Integrated Optimization of Timetabling and Electric Vehicle Scheduling: A Case Study of Aachen, Germany

Vladimir Stadnichuk[1], Jenny Segschneider[2], Arie Koster[3], Grit Walther[4]

[1,4] School of Business and Economics, RWTH Aachen University
[2, 3] Research Group on Discrete Optimization, RWTH Aachen University

vladimir.stadnichuk@om.rwth-aachen.de [1]
segschneider@math2.rwth-aachen.de [2]
koster@math2.rwth-aachen.de [3]
walther@om.rwth-aachen.de [4]

September 27, 2024

1 Introduction

Many cities globally transition their bus fleet from traditional gasoline buses to a fully electric fleet to meet sustainability goals (Sustainable Bus, 2024). While electric buses induce zero local emissions, they face significant challenges such as longer charging times and a more limited driving range compared to their gasoline counterparts. These limitations make the transition highly complex, requiring either a substantially increased number of electric buses, which is costly, or intricate planning of the charging schedule to ensure sufficient battery capacity for all trips.

In this work, we focus on transit network planning that aims to transform gasoline bus fleets to fully electric ones in cities that employ periodic schedules. Periodic schedules ensure buses run at consistent intervals throughout the day, providing predictable and reliable service for passengers. This regularity enhances passenger satisfaction by offering a dependable and easily understandable timetable, reducing wait times and improving overall service quality. Such schedules are particularly popular in many cities across Germany, Netherlands, and Switzerland, where they have been successfully implemented in the public transportation systems.

Transit network planning is typically divided into five sequential steps: line planning, frequency setting, timetabling, vehicle scheduling, and crew scheduling/rostering (Ceder & Wilson, 1986). For periodic schedules, the line planning and frequency setting steps involve defining the layouts and frequencies of the lines, while the final periodic timetable is completed in the timetabling step by setting the departure times of the individual bus lines. In the vehicle scheduling step, vehicles are assigned to the trips specified in the timetable. This step includes replacing gasoline buses with electric buses and determining the charging schedule. This planning problem is referred to as the Electric Vehicle Scheduling Problem (EVSP) and studied excessively in scientific literature. In the final step, bus drivers are assigned to the vehicle and trip pairs during crew scheduling/rostering.

In the context of electrification, it would be advantageous to redesign the entire planning process and integrate the planning steps. But in Germany, the first two steps are heavily regulated (cf. SWA, 2019; The Federal Office of Justice Germany, 2024). Consequently, bus operators are generally reluctant to alter line layouts or frequencies. Personnel planning also involves complex challenges with numerous regulatory constraints. Thus, personnel costs are typically estimated at a flat rate, added to the vehicle and operational costs in the earlier planning steps. However, bus operators might consider minor adjustments in the timetabling step, such as to line departure times, if the benefits are substantial.

Against this background, our research aims to evaluate the potential benefits of an integrated approach that allows for minor adjustments to an existing periodic timetable before solving the underlying EVSP, which models the transition from a gasoline to an electric bus fleet. We begin by reviewing the scientific literature on timetabling and vehicle scheduling and identifying the research gap in Section 2. We then formally define the integrated planning problem and present an iterative algorithmic framework in Section 3. Our method's advantage lies in the use of individual subroutines for the timetabling and EVSP steps, enhancing flexibility and adaptability to various practical requirements. The subroutines for the timetabling and EVSP steps are detailed in Sections 4 and 5, respectively. Finally, we address our research question in Section 6, where we conduct a computational study based on realistic data from Aachen. Concluding remarks are given in Section 7.

2 Literature Review

In this section, we recap the scientific literature on the timetabling and vehicle scheduling planning steps. For periodic schedules, the timetabling step consists of finding departure times for each bus line while ensuring transfers between lines, and the scheduling step assigns trips to buses such that each trip can be served by the assigned bus. While these planning steps are usually solved sequentially, we will also highlight some recent developments that attempt to solve both planning problems in an integrated manner.

The planning problem of finding departure times for all lines with given frequencies that satisfy all transfer constraints is often modeled as a Periodic Event Scheduling Problem (PESP). Introduced by Serafini and Ukovich (1989), the PESP is often solved using one of the following two MIP models derived from an event-auxiliary graph. The initial MIP model proposed by Serafini and Ukovich (1989) directly plans the event times. However, as this formulation has a weak LP relaxation, an alternative Cycle Basis formulation was developed (Liebchen & Möhring, 2007; Liebchen & Peeters, 2009). This new formulation plans only the times between the events, taking advantage of the periodic structure. In this way, the resulting formulation aggregates multiple symmetric solutions with different event times that are just shifted by a fixed time interval, resulting in an overall stronger formulation. There also exist multiple alternative solution approaches to the PESP, including constraint programming (Gattermann et al., 2016; Großmann et al., 2012), or modulo-simplex heuristics (Borndörfer et al., 2017; Goerigk & Schöbel, 2013).

Regarding vehicle scheduling, the EVSP originated from the classical Vehicle Scheduling Problem (VSP), which, given a fixed timetable, finds a feasible vehicle rotation, i.e., an assignment of trips to gasoline buses. While the VSP with a single depot and bus type is solvable in polynomial time, there exist multiple extensions of the VSP, considering multiple bus types and depots. Although these are NP-hard problems, practical efficient exact MIP models (and heuristics) exist. For a comprehensive overview of the topic, we refer to Bunte and Kliewer (2009).

The EVSP extends the VSP to electric buses with their inherent limited range and charging requirements. Due to the capacity constraint, the EVSP is directly NP-hard even for a single bus type and depot (Sassi & Oulamara, 2016). From a practical perspective, multiple additional assumptions are often integrated into the EVSP, such as multiple bus types and depots, non-linear (re-)charging curves (Olsen & Kliewer, 2020), detailed scheduling of the charging process itself (Abdelwahed et al., 2020), or uncertainty (An, 2020; Li et al., 2021; Liu & Song, 2017; Shen et al., 2023).

Janovec and Koháni (2019) formulate and solve the EVSP with a compact MIP model. Adler and Mirchandani (2017) and Niekerk et al. (2017) propose a branch-and-price approach to solve the EVSP for both linear and concave charging functions. In their models, they use a set covering formulation to select a subset of feasible vehicle rotations to serve all trips. The vehicle rotations are generated by a dynamically solved subproblem that can cope with various charging assumptions, making their approach flexible. As their exact approach struggles to solve even medium-sized instances, they extend it with heuristics to tackle larger instances. Li (2014) also develops a branch-and-price algorithm, but assumes fast recharging or battery swapping for the recharging process. Recently, De Vos et al. (2024) formulated a new branch-and-price model that considers the limited capacity of charging stations and extend it with heuristics to solve larger instances.

Olsen et al. (2020) focus on repair heuristics that first compute an optimal solution for the VSP, disregarding battery capacities, and then repair the vehicle rotations to be battery-feasible. They conclude that the percentage of feasible vehicle rotations generated by their approach meets the requirements for an initial implementation of electric buses in practice. Other MIP models and heuristics extend the EVSP by the strategic decision of locating charging stations (Berthold et al., 2017; Liu et al., 2021; Olsen & Kliewer, 2022; Rogge et al., 2018).

Focusing on integrated approaches, the combined problem of designing a non-periodic timetable and VSP has already been extensively researched in the literature (Carosi et al., 2019; Desfontaines & Desaulniers, 2018; Ibarra-Rojas et al., 2014; Laporte et al., 2017). Van Lieshout and Bouman (2018) study the problem from a theoretical side and provide complexity results for the problem of estimating the cost of a vehicle schedule for a given timetable. Xu et al. (2023) extend these works to an electric bus fleet. They develop a MIP model and extend it with a Lagrangian relaxation heuristic to solve instances of larger size. However, the assumption of a periodic timetable brings challenges distinct from the non-periodic schedules assumed there. While in non-periodic timetables, the departure time of individual trips can be adjusted to better fit the underlying EVSP, only the departure times of whole lines, but no individual trips, can be adjusted in periodic timetables.

To the best of our knowledge, only Van Lieshout (2021) combine periodic timetabling with the VSP. They study the structural properties of the solution space and develop an efficient MIP formulation for the integrated problem. However, their model relies on specific structural properties of the solution space resulting from the VSP and cannot easily be extended to the EVSP. Furthermore, the discussed literature models focus on planning the timetable and vehicle rotation from scratch, while we assume that an existing periodic timetable is already given, and only slight modifications are allowed before the EVSP is solved.

Therefore, our contribution to the scientific literature is the development of an iterative approach that solves the integrated periodic timetabling and EVSP constrained to an existing periodic timetable that may only be slightly modified. The presented approach uses modular subroutines for solving both the PESP and EVSP, making our model very flexible and easy to implement. The modular parts can be replaced to fit current needs and therefore be easily adjusted to consider new constraints.

While we allow for arbitrary EVSP as a subroutine, we also make a direct contribution to the EVSP literature. We link the EVSP with the Bin Packing with Conflicts (BPwC), a theoretical optimization problem closely related to the basic EVSP without recharging (Huang et al., 2023; Jansen & Öhring, 1997). We not only show that the existing approximation ratios from BPwC can be directly transferred to a specified EVSP variant, but also extend these results by presenting new approximation ratios for the EVSP with linear recharging. To the best of our knowledge, we are the first to provide approximation bounds for a variant of the EVSP with recharging. Additionally, our theoretical results offer general insights into the quality of repair heuristics that first solve vehicle scheduling without battery restrictions to optimality and then repair the solution to make it battery-feasible.

