CIRRELT

CIRRELT-2026-06

The Two-Echelon Multicommodity Location-Routing Problem with Stochastic and Correlated Demands

David Escobar-Vargas
Teodor Gabriel Crainic
Walter Rei
Stein W. Wallace

February 2026

Bureau de Montréal
Université de Montréal
C.P. 6128, succ. Centre-Ville
Montréal (Québec) H3C 3J7
Tél : 1-514-343-7575
Télécopie : 1-514-343-7121

Bureau de Québec
Université Laval,
2325, rue de la Terrasse
Pavillon Palasis-Prince, local 2415
Québec (Québec) G1V 0A6
Tél : 1-418-656-2073
Télécopie : 1-418-656-2624

The Two-Echelon Multicommodity Location-Routing Problem with Stochastic and Correlated Demands†

David Escobar-Vargas[1,2], Teodor Gabriel Crainic[1,3,*], Walter Rei[1,3], Stein W. Wallace[4]

[1] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

[2] Department of Computer Science and Operations Research, Université de Montréal

[3] School of Management, Université du Québec à Montréal

[4] Centre for shipping and logistics, Norwegian School of Economics (NHH), Bergen, Norway

Abstract. This study introduces the two-echelon multicommodity location routing problem with stochastic and correlated demands. We propose a two-stage stochastic programming formulation, with second-echelon facilities location and commodity allocation decisions defining the first stage, while recourse decisions, which are made in the second stage, establish how the observed demands are distributed. The overall objective is to minimize the cost of the first-stage design decisions plus the total expected routing cost incurred in the second stage. We propose a Progressive Hedging matheuristic that heuristically solves the formulation, incorporating a series of algorithmic enhancements to accelerate the exploration of the solution space. These enhancements include: 1) population structures to obtain alternative and diverse solutions for the scenario subproblems that need to be solved throughout the search process; 2) alternative strategies to define the reference solutions which are used to guide and accelerate the overall search; and 3) a reset procedure that reduces the risk of the method becoming trapped in local optima. We assess the efficiency and effectiveness of all proposed strategies through extensive computational experiments, evaluating their capability to generate high-quality solutions across various problem characteristics and demand correlations.

Keywords: Two-echelon location-routing problem, stochastic demand, multicommodity, origin-destination demands, progressive hedging

Acknowledgements. While working on the project, the second author held the UQAM Chair in Intelligent Logistics and Transportation Systems Planning and was Adjunct Professor, Department of Computer Science and Operations Research, Université de Montréal, Canada, while the third author held the Canada Research Chair in Stochastic Optimization of Transport and Logistics Systems. We gratefully acknowledge the financial support provided by the Natural Sciences and Engineering Council of Canada (NSERC) through its Discovery and Collaborative Research & Development grant programs, as well as by the Fonds de recherche du Québec through their Teams and CIRRELT infrastructure grants.

†Revised version of CIRRELT-2023-38

---

* Corresponding author: teodorgabriel.crainic@cirrelt.net

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2026

1 Introduction

The two-echelon location-routing problem (2ELRP) is an important class of combinatorial optimization problems with a wide range of applications, particularly in the freight transportation industry. At its core, it involves designing a freight transportation system that enables indirect deliveries from distribution centers (platforms) to customers via a set of intermediate facilities called satellites. The 2ELRP has been recognized as a suitable methodology for jointly addressing location decisions at one or two levels of facilities (platforms and/or satellites) together with the design of a limited set of routes at both echelons, in order to effectively serve all customer demands (Cuda et al., 2015). Despite the growing number of scientific contributions and advances in this field, most research on 2ELRP has focused on models and solution methods for "classic" problem variants and deterministic cases, while uncertain factors are often overlooked (Gendreau et al., 2014).

Considering demand uncertainty and correlations is of great significance when planning decisions are involved. In logistics planning, which encompasses strategic and tactical choices in distribution network design, obtaining accurate information about customer demand variations is essential for effective long-term planning (Lium et al., 2009). Sources of demand uncertainty include fluctuations in volume, inaccuracies in forecasts, and unexpected variations between specific origin-destination pairs (Crainic et al., 2011). While studies on LRPs with stochastic demands often assume statistically independent request fluctuations, correlations are frequently observed in many logistics contexts (Verma and Campbell, 2019). In practice, demand values may exhibit varying degrees of positive or negative correlation across different customers (Bucci et al., 2006). Seasonal demand variations provide an example; although generally predictable, they still illustrate systematic dependence patterns. More generally, correlations and more complex covariation structures become particularly relevant for planning when systematic relationships exist among customer demands (e.g., across regions, product types, or time periods) (Heath and Jackson, 1994; Thapalia et al., 2012; Verma and Campbell, 2019; Mirhedayatian et al., 2019). Consequently, demand correlation can be assumed to exhibit a *mixed* nature, rather than being purely positive or simply independent. To the best of our knowledge, 2ELRPs that account for correlated and uncertain demands (specifically non-substitutable demands with a known origin and destination) have not yet been explored. This study aims to advance the understanding of how the integrated treatment of uncertain and correlated non-substitutable demands affects location and routing decisions. Our objective is to develop a methodology that addresses both the modelling and algorithmic challenges, thereby contributing to filling the gaps in the literature.

We thus introduce the *Two-Echelon Multicommodity Location-Routing Problem with*

Stochastic and Correlated Demands (*2E-MLRPSCD*) as a unified framework that integrates the considered attributes, namely stochastic multicommodity origin-destination demands with correlations. The problem focuses on location decisions related to the selection of satellite facilities and the allocation of multicommodity OD demands to these satellites, while also involving the design of a limited set of routes at both echelons to efficiently serve demand. To address uncertainty, we propose a stochastic programming approach aimed at devising a unified system design that remains cost-effective under diverse demand realizations. Specifically, we present a two-stage stochastic programming formulation, where satellite facility location and multicommodity allocation decisions define the first stage, while recourse decisions in the second stage determine how the observed demands are served. Demand uncertainty is represented through a finite set of scenarios that approximate the variability inherent in the planning context. However, scenario-based uncertainty modelling often results in large-scale formulations that may prove impractical to solve with standalone exact methods (King and Wallace, 2012).

In stochastic optimization, various methods can be employed to support decision-making under uncertainty. A widely used approach is the Sample Average Approximation (SAA) method, which estimates the expected objective by evaluating a set of sampled scenarios. SAA provides a simple and consistent way to approximate expectations, and under appropriate sampling, it converges to the true problem solution as the sample size increases (Kleywegt et al., 2002). However, the quality of its results depends critically on the scenario generation process. If correlations among uncertain parameters are not adequately represented in the sampled scenarios, important structural dependencies may be overlooked, potentially leading to less reliable or suboptimal solutions.

In contrast, decomposition-based algorithms such as Progressive Hedging (PH) provide a computational framework for tackling large-scale scenario-based problems. PH decomposes the stochastic program into scenario-specific subproblems and iteratively enforces consistency across scenarios, thereby improving tractability and solution quality. While PH itself does not generate scenarios, it can effectively exploit scenario sets that preserve correlations, and it can be combined with acceleration strategies to reduce computational overhead. These features make PH particularly well suited for complex systems with correlated uncertainties, such as those considered in this study.

We thus introduce a novel PH-based matheuristic to heuristically solve the 2E-MLRPSCD, building on the work in (Crainic et al., 2011) applied to the network design problem. Methodologically, the classical PH algorithm iteratively solves deterministic subproblems obtained from the scenario-based decomposition of the multi-scenario optimization model. At each iteration, the PH matheuristic tackles these scenario subproblems, generating scenario-specific solutions that indicate how decisions should be made for each scenario individually. Using these solutions, an aggregate point is computed to assess the current level of consensus across subproblem solutions. Since the objective is to reach a consensus solution (i.e., all scenario subproblems agree on the first-stage decision values), the PH matheuristic iteratively adjusts the subproblem formulations to incentivize agreement based on the aggregate point obtained at each iteration. As each iteration requires solving a set of complex discrete optimization models (specifically, a deterministic 2E-MLRP for each scenario), it is crucial to minimize the number of iterations for the algorithm to remain computationally efficient. To this end, we propose a series of novel enhancements designed to accelerate convergence toward a high-quality consensus.

First, we introduce a solution population structure for each scenario subproblem, equipping the algorithm with a diverse set of alternative efficient solutions. This diversity facilitates the identification of potential consensus decisions across scenarios. Second, we propose alternative strategies for defining the aggregate point by developing general methods that extract the most relevant information from scenario-specific solutions to better guide the search toward consensus. Third, we introduce a reset procedure to reduce the risk of the method becoming trapped in local optima. Although these strategies are proposed here in the context of solving the 2E-MLRPSCD, their definition is general, as they are embedded in the core of the PH search process and are therefore broadly applicable. In the computational study, we analyze the cost sensitivity, infrastructure usage, and a comparison between the uncertain and deterministic definition of the demand to derive insights into the effectiveness of the proposed solution method.

The remainder of the paper is organized as follows. Section 2 is dedicated to describing the problem definition. An overview of the related scientific literature is provided in Section 3. Section 4 presents the system modelling and the proposed mathematical formulation. The solution method we developed is described in Section 5. Computational results are then presented and analyzed in Section 6 and concluding remarks are made in Section 7.

2 Problem Definition

This section introduces the 2E-MLRPSCD, which involves solving a 2ELRP with stochastic and correlated multicommodity OD demands. The section is divided into two parts. Subsection 2.1 presents the physical problem setting of the 2E-MLRPSCD. Subsection 2.2 outlines the representation of the stochastic and correlated demands, as well as the problem's objective and key requirements.

2.1 The 2E-MLRPSCD Setting

The two-echelon system consists of three main components: *platforms* (primary facilities serving as demand origins), *satellites* (intermediate facilities), and *customers* (demand destinations).

This problem is strongly motivated by the context of two-tier city logistics systems. Tactical planning in such systems often requires the strategic deployment of intermediary infrastructure, such as satellite facilities. These facilities are temporary spaces rented at designated locations to support freight distribution operations. They serve to transfer freight between vehicles assigned to each echelon (or tier) during specific planning horizons, such as peak seasons or periods of elevated demand.

The commodities involved are external-to-city (E2C) demands, characterized by fluctuating volumes and diverse types, each with distinct demand patterns. Correlations (both positive and negative) are often observed between commodities destined for urban areas, whether they are of similar or different types. These correlations introduce additional challenges in balancing resource allocation and maintaining timely delivery schedules.

Formally, the 2E-MLRPSCD is represented as a complete weighted directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, with vertices $\mathcal{V} = \mathcal{P} \cup \mathcal{Z} \cup \mathcal{C}$, divided into three disjoint sets: platforms $\mathcal{P}$, satellites $\mathcal{Z}$, and customers $\mathcal{C}$. Platforms are large-sized facilities with a known set of commodities to be distributed to customers. Satellites are medium- to small-sized multimodal infrastructures that serve as intermediate facilities, enabling the consolidation and sorting of freight between the two transportation echelons involved in delivering goods to customers. Each satellite location $z \in \mathcal{Z}$ is associated with a limited storage capacity Q_z and a fixed opening cost F_z.

Demand is defined between platforms and customers, with each individual demand characterized by an origin, a destination, and a requested volume to be delivered. Let $\mathcal{K}$ denote the set of OD demands. In the deterministic version of the 2E-MLRPSCD, each OD demand $k \in \mathcal{K}$ is defined by a volume vol_k, an origin $O(k)$ corresponding to a platform node in $\mathcal{P}$, and a destination $D(k)$ corresponding to a customer node in $\mathcal{C}$. Additionally, a fixed allocation cost Δ_{zk} represents the cost of serving OD demand k through satellite $z \in \mathcal{Z}$.

Each arc $(i, j) \in \mathcal{A} = A^1 \cup A^2$ is associated with a non-negative travel cost ζ_{ij} incurred by a vehicle moving from node i to node j. The set A^1 denotes the arcs of the first echelon, representing connections between platforms $\mathcal{P}$ and satellites $\mathcal{Z}$, as well as between satellites themselves. The set A^2 includes the arcs of the second echelon, which correspond to connections from satellites to customers $\mathcal{C}$ and between customers.

Freight delivery is carried out by two homogeneous fleets of vehicles, $H = H^1 \cup H^2$, with limited load capacities cap_1 and cap_2, respectively assigned to the first and second echelons. Both fleets can transport any type of demand. Vehicles are assumed to be available at each existing facility for their respective echelon, and they start and end their routes at these facilities.

The problem considered involves the selection of satellite facilities, the allocation of OD demands to satellites, and the routing of vehicles at each echelon to deliver freight

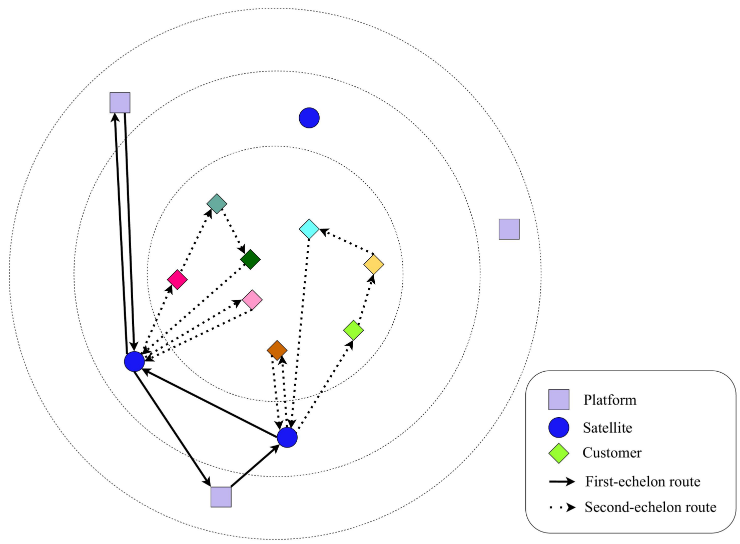

Figure 1: Topology of the 2E-MLRPSCD.

from platforms to customers via these satellite facilities. As illustrated in Figure 1, each OD demand originates at a platform and is first transported by a first-echelon vehicle to a selected satellite. There, the freight is transferred to a second-echelon vehicle. The loads arriving at satellites are thus transshipped and consolidated into second-echelon vehicles, which then carry out the final deliveries to customer destinations.

