

Euler, Ricardo; Lindner, Niels; Borndörfer, Ralf

## Article

# Price optimal routing in public transportation

EURO Journal on Transportation and Logistics (EJTL)

## Provided in Cooperation with:

Association of European Operational Research Societies (EURO), Fribourg

*Suggested Citation:* Euler, Ricardo; Lindner, Niels; Borndörfer, Ralf (2024) : Price optimal routing in public transportation, EURO Journal on Transportation and Logistics (EJTL), ISSN 2192-4384, Elsevier, Amsterdam, Vol. 13, Iss. 1, pp. 1-15,  
<https://doi.org/10.1016/j.ejtl.2024.100128>

This Version is available at:

<https://hdl.handle.net/10419/325202>

## Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

## Terms of use:

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



<https://creativecommons.org/licenses/by-nc-nd/4.0/>



# Price optimal routing in public transportation

Ricardo Euler<sup>\*</sup>, Niels Lindner, Ralf Borndörfer

Zuse Institute Berlin, Takustraße 7, Berlin, 14195, Berlin, Germany

## ARTICLE INFO

### Keywords:

Multi-objective shortest path  
Fare structure  
Public transportation  
Monoid  
Conditional fare network  
Ticket graph  
RAPTOR

## ABSTRACT

We consider the *price-optimal earliest arrival problem* in public transit (POEAP) in which we aim to calculate the Pareto-set of journeys with respect to ticket price and arrival time in a public transportation network. Public transit fare structures are often a combination of various fare strategies such as, e.g., distance-based fares, zone-based fares or flat fares. The rules that determine the actual ticket price are often very complex. Accordingly, fare structures are notoriously difficult to model, as it is in general not sufficient to simply assign costs to arcs in a routing graph. Research into POEAP is scarce and usually either relies on heuristics or only considers restrictive fare models that are too limited to cover the full scope of most real-world applications. We therefore introduce *conditional fare networks* (CFNs), the first framework for representing a large number of real-world fare structures. We show that by relaxing label domination criteria, CFNs can be used as a building block in label-setting multi-objective shortest path algorithms. By the nature of their extensive modeling capabilities, optimizing over CFNs is NP-hard. However, we demonstrate that adapting the multi-criteria RAPTOR (McRAP) algorithm for CFNs yields an algorithm capable of solving POEAP to optimality in less than 400 ms on average on a real-world dataset. By restricting the size of the Pareto-set, running times are further reduced to below 10 ms.

## 1. Introduction

The desired shift to more sustainable means of transportation necessitates an increase in the modal share of public transportation systems. Recent studies show that discounted or even free fares significantly increase such a system's adaptation (Brough et al., 2022; Bull et al., 2021; Chen et al., 2020), indicating that, from a traveler's perspective, the design of fares and ticket prices are key parameters to foster its attractiveness. This effect, however, varies between socioeconomic groups: While many riders may value fast connections with few transfers, low-income riders are more likely to choose a more affordable mode of public transportation or to even abstain from using transportation at all if they deem fares to be too expensive (Blumenberg and Agrawal, 2014; Rosenblum, 2020). This underlines the need for routing algorithms that enable passengers to select journeys optimal with respect to their individual needs. Since these needs are rarely known, we present an approach that calculates Pareto sets of optimal journeys with respect to the earliest arrival time, the number of transfers and cost. It is then on the user to choose from the options presented the one most suited for their personal needs.

Ticket prices are determined by the public transit providers' fare structure. Following Fleishman et al. (1996) a *fare structure* is “the combination of one or more fare strategies with specific tickets” while the term *fare strategy* refers to a “general fare collection and payment

structure approach”. This can be, e.g., a distance-based or zone-based fare, a short-distance discount, a flat fare or a surcharge. Until now, a unifying framework for the algorithmic treatment of public transit fare structures has been lacking. Some fare strategies, for example distance-based fares, can be addressed easily using label-setting shortest path algorithms. Finding a journey that crosses the least amount of fare zones, however, is NP-hard (Blanco et al., 2016) and cannot be modeled using real-valued arc weights. In general, the subpath optimality principle does not hold for fare structures and thus label-setting algorithms cannot be directly applied to POEAP. Consider the following example: A traveler takes a detour that passes through an additional fare zone to avoid paying the surcharge of a special connection (e.g., a ferry). It is not unlikely that, at a later point, the surcharge has to be paid regardless (e.g., because the target stop can only be reached via a ferry). In that case, taking the detour was a suboptimal decision and using the corresponding label to prune other partial journeys breaks the optimality guarantee of shortest path algorithms.

### 1.1. Our contribution

We devise algorithms that solve the price-optimal earliest arrival problem (POEAP) efficiently in practice. To this end, we build upon

<sup>\*</sup> Corresponding author.

E-mail address: [euler@zib.de](mailto:euler@zib.de) (R. Euler).

the state-of-the-art public transport routing algorithms McRAP (Delling et al., 2015) and Tight-BMRAP (Delling et al., 2019), which combine an intelligent enumeration scheme with dominance checks. Our domination rules are based on *conditional fare networks* (CFN), a novel and flexible framework for modeling fare structures of public transportation providers capable of taking most underlying fare strategies into account. A CFN models fare structures as a *ticket graph* representing *relations* between tickets. Transitions between different tickets are modeled as directed arcs and usually depend on a number of additional parameters such as, e.g., fare zones or the traveled distance. These are modeled via partially ordered monoids and events. Fare strategies that can be expressed via CFNs include (but are not limited to): zone-based fares, distance-based fares, surcharges for special vehicles or night liners, discounted short-distance fares, transfer fares and all combinations thereof. We develop domination rules for CFNs based on path relations in the ticket graph instead of the price alone. This allows us to retain subpath optimality and prove that using these rules in label-setting MOSP algorithm does in fact yield lowest-price journeys. By further restricting the size of the Pareto-set, we solve POEAP in less than ten milliseconds over the intricate fare structure of a mid-sized public transit provider from Germany.

## 1.2. Related literature

There is ample research on routing problems in public transit networks. For an overview, see Bast et al. (2016). Delling et al. (2015) introduced the RAPTOR algorithm for very fast public transit routing. We refer occasionally to typical label-setting multi-objective shortest path algorithms, by which we mean, e.g., McRAP (Delling et al., 2015), Martins' algorithm (Martins, 1984), and recently Multi-Objective Dijkstra (Maristany de las Casas et al., 2021). Delling et al. (2019) introduce a version of McRAP, Tight-BMRAP, for computing restricted Pareto-sets.

The concept of relaxed subpath optimality is discussed in Berger and Müller-Hannemann (2009). Lastly, we refer to Disser et al. (2008) on how to model public transit systems with transfers in time-dependent graphs.

In contrast to the general activity of the field, literature on price-optimal routing is generally rather scarce. This is certainly due to the usually intricate nature of public transit fare structures. Most approaches deal with fares on a heuristic basis or only consider a very narrow set of fare strategies. Most notably, Müller-Hannemann and Schnee (2005) study fare structures that entail distance- and relation-based prices, i.e., structures that are usually associated with long-distance public transportation. They approximate fares by assigning a fixed price to every arc. This approach, however, does not account for fare strategies such as, e.g., fare zones and short-distance discount tickets. Both are usually more prominent in local public transportation.

Reinhardt and Pisinger (2011) consider optimizing the number of fare zones as a special case in their study of non-additive objective functions in (multi-criteria) shortest path problems. However, their approach relies on target pruning as the sole domination technique, so that partial paths cannot be pruned until a  $s, t$ -path is known. In contrast, our approach applies to more general fare structures, includes target pruning, but also allows for pruning at earlier stages.

Schöbel and Urban (2021) identified conditions under which price-optimized routing is tractable for zone- and distance-based fare strategies. Additionally, they identify the no-elongation and no-stopover properties as desirable properties for fare structures. This follows a line of research concerned with the design of fare structures. For a review of recent work, see Schöbel and Urban (2021). Blanco et al. (2016) showed zone-based fares to result in NP-hard routing problems if reentering a zone does not entail additional costs. The proof relies on a reduction of the problem of finding a path with a minimum number of colored edges (Broersma et al., 2005). It was given in the context of flight trajectory optimization with overflight costs (Blanco et al., 2017).

A version adapted for public transport is given by Schöbel and Urban (2021). Delling et al. (2015) used RAPTOR to compute journeys that touch the smallest number of fare zones. Recently, Gündling (2020) considered price-optimized routing in intermodal transportation. They, however, only considered mileage-based and flat fares.

Our approach combines ideas from automata theory and optimization over monoids. For an overview of automata theory, see Hopcroft and Ullman (1979). The ticket graph concept is inspired by the application of finite automata to the language-constrained shortest path problem (Barrett et al., 2000). It is different, however, in that it serves to evaluate paths instead of restricting the set of feasible paths. Furthermore, our approach also covers fares based on numerical parameters that are not expressed as part of a formal language. Finally, it has been known for a while that shortest path algorithms can be generalized to ordered monoids (Zimmermann, 1981) and semirings (Mohri, 2002) in a straightforward fashion. Recently, monoids were also proposed as a general constraint model for (single-criteria) resource constrained shortest path problems (Parmentier, 2019).

This paper is an extended and improved version of work presented at the ATMOS'19 conference (Euler and Borndörfer, 2019). Apart from streamlining the presentation and proofs, the following additions were made: We now better motivate the interplay between the monoid and the ticket graph. We propose a superior approach for dealing with overlap areas. We provide an adaption of the recent Tight-BMRAP algorithm to our use case. This leads to an improvement in algorithmic performance of up to two orders of magnitude compared to the previous results in Euler and Borndörfer (2019). Finally, we provide a complexity analysis and investigate the relation of our approach to automata theory in Appendix A.

Ticket graphs of various German public transit providers can be found in Borndörfer et al. (2018, 2021).

## 1.3. Overview

In Section 2, we introduce the fare structure of MDV, an association that is responsible for the public transit fares for various operators in the Leipzig–Halle region of Germany. This fare structure will serve as a running example for the rest of the paper. We present conditional fare networks in detail in Section 3 and show how they can be used to model various aspects of fare structures. The algorithmic treatment of fares and domination rules is laid out in Section 4. Section 5 discusses how the multi-criteria RAPTOR and Tight-BMRAP algorithms can be modified to use CFNs for price-optimal search. An evaluation of the framework's performance is conducted in Section 6 using the network and fare structure of MDV. Section 7 concludes the paper with some closing remarks. In Appendix A, we provide a supplementary complexity analysis of POEAP and explore links to automata theory.

## 2. Running example: MDV

We introduce the reader to some intricacies of fare structures in public transit using the example of *Mitteldeutscher Verkehrsverbund* (MDV) (*Mitteldeutscher Verkehrsverbund GmbH*, 2019a). Throughout Sections 3 and 4, the MDV fare structure will serve as a running example to illustrate our core concepts. A schematic depiction of MDV's fare plan is given in Fig. 1.

**Example 2.1 (The fare system of MDV).** MDV's area of operations covers large rural areas in Eastern Germany, as well as the conurbation of Halle and Leipzig. As of 2019, this area is divided into a set of 56 pairwise disjoint fare zones. In most cases, the price depends on the number of visited fare zones: there are price levels for one to six fare zones. We denote the respective tickets by  $Z_i$  with  $i \in [6]$ . For example, traveling from station  $A$  to station  $L$  in Fig. 1 requires ticket  $Z_4$ . For all paths covering more than six fare zones, a ticket for MDV's whole area of operations has to be purchased, which we denote by  $M$ . The two

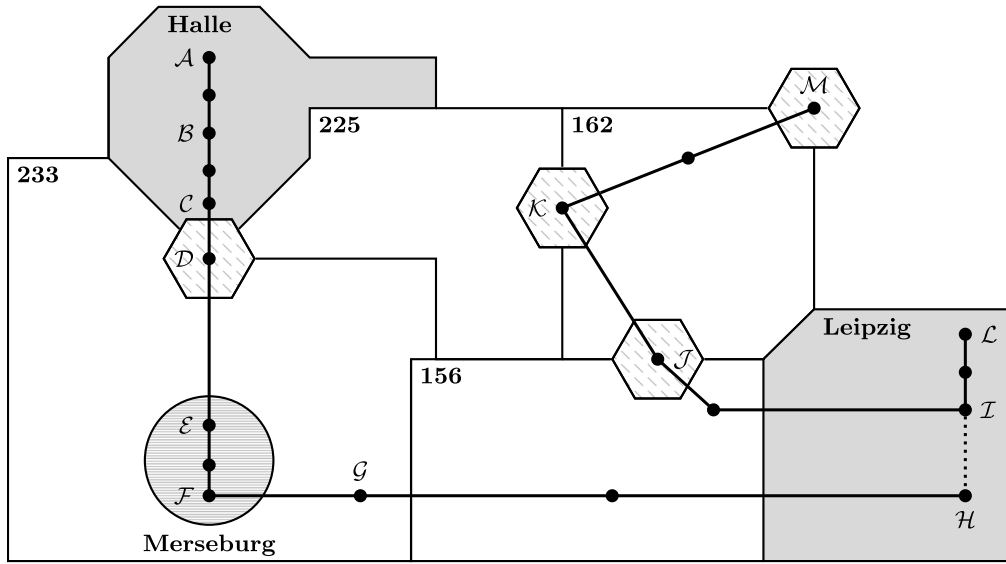


Fig. 1. A section of MDVs fare plan with two lines and six fare zones. Two of the fare zones, colored in light gray, are the cities of Halle and Leipzig. Vertically hatched hexagons represent overlap areas that can be counted as either of the neighboring zones. The horizontally hatched circle represents the small city Merseburg in which a special discounted fare is applicable. Small black nodes represent public transit stops. Footpaths are indicated by dotted lines.

larger cities, Halle and Leipzig, each form a single fare zone. Traveling in these zones requires special tickets more expensive than  $Z_1$ . These, we denote by  $H$  and  $L$ , respectively.

For all paths that pass through multiple fare zones, they, however, count as normal zones, i.e., one of the tickets  $Z_2, \dots, Z_6, M$  is applied. Hence, the paths  $A-C$  and  $H-L$  incur tickets  $H$  and  $L$ , respectively, while the path  $A-G$  incurs  $Z_2$ . Several smaller cities are part of larger fare zones, but allow for discounted fares (city fares) when traveling only in that city. For each city  $c$  in the set of such cities  $C$ , we denote the ticket by  $C_c$ . The path  $E-F$  in Merseburg ( $m$ ), hence, requires the ticket  $C_m$ . When extending the path to stop  $G$ , the ticket  $Z_1$  becomes applicable. As of 2019, there are 17 cities with city fares and two price levels (which we denote by  $C_1$  and  $C_2$ ). For paths starting in Halle and Leipzig, there are discounted tickets for short trips ( $D_H$  and  $D_L$ ), which can be used for a maximum number of four stops without transfers. Hence, paths  $A-B$  and  $I-L$  are admissible for discounted tickets  $D_H$  and  $D_L$ , respectively, while paths  $A-C$  and  $H-L$  are not. Discounted tickets also exist for other zones ( $D$ ). These are a little cheaper and depend on the length of the journey (4 km maximum) instead of the number of visited stops. Sometimes it is possible to choose between city fares and length-based discounts. In this case, the city fare is applied because it is cheaper. To not unduly burden people living at the borders of fare zones, MDV uses overlap areas. These can be counted as part of either of their adjacent fare zones, whichever is most benevolent to the traveler. For example, when traveling from  $J$  to  $M$ , all stops are counted as part of fare zone 162 and thus ticket  $Z_1$  would be applicable. When traveling from  $E$  to  $D$ ,  $D$  counts as part of the fare zone 233 but in the path  $A-D$  it counts as part of Halle. Hence, tickets  $Z_1$  and  $H$  are applicable, respectively.