We evaluate the proposed model and developed EVSP algorithms within a computational study. We provide insights into the computational performance of the proposed algorithms and address our main research question by measuring the potential gains from the option of slightly adjusting the timetable to better fit the new electric bus fleet schedule using realistic data from the Aachen bus network.

3 Problem Setting and Iterative Solution Approach

We consider the following problem setting with an given periodic timetable $\mathcal{S}$ with a cycle time T for the current gasoline bus fleet. The timetable consists of a set of bus lines $\mathcal{L}$. Each line $l \in \mathcal{L}$

operates in a time interval $[d_l, e_l]$, i.e., the first trip starts at d_l and the last trip at e_l. The line is periodically operated with frequency f_l, resulting in a headway $h_l = T/f_l$.

We now want to transform our bus fleet to a fully electric one. To support this transformation, we are allowed to shift the start time d_l of operations of a bus line l by a small amount. Due to the periodic timetable, this induces a shift in starting times of all trips of the respective bus line, e.g., a shift in 5 minutes will results in all trips of this bus line starting 5 minutes later. For each bus line l, we define PTP_l as the set of all feasible starting times that would be acceptable by the operator.

Our objective is to select the bus line start times $s' = (s'_1, \ldots, s'_n)$, that result in the cost-minimal timetable, where n is the number of bus lines and $s'_l \in PTP_l$. To this end, we propose an iterative, Benders-like procedure presented in Figure 1. The core of our approach consists of the *Timetabling Master Problem* and *EVSP Subproblem* blocks that correspond to solving the PESP and EVSP, respectively. While we present exemplary implementations of these blocks in Sections 4 and 5, which are used for our case study later, we highlight here that especially the EVSP subroutines can be solved with any EVSP algorithms. This makes our procedure easily adjustable to varying constraints and runtime requirements from practice.

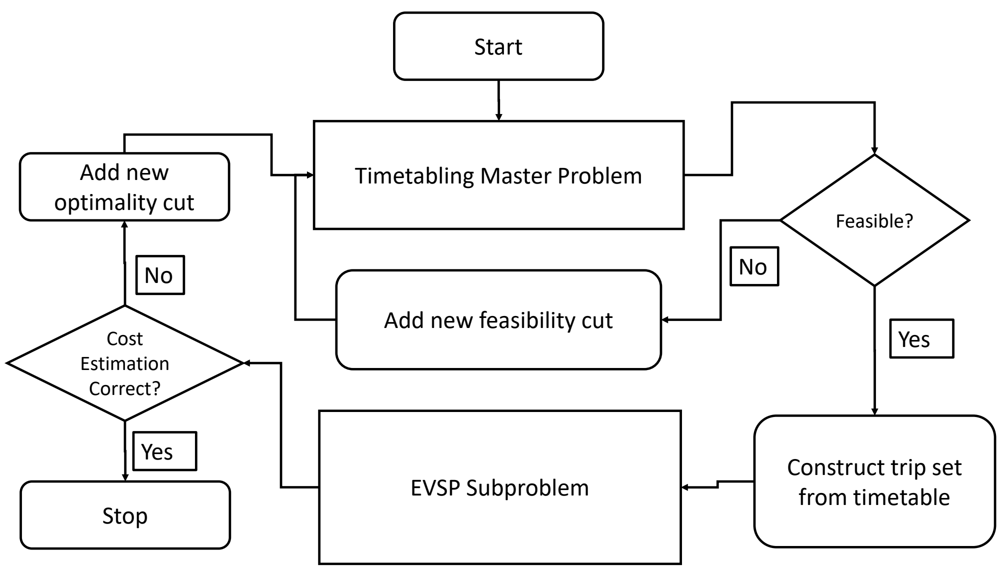

Figure 1: Iterative Approach for the integrated timetabling and vehicle scheduling problem.

The task of the *Timetabling Master Problem* is to find a good guess for a timetable by selecting start times $s' = (s'_1, \ldots, s'_n)$. As a timetable is subject to some (complex) transfer constraints, not all combinations of feasible starting times result in a feasible timetable. Hence, we define AS as the set of all feasible timetables. If the guessed start times are infeasible, i.e., $s' \notin AS$, we add what we call a feasibility cut which ensures that s' can no longer be selected. Otherwise, the *Timetabling Master Problem* assigns a cost estimation to the found feasible timetable, computes the set of trips $\mathcal{T}$ that must be executed to operate this timetable, and passes it to the EVSP subroutine.

The *EVSP Subproblem* finds the assignment of electric vehicles to the trips in $\mathcal{T}$ and computes the cost of the resulting schedule. Afterwards, we compare the found cost with the estimation made by the *Timetabling Master Problem*. If the estimation is wrong, i.e., if the estimated costs were too low, we add an optimality cut ensuring that the schedule is assigned the correct cost. Otherwise, we terminate and have found an optimal timetable if the solution to the EVSP Subproblems are optimal.

The main design challenge of our iterative procedure is that the guessed costs have to be a lower bound on the real cost of a timetable. Therefore, the added optimality cuts can be interpreted as an iterative improvement on the lower bound, and the *Timetabling Master Problem* as a heuristic search for good timetable candidates. Next, we present how we implement the *Timetabling Master Problem* for our case study.

4 Timetabling Master Problem

We model the Timetabling Master Problem as a MIP with two types of decision variables β and γ_{li}. The binary decisions γ_{li} represent whether line l starts at time $i \in PTP_l$, and β captures the costs of the vehicle rotation for the selected timetable. For a timetable $S \in AS$ and line $l \in \mathcal{L}$, we denote the times $i_{l,S} \in PTP_l$ as the starting time of line l in schedule S. The MIP formulation can then be expressed as follows:

$$\begin{aligned}
\min \quad & \beta && && \text{(1a)}\\
\text{s.t.} \quad & \sum_{i \in PTP_l} \gamma_{li} = 1 && \forall l \in \mathcal{L} && \text{(1b)}\\
& \gamma_{l_1 i_1} + \gamma_{l_2 i_2} \leq 1 && \text{if } ((l_1, i_1),(l_2, i_2)) \in TP && \text{(1c)}\\
& \beta \geq \beta(S) - \sum_{l \in \mathcal{L}} M_l^S \cdot (1 - \gamma_{l i_{l,S}}) && \forall S \in AS && \text{(1d)}\\
& \gamma_{li} \in \{0; 1\} && \forall l \in \mathcal{L}, i \in PTP_l && \text{(1e)}\\
& \beta \geq 0 && && \text{(1f)}
\end{aligned}$$

The objective (1a) is to minimize the cost of the selected timetable. Constraints (1b) ensure that exactly one starting time is selected for each bus line. Constraints (1c) represent the feasibility constraints. In this framework, a timetable violation, specifically a transfer constraint violation, is assumed to occur if a line l_1 starts at i_1 and a second line l_2 starts at i_2. For all such violating pairs TP, the MIP is instructed not to select the violating combination of line start times. Constraints (1d) are the optimality constraints. These constraints ensure that for each feasible schedule $S \in AS$, the cost of the timetable corresponds to the optimal cost derived from the EVSP subproblem, $\beta(S)$. However, if the starting time of a bus line $l \in \mathcal{L}$ is selected to be different from that in S, it is possible to save up to M_l^S cost units with the new schedule.

One type of feasible but undesired solution for Model (1) occurs when all bus lines are shifted by the same interval, such as the entire original schedule being shifted by a fixed time. To avoid these unnecessary solutions, we include the following symmetry-breaking constraints in our model

$$\sum_{l \in \mathcal{L}} \gamma_{l i_{l,S}} \leq |\mathcal{L}| - 1$$

for all schedules S that are simply the original schedule shifted by a fixed time frame. Since each line has a unique starting time in schedule S, the above constraint ensures that such a schedule cannot be selected by model (1).

Both feasibility and optimality cuts are based on certain assumptions that warrant discussion. For feasibility cuts, we opted for a rule-based approach where a rule-based method is used to check for violated transfer constraints after a timetable is fixed. This approach is simpler to implement in practice than directly integrating the PESP into the above MIP. Additionally, it offers flexibility to adjust to various additional requirements that may arise in practice. For instance, if some transfer constraints are defined over a triple rather than a pair of bus lines, the rule-based method and the constraints (1c) can be easily adapted to this setting, whereas modifying an integrated model would pose significant challenges.

For optimality cuts, the presented cuts are independent of the EVSP solver used. This independence allows for easy adjustments to the underlying modeling constraints, such as recharging assumptions, without the need to overhaul the framework itself. Furthermore, the constraints partition the overall cost of a schedule $\beta(S)$ into components associated with each bus line's starting time (γ_{li}). As a result, these cuts facilitate a smart enumeration process. Instead of enumerating all feasible schedules, the search direction is guided towards bus lines that significantly impact the overall cost. Additionally, by integrating our iterative framework into a MIP, we obtain the classic benefits of a primal-dual bound, i.e., an optimality gap, at each iteration step.

Note that the quality of the found solution of Model (1) directly depends on the employed EVSP solver. If we solve the EVSP to optimality, i.e., if $\beta(S)$ is the cost of the optimal vehicle schedule for schedule S, then the found solution is also optimal. Additionally, any bound on the quality of $\beta(S)$ directly transfers to that of the integrated problem. If the EVSP is solved heuristically, Model (1) also becomes a heuristic, without providing any indication of the quality of the found solution. This motivates us to focus on approximation algorithms that provide a trade-off between runtime and solution quality for the EVSP in the next section.