2.2 The stochastic setting

The 2E-MLRPSCD involves uncertainty in demand volumes arising from random fluctuations affecting correlated OD pairs. We assume that probability distributions are available to characterize these random variations. Correlations may be either positive or negative: to capture this structure, we define two sets of OD pairs. Pairs within the same set are positively correlated, whereas those across sets are strongly negatively correlated (i.e., low demand in one set typically coincides with high demand in the other).

The problem setting encompasses both strategic and tactical planning decisions relevant to multiple application domains. At the planning stage, location and allocation decisions must be made based on an evaluation of their impact on operations, including the available recourse actions to adapt the plan once demands are realized. In this context, recourse actions involve determining optimal vehicle routes to fulfill observed customer demands and, when necessary, resorting to outsourced services, which incur high additional operational costs.

The 2E-MLRPSCD consists of selecting the locations of satellite facilities, allocating OD demands to satellites, and constructing a limited set of routes for both first- and second-echelon vehicles. The solution must ensure that each platform's demand is assigned to an open satellite; first-echelon routes start and end at the same platform; second-echelon routes start and end at the same satellite; all customer demands are satisfied (either by the system or via outsourcing); vehicle capacities are not exceeded; each customer served by the system is visited by exactly one vehicle; satellite capacities are respected; and the total cost, comprising fixed location and allocation costs and expected routing (recourse) costs, is minimized.

3 Literature Review

The 2E-MLRPSCD belongs to the class of *Location-Routing Problems* (*LRPs*), a well-established and active area of research with numerous contributions in the literature. LRPs arise in planning processes where one or more facilities must be selected from a set of predefined locations, customers must be assigned to these facilities, and vehicle routes must be defined to satisfy the demand of each customer. Recent studies on LRPs, and particularly on two-echelon LRPs (2E-LRPs), are increasingly focusing on more realistic, multi-attribute problem settings (Escobar-Vargas and Crainic, 2024). This section aims to situate the 2E-MLRPSCD within the broader LRP and 2E-LRP literature, with a particular emphasis on the current knowledge gaps related to handling *stochastic demands* in such settings. A brief discussion of the *Progressive Hedging* strategy is also included, highlighting the challenges and limitations associated with its application to *integer programming problems*. Note that we restrict our review to works addressing demand uncertainty. Studies dealing with deterministic versions of the 2E-LRP or LRP, or considering other sources of uncertainty, are beyond the scope of this work. For comprehensive overviews of those areas, we refer the interested reader to the recent surveys (Cuda et al., 2015; Schiffer et al., 2019; Mara et al., 2021).

Due to its practical relevance, the LRP has attracted significant attention from the research community, resulting in a wide variety of high-quality solution approaches for its deterministic versions since its introduction in Maranzana (1964). While studies explicitly addressing demand uncertainty remain relatively scarce, interest in this variant has grown in response to the need for solving more realistic distribution planning problems (Cuda et al., 2015; Escobar-Vargas and Crainic, 2024). Given the complexity of incorporating demand uncertainty into LRP models, most studies have focused on proposing heuristic methods. The literature is notably characterized by the extensive use of local-search-based matheuristic frameworks to solve the underlying transportation subproblems and to guide two- or multi-stage heuristics, where location, allocation, and routing decisions are treated by different heuristics at separate stages (see, Albareda-Sambola et al., 2007; Huang, 2015; Marinakis, 2015; Marinakis et al., 2016; Zhang et al., 2019). A different approach is proposed in (Quintero-Araujo et al., 2019), where a simheuristic algorithm is developed to address the LRP with stochastic demands. This approach hybridizes a Monte Carlo simulation procedure with an iterated local search metaheuristic.

Despite these advances, the literature on LRPs under demand uncertainty remains quite limited, particularly regarding the modelling of *non-substitutable* demands. Most existing studies assume statistically independent stochastic demands, which remains the prevailing modelling approach. Actually, to the best of our knowledge, the only study explicitly addressing demand uncertainty in a two-echelon capacitated LRP is (Snoeck et al., 2018), which introduces a stochastic mixed-integer linear programming formulation motivated by a real-world application. However, research on incorporating correlated demand structures and analyzing their implications for decision-making remains scarce. Significant contributions are also needed to better understand how richer problem settings (particularly those involving complex forms of uncertainty) affect location decisions and overall planning outcomes.

Beyond the modelling considerations, there is also a fundamental need for more effective solution methods for the 2E-LRP under uncertainty. Towards this end, in both exact and approximate solution frameworks, decomposition-based methods have shown promising results for two- and multi-stage stochastic optimization models (Atakan and Sen, 2018). However, the effectiveness of such methods depends critically on how the stochastic problem can be decomposed.

Two general decomposition strategies are typically applied. The first decomposes the model by scenario, leveraging the structure induced by the representation of uncertainty. The second separates the model by decision stages, reflecting the temporal structure of the decisions. Among dual decomposition approaches, the Progressive Hedging (PH) algorithm is one of the most widely used. Originally introduced in (Rockafellar and Wets, 1991) for convex stochastic programs, PH decomposes the problem by scenario, solves each scenario subproblem independently, and iteratively seeks consensus across subproblem solutions through averaging.

However, applying PH to mixed-integer stochastic programs presents significant computational challenges due to the non-convexity of the feasible region (Atakan and Sen, 2018). To address this, various heuristic frameworks have been developed to adapt PH for integer programming contexts (see, e.g., Løkketangen and Woodruff, 1996; Haugen et al., 2001; Crainic et al., 2011; Lamghari and Dimitrakopoulos, 2016; Alvarez et al., 2021). These PH-based metaheuristics typically rely on generating high-quality (though not necessarily optimal) solutions for each scenario subproblem to guide the global search process. While this strategy provides effective guidance, relying exclusively on the best solution per scenario may limit the diversity of explored solutions, an element that can be critical to the performance of PH iterations and may hinder the ability to efficiently reach decision consensus.

The present work addresses these gaps by proposing a PH-based matheuristic tailored to the 2E-LRP. This approach is enhanced through specialized heuristics that generate diverse alternative solutions for each scenario subproblem, along with novel techniques to accelerate consensus-building across scenarios.

4 Modelling

Section 4.1 introduces the initial outline of the modelling approach, followed by the proposed mathematical formulation in Section 4.2.

4.1 Modelling uncertainty

The 2E-MLRPSCD is formulated as a two-stage stochastic program to account for strategic planning decisions. The proposed model includes a first stage, in which the locations of satellite facilities and the allocation of OD demand to satellites are determined, and a second stage, in which vehicle routes for both echelons are planned once customer demands are realized. Additionally, the second stage allows for the use of ad-hoc, outsourced capacity when necessary. This enables a portion of the commodity volumes to be delivered through methods outside the system's primary routing framework. An operational cost R is associated with the percentage of demand volume served via outsourced services, reflecting the additional expense incurred by relying on external delivery methods.

It is assumed that demands follow known probability distributions and that a corresponding correlation matrix is available. Based on these distributions and correlations, the problem is modeled by generating a set of scenarios, each representing a possible realization of the demands.

Let $\mathcal{S}$ denote the set of scenarios, where each scenario $s \in \mathcal{S}$ represents a possible realization of the random events that determine customer demand values and their correlations. Let ρ_s be the probability of occurrence of scenario s, such that $\sum_{s\in\mathcal{S}} \rho_s = 1$. For a given $s \in \mathcal{S}$, the demand volume for each customer $k \in \mathcal{K}$ is denoted by $vol_k(s)$, with $vol_k(s) \geq 0$.

4.2 Two-stage formulation for the 2E-MLRPSCD

This section presents the Mixed-Integer Programming (MIP) formulation for the 2E-MLRPSCD as a two-stage stochastic programming problem using a three-index vehicle-flow formulation. Two sets of decision variables are defined. First-stage variables address the satellite location and the allocation of OD demand to satellites. Vehicle-routing decisions at both echelons are made in the second stage. Following the general trend in the literature, we omit the compact form and present the formulation directly in terms of the set of scenarios $\mathcal{S}$. This results in second-stage variables being indexed by scenario, while first-stage variables are not, as they are fixed before any scenario is realized. The following definitions describe the decision variables that constitute the extensive form of the proposed two-stage formulation:

- $y_k \in \{0,1\},\ \forall k \in \mathcal{Z}$: binary location variable; equal to 1 if a satellite is located at site k, and 0 otherwise.
- $f_{zk} \in \{0,1\},\ \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K}$: binary allocation variable; equal to 1 if satellite z is assigned to serve the OD demand k, and 0 otherwise.
- $u^s_{zkh} \in \{0,1\},\ \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K},\ \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S}$: binary vehicle allocation variable; equal to 1 if vehicle h is assigned to serve satellite z for OD demand k in scenario s, and 0 otherwise.
- $v^s_{zch} \in \{0,1\},\ \forall z \in \mathcal{Z},\ \forall c \in \mathcal{C},\ \forall h \in \mathcal{H}^2,\ \forall s \in \mathcal{S}$: binary vehicle allocation variable; equal to 1 if vehicle h is assigned to serve customer c from satellite z in scenario s, and 0 otherwise.
- $x^s_{ijh} \in \{0,1\},\ \forall (i,j) \in \mathcal{A},\ \forall h \in \mathcal{H},\ \forall s \in \mathcal{S}$: binary vehicle flow variable; equal to 1 if arc (i,j) is traversed by vehicle h in scenario s, and 0 otherwise.
- $w^s_{zkh} \geq 0,\ \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K},\ \forall h \in \mathcal{H}^2,\ \forall s \in \mathcal{S}$: continuous variable; percentage of OD demand k served by satellite z using vehicle h in scenario s.
- $o^s_k \geq 0,\ \forall k \in \mathcal{K},\ \forall s \in \mathcal{S}$: continuous variable; percentage of OD demand k that is outsourced in scenario s.
- $b^s_{kh} \geq 0,\ \forall k \in \mathcal{K},\ \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S}$: continuous variable; percentage of OD demand k dispatched using vehicle h in scenario s.
- $L^s_{zh} \in \mathbb{Z}_{\geq 0},\ \forall z \in \mathcal{Z},\ \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S}$: integer variable; position of satellite z in the route assigned to first-echelon vehicle h in scenario s.
- $N^s_{ch} \in \mathbb{Z}_{\geq 0},\ \forall c \in \mathcal{C},\ \forall h \in \mathcal{H}^2,\ \forall s \in \mathcal{S}$: integer variable; position of customer c in the route assigned to second-echelon vehicle h in scenario s.

The extensive two-stage formulation of the 2E-MLRPSCD then becomes:

$$\min \sum_{s\in\mathcal{S}} \rho_s \left(\sum_{h\in\mathcal{H}} \sum_{(i,j)\in\mathcal{A}} \zeta_{ij} x^s_{ijh} + \sum_{k\in\mathcal{K}} R\, o^s_k \right) + \sum_{z\in\mathcal{Z}} F_z\, y_z + \sum_{z\in\mathcal{Z}} \sum_{k\in\mathcal{K}} \Delta_{zk}\, f_{zk} \tag{1}$$

subject to

$$\sum_{i\in P} \sum_{j\in\mathcal{Z}} x^s_{ijh} \leq 1 \qquad \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S} \tag{2}$$

$$\sum_{\substack{i\in(P\cup\mathcal{Z})\\ i\neq j}} x^s_{ijh} - \sum_{\substack{i\in(P\cup\mathcal{Z})\\ i\neq j}} x^s_{jih} = 0 \qquad \forall j \in (P\cup\mathcal{Z}),\ \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S} \tag{3}$$

$$L^s_{ih} - L^s_{jh} + |\mathcal{Z}|\, x^s_{ijh} \leq |\mathcal{Z}| - 1 \qquad \forall i \neq j \in \mathcal{Z},\ \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S} \tag{4}$$

$$\sum_{h\in\mathcal{H}^2}\sum_{\substack{j\in(\mathcal{Z}\cup\mathcal{C})\\ i\neq j}} x^s_{ijh} = 1 \qquad \forall i\in\mathcal{C},\ \forall s\in\mathcal{S} \tag{5}$$

$$\sum_{i\in\mathcal{Z}}\sum_{j\in\mathcal{C}} x^s_{ijh} \leq 1 \qquad \forall h\in\mathcal{H}^2,\ \forall s\in\mathcal{S} \tag{6}$$

$$\sum_{\substack{i\in(\mathcal{Z}\cup\mathcal{C})\\ i\neq j}} x^s_{ijh} - \sum_{\substack{i\in(\mathcal{Z}\cup\mathcal{C})\\ i\neq j}} x^s_{jih} = 0 \qquad \forall j\in(\mathcal{Z}\cup\mathcal{C}),\ \forall h\in\mathcal{H}^2,\ \forall s\in\mathcal{S} \tag{7}$$

$$\sum_{h\in\mathcal{H}^2}\sum_{j\in\mathcal{C}} x^s_{ijh} \leq |\mathcal{H}^2|\, y_i \qquad \forall i\in\mathcal{Z},\ \forall s\in\mathcal{S} \tag{8}$$

$$N^s_{ih} - N^s_{jh} + |\mathcal{C}|\, x^s_{ijh} \leq |\mathcal{C}| - 1 \qquad \forall i\neq j\in\mathcal{C},\ \forall h\in\mathcal{H}^2,\ \forall s\in\mathcal{S} \tag{9}$$

$$\sum_{\substack{j\in(\mathcal{Z}\cup\mathcal{C})\\ i\neq j}} x^s_{ijh} + \sum_{\substack{j\in(\mathcal{Z}\cup\mathcal{C})\\ l\neq j}} x^s_{ljh} - v^s_{lih} = 0 \qquad \forall i\in\mathcal{C},\ \forall l\in\mathcal{Z},\ \forall h\in\mathcal{H}^2,\ \forall s\in\mathcal{S} \tag{10}$$

