

27th Euro Working Group on Transportation Meeting (EWGT 2026)

Innovative Logistics of Spent Coffee Grounds Reuse: An IoT-Based Rolling Horizon Approach

Pasquale Carotenuto^a, Alessia Durante^b, Stefano Giordani^b, Alessio Salvatore^{b,c,*}

^a *Istituto per le Applicazioni del Calcolo “M. Picone” – Consiglio Nazionale delle Ricerche, via dei Taurini 19, Roma – 00185, Italy*

^b *Università di Roma “Tor Vergata” – Dipartimento di Ingegneria dell’Impresa, via del Politecnico 1, Roma – 00133, Italy*

^c *Istituto di Analisi dei Sistemi ed Informatica “Antonio Ruberti” – Consiglio Nazionale delle Ricerche, via dei Taurini 19, Roma – 00185, Italy*

Abstract

Coffee is one of the most consumed beverages worldwide, driving an industry worth billions of euros annually. However, its life cycle does not end with consumption; it generates a significant amount of by-products, including spent coffee grounds (SCG), often treated as waste. If managed through innovative recovery and reuse practices, these by-products can be recycled and transformed into valuable resources, reducing environmental impact, and optimizing logistics processes. This research explores the recovery of SCG as secondary raw materials in other production chains, leveraging IoT technologies and web-based platforms for real-time monitoring and informed decision-making. The study proposes a heuristic approach that simultaneously plans coffee deliveries and SCG collections (pick-ups) and related routing for a set of coffee shops, while dynamically adapting to evolving requests through a Rolling Horizon technique. Performance results on test cases based on extensive territorial and realistic consumption data from a selected area within Milan are promising, highlighting the potential of this approach for sustainable supply chain management.

© 2026 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

Peer-review under responsibility of the scientific committee of the Euro Working Group on Transportation Annual Meeting 2025 - EWGT2025.

Keywords: Sustainability; Coffee supply chain; Reuse and recycling; Inventory routing problem; Rolling horizon heuristics.

1. Introduction

In today’s logistics systems, many supply chains still operate as closed networks focused on the movement of a single product. This often results in inefficiencies such as empty vehicle trips, suboptimal routing, and longer delivery time, as well as negative environmental effects due to increased congestion and emissions. These issues raise operational costs and hinder sustainability. In contrast, circular economy models aim to reduce waste and improve resource efficiency by promoting reuse and better data exchange among stakeholders. However, enabling circularity in practice requires suitable infrastructure for real-time data sharing and coordination.

Technological innovations such as the Internet of Things (IoT) and real-time data analytics are playing a key role in transforming logistics operations, particularly in sectors like waste collection, recycling, and reuse [Kannan

* Corresponding author. email: alessio.salvatore@cnr.it

2352-1465 © 2026 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

Peer-review under responsibility of the scientific committee of the Euro Working Group on Transportation Annual Meeting 2025 - EWGT2025.

et al., 2023]. Optimizing these processes can lead to substantial environmental and economic benefits [Massaro Sousa and Ferreira, 2019]. A notable example is the recovery of spent coffee grounds (SCG), which can be repurposed in industries such as biofuel production, cosmetics, and agriculture [Forcina et al., 2023]. IoT platforms enable real-time monitoring of SCG accumulation in smart bins, allowing for more efficient collection, treatment, and delivery for reuse. At the same time, high and regular demand for coffee requires timely replenishment at coffee shops to ensure service continuity. Efficient management of both flows (delivery of coffee and pick-up of SCG) is therefore essential.

To support timely and efficient decision-making in such dynamic settings, with stochastic demands, Rolling Horizon (RH) heuristics have proven effective. These methods divide the planning horizon into smaller time windows and optimize sequentially, updating decisions as new information becomes available. RH approaches have been applied in various logistics domains. For example Bostel et al. [2008] address the problem of scheduling and routing field visits by technicians to customers, for maintenance or logistics activities. The schedule was designed on a continuous multiperiod horizon and updated daily; the procedures were then adapted to handle a continuous horizon, when a new plan had to be determined after the execution of each first daily period of the previous plan. Cordeau et al. [2015] introduce a rolling horizon algorithm to schedule vehicle deliveries to car dealerships by a heterogeneous fleet of truck drivers. The problem consists of scheduling deliveries over a multi-day planning horizon during which transportation requests arrive dynamically. The objective is to minimize the sum of the distances traveled, the trucker's fixed costs, the service costs, and the late delivery penalties. Finally, Kim et al. [2023] use the RH approach for solving a class of pick-up and delivery problems (PDPs) that have narrow time windows and compare their results with an offline heuristic algorithm using Google OR-Tools.

In the waste management context, Spinelli et al. [2024] apply a Rolling Horizon heuristic within a multi-stage stochastic optimization model to solve an inventory routing problem for municipal waste collection. Their goal is to maximize the total expected profit by selecting which bins to serve and generating efficient routing plans over a finite horizon. Although focused on a single product and pick-up operations only, their framework highlights the potential of RH-based approaches in dynamic and uncertain waste management settings, similar in nature to our case study, which involves two interdependent product flows (coffee delivery and SCG pick-up) and the need to coordinate both services over time.