While we cover the most important features of the fare structure, we do ignore some edge cases and explicit exceptions. These are among other things: Slightly different discount rules for specific trains, counting stations that are passed without a stop for discounted tickets and exceptions for a specific tunnel. This is done in part because they are not properly reflected in our dataset, and in part to simplify presentation.

### 3. A formal framework for fare structures

Consider a (directed) routing graph  $G = (V, A)$ , in which arcs represent either public transport connections, footpaths, or transfers

between lines and/or modes of transportation. Public transit journeys can then be interpreted as paths in  $G$ . In the following, we will consider a time-dependent formulation as presented, for example, by Disser et al. (2008), i.e., we are given a time-dependent FIFO travel time function  $c(a) : I \rightarrow I$  on each arc  $a \in A$ , where  $I$  is the set of time points.

In  $G$ , every path  $p$  is associated with a ticket  $\tau \in T$  from a ticket set  $T$  that has to be bought to use  $p$ . Each ticket has a corresponding price  $\pi(\tau) \in \mathbb{Q}^+$ . In the following, we might also write  $\pi(p)$  instead of  $\pi(\tau)$  if  $\tau$  is the ticket associated with  $p$ . The ticket of a path is determined by the fare structure.

We aim to solve the *price-optimal earliest arrival problem* (POEAP). We refer to Definition 3.6 for a precise definition, but the essence is that, for given  $s, t \in V$ , we want to find a Pareto-set of  $s, t$ -paths  $P_{s,t}^* \subseteq P_{s,t}$  with respect to arrival time and ticket price in  $G$ . Here,  $P_{s,t}$  denotes the set of all  $s, t$ -paths in  $G$ .

Ideally, we want to solve POEAP by taking advantage of the existing literature on label-setting MOSP algorithms. Ticket prices, however, usually cannot be modeled via real-valued FIFO functions on arcs. Hence, a framework for fare structures is needed that allows us to label paths in a way that (a) labels can be updated quickly when a new arc is relaxed, (b) dominance relationships between labels can be established that respect the subpath optimality property.

#### 3.1. Modeling with monoids

Note that in Example 2.1, the price of a path depends on several parameters: the number of visited stations, the total distance traveled, the set of visited fare zones and on indicators reporting whether transfers were made or whether the path crossed city borders.

All these parameters share several key properties: First, there is a natural partial order on them, indicating which configuration requires a more expensive ticket. For example, for fare zones  $A, B, C$  we have  $\{A, B\} \subset \{A, B, C\}$ ; for distances and transfers it is the canonical order on  $\mathbb{N}$ . The parameters can be summed up along a path using an appropriate notion of addition. For distances, this is the normal addition of natural numbers; for fare zones, it is the union of sets; we can use the logical OR ( $\vee$ ) on the set  $\{0, 1\}$  for indicators. Finally, we can assume the existence of a neutral element for every parameter.

The above properties suggest that the structure of a *partially ordered positive monoid* is an appropriate model for a large number of fare-relevant parameters.

**Definition 3.1 (Partially ordered monoid).** A monoid  $(H, +)$  is a set  $H$  together with an associative operation  $+$  (called addition) and a neutral element  $0 \in H$ , i.e.,  $h + 0 = h \forall h \in H$ . We call  $(H, +, \leq)$  a *partially ordered monoid* if  $\leq$  is a partial order on  $H$  that is translation-invariant with respect to the monoid operation  $+$ , i.e.,  $h_1 \leq h_2 \Rightarrow h_1 + x \leq h_2 + x \forall h_1, h_2, x \in H$ . If additionally  $0 \leq h \forall h \in H$ , we call  $(H, +, \leq)$  a *partially ordered positive monoid*.

Note that we can define the cross-product of two partially ordered monoids  $(H_1, +_1, \leq_1)$  and  $(H_2, +_2, \leq_2)$  by  $(H_1 \times H_2, +_{1,2}, \leq_{1,2})$ , where  $(h_1, h_2) +_{1,2} (i_1, i_2) := (h_1 +_1 i_1, h_2 +_2 i_2)$  and  $(h_1, h_2) \leq_{1,2} (i_1, i_2)$  if and only if  $h_1 \leq_1 i_1$  and  $h_2 \leq_2 i_2$  for  $h_1, i_1 \in H_1$  and  $h_2, i_2 \in H_2$ . The cross-product of two partially ordered positive monoids is again a partially ordered positive monoid. This construction allows us to represent all the fare-relevant parameters in Example 2.1 above as a single partially ordered positive monoid  $(H, +, \leq)$ .

Price-optimal paths can then be found in the following way: We label each arc  $a \in A$  with a weight in  $H$  representing the relevant parameters on this arc. The weight of a path  $p$ , denoted by  $w(p)$ , lives in  $H$  as well and can be obtained by summing up the weights of the arcs of  $p$ . The ticket  $\tau$  for  $p$  and its price can then be derived from  $w(p)$  using the rules of the fare structure. Finding a price-optimal  $s, t$ -path with  $s, t \in V$  can now be achieved by finding the Pareto-set of  $s, t$ -paths with regard to the partial order of  $(H, +, \leq)$ . Here, we can apply a label-setting MOSP algorithm such as, e.g., Martins' algorithm, by using elements of  $H$  as labels and the partial order of  $(H, +, \leq)$  to establish dominance between labels (Parmentier, 2019). Note, however, that this set will likely still contain many dominated paths with respect to price. This necessitates the filtering out of superfluous paths in a post-processing step.

**Example 3.1 (MDV).** For MDV, we can construct a monoid in the following way: for each city  $c \in C$  we define the monoid  $(H_c := \{0, 1\}, \vee, \leq)$  as an indicator whether our path started in  $c$  and then left the city. Hence, all arcs representing a connection leaving the city carry the weight  $1 \in H_c$ . Furthermore, we represent the distance traveled by the monoid  $(H_{dist} := \mathbb{N}, +, \leq)$ , the number of visited stations by  $(H_{stop} := \mathbb{N}, +, \leq)$ , the set of fare zones by  $(H_{zone} := 2^Z, \cup, \subseteq)$  and finally the transfers by  $(H_{tran} := \{0, 1\}, \vee, \leq)$ . Price-optimal paths can then be computed by finding the Pareto-set over the monoid  $(H_{dist} \times H_{stop} \times H_{tran} \times H_{zone} \times \prod_{c \in C} H_c, +, \leq)$  and filtering out dominated paths in a post-processing step. Here,  $+$  and  $\leq$  are induced from the component monoids.

While the above modeling approach covers a reasonable set of real-world applications, it is not expressive enough to be of much use for complex fare structures. First, note that it requires an order-preserving relationship between the monoid and the ticket prices. This assumption does not hold in general: For example, some public transit associations (e.g., in the city of Bremen, Germany, before 2020 (Verkehrsverbund Bremen/Niedersachsen GmbH, 2019)) apply night surcharges on selected lines. This means that, e.g., in the early morning, it can be beneficial to start a journey later because it becomes cheaper. In particular, it is necessary to capture time in the monoid, but the mapping between time and price does not preserve order. Second, monoids often cannot express logical pricing conditions: For example, a price depending on the order on which stops are visited cannot be modeled using the monoid-based approach. Third, shortest path search over the monoid is agnostic to the fare structure and therefore tends to consider unnecessarily large search trees. In the MDV case, if we already know an  $s, t$ -path for which, e.g., the ticket  $D_L$  is applicable, all partial paths starting in  $s$  that require a more expensive ticket can be pruned even though they might not be dominated w.r.t to  $\leq$ . To do this, however, we must find a way to quickly obtain the corresponding ticket to labels from  $H$ .

Finally, note that the labels of a label-setting MOSP algorithm live in  $H$  and might be quite large. In most cases, however, it is not necessary

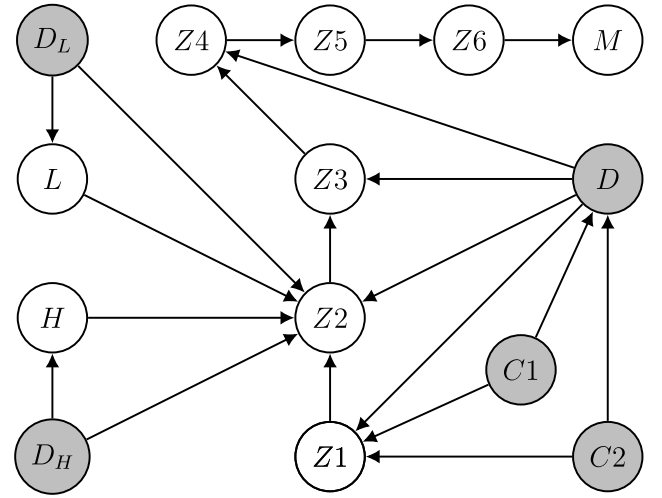


Fig. 2. Ticket graph associated with the MDV public transit network. To simplify the presentation, all tickets for city fares are collapsed to  $C_1$  and  $C_2$  representing the two price levels of city fares. Possible starting tickets are highlighted in light gray.

to carry the whole label along. Consider again the MDV case. If a path has left a city  $c$ , all city fares become unavailable and labels for all monoids  $H_c, c \in C$ , need no longer be considered. This information remains unavailable to an algorithm using the monoid-based model.

Computational results in Section 6 reveal that purely monoid-based modeling quickly becomes intractable, even when only considering the fare zone monoid  $(H_{zone}, \cup, \subseteq)$ .

### 3.2. Ticket graphs and fare events

To overcome the challenges laid out in the previous section, we extend our modeling in two directions: First, we want to take the path's ticket into account. To do so, we develop a model to represent tickets and their relationships. Second, we introduce *events* that model logical fare strategies. They also serve to reduce the size of the monoid. Addressing the first point, we notice that there is a natural progression of the applicable ticket along a path  $p$ . If  $p$  is short, a short-distance ticket might suffice. When  $p$  is extended by adding another stop  $v$  at its end, this ticket might no longer be applicable and now, e.g., a zone ticket might apply. When traveling even further, a ticket covering two zones might be needed. Hence, we can relate tickets to each other via their ability to transition into one another along paths in the routing graph. We formalize this observation by introducing a *ticket graph*  $\mathcal{T} = (T, E)$  that contains an arc  $e = (\tau_1, \tau_2) \in E$  if ticket  $\tau_1$  can transition into ticket  $\tau_2$ . Transitions depend on the weight  $w(p)$  and do not occur at every stop. Hence, we introduce a ticket transition function that checks  $w(p)$  and selects the appropriate ticket from the neighborhood of the ticket of  $p$  in  $\mathcal{T}$ . The ticket graph provides crucial advantages over the purely monoid-based approach: First, when we want to use price-based target-pruning (cf. Section 5.4), we must compute the current ticket price. The ticket graph allows doing this by updating tickets along a path, thereby only checking a few transition conditions on the outgoing edges. Without the ticket graph, all rules of the fare structure would need to be checked whenever a vertex is relaxed. Second, the ticket graph also carries information about the possible further ticket price development of a partial path: This knowledge is useful to design dominance rules.

**Example 3.2 (Running example: Ticket graph for MDV).** Consider the ticket graph in Fig. 2. All possible tickets introduced in Example 2.1 are represented as nodes. Whenever, a ticket can transition into another one, we introduce an arc. Note that, e.g., a discounted ticket for Halle



$D_H$  can never transition into a Leipzig ticket  $L$ . Now, consider again the path  $A - C$  in Fig. 1. This path requires a ticket covering one fare zone. Appending the station  $D$  will require a ticket covering two fare zones. This is modeled by performing a ticket transition in the ticket graph along the edge  $(Z_1, Z_2)$  induced by the ticket transition function for  $Z_1$ . When calculating the ticket price of a path that contains a subpath with ticket  $Z_3$  now only a condition on the number of fare zones needs to be checked to determine whether  $Z_3$  or  $Z_4$  is applicable. Without using the ticket graph, e.g., the applicability of all discounts would need to be recalculated. Furthermore, when comparing two paths having, say, tickets  $Z_2$  and  $Z_3$ , respectively, a dominance check can safely mark the path with ticket  $Z_3$  as dominated even if it is shorter in distance. In the purely monoid-based approach, both paths would be nondominated.

To reduce the dimension of the monoid, note that the parameters determining the ticket of a path can be broadly categorized as (*fare*) *states*, e.g., number of stops, and (*fare*) *events*, e.g., a transfer or the boarding of a train that requires a surcharge. Up to now, events were included in the monoid via indicators. However, this is not strictly necessary: A transfer arc could cause a ticket transition in the ticket graph. Thereafter, transfers might be ineffectual and hence it is unnecessary to record them in the monoid. To do so, we need to annotate arcs in the routing graph not only with elements of a monoid but also with events. The distinction between *states* and *events* introduces some flexibility to the modeling, as sometimes aspects of a fare structure can be modeled as both. However, we naturally aim to keep the monoid as low-dimensional as possible. Using events, prices may now also depend on logical conditions, e.g., the order of events, which is not possible in the monoid-based approach.

**Example 3.3 (Running example: Fare events for MDV).** The monoid in Example 3.1 was introduced as  $(H_{dist} \times H_{stop} \times H_{tran} \times H_{zone} \times \prod_{c \in C} H_c, +, \leq)$ . We can see any transfer as an event *tra* occurring on a transfer arc of the routing graph. In the same way, leaving any city  $c \in C$  can be seen as an event *city* occurring on arcs crossing the city's borders. Hence, we can shrink the monoid to  $(H_{dist} \times H_{stop} \times H_{zone}, +, \leq)$ , reducing the size of a label in a MOSP algorithm by 17 entries for cities and by one entry for transfers. Finally, we introduce events *hal* and *lei* for public transit arcs ending in the special fare zones of Halle and Leipzig, respectively. This is purely a design decision, since the information could also be read from the state of  $H_{zone}$ . We obtain a set of events  $S = \{city, tra, hal, lei, s_0\}$  where  $s_0$  is a dummy event with no effect, see also Example 3.4.

### 3.3. Conditional fare networks

We now combine the three core ideas of a ticket graph, fare events and modeling with monoids to a formal model of public transit fare structures.

Again, let  $G = (V, A)$  be a routing graph. Additionally, let  $(H, +, \leq)$  be a positive, partially ordered monoid,  $\mathcal{T} = (T, E)$  a ticket graph and  $S$  a set of fare events. We label each arc  $a \in A$  with a weight  $w(a) \in H$  and a fare event  $e(a) \in S$ . By collecting these weights and events along a path  $p$  in  $G$ , we build its *fare state*  $f(p)$ .

**Definition 3.2 (Fare state).** A fare state  $f \in T \times H$  is a pair of a ticket  $\tau(f)$  and a weight  $w(f)$ . We write  $F := T \times H$  for the space of all fare states.

Every vertex  $v \in V$  is labeled with an *initial fare state*  $\mu(v) \in F$ . Fare states will serve as path labels for shortest-path algorithms. They contain all the information necessary to decide domination between paths. In contrast to common (multi-objective) shortest-path applications, the arc labels  $H \times S$  live not in the same space as the path labels  $F$ .

