An Oracle-based Approach for Price-setting Problems in Logistics *

N. Pommerening, M. Hügging, D. Henke, C. Buchheim, U. Clausen

February 6, 2025

Abstract

We study a bilevel hub location problem where on the upper level, a shipment service provider – the leader – builds a transportation network and sets the prices of shipments on each possible transportation relation. Here, the leader has to take into account the customers' reaction – the follower – who will only purchase transport services depending on their individual budget. The objective is to maximize the generated profit from the leader's perspective.

Since the classical hub location problem is well known to be NP-hard, we focus on the additional complexity induced by the bilevel structure. We propose an oracle-based approach that assumes availability of an algorithm for the prize-collecting variant of the hub location problem. Our main contribution is an efficient reduction of our bilevel price-setting problem to the latter, showing that the bilevel problem is not significantly harder than this. We achieve the reduction through a novel Lagrangian decomposition approach in which the optimal multipliers can be determined analytically. Moreover, we prove that strong duality holds.

Computational experiments highlight the advantages of our method, that significantly outperforms a more standard single-level reformulation. Furthermore, the proposed algorithm can be adapted to numerous price-setting variants of classical combinatorial optimization problems, even beyond logistics.

*We wish to thank Amine Bahij for his committed efforts in executing the computational experiments and making essential modifications to the data set utilized. This work was partially supported by Deutsche Forschungsgemeinschaft (DFG) under grant nos. BU 2313/9-1 and CL 318/38-1.

1 Introduction

The hub location problem, in all its variants, is a well-studied topic in the logistics and optimization literature; see, e.g., the survey papers [1, 10] and the references therein. From a finite set of candidate locations, some have to be selected to become so-called *hubs* that serve as consolidation centers. Furthermore, an assignment of the remaining locations to the hubs has to be determined in order to form a connected hub network. This selection has to ensure that the transport of goods through the network becomes as cheap as possible with respect to the occurring transportation cost. The construction of hubs is beneficial in this model, as transport between hubs comes with discounted costs. The popularity and relevance of the hub location problem stem mainly from its strong applicability in the real world, as problems with a similar structure frequently arise in logistics, transportation, and traffic management.

In practical applications, customer behavior often plays an important role. In the classical hub location problem, the customers' decisions are only represented indirectly by a given amount of transportation volume for every pair of locations. In many applications, this simply does not capture the flexibility of the customers' behavior, and in particular the reciprocal interference of the network structure and the customers. In the present work, our aim is to remedy this by proposing a model that much better represents these mutual dependencies. Our approach uses bilevel optimization, in which two decision makers take decisions in a hierarchical order.

Surprisingly, as of now, there is not much literature on modeling customer behavior in hub location problems. However, there are some papers dealing more explicitly with the presence and autonomy of customers in such problems. Alumur and Kara [2] propose to select the hubs not only based on minimizing costs, but under the constraint of not exceeding a maximum travel time, which can be interpreted as taking into account the needs of customers in a certain sense. Contreras et al. [7] deal with a situation in which customers spontaneously decide whether they want to ship their commodity. In particular, the shipment quantities are now stochastic in order to realize the uncertainty that comes with customers acting autonomously. Mahmutoğulları and Kara [16] also use a bilevel structure to model a hub location problem. Their problem considers two competing transportation companies, which are represented as leader and follower in the bilevel setting. The goal is, from the perspective of the leader, to convince as many customers as possible to decide for the services of the leader company.

In our new approach, we propose a bilevel hub location problem in which there is a single shipment service provider or transportation company (the leader) that builds a transportation network by establishing a classical single-allocation hub network. In addition, this company controls the prices that customers have to pay in order for their products to be shipped through the network. The customers (the follower) can only book the shipment service if the demanded price fits into their budget, so the company's choice has a direct impact on the customers' ability to book its service, which in turn impacts the profit that the company generates. The customers' behavior is thus captured in the follower's problem: A shipment is booked as long as it is affordable and, in particular, customers cannot ship a fraction of their goods. The objective of the shipment service provider is to maximize its profit, where, as usual, the cost of transportation between the hubs is discounted by a certain factor $\alpha \in [0, 1]$.

In a rather recent work, Erdoğan et al. [9] investigate a similar problem. They consider a complete connected graph that can be seen as a pre-existing transportation network. For each arc, a tariff is known in advance that corresponds to the fee for transporting one unit of commodity in this existing network. Moreover, the cost of transporting one unit of commodity is known for each arc. Now, a newcomer enters the market and plans to erect a fixed number of hubs and to choose prices for the transport of commodities through their intra-hub arcs. The customers correspond to commodities that are routed either through the existing network (via a direct connection) or along a route using one of the newly built intra-hub arcs, depending on which route costs less. Here, a commodity is only routed using an intra-hub arc if the total fee for the route, i.e., the sum of the prices or tariffs of the arcs in that route, is lower than the tariff of the direct connection in the existing network by at least a specified factor $\gamma \in [0, 1)$. The authors call this the *saving threshold condition*. The problem consists in locating a fixed number of hubs and determining prices optimally such that the newcomer's profit is maximized. Different from our setting, the authors of [9] do not consider a single-allocation hub location problem, but a model in which the remaining non-hub locations can be assigned to any hub. Furthermore, there are discounted prices for transportation via the hub network only for the customers, but not for the owner of the hub network, as is usually the case in hub location problems and the main motivation for erecting hubs in the first place. Another difference concerns the structure of the follower's problem: The commodities in [9] are shipped such that the resulting prices are minimal, which is modeled as a minimum cost flow problem. This differs from our understanding that customers ship their product as long as they can afford it, without taking the prices of alternative routes into account.

The main difference of our approach, however, lies in the novel solution method and the underlying theoretical results. The authors of [9] investigate their problem from a classical bilevel perspective, exploiting LP duality and optimality criteria for the follower's problem in order to obtain a (nonlinear) single-level reformulation of their original bilevel problem. As their problem is NP-hard, they aim to obtain upper bounds on the optimal value of the original problem by studying the LP relaxation of the (linearized) single-level reformulation. In particular, they present valid inequalities to tighten their LP relaxation. In the following, we not only devise a different solution approach, but also adopt a different paradigm. Taking into account that the underlying classical hub location problem is already NP-hard, on the one hand, but that it is very well studied, on the other hand, we argue that the presence of an algorithm for the underlying hub location problem may be assumed, in the sense of an oracle. This not only allows us to better understand the additional complexity induced by the chosen bilevel structure with respect to the classical problem, and to devise oracle-based algorithms for the bilevel problem, but also opens the door for extending the approach to other underlying combinatorial problems.

Our first step to achieve such an oracle-based algorithm is to forgo a single-level reformulation and instead address the problem by Lagrangian decomposition, maintaining as much of the problem structure as possible. Through our decomposition, we obtain a prize-collecting variant of the hub location problem as one subproblem, which is an interesting optimization problem in its own right. The second subproblem only deals with the customers' decisions for given prices and is thus trivial. Our main result is that the solution of the Lagrangian dual requires no additional effort beyond solving each subproblem once, as we are able to identify optimal Lagrangian multipliers analytically and there is no duality gap. Consequently, the solution of our problem boils down to the solution of the *prize-collecting hub location problem.* In other words, our approach gives rise to an oracle-based approach to the price-setting problem based on a subroutine for the prize-collecting problem. Note the subtle difference between *price* and *prize.* The price to be paid by the customers, as determined in the price-setting problem, roughly corresponds to the prize collected by the company in the prize-collecting subproblem. Some of our results for a special case of the considered price-setting problem have already been presented in a recent master's thesis [18] that was supervised by two of the authors.

The remainder of this paper is organized as follows. In Section 2, we give a brief summary of the necessary concepts of bilevel optimization and then present our price-setting optimization model.

In Section 3, we decompose our optimization problem using Lagrangian relaxation and focus on the details and characteristics of the resulting subproblems, such as their complexity. Afterwards, we show our main result concerning the optimal choice for the Lagrangian multipliers. As a by-product, we prove that there is no duality gap in this approach. In Section 4, we discuss possible extensions of our approach and our results to bilevel price-setting problems building on different underlying problems. In Section 5, we present experimental results showing that our decomposition approach is a very powerful tool for solving the presented bilevel problem and that it significantly outperforms the use of a more classical single-level reformulation of the bilevel problem. Finally, we summarize our most important findings and discuss open questions and future research in Section 6.

2 The price-setting bilevel optimization problem

2.1 Bilevel optimization

In this paper, we investigate a *price-setting bilevel hub location problem*, or (PS-BHLP) for short. For this, we first briefly recall the fundamental aspects of bilevel optimization provided, e.g., in [8]. Bilevel optimization problems involve two decision makers who interact in a hierarchical order, in the sense that an optimization problem contains a second optimization problem in its constraints. Let the latter, called *lower-level problem* or *follower's problem*, be given by

$$\min_{y \in Y} \{f(x,y) : g(x,y) \leq 0, h(x,y) = 0\} ,$$

where f, g, and h are functions of appropriate dimensions, parameterized in $x \in X$. For all $x \in X$, let $\Psi(x)$ denote the set of optimal solutions to the lower-level problem and let $y(x)$ be an element of $\Psi(x)$. Then, the objective of the *bilevel problem* or *leader's problem* is to select the variable x such that it satisfies the so-called *upper-level constraints* $G(x, y(x)) \leq 0$ and $H(x, y(x)) = 0$ and optimizes the *upper-level objective function* $F(x, y(x))$, with G, H, and F having appropriate dimensions. The bilevel problem is thus given as

$$\min_{x \in X} \{F(x,y(x)) : G(x,y(x)) \leq 0, H(x,y(x)) = 0, y(x) \in \Psi(x)\} . \quad (1)$$

This problem thus models the interaction between two decision makers as follows: First, the *leader* takes her decision $x \in X$, and based on this, the *follower* responds with his choice $y(x) \in \Psi(x)$.

When optimizing her objective function, the leader must take into account the follower's response, as this affects her objective value and possibly the feasibility of the joint solution $(x, y(x))$.

Note that the above problem description is imprecise in case the set $\Psi(x)$ contains more than one lower-level optimal solution. There are two common approaches to deal with this, called the *optimistic* and the *pessimistic view*: If the follower has multiple optimal solutions, in the first case, he chooses a solution that is most beneficial from the leader's point of view; otherwise, he acts as an adversary and chooses a solution causing as much damage as possible to the leader. Formally, the leader's problem as stated in (1) represents the optimistic view, as the leader can choose $y(x)$ from $\Psi(x)$. Bilevel optimization is strongly NP-hard in general, even when all the functions involved are linear [13, 12]. The standard solution approach is to recast the bilevel problem as a single-level problem using optimality conditions for the lower level; see, e.g., [6] and the references therein.

An important class of bilevel problems, including our model presented below, consists of price-setting problems, where the leader chooses prices for a service to be paid by the follower [15].

2.2 The price-setting hub location problem

Given a complete directed graph $G = (V, E)$, let the nodes in V represent locations and let the arcs in E represent the network of routes connecting these locations. We consider a shipment service that can be booked by customers to transport their products from origin locations to destination locations through this transportation network. However, customers can only book the shipment service if they can afford to pay for it. Conversely, the shipment service provider has control over the transport prices, thus impacting the customer's ability to book its service.