5 EVSP Subproblem

In this section, we focus on our implementation of an EVSP solver. Because the EVSP is an NP-hard optimization problem (Sassi & Oulamara, 2016) that we have to solve in every iteration of the *Timetabling Master Problem*, we focus on heuristics. We recap some of the basic heuristics from the literature and provide theoretical insights into the quality of the found solutions.

In the following, we focus on a very basic version of the EVSP that includes many simplifying assumptions, such as a homogeneous bus fleet and a linear charging curve. While all the presented algorithms can easily be adjusted to model more complex problem settings, our theoretical results—especially the found approximation ratios—depend on these assumptions. As the presented basic setting is a special case of many problem variants with more complex assumptions, we provide general insights into the performance of the discussed heuristics for various EVSP variants.

5.1 Problem Setting

The EVSP is formally defined as follows: Given a set of trips $\mathcal{T}$, where each trip t has a start time d_t, end time a_t and energy requirement e_t, the problem consists of finding a feasible vehicle rotation, i.e., assignment of trips to electric buses. As the primary cost driver is the acquisition of additional buses, our objective is to minimize the number of electric buses utilized. We represent each electric bus b by the set of trips assigned to it, i.e., $b \subseteq \mathcal{T}$. We call a bus feasible if the temporal constraints are satisfied, i.e., no two trips in the bus take place at the same time. Thus, we can represent a vehicle rotation by a set of buses $\mathcal{B}$, and finding a feasible vehicle rotation for all trips is then equivalent to finding a partition of $\mathcal{T}$ into electric buses. Each bus has a battery capacity of $D \in \mathbb{R}$ and recharges linearly with a rate r. In some settings, the buses can recharge between two trips. This is called opportunity recharging. The current battery of bus b at time point $\tau \leq T$ is denoted by $e^b(\tau)$. All buses start fully charged with $e^b(0) = D$. If a bus b serves a trip t, for each time point during the trip's execution $\tau \in [d_t, a_t]$, the current battery capacity is given by

$$e^b(\tau) = e^b(d_t) - e_t \cdot \left(\frac{\tau - d_t}{a_t - d_t}\right).$$

If opportunity recharging between the end of a trip a_t and a time point $\tau \geq a_t$ is possible, the battery capacity is given by

$$e^b(\tau) = e^b(a_t) + r \cdot (\tau - a_t).$$

For our theoretical analysis later, we also make the following assumptions. We assume that all trip are round trips starting and ending at the same bus hub, which has enough charging capacity for all buses. Charging is either performed over night, or as opportunity charging between the trips. The trip's times already include the required changeover times, allowing any time window between two trips to be entirely spent with charging. Additionally, we assume that the charging speed of the battery is at least as fast as the battery consumption through driving. In other words, the time required for a full charge does not exceed the operational time of a bus. Given current advancements in charging technology and the efficiency of electric buses, this assumption holds true for most practical scenarios.

5.2 EVSP without opportunity recharging

The EVSP without recharging consists of finding a feasible vehicle rotation for a given set of trips $\mathcal{T}$. Here, a vehicle rotation is feasible if for each electric bus, the culminated energy consumption over all trips does not exceed the battery capacity. Thus, no recharging between trips is allowed. This constraint is similar to a Bin Packing Problem, where items are packed into as few bins of fixed size as possible. Actually, there is a variant of the Bin Packing Problem, the so-called Bin Packing with Conflicts (BPwC), which can be directly transformed into the EVSP without recharging. In the following, we will introduce the problem and discuss its connection to the EVSP.

The general bin packing problem asks, given a set of items with sizes and a bin capacity, how many bins are necessary to fit all items. A set of items fits in the same bin, if the sum of the sizes of these items is smaller or equal to the bin capacity. The BPwC extends the bin packing problem to additional constraints ensuring that certain pairs of items are not allowed in the same bin. Thus, in the BPwC, there is an additional conflict graph given with items as vertices and arcs marking pairs of items, which may not be put in the same bin. The items in the same bin have to be an independent set in the conflict graph.

For the special case of the BPwC where the conflict graph is colorable in polynomial time, Jansen and Öhring (1997) introduced a new heuristic. Since the EVSP is part of this special case, we briefly review their proposed heuristic. BPwC is a combination of two well-researched combinatorial problems, the graph coloring problem on the conflict graph and the bin packing problem. The authors propose to solve these two problems sequentially, by first computing a coloring on the conflict graph and then solving the Bin Packing Problem on each color class separately. Since there can be no conflicts between items of the same color class, we only have to solve instances of the regular Bin Packing problem without conflict for each color class. However, since this problem is still NP-hard, the authors test multiple, well-known heuristics for Bin Packing.

Jansen and Öhring (1997) showed that the proposed heuristic is an approximation algorithm. Approximations algorithms are heuristics that also provide a quality measure on the found solution called the approximation ratio. The approximation ratio of an algorithm is a measure of how close the solution provided by the algorithm is to the optimal solution. Specifically, for an optimization problem, the approximation ratio R is defined as the maximum ratio between the cost of the solution found by the algorithm and the cost of the optimal solution taken over all possible instances of the problem. Formally, for a minimization problem, it is given by $R = \max_I \left(\frac{A(I)}{OPT(I)}\right)$, where $A(I)$ is the cost of the solution produced by the algorithm for instance I, and $OPT(I)$ is the cost of the optimal solution for instance I. An algorithm with an approximation ratio close to 1 is considered to provide solutions that are close to optimal.

Depending on the Bin Packing heuristics used, the authors showed an approximation ratio of 3 for the NextFit Heuristic, 2.7 for the FirstFit Heuristic and 2.5 for the FirstFitDecreasing Heuristic.

There is a direct transformation from BPwC to the EVSP without recharging. We can associate the items with trips and the bins with electric buses. Then, the size of the items become the energy consumption of the trips, and the bin size is equal to the battery capacity. Finally, the temporal constraints can be expressed through the conflict graph. Each trip is then associated with a vertex in the conflict graph, and there is an edge between two trips if the two trips take place at the same time. Then, a bus b satisfies the temporal constraints, if and only if there are no arcs in the conflict graph between trips in b.

Finding the optimal coloring for the transformed EVSP instance is equivalent to solving the VSP. Because the conflict graph does not contain any battery constraints, the VSP boils down to finding a partition of the vertices of the conflict graphs (i.e. trips) such that no two vertices in the same partition are adjacent to each other, i.e. the trips can be assigned to the same bus. Since the conflict graph represents the temporal constraints, it is an interval graph and the VSP problem can be solved in polynomial time within our setting and thus,the coloring problem on the conflict graph is solvable in polynomial time despite the fact that the coloring problem on arbitrary graphs is NP-hard. Thus, we can apply the approximation algorithm proposed by Jansen and Öhring (1997) for the special case where the conflict graph is colorable in polynomial time for the EVSP without opportunity recharging.

5.3 EVSP with opportunity recharging

In this section, we extend the previous approximation results to the EVSP with opportunity recharging. To this end, we recap two basic heuristics often used for EVSP, link them to Bin Packing problem variants, and provide theoretical insights into their solution quality.

We start with an adaptation of the classical First Fit (FF) Heuristic presented in Algorithm 1. Given an ordering of the trips (line 3), the algorithm selects the trips according to this order and places them in the first available bus while respecting all underlying feasibility constraints (line 6-7) Due to its simplicity, this heuristic has often been applied to find some (initial) vehicle rotation, e.g., in Adler and Mirchandani (2017) and Dirks et al. (2022).

One advantage of this heuristic is that the point of checking $b_i \cup \{t^i\}$ for feasibility in line 6 is left undefined. The heuristic can be adapted to different assumptions by changing this testing procedure accordingly. In the case of the EVSP with opportunity recharging, we only need to ensure that each trip in $b_i \cup \{t^i\}$ is battery feasible and does not co-occur as an already assigned trip.

The quality of the solution produced by Algorithm 1 strongly depends on the chosen ordering. Assuming we know the optimal solution to the EVSP, we can choose an ordering based on this optimal solution that groups the trips within a single bus together. Using this ordering, our First Fit (FF) Heuristic would not use more buses than in the optimal solution; hence, solving the problem to optimality. Conversely, finding such an ordering that ensures optimality is NP-hard. Hence, simple orderings based on start time or energy consumption of the trips are common.

Algorithm 1 First Fit (FF) Heuristic for EVSP

1: **Input:** A list of trips, each with a starting time, end time and energy consumption, ordering on trips $<^t$
2: **Output:** Assignment of trips to electric buses $\mathcal{B}$

3: $t^1, \ldots, t^n \leftarrow$ Sort all trips according to $<^t$

4: Initialize $b_1 = \{\}$, $k = 2$ and the ordered list $\mathcal{B} = [b_1]$, containing one electric bus with no assigned trips

5: **for** $i = 1, \ldots, n$ **do**
6:     find the smallest $j \in \{1, \ldots, k\}$ such that $b_j \cup \{t^i\}$ is feasible
7:     $b_j \leftarrow b_j \cup \{t^i\}$
8:     **if** no feasible insertion was found **then**
9:         Add a new empty electric bus b_k to $\mathcal{B}$.
10:         $k \leftarrow k + 1$
11:     **end if**
12:     goto line 6
13: **end for**

Instead of considering all the feasibility requirements at once, a commonly used approach is to first solve the VSP and then apply a repair step that makes the found schedule feasible for electric buses (Olsen et al., 2020). While this approach enables us to employ efficient existing VSP solvers in the first step, the found VSP solution limits the flexibility of the subsequent repair step, resulting in potentially poor solution quality.

