

Meta Partial Benders Decomposition for the Logistics Service Network Design Problem

Simon Belieres¹, Mike Hewitt², Nicolas Jozefowicz³, and Frédéric Semet⁴

¹*CNRS, LAAS, 7 Avenue du Colonel Roche, 31077 Toulouse Cedex 4, France*

²*Quinlan School of Business, Loyola University, 16 E. Pearson Ave., IL, Chicago 60611, USA*

³*LCOMS EA 7306, Université de Lorraine, Metz 57000, France*

⁴*Université de Lille, CNRS, Centrale Lille, Inria, UMR 9189 - CRISTAL, F-59000 Lille, France*

Abstract

Supply chain transportation operations often account for a large proportion of product total cost to market. Such operations can be optimized by solving the Logistics Service Network Design Problem (LSNDP), wherein a logistics service provider seeks to cost-effectively source and fulfill customer demands of products within a multi-echelon distribution network. However, many industrial settings yield instances of the LSNDP that are too large to be solved in reasonable run-times by off-the-shelf optimization solvers. We introduce an exact Benders decomposition algorithm based on partial decompositions that strengthen the master problem with information derived from aggregating subproblem data. More specifically, the proposed Meta Partial Benders Decomposition intelligently switches from one master problem to another by changing both the amount of subproblem information to include in the master as well as how it is aggregated. Through an extensive computational study, we show that the approach outperforms existing benchmark methods and we demonstrate the benefits of dynamically refining the master problem in the course of a partial Benders decomposition-based scheme.

Keywords— Logistics, Service Network Design, Supply Chain, Benders Decomposition

1 Introduction

Supply chains are complex networks of partner stakeholders that create products and distribute them to a consumer market. These stakeholders can be decomposed into echelons, wherein each echelon group stakeholders that serve similar roles. A major share of the supply chain operating costs is due to freight transportation within the multi-echelon distribution network that links stakeholders. As a result, cost-effective transportation operations are essential to ensure profitability. An increasingly common trend to reduce distribution costs is to outsource supply chain management activities to third-party logistics (3PL) service providers [Marasco, 2008] that specialize in the integration of warehousing and transportation services. In this paper, we present a new algorithmic strategy for solving a transportation problem encountered by a 3PL partner in the management of a restaurant supply chain.

In the problem studied, restaurants place orders of products for the coming weeks and the 3PL must design a transportation plan to fulfill these orders. As the products ordered by the restaurants are generic, i.e. they can be shipped by different suppliers of the supply chain, designing the transportation plan requires sourcing the products ordered by the restaurants and routing these products through a multi-echelon distribution network. This planning process can be assisted through the solution of the Logistics Service Network Design Problem (LSNDP). The LSNDP is a tactical transportation problem with few dedicated studies. Dufour et al. [Dufour et al., 2018] present a case study of the United Nations Humanitarian Response Depot, a humanitarian logistics service provider. They study a supply chain that aims to distribute relief goods and equipments in East Africa and they present a logistics service network design model for optimizing the transportation operations. They solve instances based on different network configurations and they demonstrate that substantial savings can be achieved by integrating a new distribution center. Belieres et al. [Belieres et al., 2020a] describe the LSNDP studied in this article. They investigate how the distribution strategy used by the 3PL can negatively impact overall logistics costs and they assess the savings that can be generated by extending this distribution strategy.

Belieres et al. [Belieres et al., 2020b] propose an effective Benders decomposition-based algorithm to solve this problem. Their strategy includes valid inequalities, heuristic solutions, as well as an advanced decomposition strategy inspired by the recently-proposed *Partial Benders Decomposition* (PBD). Introduced in the context of solving two-stage stochastic programs, the PBD retains some subproblem structure into the master problem to reduce the number of iterations and accelerate the algorithm convergence. In a similar fashion, Belieres et al. [Belieres et al., 2020b] propose to reinforce the master problem by including variables and constraints that model the need to satisfy customer demands of a “super-product” that aggregates all the products to be routed. The authors demonstrate that strengthening the master problem with this aggregated information enables an enhanced Benders decomposition-based algorithm to produce provably high-quality solutions in much less time than a general-purpose mixed-integer programming solver. In this paper, we present multiple PBD-related enhancements to the algorithm presented in [Belieres et al., 2020b], yielding an algorithm that exhibits much better computational performance.

First, whereas Belieres et al. [Belieres et al., 2020b] consider a single super-product, we consider multiple. Specifically, we propose determining a set of super-products and then adding variables and constraints to the master problem related to routing each super-product. To determine a set of super-products of a given cardinality, we propose partitioning the set of products and associating a super-product with each set of the partition. In addition, we characterize settings wherein the set of products can be partitioned to form super-products that yield a master problem provably equivalent to the original problem. Namely, a master problem that enables a Benders decomposition-based algorithm to converge in a single iteration. We use that characterization as the basis of a strategy for partitioning the set of products into a given number of subsets, and resulting super-products, in more general settings.

One would expect a trade-off associated with using multiple super-products to formulate the master problem. On the one hand, a greater number of super-products should yield a stronger master problem and enable the algorithm to converge in fewer iterations. For example, considering more super-products will increase the likelihood that the resulting master problem is provably equivalent to the original problem. On the other hand, a greater number of super-products will likely yield a master problem that is harder to solve, and thus the algorithm will spend more time executing each iteration. However, it is not yet known how to accurately estimate the impact of either on the execution of the resulting Benders decomposition-based algorithm.

Thus, we propose a new algorithmic strategy that exploits this feature by intelligently varying the amount of subproblem information used to formulate the master problem. We refer to this strategy as Meta Partial Benders Decomposition (Meta-PBD). Meta-PBD operates in two phases, with the first aiming to explore different areas within the feasible region of the original problem and to quickly determine high-quality solutions. To do so, the number of super-products used to formulate the master problem changes dynamically while respecting a threshold value to ensure computational tractability. The second phase aims to close the optimality gap. To do so, it drastically increase the amount of subproblem information used to formulate the master problem, such that the latter becomes provably equivalent to the original problem. While this master problem is computationally challenging, high-quality solutions generated through the first phase enable many nodes of the resulting branch-and-bound search tree to be pruned, accelerating convergence.

Our contribution is twofold. First, we provide theoretical results on how to aggregate subproblem information to determine a strong partial Benders decomposition for the LSNDP. While the results presented are specific to the LSNDP, they can easily be extended to other multi-product network design problems. We also investigate how the level of aggregation of the subproblem information impacts the resulting partial Benders decomposition. To the best of our knowledge, this has not been studied in the literature. Second, we introduce a new exact algorithmic strategy for solving the LSNDP. Through an extensive series of experiments performed on random instances, we show that Meta-PBD determines provably high-quality solutions in limited running times and strictly outperforms the algorithm proposed by Belieres et al. [Belieres et al., 2020b], including solving more instances to optimality. We also demonstrate the effectiveness of our approach on a set of industrial instances [Belieres et al., 2020a].

The remainder of the paper is organized as follows. In Section 2, we review the relevant literature. In Section 3, we present a description and formulation of the Logistics Service Network Design Problem. In Section 4, we present a Benders decomposition-based algorithm for this problem as well as different super-product-based master problems. In Section 5, we present the Meta-PBD algorithm. In Section 6, we assess the performance of Meta-PBD with a series of computational experiments. In Section 7, we draw conclusions and discuss future work.

2 Literature review

In this section, we first review the literature relevant to the Logistics Service Network Design Problem (LSNDP). Then, we review the literature documenting advanced decomposition strategies for enhancing Benders decomposition-based algorithms.

Since the LSNDP [Dufour et al., 2018, Belieres et al., 2020b] has received little attention to date, we position it regarding existing network design problems. The LSNDP is similar to the Service Network Design Problem (SNDP) [Crainic, 2000, Wieberneit, 2008] in that both problems focus on transportation planning decisions within a facility network. However, the LSNDP considers supply chains wherein products are made by multiple suppliers and thus fulfilling customer orders involves a sourcing decision. This is in contrast to most variants of the SNDP studied in the literature, wherein the origin of goods to be transported is presumed to be given. Besides, the SNDP makes no presumptions regarding the structure of the network, while the LSNDP presumes a multi-echelon structure. On the other hand, the LSNDP is similar to supply chain optimization problems [Beamon, 1998] such as the Logistics Network Design Problem (LNDP) [Srivastava, 2008] and the Supply Chain Network Design Problem (SCNDP) [Melo et al., 2009], which also aim to optimize the sourcing and fulfillment of orders through a multi-echelon distribution network. Nevertheless, supply chain optimization problems primarily focus on

strategic planning decisions such as facility location [Cordeau et al., 2006, Cheong et al., 2007], whereas the LSNDP assumes the facilities in a network, and their capacities, have already been established. In addition, the LSNDP seeks to optimize transportation costs that are generally not estimated accurately in supply chains optimization problems, since these problems usually cover long-term planning horizons that are discretized in monthly or yearly periods. For more details regarding the positioning of the LSNDP with regard to the SNDP and the LNDP, we refer the interested reader to the literature review proposed in [Belieres et al., 2020b]

To solve the LSNDP, we propose a solution approach based on Benders decomposition. Proposed by J.F. Benders [Benders, 1962] in 1962, Benders decomposition is designed for large-scale optimization problems that possess a certain feature. Namely, that the problem involves variables, often referred to as complicating variables, that, when fixed, yield (sub)problems that are significantly less difficult to solve computationally-speaking. Instead of solving one large Mixed Integer Linear Programs (MILP), Benders decomposition separates the problem into a master problem that is solved to determine the values of the complicating variables, and one or multiple subproblems formulated by fixing the complicating variables to the values of the most recent master problem solution. When subproblems are feasible, and variable values from solutions to the master problem and subproblems can be used to form a provably optimal solution to the original problem, the algorithm stops. When they cannot, Benders cuts derived from the subproblems are used to strengthen the master problem, and the process repeats.

The Benders decomposition algorithm has been used to solve a wide range of optimization problems, including network design problems [Costa, 2005]. However, only implementing the steps of Benders decomposition that guarantee convergence often yields an algorithm that requires a substantial amount of time and memory before converging. As a result, many enhancement strategies have been proposed to accelerate the method. Instead of reviewing those here, we refer to [Rahmaniani et al., 2017], which provides a taxonomy of these acceleration techniques. That taxonomy includes advanced algorithms for solving the master problem and the subproblem (e.g. [Benoist et al., 2002, Cordeau et al., 2001]), approaches to generate solutions and constraints from those solutions (e.g. [Costa et al., 2012, Magnanti and Wong, 1981, Rei et al., 2009]), and advanced decomposition strategies, i.e. (e.g. [Gendron et al., 2016]). Contrary to standard decompositions, where all the linking constraints and non complicating variables are projected out, advanced decomposition strategies retain part of the subproblem information in the master problem, which can considerably improve the convergence of the algorithm. Indeed, solution algorithms based on standard Benders decomposition often struggle computationally-speaking, as partitioning complicating and non-complicating variables can lead to hide many of the underlying structures inherent to the original optimization problem. These algorithms instead have to iteratively re-discover representations of that structure by repeatedly solving master and sub-problems. In addition, many optimization solvers have effective techniques for detecting and exploiting known structures in optimization problems. On the other hand, advanced decomposition strategies retain subproblem information in the master problem, which reinforces the bounds it produces and allows to reduce the number of Benders iterations.

The exact Benders decomposition algorithm introduced in this study relies on such an advanced decomposition strategy: the *Partial Benders Decomposition* proposed by Crainic et al. [Crainic et al., 2014, Crainic et al., 2016]. Proposed in the context of solving two-stage stochastic programs, the partial decomposition technique consists in including explicit information from the scenario sub-problems in the master problem. Through an extensive computational study, the authors demonstrate that Benders schemes based on partial decompositions rather than standard decompositions yields significant improvements in terms of stability of the solution process and computational time. Recently, partial

decomposition techniques has also been applied in the context of solving deterministic programs. In that case, the information used to strengthen the master problem is obtained by aggregating data related to the subproblem. Fontaine et al. [Fontaine et al., 2017] address a multi-modal variant of the SNDP and propose a Benders decomposition algorithm wherein the master problem is reinforced by considering an aggregation of capacity over compartments of each service. Belieres et al. [Belieres et al., 2020b] propose a similar type of enhanced Benders decomposition algorithm for solving the LSNDP. In this approach, the master problem is reinforced with variables and constraints that model the need to route a super-product derived from the aggregation of all the products to be transported. An extensive computational study shows this new master problem yields significantly stronger lower bounds.

3 Problem definition and mathematical model

In this section, we first define the problem considered in this paper. We then present our mathematical formulation of that problem.

3.1 Problem definition

We study a tactical transportation planning problem that is inspired by the operations of a distribution company supporting restaurant supply chains. Specifically, the distribution company plans product distribution from suppliers to customers over a fixed planning horizon. For more details on the practical application, we refer the interested reader to [Belieres et al., 2020a].