$$\sum_{h\in\mathcal{H}^2}\sum_{i\in\mathcal{Z}} v^s_{ijh} = 1 \qquad \forall j\in\mathcal{C},\ \forall s\in\mathcal{S} \tag{11}$$

$$\sum_{h\in\mathcal{H}^1} u^s_{iO(k)h} = \sum_{h\in\mathcal{H}^2} v^s_{iD(k)h} \qquad \forall i\in\mathcal{Z},\ \forall k\in\mathcal{K},\ \forall s\in\mathcal{S} \tag{12}$$

$$\sum_{h\in\mathcal{H}^2}\sum_{i\in\mathcal{Z}} w^s_{ijh} + o^s_j = 1 \qquad \forall j\in\mathcal{K},\ \forall s\in\mathcal{S} \tag{13}$$

$$w^s_{ijh} \leq v^s_{iD(j)h} \qquad \forall i\in\mathcal{Z},\ \forall j\in\mathcal{K},\ \forall h\in\mathcal{H}^2,\ \forall s\in\mathcal{S} \tag{14}$$

$$b^s_{kh} \geq \sum_{i\in\mathcal{Z}} w^s_{ikl} - \left(2 - \sum_{i\in\mathcal{Z}} v^s_{iD(k)l} - \sum_{i\in\mathcal{Z}} u^s_{iO(k)h}\right) M \qquad \forall h\in\mathcal{H}^1,\ \forall l\in\mathcal{H}^2,\ \forall k\in\mathcal{K},\ \forall s\in\mathcal{S} \tag{15}$$

$$\sum_{k\in\mathcal{K}} vol_k(s) \sum_{h\in\mathcal{H}^2} w^s_{ikh} \leq Q_i \qquad \forall i\in\mathcal{Z},\ \forall s\in\mathcal{S} \tag{16}$$

$$\sum_{k\in\mathcal{K}} vol_k(s) \sum_{i\in\mathcal{Z}} w^s_{ikh} \leq \mathrm{cap}_2 \qquad \forall h\in\mathcal{H}^2,\ \forall s\in\mathcal{S} \tag{17}$$

$$\sum_{k\in\mathcal{K}} b^s_{kh} \leq \mathrm{cap}_1 \qquad \forall h\in\mathcal{H}^1,\ \forall s\in\mathcal{S} \tag{18}$$

$$\sum_{h\in\mathcal{H}^1} u^s_{iO(k)h} = f_{iO(k)} \qquad \forall i\in\mathcal{Z},\ \forall k\in\mathcal{K},\ \forall s\in\mathcal{S} \tag{19}$$

$$\sum_{h\in\mathcal{H}^2} v^s_{zD(k)h} = f_{zO(k)} \qquad \forall z\in\mathcal{Z},\ \forall k\in\mathcal{K},\ \forall s\in\mathcal{S} \tag{20}$$

$$y_z \in \{0,1\} \qquad \forall z\in\mathcal{Z} \tag{21}$$

$$f_{zk} \in \{0,1\} \qquad \forall z\in\mathcal{Z},\ \forall k\in\mathcal{K} \tag{22}$$

$$f_{zk} \in \{0,1\} \qquad \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K} \tag{23}$$
$$u^s_{zkh} \in \{0,1\} \qquad \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K},\ \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S} \tag{24}$$
$$v^s_{zch} \in \{0,1\} \qquad \forall z \in \mathcal{Z},\ \forall c \in \mathcal{C},\ \forall h \in \mathcal{H}^2,\ \forall s \in \mathcal{S} \tag{25}$$
$$x^s_{ijh} \in \{0,1\} \qquad \forall (i,j) \in \mathcal{A},\ \forall h \in \mathcal{H},\ \forall s \in \mathcal{S} \tag{26}$$
$$w^s_{zkh} \geq 0 \qquad \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K},\ \forall h \in \mathcal{H}^2,\ \forall s \in \mathcal{S} \tag{27}$$
$$o^s_k \geq 0 \qquad \forall k \in \mathcal{K},\ \forall s \in \mathcal{S} \tag{28}$$
$$b^s_{kh} \geq 0 \qquad \forall k \in \mathcal{K},\ \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S} \tag{29}$$
$$L^s_{zh} \in \mathbb{Z}_{\geq 0} \qquad \forall z \in \mathcal{Z},\ \forall h \in \mathcal{H}^1,\ \forall s \in \mathcal{S} \tag{30}$$
$$N^s_{ch} \in \mathbb{Z}_{\geq 0} \qquad \forall c \in \mathcal{C},\ \forall h \in \mathcal{H}^2,\ \forall s \in \mathcal{S}. \tag{31}$$

The objective function (1) seeks to minimize the sum of the expected total routing and outsourcing costs, as well as the total fixed costs associated with locating satellites and allocating them to the OD demands. Constraints (2) ensure that each first-echelon vehicle is assigned to at most one platform. Constraints (3) enforce flow conservation at both platforms and satellite facilities. Constraints (4) are Miller-Tucker-Zemlin (MTZ)-type sub-tour elimination constraints for the first-echelon routes. Constraints (5) ensure that each customer is visited by exactly one second-echelon vehicle. Constraints (6) ensure that each second-echelon vehicle is assigned to at most one satellite. Constraints (7) enforce flow conservation at satellites and customer locations. Constraints (8) ensure that second-echelon vehicles can only be operated from located satellites. Constraints (9) are MTZ-type sub-tour elimination constraints for second-echelon routes. Constraints (10) link routing and allocation variables. Constraints (11) guarantee that each customer is assigned to exactly one satellite. Constraints (12) enforce flow conservation for each OD demand k at each satellite z. Constraints (13) ensure that the sum of the portions of customer demand served by a satellite and by outsourced services meets the total demand for each customer. Constraints (14) ensure that satellites only serve customers to whom they have been assigned. Constraints (15) ensure that, for each OD demand k, the portion of the demand served via a located satellite corresponds to the inbound volume originating from the associated platform. Constraints (16) impose that the total flow dispatched from an open satellite z does not exceed its storage capacity. Constraints (17) and (18) ensure that the commodity flow carried by each vehicle, in the second and first echelons respectively, does not exceed the vehicle's capacity. Constraints (19) and (20) link the allocation and vehicle assignment variables for the first- and second-echelon vehicles, respectively. Finally, constraints (21)–(31) define the integrality and non-negativity domains for all decision variables in the model.

5 A PH-based Matheuristic for the 2E-MLRPSCD

As its name implies, the PH methodology is derived from the algorithm introduced in (Rockafellar and Wets, 1991) for multi-stage stochastic optimization problems. From a methodological standpoint, the *classical* PH algorithm iteratively solves the set of deterministic subproblems obtained from the scenario-based decomposition of the extensive formulation. At each iteration, PH solves each scenario-specific deterministic subproblem independently, resulting in a set of solutions that may differ from one another. The search then proceeds by computing a reference solution (typically the expected value of the scenario-specific solutions) which is also used to assess the overall level of consensus among them. The formulations of the scenario subproblems are then adjusted to incentivize agreement, i.e., to encourage convergence toward a common implementable solution. This process is repeated until either a consensus solution is reached or a stopping criterion is satisfied (e.g., a time limit).

It is well known that PH does not necessarily converge to an optimal solution when applied to mixed-integer programs, such as the 2E-MLRPSCD. An additional algorithmic challenge arises from the computational burden of solving a series of NP-hard problems (one for each scenario) at every PH iteration. There is thus a clear need for an efficient guiding strategy and supporting procedures to accelerate the search for a consensus solution. Inspired by the work in (Crainic et al., 2011) on the stochastic network design problem, we introduce a PH matheuristic that incorporates several algorithmic and methodological enhancements aimed at speeding up the identification of a high-quality implementable solution: (1) a population-based structure to generate diverse alternative solutions for each scenario subproblem; (2) novel scenario-selection strategies that extract key insights from subproblem solutions to guide the search toward consensus; (3) a specialized heuristic to define a high-quality reference solution at the first PH iteration; and (4) a reset procedure designed to prevent the PH matheuristic from becoming trapped in local optima.

5.1 General structure

The proposed PH matheuristic is illustrated in Figure 2. The algorithm begins with a scenario-based decomposition of the extensive formulation (Section 4.2), which yields a set of subproblems, each corresponding to a deterministic 2E-MLRPSCD instance for a given scenario $s \in \mathcal{S}$. In contrast to the algorithmic structure introduced in (Crainic et al., 2011), the proposed matheuristic generates a set of alternative solutions for each scenario subproblem rather than retaining only the single best one. This strategy is intended to broaden the design space (particularly for location and allocation decisions) and enhance the algorithm's ability to identify high-quality consensus solutions.

We introduce two population structures to manage the sets of alternative solutions generated for each scenario subproblem: a set of *local populations*, one per scenario subproblem, and a *global population* that supports the overall PH search.

Local populations are updated at each iteration of the PH algorithm. The objective is to retain the best-performing solutions for the corresponding scenario subproblems. To this end, a *ranking* measure is defined to evaluate both the *quality* (in terms of objective value) and the *diversity* of the solutions generated for each subproblem at every iteration. The ranking of a candidate solution is determined relative to the solutions already present in the local population and is used to decide whether it should be included in the set (potentially replacing the current worst-ranked solution). The goal is to maintain a local population composed of solutions that are not only high-quality with respect to the objective function but also diverse in terms of their first-stage decision variable values.

The global population is constructed at each iteration of the PH algorithm from the *elite* subset, i.e., the best solutions extracted from each local population. A reference solution is then computed based on a selected subset of solutions from the global population, using one of the proposed scenario selection strategies (Section 5.4). This reference solution is used to guide the search by modifying the cost terms in the objective function of each scenario subproblem, with the aim of steering the search toward consensus on the first-stage decisions across all scenarios.

The algorithm terminates either when a consensus is reached on the first-stage decisions or when an external stopping criterion is met. Throughout the process, the best feasible solution identified at each iteration of the PH algorithm is recorded. We now provide a detailed description of each step of the proposed PH matheuristic.

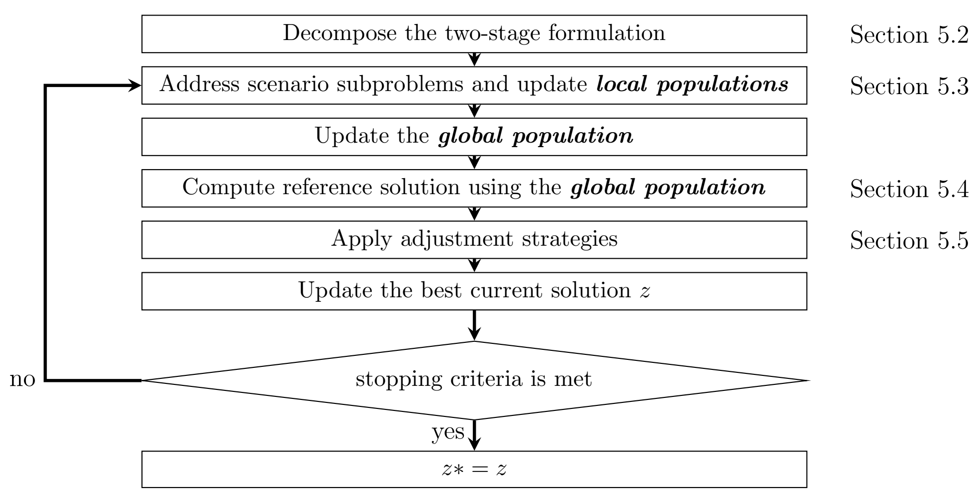

Figure 2: Progressive Hedging-based matheuristic for the 2E-MLRPSCD

5.2 Scenario decomposition for the 2E-MLRPSCD

The decomposition strategy applied to the extensive formulation requires that the first-stage decisions be expressed as scenario-dependent. To ensure consistency across scenarios, nonanticipativity constraints are explicitly introduced to enforce that the first-stage decision variables take the same values in all scenario subproblems (see Appendix A for details). Let y_z^s, for $z \in \mathcal{Z}$, and f_{zk}^s, for $z \in \mathcal{Z}$ and $k \in \mathcal{K}$, denote the scenario-specific location and allocation decision variables, respectively, for each scenario $s \in \mathcal{S}$. Let $\bar{y}_z$ and $\bar{f}_{zk}$ represent the corresponding first-stage decision variables in the current reference solution. Constraints (8), (19), and (20) are thus expressed using the scenario-specific variables, and the following nonanticipativity constraints are added to ensure alignment with the reference solution:

$$y_z^s = \bar{y}_z \quad \forall z \in \mathcal{Z},\ \forall s \in \mathcal{S} \tag{32}$$
$$f_{zk}^s = \bar{f}_{zk} \quad \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K},\ \forall s \in \mathcal{S} \tag{33}$$
$$\bar{y}_z \in \{0,1\} \quad \forall z \in \mathcal{Z} \tag{34}$$
$$\bar{f}_{zk} \in \{0,1\} \quad \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K} \tag{35}$$

Constraints (32) and (33) are relaxed using an augmented Lagrangian approach, with penalty parameter γ and Lagrange multipliers λ_z^s and μ_{zk}^s. This yields the following relaxed reformulation of the extensive model:

$$\begin{aligned}\min \sum_{s\in\mathcal{S}} \rho_s \Bigg(\sum_{h\in\mathcal{H}} \sum_{(i,j)\in\mathcal{A}} \zeta_{ij} x_{ijh}^s + \sum_{k\in\mathcal{K}} R\, o_k^s + \sum_{z\in\mathcal{Z}} \left(F_z + \lambda_z^s + \frac{1}{2}\gamma + \gamma \bar{y}_z \right) y_z^s \\ + \sum_{z\in\mathcal{Z}} \sum_{k\in\mathcal{K}} \left(\Delta_{zk} + \mu_{zk}^s + \frac{1}{2}\gamma + \gamma \bar{f}_{zk} \right) f_{zk}^s \Bigg)\end{aligned} \tag{36}$$

subject to constraints (2)–(7), (9)–(18), (34)–(35) and:

$$\sum_{h\in\mathcal{H}^2} \sum_{j\in\mathcal{C}} x_{ijh}^s \le |\mathcal{H}^2|\, y_i^s \quad \forall i \in \mathcal{Z},\ \forall s \in \mathcal{S} \tag{37}$$
$$\sum_{h\in\mathcal{H}^1} u_{iO(k)h}^s = f_{ik}^s \quad \forall i \in \mathcal{Z},\ \forall k \in \mathcal{K},\ \forall s \in \mathcal{S} \tag{38}$$
$$\sum_{h\in\mathcal{H}^2} v_{zD(k)h}^s = f_{zk}^s \quad \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K},\ \forall s \in \mathcal{S} \tag{39}$$
$$y_z^s \in \{0,1\} \quad \forall z \in \mathcal{Z},\ \forall s \in \mathcal{S} \tag{40}$$
$$f_{zk}^s \in \{0,1\} \quad \forall z \in \mathcal{Z},\ \forall k \in \mathcal{K},\ \forall s \in \mathcal{S} \tag{41}$$

For all $s \in \mathcal{S}$, and considering the scenario-specific first-stage variables, constraints (37) ensure that second-echelon vehicles can only operate from located satellites. Constraints (38) and (39) link the facility allocation variables to the corresponding vehicle assignment variables. Finally, constraints (40) and (41) enforce the binary nature of the scenario-specific location and allocation decision variables.