We study these types of heuristics that first solve the VSP and then repair the solution in detail. Algorithm 2 shows our First Fit Coloring (FFC) Heuristic that first solves the VSP and then applies the FF Heuristic on the vehicle rotation of each bus independently to make the schedule feasible for electric buses. Hence, it follows the same idea as the heuristic from Jansen and Öhring (1997) by first partitioning the trips into color classes by solving the VSP, and ensuring that battery capacities are respected within each color class.

We proceed by proving approximation bounds for the FFC Heuristic. We then show that the found approximation ratio is tight for all Heuristics using this approach, i.e., any algorithm that first solves the VSP and then operates on the found vehicle rotations independently cannot achieve a better approximation ratio than Algorithm 2.

Our approach to proving the approximation ratio of Algorithm 2 consists of bounding the maximum number of buses required to replace a single gasoline bus. To obtain this bound, we begin by making a structural observation about the partial solutions $\mathcal{B}_g$ for one gasoline bus g: for each pair of electric buses in the partial solution and at each point in time, the sum of their batteries is at least the battery capacity.

Lemma 1. Let $\mathcal{T}$ be a set of trips, $g \subseteq \mathcal{T}$ a gasoline bus and $b^1, b^2 \subseteq g$ with $b^1 \cup b^2 \subseteq g$, $b^1 \cap b^2 = \emptyset$

Algorithm 2 First Fit Coloring (FFC) Heuristic

1: **Input:** A list of trips $\mathcal{T}$, each with a starting time, end time and energy consumption, ordering on trips $<^t$
2: **Output:** Assignment of trips to electric buses $\mathcal{B}$.

3: Compute a solution $\mathcal{G}$ to the VSP with gasoline buses.

4: **for** each color gasoline bus $g \in \mathcal{G}$ **do**
5: Apply the First Fit Heuristic in Algorithm 1 using $<^t$ to solve the EVSP, obtaining a feasible vehicle rotation $\mathcal{B}_g$ for the current gasoline bus.
6: $\mathcal{B} = \mathcal{B} \cup \mathcal{B}_g$
7: **end for**

be two battery-feasible electric buses serving trips in g. Then, at each point in time τ, the sum of the batteries $e^{b^1}(\tau) + e^{b^2}(\tau) \geq D$.

Proof. Since all trips can be serviced by one gasoline bus, at each point in time there is at most one trip that needs to be served. Thus, at most one electric bus is driving at any point in time τ and the other electric bus is either charging or already fully charged. If one bus is fully charged, the assumption trivially holds. Let t be the trip in g with latest start time and $d_t \leq \tau$. Since we assume that charging a battery is at least as fast as discharging through driving, the sum over the battery capacities does not get lower if only one bus charges. Thus, $e^{b^1}(\tau) + e^{b^2}(\tau) \geq e^{b^1}(d_t) + e^{b^2}(d_t)$. It only remains to be shown that the assumption holds for the beginning of each trip $t \in g$. We show this by induction over the number of trips $n = |g|$. The assumption holds trivially for $n = 1$. For $n \geq 2$, let $t_n \in g$ be the last trip. By induction, we know that the assumption holds for $g \setminus \{t_n\}$. Additionally, it obviously holds

$$e^{b^1}(d_{t_n}) + e^{b^2}(d_{t_n}) \geq e^{b^1}\left(d_{t_{n-1}}\right) + e_{t_{n-1}} + e^{b^2}\left(d_{t_{n-1}}\right) - e_{t_{n-1}} \geq D$$

Thus, the assumption holds for the beginning of each trip. This concludes the proof. □

This Lemma already yields an approximation ratio for instances with large battery capacities, where all trips have an energy requirement of at most half the battery capacity.

Corollary 1. If each trip t satisfies $e_t \leq \frac{D}{2}$, then the FFC Heuristic is a 2-approximation.

Proof. Since the sum of the current battery of two electric buses is always at least the battery capacity, one of the two buses has at least half the battery capacity. Thus, trips smaller than half the battery capacity can always be inserted in one of two electric buses. Thus, the partial solution for each gasoline bus consists of two electric buses. Since the optimal solution of the VSP is also a lower bound on the optimal solution of the EVSP, the solution computed by the heuristic takes at most double the number of buses compared to the optimal solution. □

For small battery capacities, we obtain a worse approximation ratio.

Theorem 1. The FFC Heuristic with items sorted by starting time is a 3-approximation.

Proof. Let $g = \{t_1, \ldots, t_n\} \in \mathcal{G}$ be a gasoline bus from the optimal solution of the VSP with the trips sorted by starting time. We show by induction over the number of trips n that in the result computed by the heuristic, each gasoline bus is replaced by at most three electric buses and at the end time of each trip a_t, one of the three electric buses is fully charged. For $n \leq 2$, the assumption obviously holds. For $n > 2$, the trip t_n would be inserted last by the Heuristic. By induction, we know that the assumption holds for $g \setminus \{t_n\}$. Thus, applying the heuristic to the set of trips $g \setminus \{t_n\}$ computes (at most) three electric buses $b^1, b^2, b^3 \subseteq g$. Additionally, at the end of trip t_{n-1} (resp. at the beginning of trip t_n), at least one of the three buses is fully charged. Thus, t_n can be inserted in that bus. It only remains to be shown that at the end of trip t_n, one bus is again fully charged. If there is a bus which is fully charged at the end of t_{n-1} and t_n is not inserted in this bus, this holds trivially. Thus, we now look at the case of t_n being inserted in the only bus which is fully charged. Here, we have to differentiate between two cases.

Case 1: b^2 or b^3 are fully charged and t_n is inserted in that bus. Since the Heuristic always inserts the trip in the first bus where the trip fits, this can only be the case if t_n does not fit into b^1. Thus, at the beginning of t_n, bus b^1 has battery capacity $e^{b^1}(d_{t_n}) < e_{t_n}$. With Lemma 1, we get that both buses b^2 and b^3 must have battery $e^{b^j}(d_{t_n}) > D - e_{t_n}$ for $j \in \{2,3\}$. Since only one of them serves trip t_n, the other bus charges at least e_{t_n} during the trip and is fully charged by the end of it.

Case 2: b^1 is the only fully charged bus. This can only be the case if t_{n-1} did not fit into b^1, else it would be inserted into b^1, and b^1 could not be the only fully charged bus. We can now use the same argumentation as in case 1 to show that either b^2 or b^3 have to be fully charged by the end of trip t_{n-1}, which contradicts the assumption that b^1 is the only fully charged bus.

Thus, g can be serviced by three electric buses. Since the number of gasoline buses required for the schedule is an upper bound on the number of electric buses required, this concludes the proof. $\square$

Both bounds carry over to the Integrated Timetabling Approach.

Corollary 2. For the Integrated Timetabling Approach, the solution found by Model (1) when using FFC to solve the EVSP subproblem is at least a 3-approximation. If each trip t satisfies $e_t \leq \frac{D}{2}$, the solution is at least a 2-approximation.

For Algorithm 1, we discussed that the quality of the found solution is strongly dependent on the chosen ordering. This is not the case for Algorithm 2. In contrary, the approximation ratio is independent of the selected algorithms for the repair step, showing that the VSP is in general to restrictive to build upon.

Theorem 2. A heuristic for the EVSP which first solves the VSP optimally and then solves the EVSP on each gasoline bus separately cannot be better than a 3-approximation.

Proof. Let $n \geq 2$, $\varepsilon > 0$ small, and the battery capacity $D = 1$. Furthermore, the buses charge at a rate of 1 per hour and require one battery of $D = 1$ to drive an hour. The time horizon is $\frac{n+1}{2}$ hours.

We begin by looking at a small instance with only three trips t_1, t_2 and t_3. The trips start right after each other in that order and have energy requirement $e_{t_1} = 1, e_{t_2} = \frac{1}{2}$ and $e_{t_3} = \frac{1}{2} + \varepsilon$ respectively. This instance could be served by one gasoline bus; however, even in an optimal solution, three electric buses are required. The first two trips and do not fit in the same bus and require two buses. At the end of the second trip, both buses have battery capacity $\frac{1}{2}$, thus the third trip fits in neither of the two buses and requires a third bus. We now present an instance of the EVSP, where the optimal solution to the VSP computed in the heuristic requires n gasoline buses which each include three trips of the form introduced above. We identify one trip $t_i = (d_i, a_i, e_i)$ by its starting time d_i, end time a_i and energy requirement e_i. Then, let the set of all trips be given by

$$\begin{aligned}
S_j &= \{t_1^j = (\frac{j}{2}, 1 + \frac{j}{2}, 1), t_2^j = (1 + \frac{j}{2}, 1 + \frac{j+1}{2}, \frac{1}{2}),\\
&\quad t_3^j = (\frac{j+1}{2}, \frac{j+2}{2} + \varepsilon, \frac{1}{2} + \varepsilon)\}\\
S_j^* &= \{(k, k + \frac{1}{2}, \frac{1}{2}) \colon k \in \{\left\lceil \frac{j+2}{2} + \varepsilon \right\rceil, \ldots, \left\lfloor \frac{n}{2} + 2 \right\rfloor\}\}\\
\mathcal{S} &= \bigcup_{j=1}^{n} S_j \cup S_j^*
\end{aligned}$$

The trips in S_j are the three special trips requiring three buses. The trips in S_j^* are dummy trips to ensure that for each pair $j \neq j'$, the trips from S_j and the trips from $S_{j'}$ are not all served by the same gasoline bus. Both the solution using gasoline buses and the optimal solution using electric buses are visualized in Figure 2 for even n. Each trip is visualized by a colored rectangle with their position on the x-axis marking start and end times. Each row represent a bus. Trips from t_1^j are drawn in blue, t_2^j in green, t_3^j in violet, and all trips from S_j^* in yellow. The upper half of the Figure shows a solution of the VSP using gasoline buses. The lower half shows an optimal solution of the EVSP using electric buses. In this instance, an optimal solution with gasoline buses would require at least n buses since there are n yellow trips of the form $(\frac{n}{2} + 2, \frac{n+1}{2} + 2, \frac{1}{2}) \in S_j^*$ for each $j \in \{1, \ldots, n\}$. Then, an optimal solution using gasoline buses is given by each bus g^j serving all trips $S_j \cup S_j^*$ for each $j \in \{1, \ldots, n\}$.