We now want to enable the tracking of fare states along paths in  $G$ . To do so, we formalize the notion of the *ticket transition function* on tickets  $\tau \in T$  in the ticket graph  $\mathcal{T} = (T, E)$ . A ticket transition function

**Table 1**

MDV fare monoid.

Name	Represents	Ground set	Operator	Partial order	Neutral element
$H_{dist}$	Distance	$\mathbb{N}$	$+$	$\leq$	0
$H_{stop}$	Stops	$\mathbb{N}$	$+$	$\leq$	0
$H_{zone}$	Fare zones	$2^Z$	$\cup$	$\subseteq$	$\emptyset$

returns the ticket  $\tau_2 \in T$  a ticket  $\tau_1 \in T$  transitions into given an accumulated weight  $h \in H$  and a fare event  $s \in S$ . Possible candidates are the neighborhood of  $\tau_1$  in  $\mathcal{T}$  as well as  $\tau_1$  itself.

**Definition 3.3 (Ticket transition function).** The ticket transition function  $\Gamma : T \times H \times S \rightarrow T$  of  $\mathcal{T}$  is a function that, given a weight  $h \in H$  and event  $s \in S$ , maps each ticket  $\tau \in T$  into its closed out-neighborhood  $\delta^+(\tau) \cup \{\tau\}$ .

The definition is intentionally kept as general as possible to capture a large number of possible transition conditions. We use the notion of ticket transition functions to define the update of a fare state when relaxing an arc of the routing graph.

**Definition 3.4 (Fare update function).** Let  $f \in F$  and  $a \in A$ . Then, the fare update function  $\text{Up} : F \times A \rightarrow F$  is given by  $g := \text{Up}(f, a)$  with

$$w(g) := w(f) + w(a)$$

$$\tau(g) := \Gamma(\tau(f), w(g), e(a)).$$

The fare state of a path  $p = (v_1, \dots, v_n)$  can now be tracked by letting  $f_1 := \mu(v_1)$  and  $f_i := \text{Up}(f_{i-1}, (v_{i-1}, v_i)) \forall i = 2, \dots, n$ . In particular, when calculating  $f_i$  we need only consider the out-neighborhood of  $\tau(f_n)$  instead of reevaluating all fare strategies.

Combining all the above definitions, we arrive at the notion of *conditional fare networks* which can precisely describe a fare structure.

**Definition 3.5 (Conditional fare network).** Let  $G = (V, A)$  be a routing graph and let the following be given:

1. a directed acyclic ticket graph  $\mathcal{T} = (T, E)$  with transition function  $\Gamma$ ,
2. arc weights  $w : A \rightarrow H$  from a partially ordered, positive monoid  $(H, +, \leq)$ ,
3. arc events  $e : A \rightarrow S$ ,
4. initial fare states  $\mu : V \rightarrow F$  and
5. a price function  $\pi : T \rightarrow \mathbb{Q}_+$  that is monotonously non-decreasing along directed paths in  $T$ , i.e., if there is a directed  $\tau_1 - \tau_2$ -path in  $\mathcal{T}$  for  $\tau_1, \tau_2 \in T$ , then  $\pi(\tau_1) \leq \pi(\tau_2)$ . We write  $\pi(p)$  instead of  $\pi(\tau(f(p)))$  for a path  $p \in P$ .

We call the six-tuple  $(\mathcal{T}, \Gamma, w, e, \mu, \pi)$  a *conditional fare network*  $\mathcal{N}$  of  $G$ .

Note that cycle-freeness in  $\mathcal{T}$ , the monotonicity condition on  $\pi$  and the positivity of  $H$  ensure that no price-decreasing cycles exist in  $G$ . We consider those assumptions natural enough that any reasonable fare structures should satisfy them.

**Example 3.4 (Running example: Conditional fare network for MDV).** We can now give the complete conditional fare network  $\mathcal{N}$  for the fare structure of MDV. We have already introduced the ticket graph (Example 3.2). As in Example 3.3, we define the monoid  $(H, +, \leq)$  as  $(H_{dist} \times H_{stop} \times H_{zone}, +, \leq)$ . The components of  $(H, +, \leq)$  are summarized in Table 1.

The set of fare events is  $S = \{city, tra, hal, lei, s_0\}$ . The meaning of these events is summarized in Table 2. Vertices  $v \in V$  in the routing graph can now be annotated with an initial fare state from  $F = T \times H$ . For example, for station  $G$  in Fig. 1 we have  $\mu(v) = (D, (0, 0, \{156\}))$ , i.e., we start with the short-distance ticket, and the weight is composed of 0 m of distance, 0 visited stops and the fare zone 156. Arcs  $a \in A$  are

$$\begin{aligned}
\Gamma(M, h, s) &= M \\
\Gamma(Z_i, h, s) &= \begin{cases} Z_{i+1} & |h_{zone}| = i + 1 \\ Z_i & \text{otherwise} \end{cases} & \Gamma(D, h, s) &= \begin{cases} Z_1 & |h_{zone}| = 1 \wedge (h_{dist} > 4 \vee s = tra) \\ Z_2 & |h_{zone}| = 2 \wedge (h_{dist} > 4 \vee s = tra) \\ Z_3 & |h_{zone}| = 3 \wedge (h_{dist} > 4 \vee s = tra) \\ D & \text{otherwise} \end{cases} \\
\Gamma(L, h, s) &= \begin{cases} L & s = lei \vee s = tra \\ Z_2 & \text{otherwise} \end{cases} & \Gamma(D_L, h, s) &= \begin{cases} Z_2 & s \neq lei \wedge h_{stop} > 4 \\ L & s = tra \vee (s = lei \wedge h_{stop} > 4) \\ D_L & \text{otherwise} \end{cases} \\
\Gamma(H, h, s) &= \begin{cases} H & s = hal \vee s = tra \\ Z_2 & \text{otherwise} \end{cases} & \Gamma(D_H, h, s) &= \begin{cases} Z_2 & s \neq hal \wedge h_{stop} > 4 \\ H & s = tra \vee (s = hal \wedge h_{stop} > 4) \\ D_H & \text{otherwise} \end{cases} \\
\Gamma(C_1, h, s) &= \begin{cases} Z_1 & s = city \wedge h_{dist} > 4 \\ D & s = city \wedge h_{dist} \leq 4 \\ C_1 & \text{otherwise} \end{cases} & \Gamma(C_2, h, s) &= \begin{cases} Z_1 & s = city \wedge h_{dist} > 4 \\ D & s = city \wedge h_{dist} \leq 4 \\ C_2 & \text{otherwise} \end{cases}
\end{aligned}$$

Box I.

**Table 2**  
MDV fare events.

Fare event	Represents
city	Leaving city with city ticket $C_1$ or $C_2$
tra	Transfer
hal	Head of Arc is in Halle
lei	Head of Arc is in Leipzig
$s_0$	Nothing

now annotated with weights from  $H$  and events from  $S$ . For example, if  $a_1 = (F, G)$  had a length of 231 m, it would be annotated with  $w(a_1) = (231, 1, \{233\})$  and  $e(a_1) = city$ . The arc  $a_2 = (I, H)$  represents a footpath and is annotated with  $w(a_2) = (0, 0, \emptyset)$  and  $e(a_2) = tra$ .

Hence, to construct a conditional fare network for MDV, we now only need to give the transition functions for  $\mathcal{T}$ . Let  $s \in S$  and  $h = (h_{dist}, h_{stop}, h_{zone}) \in H$ . Then, the ticket transition function  $\Gamma$  of  $\mathcal{T}$  is defined by the equations given in Box I. For example, the conditions for  $\Gamma(D_L, h, s)$  mean that we need to transition to the two zones ticket  $Z_2$  whenever we leave Leipzig and travel for more than four stops, and that we transition to the standard Leipzig ticket  $L$  whenever we transfer or surpass the four stops limit within Leipzig. In all other cases, we can stick with the discounted Leipzig ticket  $D_L$ . Note that we did not yet cover MDV's overlap areas. We discuss in Section 3.4 why it is best to address these in a preprocessing step.

We can now finally formalize the price-optimal earliest arrival problem.

**Definition 3.6 (Price-optimal earliest arrival problem (POEAP)).** Let a public transportation network be given as a directed graph  $G = (V, A)$  together with a conditional fare network  $(\mathcal{T}, \Gamma, w, e, \mu, \pi)$  and a time-dependent FIFO travel time function  $c(a) : I \rightarrow I \forall a \in A$ . Then, the price-optimal earliest arrival problem (POEAP) asks to find a Pareto-set (w.r.t. price and arrival time) of  $s, t$ -paths  $P_{s,t}^* \subseteq P_{s,t}$  in  $G$ , i.e.,

$$\forall p^* \in P_{s,t}^* \nexists p \in P_{s,t} : \pi(p) \leq \pi(p^*) \wedge c(p) \leq c(p^*) \wedge (\pi(p) < \pi(p^*) \vee c(p) < c(p^*)) \quad (1)$$

$$\forall p \in P_{s,t} \exists p^* \in P_{s,t}^* : \pi(p^*) \leq \pi(p) \wedge c(p^*) \leq c(p). \quad (2)$$

### 3.4. Some hints on modeling with CFNs

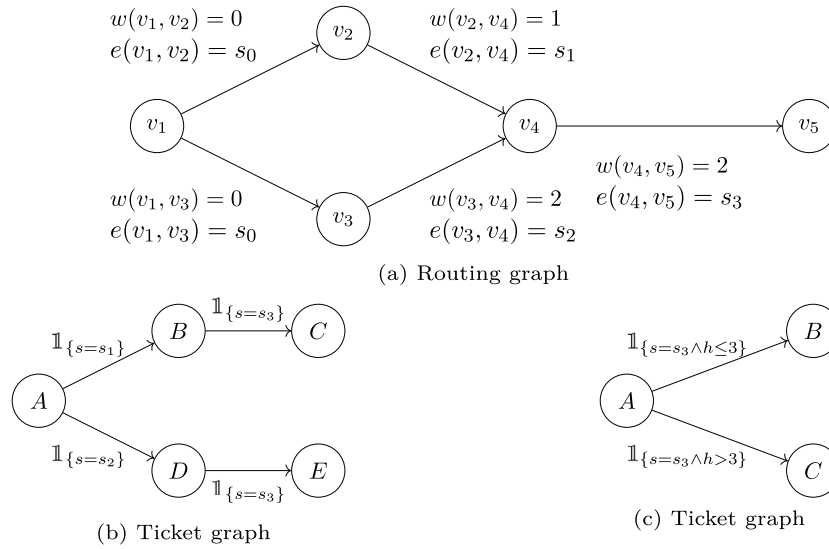
In the following, we elaborate on some common features of fare structures and how they can be modeled using conditional fare networks.

**Transfer penalties, footpaths, and surcharges.** When footpaths have no influence on the ticket, they can be modeled as arcs with a weight of  $0 \in H$  and an event  $s_0$  that cannot activate a ticket transition. This way, a footpath does not change the current fare state. The transition from a footpath to a public transportation vehicle requires some care. Assume we walk from stop  $v_0$  to  $v_1$  along arc  $a_0 = (v_0, v_1)$  to take a vehicle along  $a_1 = (v_1, v_2)$  to reach  $v_2$ . Some fare structures use the number of stops a path touches to calculate prices. Here, this number would be two. Counting a stop when relaxing  $a_0$  is a mistake if the optimal path would be to continue on foot. Counting both  $v_1$  and  $v_2$  when relaxing  $a_1$  is also wrong, since this would overcount the number of stops for every journey that reaches  $v_1$  via a vehicle. Hence, the graph model needs to be extended by splitting up stops into vertices for every route and a vertex that is connected to footpaths. These vertices are then connected via transfer arcs and boarding arcs. Placing weights and events different from  $0$  and  $s_0$  on transfer arc allows us to make the applicable ticket dependent on the number of transfers, while events on arcs representing boarding can be used to model surcharges for the boarded route. For more details on how to build these expanded graphs, we refer to Disser et al. (2008).

**Overlap areas.** Some fare structures that are based on fare zones contain overlap areas. Stations in an overlap area can be counted as part of either of its neighboring zones, whichever is cheapest for the customer. This is meant to mitigate sharp price increases for short journeys at fare zone borders. MDV uses them as well as several other German railway companies (e.g., Verkehrsverbund Bremen/Niedersachsen GmbH (Verkehrsverbund Bremen/Niedersachsen GmbH, 2023)).

At a first glance, one might be tempted to represent overlap areas as tickets in the ticket graph. A label propagated along a path starting in an overlap area then keeps this ticket until a regular fare zone is picked up along the path and transitions in the zone ticket for this fare zone. This approach, however, becomes cumbersome when several overlap areas border each other. In this case, a ticket for each combination of overlap areas needs to be introduced.

Alternatively, overlap areas can be incorporated by label duplication: Assume an overlap area neighbors  $n$  fare zones. We associate each arc



**Fig. 3.** Example of a routing graph (a) with two possible conditional fare networks (b) and (c). For both networks, the underlying partially ordered monoid is  $(\mathbb{R}, +, \leq)$ , the fare events are  $S = \{s_0, s_1, s_2, s_3\}$  and the initial fare state for all vertices  $v_i$  with  $i = 1, \dots, 5$  is  $\mu(v_i) = (A, 0)$ . We set prices for the tickets as  $\pi(A) = 0, \pi(B) = 2, \pi(C) = 3, \pi(D) = 1$  and  $\pi(E) = 5$ . The value of the transition function  $\Gamma$  for a given weight  $h$  and event  $s$  is given via indicator functions on the fare arcs. Using the ticket graph (b), the upper  $v_1, v_5$ -path yields ticket  $C$ , while the lower path yields ticket  $E$ . Using ticket graph (c), the upper path yields ticket  $B$ , the lower path yields ticket  $C$ .

$a$  whose  $head(a)$  represents a stop in the overlap area with  $n$  different weights, one for each fare zone it could possibly be part of. When settling the vertex in a shortest path search, the current fare state is updated once for each weight, thereby creating  $n$  new labels. This, however, leads to an increased need for dynamic memory allocation, which should be avoided.

Hence, we propose simply route duplication as the most convenient model for overlap areas. Whenever a route of the timetable contains a stop  $v$  in an overlap area neighboring  $n$  fare zones, we simply introduce  $n$  routes each with a single fare zone at  $v$ . In the routing graph, this corresponds to introducing parallel arcs with each storing a different fare zone in its weight. To avoid creating unnecessary duplicates, this is done block-wise, i.e., only for each consecutive sequence of stops along a route that are in the same overlap area. Hence, overlap areas are taken care of in a preprocessing step and are not represented in the conditional fare network.

#### 4. CFNs in routing algorithms

Label-setting MOSP algorithms rely on dynamic programming and the subpath optimality condition (Berger and Müller-Hannemann, 2009). That is, every subpath of an optimal  $s, t$ -path is in itself an optimal path. For POEAP, when comparing paths in  $G$  naively by the optimal price  $\pi$ , the subpath optimality condition is usually violated. Consider taking a local detour to avoid a fare zone: Later on, travelers may be forced to cross the zone due to the infrastructure, turning the locally dominant detour into a suboptimal choice. On the other hand, a locally dominated subpath might still lead to an optimal  $s, t$ -path. This type of problem persists in CFNs: the transition between tickets depends on the weights and events already collected, but also on the structure of the reachable ticket graph. Example 4.1 highlights that problems can already arise even in simple cases.