On the supply side, products are provided by a set of suppliers in different locations, and are classified into product families. Examples of product families in restaurant supply chains are meats, beverages, fruits, vegetables, and paper products. Each supplier specializes in a subset of product families and only provides products from those families in which it specializes. However, a supplier that specializes in a product family is not required to produce all the products in that family. For example, the fruit family may contain both bananas and apples but a fruit supplier may supply apples but not bananas. In the industrial setting that inspired this research there are seven product families and each supplier specializes in at most three of them.

On the demand side, product deliveries are requested by customers, which are restaurants in the industrial problem that motivated this research. To facilitate their inbound logistics operations, each customer has a periodic schedule that specifies the days and time windows during which they can receive products. For example, restaurant A may request delivery of products every Monday between 6 and 7 a.m. Note that deliveries must occur during the specified time window, as early or delayed deliveries are not allowed. The types and quantities of products ordered do not have to be the same from one request to another. For example, restaurant A may request one pallet of bananas and one pallet of wine for the first Monday of the month, and one pallet of broccoli for the second Monday of the month.

Restaurants place product orders that specify the delivery day and time window but not the supplier for each product. Thus, the distribution company must determine how to source orders. This in turn implies that a single customer order that consists of multiple products may be sourced from multiple suppliers. For example, to fulfill restaurant A's order for the first Monday of the month, the distribution company may decide to ship one pallet of bananas from supplier 1 and one pallet of wine from supplier 2. We presume that restaurant orders are known far enough in advance that suppliers can design production plans to avoid stockouts.

The distribution company can ship products directly from suppliers to customers. However, products are small relative to vehicle capacity. Thus, to consolidate orders and increase vehicle utilization, the distribution company may instead transport products through a distribution network that links suppliers with customers. Intermediate terminals within this distribution network are referred to as *Warehouses*. Each warehouse can store a limited amount of product, but doing so incurs a per-unit, per-unit-of-time cost. A vehicle dispatched from a supplier to a warehouse or from a warehouse to another warehouse can carry products destined for different customers. However, for this industrial partner, a vehicle dispatched to a customer can only transport products intended for that customer. Figure 1 represents an example of such a distribution network, where S_x nodes are supplier facilities, C_x nodes are customer locations, and W_x nodes are warehouses within the distribution network.




Figure 1: Distribution network

Transportation between a pair of terminals is denoted as a *service*. A service is defined by a departure terminal, a departure time, an arrival terminal, and an arrival time. Using a service requires the assignment of transportation units to that service. Each unit has a fixed capacity and a cost. Thus, the distribution company determines the services to execute and the number of vehicles to dispatch on each service.

We focus on a situation wherein the distribution company is an asset-less logistics service provider and outsources transportation operations to a third party carrier. The distribution company develops a transportation plan for moving products and communicates the resulting needs for point-to-point transportation moves to the carrier. Therefore, the distribution company does not make vehicle fleet management decisions such as planning empty moves. Similarly, the distribution company assumes that the carrier's fleet is large enough to satisfy the services it wants to have performed. In total, the third-party logistics company seeks to determine the shipment origin of each product requested, as well as the services and transportation capacities needed to support deliveries. Their objective is to minimize the overall cost of product transportation and storage.

3.2 Mathematical model

We model the supply chain with a directed network, $\mathcal{G} = (\mathcal{N}, \mathcal{A})$ with node set \mathcal{N} and arc set \mathcal{A} . The node set \mathcal{N} is composed of nodes that model suppliers \mathcal{S} , customers \mathcal{C} , and warehouses \mathcal{W} . The arc set \mathcal{A} is composed of arcs that model transportation between pairs of nodes. Due to the multi-echelon structure of the supply chain, \mathcal{A} does not contain arcs that model transportation to a supplier, or arcs that model transportation from a customer. A transportation arc can only be defined in one of the following situations: i) from a supplier to a warehouse; ii) from a supplier to a customer; iii) from a warehouse to another warehouse; iv) from a warehouse to a customer. With each arc $a = (i, j) \in \mathcal{A}$ is associated a travel time $t_{ij} \in \mathbb{R}^{+*}$, a per unit of flow cost $c_{ij} \in \mathbb{R}^{+*}$, a vehicle capacity \hat{u} , and a fixed cost per vehicle, $f_{ij} \in \mathbb{R}^{+*}$.

The set of all products ordered by customers is denoted by \mathcal{P} . The set of products manufactured by supplier i is denoted by \mathcal{P}^i . We denote the set of product families as $\mathcal{F} = \{\mathcal{F}_1, \dots, \mathcal{F}_{|\mathcal{F}|}\}$. As each product belongs to a single family, product families form a partition of the product set. Formally, we have that $\forall (i, j) \in \{1, \dots, |\mathcal{F}|\}^2, i \neq j, \mathcal{F}_i \cap \mathcal{F}_j = \emptyset$, and $\bigcup_{i \in \{1, \dots, |\mathcal{F}|\}} \mathcal{F}_i = \mathcal{P}$.

The distribution company determines a transportation plan over a time horizon of length \mathcal{T} . To model the time dimension, we consider a time-expanded network $\mathcal{G}_{\mathcal{T}} = (\mathcal{N}_{\mathcal{T}}, \mathcal{H}_{\mathcal{T}} \cup \mathcal{A}_{\mathcal{T}})$ derived from the network \mathcal{G} . $\mathcal{G}_{\mathcal{T}}$ contains a time-expanded node $(i, t) \in \mathcal{N}_{\mathcal{T}}$, for each $i \in \mathcal{N}$ and $t \in \mathcal{T}$. Thus, time-expanded nodes model either time-expanded suppliers $\mathcal{S}_{\mathcal{T}}$, time-expanded customers $\mathcal{C}_{\mathcal{T}}$ or time-expanded warehouses $\mathcal{W}_{\mathcal{T}}$. Arcs in $\mathcal{H}_{\mathcal{T}}$ model product storage at warehouses. Thus, for each $i \in \mathcal{W}$ and each $t \in [1, |\mathcal{T}| - 1]$, there is a time-expanded arc $((i, t), (i, t + 1))$ in $\mathcal{H}_{\mathcal{T}}$ with a per-unit-of-flow cost c_{ii} and a storage capacity $wlim_i$. Arcs in $\mathcal{A}_{\mathcal{T}}$ model transportation between different locations, taking account of the travel time. To construct these arcs, for each $(i, j) \in \mathcal{A}$ and each time $t \in \mathcal{T}$ such that $t + t_{ij} < |\mathcal{T}|$, we build a time-expanded arc $((i, t), (j, t + t_{ij}))$. Therefore, an arc $((i, t), (j, t + t_{ij}))$ in $\mathcal{A}_{\mathcal{T}}$ models the flow of products departing from i at time t , and arriving to j at time $t + t_{ij}$. We note that before creating the time-expanded graph $\mathcal{G}_{\mathcal{T}}$, arc travel times t_{ij} may need to be modified to ensure that arcs $(i, j) \in \mathcal{A}$ can be mapped to arcs of the form $((i, t), (j, t + t_{ij}))$. As an example, for an arc (i, j) with $t_{ij} = 5h$, a 3-hours discretization requires to round up (i, j) travel time to the nearest multiple of the time-step, i.e. 6 hours. Consequently, discretizing the planning horizon may introduce travel time approximations that tend to reduce with finer time-steps.

Thus, we formulate the Logistics Service Network Design problem defined over a time-expanded graph $\mathcal{G}_{\mathcal{T}}$. For each transportation arc $((i, t), (j, t')) \in \mathcal{A}_{\mathcal{T}}$, the integer variable $y_{ij}^{tt'}$ corresponds to the number of trucks dispatched. For each transportation/holding arc $((i, t), (j, t')) \in \mathcal{H}_{\mathcal{T}} \cup \mathcal{A}_{\mathcal{T}}$, the continuous variable $x_{ij}^{ptt'}$ models the flow of product p on that arc. Note that variable $x_{ij}^{ptt'}$ is not defined if $((i, t), (j, t'))$ is a transportation arc originating from a supplier i that does not manufacture product p (i.e. $p \notin \mathcal{P}^i$). Customer demands are represented by d_{ct}^p , which is the amount of product $p \in \mathcal{P}$ requested by customer $c \in \mathcal{C}$ to be delivered at time $t \in [1, |\mathcal{T}|]$.

It should be noted that the formulation of the LSNDP below explicitly references product families. This is technically not necessary. However, as we leverage this problem characteristic in the Benders decomposition-based algorithm we present later in this paper, we include them in this formulation. Ultimately, the LSNDP can be stated as:

$$\text{minimize} \quad \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} f_{ij} y_{ij}^{tt'} + \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} \sum_{p \in \mathcal{P}} c_{ij} x_{ij}^{ptt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_{\mathcal{T}}} \sum_{p \in \mathcal{P}} c_{ii} x_{ii}^{ptt+1} \quad (1)$$

subject to the following constraints :

$$\sum_{((i,t),(j,t')) \in \mathcal{A}_T \cup \mathcal{H}_T} x_{ij}^{ptt'} - \sum_{((j,t'),(l,t'')) \in \mathcal{A}_T \cup \mathcal{H}_T} x_{jl}^{pt't''} = 0, \quad \forall (j,t') \in \mathcal{W}_T, \forall p \in \mathcal{F}_i, \forall \mathcal{F}_i \in \mathcal{F} \quad (2)$$

$$\sum_{((i,t),(j,t')) \in \mathcal{A}_T} x_{ij}^{ptt'} \geq d_{jt'}^p, \quad \forall (j,t') \in \mathcal{C}_T, \forall p \in \mathcal{F}_i, \forall \mathcal{F}_i \in \mathcal{F} \quad (3)$$

$$\sum_{p \in \mathcal{P}} x_{ii}^{ptt+1} \leq \text{lim}_i^{\mathcal{W}}, \quad \forall ((i,t),(i,t+1)) \in \mathcal{H}_T \quad (4)$$

$$\sum_{p \in \mathcal{P}} x_{ij}^{ptt'} \leq \hat{u} y_{ij}^{tt'}, \quad \forall ((i,t),(j,t')) \in \mathcal{A}_T \quad (5)$$

$$x_{ij}^{ptt'} \in \mathbb{R}^+, \quad \forall ((i,t),(j,t')) \in \mathcal{A}_T \cup \mathcal{H}_T, \forall p \in \mathcal{P}^i \quad (6)$$

$$y_{ij}^{tt'} \in \mathbb{N}^+, \quad \forall ((i,t),(j,t')) \in \mathcal{A}_T \quad (7)$$

The LSNDP seeks to minimize the transportation plan overall cost (1), which is the sum of trucks fixed costs on transportation arcs (first term), and flow linear costs on transportation arcs (second term) and holding arcs (third term). Flow feasibility is ensured by the two first constraints. Constraints (2) enforce the flow conservation at each warehouse. Constraints (3) enforce that each customer demands are fulfilled. Constraints (4) limit the total amount of product stored by each warehouse. Constraints (5) ensure that a sufficient number of trucks are allocated to transport products. Variable domains are defined by (6) and (7).

4 Benders decompositions for the LSNDP

In this paper, we present a Partial Benders Decomposition-based algorithm that intelligently switches between master problems strengthened with different aggregated subproblem informations. As such, this section focuses on different master problems for solving the LSNDP with a Benders decomposition-based algorithm. We begin this section by presenting a Benders decomposition for the problem based upon the standard master problem. We then introduce a new master problem that is reinforced with variables and constraints that model the need to route \mathcal{K} super-products derived from a \mathcal{K} -partition of the product set. We study the theoretical properties of such a master problem, including establishing that it is a relaxation of the original problem. We demonstrate that the way products are partitioned impacts the strength of the resulting master problem. Lastly, we discuss the effect that the number of super-products \mathcal{K} has on the strength of the resulting master problem.

4.1 Standard Benders decomposition for the LSNBP

In the context of the LSNBP, a standard Benders decomposition yields a master problem that allocates vehicles on transportation arcs. The resulting subproblem aims to route products from suppliers to customers, while ensuring that the total flow on each transportation arc does not exceed the capacity defined by the solution to the master problem. Based on the set of extreme rays of the subproblem dual polyhedron, Ω , and the set of extreme points of the subproblem dual polyhedron, Γ , the standard master problem, **SMP**, is formulated as follows:

$$\min \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} f_{ij} y_{ij}^{tt'} + z \quad (8)$$

$$0 \geq \sum_{(c,t) \in \mathcal{C}_{\mathcal{T}}} \sum_{p \in \mathcal{P}} d_{ct}^p \rho_{ct}^p + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_{\mathcal{T}}} \lim_i^{\mathcal{W}} \rho_{ii}^{tt'} - \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} \hat{u} \rho_{ij}^{tt'} y_{ij}^{tt'}, \quad \forall \rho \in \Omega \quad (9)$$

$$z \geq \sum_{(c,t) \in \mathcal{C}_{\mathcal{T}}} \sum_{p \in \mathcal{P}} d_{ct}^p \pi_{ct}^p + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_{\mathcal{T}}} \lim_i^{\mathcal{W}} \pi_{ii}^{tt'} - \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} \hat{u} \pi_{ij}^{tt'} y_{ij}^{tt'}, \quad \forall \pi \in \Gamma \quad (10)$$

$$y_{ij}^{tt'} \in \mathbb{N}^+, \quad \forall ((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}} \quad (11)$$

$$z \in \mathbb{R}^+ \quad (12)$$

The objective function, (8), computes costs related to the allocation of the vehicle fleet as well as an estimate of product routing costs. Constraints (9) and (10) are respectively the feasibility and optimality standard Benders cuts added dynamically after solving the subproblem.