The objective of the shipment service provider is to maximize its profit. In the context of bilevel optimization, we deal with the following hierarchical decision structure: The leader (the shipment service provider) takes her decisions first and solves a classical κ-hub location problem with single allocation, represented by binary variables x with

$$x \in \mathrm{HLP}_{\text{single-alloc}} := \left\{x \in \{0,1\}^{V \times V} : \textstyle\sum_{j \in V} x_{jj} = \kappa,\ \sum_{j \in V} x_{ij} = 1\ \forall i \in V,\ \ x_{ij} \le x_{jj}\ \forall i, j \in V\right\}.$$

Here, the x_{ij} both model whether $j \in V$ is chosen as a hub (if $x_{jj} = 1$) and whether a node $i \in V$ is allocated to a hub j (if $x_{ij} = 1$). The first constraint makes sure that exactly κ hubs are chosen, while the second class of constraints models a unique allocation, and the third class of constraints forbids to allocate nodes to non-hub nodes. In addition, the leader chooses a per unit price $p_{ij} \geq 0$ for each pair $i, j \in V$ that must be paid for the customer's product to be transported from i to j.

Let the leader's cost of transporting one unit of the product along arc (i, j) be given by $c_{ij} \geq 0$, and between two hubs, let the cost be discounted by a factor $\alpha \in [0, 1]$. Throughout, we will use the shorthand notation $\tilde{\mathbf{c}}_{ij}(x)$ for the cost (arising for the shipment service provider) of transporting one unit of the product from node i to node j through the hub network defined by x, i.e.,

$$\tilde{\mathbf{c}}_{ij}(x) := \sum_{k \in V} c_{ik} x_{ik} + \alpha \sum_{k,l \in V} c_{kl} x_{ik} x_{jl} + \sum_{l \in V} c_{lj} x_{jl} \, .$$

However, in order to maximize her profit, the leader also has to take into account the revenues received from the customers, i.e., she has to anticipate the follower's reaction when taking her decision. Here, the follower's problem models the behavior of all customers simultaneously with respect to the decision taken by the leader. For each pair $i, j \in V$, let the set $M_{ij} := \{1, ..., m_{ij}\}$ consist of all customers who wish to ship their product from i to j. A customer $z \in M_{ij}$ has a budget $b_{ij}^z \geq 0$, their shipment volume is given by $a_{ij}^z > 0$ and they can only afford to book the shipment of their product on the route if $a_{ij}^z p_{ij} \leq b_{ij}^z$. In particular, we assume that the customer cannot ship a fraction of their product in case their budget is too small. In the follower's problem, the customers' behavior is thus modeled with the help of binary variables y_{ij}^z that indicate whether customer $z \in M_{ij}$ can afford the shipment of their goods and thus books it ($y_{ij}^z = 1$) or not ($y_{ij}^z = 0$). We assume that, whenever a customer can afford to ship their goods, they do not decide against it. Therefore, the follower's objective is to maximize the sum of all y variables.

Combining the different perspectives, the full bilevel optimization problem can now be formulated as follows:

$$\left.\begin{aligned}
\max \quad & \sum_{i,j \in V} \sum_{z \in M_{ij}} a_{ij}^z y_{ij}^z (p_{ij} - \tilde{\mathbf{c}}_{ij}(x)) \\
\text{s.t.} \quad & x \in \text{HLP}_{\text{single-alloc}} \\
& p_{ij} \geq 0 \quad \forall i, j \in V \\
& \left.\begin{aligned}
y \in \arg\max \quad & \sum_{i,j \in V} \sum_{z \in M_{ij}} y_{ij}^z \\
\text{s.t.} \quad & a_{ij}^z p_{ij} y_{ij}^z \leq b_{ij}^z \quad \forall i, j \in V, z \in M_{ij} \\
& y_{ij}^z \in \{0, 1\} \quad \forall i, j \in V, z \in M_{ij}
\end{aligned}\right\} \text{(FP)}
\end{aligned}\right\} \text{(PS-BHLP)}$$

Herein, the lower-level problem (FP) models the behavior of all customers at the same time. Note that (FP) decomposes into separate optimization problems for each $i, j \in V$, and additionally for each $z \in M_{ij}$ if p_{ij} is fixed. In contrast to the general bilevel optimization setting, we do not have to distinguish between the pessimistic and the optimistic view in our case. Indeed, the follower's problem always has a unique optimal solution that can be determined easily as follows.

Proposition 1. *Let $p_{ij} \geq 0$ for all $i, j \in V$. Then the follower's problem (FP) has a unique optimal solution y^*, which is given by*

$$(y^*)_{ij}^z = \begin{cases} 1 & \text{if } p_{ij} \leq b_{ij}^z / a_{ij}^z \\ 0 & \text{otherwise} \end{cases}$$

for all $i, j \in V$ and $z \in M_{ij}$.

For the following, we may assume that the customers are sorted such that

$$\frac{b_{ij}^1}{a_{ij}^1} > \frac{b_{ij}^2}{a_{ij}^2} > \cdots > \frac{b_{ij}^{m_{ij}}}{a_{ij}^{m_{ij}}} \tag{2}$$

holds for each pair $i, j \in V$. This assumption can be made without loss of generality. In fact, we can first sort the customers for all $i, j \in V$ and thus only need to argue why strict inequalities can be assumed. But, as soon as several customers for the same pair $i, j \in V$ have the same ratio, they will always take the same decisions according to Proposition 1. It is thus possible to merge these customers and represent them by one customer whose budget and transport volume are the sum of the corresponding values of the customers being merged.

We now turn our attention to the leader's problem and first show that an optimal solution always exists, which is not clear a priori due to the bilevel structure and the continuous leader's variables p_{ij}.

Theorem 1. *Problem* (PS-BHLP) *admits an optimal solution.*

Proof. Since the x variables can only achieve finitely many feasible values, it suffices to show that there exist values for p_{ij}, $i, j \in V$, that maximize the objective of (PS-BHLP) for fixed x. This can be done independently for each pair $i, j \in V$, so that the task reduces to showing that

$$\varphi_{ij}(p_{ij}) := \sum_{z \in M_{ij}} a_{ij}^z y_{ij}^z (p_{ij} - \tilde{\mathbf{c}}_{ij}(x))$$

attains its maximum over $p_{ij} \geq 0$, where each y^z_{ij} is determined by p_{ij} as described in Proposition 1. Note that the function φ_{ij} is piecewise linear and continuous from the left with possible breakpoints in b^z_{ij}/a^z_{ij} for $z \in M_{ij}$. Also, note that $\varphi_{ij}(p_{ij}) = 0$ for $p_{ij} > b^1_{ij}/a^1_{ij}$, so that φ_{ij} cannot be unbounded. Moreover, each linear piece of φ_{ij} is monotonously increasing, so that φ_{ij} is upper semi-continuous. This concludes the proof. □

The structure of the function φ_{ij} is illustrated in Figure 1. Intuitively, for a fixed network structure and within the range where the same customers can afford to order the transport of their goods, it is always optimal to push the price to the upper limit of what these customers can pay.

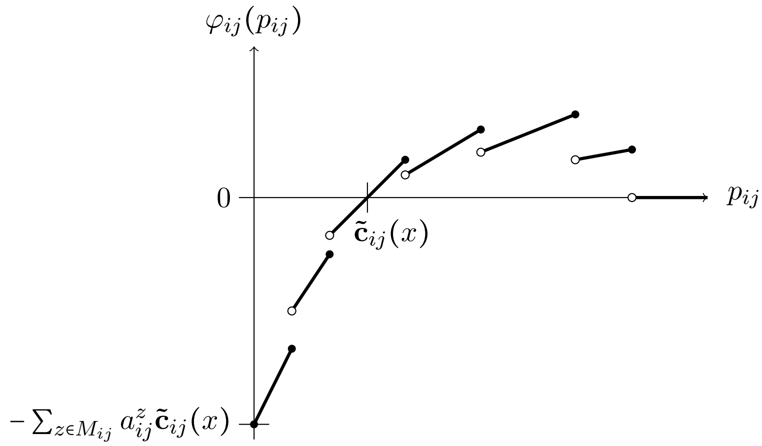

Figure 1: Typical structure of the function $\varphi_{ij}(p_{ij})$ defined in the proof of Theorem 1

Considering the close relation to the hub location problem and the additional complexity introduced by the bilevel structure, it is not surprising that the problem (PS-BHLP) is NP-hard.

Theorem 2. *The problem* (PS-BHLP) *is NP-hard, even if* $m_{ij} = 1$ *for all* $i, j \in V$*, i.e., even if there is only one customer per relation.*

Proof. We polynomially reduce the classical hub location problem, which is known to be NP-hard [14], to our bilevel variant. Let an instance $I = (V, \kappa, \alpha, a, c)$ of the classical κ-hub location problem be given, where a_{ij} is the fixed shipping volume from node i to node j, i.e., consider the problem

$$\left.\begin{aligned} \min \quad & \sum_{i,j \in V} a_{ij} \tilde{\mathbf{c}}_{ij}(x) \\ \text{s.t.} \quad & x \in \text{HLP}_{\text{single-alloc}} \end{aligned}\right\} \qquad (3)$$

We then define an instance $\tilde{I} = (V, \kappa, \alpha, a, c, b)$ of (PS-BHLP) by setting $m_{ij} := 1$ for all $i, j \in V$; thus we may drop the superscript z. Furthermore, we define $b_{ij} := M$ for all $i, j \in V$, where $M > 0$ is big enough, e.g., $M := \sum_{i,j \in V} c_{ij} + 1$, so that we have $M > \max \{c_{ih} + \alpha c_{hl} + c_{lj} : i, j, h, l \in V\}$. Now, if we consider any feasible solution p, x, y of (PS-BHLP), the leader has exactly two choices for each pair i, j. If she chooses to make the transport too expensive for the customer, i.e., $p_{ij} > b_{ij}$ and therefore $y_{ij} = 0$, then the contribution to the objective function is zero. If, instead, she chooses to make the transport affordable for the customer, i.e., $p_{ij} \leq b_{ij}$ and hence $y_{ij} = 1$, then there is no incentive for her to choose any price other than $p_{ij} = b_{ij}$, as this is the maximum price the customer is able to pay. By definition of M, the pair i, j contributes positively to the objective value, as $a_{ij}(b_{ij} - \tilde{\mathbf{c}}_{ij}(x)) = a_{ij}(M - \tilde{\mathbf{c}}_{ij}(x)) > 0$. Since M always exceeds $\tilde{\mathbf{c}}_{ij}(x)$, for any feasible x and any pair $i, j \in V$, an optimal solution p^*, x^*, y^* of (PS-BHLP) satisfies $p^*_{ij} = M$ and $y^*_{ij} = 1$ for all $i, j \in V$, and the optimal value of (PS-BHLP) is given by

$$\sum_{i,j \in V} a_{ij} y^*_{ij} (p^*_{ij} - \tilde{\mathbf{c}}_{ij}(x^*)) = \underbrace{\sum_{i,j \in V} a_{ij} M}_{\text{constant}} - \sum_{i,j \in V} a_{ij} \tilde{\mathbf{c}}_{ij}(x^*) \, .$$

Since x^* maximizes the objective function of (PS-BHLP) for $p^*_{ij} = M$ and $y^*_{ij} = 1$, $i, j \in V$, it minimizes the objective function of (3). As the instance $\tilde{I}$ can be constructed in polynomial time, the classical hub location problem polynomially reduces to (PS-BHLP) and thus, the claim follows from the NP-hardness of the former. □

However, it follows from our results presented in the following section (or from the single-level reformulation devised in Section 5.1) that (PS-BHLP) is NP-easy and hence NP-equivalent.