The optimal solution of the EVSP using the trips of the gasoline bus g^j requires 3 electric buses due to the trips in S_j. Thus, each Heuristic which separates the gasoline buses requires at least 3 electric buses for each coloring class and thus returns a solution with at least $3n$ electric

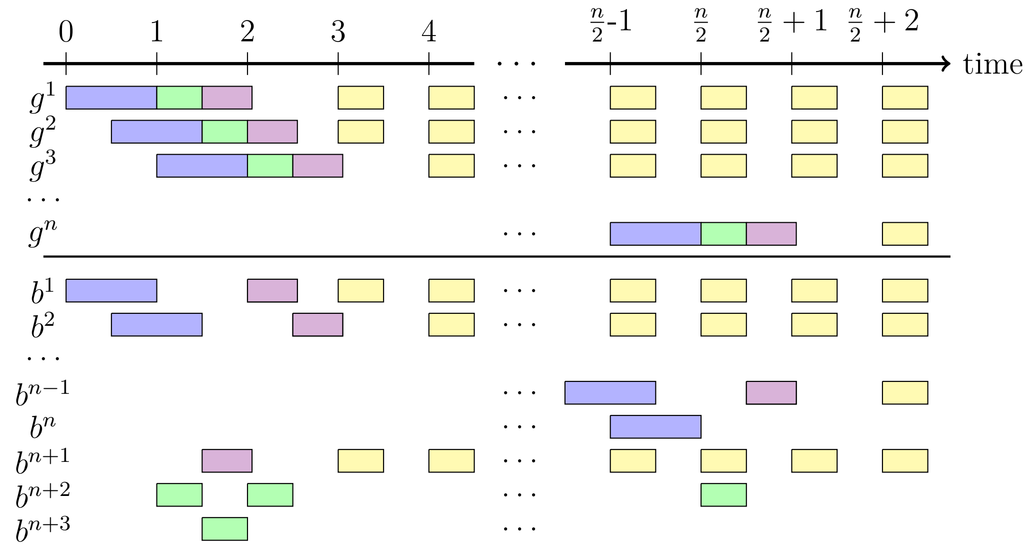

Figure 2: Visualization of the solutions for the VSP and EVSP for even values of n

buses. However, an optimal solution with electric buses only requires $n+3$ electric buses. Then, a solution with $n+3$ electric buses is given by:

$$\begin{aligned}
b_i &= \{t_1^i, t_3^{i+1}\} \cup S_{i+1}^* && i = 1, \dots, n-1 \\
b_n &= \{t_3^1\} \cup S_1^* \\
b_{n+1} &= \{t_1^n\} \\
b_{n+2} &= \{t_2^{2j} \colon j = 1, \dots, \frac{n}{2}\} \\
b_{n+3} &= \{t_2^{2j-1} \colon j = 1, \dots, \frac{n}{2}\}
\end{aligned}$$

Thus, the Heuristic is at best a 3-approximation for large values of n. □

Algorithm 3 First Fit ordered by Coloring (FFobC) Heuristic

1: **Input:** A list of trips $\mathcal{T}$, each with a starting time, end time and energy consumption, ordering on gasoline buses $<^{\mathcal{G}}$, ordering on trips $<^t$
2: **Output:** Assignment of trips to electric buses $\mathcal{B}$.

3: Compute a solution $\mathcal{G} = [g^1, \dots, g^k]$ to the VSP with gasoline buses ordered using $<^{\mathcal{G}}$
4: Sort the trips in each bus $g^i = [t^{i,1}, t^{i,2}, \dots, t^{i,|g^i|}]$ according to $<^t$
5: $\mathcal{B} \leftarrow$ First Fit Heuristic with ordering according to ordered list $g^1 \cup g^2 \cup \dots \cup g^k$

Theorem 2 clearly shows that heuristics based on first solving the VSP and then repairing the found vehicle rotations are prone to getting stuck in bad local optima after the VSP step. To escape such local optima, we propose combining these iterative heuristics with the global approach of the FF Heuristic. Algorithm 3 presents our combined First Fit ordered by Coloring (FFobC)

Heuristic, which is a combination of the FF and FFC Heuristics. We first solve the VSP, but then use this solution to derive an ordering for the FF Heuristic from the found VSP solution. While it is evident that FFobC will always be at least as good as the FFC Heuristic, it also has the chance of escaping bad local optima. Looking at the counter-example from Theorem 2, the new heuristic would return an optimal solution to this instance. Also, the theoretical results from Corollary 1 and Theorem 1 hold for FFobC Heuristic. It is an open question whether a better approximation ratio can be shown.

6 Case Study

We present a computational study to demonstrate the quality and effectiveness of our Integrated Timetabling and Vehicle Scheduling approach. To this end, we used an instance based on real live data from the bus network of the city of Aachen, Germany. The experiments were carried out using a Intel(R) Core(TM) i5-8265U CPU@1.60GHz and 4 physical cores. The algorithms were implemented using Java 21.0.1 and GUROBI 10.0.3 with Windows 10 as operating system.

6.1 Problem Instance and Parameters

The city of Aachen, located in western Germany with approximately 250,000 inhabitants, operates a comprehensive public transport system that relies solely on bus services, which is typical for mid-sized cities in Germany.

The bus network of Aachen is structured around a central hub known as the "Bushof". The Bushof serves as the primary interchange point where most of the bus lines converge, facilitating easy transfers between different lines. This central hub model is designed to optimise connectivity and minimise transfer times for passengers traveling across various parts of the city, and can be found in several cities across Germany.

Our instance includes 8 bus lines with a total of 347 trips and about 280 hours total driving time, which could be served by 19 gasoline buses. All data was taken from the public timetable information of the AVV. Four of the lines have a headway of 15 minutes, two have a headway of 30 minutes and the remaining two have a headway of one hour. The first trip starts at 5am and the last trip ends just before 1am the next day, resulting in an operational time window of 20 hours. We slightly adapted the trips to fit our assumption of having only round-trips starting from the same depot by merging two back and forth trips into one round-trip. Considering that the start times match perfectly, this is also implemented in practice. For transfer constraints, we define that if a line l_2 departs within a 5 to 15-minutes window after the arrival of line l_1, passengers commonly transfer from line l_1 to line l_2. Therefore, in any feasible schedule, l_2 must depart within this time window relative to the arrival of l_1. The complete list of bus lines and transfer constraints used can be found in A.

The electric buses currently in use in Aachen have a battery capacity of 292kW and a range of 150-200km (electrive, 2021). Since our formulation expresses battery capacity in terms of maximum driving time, we conservatively estimate a driving time of 3h by using the speed limit of 50km/h in German cities. To assess the impact of larger battery capacities, we conducted experiments with driving times of 3h, 7h and 16h. The charging stations currently available in Aachen are used for over-night charging and supply approximately 75kWh. For opportunity charging between trips, we assume the presence of fast-charging stations, which charge a 300kW battery to full capacity in 3 hours. Additionally, we conducted some tests for the case without opportunity recharging between trips.

For the integrated timetabling approach, we allow a shift in starting times of up to 10 minutes. We discretize the operational time window into 5-minute intervals. Thus, we permit time shifts of 5 or 10 minutes in either direction for each line. Finally, we set M_l^S to the number of buses that serve trips from line l in the heuristic solution of the EVSP for schedule S.

To solve the EVSP subproblem, we tested the three heuristics from Section 5: the (regular) First Fit (FF) heuristic 1, the First Fit Coloring (FFC) heuristic 2, and the First Fit ordered by Coloring (FFobC) heuristic 3. For each heuristic, we tested ordering the trips by start time, which is denoted by _start, as well as sorting the trips by energy consumption, which is denoted by _energy, leaving us with a total of 6 heuristics.

For each heuristic, the Integrated Timetabling Approach computed a solution in under two minutes. There were no significant differences in computation time between the heuristics or across different battery capacities. In each computation, the EVSP was solved up to 6,000 times, which is a substantial improvement over simple enumeration of all shifted time tables, which would require solving the EVSP for up to 390,000 different schedules ignoring transfer constraints.

6.2 Results without opportunity recharging

We begin by examining the results allowing only overnight charging. While solving this instance, we track all the computed timetables and report our results in the interval form shown in Figure 3. The range indicated for each heuristic represents the range of EVSP solutions computed during all

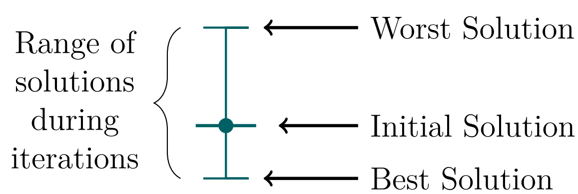

Figure 3: Legend Solution Presentation

iterations, while the marked solution indicates the solution obtained from applying the heuristic to the initial schedule currently used in Aachen. The lower bound of the interval is the final result. A large interval suggests significant potential gains from shifting the start times if a suboptimal initial schedule is currently used. The gap between the lower bound of the interval and the initial solution reflects the improvement from the current schedule through the integrated timetabling approach.