For a given reference solution $\bar{y}_i$ and $\bar{f}_{ik}$ (Section 5.5), the relaxed reformulation is then decomposed by scenario, resulting in a set of individual subproblems. Each scenario subproblem corresponds to a deterministic, scenario-specific instance in which the fixed costs are modified by the associated Lagrange multipliers λ_i^s and μ_{ik}^s. Together with the penalty parameter γ, these multipliers penalize discrepancies between the location and allocation decisions made in the local design and those defined in the current reference solution.

5.3 Subproblem algorithm

This section presents the method proposed to update the local population of each scenario subproblem at every iteration of the PH matheuristic. The objective is to maintain, throughout the PH search, local populations that preserve both the quality and diversity of the best solutions found while solving the corresponding scenario subproblems.

The MIP for each scenario subproblem $s \in \mathcal{S}$ consists of the objective function (36), constraints (2)–(7), (9)–(18), and (37)–(41), along with a complete *a priori* enumeration of the subtour elimination constraints. (For simplicity of presentation, variables are not marked with the PH iteration index.) Each feasible solution identified while solving the MIP is evaluated and ranked to determine whether it should be included in the local population (replacing the one currently with the lowest rank) or discarded. Each local population Γ_s, for $s \in \mathcal{S}$, is characterized by a fixed size ψ_T, which includes a reduced number ψ_E of elite solutions.

The rank of a solution Sol is determined by a *fitness* measure $F(Sol)$ that combines its rank based on objective quality, $RKQ(Sol)$, and its rank based on contribution to the local population diversity, $RKD(Sol)$. The combined fitness is computed as:

$$F(Sol) = RKQ(Sol) + \left(1 + \frac{\psi_E}{\psi_T}\right) RKD(Sol), \tag{42}$$

where ψ_E and ψ_T denote the number of elite solutions and the total population size, respectively.

Let $\Xi(Sol)$ denote the *diversity contribution* of candidate solution Sol, defined relative to its average dissimilarity from the ψ_T solutions already present in the population. Let $\phi_k(Sol)$, for $k \in \mathcal{K}$, denote the *negative correlation score* of commodity k in solution Sol, defined as the number of other commodities (i.e., $k' \neq k$) that are assigned to the same satellite and share a negative correlation with k. Note that $\phi_k(Sol)$ highlights an opportunity for improving system efficiency through the consolidation of negatively correlated demands (King and Wallace, 2012). Inspired by the methodological developments proposed in (Vidal et al., 2012), we propose a normalized Hamming distance $\sigma(Sol, Sol')$ to measure the dissimilarity between two distinct solutions $Sol \neq Sol'$, based on both their satellite allocation decisions $\xi_k(Sol)$ and their corresponding negative correlation scores $\phi_k(Sol)$, for all $k \in \mathcal{K}$. Let $\mathbf{1}(\text{cond})$ be the indicator function, which returns 1 if the condition cond is true and 0 otherwise. The proposed Hamming distance and the resulting diversity contribution of solution Sol are defined by equations (43) and (44), respectively:

$$\sigma(Sol, Sol') = \frac{1}{2|\mathcal{K}|} \sum_{k \in \mathcal{K}} \left[\mathbf{1}(\xi_k(Sol) \neq \xi_k(Sol')) + \mathbf{1}(\phi_k(Sol) \neq \phi_k(Sol')) \right], \quad \forall Sol' \in \Gamma_s, \tag{43}$$

$$\Xi(Sol) = \frac{1}{|\Gamma_s|} \sum_{Sol' \in \Gamma_s} \sigma(Sol, Sol'). \tag{44}$$

Recall that: 1) the solution ranking and local population updating processes occur while solving each scenario subproblem; 2) local populations are updated rather than rebuilt from scratch at each PH iteration. This continuous learning process turns local populations into a form of memory for the PH algorithm, allowing them to retain well-ranked solutions from previous iterations and to provide diversified elite solutions for computing the next reference solution.

5.4 Defining the reference solution

The new reference solution is computed from the solutions in the global population, which is constructed from scratch after all scenario subproblems have been processed. The global population gathers the elite solutions from all local populations and has a total size of $\psi_G = \psi_E \times |\mathcal{S}|$.

Let Λ_s denote the set of elite solutions for scenario $s \in \mathcal{S}$ included in the global population constructed at iteration ν of the proposed PH matheuristic. Let y^{ν}_{az} and f^{ν}_{azk} be the values of the first-stage variables in solution $Sol_a \in \Lambda_s$, while $\bar{y}^{\nu}_{z}$ and $\bar{f}^{\nu}_{zk}$ denote the corresponding variables in the reference solution constructed at iteration ν.

We present four *selection strategies* for determining the solutions involved in computing the new reference solution. These include a *classic* strategy, adapted from the procedure traditionally used in PH-based methods (Crainic et al., 2011), as well as three novel strategies. A comprehensive description of these strategies is provided in the following sections, followed by a presentation of the specialized heuristic proposed to define the reference solution for the first PH iteration.

5.4.1 Classic strategy

This strategy follows the approach proposed in (Crainic et al., 2011), which is based on the guidelines originally defined for the Progressive Hedging (PH) method in (Rockafellar and Wets, 1991). The reference solution is computed by aggregating, into a weighted average, the best solutions obtained for each scenario subproblem $s \in \mathcal{S}$. The probability of occurrence ρ_s associated with each scenario $s \in \mathcal{S}$ is used as the corresponding weight.

We implement this strategy by setting the size of the elite population for all scenario subproblems to $\psi_E = 1$. The new reference solution is then computed using equations (45) and (46):

$$\bar{y}_z^{\nu} = \sum_{s \in \mathcal{S}} \sum_{a \in \Lambda_s} \rho_s \, y_{az}^{s\nu}, \qquad \forall z \in \mathcal{Z}, \tag{45}$$

$$\bar{f}_{zk}^{\nu} = \sum_{s \in \mathcal{S}} \sum_{a \in \Lambda_s} \rho_s \, f_{azk}^{s\nu}, \qquad \forall z \in \mathcal{Z},\ k \in \mathcal{K}. \tag{46}$$

Notice that $\bar{y}_z^{\nu} \in \{0,1\}$, $\forall z \in \mathcal{Z}$, and $\bar{f}_{zk}^{\nu} \in \{0,1\}$, $\forall z \in \mathcal{Z},\ k \in \mathcal{K}$, at a given iteration ν, indicate that the method has reached consensus for the first-stage decision variables, and that the PH matheuristic has found an implementable solution for the stochastic problem. However, in most cases, the integrality of some first-stage variables is not enforced, i.e., $0 < \bar{y}_z^{\nu} < 1$ and $0 < \bar{f}_{zk}^{\nu} < 1$, meaning that the current reference solution is infeasible. Yet, these fractional values reflect trends in the decision-making process. Hence, one can interpret $\bar{y}_z^{\nu} \approx 0$ as a tendency not to select facility i, while $\bar{y}_z^{\nu} \approx 1$ indicates the opposite (that is, a trend toward selecting facility i). The same interpretation applies to the reference solution values associated with the allocation decisions $\bar{f}_{zk}^{\nu}$.

5.4.2 Probabilistic strategy

The aggregation operators for the three proposed strategies take into account the complete elite populations Λ_s, $s \in \mathcal{S}$, of the scenario subproblems (i.e., $\psi_E = |\Lambda_s| \geq 1$). This allows for a broader range of options to be considered for the first-stage decision variables.

The probabilistic strategy generalizes the classic aggregation operator by distributing the scenario probability equally among the solutions in its elite set Λ_s. The reference solution at iteration ν is then computed using equations (47) and (48):

$$\bar{y}_z^{\nu} = \sum_{s \in \mathcal{S}} \frac{\rho_s}{|\Lambda_s|} \sum_{a \in \Lambda_s} y_{az}^{s\nu} \qquad \forall z \in \mathcal{Z}, \tag{47}$$

$$\bar{f}_{zk}^{\nu} = \sum_{s \in \mathcal{S}} \frac{\rho_s}{|\Lambda_s|} \sum_{a \in \Lambda_s} f_{azk}^{s\nu} \qquad \forall z \in \mathcal{Z},\ k \in \mathcal{K}. \tag{48}$$

5.4.3 Social strategy

The idea of this strategy is to favour solutions with the highest *social scores* within the global population, i.e., solutions that share the most similarities in the first-stage location and allocation decisions.

Let $\pi(Sol_a, Sol_b)$ denote the similarity measure between two distinct solutions, $Sol_a = (\tilde{y}, \tilde{f})$ and $Sol_b = (\hat{y}, \hat{f})$, computed by equation (49). Here, $\kappa_k(\cdot)$ returns the allocation of commodity $k \in \mathcal{K}$ within a given solution. The *social score* of a solution Sol_a with respect to all the other solutions in the global population Λ is then defined as

$$\text{Score}(Sol_a) = \sum_{Sol_b \in \Lambda,\, b \neq a} \pi(Sol_a, Sol_b),$$

$$\pi(Sol_a, Sol_b) = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \mathbf{1}\Big(\kappa_k(\tilde{f}) = \kappa_k(\hat{f})\Big) + \frac{1}{|\mathcal{Z}|} \sum_{z \in \mathcal{Z}} \mathbf{1}(\tilde{y}_z = \hat{y}_z)\,. \tag{49}$$

The solutions in the global population are then ranked based on their social scores. A reduced set of elite solutions with the highest scores is selected, the sets Λ_s are adjusted accordingly, and the reference solution is computed using equations (47) and (48). It should be mentioned that using only the single top-ranked elite solution and setting its first-stage decisions as the reference solution proved unproductive in preliminary experiments.

5.4.4 Decision-based scenario clustering strategy

This strategy identifies scenario sets that lead to mutually acceptable solutions, i.e., solutions that remain efficient when evaluated across all subproblems associated with the scenarios in the set. It relies on a *dissimilarity function*, inspired by the opportunity cost measure proposed in (Hewitt et al., 2022), to quantify the impact of implementing the first-stage decisions associated with one scenario when another scenario is realized. We extend this idea by leveraging the elite solutions of the scenario subproblems in the global population.

Let $g((y^\nu_{ai}, f^\nu_{ai}), s_j)$ denote the value of the objective function (36) evaluated with the first-stage decision variables y^ν_{ai} and f^ν_{ai} of solution $Sol_a \in \Lambda_{s_i}$ at iteration ν, when scenario s_j occurs. We define the *opportunity cost* $\theta(s_i \mid s_j)$ as the decision value associated with scenario s_i when scenario s_j is realized. Formally, $\theta(s_i \mid s_j)$ is calculated by equation (50) as the minimum value obtained by evaluating all combinations of elite solutions associated with each pair of distinct scenarios $s_i, s_j \in \mathcal{S}$, $i \neq j$. The *opportunity cost dissimilarity function* for that same pair of scenarios, equation (51), represents the loss incurred by optimizing under the assumption that one scenario occurs when in fact the other materializes:

$$\theta(s_i \mid s_j) = \min_{\substack{a\in\Lambda_{s_i}\\ b\in\Lambda_{s_j}}} \{g((y_a^\nu, f_a^\nu), s_j) - g((y_b^\nu, f_b^\nu), s_j)\}, \tag{50}$$

$$d(s_i \mid s_j) = \theta(s_i \mid s_j) + \theta(s_j \mid s_i). \tag{51}$$

Then, a *Normalized Spectral Clustering* procedure is applied to determine which scenarios are close to each other in terms of the *opportunity cost dissimilarity function* (Hewitt et al., 2022). This yields a set of scenario clusters $CL = \{cl_1, cl_2, \dots, cl_{|CL|}\}$. We define the set of *representative scenarios* Υ, where each scenario $s \in \Upsilon$ corresponds to the *medoid* of its cluster, i.e., the scenario with the minimum average opportunity cost dissimilarity to all other scenarios within the cluster. The probability η_i associated with each representative scenario $s_i \in \Upsilon$ is then computed as the sum of the probabilities ρ_s of all scenarios in the same cluster (equation (52)). Finally, the reference solution is computed using equations (53) and (54):

$$\eta_i = \sum_{s\in cl_i} \rho_s, \qquad \forall cl_i \in CL, \tag{52}$$

$$\bar{y}_z^\nu = \sum_{s\in\Upsilon} \frac{\eta_s}{|\Lambda_s|} \sum_{a\in\Lambda_s} y_{az}^{s\nu}, \qquad \forall z \in \mathcal{Z}, \tag{53}$$

$$\bar{f}_{zk}^\nu = \sum_{s\in\Upsilon} \frac{\eta_s}{|\Lambda_s|} \sum_{a\in\Lambda_s} f_{azk}^{s\nu}, \qquad \forall z \in \mathcal{Z},\ k \in \mathcal{K}. \tag{54}$$

5.4.5 First iteration reference solution

An initial reference solution $(\bar{y}_z, \bar{f}_{zk}), z \in \mathcal{Z}, k \in \mathcal{K}$ must be defined after the first resolution of the scenario subproblems. At this stage, however, no *memory* is available to assess the quality of these solutions, since the populations are still empty. Defining a high-quality reference solution at the end of the first iteration is crucial, as the consensus search proceeds by successively adjusting the objective function costs of the scenario subproblems to gradually encourage agreement. The quality of the first-iteration decisions therefore has a strong influence on all subsequent iterations.

