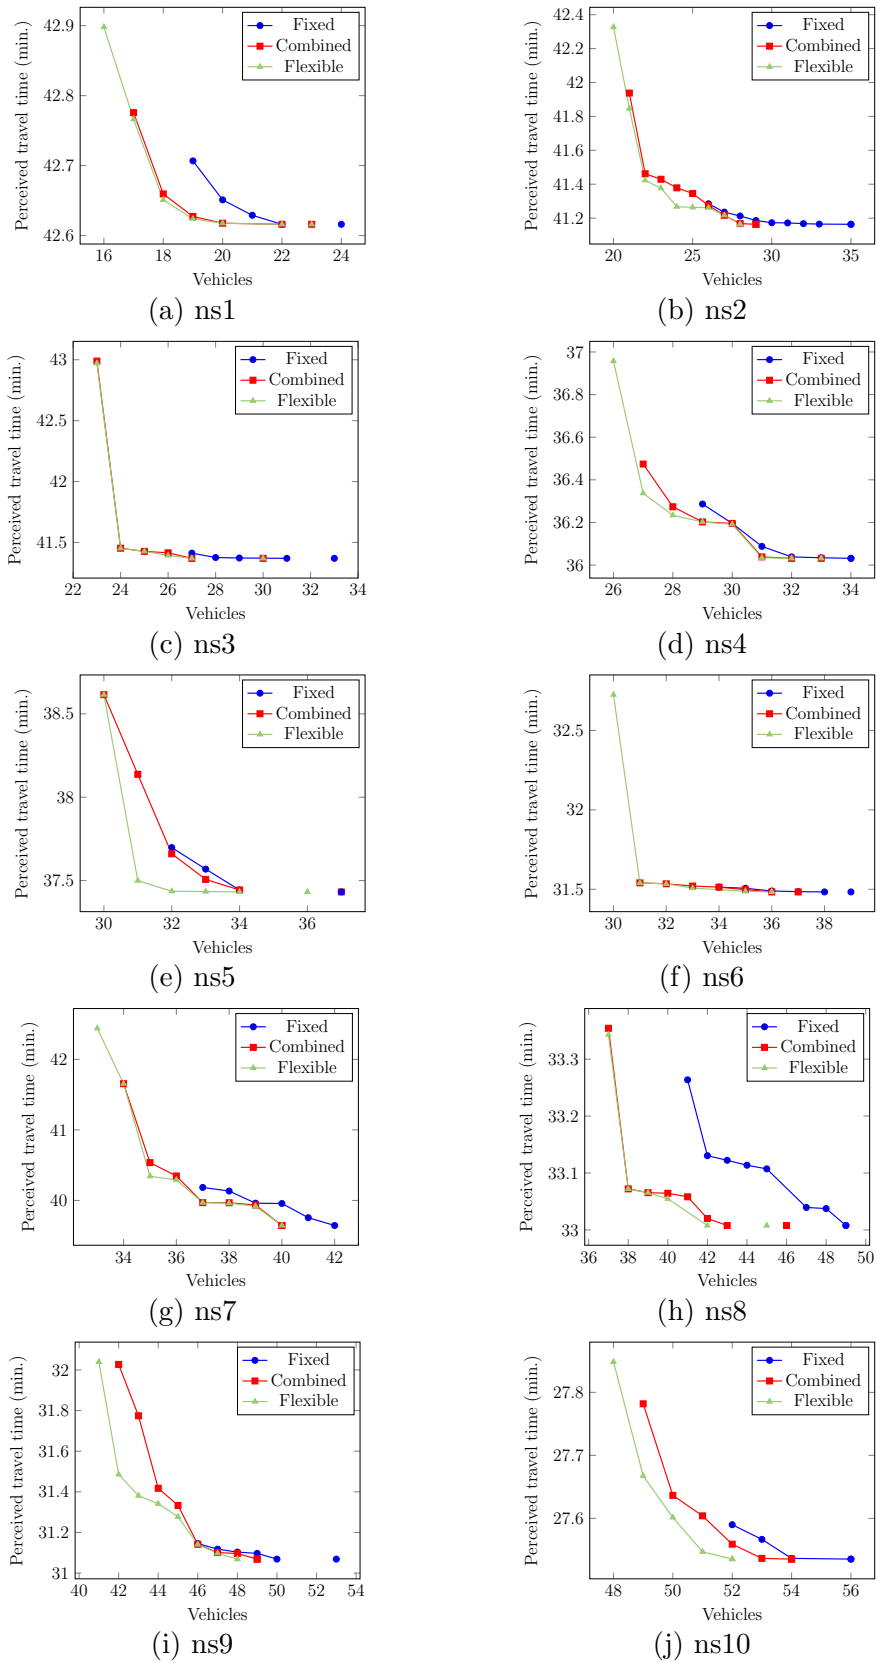


Figure 9: Pareto-efficient solutions and the sequential solutions, for fixed, combined and flexible circulations. In some cases, the sequential solutions are not Pareto-efficient, these points are disconnected from the curve.

To further analyze the benefit of integrating timetabling and circulation scheduling, Figures 10a-10c depict the increase in travel time and decrease in number of vehicles for all Pareto-efficient solutions, relative to the sequential solution. The key insight is that substantial decreases in the number of required vehicles can be achieved at the cost of very limited increases in travel time. Even more, for combined circulations it is possible to decrease the number of vehicles without any increases in travel time for six out of the ten instances. For fixed and flexible circulations, this is possible in five of the instances. If a 0.01 percent travel time increase is permitted, significant savings in vehicles can be obtained for the majority of the instances. This clearly shows that the timetable obtained by sequentially optimizing the travel time and the number of vehicles is strongly inferior to timetables obtained by jointly optimizing the travel time and the number of vehicles. All in all, our results indicate that the number of vehicles can typically be reduced considerably without doing any serious harm to the passengers' perspective.