Figure 4 shows the results for all battery capacities and heuristics without opportunity recharging between trips. The results are presented for the different battery capacities. Additionally, the dashed blue line in each graph represents a lower bound on the number of electric buses. In the schedule, trips with a total driving time of ≈ 280 hours have to be served. Thus, even if the trips match perfectly, a schedule requires at least 94 buses with battery capacities of 3h, 40 buses with battery capacities if 7h or 18 buses with battery capacities of 16h. The lower bound for 16h battery capacities can be raised to 19 buses because the initial schedule can be served by 19 gasoline buses. The number of gasoline buses is also optimal under time shifts, thus, 19 is also a lower bound on the number of electric buses needed.

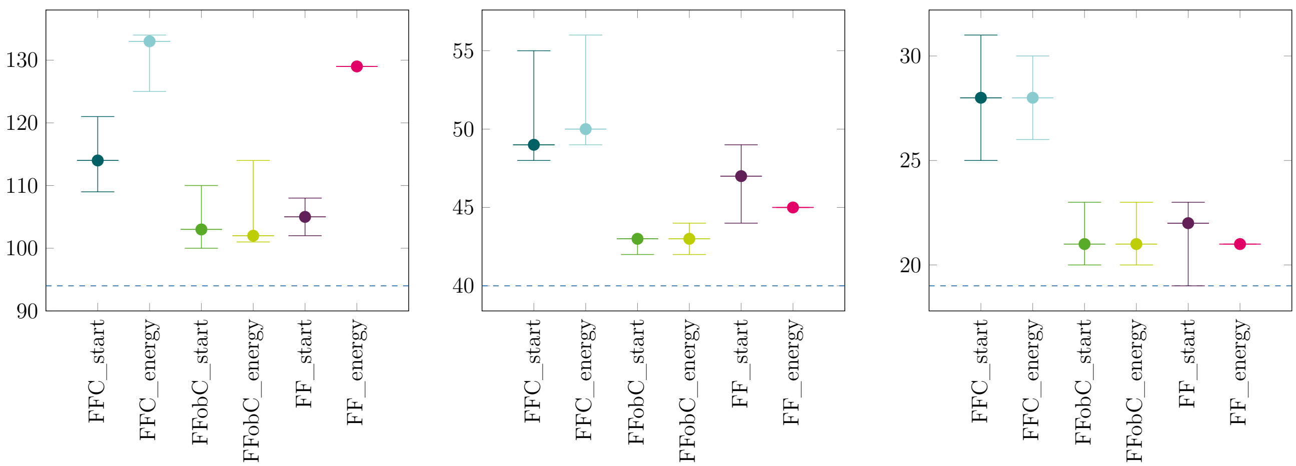

(a) 3h battery **(b)** 7h battery **(c)** 16h battery

Figure 4: start solution and gap between worst and best computed solution to the Integrated Timetabling and Vehicle Scheduling Problem over all iterations, instances without recharging

As can be seen, FF_energy shows no improvement through the integrated approach and computed the same result in each iteration. This may be due to the fact that the ordering of the trips is not affected by shifts in start times. In contrast, all other heuristics either partially sort by start times or are influenced by the gasoline bus plan, which inherently depends on start times. These results highlight the previously discussed high dependency of the selected ordering for the solution quality.

The results show that FFobC generally provides the best results, which confirms our expectations from the theoretical analysis. The only exception to this is FF_start, which computes a solution requiring one less electric bus for the instances with 16h battery capacity. Compared to FFC, especially for instances with medium and large batteries, the worst solutions computed by FFobC are still better than the best solution computed by FFC. Furthermore, the best solutions computed by FFobC are not far from the lower bounds. Since it is not likely that these lower bounds can be reached for small and medium batteries, this indicates that the computed solutions are close to optimal with an optimality gap of at most 10% on all three instances. The range of solutions during the iterations is small, especially for medium and large battery capacities, which implies less room for improvement when using the integrated approach. Additionally, we note that sorting by start time or energy consumption does not have a large effect on the results, the only difference being one electric bus. This may be caused by the ordering in the FFobC primarily depending on the gasoline schedule and only secondarily on the start time or energy consumption.

Upon examining the solutions returned by FFC, we observe that it includes many buses with large shares of their battery capacity left unused. This shows that because the VSP solution ignores battery restrictions, the found vehicle rotations cannot easily be executed by a subfleet of electric buses, resulting in the heuristic getting stuck in bad local optima. In contrast, FFobC can escape these parts of the solution space by reassigning buses between the given vehicle rotations from the VSP solution, which explains the difference in solution quality between the two heuristics. Additionally, it is noteworthy that there is a significant difference between the two sorting methods for the instance with small battery capacity. Here, sorting by energy performs much worse than sorting by start time. This is surprising because sorting by decreasing size is usually a good heuristic for Bin Packing. This discrepancy may be due to large trips blocking buses, which are better distributed when sorting by start time.

Finally, FF_start is overall slightly worse than FFobC for small and medium batteries and slightly better for large batteries, where it reaches the lower bound and computes an optimal solution. Notably, the range of solutions becomes larger with increasing battery capacity. This may be because the instance resembles a coloring problem more than a bin-packing problem, especially for medium to large battery capacities. When comparing the two sorting methods, sorting by start time appears to be slightly more effective than sorting by energy consumption across all heuristics.

6.3 Results with opportunity recharging

Next, we examine the influence of opportunity recharging between trips on the integrated approach. In Figure 5, we see the results for the case that recharging between trips is allowed. In addition to the lower bounds of 19 gasoline buses represented by the dashed, blue lines, the dotted, green lines represent upper bounds for the FFC and FFobC heuristics gained from the approximation ratios proven in Section 5. There, we showed that the heuristics require at most 3 electric buses for each gasoline bus (or 2 electric buses if the battery capacity is larger than double the longest trip). The longest trip in the instance takes 130 minutes. Thus, for 7 and 16 hour batteries, we get an upper bound of 38 electric buses for the initial and optimal solutions by Corollary 1. For battery capacities of 3h, we get an upper bound of 57 buses by Theorem 1.

As can be seen, incorporating opportunity recharging between trips returns results that are generally similar to those observed without recharging. In particular, the performance of the heuristics relative to each other is similar to the previous section, with FFobC yielding the best results. However, for FF_energy with small and medium batteries, we observe a slight improvement through

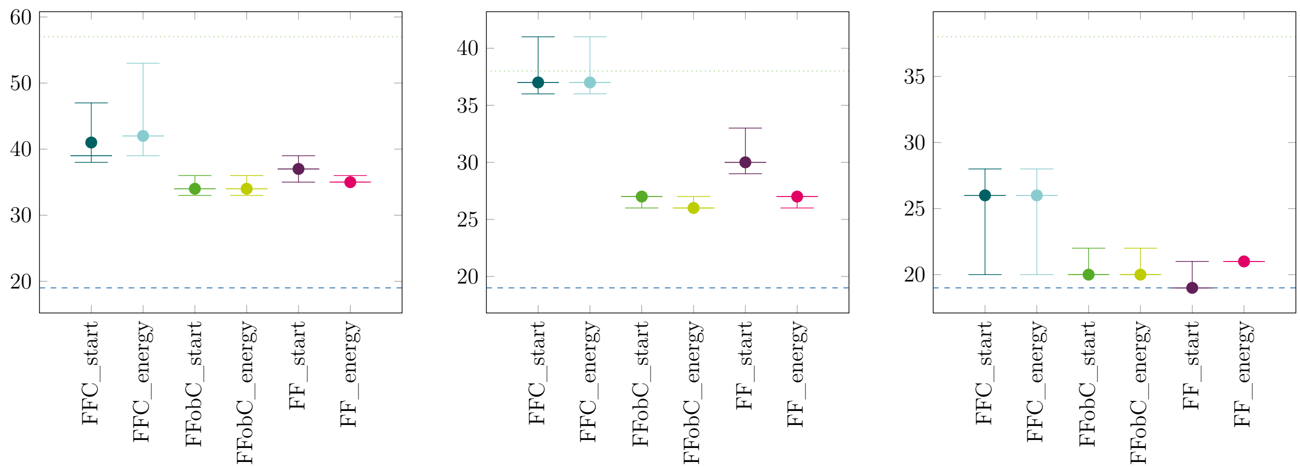

(a) 3h battery **(b)** 7h battery **(c)** 16h battery

Figure 5: start solution and gap between worst and best computed solution to the Integrated Timetabling and Vehicle Scheduling Problem over all iterations, instances with opportunity recharging

integrated timetabling. Notably, FFC_energy performs better on the instance with a small battery compared to the instance without recharging. The option to recharge likely compensates for scheduling errors made early in the timetable. For battery capacities of 16h, we also note that the best solutions computed by the FFobC and FF heuristics did not change. For FF_start, this is due to the fact that the best solution in the instance without opportunity recharging already reached the lower bound of 19 gasoline buses. Thus, no better solution is possible. For the other heuristics, this indicates that adding opportunity recharging is not very effective for buses with large battery capacities.