Based on an allocation of vehicles \bar{y} , the subproblem **SP**(\bar{y}) is formulated as follows:

$$\min \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} \sum_{p \in \mathcal{P}} c_{ij} x_{ij}^{ptt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_{\mathcal{T}}} \sum_{p \in \mathcal{P}} c_{ii} x_{ii}^{ptt+1} \quad (13)$$

$$(2)-(3)-(4)$$

$$\sum_{p \in \mathcal{P}} x_{ij}^{ptt'} \leq \hat{u} \bar{y}_{ij}^{tt'}, \quad \forall ((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}} \quad (14)$$

$$x_{ij}^{ptt'} \in \mathbb{R}^+, \quad \forall ((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}} \cup \mathcal{H}_{\mathcal{T}}, \quad \forall p \in \mathcal{P}^i \quad (15)$$

Given a vehicle allocation \bar{y} , the subproblem seeks to route products from suppliers to customers at a minimal cost, while ensuring that the flow variables respect constraints (2)-(3)-(4) from the original problem. Constraint (14) is defined on each transportation arc and imposes that the total flow cannot exceed the available capacity.

4.2 A master problem based on super-products

The principle behind *Partial Benders Decomposition* [Crainic et al., 2014, Crainic et al., 2016] is to strengthen the master problem with information derived from the subproblem. Belieres et al. [Belieres et al., 2020b] apply this principle in the context of solving the LSNDP by reinforcing the standard master problem with variables and constraints related to routing a single super-product obtained by aggregating all products $p \in \mathcal{P}$. We extend this approach to multiple super-products. In particular, given a partition of the product set into subsets, we propose a master problem obtained by aggregating each product subset into its own super-product. We next present this master problem.

Let $\{\mathcal{P}_1, \dots, \mathcal{P}_K\}$ be a \mathcal{K} -partition of the product set \mathcal{P} . We create the super-product χ_k by aggregating all products included in subset \mathcal{P}_k . For a given customer c at time t , the demand of super-product χ_k is the sum of the demands for all products in \mathcal{P}_k , i.e. $D_{ct}^{\chi_k} = \sum_{p \in \mathcal{P}_k} d_{ct}^p$. A super-product flow variable

$x_{ij}^{\chi_k tt'}$ is defined for each arc $((i, t), (j, t'))$ and each product $p \in \mathcal{P}_k$ such that a flow variable $x_{ij}^{p tt'}$ is defined in the LSNDP. Thus, for a supplier i that does not supply any product of \mathcal{P}_k (i.e. $\mathcal{P}^i \cap \mathcal{P}_k = \emptyset$), and for all arcs originating from i , the continuous variable $x_{ij}^{\chi_k tt'}$ is not defined. We let Ξ denote the set of all such super-products and refer to the resulting master problem as the **K-EMP**.

Figures 2 and 3 illustrate an example. The original problem is depicted in Figure 2. A single customer requests one unit of p_1, p_2, p_3 and p_4 . Supplier s_1 manufactures products p_1 and p_2 , while s_2 manufactures p_2 and p_3 , and s_3 manufactures p_3 and p_4 . Given a 2-partition of the products: $\{\mathcal{P}_1 = \{p_1, p_2\}, \mathcal{P}_2 = \{p_3, p_4\}\}$, we aggregate the products into two super-products χ_1 and χ_2 . The aggregated problem is depicted in Figure 3. As c requests one unit of p_1 and one unit of p_2 in the original problem, c request two units of super-product χ_1 in the aggregated problem. As s_1 and s_2 manufacture at least one product of \mathcal{P}_1 in the original problem, they manufacture χ_1 in the aggregated problem as well. On the contrary, as s_3 does not manufacture any product of \mathcal{P}_1 in the original problem, it does not manufacture χ_1 in the aggregated problem. A similar reasoning applies to super-product χ_2 and products p_3 and p_4 .

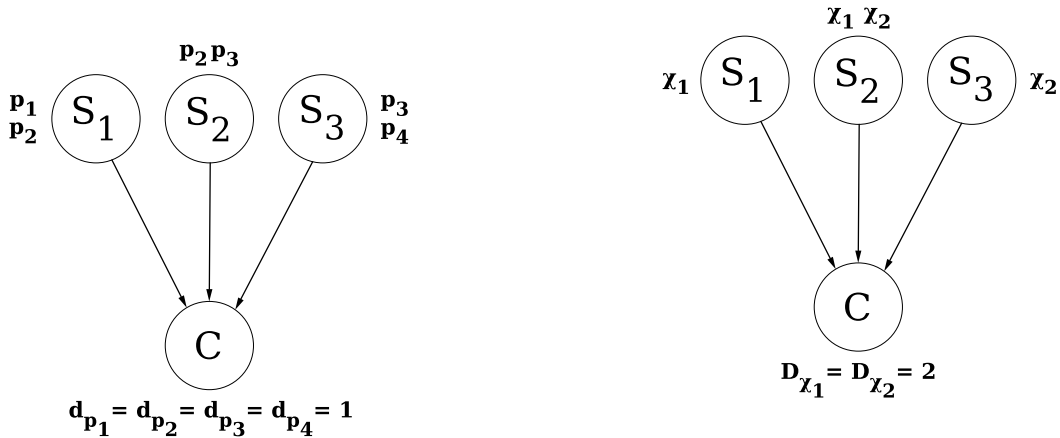


Figure 2: Before the aggregation of products

Figure 3: After aggregating p_1 with p_2 , and p_3 with p_4

Note that the aggregation of products is an approximation of the original problem. For example, a supplier may manufacture a super-product χ_k in the **K-EMP** whereas it does not manufacture all

products in \mathcal{P}_k in the original problem. In our example, s_2 in the aggregated problem manufactures all super-products and thus can satisfy all customer demands. However, in the original problem s_2 only manufacture products p_2 and p_3 , and cannot satisfy demands for products p_1 and p_4 . Thus, not every feasible solution to the **K-EMP** can be converted to a feasible solution for the LSNDP.

The **K-EMP** allocates vehicle capacity on transportation arcs in order to satisfy the routing of the K super-products. It is formulated as follows:

$$\min \sum_{((i,t),(j,t')) \in \mathcal{A}_T} f_{ij} y_{ij}^{tt'} + z \quad (16)$$

$$\sum_{((i,t),(j,t')) \in \mathcal{A}_T \cup \mathcal{H}_T} x_{ij}^{\chi_k tt'} - \sum_{((j,t),(l,t'')) \in \mathcal{A}_T \cup \mathcal{H}_T} x_{jl}^{\chi_k t' t''} = 0, \quad \forall (j, t') \in \mathcal{W}_T, \forall \chi_k \in \Xi \quad (17)$$

$$\sum_{((i,t),(j,t')) \in \mathcal{A}_T} x_{ij}^{\chi_k tt'} \geq D_{jt'}^{\chi_k}, \quad \forall (j, t') \in \mathcal{C}_T, \forall \chi_k \in \Xi \quad (18)$$

$$x_{ii}^{\chi_k tt+1} \leq \text{lim}_i^{\mathcal{W}}, \quad \forall ((i, t), (i, t+1)) \in \mathcal{H}_T \quad (19)$$

$$\sum_{\chi_k \in \Xi} x_{ij}^{\chi_k tt'} \leq \hat{u} y_{ij}^{tt'}, \quad \forall ((i, t), (j, t')) \in \mathcal{A}_T \quad (20)$$

$$z \geq \sum_{((i,t),(j,t')) \in \mathcal{A}_T} \sum_{\chi_k \in \Xi} c_{ij} x_{ij}^{\chi_k tt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_T} \sum_{\chi_k \in \Xi} c_{ii} x_{ii}^{\chi_k tt+1} \quad (21)$$

(9)-(10)

$$x_{ij}^{\chi_k tt'} \in \mathbb{R}^+, \quad \forall ((i, t), (j, t')) \in \mathcal{A}_T \cup \mathcal{H}_T, \forall \chi_k \in \Xi, \mathcal{P}^i \cap \mathcal{P}_k \neq \emptyset \quad (22)$$

$$y_{ij}^{tt'} \in \mathbb{N}^+, \quad \forall ((i, t), (j, t')) \in \mathcal{A}_T \quad (23)$$

$$z \in \mathbb{R}^+ \quad (24)$$

The objective function is the same as for the **SMP**. Constraints (17) enforce the flow conservation of each super-product at each warehouse. Constraints (18) ensure that, for each customer, each super-product demand is fulfilled. Constraint (19) limits the total amount of super-product stored by each warehouse. Constraints (20) ensure that enough vehicle capacity is allocated to support the flows of super-products. Constraint (21) bounds the flows cost approximation z . Constraints (9) and (10) are the standard Benders cuts added dynamically after solving the subproblem. Constraints (22), (23), and (24) define the decision variables and their domain.

For Benders to converge it must solve a master problem that is a relaxation of the original problem. While [Belieres et al., 2020b] prove that **K-EMP** is a relaxation when $\mathcal{K} = 1$, the following is true for general \mathcal{K} . We defer the proof to the Appendix, Section 8.

Theorem 4.1. *The \mathcal{K} -enhanced master problem, **K-EMP**, is a relaxation of the Logistics Service Network Design problem, LSNDP.*

4.3 Creating a strong super-product based master problem

For a given \mathcal{K} , there are different ways to partition the set of products into \mathcal{K} subsets of products, and thus different master problems, **K-EMP**. Different partitions of the product set may result in master problems that approximate the original problem to different degrees. We first identify a specific case where an appropriate partitioning of the product set gives a master problem that is equivalent to the original problem. Based on this observation, we define a metric that evaluates the potential benefit of aggregating a pair of products on the resulting Benders algorithm, and thus the potential benefit of including those products in the same subset of the \mathcal{K} -partition.

We first explain this specific case and then discuss the intuition behind why it yields a **K-EMP** that is equivalent to the original problem. Consider an instance of the original problem wherein each supplier that provides any product $p \in \mathcal{P}_k$ can also provide all other products included in \mathcal{P}_k . In this case, the products of \mathcal{P}_k can be aggregated without loss of information. Furthermore, given a \mathcal{K} -partition of the products $\{\mathcal{P}_1, \dots, \mathcal{P}_K\}$, if for all subsets of \mathcal{P}_k , each supplier that provides any product $p \in \mathcal{P}_k$ can also provide all other products included in \mathcal{P}_k , then all product subsets of the \mathcal{K} -partition can be aggregated without loss of information. More precisely, one can prove that in this specific case the resulting **K-EMP** is equivalent to the original problem. We next present the intuition behind a proof of this claim.

Figures 4 and 5 illustrate an example of equivalence between the **K-EMP** and the LSNDP. The original problem is depicted in Figure 4. Both products p_1 and p_2 are offered by s_1 and s_2 . Aggregating p_1 and p_2 into super-product χ_1 induces no loss of information. Indeed, if we aggregate p_1 with p_2 (see Figure 5), the demand of super-product χ_1 can be satisfied by s_1 or s_2 . As the demand of super-product χ_1 sums the demands for p_1 and p_2 , the **K-EMP** assumes that both s_1 and s_2 can satisfy the demands for p_1 and p_2 . Yet, in the original problem, s_1 and s_2 provide both p_1 and p_2 . Thus, we can aggregate p_1 and p_2 without loss of information. Similarly, the aggregation of p_3 and p_4 into super-product χ_2 induces no loss of information. Therefore, the 2-partition $\{\mathcal{P}_1 = \{p_1, p_2\}, \mathcal{P}_2 = \{p_3, p_4\}\}$ yields a **K-EMP** that is equivalent to the LSNDP.




Figure 4: Before the aggregation of products




Figure 5: After aggregating p_1 with p_2 , and p_3 with p_4

Formally, we have the following theorem, whose proof is included in the Appendix, Section 8.

Theorem 4.2. Let $\{\mathcal{P}_1, \dots, \mathcal{P}_\mathcal{K}\}$ be a \mathcal{K} -partition of the products, and $\chi_k \in \Xi, k \in \{1, \dots, \mathcal{K}\}$, the associated super-products. If for each $k \in \{1, \dots, \mathcal{K}\}$, each supplier $s \in \mathcal{S}$ that provides any product $p \in \mathcal{P}_k$ can also provide all other products included in \mathcal{P}_k , then the **K-EMP** is equivalent to the LSNDP.