3 Lagrangian decomposition

The first step in our solution approach to the bilevel problem (PS-BHLP) introduced in the previous section is to apply Lagrangian decomposition. The latter is a special case of Lagrangian relaxation, which is covered in many textbooks; see, e.g., [17]. To the best of our knowledge, this is the first time Lagrangian relaxation is used to address a bilevel optimization problem. The intuitive idea is to decouple the customers' decisions from the hub location problem by relaxing the connection between the two while still maintaining as much of the problem structure as possible. More precisely, we first introduce a copy $\bar{y}^z_{ij} \in \{0,1\}$ of each follower's variable y^z_{ij} and enforce equality between them using constraints $\bar{y}^z_{ij} = y^z_{ij}$ for all $i,j \in V$ and $z \in M_{ij}$, where all new variables and constraints are added to the leader's problem. Furthermore, we replace the variables y by the copies $\bar{y}$ in the part of the objective function that deals with the transportation cost.

In addition, we introduce *precedence constraints* of the form $\bar{y}^z_{ij} \le \bar{y}^l_{ij}$ for the copies $\bar{y}$, for $i,j \in V$ and $l < k \in M_{ij}$. We point out that, due to Assumption (2), these precedence constraints are redundant for Problem (PS-BHLP), as they are necessarily satisfied for an optimal solution of the follower's problem; see Proposition 1. Hence, including them in Problem (PS-BHLP) does not change the feasible set and, thus, the following problem is equivalent to (PS-BHLP):

$$
\begin{aligned}
\max \quad & \sum_{i,j \in V} \sum_{z \in M_{ij}} a^z_{ij} y^z_{ij} p_{ij} \quad - \sum_{i,j \in V} \sum_{z \in M_{ij}} a^z_{ij} \bar{y}^z_{ij} \tilde{\mathbf{c}}_{ij}(x) \\
\text{s.t.} \quad & x \in \mathrm{HLP}_{\text{single-alloc}} \\
& p_{ij} \ge 0 && \forall i,j \in V \\
& \bar{y}^z_{ij} \in \{0,1\} && \forall i,j \in V,\ z \in M_{ij} \\
& \bar{y}^z_{ij} = y^z_{ij} && \forall i,j \in V,\ z \in M_{ij} \\
& \bar{y}^z_{ij} \le \bar{y}^l_{ij} && \forall i,j \in V,\ l,z \in M_{ij},\ l < z \\
& \text{(FP)}
\end{aligned}
$$

Now we relax the equality constraint $y^z_{ij} = \bar{y}^z_{ij}$ for all $i,j \in V$ and $z \in M_{ij}$ and introduce an associated Lagrangian multiplier $\lambda^z_{ij} \in \mathbb{R}$. The resulting Lagrangian relaxation has the same constraints as the previous problem except for the relaxed equality constraints, and the objective function now reads

$$
\max \quad \sum_{i,j \in V} \sum_{z \in M_{ij}} a^z_{ij} y^z_{ij} p_{ij} - \sum_{i,j \in V} \sum_{z \in M_{ij}} a^z_{ij} \bar{y}^z_{ij} \tilde{\mathbf{c}}_{ij}(x) + \sum_{i,j \in V} \sum_{z \in M_{ij}} \lambda^z_{ij} (\bar{y}^z_{ij} - y^z_{ij}) \,.
$$

Through sharp observation it becomes apparent that the Lagrangian relaxation decomposes into two separate subproblems. We call the first problem the *customers' decision problem.* It contains the variables p deciding the prices as well as the follower's problem in the variables y. In particular, it is still a bilevel optimization problem:

$$\left.\begin{aligned}\mathcal{L}_{p,y}(\lambda) = \quad \max \quad & \sum_{i,j \in V} \sum_{z \in M_{ij}} y_{ij}^z (a_{ij}^z p_{ij} - \lambda_{ij}^z) \\ \text{s.t.} \quad & p_{ij} \geq 0 \qquad \forall i,j \in V \\ & \text{(FP)}\end{aligned}\right\} \tag{CDP}$$

The second problem is called *precedence prize-collecting hub location problem*. It contains the variables x concerning the hub location as well as the copied variables $\bar{y}$ and the precedence constraints:

$$\left.\begin{aligned}\mathcal{L}_{x,\bar{y}}(\lambda) = \quad \max \quad & \sum_{i,j \in V} \sum_{z \in M_{ij}} \bar{y}_{ij}^z (\lambda_{ij}^z - a_{ij}^z \tilde{\mathbf{c}}_{ij}(x)) \\ \text{s.t.} \quad & x \in \text{HLP}_{\text{single-alloc}} \\ & \bar{y}_{ij}^z \leq \bar{y}_{ij}^l \qquad \forall i,j \in V,\ l,z \in M_{ij},\ l < z \\ & \bar{y}_{ij}^z \in \{0,1\} \qquad \forall i,j \in V,\ z \in M_{ij}\end{aligned}\right\} \tag{PPC-HLP}$$

Our decision to include the (redundant) precedence constraints in Problem (PS-BHLP) only to shift them into the prize-collecting subproblem might seem counterintuitive at first; the motivation for this will become clear in Section 3.2, where we show that the precedence constraints allow us to make a simple guess for optimal Lagrangian multipliers and help us to prove their optimality. We will then show that instead of (PPC-HLP), one may also consider a pure prize-collecting hub location problem (without any kind of precedence constraints) by defining the prizes in the latter in a particular way depending on the optimal Lagrangian multipliers for (PPC-HLP).

Now, one can easily show the following result, which can also be derived from the general theory of Lagrangian relaxation; see, e.g., [17, Prop. 6.1].

Proposition 2. *For any fixed λ, an upper bound on the optimal value of* (PS-BHLP) *is given by the sum of the optimal values of the customers' decision problem* (CDP) *and the precedence prize-collecting hub location problem* (PPC-HLP)*. Formally, for $m = \sum_{i,j \in V} m_{ij}$, we thus have*

$$\text{OPT(PS-BHLP)} \leq \mathcal{L}_{p,y}(\lambda) + \mathcal{L}_{x,\bar{y}}(\lambda) \quad \forall \lambda \in \mathbb{R}^m .$$

3.1 Complexity of the subproblems

We now show that the NP-hardness of the problem (PS-BHLP) transfers to the precedence prize-collecting hub location subproblem, while the customers' decision problem is polynomially solvable. We emphasize that this is a rather atypical use of Lagrangian relaxation, which usually aims at producing subproblems that are all efficiently solvable. In contrast, our goal is an oracle-based approach. Furthermore, while Lagrangian relaxation typically only leads to dual bounds, we will see in Section 3.2 below that our Lagrangian dual problem leaves no duality gap.

It is easy to see that the customers' decision problem (CDP) actually decomposes into $\mathcal{O}\left(|V|^2\right)$ problems, one for each relation $i, j \in V$, which have the following form:

$$\left.\begin{aligned} \max \quad & \sum_{z \in M_{ij}} y_{ij}^z (a_{ij}^z p_{ij} - \lambda_{ij}^z) \\ \text{s.t.} \quad & p_{ij} \geq 0 \\ & y \in \arg\max \quad \sum_{z \in M_{ij}} y_{ij}^z \\ & \qquad \text{s.t.} \quad a_{ij}^z p_{ij} y_{ij}^z \leq b_{ij}^z \quad \forall z \in M_{ij} \\ & \qquad\qquad\; y_{ij}^z \in \{0,1\} \qquad \forall z \in M_{ij} \end{aligned}\right\} \qquad (\text{CDP}_{ij})$$

For each resulting subproblem, an optimal solution can be computed very efficiently.

Lemma 1. *For any given* λ*, the problem* (CDP_{ij}) *can be solved in time* $\mathcal{O}(m_{ij})$ *and an optimal solution* (p_{ij}^*, y_{ij}^*) *is either given by*

$$p_{ij}^* = \frac{b_{ij}^k}{a_{ij}^k}\,, \qquad (y^*)_{ij}^z = \begin{cases} 1 & \text{for } z = 1, \dots, k\,, \\ 0 & \text{otherwise} \end{cases}$$

for some critical index $k \in M_{ij}$*, or by any* $p_{ij}^* > b_{ij}^1 / a_{ij}^1$ *and* $(y^*)_{ij}^z = 0$ *for all* $z \in M_{ij}$*.*

Proof. Observe that the objective function of (CDP_{ij}) is very similar to the objective function φ_{ij} from the proof of Theorem 1. It can be described as a function of $p_{ij} \geq 0$, since y_{ij}^z for all $z \in M_{ij}$ is determined by p_{ij} as described in Proposition 1. If $p_{ij} \in \left(b_{ij}^{k+1}/a_{ij}^{k+1}, b_{ij}^k/a_{ij}^k\right]$ for $k \in \{1, \dots, m_{ij} - 1\}$, this leads to the follower's response y_{ij} in which k is the critical index, i.e., the largest index z with $y_{ij}^z = 1$. If $p_{ij} \in (b_{ij}^1/a_{ij}^1, \infty)$, this leads to the follower's response $y_{ij}^z = 0$ for all $z \in M_{ij}$ which trivially produces an objective value of zero. Each of the previous intervals for p_{ij} corresponds to a linear segment of the objective function that is monotonously increasing and continuous from the left. Consequently, an optimal choice of the price p_{ij}^* will be either one of the breakpoints b_{ij}^z/a_{ij}^z or any value larger than b_{ij}^1/a_{ij}^1. □

By the decomposition of (CDP) into independent subproblems (CDP$_{ij}$), which in turn can be solved in time $\mathcal{O}(m_{ij})$, we obtain:

Theorem 3. *For any given λ, the customers' decision problem* (CDP) *admits an optimal solution, and the problem can be solved efficiently in time $\mathcal{O}\left(|V|^2 \max_{i,j\in V} m_{ij}\right)$.*

For the precedence prize-collecting subproblem, the existence of an optimal solution is obvious since all variables are binary, the problem is always feasible, and the problem is a single-level problem. However, the problem remains as hard as the original bilevel problem.