On instances with small and medium battery, the best solution of all heuristics still needs 25% more electric buses than the lower bound given by the number of gasoline buses required. However, this does not necessarily indicate bad performance of the heuristics but rather reflects on the quality of the lower bound on instances with smaller batteries. For large battery capacities, the lower bound is again reached by FF_start even in the initial solution. This suggests that with a battery capacity of 16 hours and opportunity recharging, the operational differences between electric and gasoline buses become negligible. This also explains the lack of improvement over the initial schedule in this case, as it was already optimized for gasoline buses.

Finally, we note that the upper bounds computed from the approximation ratio is not tight for the FFobC heuristics and can probably be improved. For the FFC heuristics, the upper bound is nearly reached by the initial solution with battery capacities of 7h, which shows that the bound is tight even in applications with medium battery capacities for the FFC heuristics. In these computations, the upper bound is broken by some solutions during the iterations. The upper bound only holds for schedules that can be served by 19 gasoline buses. This is not necessarily the case for all schedules computed during the iterations, thus, the worst solution computed needing more than 38 electric buses does not invalidate the approximation ratio.

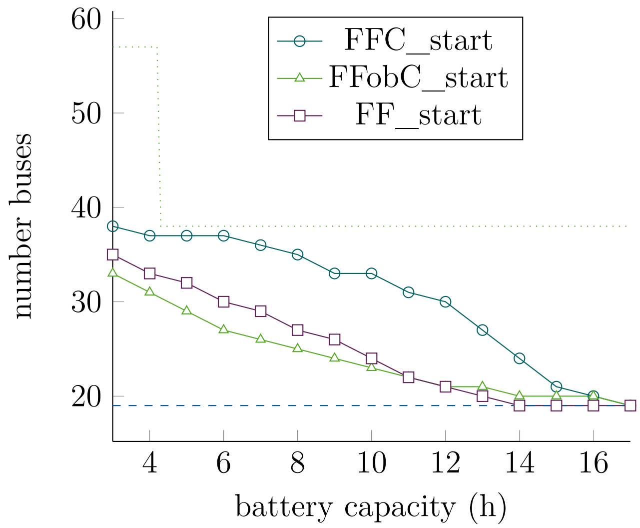

Figure 6: Best Solutions to the Integrated Timetabling and Vehicle Scheduling Problem with opportunity recharging

To further investigate the influence of the battery capacity on the solution, we computed the solution ranges for all battery capacities from 3h to 17h exemplary for FFC_start, FFobC_start and FF_start with opportunity recharging. The lower bound of 19 gasoline buses is again depicted as a dashed, blue line, and the corresponding upper bound as a dotted, green line.

Comparing the three heuristics, we can see that FFC_start returns the worst results across all battery capacities and nearly reaches the upper bound given by the approximation ratio for battery capacities of 5h and 6h. This shows that the approximation ratio found in Corollary 1 is nearly tight on this instance for FFobC_start. The other two heuristics perform better, with FFobC_start performing best for small and medium battery capacities and FF_start performing best on instances with large battery capacities of at least 13h. Specifically, the lower bound of 19 buses is reached for battery capacities of 14h when using FF_start, which is earlier than the other two heuristics that only reach the lower bound for battery capacities of 17h. This confirms our previous results.

Furthermore, for FFC_start, the number of buses barely changes for small batteries of up to 9h and then reduces fast until reaching the lower bound with a battery capacity of 17h. The opposite holds for FFobC_start, where the reduction in the number of electric buses needed is faster for small batteries and slows down for battery capacities greater than 12h. This is due to FFobC_start being more flexible and spreading trips from one electric bus across all other buses when increasing the battery capacity. This is not possible for FFC_start, where the color classes are solved separately with up to 3 electric buses each and trips can only be shifted between these 3 buses when increasing the battery capacity.

Finally, we note that since the lower bound is first reached by FF_start for battery capacities of 14h, using buses with capacities larger than 14h when opportunity recharging is possible is unnecessary in this instance. In additional experiments where opportunity recharging is not possible, all heuristics returned at least 20 electric buses for battery capacities of 15h. Thus, in this case, using battery capacities of up to 16h can be beneficial.

6.4 Managerial Insights

Finally, we look at the potential savings by implementing the changes suggested in this study. The spider plot in Figure 7 provides a visual representation of the improvements achieved by our proposed changes: Opportunity recharging, using larger batteries and integrated timetabling. Additionally, the plot shows the influence of combining these changes. All heuristics show very

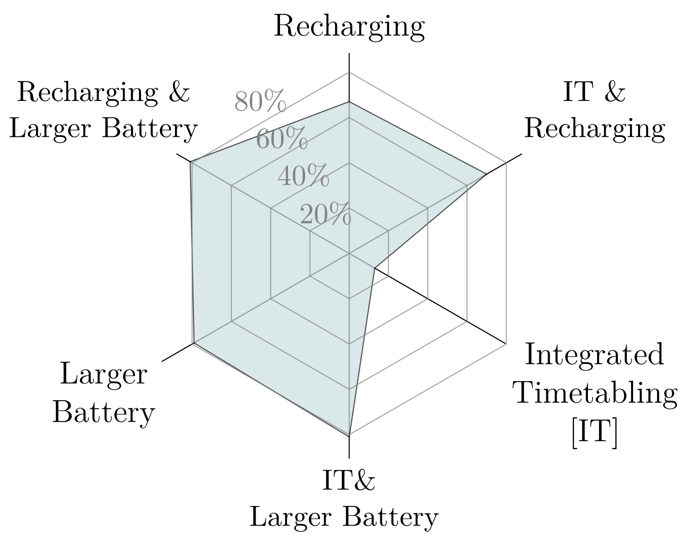

Figure 7: Improvements through opportunity recharging, larger battery capacity and Integrated Timetabling for FFC_start

similar behavior in this context, thus, we exemplarily show the results of FFC_start. The plot shows the percentages of electric buses saved for each of the proposed changes. As a basis for the comparisons, we used the initial solution without opportunity recharging between trips and 3h battery capacity.

Using buses with larger batteries has the largest impact, followed by opportunity recharging that still achieves saves around 60% of electric buses compared to our comparison scenario. Integrated timetabling archives the lowest overall savings of only up to 15%. However, integrated timetabling is also the easiest to implement option. While both recharging and buses with larger battery capacity involve significant initial investments into the fleet and infrastructure, adjusting the timetable can be implemented faster and does not require any additional direct investments. Therefore, integrated timetabling can be used as an intermediate solution that enables the (partial) electrification of a larger part of the fleet at an earlier stage. Combining the integrated approach with either larger battery capacities or opportunity recharging results in slight improvements, reinforcing the potential of our integrated timetabling strategy as an economical and efficient solution for enhancing the scheduling of electric buses.

7 Conclusion

In this study, we introduce a novel integrated timetabling approach that facilitates minor adjustments to existing periodic timetables to better implement the transition from gasoline buses to electric buses. We employ an iterative framework that incorporates optimality cuts derived from EVSP solutions into a Timetabling Master MIP model. To effectively address the EVSP, we study three heuristics. The first heuristic is an adaptation of the classical First Fit heuristic for the Bin Packing Problem tailored to the EVSP. The second heuristic involves initially solving the VSP and then repairing the solution. The third heuristic combines the ideas of the previous two. We study the theoretical properties of the discussed heuristics, showing that the second and third heuristic provides an approximation guarantee of at most 3. We also provide general insights into iterative heuristics that first solve the VSP and then adjust the found vehicle rotations to be executed by electric buses, showing that there are instances for which no such algorithm can be better than selecting 3 times more electric buses than the optimal solution.

We validate our approach through a case study conducted on the public transportation system in Aachen, Germany. The results confirm our theoretical insights and show worse computational performance of the repair heuristic compared to the other heuristics. We also derive managerial insights based on our analysis.

Future research directions are manifold. One promising approach is the development of an integrated approach for the EVSP. For the integrated timetabling and EVSP, expanding the case study to relax some of the assumptions made in this research would provide a more comprehensive understanding of the problem. Specifically, future studies could examine heterogeneous bus fleets with different types of electric buses, incorporating non-linear charging profiles, uncertainties such as variable energy consumption due to traffic conditions or weather, and incorporating the costs associated with developing and maintaining charging infrastructure. Finally, applying the developed methodologies and models to case studies in other cities would be valuable.

References

Abdelwahed, A., van den Berg, P. L., Brandt, T., Collins, J., & Ketter, W. (2020). Evaluating and optimizing opportunity fast-charging schedules in transit battery electric bus networks. *Transportation Science*, *54*(6), 1601–1615. DOI: 10.1287/trsc.2020.0982.

Adler, J. D., & Mirchandani, P. B. (2017). The vehicle scheduling problem for fleets with alternative-fuel vehicles. *Transportation Science*, *51*(2), 441–456. DOI: 10.1287/trsc.2015.0615.

An, K. (2020). Battery electric bus infrastructure planning under demand uncertainty. *Transportation Research Part C: Emerging Technologies*, *111*, 572–587. DOI: 10.1016/j.trc.2020.01.009.

Berthold, K., Förster, P., & Rohrbeck, B. (2017). Location planning of charging stations for electric city buses. In *Operations Research Proceedings* (pp. pp. 237–242) (pp. 237–242). Springer International Publishing. DOI: $10.1007/978-3-319-42902-1_32$.