Based on this theorem, we define a statistic that measures the potential benefit of including two products in the same product subset. For a product p , we denote \mathcal{S}^p as the set of suppliers that manufacture p . We denote the matching rate of a pair of products p_i and p_j as:

$$m(p_i, p_j) = \frac{|\mathcal{S}^{p_i} \cap \mathcal{S}^{p_j}|}{|\mathcal{S}^{p_i} \cup \mathcal{S}^{p_j}|}$$

This matching rate measures the fraction of suppliers of p_i or p_j that offer both products. We view this matching rate as an indicator of whether or not p_i and p_j should be aggregated. At one extreme, two products with a matching rate of one are provided by the same suppliers and thus can be aggregated without loss of information. At the other extreme, two products with a matching rate of zero have no common supplier and thus there is little to no perceived benefit on the Benders algorithm in aggregating them. We also define the matching rate of a product set \mathcal{P}_i as the average matching rate of all pairs of products in that set, ie:

$$m(\mathcal{P}_i) = \frac{2}{|\mathcal{P}_i| \times (|\mathcal{P}_i| - 1)} \sum_{i=1}^{|\mathcal{P}_i|} \sum_{j=i+1}^{|\mathcal{P}_i|} m(p_i, p_j)$$

Similarly, this value comprised between zero and one indicates whether or not products of \mathcal{P}_i should be aggregated. We note that, given $\{\mathcal{P}_1, \dots, \mathcal{P}_\mathcal{K}\}$ a \mathcal{K} -partition of the product set, the Theorem 4.2 holds as $m(\mathcal{P}_k) = 1, \forall k \in \{1, \dots, \mathcal{K}\}$. In the second phase of our Meta-PBD strategy, we propose an algorithm for determining such an "optimal" \mathcal{K} -partition of the product set.

4.4 Relationship between \mathcal{K} and the resulting K-EMP strength

The number of super-products \mathcal{K} used to formulate the **K-EMP** can vary between 0 and $|\mathcal{P}|$. When $\mathcal{K} = 0$, the **K-EMP** corresponds to the standard Benders master problem, while when $\mathcal{K} = |\mathcal{P}|$ the master is strengthened with $|\mathcal{P}|$ super-products, with a one to one correspondance between super-products and products. In that case, the **K-EMP** is clearly equivalent to the original problem. Similarly, the greater the number of super-products, the easier it is to meet the condition for applying Theorem 4.2. Thus, one could conclude that the higher the value of \mathcal{K} , the stronger the bound produced by solving the resulting super-product-based master problem.

However, this is not always true. In fact, solving a **(K+1)-EMP** may yield a weaker bound than solving a **K-EMP**. For example, consider again the LSNDP depicted in Figure 4. We observed that the 2-partition $\{\mathcal{P}_1 = \{p_1, p_2\}, \mathcal{P}_2 = \{p_3, p_4\}\}$ results in a **2-EMP** that is equivalent to the original problem. Let us consider a 3-partition $\{\mathcal{P}'_1 = \{p_1\}, \mathcal{P}'_2 = \{p_2, p_3\}, \mathcal{P}'_3 = \{p_4\}\}$ that yields 3 super-products: χ_1, χ_2 and χ_3 . As the matching rate of p_2 and p_3 is null, this 3-partition induces a loss of information. Thus, the 3-partition defines a **3-EMP** that will produce a weaker bound than the **2-EMP** defined above.

A master problem with \mathcal{K} super-products is not necessarily weaker than all master problems with $\mathcal{K}+1$ super-products. On the other hand, we have the following theorem that there always exists at least one master problem with $\mathcal{K}+1$ super-products such that the master problem based on \mathcal{K} super-products is a relaxation. While we leave the proof of Theorem 4.3 to the Appendix (Section 8), we note that the

procedure for creating such a $\mathcal{K} + 1$ master problem is straightforward. Namely, one needs to partition one of the product sets in the existing \mathcal{K} -partition.

Theorem 4.3. *Given a \mathcal{K} -enhanced master problem, $\mathcal{K} \in \{1, \dots, |\mathcal{P}| - 1\}$, there always exist a $\mathcal{K} + 1$ -enhanced master problem such that **K-EMP** is a relaxation of **K+1-EMP***

5 Meta Partial Benders Decomposition

In the previous section, we established a condition wherein a **K-EMP** is equivalent to the original problem. It is noteworthy that using such a **K-EMP** as the master problem in a Benders decomposition-based algorithm would allow the algorithm to converge in a single iteration. However, it may require using a large number of super-products to meet this condition, making the resulting master problem computationally challenging to solve. Ultimately, the number of super-products used to formulate a **K-EMP** has both positive and negative impacts on the performance of a Benders decomposition-based algorithm. On the one hand, a larger number of super-products should yield a stronger master problem and allow the algorithm to converge in fewer iterations. On the other hand, a larger number of super-products should result in a master problem that is harder to solve and therefore the algorithm will spend more time executing each iteration. However, the magnitude of either impact cannot be accurately estimated before the execution of the resulting Benders decomposition-based algorithm.

Thus, we propose an algorithm that executes a Benders decomposition-based algorithm on different master problems with different numbers of super-products. We refer to this algorithm as *Meta Partial Benders Decomposition* (Meta-PBD). Meta-PBD operates in two phases. Phase I has a time limit set to t_1^{max} units of time. It explores different areas of the original problem feasible region by dynamically changing the number of super-products used to formulate the master problem. Note that in Phase I, the number of super-products is limited by a threshold value \mathcal{K}_{max} to ensure computational tractability. Phase II has a time limit set to t_2^{max} units of time and it aims to close the optimality gap. To do so, it determines a **K-EMP** that is equivalent to the original problem, initiates it with the best bounds found in Phase I, and solves it. We first present a clustering-based strategy used to partition the product set prior to the algorithm. We then present the first phase and the second phase of Meta-PBD. The whole algorithm is described in Algorithm 1.

5.1 Partitioning the product set

In the first phase, the number of super-products used to formulate the master problem can range from 1 to a threshold \mathcal{K}_{max} . In our implementation, the partitioning of the product set is based in part on the K-medoids [Jain and Dubes, 1988, Berkhin, 2006] method, a greedy algorithm that partitions N objects into K clusters. Like K-means, K-medoids seeks to put objects into K clusters such that the sum of distances from the objects to the centers of clusters to which they are assigned is minimized. Unlike K-means, in K-medoids the center of each cluster is one of the objects in that cluster.

K-medoids requires a distance measure between each pair of objects. In the context of Meta-PBD, and due to Theorem 4.2, we seek to partition product set into product subsets with high matching rates. As the K-medoids algorithm seeks to minimize total distance, we set a proximity measure between each pair of products $(p, p') \in \mathcal{P}^2$ as $d(p, p') = 1 - m(p, p')$.

For the sake of simplicity, we compute the partitions of the product set that may be used in Phase 1 in a preprocessing fashion. We start with $\mathcal{K} = 1$ and we use an incremental approach to determine the

Algorithm 1: Meta Partial Benders Decomposition

Data: Maximum number of super-products, \mathcal{K}_{max} , Time limit on bounds improvement, t^{bounds} , Minimal improvement required for the bounds, $impr^{bounds}\%$, Threshold on the number of master solutions explored, $msols_{max}$, Time limit for phase I, t_1^{max} , Time limit for phase II, t_2^{max}

1-partition \leftarrow whole product set \mathcal{P}

for \mathcal{K} ranging from 1 to $\mathcal{K}_{max} - 1$ **do**

 Identify largest product subset from the \mathcal{K} -partition

 Partition it by applying the 2-medoids algorithm, using distance $1 - m(p, p')$ for each $(p, p') \in \mathcal{P}$

$\mathcal{K} + 1$ -partition \leftarrow combine the obtained subsets with all other product subsets from the \mathcal{K} -partition unchanged

Launch phase I using the obtained partitions

Launch phase II using the best bounds found in Phase I

Result: Final solution, UB

different partitions of the product set. At each iteration, we build a $\mathcal{K} + 1$ -partition of the product set by partitioning the largest product subset of the current \mathcal{K} -partition into two product subsets, leaving the other product subsets (and resulting super-products) unchanged. The largest product subset is partitioned by applying the 2-medoids algorithm. Given Theorem 4.3, this strategy ensures that, for $(k_1, k_2) \in \{1, \dots, \mathcal{K}\}^2$, if $k_2 > k_1$ the master problem K_2 -EMP yields bounds at least as strong as those obtained solving the master problem K_1 -EMP. This preprocessing phase terminates as a \mathcal{K} -partition of the product set for each value of \mathcal{K} ranging from 1 to \mathcal{K}_{max} .

5.2 Phase I

The first phase iteratively explores different areas of the original problem feasible region and it aims to quickly determine high-quality solutions. At each iteration, it changes the number of super-products and it solves the resulting master problem. It uses solution progress criteria and an assessment of the master computational tractability to determine whether the master should be changed as well as whether the number of super-products should decrease or increase. Based on these decisions, and if the number of super-products can decrease/increase, it applies an integral bisection search to identify a more promising number of super-products. At each iteration, the master problem is solved using the Benders decomposition-based solution algorithm proposed by Belieres et al. [Belieres et al., 2020b]. Although that approach is presented in the context of solving a master problem based on a single super-product, it can be easily extended to a master problem based on multiple super-products.

The first phase is initiated with $\mathcal{K} = 1$, and it starts by solving the **1-EMP**. The reason behind this choice is that the **1-EMP** is the most tractable master problem, and thus a good candidate to quickly improve the initial bounds. At each iteration, the algorithm must (i) detect when the progress of the Benders decomposition-based algorithm for the current **K-EMP** has slowed and identify the reason for this limited progress and (ii) suggest a more adequate number of super-products. We next describe how the algorithm performs each step.

The algorithm tracks whether the Benders decomposition-based algorithm applied with the current **K-EMP** was able to improve either the primal or dual bound on the optimal objective function value

of the original problem. If t^{bounds} units of time have passed in which neither the primal bound nor the dual bound have improved by at least $impr^{bounds}$, the algorithm counts the number of integer master solutions that have been produced in the last t^{bounds} units of time. If this number is lower than $msols_{max}$, the algorithm judges that the slow progress is due to a lack of computational tractability and it indicates that the number of super-products should decrease. On the other hand, if the number of integer master solutions produced is equal to or greater than $msols_{max}$, the algorithm judges that the current master problem should be strengthened and thus that the number of super-products should increase. As the number of super-products for the next iteration is obtained via an integral bisection search, before stopping the current solution the algorithm first determines if the interval to be bisected is not empty. In that case only, the solution is stopped and the next number of super-products is computed. Figures 6 and 7 illustrate two examples of bisection search with $\mathcal{K}_{max} = 5$.




Figure 6: First example of bisection search




Figure 7: Second example of bisection search

To determine if the number of super-products can decrease or increase, let K^- be the largest number of super-products used previously that is smaller than \mathcal{K} and let K^+ be the smallest number of super-products used previously that is larger than \mathcal{K} . Then, \mathcal{K} cannot decrease if $\mathcal{K} = 1$ or $K^- = \mathcal{K} - 1$. Similarly, \mathcal{K} cannot increase if $\mathcal{K} = \mathcal{K}_{max}$ or $K^+ = \mathcal{K} + 1$. Outside of these cases, the number of super-products to be considered at the next iteration is obtained by an integral bisection search. If the algorithm aims to increase (decrease) the number of super-products, we compute the midpoint between \mathcal{K} and K^+ (K^-) and we round it up if it is not integer. The obtained value defines the number of super-products to be considered at the next iteration. Note that it may happen that the algorithm seeks to increase \mathcal{K} while $\mathcal{K} < \mathcal{K}_{max}$ and no number of super-products larger than \mathcal{K} was studied previously. In that case, the number of super-products to be considered at the next iteration is set to $\lceil \frac{\mathcal{K}_{max} - \mathcal{K}}{2} \rceil$.

We note that at each iteration we initiate the current Benders decomposition-based algorithm with the best upper bound and the best lower bound found in previous iterations. In addition, we strengthen the **K-EMP** at a given iteration with all the Benders cuts generated in previous iterations. As the Benders cuts (9) and (10) only involve y and z variables, they can be re-used from one iteration to the next. Figure 8 illustrates the steps taken at each iteration of Meta-PBD. Detailed pseudo-code for the first phase is presented in Algorithm 2, which is in the Appendix (Section 8).




Figure 8: Iteration of Phase I

5.3 Phase II

In the second phase, we formulate a master problem that is equivalent to the original problem. To determine a partition of the product set that satisfies Theorem 4.2, we start with $\mathcal{K} = 1$ and we use an incremental approach. At each iteration, we identify a product subset of the current \mathcal{K} -partition with a matching rate lower than 1. We then build a $\mathcal{K} + 1$ -partition of the product set by partitioning the identified product subset into two product subsets, which is done by applying the 2-medoids algorithm. If no product subset is identified, the current \mathcal{K} -partition satisfies Theorem 4.2 and we can formulate the master problem.