We propose a heuristic to define the reference solution for the first iteration of the PH matheuristic and to initialize the global population. Recall that the global population is composed of at least one elite solution selected from the local population of each scenario subproblem. The *quality* of the initial reference solution therefore depends both on the quality of the solutions present in the initial global population and on the specific selection strategy used to obtain the point.

The proposed heuristic proceeds by generating two initial global populations, GP_1 and GP_2, obtained by applying the probabilistic and classic strategies, respectively. A reference solution is then computed for each population. If a reference solution contains decision variables with continuous rather than integer values, the heuristic *approximates* it by rounding each continuous value to the nearest discrete value. Finally, each reference solution is evaluated in the extensive formulation, and the one yielding the best objective function value is selected.

Note that when GP_2 is selected as the best initial global population, the local populations are reduced to their best elite solutions. All subsequent updates then proceed as described previously. It is also worth noting that preliminary experiments showed that applying the proposed heuristic at every iteration of PH is not effective, as the method tends to become trapped in local optima.

5.5 Consensus procedure

We propose two heuristics to adjust the costs of the scenario subproblems, with the aim of guiding the PH method toward consensus on the first-stage solutions across all scenario subproblems. Inspired by the work in (Crainic et al., 2011), these two adjustment heuristics modify the location and allocation costs, respectively, both globally to influence the overall search and locally to steer the search within each scenario subproblem.

The proposed global adjustment begins with the reference solution, defined by $\bar{y}_z^\nu$ and $\bar{f}_{zk}^\nu$ at iteration ν, to identify trends across the scenario solutions. The costs are defined according to the objective function (36). Locally, we define the location and allocation costs of each scenario subproblem as

$$\bar{B}_z^{s\nu} = \left(F_z + \lambda_z^s + \frac{1}{2}\gamma + \gamma\bar{y}_z\right) \quad \text{and} \quad \bar{E}_{zk}^{s\nu} = \left(\Delta_{zk} + \mu_{zk}^s + \frac{1}{2}\gamma + \gamma\bar{f}_{zk}\right), \forall z \in \mathcal{Z},\, k \in \mathcal{K}$$

respectively.

As mentioned previously, low values of $\bar{y}_z^\nu$ and $\bar{f}_{zk}^\nu$ indicate that most scenario solutions share the decision to exclude the facility at location z and the associated allocation (i.e., allocating commodity k to satellite z). Conversely, high values imply that the opposite decision prevails in the majority of the scenario solutions. To operationalize this observation, we introduce a parameter $\beta > 1$ as the cost adjustment rate, along with threshold parameters $0 \leq \epsilon^y \leq 0.5$ and $0 \leq \epsilon^f \leq 0.5$, which determine when the values of $\bar{y}_z^\nu$ and $\bar{f}_{zk}^\nu$ should be interpreted as low or high. Specifically, when $\bar{y}_z^\nu$ and $\bar{f}_{zk}^\nu$ fall below ϵ^y and ϵ^f, respectively, the corresponding fixed costs are increased to discourage locating the facility and performing the associated allocation. On the other hand, when these values exceed $1 - \epsilon^y$ and $1 - \epsilon^f$, the fixed costs are reduced to encourage including the facility in the network design and performing the associated allocation. This adjustment procedure is formalized by the following update rules:

$$\bar{B}_z^{\nu} = \begin{cases} \beta B_z^{\nu-1}, & \text{if } \bar{y}_z^{\nu-1} < \epsilon^y, \\ \frac{1}{\beta} B_z^{\nu-1}, & \text{if } \bar{y}_z^{\nu-1} > 1 - \epsilon^y, \\ B_z^{\nu-1}, & \text{otherwise;} \end{cases} \tag{55}$$

$$\bar{E}_{zk}^{\nu} = \begin{cases} \beta \bar{E}_{zk}^{\nu-1}, & \text{if } \bar{f}_{zk}^{\nu-1} < \epsilon^f, \\ \frac{1}{\beta} \bar{E}_{zk}^{\nu-1}, & \text{if } \bar{f}_{zk}^{\nu-1} > 1 - \epsilon^f, \\ \bar{E}_{zk}^{\nu-1}, & \text{otherwise.} \end{cases} \tag{56}$$

The second adjustment strategy is applied at the level of each scenario subproblem $s \in \mathcal{S}$. In this case, the costs of variables that significantly deviate from the current reference solution at iteration ν are further adjusted using equations (57) and (58). Let $0.5 < \delta^y < 1$ and $0.5 < \delta^f < 1$ denote the thresholds that determine when a local adjustment is triggered for the location and allocation variables, respectively:

$$\bar{B}_z^{s\nu} = \begin{cases} \beta B_z^{\nu}, & \text{if } \left| y_z^{s(\nu-1)} - \bar{y}_z^{\nu-1} \right| \geq \delta^y \text{ and } y_z^{s(\nu-1)} = 1, \\ \frac{1}{\beta} B_z^{\nu}, & \text{if } \left| y_z^{s(\nu-1)} - \bar{y}_z^{\nu-1} \right| \geq \delta^y \text{ and } y_z^{s(\nu-1)} = 0, \\ B_z^{\nu}, & \text{otherwise;} \end{cases} \tag{57}$$

$$\bar{E}_{zk}^{s\nu} = \begin{cases} \beta \bar{E}_{zk}^{\nu}, & \text{if } \left| f_{zk}^{s(\nu-1)} - \bar{f}_{zk}^{\nu-1} \right| \geq \delta^f \text{ and } f_{zk}^{s(\nu-1)} = 1, \\ \frac{1}{\beta} \bar{E}_{zk}^{\nu}, & \text{if } \left| f_{zk}^{s(\nu-1)} - \bar{f}_{zk}^{\nu-1} \right| \geq \delta^f \text{ and } f_{zk}^{s(\nu-1)} = 0, \\ \bar{E}_{zk}^{\nu}, & \text{otherwise.} \end{cases} \tag{58}$$

The proposed PH matheuristic is designed to terminate once either a consensus solution is found or another stopping criterion is met (e.g., a computation time limit). A consensus solution is said to be reached when all first-stage decisions $\bar{y}_z^{\nu}$ and $\bar{f}_{zk}^{\nu}$, $\forall z \in \mathcal{Z}$, $k \in \mathcal{K}$, have converged at iteration ν. However, full consensus across all first-stage decisions may not be achieved at the end of each iteration of the PH matheuristic. In such cases, the PH is designed to produce a feasible solution for the 2E-MLRPSCD by using the extensive formulation introduced in Section 4.2. This solution procedure fixes the location and allocation variables for which consensus has been obtained and then solves the corresponding restricted mixed-integer program, as defined by the extensive formulation. The resulting solution is guaranteed to be feasible for all design decisions.

5.6 Reset procedure

As described previously, the proposed PH algorithm relies on the global population to determine the reference solution at each iteration. This global population is composed of the elite solutions drawn from the local populations of each scenario subproblem. As the search progresses, certain solutions may persist in the local populations of some scenario subproblems for multiple PH iterations. As a result, the global population may also remain unchanged over several consecutive iterations.

If the global population does not evolve, the reference solution may become trapped in a set of values that hinder progress toward consensus. To mitigate this issue, we propose a reset procedure that partially reinitializes the overall search process. Specifically, the reset is triggered when the reference solution remains unchanged for ι consecutive iterations. When triggered, the procedure clears all local populations and repopulates them with solutions obtained from their respective scenario subproblems during the current iteration. It is important to note that while this reset defines a new set of candidate solutions for the global population, the current cost coefficients for the first-stage decisions in each scenario subproblem (updated over the previous PH iterations) are preserved.

6 Computational Results

The computational experiments conducted assess: (1) the stability of the scenario generation procedure (Section 6.2); (2) the performance of the proposed PH-based matheuristic (Section 6.3); (3) the effectiveness of the proposed acceleration procedures for the 2E-MLRPSCD; and (4) the importance of explicitly considering stochastic demands and the impact of correlations (Sections 6.5 and 6.4). Section 6.1 first introduces the instances and the scenario generation procedure.

All experiments were carried out on a single machine equipped with an Intel(R) Core(TM) i7-7800X processor and 128 GB of RAM running Linux. The mathematical formulation and the proposed solution method were implemented in `C++` using IBM ILOG CPLEX Concert Technology 20.1. The MIPs within the solution method were solved with an optimality gap tolerance of 1% as the stopping criterion. Computation times are reported in seconds. The tables present summarized results, while more detailed results are provided in Appendix B (supplementary material).

6.1 Instances

We define our testbed based on the instances introduced in (Dellaert et al., 2019) for the 2EVRPTW, since no instances were available in the literature involving the integrated treatment of all the attributes considered in the 2E-MLRPSCD. The instances proposed in (Dellaert et al., 2019) simulate an urban area composed of platforms, satellites, and customers. The original instances do not consider stochastic correlated OD demands, which are explicitly included in the 2E-MLRPSCD. Furthermore, the original instances include delivery time windows, which are not considered in the present setting. Therefore, adjustments were made to the original instances to construct the testbed for the present study. These adjustments involved introducing stochastic and correlated OD demands and removing the temporal components from the original instances.

Our instance set consists of 60 instances, each with 15 OD demands. We randomly assigned to each platform facility a unique set of OD demands. The same vehicle load capacities, as defined in (Dellaert et al., 2019), were used. First-level vehicles therefore have a capacity of $cap_1 = 200$, and second-level vehicles have a capacity of $cap_2 = 50$. Travel costs are computed as the ceiling of the Euclidean distances.

category	distribution	mean	standard deviation
CA	left-skewed lognormal	2.7	0.4
CB	symmetrical lognormal	2.7	0.1
CC	left-skewed lognormal	3.25	0.4
CD	symmetrical lognormal	3.25	0.1

Table 1: Instance category description

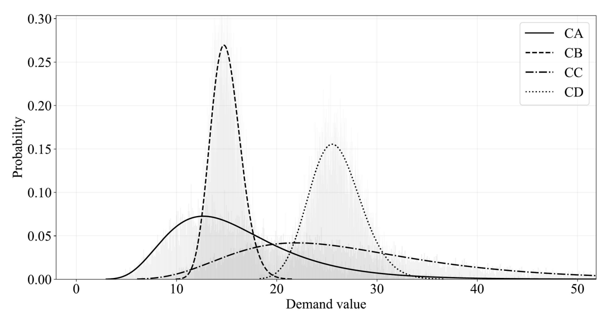

Figure 3: Instance category distribution for scenario generation.

Scenarios are generated using the copula-based method proposed in (Kaut, 2014) to preserve the statistical properties defined for the stochastic OD demands. This procedure requires, as input, the target marginal distribution for each OD demand (which can be specified using one of the marginal distribution types available in the method) and the correlation matrix between OD pairs.

To determine the marginal distributions for all OD demands, we considered the original demand values in the instance set defined in (Dellaert et al., 2019). We identified the distribution that provided the closest fit (among the marginal distributions available in the copula-based method) to the OD demand values across the complete instance set. This analysis led us to select a lognormal distribution (with similar mean and standard deviation values) as the best fit for the given demand data. We then defined a set of four lognormal distributions with different mean and standard deviation values to capture the impact of variations in the marginal distributions representing demand. Table 1 introduces the proposed instance set, categorized into four groups: CA, CB, CC, and CD.

As illustrated in Figure 3, a lognormal distribution with consistent mean and standard deviation values, as defined in Table 1, is used for each instance category.

To define demand correlation, two testbed instances are proposed: one considering demand correlation and the other without it. In the first case, correlation matrices are randomly generated. Correlations between OD pairs are determined using a standard normal distribution, with values restricted to the range $[-0.6, 0.6]$ for each correlation coefficient.

A properly defined correlation matrix must be positive semidefinite (Xu and Evers, 2003). Before applying the scenario generation method, this condition is verified for each correlation matrix obtained; matrices that do not satisfy this property are ignored by the copula-based method. Scenarios are generated once the positive semidefinite condition is confirmed. The copula-based heuristic uses the mean and standard deviation of the distributions for each instance category, along with the correlation matrix defined for each instance, to generate a predefined number of scenarios $|S|$ with equal probability. Thus, the probability of occurrence ρ_s for scenario $s \in \mathcal{S}$ is given by $\rho_s = 1/|S|$.

6.2 Scenario stability

This section presents the computational experiments conducted to assess the stability of the chosen scenario generation procedure, both with and without demand correlation. The goal of assessing scenario stability is to ensure that the scenario trees used do not significantly influence the results obtained when solving the stochastic problem (Kaut and Wallace, 2007). In our case, we employed the copula-based method presented in (Kaut, 2014) to generate the scenario sets. Unlike other approaches (e.g., sampling methods), this method has a high probability of producing identical scenario trees across repeated runs when the same correlation structure and distributional inputs are used. The use of 'standard' in-sample and out-of-sample stability tests (see, e.g., Kaut and Wallace, 2007) is not appropriate in this context, as these tests may overestimate the quality of the scenario generation method (Guo et al., 2019). Instead, we build on the work of (Zhang et al., 2021) to develop a valid variant of the standard approach that is better suited to our problem setting.