**Example 4.1 (Label dominance in Fig. 3).** Consider the routing graph (a) together with the conditional fare network (b). Examining the paths  $p_1 = (v_1, v_2, v_4)$  and  $p_2 = (v_1, v_3, v_4)$ , we find their respective fare states are  $f(p_1) = (B, 1)$  and  $f(p_2) = (D, 2)$ . Extending them by  $v_5$  to  $p'_1$  and  $p'_2$  yields  $f(p'_1) = (C, 3)$  and  $f(p'_2) = (E, 4)$ . Comparing fare states by price would indicate that  $p_1$  could be pruned at  $v_4$  since  $\pi(B) > \pi(D)$ . This is a suboptimal choice as  $p'_1$  dominates  $p'_2$  since  $\pi(C) < \pi(E)$ . Hence, price cannot be used as dominance criterion for fare states. A

natural alternative would be to use the partial order defined by paths in the ticket graph, instead. A ticket  $\tau_1$  then dominates a ticket  $\tau_2$  if there is a  $\tau_1, \tau_2$ -path. This would render the tickets  $B$  and  $D$  and the tickets  $C$  and  $E$  mutually incomparable. The idea, however, comes with problems of its own. To see this, consider the conditional fare network (c). At  $v_4$ , we have  $f(p_1) = (A, 1)$  and  $f(p_2) = (A, 2)$  and hence both paths are equivalent and it would be sensible to keep only one of them based on the relation between  $w(f(p_1))$  and  $w(f(p_2))$ . By relaxing  $(v_4, v_5)$ , we obtain  $f(p'_1) = (B, 3)$  and  $f(p'_2) = (C, 4)$ , which are incomparable, i.e., the fare states of  $p'_1$  and  $p'_2$  diverged from comparable to incomparable. Consequently, any dominance rule pruning either  $p_1$  or  $p_2$  would be defective.

To mitigate these and similar problems, we might assume a general incomparability of fare states. This comes down to enumerating all  $s, t$ -paths and simply sorting them by price. However, in a sensibly designed fare structure, it is usually clear which ticket is better, and taking a cheaper subpath should usually not turn out more expensive overall. In the remainder of this section, we propose a more tailored approach. It bases domination rules on path relationships but adds exceptions to cover cases in which it is not safe to do so.

##### 4.1. Dominance for fare states

We want to define a partial order for fare states that restores subpath optimality while not relaxing dominance too generously.

To do so, we partition the ticket set  $T$  into three disjoint *comparability groups*:  $C_F$  (full comparability),  $C_P$  (partial comparability),  $C_N$  (no comparability). Based on the partition  $C = (C_F, C_P, C_N)$ , we define the partial order.

**Definition 4.1 (Comparability of fare states).** Let  $f_1 = (\tau_1, h_1), f_2 = (\tau_2, h_2)$  be fare states. We say  $f_1 \leq_C f_2$  if and only if  $\tau_1 \notin C_N, h_1 \leq h_2$  and

$$\tau_1 = \tau_2 \quad \text{if } \tau_1 \in C_P \quad (3)$$

$$\exists \tau_1, \tau_2\text{-path in } \mathcal{T} \quad \text{if } \tau_1 \in C_F. \quad (4)$$

If  $f_1 \leq_C f_2$  and either  $h_1 < h_2$  or  $\tau_1 \neq \tau_2$ , we say that  $f_1$  is *strictly less* than  $f_2$ , i.e.,  $f_1 <_C f_2$ .



We denote by  $P_{s,t}^f$  the set of all Pareto-optimal paths with respect to  $\leq_C$ , i.e.,

$$p^* \in P_{s,t}^f \Rightarrow \nexists s, t\text{-path } p : f(p) <_C f(p^*). \quad (5)$$

We call paths in  $P_{s,t}^f$  *state-optimal*. They form a superset of the set of price-optimal paths. MOSP algorithms on graphs with weights from partially ordered monoids rely on the monoid operation being translation-invariant with respect to the partial order. Similarly, for CFN's, we need the update function to be monotone along all arcs  $a \in A$ , i.e.,

$$\forall f_1, f_2 \in F : f_1 \leq_C f_2 \implies \forall a \in A : \text{Up}(f_1, a) \leq_C \text{Up}(f_2, a). \quad (6)$$

This condition is enough to ensure that a weaker form of subpath optimality holds.

**Proposition 4.1 (Weak subpath optimality).** *Let  $G = (V, A)$  be a routing network and  $\mathcal{N} = (\mathcal{T}, \Gamma, w, e, \mu, \pi)$  be its conditional fare network. Let  $p^* \in P_{s,t}^f$  be a state-optimal  $s, t$ -path in  $G$  for some  $s, t \in V$ . Then, there is a path  $p' = (s = v_0, v_1, \dots, v_{n-1}, v_n = t) \in P_{s,t}^f$  with  $f(p^*) = f(p')$ , such that every subpath  $p'' = (v_0, \dots, v_l)$ ,  $l < n$ , of  $p'$  is a state-optimal  $v_0, v_l$ -path.*

**Proof.** Let  $p^* = (s = v_0, v_1, \dots, v_{n-1}, v_n = t) \in P_{s,t}^f$  be a state-optimal  $s, t$ -path with fare states  $(f_0, \dots, f_n)$ . Assume there is a  $s, t$ -path  $\tilde{p} = (s = u_0, u_1, \dots, u_{l-1}, u_l = t)$  with fare states  $(g_0, g_1, \dots, g_l)$  and  $g_0 = f_0$ . By (6), we can choose  $\tilde{p}$  to be a simple path. Let  $k$  be the largest integer such that  $v_{n-k} = u_{l-k}$ , i.e., the paths  $(v_{n-k}, \dots, v_n)$  and  $(u_{l-k}, \dots, u_l)$  are equal. Now assume  $\tilde{f}_{l-k-1} <_C f_{n-k-1}$ . By definition,  $f_{n-k} = \text{Up}(f_{n-k-1}, (v_{n-k-1}, v_{n-k}))$  and  $g_{l-k} = \text{Up}(g_{l-k-1}, (u_{l-k-1}, u_{l-k}))$ . We apply (6) to obtain  $\tilde{f}_{l-k} \leq_C f_{n-k}$ . By repeating the process for  $i \in \{k-1, \dots, 0\}$ , we find  $g_i \leq_C f_n$ . Since  $p^*$  was state-optimal, it follows that  $g_l = f_n$ , and consequently,  $\tilde{p}$  is also state-optimal. Since the number of paths in  $G$  is finite, we can repeat this procedure to find the path  $p'$ .  $\square$

Proposition 4.1 does not imply that every subpath of a state-optimal path is state-optimal. We can, however, discard all state-optimal paths without this property since a path with an equal fare state still remains in  $P_{s,t}^f$ . Hence, label-setting MOSP algorithms can still be applied.

#### 4.2. The comparability partition

In choosing  $C_F$ ,  $C_P$  and  $C_N$ , there is some degree of freedom. We want  $C_F$  to be as big and  $C_N$  as small as possible while still fulfilling (6). It is clear that the best choice does not only depend on the ticket graph  $\mathcal{T}$  and the transition function  $\Gamma$ , but also on  $G$  and the arc weights and events. Such an approach, however, mostly likely requires extensive computations on  $G$ . We propose a solution that depends only on  $\mathcal{T}$  and  $\Gamma$  and needs no recomputation when changes in the routing network occur.

First, we introduce some notation. If there is a directed path in  $\mathcal{T}$  between  $\tau_1, \tau_2 \in T$ , we write  $\tau_1 \rightarrow \tau_2$ . This includes the case  $\tau_1 = \tau_2$ . The *reach*  $R(\tau)$  of a vertex  $\tau \in T$  is the subgraph induced by all vertices reachable from  $\tau$ , i.e.,  $R(\tau) := \mathcal{T}[\{k \in T : \tau \rightarrow k\}]$ .

**Definition 4.2 (No-overtaking property).** Let  $\tau \in T$  be a ticket. We say its reach  $R(\tau)$  has the *no-overtaking property* if for all tickets  $k, l \in R(\tau)$  with  $k \rightarrow l$  and  $(h, s) \in W$  it holds that

$$\forall \bar{h} \in H : h \leq \bar{h} \implies \Gamma(k, h, s) \rightarrow \Gamma(l, \bar{h}, s). \quad (7)$$

The no-overtaking property bears some resemblance to the FIFO (first-in, first-out) property: A worse fare state, i.e., a worse weight or ticket, cannot give rise to a better fare state when relaxing the same arc in the routing graph. Note that the no-overtaking property has to be fulfilled not only for the neighborhood of a ticket  $\tau$  but for the reach  $R(\tau)$ . Subgraphs with the no-overtaking property allow for the strictest domination rules. We use them as comparability group  $C_F$ .

**Definition 4.3 (Comparability partition).** Let  $G = (V, A)$  be a routing network with a CFN  $\mathcal{N} = (\mathcal{T}, \Gamma, w, e, \mu, \pi)$ . We define

$$C_F := \{\tau \in T : R(\tau) \text{ traceable and has the no-overtaking property}\} \quad (8)$$

$$C_P := \{\tau \in T \setminus C_F : \forall k \in R(\tau) \forall s \in S \forall h_1, h_2 \in H : \Gamma(k, h_1, s) = \Gamma(k, h_2, s)\} \quad (9)$$

$$C_N := \{\tau \in T \setminus (C_F \cup C_P)\}. \quad (10)$$

It is not enough to fulfill (7) for  $\tau \in T$  to be in the set  $C_F$ . Its reach  $R(\tau)$  has also to be *traceable*, i.e., contain a Hamiltonian path. This condition is needed to avoid the divergence seen in Example 4.1. If a ticket has non-traceable reach or does not have the no-overtaking property, it is placed in  $C_P$ . For tickets  $\tau \in C_P$ , the transition functions of tickets  $k \in R(\tau)$  must be independent of  $(H, +, \leq)$ . This, again, is necessary to ensure that comparable fare states do not diverge in an incomparable state after an update, i.e., all tickets that can be reached from a ticket in  $C_F$  themselves need to be in  $C_F$ . All remaining tickets are added to  $C_N$ . Fare states containing tickets from  $C_N$  can never be dominated.

**Example 4.2 (Dominance for MDV Fares).** In the graph in Fig. 2, all nodes have traceable reach, and it is easy to verify that the no-overtaking property does indeed hold for all tickets. Hence, we can set the comparability partition to  $C_F = T$ ,  $C_P = C_N = \emptyset$ .

**Example 4.3 (Dominance for Example 4.1).** For Ticket Graph (b), we have  $C_F = \{B, C, D, E\}$ ,  $C_P = \{A\}$  and  $C_N = \emptyset$ . For Ticket Graph (c), we have that  $C_F = \{B, C\}$ ,  $C_P = \emptyset$  and  $C_N = \{A\}$ .

**Proposition 4.2 (Monotonicity of the comparability partition).** *The partial order  $\leq_C$  defined by Definitions 4.1 and 4.3 fulfills the monotonicity condition (6).*

**Proof.** Let  $a \in A$  and  $f_1, f_2 \in F$  such that  $f_1 \leq_C f_2$ . For  $i \in \{1, 2\}$ , we write  $g_i := \text{Up}(f_i, a)$ , i.e.,  $w(g_i) = w(f_i) + w(a)$  and  $\tau(g_i) = \Gamma(\tau(f_i), w(g_i), e(a))$ . By positivity of the monoid  $(H, +, \leq)$ ,  $w(f_1) \leq w(f_2)$  directly implies  $w(g_1) \leq w(g_2)$ . It remains to show that  $\tau(g_1) \rightarrow \tau(g_2)$ . To do so, we need to distinguish the cases  $\tau(f_1) \in C_P$  and  $\tau(f_1) \in C_F$ .

First, assume that  $\tau(f_1) \in C_P$  and hence  $\tau(f_1) = \tau(f_2)$ . By the definition of  $C_P$ , we obtain

$$\tau(g_1) = \Gamma(\tau(f_1), w(g_1), e(a)) = \Gamma(\tau(f_2), w(g_2), e(a)) = \tau(g_2).$$

Thus,  $\tau(g_1) = \tau(g_2)$ . Note that the definitions of  $C_P$  and  $C_F$  imply that  $\tau(g_1) \in C_P \cup C_F$  since  $\tau(g_1) \in R(\tau(f_1))$  and hence  $g_1 \leq_C g_2$ . Now, assume  $\tau(f_1) \in C_F$ . Note that  $R(\tau(f_1)) \subset C_F$ . This allows us to apply (7) to obtain

$$\tau(g_1) = \Gamma(\tau(f_1), w(g_1), e(a)) \rightarrow \Gamma(\tau(f_2), w(g_2), e(a)) = \tau(g_2),$$

which concludes the proof.  $\square$

Propositions 4.1 and 4.2 allow us to apply label-setting MOSP algorithms to POEAP using the comparability partition from Definition 4.3. However, we obtain only the set of state-optimal paths. It remains to show that this set contains the cheapest path.

**Proposition 4.3 (Correctness).** *Let  $\pi^* := \min_{P_{s,t}^f} \pi(p)$ . Then, there is at least one  $s, t$ -path  $p^*$  with  $\pi^* = \pi(p^*)$  and  $p^* \in P_{s,t}^f$ .*

**Proof.** Consider a path  $p \in P_{s,t}^f$  with  $\pi(p) = \pi^*$ . If there is  $p' \in P_{s,t}^f$  with  $\tau(f(p')) = \tau(f(p))$ , we are done. If not, all such paths must be dominated w.r.t. to  $<_C$  and hence there is a path  $p' \in P_{s,t}^f$  with  $\tau(f(p')) \rightarrow \tau(f(p))$ . This implies  $\pi(\tau(f(p'))) \leq \pi(\tau(f(p)))$  and hence a path of the same price as  $p$  is present in  $P_{s,t}^f$ .  $\square$

## 5. Price-optimal RAPTOR

In this section, we will discuss how to use conditional fare networks to implement a price-optimal version of the multi-criteria RAPTOR algorithm (abbr. McRAP) (Delling et al., 2015). Since RAPTOR implicitly optimizes the number of trips, we obtain an multi-criteria algorithm that optimizes for travel time, number of trips and price. So far, we have presented our framework in a graph-based context. RAPTOR, however, does not use a graph model but works directly on the timetable. The adaption for RAPTOR is straightforward.

This section is structured as follows. A review of the RAPTOR algorithm is provided in Section 5.1. Then, in Section 5.2, we lay out how McRAP can be modified to use conditional fare networks for price optimization. In Section 5.3, we briefly recap how improvements in run times can be achieved by calculating a restricted Pareto-set using the recently introduced Bounded-RAPTOR-algorithm (BMRAP) (Delling et al., 2019). The algorithm excludes all journeys that need significantly more transfers or take significantly more time than the journeys found with an (earliest arrival) RAPTOR query. Finally in Section 5.4, we introduce two speed-up techniques that are tailored to our application.

Sections 5.1 and 5.3 give succinct summaries of the RAPTOR and BMRAP algorithms. For a thorough presentation, see the original research in Delling et al. (2015) and Delling et al. (2019), respectively. We provide the pseudo-code of BMRAP in Appendix B.

### 5.1. Multi-criteria search with McRAP

We largely adhere to the notation of Delling et al. (2015) albeit with minor modifications to avoid the reuse of variables.