The resulting master problem is solved using a generic branch-and-cut implementation rather than the Benders decomposition-based solution algorithm proposed by Belieres et al. [Belieres et al., 2020b]. Indeed, since the master problem is equivalent to the original problem, all integer master solutions found in the branch-and-bound tree are feasible, such that using the valid inequalities and the heuristic solutions proposed by Belieres et al. [Belieres et al., 2020b] do not bring added value. To reduce the branch-and-bound search tree and accelerate convergence, we initiate the branch-and-cut with the best upper bound and the best lower bound found in the first phase. Detailed pseudo-code for the second phase is presented in Algorithm 3, which is in the Appendix (Section 8).

6 Computational study

The goal of this computational study is to study the effectiveness of Meta-PBD at solving instances of the LSNDP. We first describe the instances used in this computational study. Specifically, we describe a set of random instances used to analyze the performance of Meta-BPD according to certain parameters, and a set of industrial instances used to assess the performance of our method on realistic cases. Second, we describe how this computational study was conducted. Next, we analyze how the number of super-products impacts the bounds produced by the master problem as well as the time required to solve the master problem. Finally, we study the performance of Meta-PBD on both the random instances and the industrial instances.

6.1 Instances

We first describe the set of random instances generated for this computational study. Then, we briefly present the set of industrial instances taken from [Belieres et al., 2020a].

6.1.1 Random instances

These instances are inspired by the operations of our industrial partner and produced by a random generator similar to the one described in [Belieres et al., 2020b]. This random generator produces a static graph \mathcal{G} and a time dimension based on the following parameters: the size of the node-set \mathcal{N} , the connectivity radius α , the number of days in the planning horizon D , and the number of time-periods per day Δ . The connectivity radius is a threshold on how close two nodes must be to each other for transportation to be an option. See [Belieres et al., 2020b] for details regarding how the network \mathcal{G} and the time dimension are constructed. Likewise, customer requests, vehicle capacities, and warehouse capacities are defined in a similar way to what is done in [Belieres et al., 2020b].

Our instances differ from those presented in [Belieres et al., 2020b] as we generate products and assign products to suppliers in a different fashion. Namely, [Belieres et al., 2020b] did not explicitly recognize the concept of product families. Thus, the generator used in this study accepts two more parameters: (i) the cardinality of the product families set \mathcal{F} , and, (ii) a supply probability ϕ . The generator randomly assigns each product to one product family and one, two, or three product families to each supplier. Then the generator randomly distributes products to the suppliers. Specifically, each supplier that manufactures a given product family has a probability ϕ to offer each product in that product family.

The random instances are based on combinations of the following parameter values: $|\mathcal{N}| = \{50\}$, $\alpha = \{10, 30\}$, $D = 30$ days, $\Delta = \{2, 3, 4\}$, $\mathcal{F} = \{7\}$, $|\mathcal{P}| = \{100\}$, and $\phi = \{25\%, 50\%, 75\%\}$. In addition, due to the randomness of the instance generator, we generated 5 instances for each of the 18 possible combinations of parameter values, yielding 90 instances in total.

Recall that for a pair of products (p_i, p_j) , the matching rate $m(p_i, p_j)$ measures the fraction of suppliers that manufacture both products. Thus, matching rates are directly impacted by the value of the parameter ϕ . As a result, for each generated instance we compute the matching rate of each product family, which indicates whether or not products of that family should be aggregated. We report in Table 1 the average product family matching rates, averaged over instances generated with the same value of ϕ . Not surprisingly, the larger the value of ϕ the larger the resulting matching rate.

ϕ	25%	50%	75%
Matching rate	46.51%	57.42%	76.77%

Table 1: Average matching rate of the product families

6.1.2 Industrial instances

These instances are taken from the article [Belieres et al., 2020a] and they reflect the operations performed by a 3PL in the supply chain management of a French restaurant chain. These instances are constructed in part from real data provided by our industrial partner. These data include: geographic positions of the stakeholders, travel times between each pair of stakeholders, transportation costs, storage

costs, product families manufactured by each supplier, customers demand schedules and product demand seasonalities. We refer the interested reader to [Belieres et al., 2020a] for more details on the real data and/or the generation of the industrial instances.

In this study, we focus on the second set of industrial instances described in [Belieres et al., 2020a]. These instances are referred to as *difficult* instances as they involve time-expanded networks of significantly larger scale than those involved in the first set of industrial instances. These instances are based on combinations of the following parameter values: $|\mathcal{N}| = \{67\}$, $\alpha = \{0, 20, 40, 60\}$, $D = 14$ days, $\Delta = \{2, 4, 6\}$, $\mathcal{F} = \{3\}$, $|\mathcal{P}| = \{131\}$, and 2 demand seasonalities (summer and winter).

6.2 Setup of study

To assess the effectiveness of **Meta-PBD**, we compare its performance against multiple benchmarks. For the first, **CPLEX**, the LSNDP is solved with CPLEX branch-and-cut solver with all parameter values left at their defaults. The next three benchmarks are direct applications of the Benders decomposition-based algorithm proposed in [Belieres et al., 2020b] on different partial decompositions of the LSNDP. These approaches are referred to as *static* PBD-approaches since the considered master problem solved remains unchanged through the optimization process. In the method **Single**, the master problem is formulated using a single super-product that aggregates all the products in the instance. On the other hand, the next two methods involve master problems based on multiple super-products. The method **Families** partitions the set of products based on the product families to which they belong. Thus, the resulting master problem is formulated with one super-product for each product family. The method **Random** randomly partitions the set of products into $|\mathcal{F}|$ subsets, and thus randomly determines how super-products and the resulting master problem are formulated.

Each method is seeded with the same heuristic solution (x_h, y_h) , which is obtained by solving the linear relaxation of the LSNDP and then rounding up the value of each vehicle variable $y_{ij}^{tt'}$ in its optimal solution. All algorithms are coded in C++ and executed on an Intel Xeon E5-2695 processor with 16 GB of memory under Linux 16.04. Linear and integer programs were solved using Cplex 12.7. All algorithms are executed with a stopping criteria of a proven optimality gap of 1%. Regarding parameters for Meta-PBD, we set \mathcal{K}_{max} to 10, t^{bounds} to 1/18 of the global time limit, $impr^{bounds\%}$ to 1%, $msols_{max}$ to 1, t_1^{max} to 1/3 of the global time limit, and t_2^{max} to 2/3 of the global time limit. Note that these values were determined through a set of tuning experiments.

6.3 Impact of \mathcal{K} on K-EMP

We first study the impact that the number of super-products has on the resolution of the resulting master problem, both in terms of the bound produced and of the computational time needed. Note that we use the random instances, for which 7 product families are involved.

For each instance and all values of \mathcal{K} from 1 to 7, we apply the K-medoids algorithm to partition the set of products \mathcal{P} . Based on this partition, we formulate the resulting **K-EMP** and we solve its linear programming relaxation. In addition, for each instance we solve the linear programming relaxation of the LSNDP. For each linear programming relaxation (LPR) we consider two statistics: (i) the objective function value of the optimal solution, r_{LPR}^{sol} , and (ii) the computational time required for its solution, r_{LPR}^{time} .

To measure the impact that \mathcal{K} has on the quality of the bounds produced by the master problem as well as on the time required for solving the master problem, for each value of \mathcal{K} and for each instance

we compute the following indicators:

$$\text{lb-root-gap} = \frac{r_{LSNDP}^{sol} - r_{K-EMP}^{sol}}{r_{LSNDP}^{sol}} \quad \text{root-time-ratio} = \frac{r_{K-EMP}^{time}}{r_{LSNDP}^{time}}$$

lb-root-gap indicates the gap between the optimal solution of the **K-EMP** linear programming relaxation and the optimal solution of the LSNDP linear programming relaxation. root-time-ratio indicates the ratio between the computational time required for solving the **K-EMP** linear programming relaxation and the computational time required for solving the LSNDP linear programming relaxation. In Figure 9, we display for different values of \mathcal{K} the average lb-root-gap and root-time-ratio, averaged over all instances.




Figure 9: Root-gaps and computational time ratios for the **K-EMP**

We see that the **K-EMP** based on a single super-product gives a root-gap of 35.7%. However the average *root-gap* decreases as the number of super-products used to formulate the master increases. This confirms that considering more super-products in the master leads to better approximate the original problem. This improvement is significant, as the average *root-gap* decreases by more than 20% when 7 super-products are used to formulate **K-EMP**. As expected, the computational time for solving the master problem linear programming relaxation increases with the number of super-products. However, in the worst case, the linear programming relaxation of **K-EMP** is solved in less than 0.5% of the time needed to solve the LSNDP root relaxation.

Belieres et al. [Belieres et al., 2020b] propose three classes of valid inequalities for strengthening a **K-EMP** based on a single super-product. Since the same valid inequalities are used in this study, we ran a similar experiment to the one just described, albeit with these constraints added to the formulation of each **K-EMP**. We then recalculated the *root-gap* and *root-time* statistics in the same manner. We observed the same trends in those statistics as in Figure 9. Namely, the larger the value of \mathcal{K} , the smaller the value of *root-gap* and the larger the value of *root-time*. However, the magnitudes were much different, as the *root-gap* equals 0.63% for $K = 1$ and 0.38% for $K = 7$. We thus observe that for $K = 1$, the optimal solution of the **K-EMP** linear relaxation already provides a very tight bound on the LSNDP linear relaxation as the valid inequalities are added.

6.4 Results obtained on the random instances

We analyze the results obtained on the random instances with a time limit of 3 hours. We first compare our approach to various benchmark methods and we computationally demonstrate that it

strictly outperforms the algorithm proposed by Belieres et al. [Belieres et al., 2020b], i.e. **Single**. We then analyze the performance of our approach in more detail.

6.4.1 Benchmarking performance of Meta-PBD

We compare the performance of **Meta-PBD** variants with that of the following benchmark methods: **CPLEX**, **CPLEX-Benders**, **Single**, **Families** and **Random**. We do so based in part on the objective function value, UB_x , of the best primal solution produced by method x , as well as the strongest dual bound, LB_x , it produced. Let UB_{Best} denote the best objective function value over all methods. Similarly, we let LB_{Best} denote the best dual bound produced by all methods. For each instance and each method x , we compute two performance indicators:

$$gap_{UB} = \frac{UB_x - UB_{Best}}{UB_x} \quad gap_{LB} = \frac{LB_{Best} - LB_x}{LB_{Best}}$$

Both indicators are non-negative and equal 0 if method x produced the primal/dual solution with best objective function value over all methods. We also compute the average optimality gap at termination for method x over all instances as $(UB_x - LB_x)/UB_x$.

In Table 2 we present averages of the optimality gap and each performance indicator for each method. We also report the number of instances solved as well as $nb.LB_{Best}$ and $nb.UB_{Best}$, i.e. the number of instances for which method x found the best lower/upper bound over all methods. Note that for one instance, the best upper bound was found by both **Meta-PBD** and **Single**. Best values are noted in bold.

Method	Optimality Gap	Solved	gap_{LB}	$nb.LB_{Best}$	gap_{UB}	$nb.UB_{Best}$
CPLEX	10.47%	2	0.06%	1	7.17	2
Families	5.33%	0	0.50%	0	2.09%	16
Meta-PBD	3.06%	11	0.00%	88	0.23%	65
Single	5.02%	0	0.66%	1	1.60%	7
Random	8.45%	0	0.64%	0	5.20%	1

Table 2: Comparing **META-PBD** to the benchmarks

Meta-PBD significantly outperforms all the benchmarks regardless of the indicator considered. In particular, we observe that **Meta-PBD** provides primal solutions much better than those produced by the other benchmarks. It is also noteworthy that **Meta-PBD** provides strict improvements over the algorithm proposed by Belieres et al. [Belieres et al., 2020b], i.e. **Single**. Overall, our approach provides a gap at termination that is more than 2 percent better than that of **Single**. Similarly, we observe that **Meta-PBD** solves 11 instances while **Single** never closes the optimality gap within 1%. The only benchmark that solves some instances to optimality is **CPLEX**. Nevertheless, the performance of **CPLEX** is quite poor, as it yields the largest optimality gap on average. Although **CPLEX** produces strong dual bounds, it encounters difficulties on the primal side. Ultimately, all the Partial Benders Decomposition-based strategies provide better primal solutions than **CPLEX** overall.

We next turn our attention to the Partial Benders Decomposition-based strategies, i.e. **Meta-PBD** and the static PBD-approaches **Single**, **Families**, and **Random**. We first observe that all strategies produce strong dual bounds, and that the quality of these bounds does not vary significantly from one

strategy to another. This result is highly anticipated since we observed in subsection 6.3 that, when valid inequalities are added, the linear relaxation of the master problem with a single super-product already provides a very tight bound on the LSNDP linear relaxation. As a result, the improvement of the bounds produced by the master problem is of small magnitude as the number of super-products increases. On the other hand, we observe greater differences in the quality of the primal solutions produced.