Based on the guidelines proposed in (Zhang et al., 2021), stability tests require generating and evaluating a subset of scenario trees for each problem instance, using a fixed scenario tree size. To test the stability of a scenario tree of size $|S|$, a set of $2m + 1$ scenario trees is defined with sizes $|S| - m, |S| - (m - 1), \ldots, |S|, \ldots, |S| + m$, where m is a positive integer. Let $Z_{|S|+i}$ denote the optimal (or best-known) solution for each $i \in [-m, m]$ of the $2m + 1$ scenario trees obtained for a given problem instance. The proposed PH matheuristic is used to solve the 2E-MLRPSCD corresponding to each of the $2m+1$ scenario trees, resulting in $2m+1$ solutions $Z_{|S|+i}$, one per scenario tree. Each solution $Z_{|S|+i}$ is then evaluated by computing its objective function value $F(Z_{|S|+i})$ on each of the $2m+1$ scenario trees, yielding $2m+1$ objective function values for every $Z_{|S|+i}$. For each solution $Z_{|S|+i}$, the maximum ($F^+(Z_{|S|+i})$), minimum ($F^-(Z_{|S|+i})$), and variance ($\sigma_{|S|+i}$) of these values are recorded. Stability is then assessed by computing the relative difference (RD) and the variance (VAR) as follows:

$$RD = \max_{i \in [-m,m]} \left\{ \frac{F^+(Z_{|S|+i}) - F^-(Z_{|S|+i})}{F^+(Z_{|S|+i})} \times 100\% \right\} \tag{59}$$

$$VAR = \max_{i \in [-m,m]} \left\{ \sigma_{|S|+i} \right\} \tag{60}$$

Table 2 presents a summary of the RD and VAR values obtained for the $2m+1$ scenario trees defined for each problem instance. The table reports the number of scenarios for each scenario tree ($|S|$), along with the minimum (MIN), average (AVR), and maximum (MAX) values of the relative difference and variance. To assess stability, a threshold criterion of RD $\leq$ 2% is applied. Two additional performance measures are introduced: the number of instances satisfying the stability criterion, referred to as the **valid stability requirement (VSR)**, and the average RD of the instances that fail to meet the criterion, denoted as the **average invalid stability requirement (AISR)**. Experiments are conducted using multiple scenario trees with varying numbers of scenarios ($|S|$), with the parameter m set to 4, following the approach proposed in (Guo et al., 2019). To reduce noise in the results, only the best objective function value obtained for each instance by the proposed PH matheuristic is considered in the stability tests.

The results reported in Table 2 show the expected decrease in the RD values as the number of scenarios increases. It is worth noting that, due to the inherent randomness and heuristic nature of the scenario generation procedure, small fluctuations in the relative error are observed for some instances. However, these do not compromise the validity of the proposed stability tests. Overall, using 30 scenarios offers the best trade-off between solution stability and scenario size, yielding a relative error below 2% for 53 out of 60 instances, and an average relative error of 2.6% for the remaining instances. While increasing the number of scenarios would further reduce the relative error, solving the corresponding subproblems for all instances using CPLEX at each iteration of the PH matheuristic becomes prohibitively time-consuming.

Figure 4 displays the RD values for each instance type as a function of the number of scenarios used in the stability tests. It can be observed that the dispersion of the demand distributions has a significant impact on solution stability. In particular, instances of types CA and CC, which are characterized by more dispersed demand distributions, exhibit larger fluctuations in RD values. This behavior can be attributed to the copula-based method generating a more diverse set of scenarios, which increases the volatility of the objective function values due to greater variability in the corresponding recourse actions.

In conclusion, the results presented in Table 2 indicate that using $|S| = 30$ provides the best trade-off between solution stability and scenario tree size. Moreover, Figure 4

	RD (%)					VAR		
$\|S\|$	VSR	AISR	MIN	AVERAGE	MAX	MIN	AVERAGE	MAX
10	38	3.96	0	1.79	8.58	0	1041.89	15222.89
20	47	3.17	0	1.13	5.93	0	335.62	3351.09
30	54	2.65	0	0.82	3.55	0	190.04	1944.3
40	53	2.25	0	0.58	2.25	0	108.05	1142.38
50	53	2.01	0	0.45	2.01	0	74.62	933.34
100	60	n.a	0	0.28	1.44	0	32.52	525.93

Table 2: Stability tests: summarized results of relative difference and variance for different scenario sizes.

shows that increasing the number of scenarios beyond $|S| = 30$ yields only marginal improvements in relative difference values. This suggests that employing scenario trees with $|S| > 30$ is not advantageous, as the associated computational burden becomes significantly higher at each iteration of the PH matheuristic.

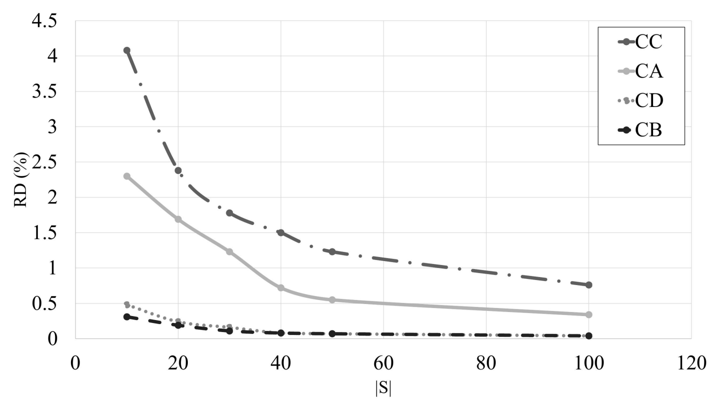

Figure 4: Stability test: Relative difference for each instance type vs scenario size.

6.3 Performance of the PH matheuristic

We compared the performance of the PH matheuristic with that of CPLEX for solving the complete stochastic model. The results presented in this section focus on the quality of the upper bounds obtained and the computational time required by each solution method. All methods were subject to a maximum time limit of 2 hours. In addition, the PH matheuristic was restricted to a maximum of 60 iterations. CPLEX was run with default parameter settings, using a thread limit of 6 for solving the full stochastic model, and a thread limit of 1 for solving the scenario subproblems within the PH matheuristic.

To address the stochastic problem, computational tests were conducted by solving the complete two-stage stochastic model with CPLEX or by employing the proposed PH matheuristic. Furthermore, experiments were performed using a *classic strategy* as a baseline for the PH matheuristic. This approach replicates the steps of the 'classic' PH matheuristic proposed in (Crainic et al., 2011) (described in Section 5.4.1). In all tables, the results obtained using CPLEX and the PH matheuristic are labeled CPLEX and PH, respectively. The PH matheuristic results are further differentiated by the specific version of the strategy used: classic (CL), probabilistic (PS), social (SS), and decision-based clustering (DCS). Each table reports the average optimality gap expressed as a percentage (OG), the average computational time in seconds, and the average upper bound differences between PH and CPLEX (Diff. UB).

Table 3 presents the results for instances with no demand correlation. It can be observed that solving the full stochastic problem using CPLEX is challenging: CPLEX achieves an average optimality gap of 21% within the 2-hour time limit. The PH matheuristic outperforms CPLEX. In particular, the classic approach (CL) achieves an average improvement of 14.5% in solution quality over CPLEX. Moreover, the classic approach reaches consensus and generates high-quality upper bounds for 53 out of the 60 instances within the time limit. However, the exclusive use of the best-quality solutions from each scenario subproblem to define the reference solution is not sufficiently effective for reaching consensus across the entire set of first-stage decisions. This adverse effect is especially pronounced in the CA and CC instances, where the demand distributions are more dispersed. As a result, the corresponding scenario subproblems tend to produce more diverse first-stage decisions, hindering the consensus process.

Compared to the classic approach, the proposed strategies are able to achieve consensus for the entire set of instances. The probabilistic strategy, which builds on the classic approach, yields significant improvements in both solution quality and runtime. Specifically, it achieves an average improvement of 16.2% over CPLEX and a 35.3% reduction in average runtime compared to the classic approach. To efficiently reach consensus, the PH matheuristic thus benefits from incorporating alternative solutions for each scenario subproblem. This increases the number of complementary first-stage decisions considered in the aggregation process at each iteration. Despite the general improvements obtained with the probabilistic strategy, the social and clustering strategies are able to more effectively leverage the available alternative solutions, further enhancing the overall performance of the algorithm.

The social strategy consistently yields reduced runtimes, with the largest average decrease of 68.3% compared to the classic approach. The consensus-driven mechanism, which ranks the global population, also helps the PH matheuristic reach consensus more quickly and reduce computational time.

However, reaching consensus faster does not necessarily lead to higher-quality solutions. This is illustrated by comparing the results obtained with the decision-based clustering strategy and the social strategy. Although PH runs using the social strategy are, on average, 40% faster, the decision-based clustering strategy produces higher-quality consensus solutions. On average, the PH matheuristic achieves an optimality gap of 2.6% with the decision-based clustering strategy, compared to 6.8% when using the social strategy.

Instance type	**CPLEX**		**PH**											
			CL			**PS**			**SS**			**DCS**		
	OG (%)	**time (s)**	**diff. ub**	**OG (%)**	**time (s)**	**diff. ub**	**OG (%)**	**time (s)**	**diff. ub**	**OG (%)**	**time (s)**	**diff. ub**	**OG (%)**	**time (s)**
CA	30.91	7200	-7.02	28.24	4458.16	-13.54	23.71	1323.66	-38.65	8.09	884.99	-44.98	3.76	2295.22
CB	23.66	7200	-32.54	1.54	3568.78	-32.54	1.54	1350.34	-32.62	1.47	1081.83	-32.66	1.44	1239.97
CC	17.42	7200	-9.07	10.07	4651.69	-9.30	9.87	4243.97	-9.04	10.21	2326.02	-19.35	1.95	2442.96
CD	15.79	7200	-9.19	8.11	4520.49	-9.32	8.01	4210.59	-10.11	7.38	1151.61	-14.92	3.47	2200.99
Averages	***21.94***	***7200***	***-14.46***	***11.99***	***4299.78***	***-16.17***	***10.78***	***2782.14***	***-22.61***	***6.79***	***1361.11***	***-27.98***	***2.65***	***2044.79***

Table 3: Summarized results on instances with no demand correlation.

Instance type	**CPLEX**		**PH**											
			CL			**PS**			**SS**			**DCS**		
	OG (%)	**time (s)**	**diff. ub**	**OG (%)**	**time (s)**	**diff. ub**	**OG (%)**	**time (s)**	**diff. ub**	**OG (%)**	**time (s)**	**diff. ub**	**OG (%)**	**time (s)**
CA	28.70	7200	1.49	29.77	7200.00	-29.94	10.14	2393.54	-38.05	4.71	2178.57	-42.06	2.35	1663.38
CB	24.35	7200	-33.52	0.86	4071.12	-33.76	0.67	1384.67	-33.76	0.67	1989.33	-33.88	0.58	1920.12
CC	13.03	7200	-11.17	3.57	4622.19	-12.50	2.47	2095.32	-12.71	2.29	2167.33	-13.64	1.48	2091.91
CD	11.77	7200	-7.76	5.10	4809.67	-10.45	2.84	1609.31	-11.35	2.02	2061.57	-11.87	1.56	1781.87
Averages	***19.46***	***7200***	***-12.74***	***9.82***	***5175.74***	***-21.66***	***4.03***	***1870.71***	***-23.97***	***2.42***	***2099.20***	***-25.36***	***1.49***	***1864.32***

Table 4: Summarized results on instances with demand correlation.

Computational results for instances with demand correlation are reported in Table 4. As observed in the case of instances without demand correlation, solving the complete two-stage formulation with CPLEX results in the worst optimality gap, reaching the maximum time limit for all instances. In contrast, the classic approach shows significant performance improvements, achieving consensus for 40 out of the 60 instances within the time limit. Both the probabilistic and social strategies demonstrate considerable gains in solution quality compared to CPLEX, with average optimality gaps of 4% and 2.4%, respectively. Among all selection strategies, the decision-based clustering strategy yields the best overall performance in terms of both solution quality and runtime. It produces the best solutions for 50 out of the 60 instances and achieves an average runtime reduction of 64% compared to the classic approach.

The performance of the PH matheuristic with each proposed aggregation strategy varies significantly when tested on instances with and without demand correlation. Scenario trees generated under the assumption of entirely uncorrelated demands often include scenarios with a large number of high demand values. These high-demand scenarios tend to dominate the solution structure, as the first-stage decisions made under such conditions are more likely to remain feasible across scenarios with lower demand values. This effect is less pronounced in instances where negative demand correlation is introduced. As the degree of negative correlation increases, the likelihood of generating scenarios with consistently high demand values decreases, resulting in greater diversity among the demand values across scenarios. This increased variability leads to a more heterogeneous set of first-stage decisions, which poses additional challenges for the PH matheuristic in achieving consensus. This general phenomenon explains the improved performance of the proposed acceleration strategies in settings where demand correlation is considered.

6.4 Correlation impact

Our experiments consider three sets of instances, each corresponding to a different level of demand correlation. Scenarios are generated for instances involving 15 OD demand pairs, under weak, moderate, and strong correlation settings. These correlation levels are defined using a rescaled normal distribution, as described in Table 5. Figure 5 illustrates how the correlation structure differs across these settings for the same instance.

Computational studies are conducted using the proposed PH framework with the decision-based clustering strategy, selected for its superior overall performance. We compare the results obtained under weak and strong correlation levels to those of the moderate case. Table 6 reports the average relative differences in upper bound value (UB dif.), standard deviation (STD dif.), and run time (Time dif.). The standard deviation is used as an indicator of how dispersed the demand values are within the scenarios generated for different correlation levels. A negative value in the UB dif. or Time dif. columns indicates a reduction in the respective measure for instances with either weak or strong correlation. Likewise, a negative value in the STD dif. column implies that the scenarios generated with moderate correlation are more dispersed than those generated under the other two correlation levels.

The results reported in Table 6 show notable differences across instance categories. For instance types CA and CB, which are characterized by lower demand values, there is an average increase of 4% in the upper bound value. In contrast, instance types CC and CD, which exhibit greater demand values, show an average decrease of 14% in the upper bound value. Several factors contribute to this behavior. One is the dispersion of the scenarios: achieving consensus over a more diverse set of scenarios often results in solutions with a larger number of routes or more intricate route compositions, in order to accommodate the demands of the complete scenario set without extensive use of recourse actions. Another contributing factor is the ratio of positive to negative demand correlations. A higher ratio of positive correlation leads to scenarios containing more instances of high demand values compared to those with weaker or more negative correlation levels. It is worth noting that, although the marginal probability distributions for each instance category remain constant across the three correlation types, some correlation matrices result in scenarios that are either more dispersed or contain systematically higher demand values. These variations can significantly influence the final solution when determining consensus.