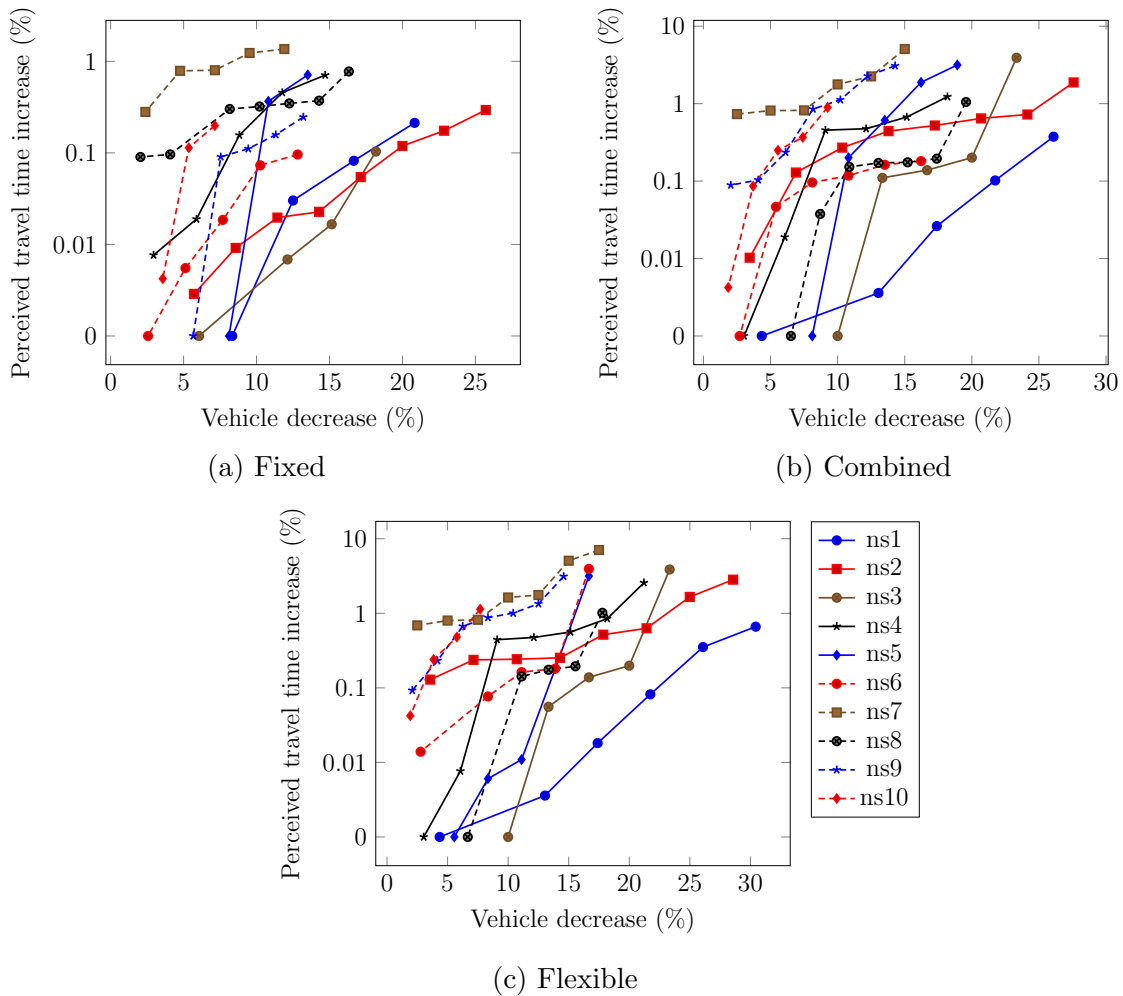


Figure 10: Increase in travel time (on a logarithmic scale) plotted against the decrease in the number of vehicles, both relative to the sequential solution.

9 Conclusion and Future Research

We presented a new periodic timetabling approach that allows to explicitly consider the trade-off between passengers' travel time and the number of required vehicles required to operate the timetable. We introduced contraction techniques, valid inequalities and symmetry-breaking constraints to cope with the complexities of this integrated problem.

Computational results based on the Dutch railway network illustrate the value of our approach. The improved formulation performs better when it comes to finding timetables requiring few vehicles compared to an existing formulation, approximately halving the average optimality gap. Furthermore, compared to sequential optimization of the travel time and the number of vehicles, we are able to find timetables requiring considerably fewer vehicles with only minimal increases of travel times. As such, our approach gives operators the opportunity to save costs without strongly decreasing the level of service, or conversely, increasing the level of service by introducing new lines or raising frequencies without strongly increasing the costs.

For further research, the presented approach could be embedded in a scheme for integrated public transport planning, where the goal is to find a line plan, timetable and vehicle schedule that are both inexpensive to operate and attractive to passengers. Next to that, our methods could be applied in different periodic scheduling problems, such as aircraft scheduling or job-shop scheduling. Finally, it would be interesting to consider a further integration between timetabling and vehicle scheduling in railway contexts, where vehicles can be coupled to meet passenger demand.

Acknowledgements This research was funded by NWO, the Netherlands Organisation for Scientific Research, as part of the research programme Complexity in Transport & Logistics (project number 439.16.111).

References

- Borndörfer, R., Hoppmann, H., Karbstein, M., & Löbel, F. (2016). The modulo network simplex with integrated passenger routing. In *Operations Research Proceedings 2016* (pp. 637–644). Springer.
- Borndörfer, R., Karbstein, M., Liebchen, C., & Lindner, N. (2018). A simple way to compute the number of vehicles that are required to operate a periodic timetable. In *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2018)*.
- Bunte, S., & Kliewer, N. (2009). An overview on vehicle scheduling models. *Public Transport*, 1(4), 299–317.