As **Random** and **Families** both involve 7 super-products, comparing these approaches in terms of primal solutions provides evidence for our hypothesis that how the product set is partitioned is critical for the resulting master problem to enable a Benders decomposition-based algorithm to perform well. The performance of **Single** is very close to that of **Families**. Overall, **Single** yields a better optimality gap and produces better primal solutions on average. Nevertheless, **Families** finds the best primal solution over all methods much more often than **Single**. This result suggests that the number of super-product to consider in a static PBD-approach depends on the nature of the instance solved. To confirm this hypothesis, we compare the performance **Single** and **Families** for different values of ϕ .

As noted, we generated instances for three different values of ϕ , the probability a supplier produces a product in a family it supplies. In addition, the larger the value of ϕ , the more likely the condition to which Theorem 4.2 holds, in which case the master problem based on product families is equivalent to the original problem. For each value of ϕ , we report in Table 3 the number of instances for which **Single** provides a better primal/dual solution than **Families**, and vice versa. Best values are noted in bold.

ϕ	Best primal solution		Best dual solution	
	Families	Single	Families	Single
25%	18	12	13	17
50%	26	4	16	14
75%	30	0	20	10

Table 3: Comparing the primal/dual solutions produced by **Families** and **Single**

As anticipated, **Families** tends to perform better than **Single** as the value of ϕ increases. This confirms that, in the case of a static PBD-approach, the number of super-products to consider depends greatly on the nature of the instance. In the context of multi-product supply chains, products are often already classified into different families, and such classifications should be used to determine the partition of the product set. Nevertheless, for a given supply chain there exist in practice multiple classifications of the products with different levels of granularity. For example, the family of fresh products could be decomposed into fruit products, vegetable products, and meat products. As a result, it is assumed that using different partitions of the product set with different levels of granularities can be beneficial, which motivates our Meta-PBD framework.

6.4.2 Analyzing the performance of Meta-PBD

As observed in Table 2, **Meta-PBD** produces primal solutions significantly stronger than those computed by the static PBD-approaches. Thus, we investigate to understand why **Meta-PBD** outperforms the benchmarks.

We assess the interest of changing dynamically the number of super-products in the first phase of **Meta-PBD**. To do so, we study the performance study of each iteration of **Meta-PBD**'s first phase. We report in Table 4 the distribution of the number of iterations performed during the first

phase of **Meta-PBD**. For example, column "3 iter." indicates the percentage of instances for which the Benders decomposition-based algorithm was executed three times during the first phase of **Meta-PBD**. We report in Table 5 the improvement obtained in the primal bound and the dual bound during each iteration of **Meta-PBD**'s first phase. We note that since **Meta-PBD**'s first phase did not execute 5 iterations for all instances, for a given number of iterations, we average over instances wherein **Meta-PBD**'s first phase executed as many iterations.

1 iter.	2 iter.	3 iter.	4 iter.	5 iter.
32.22%	34.44%	21.11%	11.11%	1.11%

Table 4: Distribution of instances by # of iterations performed during the first phase of **Meta-PBD**

Iter.	LB Impr.	UB Impr.
1	0.44%	8.54%
2	0.15%	3.52%
3	0.05%	1.30%
4	0.03%	0.06%
5	0.00%	0.00%

Table 5: Performance per iteration of the first phase of **Meta-PBD**

Most of the improvement on the primal and the dual bounds is performed during the first iteration. Nevertheless, while the improvement decrease with the number of iterations, we observe that the primal bound is significantly improved in both the second and the third iteration. For 80% of the instances, the best primal solution obtained at the end of **Meta-PBD**'s first phase was computed during the last iteration.

We now compare the primal solutions obtained at the end of the first phase, i.e. **Meta-PBD-1h**, to those computed with the static PBD-approaches in the same amount of time, i.e. **Single-1h** and **Families-1h**. For each instance, we measure a primal gap between the solution computed by **Meta-PBD-1h** and those computed by **Single-1h** and **Families-1h**. A positive gap indicates that **Meta-PBD-1h** found a better primal solution than the considered benchmark. Overall, we observe that **Meta-PBD-1h** computes primal solutions 4.22% better than **Families-1h** and 2.87% better than **Single-1h**. This clearly demonstrates the added value of changing dynamically the number of super-products used to formulate the master problem.

A critical factor in the performance of a Benders decomposition-based algorithm is its ability to quickly solve master problems, which in turn yield subproblems that may be solved as part of a process for creating feasible solutions to the original problem. Thus, we report in Table 6 the average number of subproblems generated by **Meta-PBD-1h**, **Single-1h** and **Families-1h**, as well as the average percentage of those subproblems that are feasible.

This table demonstrates how the number of super-products used to construct the **K-EMP** impacts the number and quality of master problem solutions generated. As **Single-1h** solves the most computationally tractable master problem, it yields the greatest number of subproblems. However, as that master problem has the poorest approximation of the original problem, those subproblems are the least likely to

	Families-1h	Meta-PBD-1h	Single-1h
Average number of subproblems generated	3.72	5.16	5.76
% subproblems that are feasible	47.76%	46.55%	33.01%
% subproblems that are infeasible	52.24%	53.45%	66.99%

Table 6: Subproblem generation and feasibility after 1 hour of computation

be feasible amongst the three methods. Conversely, with a master problem that is stronger, but harder to solve, **Families-1h** generates the fewest number of subproblems. However, those subproblems are the most likely to be feasible amongst the three methods. **Meta-PBD-1h** achieves a balance between the other two methods. It generates nearly as many subproblems as **Single-1h**, but the percentage of subproblems that are feasible is nearly to that of **Families-1h**.

We now illustrate in Figure 10 the distribution of occurrences of primal solution improvement over the first hour of computation. The computational time of 3,600 seconds is divided into 10 intervals of 360 seconds. Thus, a value of 10 for the first interval indicates that the incumbent primal solution was improved 10 times within the 360 first seconds of computation.




Figure 10: Distribution of the occurrences of primal solution improvement after 1 hour of computation

In Figure 10, we see that **Meta-PBD-1h** improves the incumbent throughout the optimization process. This is in contrast to **Single-1h** and **Families-1h**, both of which find the majority of the improved primal solutions early in their execution. Overall, **Meta-PBD-1h** improves the incumbent more often than **Single-1h** and **Families-1h** combined. Thus, we conclude that dynamically varying the subproblem information used to formulate the master problem enables to change the portion of the feasible region that is explored and allows to discover more promising solutions.

We now demonstrate the interest of **Meta-PBD**'s second phase, which formulates and solves a master problem equivalent to the original problem. To do so, we compare the results obtained with **Meta-PBD** to those obtained when performing the first phase alone with the same time limit, i.e. **Phase-I-3h**, or the seconde phase alone with the same time limit, i.e. **Phase-II-3h**. In Table 7, we indicate values for the same performance indicators as those reported in Table 2. Note that here, we use the best primal/dual bound found over **Meta-PBD**, **Phase-I-3h**, and **Phase-II-3h** to compute values for the performance indicators.

Phase-I-3h provides a significantly better average optimality gap at termination than **Phase-II-3h**,

Method	Optimality Gap	Solved	gap _{LB}	nb.LB _{Best}	gap _{UB}	nb.UB _{Best}
Meta-PBD	3.06%	11	0.00%	62	0.40%	53
Phase-I-3h	3.71%	0	0.50%	1	0.58%	43
Phase-II-3h	8.72%	8	0.01%	27	6.26%	6

Table 7: Comparing **META-PBD** to **Phase-I-3h** and **Phase-II-3h**

which can be attributed to the difference in the primal solutions computed. This result indicates that directly solving a master problem strengthened with a large amount of subproblem information is not an efficient strategy, as it yields a branch-and-bound search tree that is prohibitively large. On the other hand, **Phase-II-3h** manages to solve 8 instances to optimality, against 0 instances for **Phase-I-3h**. This demonstrates that, while strengthening the master problem with aggregated subproblem information allows to quickly compute high-quality primal solutions, having a loss of information compared to the original problem prevents to compute the optimal master solution. Our approach strikes a compromise as it significantly outperforms both benchmarks. Ultimately, **Meta-PBD** solves more instances than **Phase-II-3h** and it provides better primal solutions than **Phase-I-3h**. When executed after the first phase, the second phase improves the best primal bound by 1.06% on average. Similarly, it improves the best dual bound by 0.59% on average.

6.5 Results obtained on the industrial instances

In this section, we assess how **Meta-PBD** compares with the best performing benchmark, **Single**, on the set of industrial instances with a maximum run-time of 3 hours. The instances, which are described by the name “LSNDP _{α} Δ S,” vary according to three parameters: the connectivity radius (α), the number of time-periods per day (Δ), and the season considered (S). We set a maximum run-time of 5 hours for each run. In Table 8, we report the optimality gap, the dual solution, and the primal solution computed by **Meta-PBD** and **Single**. The best values are noted in bold. Note that none of the instances could be solved by either method within the time limit.

Instance	Meta-PBD			Single		
	Gap	LB	UB	Gap	LB	UB
LSNDP_0.2_S	2.20	112,069	114,585	2.00	112,212	114,506
LSNDP_0.2_W	2.29	94,867	97,090	2.45	94,902	97,284
LSNDP_0.4_S	7.74	107,709	116,740	5.50	107,857	114,130
LSNDP_0.4_W	4.04	91,320	95,163	3.77	91,342	94,925
LSNDP_0.6_S	8.27	106,210	115,780	9.02	107,245	117,884
LSNDP_0.6_W	10.85	89,941	100,887	9.86	90,913	100,855
LSNDP_20.2_S	4.56	104,291	109,270	18.34	97,984	119,985
LSNDP_20.2_W	4.11	88,562	92,354	20.60	81,652	102,832
LSNDP_20.4_S	7.53	101,169	109,413	35.18	92,411	142,560
LSNDP_20.4_W	6.40	85,464	91,310	23.75	76,590	100,441
LSNDP_20.6_S	27.17	100,771	138,356	32.58	87,284	129,466
LSNDP_20.6_W	7.79	85,433	92,655	39.10	74,748	122,730

Instance	Meta-PBD			Single		
	Gap	LB	UB	Gap	LB	UB
LSNDP_40.2_S	4.73	101,917	106,976	52.52	83,237	175,318
LSNDP_40.2_W	4.28	86,585	90,460	30.94	71,663	103,774
LSNDP_40.4_S	8.36	98,095	107,046	67.49	81,312	250,086
LSNDP_40.4_W	8.46	83,392	91,095	58.37	65,025	156,210
LSNDP_40.6_S	10.20	97,708	108,812	56.62	72,382	166,844
LSNDP_40.6_W	10.43	82,926	92,579	59.47	62,174	153,406
LSNDP_60.2_S	4.59	101,338	106,209	36.04	81,467	127,375
LSNDP_60.2_W	4.34	86,032	89,932	33.03	70,320	105,004
LSNDP_60.4_S	5.68	98,128	104,034	41.87	77,019	132,501
LSNDP_60.4_W	7.44	83,083	89,761	50.64	61,115	123,811
LSNDP_60.6_S	25.31	97,717	130,822	76.97	71,333	309,785
LSNDP_60.6_W	8.17	82,982	90,362	54.71	60,461	133,509

Table 8: Comparing **Meta-PBD** to **Single** on the industrial instances

As for the results provided in Table 2, one first observes that **Meta-PBD** strictly outperforms **Single**, especially as the size of the instance increases. Overall, **Meta-PBD** produces a better primal solution than **Single** for 19 out of the 24 instances, and a better dual solution than **Single** for 18 out of the 24 instances. The average optimality gap at termination is of 8.12% for **Meta-PBD**, against 34.20%

for **Single**. This large difference is essentially due to the quality of the primal solutions computed, as overall, **Meta-PBD** determines transportation plans 20% more cost-effective than **Single**.

7 Conclusions and future work

In this paper, we studied the Logistics Service Network Design Problem (LSNDP), a transportation planning problem that arises in the management of supply chains. We proposed a solution algorithm based on the recently-proposed *Partial Benders Decomposition* (PBD) technique, wherein information derived from subproblem(s) is used to reinforce the master problem solved in the context of a Benders decomposition-based algorithm. In the proposed approach, the master problem is strengthened with variables and constraints that model the need to route one or more “super-products” that are aggregations of a subset of products.

Existing implementations of PBD formulate and use a single master problem to solve the original problem. We proposed a solution framework called *Meta Partial Benders Decomposition* (Meta-PBD) that changes the master problem considered throughout its execution based on information regarding the progress of the optimization process. Through an extensive computational study, we demonstrated that Meta-PBD strictly outperforms partial Benders decomposition-based algorithms based on a single master problem. More specifically, by changing the master problem used during its execution, Meta-PBD diversifies the search, allowing it to find primal solutions significantly better than those computed by the benchmarks and to solve more instances to optimality.