The RAPTOR algorithm does not use a graph model, but works directly on the timetable. A timetable is a tuple  $\mathbb{T} = (\Pi, \mathcal{P}, \mathcal{R}, \mathcal{D}, \mathcal{F})$  consisting of a period of operation  $\Pi$ , a set of stops  $\mathcal{P}$ , a set of routes  $\mathcal{R}$ , a set of trips  $\mathcal{D}$  and a set of footpaths  $\mathcal{F}$ . A stop  $p \in \mathcal{P}$  is a location where a vehicle can be boarded or exited. Each stop  $p$  has a (possibly zero) transfer time  $\eta_{ch}(p) \in \mathbb{N}$  that is applied whenever a vehicle is boarded at  $p$ . A trip  $d \in \mathcal{D}$  is a sequence of stops together with arrival and departure times  $\eta_{arr}(d, p)$  and  $\eta_{dep}(d, p) \in \Pi$ . A route is a set of trips, where all trips have the same sequence of stops and no trip overtakes another one. We denote the sets of stops and trips of a route  $r$  by  $\mathcal{P}(r)$  and  $\mathcal{D}(r)$ , respectively. Finally, a footpath  $(p_1, p_2, l) \in \mathcal{F}$  is a pair of stops  $(p_1, p_2)$  combined with a walking time  $l$ .

RAPTOR operates in rounds  $k = 1, \dots, K$  on  $\mathbb{T}$  with  $K \in \mathbb{N} \cup \{\infty\}$  the maximum number of transfers to be considered. It maintains arrival time labels  $\eta_{arr}(k, p)$  for every  $p \in \mathcal{P}$  and every round  $k \in 1, \dots, K$ . Every entry of  $\eta_{arr}$  is initialized to  $\infty$ . Each round begins with a set of marked stops. All routes touching these stops are collected and then processed in an arbitrary order. When processing a route  $r \in \mathcal{R}$ , RAPTOR begins with the first marked stop and from there iterates through the stops in order of travel. For each stop  $p \in \mathcal{P}(r)$ , RAPTOR finds the earliest trip  $d$  that can be taken at  $p$  after  $\eta_{arr}(k-1, p) + \eta_{ch}(p)$ . In the same sweep, the arrival times of  $d$  are used to update  $\eta_{arr}(k, p)$ . All stops whose arrival times improved over  $\eta_{arr}(k-1, p)$  are marked.

In a second step, a footpath search is performed. All footpaths  $(p_1, p_2, l)$  starting at a marked stop are processed in an arbitrary order, and labels in  $\eta_{arr}(\cdot, p_2)$  are updated accordingly. Again, each stop  $p$  with an improved arrival time  $\eta_{arr}(k, p)$  is marked for the next round. Note that this requires the footpath set to be transitively closed. The algorithm terminates after  $K$  rounds or when no more stops can be marked.

RAPTOR can be modified slightly to allow for multi-criteria search. Instead of only the arrival times  $\eta_{arr}(k, p)$ , McRAP now maintains a label bag  $B_k(p)$  for every stop  $p$  and round  $k$ . Each label  $L \in B_k(p)$  contains an entry for every optimization criterion. When processing a route  $r \in \mathcal{R}$  at a starting stop  $p_s$ , a route bag  $B_r$  is created, and all labels from  $B(k-1, p_s)$  are updated with  $\eta_{ch}(p)$  and copied into  $B_r$ . Each label in  $B_r$  is associated with a trip  $d \in \mathcal{D}(r)$ . At each stop  $p \in \mathcal{P}(r)$  after  $p_s$ , McRAP updates all labels in  $B_r$ , merges  $B_r$  into  $B_k(p)$  and finally merges all labels from  $B(k-1, p)$  into  $B_r$  and assigns a trip to them. In each step, dominated labels are removed.

### 5.2. Using conditional fare networks in McRAP

Adapting McRAP to incorporate fares is now fairly straightforward but requires several modifications to  $\mathbb{T}$ . For every trip  $d \in \mathcal{D}$ , we additionally store two pairs of weights and events for each stop  $p$ . The first pair  $w_1(d, p)$ ,  $e_1(d, p)$  is considered when reaching the stop  $p$  while iterating along  $d$ . The second pair  $w_2(d, p)$ ,  $e_2(d, p)$  is picked up when boarding  $d$  at  $p$ . This distinction is necessary since there is no direct equivalent in  $\mathbb{T}$  to the transfer arcs used in the graph-based setting to model, e.g., surcharges. Furthermore, the definition of a route needs a slight adjustment: A route  $r \in \mathcal{R}$  is a set of trips where all trips have the same sequence of stops and the same weights and events, i.e., for  $i = 1, 2$ ,  $w_i(d_1, p) = w_i(d_2, p)$  and  $e_i(d_1, p) = e_i(d_2, p)$  for all  $d_1, d_2 \in \mathcal{D}(r)$  and  $p \in \mathcal{P}(r)$ , and no trip overtakes another one.

Now, a label  $L = (\eta, f)$  in a label bag  $B_k(p)$  consists of an arrival time  $\eta$  and a fare state  $f$ . Labels in route bags additionally hold the current trip  $d$ . When updating a label  $L = (\eta, f, d)$  from route bag  $B_r$  at a stop  $p$ , the arrival time  $\eta$  is updated to  $\eta_{arr}(d, p)$  and the fare state is updated with  $w_2(d, p)$  and  $e_2(d, p)$  as in Definition 3.4. Update steps that are associated with transfers are performed whenever labels are merged into  $B_r$ . Here,  $w_2(d, p)$  and  $w_1(d, p)$  are used for updating. Dominance of labels is checked according to the theory developed in Section 4 while also taking arrival times into account. Since walking is usually free of charge, fare states do not need to be updated in the footpath stage. Hence, footpaths are also not enriched with fare information.

Using McRAP, we obtain the Pareto-set  $\mathcal{J}^f$ , optimizing for arrival time, number of trips and fare state. The smaller set  $\mathcal{J}^* \subseteq \mathcal{J}^f$ , optimizing for price instead of fare state, can be calculated in a post-processing step.

### 5.3. Restricted Pareto-sets

By design of fare structures, the cheapest path is often among the fastest, as detours are penalized by increases in price, arrival time and transfers. At other times, a negligible reduction in price might be achievable at the expense of a significant increase in travel time. Such journeys are unlikely to be chosen by a traveler. Hence, it appears beneficial for a price-optimal search to prune all labels that are worse by a certain margin (regarding both arrival time and transfers) than the results of a normal RAPTOR query. This is achieved by the following pruning schemes, first introduced by Delling et al. (2019) for general multi-objective search with RAPTOR.

Let  $\mathcal{J}_A$  be the Pareto-set of all anchor journeys found with a RAPTOR query, i.e., optimizing only for arrival time and number of transfers. We denote the arrival time of a journey  $J$  by  $\eta_{arr}(J)$  and its number of trips by  $\text{tr}(J)$ . We aim to calculate a restricted Pareto-set  $\mathcal{J}_R$  with  $\mathcal{J}_A \subseteq \mathcal{J}_R \subseteq \mathcal{J}^f$  of journeys that do not have a significantly higher arrival time or number of trips than some journey from  $\mathcal{J}_A$ . Let  $\sigma_{arr} \in \mathbb{R}^+$  and  $\sigma_{tr} \in \mathbb{N}^+$  be the maximal acceptable slacks for arrival time and number of transfers, respectively. Then, we define

$$\mathcal{J}_R := \{J \in \mathcal{J}^f \mid \exists J_A \in \mathcal{J}_A \text{ such that } \eta_{arr}(J) \leq \eta_{arr}(J_A) + \sigma_{arr} \text{ and } \text{tr}(J) \leq \text{tr}(J_A) + \sigma_{tr}\}.$$

To obtain a two-stage pruning scheme, we can first run a normal RAPTOR query. The labels obtained in this first stage can then be used to prune the multi-criteria search. Let  $\eta_k$  be the optimal arrival time at the target stop  $p_t$  in round  $k$  of the first stage (computed with RAPTOR). During round  $k$  of McRAP, we prune every label that has an arrival time  $\eta$  with  $\eta > \eta_k + \sigma_{arr}$ . This pruning scheme is called Target-BMRAP. Note that Target-BMRAP works on the assumption that  $\sigma_{tr} = \infty$ . Also, the bound is not tight for all stops other than  $p_t$ .

The set  $\mathcal{J}_R$  can be computed with the more involved Tight-BMRAP. Here, three rounds are performed. As for Target-BMRAP, we first perform a normal RAPTOR search, obtaining the Pareto-set  $\mathcal{J}_A$ . Then, multiple reverse RAPTOR queries are performed to build bounds at all

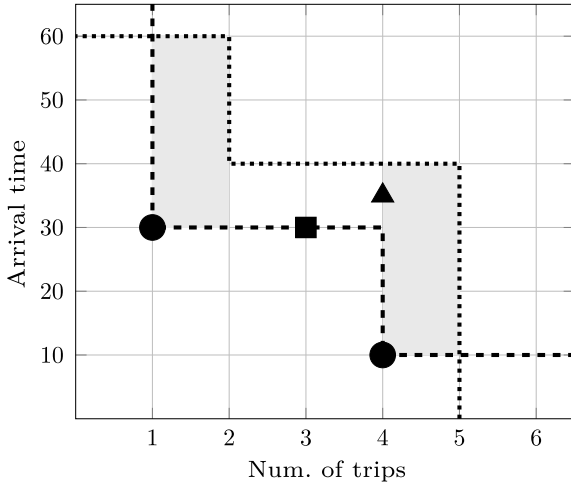


Fig. 4. Illustration of the solution space of a Tight-BMRAP search with  $\sigma_{arr} = 30$  min and  $\sigma_{tr} = 1$ . The circle marks represent the anchor journeys from  $J_A$ . Tight-BMRAP prunes all journeys to the right of the dotted line spanned by those journeys. The area to the left of the dashed line contains no Pareto-optimal journeys. The journeys from  $J^f$  that fall into the area enclosed by the dashed and dotted lines form  $J_R$ . The light gray area forms  $\tilde{J}_R$ . Note that the journey marked by the square mark is in  $J_R$  but not in  $\tilde{J}_R$ . The journey represented by the triangle mark is in  $\tilde{J}_R$  even though it uses more trips and has a later arrival time.

stops  $p \in P$ . The third stage is the actual McRAP round using the previously computed bounds for pruning.

In the following, we describe the second and third stages in more detail. Let  $m := K + \sigma_{tr}$ , where  $K$  is the maximum number of trips in any journey  $J_A \in J_A$ . A backward RAPTOR search with starting time  $\eta_{arr}(J_A) + \sigma_{arr}$  and  $n_{J_A} := \text{tr}(J_A) + \sigma_{tr}$  rounds is performed for every journey  $J_A \in J_A$ . The backward search works on departure times instead of arrival times and transfer times are not applied when boarding a vehicle but instead when disembarking.

Each backward search computes latest departure times  $\eta_{dep}(J_A, k, p)$  such that  $p_i$  can still be reached earlier than  $\eta_{arr}(J_A) + \sigma_{arr}$  while using at most  $k$  more trips onward from  $p$ . It is possible that  $\eta_{dep}(J_A, k, p)$  remains at its initialization value of  $-\infty$ . Using  $\eta_{dep}(J_A, n_{J_A} - k, p)$  for pruning labels in round  $k$  of a forward McRAP search computes

$$\{J \in J^f \mid \eta_{arr}(J) \leq \eta_{arr}(J_A) + \sigma_{arr} \text{ and } \text{tr}(J) \leq \text{tr}(J_A) + \sigma_{tr}\}.$$

Carefully overlapping the labels  $\eta_{dep}(J_A, k, p)$  results in a set of labels  $\eta_{dep}(k, p)$  with  $k = 1, \dots, m$ , i.e.,

$$\eta_{dep}(k, p) := \max_{\substack{J_A \in J_A: \\ k \geq m - n_{J_A}}} \{\eta_{dep}(J_A, k - m + n_{J_A}, p)\}.$$

The third stage is now a normal McRAP search enriched with two pruning rules. Let  $k$  be the current round.

- Labels  $L$  are not merged into to  $B_k(p)$  at stop  $p$  if  $\eta_{arr}(L) > \eta_{dep}(m - k, p)$ ,
- Labels  $L$  can be removed from  $B_r$  at stop  $p$  if  $\eta_{arr}(L) > \eta_{dep}(m - k + 1, p) + \eta_{ch}(p)$ .

The summand  $\eta_{ch}(p)$  in the second rule is required since it was factored into  $\eta_{dep}(m - k + 1, p)$  and no transfer is performed. This third stage computes exactly  $J_R$ .

The definition we gave for  $J^f$  differs from the one given by Delling et al. (2019). The original authors assign each journey  $J \in J^f$  its anchor journey  $J_A \in J_A$ , that has the maximum number of trips smaller

or equal to  $\text{tr}(J)$ . They then define the restricted Pareto-set  $\tilde{J}_R$  to be

$$\begin{aligned} \tilde{J}_R &:= \{J \in J^f \mid \text{its anch. journ. } J_A \in J_A \text{ has } \eta_{arr}(J) \leq \eta_{arr}(J_A) + \sigma_{arr} \text{ and } \text{tr}(J) \leq \text{tr}(J_A) + \sigma_{tr}\} \\ &\subseteq \{J \in J^f \mid \exists J_A \in J_A \text{ such that } \eta_{arr}(J) \leq \eta_{arr}(J_A) + \sigma_{arr} \text{ and } \text{tr}(J_A) \leq \text{tr}(J) \leq \text{tr}(J_A) + \sigma_{tr}\} \\ &\subseteq J_R. \end{aligned}$$

Note that Tight-BMRAP as presented here and in Delling et al. (2019) does, in fact, calculate  $J_R$  and not  $\tilde{J}_R$  and that  $J_R$  is a more interesting set to compute as it does not seem beneficial to impose lower bounds on the objective function. The difference between  $J_R$  and  $\tilde{J}_R$  is visualized in Fig. 4. The first fact can be seen easily by considering the case where  $\sigma_{arr} = \infty$ . Then, the pruning scheme does not prune anything but all journeys with more than  $m$  trips. Specifically, no lower bound on the number of trips is applied.

#### 5.4. Speed-up techniques

**Price-based target pruning.** In RAPTOR as well as Dijkstra's algorithm, it is possible to use *target pruning* (Delling et al., 2015) to delete labels that are worse than the labels that have already been found at the target stop. Naturally, the same speed-up technique is also possible for our algorithm. Moreover, we need not use  $\leq_C$  to compare fare states. Since the labels at the target stop are never updated and the price function  $\pi$  is non-decreasing, a partial journey already more expensive than the incumbent cheapest journey cannot be price-optimal. Hence, in round  $k$  of McRAP, we can prune all labels with a fare state  $f$  with  $\pi(\tau(f)) \geq \pi^*$ , with  $\pi^*$  being the best price at the target stop with at most  $k$  trips. We refer to this technique as *Price-Based Target Pruning (PTP)*.