Borndörfer, R., Hoppmann, H., Karbstein, M., & Löbel, F. (2017, July). The modulo network simplex with integrated passenger routing. In *Operations Research Proceedings 2016* (pp. pp. 637–644) (pp. 637–644). Springer International Publishing. DOI: $10.1007/978-3-319-55702-1_84$.

Bunte, S., & Kliewer, N. (2009). An overview on vehicle scheduling models. *Public Transport*, *1*(4), 299–317. DOI: 10.1007/s12469-010-0018-5.

Carosi, S., Frangioni, A., Galli, L., Girardi, L., & Vallese, G. (2019). A matheuristic for integrated timetabling and vehicle scheduling. *Transportation Research Part B: Methodological*, *127*, 99–124. DOI: 10.1016/j.trb.2019.07.004.

Ceder, A., & Wilson, N. H. (1986). Bus network design. *Transportation Research Part B: Methodological*, *20*(4), 331–344. DOI: 10.1016/0191-2615(86)90047-0.

De Vos, M. H., Van Lieshout, R. N., & Dollevoet, T. (2024). Electric vehicle scheduling in public transit with capacitated charging stations. *Transportation Science*, *58*(2), 279–294. DOI: 10.1287/trsc.2022.0253.

Desfontaines, L., & Desaulniers, G. (2018). Multiple depot vehicle scheduling with controlled trip shifting. *Transportation Research Part B: Methodological*, *113*, 34–53. DOI: 10.1016/j.trb.2018.05.011.

Dirks, N., Schiffer, M., & Walther, G. (2022). On the integration of battery electric buses into urban bus networks. *Transportation Research Part C: Emerging Technologies*, *139*, 103628. DOI: 10.1016/j.trc.2022.103628.

electrive. (2021). Ekoenergetyka baut 17 depot-lader in aachen [Accesssed on 06.06.2024].

Gattermann, P., Großmann, P., Nachtigall, K., & Schöbel, A. (2016). Integrating passengers' routes in periodic timetabling: A sat approach. DOI: 10.4230/OASICS.ATMOS.2016.3.

Goerigk, M., & Schöbel, A. (2013). Improving the modulo simplex algorithm for large-scale periodic timetabling. *Computers & Operations Research*, *40*(5), 1363–1370. DOI: 10.1016/j.cor.2012.08.018.

Großmann, P., Hölldobler, S., Manthey, N., Nachtigall, K., Opitz, J., & Steinke, P. (2012). Solving periodic event scheduling problems with SAT. In *Advanced Research in Applied Artificial Intelligence* (pp. pp. 166–175) (pp. 166–175). Springer Berlin Heidelberg. DOI: $10.1007/978-3-642-31087-4_18$.

Huang, Z., Zhang, A., Dósa, G., Chen, Y., & Xiong, C. (2023). Improved approximation algorithms for bin packing with conflicts. *International Journal of Foundations of Computer Science*, 1–16. DOI: 10.1142/s0129054122460054.

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. DOI: 10.1016/j.trb.2014.08.010.

Janovec, M., & Koháni, M. (2019). Exact approach to the electric bus fleet scheduling. *Transportation Research Procedia*, *40*, 1380–1387. DOI: 10.1016/j.trpro.2019.07.191.

Jansen, K., & Öhring, S. (1997). Approximation algorithms for time constrained scheduling. *Information and Computation*, *132*(2), 85–108. DOI: 10.1006/inco.1996.2616.

Laporte, G., Ortega, F. A., Pozo, M. A., & Puerto, J. (2017). Multi-objective integration of timetables, vehicle schedules and user routings in a transit network. *Transportation Research Part B: Methodological*, *98*, 94–112. DOI: 10.1016/j.trb.2016.12.018.

Li, J.-Q. (2014). Transit bus scheduling with limited energy. *Transportation Science*, *48*(4), 521–539. DOI: 10.1287/trsc.2013.0468.

Li, L., Lo, H. K., Huang, W., & Xiao, F. (2021). Mixed bus fleet location-routing-scheduling under range uncertainty. *Transportation Research Part B: Methodological*, *146*, 155–179. DOI: 10.1016/j.trb.2021.02.005.

Liebchen, C., & Möhring, R. H. (2007). The modeling power of the periodic event scheduling problem: Railway timetables — and beyond. In *Lecture Notes in Computer Science* (pp. pp. 3–40) (pp. 3–40). Springer Berlin Heidelberg. DOI: $10.1007/978-3-540-74247-0_1$.

Liebchen, C., & Peeters, L. (2009). Integral cycle bases for cyclic timetabling. *Discrete Optimization*, *6*(1), 98–109. DOI: 10.1016/j.disopt.2008.09.003.

Liu, X., Qu, X., & Ma, X. (2021). Optimizing electric bus charging infrastructure considering power matching and seasonality. *Transportation Research Part D: Transport and Environment*, *100*, 103057. DOI: 10.1016/j.trd.2021.103057.

Liu, Z., & Song, Z. (2017). Robust planning of dynamic wireless charging infrastructure for battery electric buses. *Transportation Research Part C: Emerging Technologies*, *83*, 77–103. DOI: 10.1016/j.trc.2017.07.013.

Niekerk, M. E. V. K., van den Akker, J. M., & Hoogeveen, J. A. (2017). Scheduling electric vehicles. *Public Transport*, *9*(1-2), 155–176. DOI: 10.1007/s12469-017-0164-0.

Olsen, N., & Kliewer, N. (2020). Scheduling electric buses in public transport: Modeling of the charging process and analysis of assumptions. DOI: $10.23773/2020_4$.

Olsen, N., & Kliewer, N. (2022). Location planning of charging stations for electric buses in public transport considering vehicle scheduling: A variable neighborhood search based approach. *Applied Sciences*, *12*(8), 3855. DOI: 10.3390/app12083855.

Olsen, N., Kliewer, N., & Wolbeck, L. (2020). A study on flow decomposition methods for scheduling of electric buses in public transport based on aggregated time–space network models. *Central European Journal of Operations Research*, *30*(3), 883–919. DOI: 10.1007/s10100-020-00705-6.

Rogge, M., van der Hurk, E., Larsen, A., & Sauer, D. U. (2018). Electric bus fleet size and mix problem with optimization of charging infrastructure. *Applied Energy*, *211*, 282–295. DOI: 10.1016/j.apenergy.2017.11.051.

Sassi, O., & Oulamara, A. (2016). Electric vehicle scheduling and optimal charging problem: Complexity, exact and heuristic approaches. *International Journal of Production Research*, *55*(2), 519–535. DOI: 10.1080/00207543.2016.1192695.

Serafini, P., & Ukovich, W. (1989). A mathematical model for periodic scheduling problems. *SIAM Journal on Discrete Mathematics*, *2*(4), 550–581. DOI: 10.1137/0402049.

Shen, Y., Li, Y., Chen, C., & Li, J. (2023). Electric vehicle scheduling based on stochastic trip time and energy consumption. *Computers & Industrial Engineering*, *177*, 109071. DOI: 10.1016/j.cie.2023.109071.

Sustainable Bus. (2024). Electric bus, main fleets and projects around the world [Accessed on 28.08.2024]. https://www.sustainable-bus.com/electric-bus/electric-bus-public-transport/ Accessed.

SWA. (2019, July). Linienplanung: Das steckt dahinter [Accessed on 25.05.2024]. https://www.sw-augsburg.de/magazin/detail/linienplanung-das-steckt-dahinter/ Accessed.

The Federal Office of Justice Germany. (2024). Personenbeförderungsgesetz [Accessed on 24.05.2024]. https://www.gesetze-im-internet.de/pbefg/ Accessed.

Van Lieshout, R. N. (2021). Integrated periodic timetabling and vehicle circulation scheduling. *Transportation Science*, *55*(3), 768–790. DOI: 10.1287/trsc.2020.1024.

Van Lieshout, R. N., & Bouman, P. C. (2018). Vehicle scheduling based on a line plan. In R. Borndörfer & S. Storandt (Eds.), *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2018)* (pp. 15:1–15:14) (15:1–15:14, Vol. 65). Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik GmbH, Wadern/Saarbruecken, Germany. DOI: 10.4230/OASICS.ATMOS.2018.15.

Xu, X., Yu, Y., & Long, J. (2023). Integrated electric bus timetabling and scheduling problem. *Transportation Research Part C: Emerging Technologies*, *149*, 104057. DOI: 10.1016/j.trc.2023.104057.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT to rewrite the original draft with the goal of improving readability and correct spelling/grammatical mistakes. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Funding sources

Vladimir Stadnichuk is a member of the Ph.D.-programme “ACCESS”, supported by the funding scheme “NRW Forschungskollegs” by the Ministry of Culture and Science of the German State of North Rhine-Westphalia under grant number 321-8.03.07-127598.

Declaration of Interest

Declarations of interest: none

A Lines used for the Case Study

Table 8 shows a list of all lines used in our case study together with the starting times of the first and last trip as well as the headway and trip-length in minutes. Using these departure and arrival

Line	Start Time	End Time	Headway (min)	Trip-length (min)
3A	05:27	00:42	15	40
3B	05:00	00:15	15	37
13A	07:01	19:31	15	24
13B	06:57	18:57	15	30
14	06:11	23:11	30	95
16	06:02	20:02	60	111
24	7:30	21:00	30	60
46	6:32	20:32	60	130

Figure 8: List of all lines used in our case study

times, we set three transfer constraints: from line 14 to line 24, from line 16 to line 24 and from line 24 to line 14 as described in Section 6. Since the lines 3A, 3B, 13A and 13B have a trip every 15 minutes, transfer constraints for them would be redundant.