There are several avenues for enhancing Meta-PBD, both in the context of solving the LSNDP and other network design problems. For example, while Meta-PBD currently starts with a master problem based on a single super-product, we are considering enhancing it to potentially start with a master problem based on multiple super-products. Also, the concept of aggregating products/commodities at different levels of granularity can be applied to all multi-product/commodity network design problems. Therefore, another avenue of future work is to develop a Meta-PBD-type scheme for solving general multi-product/commodity network design problems.

Acknowledgements

This work has been partially supported by the French National Research Agency through the Pi-Comodalité project under the grant ANR-15-CE22-0012. This support is gratefully acknowledged.

References

- [Beamon, 1998] Beamon, B. M. (1998). Supply chain design and analysis:: Models and methods. *International journal of production economics*, 55(3):281–294.
- [Belieres et al., 2020a] Belieres, S., Hewitt, M., Jozefowicz, N., and Semet, F. (2020a). A time-expanded network reduction matheuristic for the logistics service network design problem. *Transportation Research Part E: Logistics and Transportation Review (Major revision)*.

- [Belieres et al., 2020b] Belieres, S., Hewitt, M., Jozefowicz, N., Semet, F., and van Woensel, T. (2020b). A benders decomposition-based approach for logistics service network design. *European Journal of Operational Research*, 286:523–537.
- [Benders, 1962] Benders, J. F. (1962). Partitioning procedures for solving mixed variables programming problems. *Numerische Mathematik*, 4:238–252.
- [Benoist et al., 2002] Benoist, T., Gaudin, E., and Rottembourg, B. (2002). Constraint programming contribution to benders decomposition: A case study. In *International Conference on Principles and Practice of Constraint Programming*, pages 603–617. Springer.
- [Berkhin, 2006] Berkhin, P. (2006). A survey of clustering data mining techniques. In *Grouping multi-dimensional data*, pages 25–71. Springer.
- [Cheong et al., 2007] Cheong, M. L., Bhatnagar, R., and Graves, S. C. (2007). Logistics network design with supplier consolidation hubs and multiple shipment options. *Journal of Industrial and Management Optimization*, 3(1):51.
- [Cordeau et al., 2006] Cordeau, J.-F., Pasin, F., and Solomon, M. M. (2006). An integrated model for logistics network design. *Annals of operations research*, 144(1):59–82.
- [Cordeau et al., 2001] Cordeau, J.-F., Stojković, G., Soumis, F., and Desrosiers, J. (2001). Benders decomposition for simultaneous aircraft routing and crew scheduling. *Transportation science*, 35(4):375–388.
- [Costa, 2005] Costa, A. M. (2005). A survey on benders decomposition applied to fixed-charge network design problems. *Computers & operations research*, 32:1429–1450.
- [Costa et al., 2012] Costa, A. M., Cordeau, J. F., Gendron, B., and Laporte, G. (2012). Accelerating benders decomposition with heuristic master problem solutions. *Pesquisa Operacional*, 32:03–20.
- [Crainic, 2000] Crainic, T. G. (2000). Service network design in freight transportation. *European Journal of Operational Research*, 122:272–288.
- [Crainic et al., 2014] Crainic, T. G., Hewitt, M., and Rei, W. (2014). Partial decomposition strategies for two-stage stochastic integer programs. *CIRRELT, Centre interuniversitaire de recherche sur les réseaux d'entreprise*.
- [Crainic et al., 2016] Crainic, T. G., Rei, W., Hewitt, M., and Maffioni, F. (2016). Partial benders decomposition strategies for two-stage stochastic integer programs. *CIRRELT, Centre interuniversitaire de recherche sur les réseaux d'entreprise*
- [Dufour et al., 2018] Dufour, É., Laporte, G., Paquette, J., and Rancourt, M.-È. (2018). Logistics service network design for humanitarian response in east africa. *Omega*, 74:1–14.
- [Fontaine et al., 2017] Fontaine, P., Crainic, T., Rei, W., and Jabali, O. (2017). *Multi-modal scheduled service network design with resource management for two-tier city logistics*. CIRRELT, Centre interuniversitaire de recherche sur les réseaux d'entreprise.

- [Gendron et al., 2016] Gendron, B., Scutellà, M. G., Garroppo, R. G., Nencioni, G., and Tavanti, L. (2016). A branch-and-benders-cut method for nonlinear power design in green wireless local area networks. *European Journal of Operational Research*, 255(1):151–162.
- [Jain and Dubes, 1988] Jain, A. K. and Dubes, R. C. (1988). Algorithms for clustering data. *Englewood Cliffs: Prentice Hall, 1988*.
- [Magnanti and Wong, 1981] Magnanti, T. L. and Wong, R. T. (1981). Accelerating benders decomposition: Algorithmic enhancement and model selection criteria. *Operations Research*, 29:464–484.
- [Marasco, 2008] Marasco, A. (2008). Third-party logistics: A literature review. *International Journal of production economics*, 113(1):127–147.
- [Melo et al., 2009] Melo, M. T., Nickel, S., and Saldanha-Da-Gama, F. (2009). Facility location and supply chain management—a review. *European journal of operational research*, 196(2):401–412.
- [Rahmaniani et al., 2017] Rahmaniani, R., Crainic, T. G., Gendreau, M., and Rei, W. (2017). The benders decomposition algorithm: A literature review. *European Journal of Operational Research*, 259:801–817.
- [Rei et al., 2009] Rei, W., Cordeau, J.-F., Gendreau, M., and Soriano, P. (2009). Accelerating benders decomposition by local branching. *INFORMS Journal on Computing*, 21(2):333–345.
- [Srivastava, 2008] Srivastava, S. K. (2008). Network design for reverse logistics. *Omega*, 36(4):535–548.
- [Wieberneit, 2008] Wieberneit, N. (2008). Service network design for freight transportation: a review. *OR Spectrum*, 30:77–112.

8 Appendix

In this appendix we present proofs of results discussed in the main body of this paper as well as detailed pseudo-code for the proposed Meta Partial Benders Decomposition algorithm.

Theorem 8.1. *The \mathcal{K} -enhanced master problem, **K-EMP**, is a relaxation of the Logistics Service Network Design problem, **LSNDP**.*

Proof. We prove this claim by showing that any feasible solution for the **LSNDP** is also feasible for the **K-EMP** and has the same objective function value. To do so, we let (x, y) be a feasible solution of the **LSNDP**. Let consider a solution (x^{Ξ}, y, z) such that:

$$x_{ij}^{\chi_k tt'} = \sum_{p \in \mathcal{P}_k} x_{ij}^{p tt'}, \quad \forall ((i, t), (j, t')) \in \mathcal{A}_{\mathcal{T}} \cup \mathcal{H}_{\mathcal{T}}, \forall \chi_k \in \Xi, \mathcal{P}^i \cap \mathcal{P}_k \neq \emptyset$$

$$z = \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} \sum_{\chi_k \in \Xi} c_{ij} x_{ij}^{\chi_k tt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_{\mathcal{T}}} \sum_{\chi_k \in \Xi} c_{ii} x_{ii}^{\chi_k tt+1}$$

It is easy to prove that this solution is feasible for the **K-EMP**. By construction, for each variable $x_{ij}^{p tt'}$ in the **LSNDP**, $p \in \mathcal{P}_k$, there is a corresponding variable $x_{ij}^{\chi_k tt'}$ in the **K-EMP**. We know that for each

warehouse (j, t') and each product $p \in \mathcal{P}$: $\sum_{((i,t),(j,t')) \in \mathcal{A}_T \cup \mathcal{H}_T} x_{ij}^{ptt'} - \sum_{((j,t'),(l,t'')) \in \mathcal{A}_T \cup \mathcal{H}_T} x_{jl}^{pt't''} = 0$. If we sum this expression on the products of \mathcal{P}_k , we obtain:

$$\begin{aligned} & \sum_{((i,t),(j,t')) \in \mathcal{A}_T \cup \mathcal{H}_T} \sum_{p \in \mathcal{P}_k} x_{ij}^{ptt'} - \sum_{((j,t'),(l,t'')) \in \mathcal{A}_T \cup \mathcal{H}_T} \sum_{p \in \mathcal{P}_k} x_{jl}^{pt't''} = 0 \\ \iff & \sum_{((i,t),(j,t')) \in \mathcal{A}_T \cup \mathcal{H}_T} x_{ij}^{\chi_k t t'} - \sum_{((j,t'),(l,t'')) \in \mathcal{A}_T \cup \mathcal{H}_T} x_{jl}^{\chi_k t' t''} = 0 \end{aligned}$$

Therefore (x^{Ξ}, y, z) respects constraints (17). As (x, y) respects constraints (3)-(4)-(5), it is trivial to demonstrate (x^{Ξ}, y, z) also respects constraints (18)-(19)-(20). By construction of z , (x^{Ξ}, y, z) respects constraint (21) which makes it an admissible solution for the **K-EMP**. Let $Q(x, y)$ be the objective value of (x, y) :

$$\begin{aligned} Q(x, y) &= \sum_{((i,t),(j,t')) \in \mathcal{A}_T} f_{ij} y_{ij}^{tt'} + \sum_{((i,t),(j,t')) \in \mathcal{A}_T} \sum_{p \in \mathcal{P}} c_{ij} x_{ij}^{ptt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_T} \sum_{p \in \mathcal{P}} c_{ii} x_{ii}^{ptt+1} \\ &= \sum_{((i,t),(j,t')) \in \mathcal{A}_T} f_{ij} y_{ij}^{tt'} + \sum_{((i,t),(j,t')) \in \mathcal{A}_T} \sum_{k=1}^{\mathcal{K}} \sum_{p \in \mathcal{P}_k} c_{ij} x_{ij}^{ptt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_T} \sum_{k=1}^{\mathcal{K}} \sum_{p \in \mathcal{P}_k} c_{ii} x_{ii}^{ptt+1} \\ &= \sum_{((i,t),(j,t')) \in \mathcal{A}_T} f_{ij} y_{ij}^{tt'} + \sum_{((i,t),(j,t')) \in \mathcal{A}_T} \sum_{\chi_k \in \Xi} c_{ij} x_{ij}^{\chi_k t t'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_T} \sum_{\chi_k \in \Xi} c_{ii} x_{ii}^{\chi_k t t+1} \\ &= \sum_{((i,t),(j,t')) \in \mathcal{A}_T} f_{ij} y_{ij}^{tt'} + z = Q(x^{\Xi}, y, z) \end{aligned}$$

Solution (x^{Ξ}, y, z) that replicates solution (x, y) by an aggregation of flows, is feasible to the **K-EMP**. The two solutions have an identical objective function value. Thus, **K-EMP** is a relaxation of LSNDP. \square

We next present the proof that in the special case that given a \mathcal{K} -partition of the products $\{\mathcal{P}_1, \dots, \mathcal{P}_{\mathcal{K}}\}$, if for all subsets of \mathcal{P}_k , each supplier that provides any product $p \in \mathcal{P}_k$ can also provide all other products included in \mathcal{P}_k , then all product subsets of the \mathcal{K} -partition can be aggregated to yield a **K-EMP** that is equivalent to the original problem.

Theorem 8.2. *Let $\{\mathcal{P}_1, \dots, \mathcal{P}_{\mathcal{K}}\}$ be a \mathcal{K} -partition of the products, and $\chi_k \in \Xi, k \in \{1, \dots, \mathcal{K}\}$, the associated super-products. If for each $k \in \{1, \dots, \mathcal{K}\}$, each supplier $s \in \mathcal{S}$ that provides any product $p \in \mathcal{P}_k$ can also provide all other products included in \mathcal{P}_k , then the **K-EMP** is equivalent to the LSNDP.*

Proof. We proved in subsection 4.2 that the **K-EMP** is a relaxation of the LSNDP. We demonstrate now that, in the case stated above, the LSNDP is a relaxation of the **K-EMP**. The premise of the proof is as follows. A solution of the **K-EMP** can be seen as a set of paths transporting super-products from suppliers to customers. More specifically, for each super-product χ_k , paths from suppliers of χ_k to customer (c, t) transit $D_{ct}^{\chi_k}$ units of χ_k . Thus, for each customer, the incoming flow of χ_k can be divided into $|\mathcal{P}_k|$ parts of quantity d_{ct}^p each, $p \in \mathcal{P}_k$. By converting each part into a flow of the corresponding product p , one can construct a solution for the LSNDP. Since each suppliers of χ_k provide all products of \mathcal{P}_k , the obtained solution is feasible for the LSNDP.