**Fare-specific speed-ups.** Certain dimensions in  $(H, +, \leq)$  might only be relevant for some tickets in  $T$ . For example, many short-distance tickets depend on the number of stops visited while this number is irrelevant for all other tickets that can be reached from that ticket. We can therefore alter the comparison operator  $\leq_C$  for those tickets to ignore the number of stops. Hence, more labels become comparable, which results in a smaller Pareto-set  $J^{fss}$  with  $J^* \subseteq J^{fss} \subseteq J^f$ . When using Tight-BMRAP, this results in a set  $J_R^{fss}$  with  $J_R^* \subseteq J_R^{fss} \subseteq J_R^f$ . We refer to this technique as *fare-specific speed-up (FSS)*.

## 6. Computational results

We implemented the McRAP algorithm in C++17 compiled with gcc 9.3.0 and -O3 optimization. All tests were conducted on Dell Poweredge M620 machines with 64 GB of RAM. While the general structure of the MDV fare structure is captured in our model, our computations deviate from the prices charged by MDV in the following two cases: A list of relations, that are, contrary to the general rules, not eligible for the short-distance discount, is considered. Moreover, stops and fare zones that a route passes through without stopping are not represented in the available data and therefore cannot be considered.

Our dataset was built from the publicly available timetable data (Mitteldeutscher Verkehrsverbund GmbH, 2019b) and fare structure of MDV (Mitteldeutscher Verkehrsverbund GmbH, 2019a). Structured fare data is not public and was obtained separately via InfraDialog GmbH. We extracted a timetable spanning two days from July 1, 2019 to July 2, 2019. The resulting timetable contains 4371 stops, 36670 trips, 5576 routes, and 845 footpaths. This original footpath set was not transitively closed. Since RAPTOR requires a transitively closed footpath set (Delling et al., 2015), we computed its transitive closure and obtained 1029 footpaths. We then chose a test set of 5000 origin-destination pairs (OD pairs) uniformly at random from the set of stops. After removing all OD pairs that were not connected in the time interval starting at 08:00 a.m. on July 1, 2019, a total of 4964 OD pairs remained.

**Table 3**

Computational results. Evaluation of different RAPTOR variants on the MDV dataset. All experiments were conducted with a maximum of seven rounds. For each algorithm, the table reports the optimization criteria, the speed-up techniques employed, and the arrival and trip slacks, if applicable. We report the number of scanned routes (*#Scan*), the average running time (*Time*), the number of rounds performed (*#Rounds*), the number of journeys found (*#Jn.*) and the number of journeys that are dominated w.r.t. to the price (*#PJn.*). For each result, both the average and standard deviation are reported. The algorithms in rows 2–17 run at least one RAPTOR and exactly one McRAPquery. In this case, *#Scan* and *Time* are only given for the McRAP run, while the running time is summed up overall RAPTOR and McRAP invocations.

	Criteria				Speed-Up		Slack		#Scan		Time [ms]		#Rounds		#Jn.		#PJn.	
	Trips	Time	Zones	Fare	PTP	FSS	arr	Trip	Avg.	Sd.	Avg.	Sd.	Avg.	Sd.	Avg.	Sd.	Avg.	Sd.
RAPTOR	•	•	○	○	○	○	–	–	17 295	5883	3.27	1.31	6.64	0.79	1.53	0.66	–	–
McRAP	•	•	•	○	○	○	–	–	29 881	6118	4.675	6.172	6.92	0.48	47.00	45.95	–	–
McRAP	•	•	○	•	○	•	–	–	29 891	5025	957.54	260.73	6.97	0.26	10.05	13.27	2.62	1.43
McRAP	•	•	○	•	•	•	–	–	23 460	8263	243.09	283.08	6.86	0.55	3.00	1.80	2.62	1.43
Target-BMRAP	•	•	○	•	•	•	15	–	20 132	8502	71.16	87.10	6.72	0.72	1.72	0.82	1.66	0.74
Target-BMRAP	•	•	○	•	•	•	30	–	20 637	8412	74.86	88.75	6.75	0.70	1.80	0.89	1.72	0.80
Target-BMRAP	•	•	○	•	•	•	60	–	21 394	8311	81.45	91.57	6.78	0.66	1.95	1.00	1.85	0.89
Tight-BMRAP	•	•	○	•	○	○	15	1	3192	1963	8.44	14.46	5.17	1.08	7.52	10.28	1.63	0.72
Tight-BMRAP	•	•	○	•	○	○	30	1	3551	2071	9.40	16.47	5.17	1.08	9.06	12.29	1.68	0.78
Tight-BMRAP	•	•	○	•	○	○	60	1	4316	2384	11.97	18.17	5.17	1.08	13.29	16.02	1.78	0.86
Tight-BMRAP	•	•	○	•	○	○	15	2	4606	3167	20.79	99.60	5.84	1.07	11.90	21.05	1.66	0.74
Tight-BMRAP	•	•	○	•	○	○	30	2	5147	3154	22.75	104.12	5.90	1.03	14.09	23.87	1.71	0.80
Tight-BMRAP	•	•	○	•	○	○	60	2	6487	3359	29.47	107.53	5.97	0.99	20.98	29.70	1.83	0.88
Tight-BMRAP	•	•	○	•	•	•	15	1	2954	1828	6.22	3.14	5.17	1.08	1.68	0.79	1.63	0.72
Tight-BMRAP	•	•	○	•	•	•	30	1	3274	1919	6.69	3.29	5.17	1.08	1.75	0.85	1.69	0.78
Tight-BMRAP	•	•	○	•	•	•	60	1	3984	2209	7.82	3.67	5.17	1.08	1.88	0.97	1.78	0.86
Tight-BMRAP	•	•	○	•	•	•	15	2	4155	2905	8.11	6.28	5.47	1.08	1.71	0.81	1.66	0.74
Tight-BMRAP	•	•	○	•	•	•	30	2	4606	2888	8.76	6.32	5.53	1.0	1.78	0.88	1.71	0.80
Tight-BMRAP	•	•	○	•	•	•	60	2	5758	3055	10.68	6.73	5.62	1.05	1.93	1.00	1.83	0.88

As of 2019, the fare structure of MDV contained 56 fare zones, 17 cities with a city fare, and 30 overlap areas containing 191 stops. Overlap areas were implemented by route duplication, as lined out in Section 3.4. After route duplication, the timetable contained 49072 trips in 7835 routes. All queries were performed with a starting time of 08:00 a.m.

In a first experiment, we ran several algorithms for each OD pair: the standard RAPTOR-algorithm; a McRAP variant optimizing for the arrival time, the number of trips, and the set of fare zones; then, optimizing for the arrival time, number of trips, and the fare state, two McRAP variants as well as Target-BMRAP and Tight-BMRAP with various configurations for arrival and trip slacks. This results in a total of 19 algorithms in the first experiment. When optimizing for fare state, a postprocessing step is performed to remove all journeys that are in the Pareto-set with regard to fare state but not to price. Note that the McRAP variant for fare zones is a relaxed version of the purely monoid-based approach from Section 3.1 and provides a lower bound on its performance. The full approach was implemented but ran out of memory on most queries and is therefore omitted in the results.

The experiment was conducted with a maximum of seven rounds. This number was chosen, as we believe it represents a sufficiently generous upper bound on the maximum number of transfers a traveler is willing to undertake. The results of the experiment are reported in Table 3. A standard RAPTOR run takes on average 3.27 ms with a standard deviation of 1.31 ms. The computed Pareto-set contains 1.53 journeys on average. If only fare zones are considered as additional operation criterion, the average runtime increases to 4.67s, with a standard deviation of 6.17s. This indicates that run times of up to around 10s are not out of the ordinary.

While FSS in McRAP alone does not suffice to obtain acceptable run times, the combination of FSS and PTP produces an average run time of 243.09 ms. Although this performance is no longer prohibitive for practical application, the standard deviation remains high at 283 ms.

The Pareto-set computed with McRAP contains 3 journeys on average, of which 2.62 are also price-optimal. Turning off PTP increases the number of computed journeys to 10.5; the Pareto-set of the zone-based McRAP contains 47 journeys on average. Hence, even though all tickets of MDV are in the full-comparability set  $C_F$ , a high number of superfluous journeys will be generated when no additional techniques are employed. This effect is mainly caused by the fare zones,

as they form an only partially ordered set. The 2.62 price-optimal journeys mark an increase of 71% over the 1.53 journeys found with RAPTOR. It might, however, contain journeys with an undesirable trade-off between arrival time and number of trips and price. To obtain restricted Pareto-sets with a reasonable trade-off, we ran Target-BMRAP and Tight-BMRAP with arrival time slacks of 15 min, 30 min and 60 min and in the case of Tight-BMRAP with trip slacks of 1 or 2. Using Target-BMRAP reduces run times to between 71.16 ms and 81.45 ms while restricting the size of the Pareto-set to between 1.66 and 1.85. This corresponds to between 8.5% and 21% more journeys compared to RAPTOR. We ran Tight-BMRAP both with and without PTP and FSS. Tight-BMRAP performs reasonably well even without PTP and FSS on average with run times of up to 29.47 ms. However, a comparatively high standard deviation of up to 107.53 ms hints at high performance variability. When using both PTP and FSS run times decrease to at most 10.68 ms. Even more pronounced is the decrease in the standard deviation to at most 6.73 ms. Compared to RAPTOR there are between 6.5% and 19.6% more journeys. Consequently, conditional fare networks used within Tight-BMRAP appear well-suited to provide the user with price-optimized alternative routes while increasing run times only insignificantly.

A second experiment was conducted without an upper bound on the number of rounds (see Table 4). The table  $\eta_{arr}$  was implemented as a fixed-size array with space for 25 rounds. Since the maximum number of rounds performed was 21, no journeys were cut off due to early termination. We excluded all algorithms that had already performed poorly in the first experiment. Namely, these are the Tight-BMRAP variants without additional speed-up techniques, McRAP for fare zones and McRAP for fare states without price-based target-pruning.

For McRAP, we see a significant increase in the run time of about 60% and a even more pronounced increase in its standard deviation of 151% compared to the variant with only seven rounds. All variants of Target-BMRAP exhibit similar behavior, albeit to a lesser degree. Here, the average run time increased by about 16% and the standard deviation by about 44%. Note that for RAPTOR the run time increased by only a marginal 0.01 ms and that, when compared to RAPTOR, McRAP needs to perform significantly more rounds.

For all settings of slack variables for Tight-BMRAP, the increase of the run time remains minimal at around a millisecond. Hence, Tight-BMRAP remains highly competitive, whereas the simple McRAP



**Table 4**

Computational results. Evaluation of different RAPTOR variants on the MDV dataset. All experiments were conducted with a maximum of 25 rounds. The maximum number of rounds performed across all algorithms was 21. Hence, no query was canceled prematurely. All algorithms other than RAPTOR optimized for arrival time, number of trips and fare state. For all of those, both FFS and PTP were activated. The same performance indicators as in Table 3 are reported. For Target-BMRAP and Tight-BMRAP, #Scan and #Rounds report only on the last McRAP call whereas Time reports the overall run time.

	Slack		#Scan		Time [ms]		#Rounds			#Jn.		#PJn.	
	arr	Trip	Avg.	Sd.	Avg.	Sd.	Avg.	Sd.	Max	Avg.	Sd.	Avg.	Sd.
RAPTOR	–	–	17 496	6111	3.28	1331.9	7.31	1.35	12	1.54	0.66	–	–
McRAP	–	–	32 757	16 429	389.52	710.75	11.10	3.24	21	3.16	1.91	2.77	1.56
Target-BMRAP	15	–	23 007	12 106	82.47	125.89	8.57	2.34	17	1.73	0.84	1.68	0.77
Target-BMRAP	30	–	23 724	12 093	86.38	128.22	8.73	2.33	17	1.81	0.91	1.73	0.82
Target-BMRAP	60	–	24 980	12 186	94.53	131.83	9.00	2.35	17	1.97	1.03	1.87	0.91
Tight-BMRAP	15	1	2985	1871	6.68	3.28	5.20	1.14	9	1.69	0.81	1.63	0.73
Tight-BMRAP	30	1	3309	1973	7.19	3.44	5.20	1.14	9	1.75	0.87	1.69	0.79
Tight-BMRAP	60	1	4039	2307	8.36	3.89	5.20	1.14	9	1.89	0.96	1.79	0.89
Tight-BMRAP	15	2	4360	3145	9.02	7.27	5.52	1.16	10	1.72	0.83	1.66	0.76
Tight-BMRAP	30	2	4833	3151	9.68	7.43	5.58	1.15	10	1.79	0.90	1.72	0.81
Tight-BMRAP	60	2	6071	3378	11.87	7.97	5.67	1.16	10	1.94	1.02	1.84	0.90

implementation suffers from considerably degraded performance. It is furthermore noteworthy that in all slack settings, the multi-criteria part of Tight-BMRAP needs to both scan significantly fewer routes and perform fewer rounds than even the standard RAPTOR. This clearly speaks to the strength of the pruning scheme used in Tight-BMRAP.

## 7. Conclusion

We presented conditional fare networks, a novel framework for modeling complex fare structures of public transportation providers. It is independent of the MOSP algorithm used and can be used to solve price-optimal earliest arrival queries in real-world networks. Since fare structures are often composed of various fare strategies, this requires the optimization of several objective functions. In the MDV case study, these were the fare zones, city fares, transfers, the number of stops visited, and the length of the path in kilometers. Performing naive multi-objective queries for all these objectives results in high run times with high variance and the computation of many journeys that are not price-optimal. In contrast, using a CFN-based variant of the McRAP algorithm, we were able to mitigate these effects and reduce run times to around 400 ms on average, which we deem acceptable for commercial applications. Combining CFNs with the Tight-BMRAP algorithm reduced run times further to at most 12 ms with low variance, while still computing a reasonably sized restricted Pareto-set when choosing appropriate arrival and trip slacks. Fare structures can differ quite significantly between public transportation providers. Hence, a systematic evaluation of CFNs on other public transit networks is certainly worthwhile. As MDV operates in a largely rural area with two only medium-sized urban centers, a study of larger urban centers such as Berlin or Madrid seems especially interesting. However, while timetables are widely available, fare data is not. Especially, machine-readable mappings from stations to fare zones are generally not publicly available. When they are, the data is often incomplete and requires a significant manual polishing effort.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Ricardo Euler reports financial support was provided by German Federal Ministry of Education and Research. Ricardo Euler reports financial support was provided by German Federal Ministry of Transport and Digital Infrastructure.

## Acknowledgments

We thank MDV and InfraDialog GmbH for providing the data for this study. We owe special gratitude to our master's student Rick Grap

for implementing the Tight-BMRAP-algorithm and pointing out the difference between the sets  $\mathcal{J}_R$  and  $\mathcal{J}_R$  in Section 5.3.

## Funding

This work was supported by the Research Campus MODAL funded by the German Federal Ministry of Education and Research (BMBF) [grant number 05M20ZBM]; and the German Federal Ministry of Transport and Digital Infrastructure [grant number 19E17001C].

## Appendix A. Computational complexity and links to automata theory

In this section, we study the computational complexity of POEAP. The intractability of general multi-objective shortest path problems is well-established (Hansen, 1980). The standard argument is here, that the output might be exponential in size. Note that in POEAP the number of tickets  $|T|$  is finite. Therefore, it is always possible to find a valid set of Pareto-optimal solutions with size  $\leq |T|$ . This begs the question whether POEAP can be solved in polynomial time. The answer depends on whether the monoid  $(H, +, \leq)$  is considered as part of the encoding length.