In instances of types CA and CB, moderate correlations result in a lower ratio of positive correlations in 19 and 13 out of 30 cases, respectively, compared to those with weak and strong correlations. This, combined with the lower standard deviation observed in scenarios generated using moderate correlation, leads to an overall decrease in the objective function by at least 3% compared to the other two correlation levels. In contrast, instances of types CC and CD using moderate correlations produce scenarios with greater dispersion, resulting in a notable increase in the objective function (by at least 13%) compared to scenarios with weak and strong correlations. This increase is attributed to the need for a larger number of routes or more complex route compositions to fulfill all OD demands while limiting the use of outsourced deliveries. Regarding computational time, the PH method generally requires more time to solve instances with weak and strong correlations, particularly those with lower variability, due to the more frequent application of the reset procedure to escape local optima.

Overall, the level of correlation appears to significantly influence both the solution quality and computational runtime. Negative correlation offers opportunities to consolidate negatively correlated demands, often leading to reduced variance in the total demand within routes. This enables the construction of more cost-effective routes that can serve a greater number of OD demand pairs. However, increased variability in demand values can lead to more intricate route structures in order to accommodate the full range of scenarios, potentially resulting in higher operational costs. It can be concluded that such variability is not solely attributable to the correlation level, but rather to the interaction between the demand distributions generated by the copula-based method and the imposed correlation structure. This underscores the importance of accurately capturing

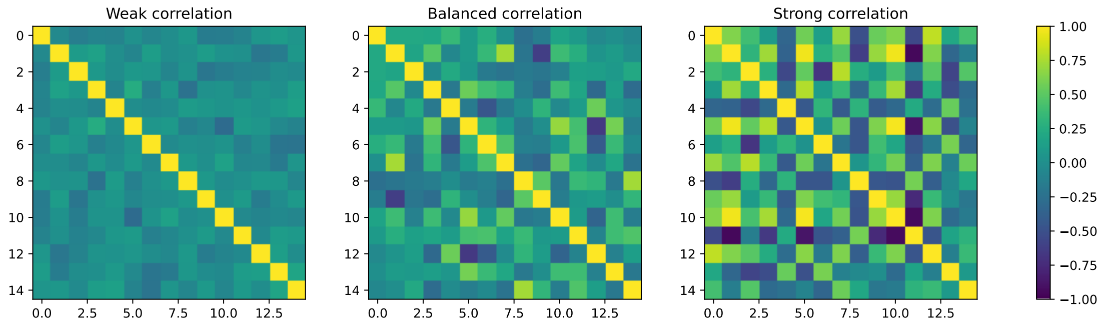

Figure 5: Correlation types for instance Cd5-6,4,15.

dependencies between OD demands when generating scenarios.

Correlation type	Distribution	Range
Weak	standard normal	[-0.3, 0.3]
Moderate	standard normal	[-0.6, 0.6]
Strong	standard normal	[-1.0, 1.0]

Table 5: Parameters for each correlation type.

Instance type	Weak correlation			Strong correlation		
	UB dif.	STD dif.	Time dif.	UB dif.	STD dif.	Time dif.
CA	5.17	1.82	2.81	6.99	2.63	18.50
CB	1.83	-24.12	7.84	2.43	-21.31	5.03
CC	-1.21	5.52	-14.41	-1.74	4.56	-7.50
CD	-26.46	-355.18	5.08	-27.07	-345.28	17.53

Table 6: Summarized results with different type of correlation types.

6.5 Value of the stochastic solution

This section reports the value of the stochastic solution (VSS), which serves as a bound to evaluate the benefit of using the stochastic model compared to the deterministic formulation for the 2E-MLRPSCD. Experiments are conducted on instances both with and without demand correlation. Following standard practice in the literature, we use a deterministic formulation (DF) based on the mean approximation of demand, where stochastic demands are replaced by their average values computed from the scenario sets. The integrated design and routing decisions are then derived from these average demand values. Results for the deterministic formulation of the 2E-MLRPSCD are obtained by solving each instance using CPLEX, with a 2-hour time limit. These results are compared to those obtained by solving the stochastic version of the 2E-MLRPSCD, as reported in Section 6.3.

Feasible solutions for all instances can be obtained using both the PH matheuristic and the deterministic formulation (DF). Accordingly, we conduct experiments comparing the DF results to those of the stochastic approach, focusing on the general percentage increase in costs. Figures 6 and 7 present the outcomes of these experiments by illustrating the percentage increase in the objective function (Cost Diff.) and the use of outsourced services (Outdiff.) for instances with and without demand correlation, respectively. The results are presented to show the minimum, average, and maximum cost increases associated with the deterministic solutions, using the stochastic approach as the baseline.

Tables 7 and 8 report density measures related to the location and allocation decisions at satellite facilities for each approach and each instance type, again distinguishing between instances with and without demand correlation. Each table includes the number of satellites ($|\mathcal{Z}|$) for each instance type and the average percentage increase in the objective function (Cost Diff.) for the DF compared to the PH matheuristic. Two additional measures are included to assess the spatial homogeneity of satellite usage and customer assignments: 1) the *satellite location density* (SLD), which represents the average number of open satellites for each instance type, and 2) the maximum and minimum *customer allocation density* (CAD), defined as the average number of customers assigned to each open satellite.

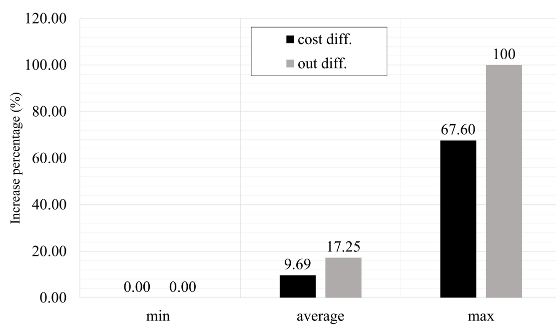

Figure 6: 2E-MLRPSCD deterministic and stochastic formulations on instances with demand correlation.

The results presented in Figure 6 and Figure 7 indicate that the stochastic formulation consistently outperforms the deterministic formulation in terms of overall solution quality. The deterministic formulation leads to significantly higher design costs and increased reliance on outsourced services. This discrepancy arises from the deterministic formulation's inability to capture key demand variations, which results in higher expected costs. The impact is particularly pronounced in design planning decisions, where the deterministic formulation often yields facility configurations that are insufficient to support the required vehicle operations across all scenarios in the second stage. This limitation is observed in both sets of instances, with the deterministic formulation relying, on average, on 17% and 6% more outsourced services for instances with and without demand correlation, respectively.

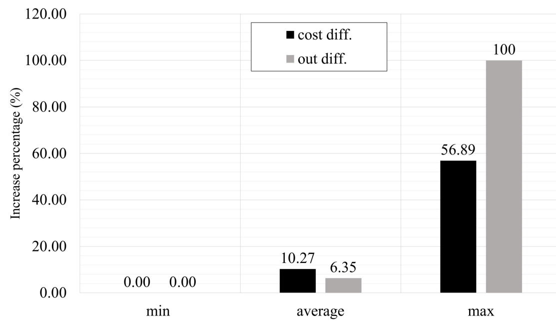

Figure 7: 2E-MLRPSCD deterministic and stochastic formulations on instances with no demand correlation.

Instance Type	$\lvert\mathcal{Z}\rvert$	Cost Diff.	PH			DF		
			SLD	CAD_{max}	CAD_{min}	SLD	CAD_{max}	CAD_{min}
CA	3	14.43	2.60	10.40	2.00	1.60	13.40	1.60
	5	47.46	2.60	5.60	3.80	1.20	12.00	3.00
	4	39.34	2.60	9.00	0.20	1.00	15.00	0.00
CB	3	0.00	1.20	12.00	3.00	1.20	12.00	3.00
	5	0.28	1.20	12.00	3.00	1.20	12.00	3.00
	4	0.00	1.00	15.00	0.00	1.00	15.00	0.00
CC	3	0.75	3.00	6.40	3.00	3.00	5.00	5.00
	5	2.07	3.60	5.40	0.40	3.20	5.00	0.40
	4	4.93	3.80	6.20	1.60	4.00	5.00	2.20
CD	3	3.53	3.00	6.20	4.00	3.00	5.00	5.00
	5	1.97	3.60	5.60	1.40	3.60	4.80	1.40
	4	4.43	3.40	6.20	1.20	4.00	5.00	2.40
Averages		*9.93*	*2.60*	*8.33*	*1.97*	*2.33*	*9.10*	*2.25*

Table 7: Location/allocation density by instance type with demand correlation

Distributions covering a wide range of low demand values, such as those in instance type CA, reveal a clear distinction between cases with and without demand correlation. When demand correlation is considered for instance type CA, there is a tendency to employ more satellite facilities to accommodate the greater diversity of scenarios. In contrast, instances without demand correlation tend to rely on fewer satellite facilities, focusing instead on addressing scenarios with high demand values, which remain relatively low compared to other instance types. For instance types CC and CD, where demand distributions yield higher values, the results generally show a high number of open satellite facilities accompanied by a more homogeneous allocation of customers. Regardless of the presence of demand correlation, high demand values drive both approaches to favor greater satellite usage in order to manage these high-value variations effectively. Interestingly, for instance type CB, which features a narrow and low-range demand distribution, the deterministic approach (based on average demand values) was able to closely approximate the actual demand distribution. As a result, the deterministic and stochastic approaches produced comparable outcomes. These observations hold even when demand correlation is not considered, with a general increase in customer allocation density across all instance types due to reduced demand variability within each scenario set.

Instance Type	$\lvert\mathcal{Z}\rvert$	Cost Diff.	PH			DF		
			SLD	CAD_{max}	CAD_{min}	SLD	CAD_{max}	CAD_{min}
CA	3	16.64	1.20	12.40	2.00	1.60	13.40	1.60
	5	37.24	1.20	13.20	1.80	1.20	12.00	3.00
	4	35.09	2.60	9.60	0.20	1.00	15.00	0.00
CB	3	0.06	1.20	13.20	3.00	1.20	13.20	3.00
	5	0.00	1.20	13.20	3.00	1.00	13.20	3.00
	4	0.05	1.00	15.00	0.00	1.00	15.00	0.00
CC	3	7.15	3.00	7.80	3.00	3.00	5.00	5.00
	5	5.42	2.80	8.60	2.20	3.20	5.00	0.40
	4	9.32	2.80	6.80	2.60	4.00	5.00	2.20
CD	3	0.39	3.00	8.60	3.00	3.00	5.00	5.00
	5	4.46	3.20	6.40	3.20	3.60	4.80	1.40
	4	7.48	3.20	7.00	2.00	4.00	5.00	2.40
Averages		*10.27*	*2.20*	*10.15*	*2.17*	*2.32*	*9.30*	*2.25*

Table 8: Location/allocation density by instance type without demand correlation

One can conclude that the stochastic approach, as addressed by the proposed PH matheuristic, is generally more cost-effective for both design and routing decisions. In contrast, the deterministic formulation tends to produce solutions that lack operational efficiency, particularly in the second stage, due to its insufficient consideration of uncertainty during the design planning phase. Unless the demand distribution is both narrow and low in magnitude, the deterministic formulation proves unsuitable for designing distribution networks under demand uncertainty, regardless of whether correlation is present. Therefore, a stochastic approach should be favored to ensure effective distribution system design involving location-routing decisions under uncertainty.

7 Conclusions

We introduced the two-Echelon Multicommodity Location-Routing Problem with Stochastic and Correlated Demands (2E-MLRPSCD). The problem is formulated as a two-stage stochastic program: the first stage determines the location of satellite facilities and the allocation of customers to satellites, while the second stage defines the vehicle routes for both echelons once customer demands are revealed. To solve the proposed model, we develop a specialized Progressive Hedging (PH)-based matheuristic incorporating several novel enhancements. These include: (1) population structures that maintain alternative and diverse solutions for the scenario subproblems; (2) strategies for defining the reference solutions that guide the overall search process; and (3) a reset procedure that mitigates the risk of convergence to local optima.

A series of numerical experiments were performed on a set of instances with varying characteristics, which computationally demonstrated that the proposed enhancements significantly improve the overall performance of the PH method developed for the 2E-MLRPSCD. Moreover, the numerical results clearly highlight the added value of explicitly accounting for demand uncertainty and its interrelations. The solutions obtained by solving the stochastic problem consistently outperformed those derived from a deterministic approximation approach.

Several interesting avenues for future research may be identified. First, there is a need to design novel heuristic and exact methods to more efficiently solve the set of scenario subproblems at each iteration of the PH matheuristic. In addition, various extensions of the problem could be explored. For instance, incorporating additional sources of uncertainty (such as uncertain travel times) would be a valuable direction with broad applicability across different domains.

Acknowledgments

While working on the project, the second author held the UQAM Chair in Intelligent Logistics and Transportation Systems Planning and was Adjunct Professor, Department of Computer Science and Operations Research, Université de Montréal, Canada, while the third author held the Canada Research Chair in Stochastic Optimization of Transport and Logistics Systems. We gratefully acknowledge the financial support provided by the Canadian Natural Sciences and Engineering Research Council (NSERC) through its Discovery and Collaborative Research & Development grant programs, as well as by the Fonds de recherche du Québec through their Teams and CIRRELT infrastructure grants.

References

M. Albareda-Sambola, E. Fernandez, and G. Laporte. Heuristic and lower bound for a stochastic location-routing problem. *European Journal of Operational Research*, 179 (3):940 – 955, 2007.

A. Alvarez, J.-F. Cordeau, R. Jans, P. Munari, and R. Morabito. Inventory routing under stochastic supply and demand. *Omega*, 102:102304, 2021.

S. Atakan and S. Sen. A progressive hedging based branch-and-bound algorithm for mixed-integer stochastic programs. *Computational Management Science*, 15(3):501–540, Oct 2018.