Theorem 4. *The precedence prize-collecting hub location problem* (PPC-HLP)*, with λ as part of the input, is NP-hard, even if there is only one customer per relation.*

Proof. We show the result by a similar reduction as in the proof of Theorem 2. More precisely, we show that the classical hub location problem can be reduced to the precedence prize-collecting hub location problem in polynomial time. Let $I = (V, \kappa, \alpha, a, c)$ be an instance of the classical hub location problem. Define an instance $\tilde{I}$ of the precedence prize-collecting problem (PPC-HLP) by keeping the input of I, setting $m_{ij} := 1$ for all $i, j \in V$ (thus dropping the superscripts z), and defining $\lambda_{ij} := M$ for all $i, j \in V$, where $M := \left(\sum_{i,j\in V} a_{ij}\right)\left(\sum_{i,j\in V} c_{ij}\right) + 1$ and hence

$$M > a_{ij} \max_{h,l\in V} \left\{c_{ih} + \alpha c_{hl} + c_{lj}\right\} \quad \forall i, j \in V \ .$$

With this choice of λ, and for all feasible x, the contribution in the objective for $\bar{y}_{ij} = 1$ is

$$1 \cdot (\lambda_{ij} - a_{ij}\tilde{\mathbf{c}}_{ij}(x)) > M - a_{ij} \max_{h,l\in V} \left\{c_{ih} + \alpha c_{hl} + c_{lj}\right\} > 0 \ ,$$

while for $\bar{y}_{ij} = 0$, the contribution to the objective value is trivially zero. Thus, for the given Lagrangian multipliers λ, any optimal solution $x^*, \bar{y}^*$ of $\tilde{I}$ satisfies $\bar{y}_{ij}^* = 1$ for all $i, j \in V$. The optimal value for instance $\tilde{I}$ is given by

$$\sum_{i,j\in V} \underbrace{\bar{y}^*_{ij}}_{=1}(\lambda_{ij} - a_{ij}\tilde{\mathbf{c}}_{ij}(x^*)) = \underbrace{\sum_{i,j\in V}\lambda_{ij}}_{\text{constant}} - \sum_{i,j\in V} a_{ij}\tilde{\mathbf{c}}_{ij}(x^*)\ .$$

As x^* maximizes the objective of the precedence prize-collecting problem for instance $\tilde{I}$, it also minimizes the objective of the classical hub location problem for instance I. □

Since the hardness result of Theorem 4 even holds for instances with only one customer per relation, we also obtain NP-hardness of the following *prize-collecting hub location problem* which does not contain any precedence constraints:

$$\left.\begin{aligned} \mathcal{Q}_{x,\bar{y}}(\lambda) = \quad \max \quad & \sum_{i,j\in V}\sum_{z\in M_{ij}} \bar{y}^z_{ij}(\lambda^z_{ij} - a^z_{ij}\tilde{\mathbf{c}}_{ij}(x)) \\ \text{s.t.} \quad & x \in \text{HLP}_{\text{single-alloc}} \\ & \bar{y}^z_{ij} \in \{0,1\} \qquad\qquad \forall i,j \in V,\ z \in M_{ij} \end{aligned}\right\} \qquad \text{(PC-HLP)}$$

Not only (PPC-HLP) might be an interesting object to study in its own right, but especially its relaxation (PC-HLP). Unlike the prize-collecting variants of other classical optimization problems such as the traveling salesman problem [4] or the vehicle routing problem [3], to the best of our knowledge, neither of the two formerly described prize-collecting hub location problems has been studied in the literature yet. The problem (PC-HLP) could be used to model situations in which the shipment service provider is not forced to accept all shipment orders, but can rather choose individually depending on whether this is profitable or not. This choice is then modeled by the binary variables $\bar{y}^z_{ij}$.

3.2 Optimal Lagrangian multipliers

As stated in Proposition 2 above, the optimal value of the Lagrangian relaxation $\mathcal{L}_{p,y}(\lambda) + \mathcal{L}_{x,\bar{y}}(\lambda)$ yields an upper bound on the optimal value of (PS-BHLP), for any choice of λ. However, we are interested in computing a *best possible*, i.e., *smallest*, such upper bound. Therefore, in this section, we address the Lagrangian dual problem

$$\mathcal{L}_D = \min_{\lambda}\ (\mathcal{L}_{y,p}(\lambda) + \mathcal{L}_{x,\bar{y}}(\lambda))\ .$$

In fact, unlike in most other Lagrangian relaxation approaches, we will be able to determine optimal multipliers λ^* analytically. We introduce them in the next definition.

Definition 1. *We define the Lagrangian multipliers* λ^* *by setting*

$$(\lambda^*)_{ij}^k := \Big(\sum_{z=1}^{k} a_{ij}^z\Big)\frac{b_{ij}^k}{a_{ij}^k} - \Big(\sum_{z=1}^{k-1} a_{ij}^z\Big)\frac{b_{ij}^{k-1}}{a_{ij}^{k-1}}$$

for all $i, j \in V$ *and* $k \in M_{ij}$.

We briefly point out that the multipliers λ^* have the following characteristic, which we will exploit from time to time:

$$\sum_{z=1}^{k}(\lambda^*)_{ij}^z = \frac{b_{ij}^k}{a_{ij}^k}\sum_{z=1}^{k} a_{ij}^z \quad \forall k \in \{1, ..., m_{ij}\}\ . \tag{4}$$

The next theorem shows that the Lagrangian multipliers λ^* from Definition 1 are indeed optimal, i.e., they minimize the gap between the optimal value of (PS-BHLP) and the bound given by the Lagrangian relaxation. In fact, these multipliers do not leave any gap.

Theorem 5. *The Lagrangian multipliers* λ^* *as defined in Definition 1 solve the Lagrangian dual problem optimally. Moreover, there is no duality gap, i.e.,* $\mathrm{OPT(PS\text{-}BHLP)} = \mathcal{L}_{p,y}(\lambda^*) + \mathcal{L}_{x,\bar{y}}(\lambda^*)$.

Proof. The first statement clearly follows from the second by Proposition 2. We prove the second statement by showing that $\mathcal{L}_{p,y}(\lambda^*) = 0$ and $\mathcal{L}_{x,\bar{y}}(\lambda^*) = \mathrm{OPT(PS\text{-}BHLP)}$. For this, it suffices to show $\mathcal{L}_{p,y}(\lambda^*) \le 0$ and $\mathcal{L}_{x,\bar{y}}(\lambda^*) \le \mathrm{OPT(PS\text{-}BHLP)}$, due to Proposition 2.

Since the customers' decision problem (CDP) decomposes into independent problems (CDP_{ij}) for each relation $i, j \in V$, we can consider each of those separately. For $p_{ij} \in \big(b_{ij}^{k+1}/a_{ij}^{k+1}, b_{ij}^k/a_{ij}^k\big]$ with $k \in \{1, \dots, m_{ij} - 1\}$, the follower responds with y_{ij} such that k is the critical index; see Lemma 1. The objective value of the feasible solution (p_{ij}, y_{ij}) in this subproblem is

$$\sum_{z=1}^{k}\underbrace{y_{ij}^z}_{=1}\big(a_{ij}^z p_{ij} - (\lambda^*)_{ij}^z\big) = p_{ij}\sum_{z=1}^{k} a_{ij}^z - \sum_{z=1}^{k}(\lambda^*)_{ij}^z = p_{ij}\sum_{z=1}^{k} a_{ij}^z - \frac{b_{ij}^k}{a_{ij}^k}\sum_{z=1}^{k} a_{ij}^z = \underbrace{\Big(p_{ij} - \frac{b_{ij}^k}{a_{ij}^k}\Big)}_{\le 0}\sum_{z=1}^{k} a_{ij}^z \le 0\ .$$

The same reasoning applies to $p_{ij} \in [0, b_{ij}^{m_{ij}}/a_{ij}^{m_{ij}}]$. Finally, if $p_{ij} > b_{ij}^1/a_{ij}^1$, we have $y_{ij}^z = 0$ for all $z \in M_{ij}$ in the follower's response, which yields an objective value of zero. In summary, we obtain $\mathcal{L}_{p,y}(\lambda^*) \le 0$.

Now consider an optimal solution $(x^*, \bar{y}^*)$ to (PPC-HLP) with $\lambda = \lambda^*$, and let k_{ij} refer to the critical index of $\bar{y}^*_{ij}$. Since $\bar{y}^*$ satisfies the precedence constraints, the optimal value $\mathcal{L}_{x,\bar{y}}(\lambda^*)$ is given by

$$\sum_{i,j\in V}\Big(\sum_{z=1}^{k_{ij}}\underbrace{(\bar{y}^*)^z_{ij}}_{=1}(\lambda^*)^z_{ij} - \tilde{\mathbf{c}}_{ij}(x^*)\sum_{z=1}^{k_{ij}}\underbrace{(\bar{y}^*)^z_{ij}}_{=1}a^z_{ij}\Big) \overset{(4)}{=} \sum_{i,j\in V}\Big(\frac{b^{k_{ij}}_{ij}}{a^{k_{ij}}_{ij}}\sum_{z=1}^{k_{ij}}a^z_{ij} - \tilde{\mathbf{c}}_{ij}(x^*)\sum_{z=1}^{k_{ij}}a^z_{ij}\Big)$$

$$= \sum_{i,j\in V}\Big(\Big(\frac{b^{k_{ij}}_{ij}}{a^{k_{ij}}_{ij}} - \tilde{\mathbf{c}}_{ij}(x^*)\Big)\sum_{z=1}^{k_{ij}}a^z_{ij}\Big).$$

If $(\bar{y}^*)^z_{ij} = 0$ for all $z \in M_{ij}$, so that there is no critical index k_{ij} for pair $i,j \in V$, we define $k_{ij} := 0$ and the corresponding term in the sum above yields zero. Set $p^*_{ij} := b^{k_{ij}}_{ij}/a^{k_{ij}}_{ij}$ if $k_{ij} > 0$ and $p^*_{ij} := b^1_{ij}/a^1_{ij} + 1$ otherwise. Then $(x^*, p^*, \bar{y}^*)$ is feasible for (PS-BHLP), as $\bar{y}^*$ is an optimal follower's response to p^*. It is straightforward to see that in (PS-BHLP), the solution $(x^*, p^*, \bar{y}^*)$ is not only feasible, but it also still achieves the same objective value $\mathcal{L}_{x,\bar{y}}(\lambda^*)$. This proves $\mathcal{L}_{x,\bar{y}}(\lambda^*) \le \text{OPT(PS-BHLP)}$. □

In the previous proof, the precedence constraints in Problem (PPC-HLP) play a very important role. Only thanks to them, we can translate an optimal solution of (PPC-HLP) into a feasible solution for (PS-BHLP) because the precedence constraints ensure that a compatible notion of a critical index is available in both (PPC-HLP) and (PS-BHLP). The same proof does not work if (PC-HLP) is used instead of (PPC-HLP), unless the precedence constraints are coincidentally satisfied. The next lemma captures such a situation, i.e., a situation in which the sorting with respect to decreasing ratios λ^z_{ij}/a^z_{ij} agrees with the sorting of the customers with respect to b^z_{ij}/a^z_{ij}.

Lemma 2. *Consider Lagrangian multipliers λ and some pair $i,j \in V$ such that the ratios λ^z_{ij}/a^z_{ij} are monotonously decreasing with increasing values of z. Then the precedence constraints for the pair i,j in* (PPC-HLP) *are redundant.*

Proof. It suffices to show that there always exists an optimal solution of (PC-HLP) that satisfies the precedence constraints, under the given monotonicity assumption. Indeed, let $(x^*, \bar{y}^*)$ be an optimal solution of (PC-HLP). Then $(\bar{y}^*)^k_{ij} = 1$ for some $k \in M_{ij}$ implies $\lambda^z_{ij} - a^z_{ij}\tilde{\mathbf{c}}_{ij}(x^*) \ge 0$, which is equivalent to $\lambda^z_{ij}/a^z_{ij} \ge \tilde{\mathbf{c}}_{ij}(x^*)$ using $a > 0$. The monotonicity of the ratios then implies that we may set $(\bar{y}^*)^z_{ij} = 1$ for all $z < k$ without losing optimality. □

In general, however, the ratios λ^z_{ij}/a^z_{ij} will not be monotonously decreasing, so that (PPC-HLP) and (PC-HLP) will not be equivalent and may lead to different optimal values. Despite this, we will now show that we can reformulate (PPC-HLP) with arbitrary multipliers λ as a (smaller) instance of (PC-HLP) with different multipliers μ. In a sense, our reduction incorporates the precedence constraints in the new multipliers μ.