Let $(x^{\bar{=}}, y, z)$ be a feasible solution of the **K-EMP**. Let Λ be the set of paths from suppliers to customers in $\mathcal{G}_{\mathcal{T}}$. By definition the flow of super-products $x^{\bar{=}}$ originates from suppliers, ends at customers, and respects flow conservation at each warehouse. Therefore we can decompose the flow of super-product $x^{\bar{=}}$ into paths of Λ , and obtain an equivalent path-solution $\gamma^{\bar{=}}$. We denote $\gamma_{\lambda}^{\chi_k}$ the quantity of k^{th} super-product transported along a path $\lambda \in \Lambda$.

Let Λ_{ct} be the set of paths from suppliers to customer (c, t) . As $(\gamma^{\bar{=}}, y, z)$ respects the **K-EMP** demand constraint, for each super-product χ_k the sum of flows along paths of Λ_{ct} sustains (c, t) demand of χ_k , i.e. $\sum_{\lambda \in \Lambda_{ct}} \gamma_{\lambda}^{\chi_k} = D_{ct}^{\chi_k} = \sum_{p \in \mathcal{P}_k} d_{ct}^p$. Thus, for each a customer (c, t) , we can partition its incoming path-flow of super-product χ_k into $|\mathcal{P}_k|$ parts $\{\dot{\gamma}^1, \dots, \dot{\gamma}^{|\mathcal{P}_k|}\}$ transiting d_{ct}^p units of χ_k each, $p \in \mathcal{P}_k$. As a result, for each product $p \in \mathcal{P}_k$ the total flow along paths in the p^{th} subset of the flow-partition equals d_{ct}^p , i.e. $\sum_{\lambda \in \Lambda_{ct}} \dot{\gamma}_{\lambda}^p = d_{ct}^p$. In addition, for each path $\lambda \in \Lambda_{ct}$, the sum of flows over the flow-partition equals the total flow of super-product χ_k , i.e. $\sum_{p \in \mathcal{P}_k} \dot{\gamma}_{\lambda}^p = \gamma_{\lambda}^{\chi_k}$.

Based on the flow-partition of $\gamma^{\bar{=}}$, we construct a path-solution γ for the LSNDP. We denote γ_{λ}^p the quantity of product p transiting along path λ . For each customer (c, t) , each $k \in \{1, \dots, \mathcal{K}\}$, and each product $p \in \mathcal{P}_k$, if $\dot{\gamma}_{\lambda}^p > 0$ then the supplier of origin manufactures χ_k in the **K-EMP**. By hypothesis, the supplier of origin manufactures all products of \mathcal{P}_k in the original problem. As a result, $\forall p \in \mathcal{P}_k$ we can set any value for γ_{λ}^p . Thus, for each customer (c, t) , each $k \in \{1, \dots, \mathcal{K}\}$, each path $\lambda \in \Lambda_{ct}$, and each product $p \in \mathcal{P}_k$, we set γ_{λ}^p value to $\dot{\gamma}_{\lambda}^p$.

Let x be the flow-solution equivalent to path-solution γ . We now demonstrate that (x, y) is feasible for the LSNDP. By definition, on a path, flow conservation is ensured. Thus, flow-solution x respects constraints (2). For each customer $(c, t) \in \mathcal{C}_{\mathcal{T}}$, each $k \in \{1, \dots, \mathcal{K}\}$, and each product $p \in \mathcal{P}_k$ we have $\sum_{\lambda \in \Lambda_{ct}} \gamma_{\lambda}^p = d_{ct}^p$. By extension, the associated flow-solution x respects demand constraints (3). As $(x^{\bar{=}}, y, z)$ is a solution of the **K-EMP**, vehicle allocation y is sufficient to route $\gamma^{\bar{=}}$. For each customer (c, t) , each $k \in \{1, \dots, \mathcal{K}\}$, and each path $\lambda \in \Lambda_{ct}$, $\sum_{p \in \mathcal{P}_k} \gamma_{\lambda}^p = \gamma_{\lambda}^{\chi_k}$ therefore vehicle allocation y is sufficient to route γ , and so x . Thus, (x, y) respects constraints (5), and is a feasible solution to the LSNDP.

Let $(x^{\bar{=}}, y, z)$ objective value be:

$$Q(x^{\bar{=}}, y, z) = \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} f_{ij} y_{ij}^{tt'} + z$$

Due to constraints (21) we have:

$$Q(x^{\bar{=}}, y, z) = \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} f_{ij} y_{ij}^{tt'} + \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} \sum_{k=1}^{\mathcal{K}} c_{ij} x_{ij}^{\chi_k tt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_{\mathcal{T}}} \sum_{k=1}^{\mathcal{K}} c_{ii} x_{ii}^{\chi_k tt+1}$$

By partitioning, for each $k \in \{1, \dots, \mathcal{K}\}$, for each product $p \in \mathcal{P}_k$ and for each path $\lambda \in \Lambda_{ct}$: $\sum_{p \in \mathcal{P}_k} \gamma_{\lambda}^p = \gamma_{\lambda}^{\chi_k}$. Regarding flow-solutions x and $x^{\bar{=}}$ associated to γ and $\gamma^{\bar{=}}$, that means on each arc, the sum of flows over the products of \mathcal{P}_k equals the flow of super-product χ_k , i.e. $\sum_{p \in \mathcal{P}_k} x_{ij}^{ptt'} = x_{ij}^{\chi_k t t'}$. Therefore:

$$Q(x^{\bar{=}}, y, z) = \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} f_{ij} y_{ij}^{tt'} + \sum_{((i,t),(j,t')) \in \mathcal{A}_{\mathcal{T}}} \sum_{k=1}^{\mathcal{K}} \sum_{p \in \mathcal{P}_k} c_{ij} x_{ij}^{ptt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_{\mathcal{T}}} \sum_{k=1}^{\mathcal{K}} \sum_{p \in \mathcal{P}_k} c_{ii} x_{ii}^{ptt+1}$$

$$= \sum_{((i,t),(j,t')) \in \mathcal{A}_T} f_{ij} y_{ij}^{tt'} + \sum_{((i,t),(j,t')) \in \mathcal{A}_T} \sum_{p \in \mathcal{P}} c_{ij} x_{ij}^{ptt'} + \sum_{((i,t),(i,t+1)) \in \mathcal{H}_T} \sum_{p \in \mathcal{P}} c_{ii} x_{ii}^{ptt+1} = Q(x, y)$$

Solution (x, y) that replicates solution $(x^{\bar{=}}, y, z)$ by partitioning of flows, is feasible to the LSNDP. The two solutions have an identical objective function value. Thus, the LSNDP is a relaxation of the **K-EMP**. As the **K-EMP** is also a relaxation of the LSNDP, we demonstrated that the **K-EMP** and the LSNDP are equivalent if for each $k \in \{1, \dots, \mathcal{K}\}$, each supplier $s \in \mathcal{S}$ that provides any product $p \in \mathcal{P}_k$ can also provide all other products included in \mathcal{P}_k . \square

We next present the proof that given a **K-EMP** for a given \mathcal{K} , one can always create a **K-EMP** based on \mathcal{K} super-products that is at least as strong.

Theorem 8.3. *Given a \mathcal{K} -enhanced master problem, $\mathcal{K} \in \{1, \dots, |\mathcal{P}| - 1\}$, there always exist a $\mathcal{K} + 1$ -enhanced master problem such that **K-EMP** is a relaxation of **K+1-EMP***

Proof. We demonstrate that for any **K-EMP** there exists a **K+1-EMP** such that any feasible solution for the **K+1-EMP** is also feasible for the **K-EMP**, and has the same objective function value. For this, we refine the \mathcal{K} -partition associated to the **K-EMP** by dividing one of its subsets into two parts.

Let consider a **K-EMP** based on a \mathcal{K} -partition of the product set: $\{\mathcal{P}_1, \dots, \mathcal{P}_{\mathcal{K}}\}$. Subsets are sorted in ascending order with respect to their number of elements; thus the last subset $\mathcal{P}_{\mathcal{K}}$ contains at least two products. Let $\{\mathcal{P}'_1, \dots, \mathcal{P}'_{\mathcal{K}+1}\}$ be a $\mathcal{K} + 1$ -partition of the product set, such that $\mathcal{P}_i = \mathcal{P}'_i$, $\forall i \in \{1, \dots, \mathcal{K} - 1\}$, and $\mathcal{P}'_{\mathcal{K}} \cup \mathcal{P}'_{\mathcal{K}+1} = \mathcal{P}_{\mathcal{K}}$, $\mathcal{P}'_{\mathcal{K}} \neq \emptyset$, $\mathcal{P}'_{\mathcal{K}+1} \neq \emptyset$.

In the resulting **K+1-EMP**, super-product χ'_k is equivalent to super-product χ_k of the **K-EMP**, $\forall k \in \{1, \dots, \mathcal{K} - 1\}$. Thus, for each $k \in \{1, \dots, \mathcal{K} - 1\}$ and for each variable $x_{ij}^{\chi'_k tt'}$ in the **K+1-EMP**, there is a corresponding variable $x_{ij}^{\chi_k tt'}$ in the **K-EMP**. In addition, as super-products $\chi'_{\mathcal{K}}$ and $\chi'_{\mathcal{K}+1}$ include products in $\mathcal{P}_{\mathcal{K}}$, for each variable $x_{ij}^{\chi'_{\mathcal{K}} tt'}$ and each variable $x_{ij}^{\chi'_{\mathcal{K}+1} tt'}$ in the **K+1-EMP**, there is a corresponding variable $x_{ij}^{\chi_{\mathcal{K}} tt'}$ in the **K-EMP**.

Let $(x^{\bar{=}}, y, z)$ be a feasible solution for the **K+1-EMP**. Let consider a solution $(x^{\bar{=}}, y, z)$ for the **K-EMP**, such that the flow value of super-product χ_k is similar to the flow value of super-product χ'_k in the **K+1-EMP**, $\forall k \in \{1, \dots, \mathcal{K} - 1\}$, and the flow value of super-product $\chi_{\mathcal{K}}$ aggregates the flow values of super-products $\chi'_{\mathcal{K}}$ and $\chi'_{\mathcal{K}+1}$ in the **K+1-EMP**. By using a similar reasoning to that in the proof of Theorem 4.1, it is trivial to demonstrate that solution $(x^{\bar{=}}, y, z)$ is feasible for the **K-EMP**, and has the same objective function value that $(x^{\bar{=}}, y, z)$. Thus, **K-EMP** is a relaxation of **K+1-EMP**. \square

Algorithm 2: Phase I

Data: Maximum number of super-products, \mathcal{K}_{max} , Time limit on bounds improvement, t^{bounds} , Minimal improvement required for the bounds, $impr^{bounds}\%$, Threshold on the number of master solutions explored, $msols_{max}$, \mathcal{K} -partitions of the product set for \mathcal{K} ranging from 1 to \mathcal{K}_{max} , Time limit, t_1^{max}

$\mathcal{K} = 1$

$K^- = \emptyset$ and $K^+ = \emptyset$

$LB = 0$ and $UB = \infty$

while *Time left AND No optimal solution found* **do**

 Formulate the **K-EMP**

 Initiate it with all Benders cuts found previously, LB, and UB

begin Solve **K-EMP** using the algorithm proposed by Belieres et al.

 [Belieres et al., 2020b]:

if *No time left OR optimal solution found* **then**

 └ Break

else if *UB/LB did not improved more than $impr^{bounds}$ in the last t^{bounds} seconds* **then**

if *Less than $msols_{max}$ produced in the last t^{bounds} seconds* **then**

if *\mathcal{K} can increase* **then**

 └ $\mathcal{K} = \lceil \frac{\mathcal{K} - K^-}{2} \rceil$

 └ Break

else

if *\mathcal{K} can decrease* **then**

if $K^+ = \emptyset$ **then**

 └ $\mathcal{K} = \lceil \frac{\mathcal{K}_{max} - \mathcal{K}}{2} \rceil$

 └ Break

else

 └ $\mathcal{K} = \lceil \frac{K^+ - \mathcal{K}}{2} \rceil$

 └ Break

 Update K^- and K^+

 Keep all Benders cuts in memory

 Update best lower bound LB and best upper bound UB

Result: Best lower bound, LB, and best upper bound, UB

Algorithm 3: Phase II

Data: Initial lower bound, LB, Initial upper bound, UB, Time limit t_2^{max}

1-partition \leftarrow whole product set \mathcal{P}

for \mathcal{K} ranging from 1 to $|\mathcal{P}| - 1$ **do**

 Identify product subset from the \mathcal{K} -partition with a matching rate lower than 1

if *No product subset identified* **then**

 Break

else

 Partition it by applying the 2-medoids algorithm, using distance $1 - m(p, p')$ for each $(p, p') \in \mathcal{P}$

$\mathcal{K} + 1$ -partition \leftarrow combine the obtained subsets with all other product subsets from the \mathcal{K} -partition unchanged

Formulate the **K-EMP** associated to the last partition found

Initiate it with LB and UB

Solve **K-EMP** with a generic branch-and-cut implementation using a time limit t_2^{max}

Result: Final solution, UB