M. J. Bucci, D. P. Warsing, M. G. Kay, R. Uzsoy, and J. R. Wilson. Modeling the Effects of Demand Correlation in Location Problems that account for Inventory Pooling, 2006.

T. G. Crainic, X. Fu, M. Gendreau, W. Rei, and S. W. Wallace. Progressive hedging-based metaheuristics for stochastic network design. *Networks*, 58(2):114–124, 2011.

R. Cuda, G. Guastaroba, and M. G. Speranza. A survey on two-echelon routing problems. *Computers and Operations Research*, 55:185–199, 2015.

N. Dellaert, F. Dashty Saridarq, T. Van Woensel, and T. G. Crainic. Branch-and-price–based algorithms for the two-echelon vehicle routing problem with time windows. *Transportation Science*, 53(2):463–479, 2019.

D. Escobar-Vargas and T. G. Crainic. Multi-attribute Two-echelon Location Routing: Formulation and Dynamic Discretization Discovery Approach. *European Journal of Operational Research*, 314(1):66–78, 2024.

M. Gendreau, O. Jabali, and W. Rei. *Chapter 8: Stochastic Vehicle Routing Problems*, pages 213–239. Society for Industrial and Applied Mathematics, 2014.

Z. Guo, S. W. Wallace, and M. Kaut. Vehicle routing with space- and time-correlated stochastic travel times: Evaluating the objective function. *INFORMS Journal on Computing*, 31(4):654–670, 2019.

K. K. Haugen, A. Løkketangen, and D. L. Woodruff. Progressive hedging as a meta-heuristic applied to stochastic lot-sizing. *European Journal of Operational Research*, 132(1):116–122, 2001.

D. C. Heath and P. L. Jackson. Modeling the evolution of demand forecast with application to safety stock analysis in production/distribution systems. *IIE Transactions*, 26(3):17–30, 1994.

M. Hewitt, J. Ortmann, and W. Rei. Decision-based scenario clustering for decision-making under uncertainty. *Annals of Operations Research*, 315(2):747–771, Aug 2022.

S.-H. Huang. Solving the multi-compartment capacitated location routing problem with pickup–delivery routes and stochastic demands. *Computers & Industrial Engineering*, 87:104 – 113, 2015.

M. Kaut. A copula-based heuristic for scenario generation. *Computational Management Science*, 11(4):503–516, Oct 2014.

M. Kaut and S. W. Wallace. Evaluation of scenario-generation methods for stochastic programming. *Pacific Journal of Optimization*, 3(2):257–271, 2007.

A. King and S. Wallace. *Modeling with Stochastic Programming*. Springer New York, NY, 01 2012.

A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502, 2002. doi: 10.1137/S1052623499363220. URL `https://doi.org/10.1137/S1052623499363220`.

A. Lamghari and R. Dimitrakopoulos. Progressive hedging applied as a metaheuristic to schedule production in open-pit mines accounting for reserve uncertainty. *European Journal of Operational Research*, 253(3):843–855, 2016.

A.-G. Lium, T. G. Crainic, and S. W. Wallace. A study of demand stochasticity in service network design. *Transportation Science*, 43(2):144–157, 2009.

A. Løkketangen and D. L. Woodruff. Progressive hedging and tabu search applied to mixed integer (0,1) multistage stochastic programming. *Journal of Heuristics*, 2(2): 111–128, Sep 1996.

S. T. W. Mara, R. Kuo, and A. M. S. Asih. Location-routing problem: a classification of recent research. *International Transactions in Operational Research*, 28(6):2941–2983, 2021.

F. E. Maranzana. On the location of supply points to minimize transport costs. *OR*, 15 (3):261–270, 2023/09/13/ 1964. Full publication date: Sep., 1964.

Y. Marinakis. An improved particle swarm optimization algorithm for the capacitated location routing problem and for the location routing problem with stochastic demands. *Applied Soft Computing Journal*, 37:680–701, 2015.

Y. Marinakis, M. Marinaki, and A. Migdalas. A hybrid clonal selection algorithm for the location routing problem with stochastic demands. *Annals of Mathematics and Artificial Intelligence*, 76(1):121–142, Feb 2016.

S. M. Mirhedayatian, T. G. Crainic, M. Guajardo, and S. W. Wallace. A two-echelon location-routing problem with synchronisation. *Journal of the Operational Research Society*, 72(1):145–160, 2019.

C. L. Quintero-Araujo, D. Guimarans, and A. A. Juan. A simheuristic algorithm for the capacitated location routing problem with stochastic demands stochastic demands. *Journal of Simulation*, 15(3):1–18, 2019.

R. T. Rockafellar and R. J.-B. Wets. Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research*, 16(1):119–147, 1991.

M. Schiffer, M. Schneider, G. Walther, and G. Laporte. Vehicle routing and location routing with intermediate stops: A review. *Transportation Science*, 53(2):319–343, 2019.

A. Snoeck, M. Winkenbach, and E. E. Mascarino. Establishing a Robust Urban Logistics Network at FEMSA through Stochastic Multi-Echelon Location Routing. In E. Taniguchi and R. Thompson, editors, *City Logistics 2*. John Wiley & Sons, Ltd, 2018.

B. K. Thapalia, S. W. Wallace, M. Kaut, and T. G. Crainic. Single source single-commodity stochastic network design. *Computational Management Science*, 9(1):139–160, Feb 2012.

A. Verma and A. M. Campbell. Strategic placement of telemetry units considering customer usage correlation. *EURO Journal on Transportation and Logistics*, 8(1):35–64, 2019.

T. Vidal, T. G. Crainic, M. Gendreau, N. Lahrichi, and W. Rei. A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research*, 60(3):611–624, 2012.

K. Xu and P. T. Evers. Managing single echelon inventories through demand aggregation and the feasibility of a correlation matrix. *Computers & Operations Research*, 30(2): 297 – 308, 2003.

D. Zhang, S. W. Wallace, Z. Guo, Y. Dong, and M. Kaut. On scenario construction for stochastic shortest path problems in real road networks. *Transportation Research Part E: Logistics and Transportation Review*, 152:102410, 2021.

S. Zhang, M. Chen, and W. Zhang. A novel location-routing problem in electric vehicle transportation with stochastic demands. *Journal of Cleaner Production*, 221:567–581, 2019.

APPENDIX

A - Decomposition strategy for the two-stage stochastic formulation

This section presents the complete set of steps used to perform the decomposition approach applied to the stochastic 2E-MLRPSCD formulation introduced in Section 4.2. This approach is based on an augmented Lagrangian strategy.

The decomposition strategy applied to the scenario-based formulation across the scenarios in $\mathcal{S}$ requires a reformulation of the first-stage decisions. Specifically, these decisions must now be defined as scenario-dependent. As a result, Constraints (8), (19), and (20) are reexpressed using scenario-specific first-stage location and allocation decisions. Accordingly, we begin with the following alternative, yet equivalent, formulation:

$$\min \sum_{s \in \mathcal{S}} \rho_s \left(\sum_{h \in \mathcal{H}} \sum_{(i,j) \in A} \zeta_{ij} x^s_{ijh} + \sum_{k \in \mathcal{K}} Ro^s_k + \sum_{z \in \mathcal{Z}} F_z y^s_z + \sum_{z \in \mathcal{Z}} \sum_{k \in \mathcal{K}} \Delta_{zk} f^s_{zk} \right) \tag{61}$$

Subject to

$$(2) - (7)$$
$$(9) - (18)$$
$$(24) - (31)$$

$$\sum_{h \in \mathcal{H}^2} \sum_{j \in C} x^s_{ijh} \leq |H^2| y^s_i \quad \forall i \in \mathcal{Z}, \forall s \in \mathcal{S} \tag{62}$$

$$\sum_{h \in \mathcal{H}^1} u^s_{iO(k)h} = f^s_{iO(k)} \quad \forall i \in \mathcal{Z}, \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \tag{63}$$

$$\sum_{h \in \mathcal{H}^2} v^s_{zD(k)h} = f^s_{zO(k)} \quad \forall z \in \mathcal{Z}, \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \tag{64}$$

$$y^s_z = \bar{y}_z \quad \forall z \in \mathcal{Z}, \forall s \in \mathcal{S} \tag{65}$$

$$f^s_{zk} = \bar{f}_{zk} \quad \forall z \in \mathcal{Z}, \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \tag{66}$$

$$y^s_z \in \{0,1\} \quad \forall z \in \mathcal{Z}, \forall s \in \mathcal{S} \tag{67}$$

$$f^s_{zk} \in \{0,1\} \quad \forall z \in \mathcal{Z}, \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \tag{68}$$

$$\bar{y}_z \in \{0,1\} \quad \forall z \in \mathcal{Z} \tag{69}$$

$$\bar{f}_{zk} \in \{0,1\} \quad \forall z \in \mathcal{Z}, \forall k \in \mathcal{K} \tag{70}$$

This reformulation now explicitly includes the set of non-anticipativity constraints, which prevent the first-stage decision variables from taking different scenario-specific values (i.e., the first-stage decisions must be implementable). Constraints (62)–(64) link the facility allocation variables with the vehicle allocation variables. Constraints (65) and (66) enforce that the first-stage solutions remain identical across all scenarios (i.e., they impose non-anticipativity), where the variables $\bar{y}_i$ and $\bar{f}_{ijk}$ serve as the reference first-stage decisions. These constraints ensure that a single set of facility location and allocation decisions is made across all scenarios, thereby preventing the implementation of scenario-specific first-stage decisions. Following the decomposition scheme originally proposed by (Rockafellar and Wets, 1991), Constraints (65) and (66) are relaxed using an augmented Lagrangian method, which results in the following objective function:

$$\begin{aligned} \min \sum_{s\in\mathcal{S}} \rho_s \Bigg(& \sum_{h\in\mathcal{H}^1} \sum_{(i,j)\in A^1} \zeta_{ij} x^s_{ijh} + \sum_{h\in\mathcal{H}^2} \sum_{(i,j)\in A^2} \zeta_{ij} x^s_{ijh} + \sum_{k\in\mathcal{K}} Ro^s_k + \sum_{z\in\mathcal{Z}} F_z y^s_z \\ & + \sum_{z\in\mathcal{Z}} \lambda^s_z (y^s_z - \bar{y}_z) + \frac{1}{2} \sum_{z\in\mathcal{Z}} \gamma (y^s_z - \bar{y}_z)^2 + \sum_{z\in\mathcal{Z}} \sum_{k\in\mathcal{K}} \Delta_{zk} f^s_{zk} \\ & + \sum_{z\in\mathcal{Z}} \sum_{k\in\mathcal{K}} \mu^s_{zk} (f^s_{zk} - \bar{f}_{zk}) + \frac{1}{2} \sum_{z\in\mathcal{Z}} \sum_{k\in\mathcal{K}} \gamma (f^s_{zk} - \bar{f}_{zk})^2 \Bigg) \end{aligned} \tag{71}$$

The objective function now includes the Lagrangian multipliers λ^s_i and μ^s_{ijk} associated with the relaxed Constraints (65) and (66), respectively, along with a penalty parameter γ. This penalty parameter applies to a quadratic term that further penalizes deviations between the scenario-specific first-stage decision variables and their corresponding reference values. Given the binary nature of the location and allocation variables, the objective function can be simplified as follows:

$$\begin{aligned} \min \sum_{s\in\mathcal{S}} \rho_s \Bigg(& \sum_{h\in\mathcal{H}} \sum_{(i,j)\in A} \zeta_{ij} x^s_{ijh} + \sum_{k\in\mathcal{K}} Ro^s_k + \sum_{z\in\mathcal{Z}} \left(F_z + \lambda^s_z + \frac{1}{2}\gamma + \gamma \bar{y}_z \right) y^s_z \\ & + \sum_{z\in\mathcal{Z}} \sum_{k\in\mathcal{K}} \left(\Delta_{zk} + \mu^s_{zk} + \frac{1}{2}\gamma + \gamma \bar{f}_{zk} \right) f^s_{zk} + \frac{1}{2} \sum_{z\in\mathcal{Z}} \gamma \bar{y}_z - \sum_{z\in\mathcal{Z}} \lambda^s_z \bar{y}_z \\ & + \frac{1}{2} \sum_{z\in\mathcal{Z}} \sum_{k\in\mathcal{K}} \gamma \bar{f}_{zk} + \sum_{z\in\mathcal{Z}} \sum_{k\in\mathcal{K}} \mu^s_{zk} \bar{f}_{zk} \Bigg) \end{aligned} \tag{72}$$

Given the objective function (72) and the constraint set: (2)-(7), (9)-(18), (24)-(31), (62)-(64), and (67)-(70), if the reference point (or solution) $\bar{y}_i$ and $\bar{f}_{ijk}$ is fixed, then the relaxed formulation can be decomposed by scenario. Specifically, for each $s \in \mathcal{S}$, a deterministic 2E-MLRP subproblem with modified fixed costs is obtained:

$$\min \sum_{s \in \mathcal{S}} \rho_s \Bigg(\sum_{h \in \mathcal{H}} \sum_{(i,j) \in A} \zeta_{ij} x^s_{ijh} + \sum_{k \in \mathcal{K}} Ro^s_k + \sum_{z \in \mathcal{Z}} \left(F_z + \lambda^s_z + \frac{1}{2}\gamma + \gamma \bar{y}_z \right) y^s_z + \sum_{z \in \mathcal{Z}} \sum_{k \in \mathcal{K}} \left(\Delta_{zk} + \mu^s_{zk} + \frac{1}{2}\gamma + \gamma \bar{f}_{zk} \right) f^s_{zk} \Bigg) \tag{73}$$

Subject to

$$(2) - (7), (9) - (18), (24) - (31), (62) - (64) \text{ and } (67) - (70)\ .$$

As previously stated, the proposed PH algorithm proceeds by solving the scenario subproblems separately, thereby obtaining scenario-specific first-stage solutions. These scenario-specific solutions are then used to compute the updated reference point. Using this reference point, the objective functions (73), for all $s \in \mathcal{S}$, are modified to incentivize agreement among the subproblem decisions (i.e., consensus). This general process is then repeated iteratively until a consensus first-stage solution is found.