Lemma 3. *Problem* (PPC-HLP) *polynomially reduces to Problem* (PC-HLP).

Proof. Let an instance of (PPC-HLP) with multipliers λ be given and let $\mathcal{L}_{x,\bar{y}}(\lambda)$ be its optimal value. In case the ratios λ_{ij}^z/a_{ij}^z are monotonously decreasing in z for all $i,j \in V$, the statement is trivially true due to Lemma 2. Therefore, assume that the ratios λ_{ij}^z/a_{ij}^z are not monotonously decreasing for $i,j \in V$ and define $k := \min\left\{z \in M_{ij} : \lambda_{ij}^z/a_{ij}^z < \lambda_{ij}^l/a_{ij}^l\right\}$ with $l := \min\{z \in M_{ij} : k < z\}$.

First, we show that any optimal solution $(x^*, \bar{y}^*)$ of (PPC-HLP) satisfies $(\bar{y}^*)_{ij}^k = (\bar{y}^*)_{ij}^l$. Since the precedence constraints forbid $(\bar{y}^*)_{ij}^k = 0$ and $(\bar{y}^*)_{ij}^l = 1$, assume $(\bar{y}^*)_{ij}^k = 1$ and $(\bar{y}^*)_{ij}^l = 0$. This means that k is the critical index and as such, its contribution $\lambda_{ij}^k - a_{ij}^k \tilde{\mathbf{c}}_{ij}(x^*)$ to the objective value must be non-negative – otherwise the solution with $y_{ij}^z = 1$ for all $z < k$ and $y_{ij}^k = 0$ would yield a strictly better objective value. By the definition of k, we derive $\lambda_{ij}^l/a_{ij}^l > \lambda_{ij}^k/a_{ij}^k \geq \tilde{\mathbf{c}}_{ij}(x^*)$. This contradicts the optimality of the solution, as a strictly better objective value could be achieved by setting $(\bar{y}^*)_{ij}^l := 1$. In summary, we have $(\bar{y}^*)_{ij}^k = (\bar{y}^*)_{ij}^l$.

As an immediate consequence of the previous observation, we can deduce that a new instance of (PPC-HLP) with multipliers μ and transport volumes $\tilde{a}$ as defined hereafter still yields $\mathcal{L}_{x,\bar{y}}(\lambda)$ as optimal value; the construction corresponds to merging the prizes k and l:

$$\mu_{ij}^k := \lambda_{ij}^k + \lambda_{ij}^l, \quad \tilde{a}_{ij}^k := a_{ij}^k + a_{ij}^l, \quad \text{and} \quad \mu_{ij}^z := \lambda_{ij}^z, \quad \tilde{a}_{ij}^z := a_{ij}^z \quad \forall z \in M_{ij} \setminus \{k,l\} .$$

Note that the sorting of the indices z does not change, although the index l no longer exists in this new instance. In case the ratios $\mu_{ij}^z/\tilde{a}_{ij}^z$ are monotonously decreasing in z, we may as well remove the precedence constraints from the problem and our claim follows thanks to Lemma 2. If not, we can repeat the reduction process. Since the number of indices z for the pair i,j decreases by one with each repetition, after at most $\sum_{i,j\in V} m_{ij}$ iterations we obtain an instance with monotonously decreasing ratios, and the claim follows from Lemma 2. □

Corollary 1. *Consider an instance of* (PPC-HLP) *with* λ^* *and construct an instance of* (PC-HLP) *with multipliers* μ^* *as in the proof of Lemma 3. Then* $\mathcal{L}_{x,\bar{y}}(\lambda^*) = \mathcal{Q}_{x,\bar{y}}(\mu^*)$ *and hence*

$$\text{OPT(PS-BHLP)} = \mathcal{L}_{p,y}(\lambda^*) + \mathcal{L}_{x,\bar{y}}(\lambda^*) = \mathcal{L}_{p,y}(\lambda^*) + \mathcal{Q}_{x,\bar{y}}(\mu^*) .$$

Note that the reduction in the proof of Lemma 3 not only allows us to remove the precedence constraints from (PPC-HLP), but in doing so, the resulting instance of (PC-HLP) even turns out to be a more compact problem. This is a significant advantage, considering the fact that the computation of $\mathcal{Q}_{x,\bar{y}}(\mu^*) = \mathcal{L}_{x,\bar{y}}(\lambda^*)$ is NP-hard, as we know from Theorems 2, 3 and 5.

Depending on the chosen solution approach, it might also make sense to solve a tractable relaxation of (PPC-HLP) or (PC-HLP), respectively, rather than the problem itself. We next show that the multipliers λ^* are still an optimal choice for the Lagrangian dual problem if we consider certain relaxations. As an auxiliary result, we first show that, in an optimal solution of Problem (CDP), any (reasonable) subset of customers can become the set of transporting customers if the price is set accordingly.

Lemma 4. *For all $i, j \in V$, choose any set $S_{ij} = \{1, \ldots, k_{ij}\}$ with $k_{ij} \in \{0, \ldots, m_{ij}\}$. Then an optimal solution to $\mathcal{L}_{p,y}(\lambda^*)$ is given by $p_{ij} := b_{ij}^{k_{ij}} / a_{ij}^{k_{ij}}$ and $y_{ij}^z := 1$ if and only if $z \in S_{ij}$.*

Proof. Again, we can argue independently for each pair $i, j \in V$. Clearly, the given choice of p_{ij} yields the given y_{ij}^z as follower's response. The contribution of pair i, j to the objective value of (CDP) is thus given by

$$\sum_{z \in M_{ij}} y_{ij}^z (a_{ij}^z p_{ij} - (\lambda^*)_{ij}^z) = \sum_{z=1}^{k_{ij}} \left(a_{ij}^z \frac{b_{ij}^{k_{ij}}}{a_{ij}^{k_{ij}}} \right) - \sum_{z=1}^{k_{ij}} (\lambda^*)_{ij}^z \overset{(4)}{=} \frac{b_{ij}^{k_{ij}}}{a_{ij}^{k_{ij}}} \sum_{z=1}^{k_{ij}} a_{ij}^z - \frac{b_{ij}^{k_{ij}}}{a_{ij}^{k_{ij}}} \sum_{z=1}^{k_{ij}} a_{ij}^z = 0 \, .$$

This is the optimal value of Problem (CDP) with λ^*, which concludes the proof. □

As the proof of Theorem 4 shows, the precedence prize-collecting hub location problem inherits its complexity from the classical hub location problem. Therefore, one might consider relaxations in which the feasible set of $\text{HLP}_{\text{single-alloc}}$ is replaced by a tractable superset. We claim that for such relaxations of (PPC-HLP), the multipliers λ^* remain optimal.

Theorem 6. *Replace the set* $\text{HLP}_{\text{single-alloc}}$ *in* (PPC-HLP) *with any superset and let $\mathcal{L}_{x,\bar{y}}^{relax}(\lambda)$ describe the optimal value of the relaxed version of the precedence prize-collecting problem with multipliers λ. Then we have $\mathcal{L}_{p,y}(\lambda^*) + \mathcal{L}_{x,\bar{y}}^{relax}(\lambda^*) \leq \mathcal{L}_{p,y}(\nu) + \mathcal{L}_{x,\bar{y}}^{relax}(\nu)$ for all ν, i.e., the best upper bound for* (PS-BHLP) *is still achieved with the multipliers λ^*.*

Proof. After reorganizing the terms, we need to show $\mathcal{L}_{x,\bar{y}}^{\text{relax}}(\nu) - \mathcal{L}_{x,\bar{y}}^{\text{relax}}(\lambda^*) + \mathcal{L}_{p,y}(\nu) - \mathcal{L}_{p,y}(\lambda^*) \geq 0$. We first consider the precedence prize-collecting problem (PPC-HLP) and its respective relaxation.

Let $(x^*, \bar{y}^*)$ be an optimal solution of the latter with the multipliers λ^*. Since the relaxation still contains both the precedence constraints and the binarity constraints of $\bar{y}$, we can define a critical index k_{ij} for any pair $i,j \in V$ with respect to $\bar{y}^*$ as before. Since the feasible set of the precedence prize-collecting problem does not depend on the Lagrangian multipliers, the solution $(x^*, \bar{y}^*)$ is also feasible for the problem with the Lagrangian multipliers ν instead of λ^* and thus gives a lower bound on the optimal value. This implies

$$\begin{aligned}
\mathcal{L}^{\text{relax}}_{x,\bar{y}}(\nu) - \mathcal{L}^{\text{relax}}_{x,\bar{y}}(\lambda^*) &\geq \sum_{i,j\in V}\sum_{z\in M_{ij}} (\bar{y}^*)^z_{ij}(\nu^z_{ij} - a^z_{ij}\tilde{\mathbf{c}}_{ij}(x^*)) - \sum_{i,j\in V}\sum_{z\in M_{ij}} (\bar{y}^*)^z_{ij}((\lambda^*)^z_{ij} - a^z_{ij}\tilde{\mathbf{c}}_{ij}(x^*)) \\
&= \sum_{i,j\in V}\sum_{z\in M_{ij}} (\bar{y}^*)^z_{ij}(\nu^z_{ij} - (\lambda^*)^z_{ij}) = \sum_{i,j\in V}\sum_{z=1}^{k_{ij}} (\nu^z_{ij} - (\lambda^*)^z_{ij}) \,.
\end{aligned}$$

Now take a closer look at the customers' decision problem (CDP). By Lemma 4, setting $y^* := \bar{y}^*$ and $p^*_{ij} := b^{k_{ij}}_{ij}/a^{k_{ij}}_{ij}$ yields an optimal solution for the customers' decision problem with multipliers λ^*. Moreover, the solution (p^*, y^*) remains feasible for the multipliers ν. This leads to

$$\begin{aligned}
\mathcal{L}_{p,y}(\nu) - \underbrace{\mathcal{L}_{p,y}(\lambda^*)}_{=0} &\geq \sum_{i,j\in V}\sum_{z\in M_{ij}} (\bar{y}^*)^z_{ij}(a^z_{ij}p^*_{ij} - \nu^z_{ij}) = \sum_{i,j\in V}\sum_{z=1}^{k_{ij}} (a^z_{ij}\frac{b^{k_{ij}}_{ij}}{a^{k_{ij}}_{ij}} - \nu^z_{ij}) \\
&= \sum_{i,j\in V}\Big(\frac{b^{k_{ij}}_{ij}}{a^{k_{ij}}_{ij}}\sum_{z=1}^{k_{ij}} a^z_{ij} - \sum_{z=1}^{k_{ij}}\nu^z_{ij}\Big) \overset{(4)}{=} \sum_{i,j\in V}\sum_{z=1}^{k_{ij}} ((\lambda^*)^z_{ij} - \nu^z_{ij}) \,.
\end{aligned}$$

The combination of the two inequalities yields the desired result. □

It remains to discuss whether Lemma 3 also remains valid when considering relaxations, i.e., whether it is still possible to remove the precedence constraints. For this, for any multiplier λ, let (PPC-HLP$^{\text{relax}}$) and (PC-HLP$^{\text{relax}}$) be the respective relaxations of (PPC-HLP) and (PC-HLP) in which the set $\text{HLP}_{\text{single-alloc}}$ has been replaced by some superset. Denote the corresponding optimal values by $\mathcal{L}^{\text{relax}}_{x,\bar{y}}(\lambda)$ and $\mathcal{Q}^{\text{relax}}_{x,\bar{y}}(\lambda)$, respectively. Then the relaxation (PPC-HLP$^{\text{relax}}$) polynomially reduces to (PC-HLP$^{\text{relax}}$) exactly the same way that (PPC-HLP) reduces to (PC-HLP), since none of the arguments in the proof of Lemma 3 depends on the specific structure of the feasible set for x. Thus, $\mathcal{L}_{p,y}(\lambda^*) + \mathcal{L}^{\text{relax}}_{x,\bar{y}}(\lambda^*) = \mathcal{L}_{p,y}(\lambda^*) + \mathcal{Q}^{\text{relax}}_{x,\bar{y}}(\mu^*) \leq \mathcal{L}_{p,y}(\nu) + \mathcal{L}^{\text{relax}}_{x,\bar{y}}(\nu)$ for all ν, where μ^* is obtained from λ^* through the reduction employed in the proof of Lemma 3.