If  $(H, +, \leq)$  is not considered part of the input then even the single-criterion version of POEAP with travel time functions  $c \equiv 0$  (denoted by  $\text{POEAP}_0$ ) is NP-hard. Blanco et al. (2016) proved finding a shortest path with respect to weights from the monoid  $(2^Z, \subseteq, \cup)$ , with fare zones  $Z$  to be NP-hard. The proof was obtained using a reduction from the minimum-color single-path problem (Broersma et al., 2005). As this is a special case of  $\text{POEAP}_0$ , the NP-hardness of  $\text{POEAP}_0$  follows immediately.

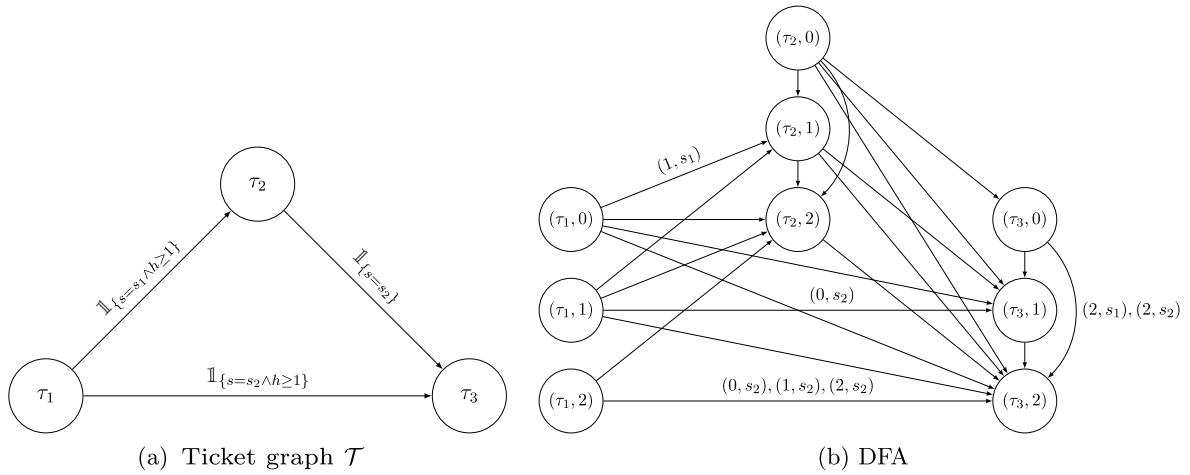
In the following, we provide an alternative reduction of the path with forbidden pairs problem to  $\text{POEAP}_0$ . Its NP-completeness was established by Gabow et al. (1976).

**Definition A.1** (Path with forbidden pairs problem). Let  $G = (V, A)$  be a directed graph and  $(a_i, b_i), i \in I$  a list of forbidden pairs. Let  $s, t \in V$ . The path with forbidden pairs problem asks whether there is a  $s, t$ -path  $p$  in  $G$  such that  $\forall i \in I : a_i \notin p \vee b_i \notin p$ . That is,  $p$  contains at most one vertex of each pair  $(a_i, b_i)$ .

**Proposition A.1** (NP-hardness of  $\text{POEAP}_0$ ). The single-criterion problem  $\text{POEAP}_0$  is NP-hard. Its canonical decision problem  $\text{POEAP}_0^{\text{DEC}}$  is NP-complete if the evaluation time of  $\Gamma$  is polynomially bounded.

**Proof.** The canonical decision problem  $\text{POEAP}_0^{\text{DEC}}$  of  $\text{POEAP}_0$  asks whether there is a path  $p$  in  $G$  with  $\pi(p) \leq k$  for some  $k \in \mathbb{Q}^+$ . We show NP-hardness by reducing the path with forbidden pairs problem





**Fig. A.5.** Transformation from CFN to DFA. The CFN  $\mathcal{N} = (\mathcal{T}, \Gamma, w, e, \mu, \pi)$  is given by the ticket graph  $\mathcal{T}$  depicted in (a). Possible ticket transitions are given as indicator functions on the arcs. The associated monoid is  $(H, +, \leq)$  with  $H := \{0, 1, 2\}$  and  $a + b := \min\{a + b, 2\}$  and  $0 \leq 1 \leq 2$  and the fare events are given as  $S = \{s_1, s_2\}$ . The definitions of  $w, e$  and  $\mu$  depend on the routing graph and are omitted in this example. Our transformation results in the DFA in (b). There is a state for each element from  $T \times H$ . Each arc depicts a possible state transformation. For each arc, there is at least one letter from  $\Sigma = H \times S$  that allows this transformation. There is no arc between  $(\tau_1, 0)$  and  $(\tau_2, 0)$  as moving from  $\tau_1$  to  $\tau_2$  in the ticket graph requires  $h \geq 1$ . There should be a loop at every state, e.g.,  $(\tau_1, 0)$  transforms into  $(\tau_1, 0)$  if letter  $(0, s_1)$  or  $(0, s_2)$  was found. We omit them in (b) to not clutter the presentation.

to POEAP. Let  $G = (V, A)$  be a directed graph,  $s, t \in V$  and  $(a_i, b_i), i \in I$ , a list of forbidden pairs. We construct a CFN  $\mathcal{N} = (\mathcal{T}, \Gamma, w, e, \mu, \pi)$  as follows. Let  $\mathcal{T} = (T, E)$  be the ticket graph with tickets  $T = \{\tau_1, \tau_2\}$  and  $E = \{(\tau_1, \tau_2)\}$ . Furthermore, we define  $(H, +, \leq)$  as follows. Let  $H = \{0, 1, 2\}^I$ . The sum of  $x, y \in H$ , is defined as  $x + y := (\min(2, x_i + y_i))_{i \in I}$ . We have  $x \leq y$  if and only if  $x_i \leq y_i$  for all  $i \in I$ . We set the ticket prices to  $\pi(\tau_1) = 0$  and  $\pi(\tau_2) = 1$ . Now, we let

$$w(v_1, v_2) = \begin{cases} e_i & v_1 = a_i \text{ or } v_2 = b_i \text{ for some } i \in I \\ 0 & \text{otherwise,} \end{cases} \quad \forall (v_1, v_2) \in A \quad (\text{A.1})$$

and

$$\Gamma(\tau_1, h, s) = \begin{cases} \tau_2 & \exists i \in I : h_i = 2 \\ \tau_1 & \text{otherwise,} \end{cases} \quad (\text{A.2})$$

where  $e_i \in \{0, 1\}^I$  is the standard unit vector with  $e_{ii} = 1$ . Let  $p$  be a simple path. W.l.o.g. we assume that the last vertex of  $p$  is not in  $\bigcup_{i \in I} \{a_i, b_i\}$ . Now, assume that  $p$  has ticket  $\tau(f(p)) = \tau_2$ . Hence, there must be an  $i \in I$  s.t.  $w(f_i) = 2$  and therefore both  $a_i \in p$  and  $b_i \in p$ . Conversely, if both  $a_i \in p$  and  $b_i \in p$ , we must have  $f_i(p) = 2$  and therefore  $\tau(f(p)) = \tau_2$ . The CFN  $\mathcal{N}$  can be built in polynomial time, as the weights  $w$  can be built in  $\mathcal{O}(|A||I|)$  time. Hence, the path with forbidden pairs problem can be polynomially reduced to POEAP<sub>0</sub>.

It remains to show that POEAP<sub>0</sub><sup>DEC</sup> is in NP. When evaluating the fare state of  $p$  as many calls to  $\Gamma$  have to be performed as there are arcs in  $p$ . Every simple  $s, t$ -path has clearly  $\leq |A|$  edges. Hence, the overall evaluation time of  $\Gamma$  for  $p$  is polynomially bounded. Finally, finding  $\pi(p)$  from  $f(p)$  requires a simple table-lookup. Hence, checking  $\pi(p) \leq k$  can be done in polynomial time.  $\square$

Now, assume that the underlying monoid  $(H, +, \leq)$  is of finite size and consider  $|H|$  a part of the encoding length. Then, POEAP can be solved in polynomial time using techniques already used by Barrett et al. (2000) for the regular language constrained shortest path problem. We begin with a short recapitulation of crucial results from automata (Hopcroft and Ullman, 1979).

**Definition A.2 (Deterministic finite automaton).** A *deterministic finite automaton* (DFA) is a 5-tuple  $(Q, \Sigma, \delta, q_0, F)$ , where  $Q$  is a finite set of states,  $\Sigma$  is a finite input alphabet,  $q_0 \in Q$  is the initial state,  $F \subseteq Q$  is the set of final states and  $\delta : Q \times \Sigma \rightarrow Q$  is the transition function.

The words accepted by a DFA are exactly the words of a regular language. Hence, we can use DFAs to define the regular language constrained shortest path problem (REG-ShP) (Barrett et al., 2000). In this problem, each arc  $a \in A$  of a directed graph  $G = (V, A)$  is associated with a letter  $\sigma(a) \in \Sigma$ . The state of a path  $p = (v_0, \dots, v_n)$  is then recursively defined via

$$q((v_0, \dots, v_i)) := \delta(q((v_0, \dots, v_{i-1})), \sigma(v_{i-1}, v_i)) \\ q(v_0) := q_0.$$

**Definition A.3 (Regular-language constrained shortest path problem (REG-ShP)).** Let a directed graph  $G = (V, A)$ , weights  $c : A \rightarrow \mathbb{Q}^+$ , a DFA  $(Q, \Sigma, \delta, q_0, F)$ , letters  $\sigma(a), a \in A$ , a source  $s \in V$  and a destination  $t \in V$  be given. Find a shortest  $s, t$ -path  $p$  such that  $q(p) \in F$ .

Conditional fare networks  $(\mathcal{T}, \Gamma, w, e, \mu, \pi)$  can be recast as DFAs for a fixed starting stop  $v \in V$  and final fare state  $f \in F$  if  $|H|$  is finite. Thus, POEAP can be tackled by solving a series of formal-language constrained shortest path problems and selecting the price-optimal path from the successful queries.

Assume the monoid  $(H, +, \leq)$  is trivial, i.e.  $H = \{e\}$  with  $e + e = e$  for some element  $e$ . Then, the transition functions  $\Gamma(\cdot, \cdot, \cdot)$  can be considered independent of  $(H, +, \leq)$ , and we can construct a DFA  $D(s) = (Q, \Sigma, \delta, q_0, F)$  in a straightforward manner by setting  $Q := T, \Sigma := S, q_0 := \tau'$  where  $(\tau', e) = \mu(s)$  is the initial fare state of  $s$  and  $F := \{\tau\}$  for some  $\tau \in T$ . The transition function  $\delta$  is defined as

$$\delta(q, \sigma) := \Gamma(q, x, \sigma) \quad \forall q \in Q, \sigma \in \Sigma.$$

As there is no direct equivalent for nontrivial  $H$  in a DFA, we incorporate it into the state set  $Q$ . We set  $Q := F \times T \times H, \Sigma := H \times S, q_0 := \mu(v)$  and  $F = \{f\}$  for some fare state  $f \in F$ . The transition function  $\delta$  is then defined by

$$\delta(q, \sigma) := (\Gamma(\tau(q), w(q) + w(\sigma), e(\sigma)), w(q) + w(\sigma)) \quad \forall q \in Q, \sigma \in \Sigma.$$

An example of this transformation can be seen in Fig. A.5.

Note that we can restrict  $\Sigma$  to those fare attributes that do really appear on arcs in  $A$ . Hence, we can assume  $\mathcal{O}(|\Sigma|) = \mathcal{O}(|A|)$ . Then, the above DFA can be created in  $\mathcal{O}(|Q| + |\Sigma||Q||B|) = \mathcal{O}(|A||T||H||B|)$  assuming the evaluation time of  $\Gamma$  to be bounded by a polynomial  $B$ . Note that constructing a DFA  $(Q, \Sigma, \delta, q_0, \{f\})$  for all  $f \in F$  is still possible in  $\mathcal{O}(|A||T||H||B|)$  time as the transition function  $\delta$  and the state space  $Q$  need only be constructed once.

REG-ShP can be solved over such a DFA in  $\mathcal{O}(|V||T||H| \log(|V||T||H|) + |A||T||H|)$  using the algorithm given by Barrett et al. (2000). Hence, a superset  $M$  of the Pareto-set of POEAP can be found in polynomial time by solving  $|F| = |T||H|$  instances of REG-ShP. As  $M$  has at most  $|T||H|$  entries it takes  $\mathcal{O}(|T||H| \log(|T||H|))$  time to extract the actual Pareto-set (Kung et al., 1975). Thus, we obtain an overall running time of  $\mathcal{O}(|V||T|^2|H|^2 \log(|V||T||H|) + |A||T|^2|H|^2 + |A||T||H||B|)$ . A slight modification of the proof for Reg-ShP in Barrett et al. (2000) allows us to obtain a tighter bound.

**Lemma A.1** (POEAP with constant travel time over finite monoids). Consider POEAP over  $\mathcal{N} = (\mathcal{T}, \Gamma, w, e, \mu, \pi)$  with  $H$  finite, i.e.,  $|H| < \infty$ . Assume all travel time functions are constants and that the evaluation time of  $\Gamma$  can be bounded by a polynomial  $B$ . Then, POEAP can be solved in  $\mathcal{O}(|V||T||H| \log(|V||T||H|) + |A||T||H||B|)$ .

**Proof.** Let  $D(s) = (Q, \Sigma, \delta, q_0, F)$  be the DFA constructed as described above but letting  $F := Q = F$ , i.e., now all states in the automaton are also accepting states. To keep consistence with automata terminology, we write  $\sigma(v_1, v_2) := (w(v_1, v_2), e(v_1, v_2))$  for  $v_1, v_2 \in V$ .

We construct a product network  $G^\times = (V^\times, A^\times)$  of  $G$  and  $D(s)$  with

$$V(G^\times) = V \times Q$$

$$E(G^\times) = \{(v_1, q_1), (v_2, q_2) | (v_1, v_2) \in E, q_2 = \delta(q_1, \sigma(v_1, v_2))\}.$$

Note that for every  $a \in A$  there are at most  $|T||H|$  edges in  $E^\times$  and hence  $|E^\times| \leq |A||T||H|$ . Thus,  $G^\times$  can be constructed in  $\mathcal{O}(|V||T||H| + |A||T||H|)$ . Using Dijkstra's algorithm, we compute a shortest path tree rooted at  $(s, q_0)$  in  $G^\times$ . In particular, we obtain a shortest  $(s, q_0), (t, q)$ -path for all  $q \in Q = F = T \times H$ . Using a Fibonacci heap, Dijkstra's algorithm has a running time of  $\mathcal{O}(|V^\times| \log(|V^\times|) + |E^\times|)$ . We can again extract the Pareto-set in  $\mathcal{O}(|T||H| \log(|T||H|))$  time, giving an overall runtime of  $\mathcal{O}(|V||T||H| \log(|V||T||H|) + |A||T||H||B|)$ .  $\square$

This result extends naturally to FIFO-travel time functions.

**Proposition A.2** (POEAP over finite monoids is polynomial time solvable in  $|H|$ ). Consider POEAP over the conditional fare network  $\mathcal{N} = (\mathcal{T}, \Gamma, w, e, \mu, \pi)$  under the assumption that  $|H| < \infty$ , that the travel time functions  $c(a) : I \rightarrow I, a \in A$  have the FIFO-property and can be evaluated in constant time, and that the evaluation time of  $\Gamma$  can be bounded by a polynomial  $B$ . Then, POEAP can be solved in  $\mathcal{O}(|V||T||H| \log(|V||T||H|) + |A||T||H||B|)$ .