- Burggraeve, S., Bull, S. H., Vansteenwegen, P., & Lusby, R. M. (2017). Integrating robust timetabling in line plan optimization for railway systems. *Transportation Research Part C: Emerging Technologies*, 77, 134–160.
- Desfontaines, L., & Desaulniers, G. (2018). Multiple depot vehicle scheduling with controlled trip shifting. *Transportation Research Part B: Methodological*, 113, 34–53.
- Fonseca, J. P., Van der Hurk, E., Roberti, R., & Larsen, A. (2018). A matheuristic for transfer synchronization through integrated timetabling and vehicle scheduling. *Transportation Research Part B: Methodological*, 109, 128–149.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability. a guide to the theory of NP-completeness. a series of books in the mathematical sciences*. WH Freeman and Co., San Francisco, Calif.
- Gattermann, P., Großmann, P., Nachtigall, K., & Schöbel, A. (2016). Integrating passengers’ routes in periodic timetabling: A SAT approach. In *16th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2016)*.
- Goerigk, M., & Liebchen, C. (2017). An improved algorithm for the periodic timetabling problem. In *17th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems* (p. 1).
- Goerigk, M., & Schöbel, A. (2013). Improving the modulo simplex algorithm for large-scale periodic timetabling. *Computers & Operations Research*, 40(5), 1363–1370.
- Großmann, P., Hölldobler, S., Manthey, N., Nachtigall, K., Opitz, J., & Steinke, P. (2012). Solving periodic event scheduling problems with sat. In *International conference on industrial, engineering and other applications of applied intelligent systems* (pp. 166–175).
- Ibarra-Rojas, O. J., Giesen, R., & Rios-Solis, Y. A. (2014). An integrated approach for timetabling and vehicle scheduling problems to analyze the trade-off between level of service and operating costs of transit networks. *Transportation Research Part B: Methodological*, 70, 35–46.
- Kaspi, M., & Raviv, T. (2013). Service-oriented line planning and timetabling for passenger trains. *Transportation Science*, 47(3), 295–311.
- Kroon, L., Huisman, D., Abbink, E., Fioole, P.-J., Fischetti, M., Maróti, G., . . . Ybema, R. (2009). The new Dutch timetable: The OR revolution. *Interfaces*, 39(1), 6–17.
- Kroon, L., Peeters, L., Wagenaar, J., & Zuidwijk, R. (2013). Flexible connections in PESP models for cyclic passenger railway timetabling. *Transportation Science*, 48(1), 136–154.
- Liebchen, C. (2008). The first optimized railway timetable in practice. *Transportation Science*, 42(4), 420–435.
- Liebchen, C., & Möhring, R. H. (2007). The modeling power of the Periodic Event Scheduling Problem: railway timetables and beyond. In *Algorithmic methods for railway optimization* (pp. 3–40). Springer.
- Liebchen, C., & Peeters, L. (2009). Integral cycle bases for cyclic timetabling. *Discrete Optimiza-*

- tion, 6(1), 98–109.
- Maróti, G. (2017). A branch-and-bound approach for robust railway timetabling. *Public Transport*, 9(1-2), 73–94.
- Nachtigall, K. (1994). *A branch and cut approach for periodic network programming*. Inst. für Mathematik.
- Nachtigall, K., & Opitz, J. (2008). Solving periodic timetable optimisation problems by modulo simplex calculations. In *8th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS 2008)*.
- Nührenberg, G. (2016). *Integrated timetabling and rolling stock scheduling* (Unpublished master's thesis). ETH-Zürich.
- Odijk, M. A. (1996). A constraint generation algorithm for the construction of periodic railway timetables. *Transportation Research Part B: Methodological*, 30(6), 455–464.
- Odijk, M. A. (1997). *Railway timetable generation*. (PhD. Thesis)
- Orlin, J. B. (1982). Minimizing the number of vehicles to meet a fixed periodic schedule: An application of periodic posets. *Operations Research*, 30(4), 760–776.
- Pätzold, J., Schiewe, A., Schiewe, P., & Schöbel, A. (2017). Look-ahead approaches for integrated planning in public transportation. In *17th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2017)*.
- Peeters, L. (2003). *Cyclic railway timetable optimization*. (PhD. Thesis. EPS-2003-022-LIS)
- Polinder, G.-J., Breugem, T., Dollevoet, T., & Maróti, G. (2019). *An adjustable robust optimization approach for periodic timetabling* (Tech. Rep.).
- Schmid, V., & Ehmke, J. F. (2015). Integrated timetabling and vehicle scheduling with balanced departure times. *OR Spectrum*, 37(4), 903–928.
- Schöbel, A. (2017). An eigenmodel for iterative line planning, timetabling and vehicle scheduling in public transportation. *Transportation Research Part C: Emerging Technologies*, 74, 348–365.
- Schrijver, A., & Steenbeek, A. (1993). *Spoorweg Dienstregelingsontwikkeling (Railway timetable development)*. Retrieved 2019-03-06, from <https://homepages.cwi.nl/~lex/sdo.pdf>
- Serafini, P., & Ukovich, W. (1989). A mathematical model for periodic scheduling problems. *SIAM Journal on Discrete Mathematics*, 2(4), 550–581.
- Van Lieshout, R. N., & Bouman, P. C. (2019). Vehicle scheduling based on a line plan. In *19th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2019)*.