3.3 Summary of the oracle-based algorithm

The results of the previous section give rise to the following solution algorithm for the price-setting hub-location problem (PS-BHLP), assuming that we have access to an oracle for (PC-HLP): First, determine the multipliers λ^* directly from the problem data, using Definition 1. Then compute the multipliers μ^* using Lemma 3 and call the oracle for (PC-HLP) with μ^*, i.e., solve $\mathcal{Q}_{x,\bar{y}}(\mu^*)$, yielding an optimal solution $(x^*, \bar{y}^*)$. Finally, compute $p^*_{ij} := b^{k_{ij}}_{ij}/a^{k_{ij}}_{ij}$, where k_{ij} again denotes the critical index with respect to $\bar{y}^*_{ij}$, for every pair $i, j \in V$. Then (x^*, p^*) is an optimal solution to (PS-BHLP). Since all these steps – except for the solution of $\mathcal{Q}_{x,\bar{y}}(\mu^*)$ – can be performed in polynomial time, the resulting algorithm represents a polynomial-time reduction from the price-setting problem (PS-BHLP) to the prize-collecting problem (PC-HLP).

If we are only equipped with an oracle for solving a relaxation of (PC-HLP), we can proceed analogously and thus implicitly solve a relaxation for (PS-BHLP). The latter can be embedded into a branch-and-bound framework.

4 Generalizations and limitations

In the previous sections, we presented an oracle-based approach to the price-setting hub location problem. While our emphasis was on the latter problem since this was our motivating application, the approach is, in fact, much more general. In particular, the oracle-based paradigm suggests a possible extension of this approach to larger problem classes. In the following, we investigate the conditions under which such extensions are possible. We first stick to (variants of) hub location problems in Section 4.1 and then consider more general optimization problems in Section 4.2.

The common observation underlying all generalizations is that none of our results actually depends on the structure of the feasible set $\text{HLP}_{\text{single-alloc}}$ of the x-variables in the price-setting problem (PS-BHLP). In particular, a closer inspection of our main results and the corresponding proofs presented in Section 3.2 reveals that all we need is a problem of the form

$$\left.\begin{aligned}
\max \quad & \sum_{s\in S}\sum_{z\in M_s} a_s^z y_s^z (p_s - \tilde{\mathbf{c}}_s(x)) - \tilde{\mathbf{c}}(x) \\
\text{s.t.} \quad & x \in X \\
& p_s \geq 0 \quad \forall s \in S \\
& y \in \arg\max \quad \sum_{s\in S}\sum_{z\in M_s} y_s^z \\
& \qquad\qquad \text{s.t.} \quad a_s^z p_s y_s^z \leq b_s^z \quad \forall s \in S, z \in M_s \\
& \qquad\qquad\qquad\;\; y_s^z \in \{0,1\} \quad\; \forall s \in S, z \in M_s
\end{aligned}\right\} \tag{PS-BP}$$

with a finite feasible set X of underlying structures, a finite set S of possible services, a set of potential customers $M_s = \{1,\dots,m_s\}$ for each $s \in S$, and an arbitrary function $\tilde{\mathbf{c}}_s \colon X \to \mathbb{R}$ defining the cost of providing one unit of service s in the structure defined by x. Additionally, we allow a function $\tilde{\mathbf{c}} \colon X \to \mathbb{R}$ describing setup costs of the structure x that are independent of the customers' decisions. The aim is then to optimize the structure $x \in X$ and the price p_s for each service $s \in S$ simultaneously, where the customers' behavior is identical to the previous sections. After Lagrangian decomposition, the customers' decision problem is thus essentially the same as before, while the second subproblem is again a precedence prize-collecting problem which can be reduced to a pure prize-collecting problem of the form

$$\left.\begin{aligned}
\max \quad & \sum_{s\in S}\sum_{z\in M_s} \bar{y}_s^z (\lambda_s^z - a_s^z \tilde{\mathbf{c}}_s(x)) - \tilde{\mathbf{c}}(x) \\
\text{s.t.} \quad & x \in X \\
& \bar{y}_s^z \in \{0,1\} \quad \forall s \in S, z \in M_s
\end{aligned}\right\} \tag{PC-P}$$

using the construction described in Lemma 3. Note that we do not require the functions $\tilde{\mathbf{c}}$ and $\tilde{\mathbf{c}}_s$ to be linear in the given bilevel problem (PS-BP), this was not even the case in our application discussed in the previous sections. However, these functions are passed on to the subproblem (PC-P), so that one needs to deal with them when addressing this subproblem. This again highlights the idea of our oracle-based approach, which is concerned with the reduction of (PS-BP) to (PC-P) rather than with solving (PC-P).

4.1 Hub location variants

We first discuss to what extent our results can be applied to other variants of the hub location problem. First, we can easily accommodate *setup costs* arising from the opening of a hub, i.e., terms of the form $\sum_{j\in V} f_{jj}x_{jj}$ in the objective function, using the function $\tilde{\mathbf{c}}$. The cardinality constraint $\sum_{j\in V} x_{jj} = k$ can of course be omitted if desired. On the other hand, it is not directly possible to include *capacity constraints* in the model (PS-BP). The reason is that the transport volumes depend on the y variables, i.e., on the customers' decisions. It is not possible to state such constraints only in terms of x, in contrast to the classical hub location problem, where the shipments are given and fixed, so that the transport volume on an arc or through a hub is uniquely determined by x. The same problem arises when considering *vehicle based costs*, i.e., piecewise constant or piecewise linear cost functions of the transport volume [19], since the objective function cannot be written as a sum of costs for each pair $i, j \in V$ then.

The situation is analogous when considering *multiple allocation* hub location problems, where our approach applies as well: While both a fixed number of hubs and setup costs fit into the framework of (PS-BP) also in the multi-allocation setting, after adapting the set X accordingly, we cannot directly apply our approach to the capacitated problem variants.

4.2 Other underlying problems

The approach devised in the previous sections is not limited to hub location problems. While in the latter, each service $s \in S$ corresponds to a pair $i, j \in V$, we may also consider problems where the service consists in delivering to a node $i \in V$. For example, the *vehicle routing problem (VRP)* [21] asks for the optimal routes of a fleet of vehicles to deliver to customers located in the nodes V, where the distances between pairs $i, j \in V$ are again given by values $c_{ij} \in \mathbb{R}_+$. In the prize-collecting variant of the VRP, each customer is willing to pay a certain price for the service and the optimizer may decide whether to deliver to a customer or not, depending on the complex interplay between the prize she can collect and the increase in costs when providing the service [3].

Now we can modify this problem and, instead of deciding for every single service demand separately, let the optimizer set a price for each service and the clients decide whether to accept this price or not. Then we are back in our setting and we can model this problem in the form of (PS-BP). For this, let X denote the feasible set of the prize-collecting VRP, i.e., each $x \in X$

corresponds to a feasible set of tours which together cover a certain set of nodes $V(x) \subseteq V$. Note that, in the prize-collecting variant, we do not necessarily have $V(x) = V$ for all $x \in X$. In terms of (PS-BP), the services $s \in S$ now correspond to nodes $v \in V$. On each node $v \in V$, a set of customers M_v is given, each of which has a certain budget for paying the delivery, and the collected prize may depend on the size of the delivery via a_s^z. We can then include all costs for the tours $x \in X$ into $\tilde{\mathbf{c}}(x)$ and set $\tilde{\mathbf{c}}_v(x) = 0$ if $v \in V(x)$ and $\tilde{\mathbf{c}}_v(x) = C$ otherwise, where C is large enough to make sure that a solution x can be optimal only if $y_v^z = 0$ for all $v \in V$ with $\tilde{\mathbf{c}}_v(x) = C$ and all $z \in M_v$. This essentially ensures that the chosen tours x can actually provide the service to all customers who decide to book it. Interestingly, the subproblem (PC-P) is a prize-collecting VRP again, with prizes depending on λ now.

In this example, we used the functions $\tilde{\mathbf{c}}_s$ to model the implicit constraints linking the chosen structure $x \in X$ to the services $s \in S$ that are feasible in this structure. Again, our model leaves a lot of room for defining $\tilde{\mathbf{c}}_s$, but ultimately one has to deal with $\tilde{\mathbf{c}}_s$ when solving the sub-problem (PC-P). However, as the example shows, the subproblem often has a lot of similarities with the underlying combinatorial optimization problem that we consider in the first place, before adding the customers' behavior to the model.

Using a very similar approach to that for the VRP, we can also address problems that do not deal with transport or logistics in the strict sense. It is easy to see that the above idea works for the prize-collecting Steiner tree problem as well. Here, the objective is to determine a tree in a given graph such that the cost of all selected edges minus the total prize collected in all nodes in the tree is minimized. In a typical application, a telecommunications company wants to connect customers to a network, but customers are willing to pay only a certain price for this. In our bilevel problem, the company does not directly decide which customer to connect but will set a price for every location and then connect all customers willing to pay the price. We believe that this is a much more realistic model than the prize-collecting model, while, as we have seen, the complexity of the problems is essentially the same and the former can be efficiently reduced to the latter.

5 Computational evaluation

As shown in Section 3, the price-setting problem (PS-BHLP) can either be reduced to a precedence prize-collecting hub location problem using the multipliers λ^* or to a more compact price-collecting hub location problem using the multipliers μ^*; both multipliers can be computed directly from the input of the problem. In the following, we present the results of an experimental evaluation comparing both approaches with each other and with a tailored single-level reformulation.

5.1 Single-level reformulation

In bilevel optimization, a classical approach is to exploit specific properties of the follower's problem in order to derive a single-level problem that is equivalent to the original bilevel problem. As the follower's problem in our situation is rather trivial to solve, it is not surprising that such an approach also works for (PS-BHLP). In fact, it turns out that the follower's problem can be replaced by rather simple upper-level constraints, as the following theorem shows:

Theorem 7. *Define* $P := \max_{i,j\in V} b^1_{ij}/a^1_{ij} + 1$ *and* $M := \max_{i,j\in V, z\in M_{ij}} a^z_{ij}\cdot P - \min_{i,j\in V, z\in M_{ij}} b^z_{ij}$. *Then, an equivalent single-level reformulation of Problem* (PS-BHLP) *is given by*

$$\left.\begin{aligned} \max \quad & \sum_{i,j\in V}\sum_{z\in M_{ij}} a^z_{ij} y^z_{ij}(p_{ij} - \tilde{\mathbf{c}}_{ij}(x)) && \\ s.t. \quad & x \in \mathrm{HLP}_{\text{single-alloc}} && \\ & 0 \le p_{ij} \le P && \forall i,j \in V \\ & a^z_{ij} p_{ij} - b^z_{ij} \le M(1 - y^z_{ij}) && \forall i,j \in V, z \in M_{ij} \quad (*) \\ & y^z_{ij} \in \{0,1\} && \forall i,j \in V, z \in M_{ij}\,. \end{aligned}\right\} \quad \text{(PS-HLP)}$$

Proof. First note that, according to Proposition 1, for any $p_{ij} > b^1_{ij}/a^1_{ij}$, we get $y^z_{ij} = 0$ for all $z \in M_{ij}$ as the follower's response, and hence, the contribution of the pair i, j in the objective function is zero, independently of the precise choice of p_{ij}. Therefore, we can include the constraint $p_{ij} \le P$ for all $i, j \in V$ in (PS-BHLP) without changing the optimal value.