**Proof.** It is well-established that time-dependent shortest path problems can be solved using a modified version of Dijkstra's algorithm for FIFO networks (Orda and Rom, 1990). The modified algorithm exhibits the same running time as the standard algorithm if the evaluation time of travel time functions is bounded by a constant. Thus, using the modified Dijkstra variant in Lemma A.1 solves POEAP with FIFO travel time functions in  $\mathcal{O}(|V||T||H| \log(|V||T||H|) + |A||T||H||B|)$ .  $\square$

## Appendix B. Pseudocode for the McRAP algorithm

We provide the pseudocode of the McRAP algorithm (Delling et al., 2015) adapted to CFNs as developed in Section 5.2. The function  $B.add((\eta, f))$  removes all labels from  $B$  that are dominated by  $(\eta, f)$ .

## References

- Barrett, C., Jacob, R., Marathe, M., 2000. Formal-language-constrained path problems. SIAM J. Comput. 30 (3), 809–837. <http://dx.doi.org/10.1137/S0097539798337716>.
- Bast, H., Delling, D., Goldberg, A.V., Müller-Hannemann, M., Pajor, T., Sanders, P., Wagner, D., Werneck, R.F., 2016. Route planning in transportation networks. In: Algorithm Engineering - Selected Results and Surveys. In: Lecture Notes in Computer Science, vol. 9220, Springer, pp. 19–80. [http://dx.doi.org/10.1007/978-3-319-49487-6\\_2](http://dx.doi.org/10.1007/978-3-319-49487-6_2).

## Algorithm 1: Price-optimal McRAP

**Data:** Time Table  $\mathbb{T} = (\Pi, \mathcal{P}, \mathcal{R}, \mathcal{D}, F)$ , CFN  $\mathcal{N} = (\mathcal{T}, \Gamma, w, e, \mu, \pi)$ , origin  $p_s \in \mathcal{P}$ , destination  $p_t \in \mathcal{P}$ , departure time  $\eta$ , number of rounds  $K$

```

1  $B_0(p_s) \leftarrow \{(\eta, \mu(p_s))\}$ 
2 mark  $p_s$ 
3 for  $k = 1$  to  $K$  do
4    $B_k(\cdot) \leftarrow \emptyset$ 
5    $Q \leftarrow \emptyset$  // FIFO queue
6   foreach marked stop  $p$  do // Find first marked stop per route
7     foreach route  $r$  with  $p \in \mathcal{P}(r)$  do
8       if  $(r, p') \in Q$  for some  $p' \in \mathcal{P}$  then
9         if  $p$  comes before  $p'$  in  $r$  then
10            $Q.remove(r, p')$ 
11            $Q.add(r, p)$ 
12       else
13          $Q.add(r, p)$ 
14       unmark  $p$ 
15   foreach  $(r, p) \in Q$  do // Traverse Routes
16      $B_r \leftarrow \emptyset$ 
17     foreach stop  $p_i$  in  $r$  beginning with  $p$  do
18       foreach  $(\eta, f, d) \in B_r$  do
19          $\eta \leftarrow \eta_{arr}(d, p)$  // Update label at new station
20          $p_i$ 
21          $w(f) \leftarrow w(f) + w_1(d, p)$ 
22          $\tau(f) \leftarrow \Gamma(\tau(f), w(f), e_1(d, p))$ 
23         if  $(\eta, f)$  not dominated by  $\cup_{j=1, \dots, k} B_j(p)$  then
24            $B_k(p).add((\eta, f))$  // Found new nondominated label
25           mark  $p_i$ 
26       foreach  $(\eta, f) \in B_{k-1}(p)$  do
27          $\eta' \leftarrow \eta + \eta_{ch}$  // Apply transfer costs
28          $w(g) \leftarrow w(f) + w_2(d, p)$ 
29          $\tau(g) \leftarrow \Gamma(\tau(f), w(g), e_2(d, p))$ 
30         if  $(\eta', g)$  not dominated by  $B_r$  then
31            $d \leftarrow \min(d : \eta_{dep}(d, p) \geq \eta')$  // Find next trip
32            $B_r.add((\eta', g, d))$  // Add nondominated label to route bag
33   foreach marked stop  $p$  do // Process Footpaths
34     foreach footpath  $(p, p', l) \in \mathcal{F}$  do
35       foreach  $(\eta, f) \in B_k(p)$  do
36         if  $(\eta + l, f)$  not dominated by  $\cup_{j=1, \dots, k} B_j(p')$  then
37            $B_k(p').add((\eta + l, f))$ 
38   if no stop is marked then // Early Termination Criterion
39     stop

```

- Berger, A., Müller-Hannemann, M., 2009. Subpath-Optimality of Multi-Criteria Shortest Paths in Time- and Event-Dependent Networks. Technical Report, Institute of Computer Science, Martin-Luther-Universität Halle-Wittenberg, URL: <http://wcm.uzi.uni-halle.de/download.php?down=10850&elem=2163494>.
- Blanco, M., Borndörfer, R., Hoang, N.D., Kaier, A., Casas, P.M., Schlechte, T., Schlobach, S., 2017. Cost projection methods for the shortest path problem with crossing costs. In: D'Angelo, G., Dollevoet, T. (Eds.), 17th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2017). In: OpenAccess Series in Informatics (OASICS), vol. 59, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, pp. 15:1–15:14. <http://dx.doi.org/10.4230/OASICS.ATMOS.2017.15>, URL: <http://drops.dagstuhl.de/opus/volltexte/2017/7893>.
- Blanco, M., Borndörfer, R., Hoang, N.-D., Kaier, A., Schlechte, T., Schlobach, S., 2016. The Shortest Path Problem with Crossing Costs. Technical Report 16–70, ZIB, Takustr. 7, 14195 Berlin, URL: [urn:nbn:de:0297-zib-61240](http://nbn-resolving.org/urn:nbn:de:0297-zib-61240).
- Blumenberg, E., Agrawal, A.W., 2014. Getting around when you're just getting by: Transportation survival strategies of the poor. J. Poverty 18 (4), 355–378. <http://dx.doi.org/10.1080/10875549.2014.951905>, URL: <https://doi.org/10.1080/10875549.2014.951905>.
- Borndörfer, R., Euler, R., Karbstein, M., 2021. Ein Graphen-basiertes Modell zur Beschreibung von Preissystemen im öffentlichen Nahverkehr. HEUREKA 21 002/127, 1–15, URL: <https://verlag.fgsv-datenbanken.de/tagungsbaende?kat=HEUREKA&subkat=FGSV+002%2F127+%282021%29&fanr=&va=&titel=>

- &text=&autor=&tagungsband=1256&titel=Ein+Graphen-basiertes+Modell+zur+Beschreibung+von+Preissystemen+im+%C3%B6ffentlichen+Nahverkehr.
- Borndörfer, R., Euler, R., Karbstein, M., Mett, F., 2018. Ein mathematisches Modell zur Beschreibung von Preissystemen im öV. Technical Report 18–47, ZIB, Takustr. 7, 14195 Berlin, URL: [urn:nbn:de:0297-zib-70564](https://nbn-resolving.org/urn:nbn:de:0297-zib-70564).
- Broersma, H., Li, X., Woeginger, G., Zhang, S., 2005. Paths and cycles in colored graphs. *Australas. J. Combin.* 31, 299–311. <http://dx.doi.org/10.1145/62.2737>.
- Brough, R., Freedman, M., Phillips, D.C., 2022. Experimental evidence on the effects of means-tested public transportation subsidies on travel behavior. *Reg. Sci. Urban Econ.* 96, 103803. <http://dx.doi.org/10.1016/j.regsciurbeco.2022.103803>, URL: <https://www.sciencedirect.com/science/article/pii/S0166046222000436>.
- Bull, O., Muñoz, J.C., Silva, H.E., 2021. The impact of fare-free public transport on travel behavior: Evidence from a randomized controlled trial. *Reg. Sci. Urban Econ.* 86, 103616. <http://dx.doi.org/10.1016/j.regsciurbeco.2020.103616>, URL: <https://www.sciencedirect.com/science/article/pii/S016604622030301X>.
- Chen, X., Ma, J., Bai, X., 2020. Mode choice behavior analysis under the impact of transfer fare discount: A case study from Beijing public transit system. In: *Resilience and Sustainable Transportation Systems*. pp. 290–299. <http://dx.doi.org/10.1061/9780784482902.033>, URL: <https://ascelibrary.org/doi/abs/10.1061/9780784482902.033>. arXiv:<https://arxiv.org/abs/https://ascelibrary.org/doi/pdf/10.1061/9780784482902.033>.
- Delling, D., Dibbelt, J., Pajor, T., 2019. Fast and exact public transit routing with restricted Pareto sets. In: *2019 Proceedings of the Twenty-First Workshop on Algorithm Engineering and Experiments. ALENEX*, pp. 54–65. <http://dx.doi.org/10.1137/1.9781611975499.5>, URL: <https://epubs.siam.org/doi/abs/10.1137/1.9781611975499.5>. arXiv:<https://arxiv.org/abs/https://epubs.siam.org/doi/pdf/10.1137/1.9781611975499.5>.
- Delling, D., Pajor, T., Werneck, R.F., 2015. Round-based public transit routing. *Transp. Sci.* 49 (3), 591–604. <http://dx.doi.org/10.1287/trsc.2014.0534>.
- Disser, Y., Müller-Hannemann, M., Schnee, M., 2008. Multi-criteria shortest paths in time-dependent train networks. In: McGeoch, C.C. (Ed.), *Proceedings of the 7th International Conference on Experimental Algorithms. WEA '08*, Springer-Verlag, Berlin, Heidelberg, pp. 347–361. [http://dx.doi.org/10.1007/978-3-540-68552-4\\_26](http://dx.doi.org/10.1007/978-3-540-68552-4_26), URL: <http://dl.acm.org/citation.cfm?id=1788888.1788914>.
- Euler, R., Borndörfer, R., 2019. A Graph- and Monoid-Based Framework for Price-Sensitive Routing in Local Public Transportation Networks. In: Cacchiani, V., Marchetti-Spaccamela, A. (Eds.), *19th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2019)*. In: *OpenAccess Series in Informatics (OASICS)*, vol. 75, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, pp. 12:1–12:15. <http://dx.doi.org/10.4230/OASICS.ATMOS.2019.12>, URL: <http://drops.dagstuhl.de/opus/volltexte/2019/11424>.
- Fleishman, D., Shaw, N., Joshi, A., Freeze, R., Oram, R., 1996. Fare Policies, Structures and Technologies. TCRP Report 10, Transport Cooperative Research Program, Transportation Research Board, Washington DC, URL: <https://www.trb.org/Publications/Blurbs/153836.aspx>.
- Gabow, H.N., Maheshwari, S.N., Osterweil, L.J., 1976. On two problems in the generation of program test paths. *IEEE Trans. Eng. SE-2* (3), 227–231. <http://dx.doi.org/10.1109/TSE.1976.233819>.
- Gündling, F., 2020. Efficient Algorithms for Intermodal Routing and Monitoring in Travel Information Systems (Ph.D. thesis). Technische Universität, Darmstadt, <http://dx.doi.org/10.25534/tuprints-00014212>, URL: <http://tuprints.ulb.tu-darmstadt.de/14212/>.
- Hansen, P., 1980. Bicriterion path problems. In: *Lecture Notes in Economics and Mathematical Systems*, vol. 177, [http://dx.doi.org/10.1007/978-3-642-48782-8\\_9](http://dx.doi.org/10.1007/978-3-642-48782-8_9).
- Hopcroft, J.E., Ullman, J.D., 1979. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley Publishing Company.
- Kung, H.T., Luccio, F., Preparata, F.P., 1975. On finding the maxima of a set of vectors. *J. ACM* 22 (4), 469–476. <http://dx.doi.org/10.1145/321906.321910>, URL: <http://doi.acm.org/10.1145/321906.321910>.
- Maristany de las Casas, P., Sedeno-Noda, A., Borndörfer, R., 2021. An improved multiobjective shortest path algorithm. *Comput. Oper. Res.* 135, <http://dx.doi.org/10.1016/j.cor.2021.105424>.
- Martins, E.Q.V., 1984. On a multicriteria shortest path problem. *European J. Oper. Res.* 16 (2), 236–245. [http://dx.doi.org/10.1016/0377-2217\(84\)90077-8](http://dx.doi.org/10.1016/0377-2217(84)90077-8).
- Mitteldeutscher Verkehrsverbund GmbH, 2019a. MDV fares. <https://www.mdv.de/tickets/befoederungsbedingungen-tarifbestimmungen/>, Accessed: 2019-08-11.
- Mitteldeutscher Verkehrsverbund GmbH, 2019b. MDV GTFS data. <https://www.mdv.de/informationen/downloads/>, Accessed: 2019-08-11.
- Mohri, M., 2002. Semiring frameworks and algorithms for shortest-distance problems. In: *J. Autom. Lang. Comb.*, J. Autom. Lang. Comb. 7 (3), 321–350, URL: <http://dl.acm.org/citation.cfm?id=639508.639512>.
- Müller-Hannemann, M., Schnee, M., 2005. Paying less for train connections with MOTIS. In: *Proceedings of the 5th Workshop on Algorithmic Methods and Models for Optimization of Railways*. In: *OpenAccess Series in Informatics*, vol. 2, p. 657. <http://dx.doi.org/10.4230/OASICS.ATMOS.2005.657>.
- Orda, A., Rom, R., 1990. Shortest-path and minimum-delay algorithms in networks with time-dependent edge-length. *J. ACM* 37 (3), 607–625. <http://dx.doi.org/10.1145/79147.214078>.
- Parmentier, A., 2019. Algorithms for non-linear and stochastic resource constrained shortest path. *Math. Methods Oper. Res.* 89 (2), 281–317. <http://dx.doi.org/10.1007/s00186-018-0649-x>.
- Reinhardt, L.B., Pisinger, D., 2011. Multi-objective and multi-constrained non-additive shortest path problems. *Comput. Oper. Res.* 38 (3), 605–616. <http://dx.doi.org/10.1016/j.cor.2010.08.003>, URL: <https://www.sciencedirect.com/science/article/pii/S0305054810001656>.
- Rosenblum, J., 2020. Expanding Access to the City: How Public Transit Fare Policy Shapes Travel Decision Making and Behavior of Low-Income riders (Ph.D. thesis). Massachusetts Institute of Technology. Department of Urban Studies and Planning, URL: <https://hdl.handle.net/1721.1/127617>.
- Schöbel, A., Urban, R., 2021. The cheapest ticket problem in public transport. <http://dx.doi.org/10.48550/arXiv.2106.10521>, pre-print. arXiv:2106.10521 [math.OC].
- Verkehrsverbund Bremen/Niedersachsen GmbH, 2019. VBN nightliner fares. <https://web.archive.org/web/20200930070236/https://www.vbn.de/tickets/ticketangebot/nachtlinienzuschlag/>, Accessed: 2023-11-10.
- Verkehrsverbund Bremen/Niedersachsen GmbH, 2023. VBN fares. <https://www.vbn.de/tickets/tarifbestimmungen>, Accessed: 2023-11-10.
- Zimmermann, U., 1981. *Linear and Combinatorial Optimization in Ordered Algebraic Structures*. In: *Annals of discrete mathematics*, vol. 10, North-Holland.