Now consider a feasible solution (x, p, y) of (PS-BHLP). Then, by Proposition 1, the solution y of the follower's problem together with p satisfies the constraints $(*)$. Hence, the feasible set of (PS-BHLP) is a subset of the feasible set of (PS-HLP) and OPT(PS-BHLP) $\le$ OPT(PS-HLP).

For the other direction, consider an optimal solution (x, p, y) of Problem (PS-HLP). We show that there is a feasible solution of Problem (PS-BHLP) with at least the same objective value. For this, we consider some fixed pair $i, j \in V$ and distinguish two cases according to the sign of $p_{ij} - \tilde{\mathbf{c}}_{ij}(x)$. If $p_{ij} > \tilde{\mathbf{c}}_{ij}(x)$, the optimality of y implies $y^z_{ij} = 1$ as long as this is not forbidden by constraint $(*)$. In other words, we have $y_{ij}^z = 1$ if and only if $p_{ij} \le b_{ij}^z / a_{ij}^z$ and thus y_{ij}^z is an optimal response to p_{ij} according to Proposition 1. In particular, (x, p, y) is a feasible solution for (PS-BHLP). If $p_{ij} \le \tilde{\mathbf{c}}_{ij}(x)$, the contribution of pair i, j to the objective function is non-positive. However, in (PS-BHLP), it is feasible to choose p_{ij} large enough, e.g., $p_{ij} = P$, and set $y_{ij}^z = 0$ for all $z \in M_{ij}$, so that we obtain an objective value of zero for the pair i, j. In summary, we have shown that OPT(PS-HLP) $\le$ OPT(PS-BHLP). □

5.2 Linearization

In order to solve either the hub location subproblems (PPC-HLP) or (PC-HLP) in our approach or the single-level reformulation (PS-HLP), we first need to linearize the corresponding objective functions. In the classical hub location problem, the only term in need of linearization is the quadratic term $x_{ik}x_{jl}$ which occurs in $\tilde{\mathbf{c}}_{ij}(x)$. However, in our situation, the customer decision variables and their copies y and $\bar{y}$ lead to the quadratic terms $p_{ij}y_{ij}^z$ and the cubic terms $y_{ij}^z x_{ik}x_{jl}$ in the objective of (PS-HLP). In the (precedence) prize-collecting problem, we only have the cubic terms $\bar{y}_{ij}^z x_{ik}x_{jl}$. Inspired by Skorin et al. [20], we introduce non-negative linearization variables $X_{ijkl}^z \mathrel{\hat{=}} \bar{y}_{ij}^z x_{ik}x_{jl}$ with corresponding linear cost coefficients $C_{ijkl}^z := c_{ik} + \alpha c_{kl} + c_{lj}$, for all $i, j, k, l \in V$ and $z \in M_{ij}$, which indicate whether customer $z \in M_{ij}$ books the shipment of their product on the relation (i, j) where i is assigned to hub k and j is assigned to hub l. To enforce that the variable X_{ijkl}^z models the desired cubic term, we include the following set of conditions:

$$\sum_{k,l \in V} X_{ijkl}^z = \bar{y}_{ij}^z \quad \forall i, j \in V, z \in M_{ij}, \qquad \sum_{l \in V} X_{ijkl}^z \le x_{ik}, \sum_{k \in V} X_{ijkl}^z \le x_{jl} \quad \forall i, j, k \in V, z \in M_{ij} \tag{5}$$

It is now easy to verify the following

Theorem 8. *Let $x \in \mathrm{HLP}_{\text{single-alloc}}$. Then the constraints (5) ensure the binarity of X_{ijkl}^z for all $i, j, k, l \in V$ and $z \in M_{ij}$. Furthermore, $X_{ijkl}^z = 1$ if and only if $x_{ik} = x_{jl} = \bar{y}_{ij}^z = 1$.*

Applying this linearization technique to either (PPC-HLP) or (PC-HLP) results in the linear objective function $\sum_{i,j \in V} \sum_{z \in M_{ij}} \bar{y}_{ij}^z \lambda_{ij}^z - \sum_{i,j \in V} \sum_{z \in M_{ij}} a_{ij}^z C_{ijkl}^z X_{ijkl}^z$. For (PS-HLP), we apply the same technique to the cubic terms $y_{ij}^z x_{ik}x_{jl}$ and thus obtain $\sum_{i,j \in V} \sum_{z \in M_{ij}} a_{ij}^z \left(p_{ij}y_{ij}^z - C_{ijkl}^z X_{ijkl}^z\right)$ as objective function. Since we use the single-level reformulation (PS-HLP) only for comparison, we leave the linearization of the remaining quadratic terms $p_{ij}y_{ij}^z$ to the IP solver. While the linearization of the nonlinear terms originating from the hub location structure is common to all three solution approaches, the additional nonlinearity in (PS-HLP) is typical for single-level reformulations, but avoided in our approaches.

5.3 Experimental setup

Our experiments were carried out on a Dell Precision T5810 with an Intel Xeon E5-1630 processor running at 3.70 GHz and 128 GB RAM. We solved the respective linearized versions of (PPC-HLP), (PC-HLP), and (PS-HLP) with Gurobi 11.0.3 [11] using default parameters, except that the number of threads (`Threads`) was set to one and the time limit (`TimeLimit`) was set to one CPU hour. Moreover, we had to decrease the integer feasibility tolerance (`IntFeasTol`) to 10^{-9}, since we obtained infeasible solutions for some instances when applying the single-level reformulation with the default value of 10^{-5}.

Our test instances were constructed on the basis of the Civil Aeronautics Board (CAB) data set, which is freely available as part of the OR library [5]. It consists of 25 U.S. cities through which about half of the air traffic in the United States passes. For instances with fewer cities, we selected randomly from these 25 cities. We used scaled costs from the CAB dataset for our test instances, dividing the given costs between two cities by 1000. We considered a discount factor $\alpha \in \{0.5, 0.7\}$ and chose the number of hubs as $\kappa = 3$. For each customer $z \in \{1, ..., m\}$ on a relation i, j, we randomly selected their transport volume a_{ij}^z from $\{1, ..., 20\}$ and generated their budget as $b_{ij}^z := a_{ij}^z \cdot c_{ij} \cdot x_r$, where x_r is a random variable uniformly distributed in $[0.2, 5.0]$.

For solving the single-level reformulation (PS-HLP), we simply call Gurobi and measure the time required to compute an optimal solution. The running times we report for the oracle-based approaches using the prize-collecting problems (PPC-HLP) or (PC-HLP) also include the times for computing the respective multipliers λ^* or μ^*, which, however, turn out to be negligible.

5.4 Results

In the first experiment, we increased the number of nodes in the network from 10 to 25 and fixed the number of customers per relation to $m = 5$; results are presented in Table 1. In each entry of the table, we state the number of instances solved within the time limit, out of 10 random instances, and the average running time in CPU seconds over the solved instances, for each solution approach. The results consistently show that both oracle-based approaches require significantly less time than

Instance				(PS-HLP)		(PPC-HLP)		(PC-HLP)	
$\lvert V\rvert$	κ	α	m	#sol	time (s)	#sol	time (s)	#sol	time (s)
10	3	0.7	5	10	70.3	10	5.0	10	4.4
15	3	0.7	5	10	871.7	10	44.9	10	22.5
20	3	0.7	5	0	—	10	283.2	10	130.2
25	3	0.7	5	0	—	10	1675.4	10	1365.0
10	3	0.5	5	10	62.4	10	5.0	10	4.2
15	3	0.5	5	9	1114.4	10	51.4	10	26.0
20	3	0.5	5	0	—	10	299.1	10	160.9
25	3	0.5	5	0	—	10	663.3	10	396.8

Table 1: Results for instances with $m = 5$ customers per relation and an increasing number of nodes $|V|$; we state the number of instances solved within the time limit (#sol) and the average running times over all solved instances

the solution of (PS-HLP), but also that the solution based on (PC-HLP) clearly outperforms the solution using (PPC-HLP) All running times increase strongly with the number of nodes, which is not surprising in view of the NP-hardness of the problem being solved. The results also suggest that a larger discount, i.e., a smaller value of α, has a negative effect on running times.

In the second experiment, we increased the number of customers per relation from 3 to 11 while keeping the number of nodes at 10, since instances with more nodes proved to be difficult to solve for (PS-HLP) even for a small number of customers. Results are presented in Table 2. The general picture is very similar to Table 1: Again, (PPC-HLP) outperforms (PS-HLP), while the former is still outperformed by (PC-HLP). In all three approaches, it can be observed that the running time increases with the number of costumers. However, it appears that the single-level reformulation suffers the most from the larger number of costumers, whereas the running times of the oracle-based approaches seem to grow only proportionally to the number of customers.

In summary, our test results show that our Lagrangian decomposition approach is a powerful tool to solve the bilevel problem (PS-BHLP), as it is significantly faster than a more standard single-level reformulation approach in all cases considered. In Figure 2, we illustrate a typical optimal solution for an instance with $|V| = 7$ and $m = 3$.

6 Conclusion

In this paper, we considered a bilevel price-setting hub location problem modeling the complex interaction of two decision-making parties: The leader decides on a hub location network and

	BAL	BOS	CIN	DAL	DEN	LA	MEM
BAL	—	$p^* = \infty$	$p^* = 1.32$	$p^* = 3.27$	$p^* = 4.63$	$p^* = 6.04$	$p^* = 2.12$
	—	~~12~~ ~~5.38~~	6 11.72	11 64.46	17 105.53	20 135.89	12 25.44
	—	~~18~~ ~~4.89~~	20 28.71	4 20.01	12 68.41	7 42.31	~~5~~ ~~8.04~~
	—	~~14~~ ~~1.15~~	5 6.61	18 58.88	16 74.05	~~10~~ ~~40.70~~	~~17~~ ~~15.27~~
BOS	$p^* = \infty$	—	$p^* = 3.41$	$p^* = 4.99$	$p^* = 6.63$	$p^* = 10.89$	$p^* = 3.78$
	~~17~~ ~~10.20~~	—	18 61.44	12 59.84	11 79.06	14 152.52	13 66.49
	~~10~~ ~~5.88~~	—	~~8~~ ~~12.80~~	~~1~~ ~~3.41~~	10 66.57	~~9~~ ~~40.72~~	10 37.84
	~~3~~ ~~0.65~~	—	~~17~~ ~~22.19~~	~~8~~ ~~7.18~~	8 53.02	~~2~~ ~~8.58~~	~~9~~ ~~19.58~~
CIN	$p^* = 0.87$	$p^* = 3.46$	—	$p^* = 3.66$	$p^* = 2.06$	$p^* = 5.68$	$p^* = 1.67$
	6 6.51	6 20.73	—	10 36.62	1 3.80	11 71.60	8 13.35
	13 11.25	~~4~~ ~~9.02~~	—	~~7~~ ~~11.71~~	4 11.73	19 123.17	~~5~~ ~~5.30~~
	~~14~~ ~~4.35~~	~~1~~ ~~0.89~~	—	~~6~~ ~~9.72~~	13 26.79	6 34.11	~~14~~ ~~6.42~~
DAL	$p^* = 3.10$	$p^* = 6.72$	$p^* = 2.54$	—	$p^* = 2.76$	$p^* = 4.04$	$p^* = 1.08$
	15 46.51	14 106.72	10 39.20	—	2 5.53	4 24.88	17 18.33
	~~19~~ ~~22.38~~	14 94.05	3 9.91	—	~~4~~ ~~4.02~~	17 68.72	~~5~~ ~~2.86~~
	~~14~~ ~~9.94~~	~~3~~ ~~18.24~~	19 48.25	—	~~3~~ ~~2.97~~	~~17~~ ~~29.19~~	~~12~~ ~~6.81~~
DEN	$p^* = 6.15$	$p^* = 4.56$	$p^* = 0.79$	$p^* = 1.50$	—	$p^* = 3.84$	$p^* = 1.29$
	20 122.95	3 14.41	11 8.74	4 7.76	—	11 42.29	16 20.70
	~~16~~ ~~57.64~~	1 4.56	~~5~~ ~~3.82~~	10 14.95	—	~~12~~ ~~4.21~~	~~1~~ ~~0.98~~
	~~5~~ ~~9.95~~	~~7~~ ~~14.06~~	~~13~~ ~~8.49~~	~~19~~ ~~12.86~~	—	~~12~~ ~~3.09~~	~~7~~ ~~4.05~~
LA	$p^* = 7.91$	$p^* = 7.59$	$p^* = 7.55$	$p^* = 5.26$	$p^* = 3.42$	—	$p^* = 6.42$
	20 158.20	19 144.12	11 102.00	4 21.05	8 33.50	—	13 83.48
	~~2~~ ~~11.16~~	~~6~~ ~~29.90~~	15 113.23	~~7~~ ~~7.52~~	20 68.38	—	~~16~~ ~~33.64~~
	~~2~~ ~~7.21~~	~~6~~ ~~12.58~~	~~14~~ ~~13.67~~	~~3~~ ~~2.26~~	~~16~~ ~~7.54~~	—	~~5~~ ~~3.04~~
MEM	$p^* = 3.81$	$p^* = 5.14$	$p^* = 1.26$	$p^* = 1.52$	$p^* = 2.43$	$p^* = 7.46$	—
	19 72.37	15 78.46	17 21.38	19 35.06	8 19.47	7 52.23	—
	~~4~~ ~~11.87~~	5 25.71	~~3~~ ~~2.76~~	12 18.22	~~2~~ ~~1.05~~	~~2~~ ~~5.85~~	—
	~~4~~ ~~3.99~~	~~15~~ ~~8.83~~	~~16~~ ~~6.14~~	~~13~~ ~~4.89~~	~~15~~ ~~6.77~~	~~15~~ ~~27.86~~	—

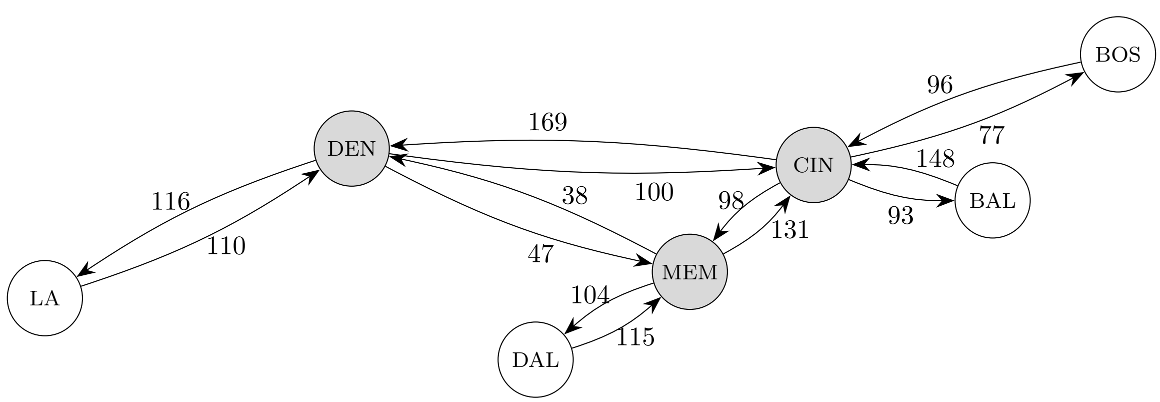

Figure 2: The optimal solution for an instance with seven cities and three customers per relation and $\alpha = 0.7$. The table (above) shows the optimal price p^*_{ij} for each pair i, j and, for each customer on the relation, their shipment volume and budget. The latter are crossed out if the respective customer cannot afford the service. The optimal hub network and the resulting shipping volumes on each arc are shown below.

Instance				(PS-HLP)		(PPC-HLP)		(PC-HLP)	
$\|V\|$	κ	α	m	#sol	time (s)	#sol	time (s)	#sol	time (s)
10	3	0.7	5	10	70.6	10	5.1	10	4.5
10	3	0.7	7	10	281.0	10	9.0	10	5.5
10	3	0.7	9	10	836.0	10	13.4	10	7.9
10	3	0.7	11	6	1740.5	10	18.4	10	8.9
10	3	0.7	15	0	—	10	29.0	10	10.2
10	3	0.7	20	0	—	10	49.9	10	14.2
10	3	0.7	50	—	—	10	161.4	10	30.1
10	3	0.7	100	—	—	10	561.5	10	71.9
10	3	0.5	5	10	62.6	10	4.9	10	4.1
10	3	0.5	7	10	299.5	10	7.6	10	5.2
10	3	0.5	9	10	947.1	10	10.9	10	6.3
10	3	0.5	11	4	2076.6	10	18.7	10	8.3
10	3	0.5	15	0	—	10	32.0	10	10.4
10	3	0.5	20	0	—	10	48.1	10	13.1
10	3	0.5	50	—	—	10	146.4	10	30.4
10	3	0.5	100	—	—	10	277.6	10	56.6

Table 2: Results for instances with $|V| = 10$ nodes and an increasing number of customers m; we state the number of instances solved within the time limit (#sol) and the average running times over all solved instances

prices that she asks for the transport per unit along each shipping route. The follower's problem models the behavior of customers that order the transport of their product along a specific shipment route as long as their individual budget allows for it.

We presented a Lagrangian decomposition approach which results in the efficiently solvable customers' decision problem on the one hand, and an NP-hard prize-collecting hub location problem variant on the other hand. From an abstract perspective, the decomposition thus produces a bilevel problem that is easy to solve and a single-level problem that inherits the complex underlying problem structure. More specifically, our decomposition approach still maintains a crucial part of the original problem structure in the prize-collecting hub location problem. While the latter problem in itself is interesting and, so far, has not been dealt with in the literature, it also allows for an oracle-based solution approach to the bilevel price-setting problem. Indeed, there are easy to compute optimal Lagrangian multipliers for the decomposition leaving no duality gap, which shows that the bilevel price-setting problem actually polynomially reduces to the prize-collecting hub location problem variant. This reduction was the key to a significantly faster solution of our problem instances compared to a more conventional single-level reformulation, as observed in our computational experiments. Beyond this, we showed that our solution approach extends to a much wider range of problems, with only minimal changes needed, due to the oracle-based paradigm.

References

[1] Sibel Alumur and Bahar Y. Kara. Network hub location problems: The state of the art. *European Journal of Operational Research*, 190:1–21, 10 2008.

[2] Sibel Alumur and Bahar Y. Kara. A hub covering network design problem for cargo applications in Turkey. *Journal of the Operational Research Society*, 60(10):1349–1359, 2009.

[3] Claudia Archetti, M. Grazia Speranza, and Daniele Vigo. Vehicle routing problems with profits. In Paolo Toth and Daniele Vigo, editors, *Vehicle Routing: Problems, Methods, and Applications*, chapter 10, pages 273–297. SIAM, 2014.

[4] Egon Balas. The prize collecting traveling salesman problem. *Networks*, 19(6):621–636, 1989.

[5] John E. Beasley. OR-Library – Hub location. `http://people.brunel.ac.uk/~mastjjb/jeb/orlib/phubinfo.html`, 2012.

[6] Christoph Buchheim. Bilevel linear optimization belongs to NP and admits polynomial-size KKT-based reformulations. *Operations Research Letters*, 51(6):618–622, 2023.

[7] Ivan Contreras, Jean-François Cordeau, and Gilbert Laporte. Stochastic uncapacitated hub location. *European Journal of Operational Research*, 212(3):518–528, 2011.

[8] Stephan Dempe. *Foundations of Bilevel Programming*. Springer New York, 2002.

[9] Güneş Erdoğan, Maria Battarra, and Antonio M. Rodríguez-Chía. The hub location and pricing problem. *European Journal of Operational Research*, 301(3):1035–1047, 2022.

[10] Reza Zanjirani Farahani, Masoud Hekmatfar, Alireza Boloori Arabani, and Ehsan Nikbakhsh. Hub location problems: A review of models, classification, solution techniques, and applications. *Computers & Industrial Engineering*, 64(4):1096–1109, 2013.

[11] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024.

[12] Pierre Hansen, Brigitte Jaumard, and Gilles Savard. New branch-and-bound rules for linear bilevel programming. *SIAM Journal on Scientific and Statistical Computing*, 13(5):1194–1217, 1992.

[13] Robert G. Jeroslow. The polynomial hierarchy and a simple model for competitive analysis. *Mathematical Programming*, 32(2):146–164, 1985.

[14] Bahar Y. Kara. *Modeling and analysis of issues in hub location problem.* PhD thesis, Bilkent University, 1999.

[15] Martine Labbé and Alessia Violin. Bilevel programming and price setting problems. *Annals of Operations Research*, 240:141–169, 2016.

[16] A. İrfan Mahmutoğulları and Bahar Y. Kara. Hub location under competition. *European Journal of Operational Research*, 250(1):214–225, 2016.

[17] George Nemhauser and Laurence Wolsey. *Integer and Combinatorial Optimization.* John Wiley & Sons, Inc., 1988.

[18] Hannah Porck. Decomposition approaches for a bilevel hub location problem. *Master's thesis, TU Dortmund University*, 2021.

[19] Borzou Rostami and Christoph Buchheim. The uncapacitated single allocation p-hub median problem with stepwise cost function. *Optimization Online*, 2015.

[20] Darko Skorin-Kapov, Jadranka Skorin-Kapov, and Morton O'Kelly. Tight linear programming relaxations of uncapacitated p-hub median problems. *European Journal of Operational Research*, 94(3):582–593, 1996.

[21] Paolo Toth and Daniele Vigo, editors. *The Vehicle Routing Problem.* Monographs on Discrete Mathematics and Applications. SIAM, Philadelphia, PA, 2002.