In our study, the primary goal of the methodology is to optimize both coffee delivery and collection of spent coffee grounds (SCG), a crucial yet often scarcely investigated aspect of the coffee supply chain. These logistics operations are dynamic, as they are subject to changing demand patterns and operational constraints. As a result, an adaptive integrated approach is required to effectively plan coffee deliveries and SCG pick-up and the related vehicle routing.

In operational terms, we suppose to act as Vendor Managed Inventory [Waller et al., 1999] for each coffee shop in the coffee replenishment process, and as a third-party logistics partner for companies that recycle spent coffee grounds (SCG). This dual role gives us access to key input data: the coffee shop demand, the current inventory levels at each coffee shop, and the amount of SCG gradually accumulated in smart bins, monitored in real time through IoT technology. The inventory management is modeled with an integrated dynamic lot sizing model with lost sales (see, e.g., Aksen et al. [2003]) for planning both coffee deliveries and SCG pick-ups for a set of coffee shops. To address these challenges, this study considers an integrated dynamic inventory and routing problem under stochastic demand. We apply a dynamic inventory control policy for both products, using a Rolling Horizon (RH) heuristic, which offers an effective and flexible framework for decision-making over time. To further optimize both the inventory management and the routing phase, we also introduce a Consolidated Inventory-Routing Policy (CIRP), aimed at combining delivery and pick-up operations to reduce the number of visits over the planning period and the related routing costs.

The remainder of the paper is organized as follows. In Section 2, we provide a detailed description of the problem and the basic integrated inventory model. In Section 3, we describe the adopted methodology, presenting the details of the rolling horizon approach, the inventory control phase and the related adopted inventory policies, and the description of the routing phase. In Section 4 we discuss the experimentation results, and finally, in Section 5 we draw some conclusion.

2. Problem and model description

The optimization problem can be modeled as a two-product dynamic inventory routing problem (DIRP) with stochastic demands, where product 1 (coffee) has to be delivered to coffee shops and product 2 (SCG) has to be picked-up from them. The planning horizon covers a set T of $|T|$ time periods (days). A fleet of K homogeneous vehicles is used to simultaneously deliver coffee to and pick-up SCG from coffee shops. Let C^1 and C^2 be the capacity of each vehicle for the delivery of coffee and the pick-up of SCG, respectively. We also consider a maximum number of working hours per vehicle per day, denoted by H . This constraint implicitly limits the maximum number of trips (visits) that each vehicle can perform daily.

We consider a complete graph $G = (V, E)$, where $V = \{0, 1, \dots, n, n + 1\}$ is the set of vertices, with $N = \{1, \dots, n\}$ being the clients (coffee shops) sites and 0 and $n + 1$ the locations of start and end depots: we differentiate these two depots also because in general product 1 (coffee) is delivered to clients from a (start) depot, while product 2 (SCG) is picked-up from clients and transported to a possible distinct (end) depot. Vehicles, capable of transporting both the coffee to be delivered and the SCG to be picked-up, are assumed to be available at the start depot. The edge set E represents the links between pairs of locations, and a travel cost c_{ij} is associated to each edge $(i, j) \in E$.

The product 1 (coffee) is consumed in coffee shop i according to daily demands $d_{i,t}^1$ for each time period (day) $t \in T$ that are not known in advance, producing at the same time the product 2 (SCG) as a waste material that can be possibly recycled. Both products are stored in dedicated containers in a coffee shop i . For both products $s = 1, 2$, (on-hand) inventories must be maintained between specific lower and upper limits defined by the interval $[\underline{U}_i^s, \bar{U}_i^s]$, otherwise a lost penalty applies.

If (net) inventory I_{it}^1 of product 1 (coffee) at the end of period t for coffee shop i is less than the lower limit \underline{U}_i^1 (e.g., $\underline{U}_i^1 = 0$), then in that period there will be a coffee shortage (loss of demand) amount $\ell_{it}^1 = \max\{\underline{U}_i^1 - I_{it}^1, 0\}$. Let I_{it}^{+1} (with $0 \leq \underline{U}_i^1 \leq I_{it}^{+1} \leq \bar{U}_i^1$) be the inventory on-hand at the end of period t : it results $I_{it}^1 = I_{it}^{+1} - \ell_{it}^1$.

Similarly, if (net) inventory I_{it}^2 of product 2 (SCG) in the end of period t is greater than the upper limit \bar{U}_i^2 , then in that period there will be an excess amount $\ell_{it}^2 = \max\{I_{it}^2 - \bar{U}_i^2, 0\}$ of SCG that cannot be stored and, hence, will be lost for recycling (discarded in another way). Let I_{it}^{+2} (with $0 \leq \underline{U}_i^2 \leq I_{it}^{+2} \leq \bar{U}_i^2$) be the inventory on-hand at the end of period t : it results $I_{it}^2 = I_{it}^{+2} + \ell_{it}^2$. Moreover, for product 2 (SCG) it is assumed $\underline{U}_i^2 = 0$.

In each time period $t \in T$, each coffee shop $i \in N$ has to satisfy a coffee demand d_{it}^1 (which is not known in advance), decreasing its initial coffee inventory $I_{i,t-1}^1$, as well as producing a SCG amount $d_{i,t}^2 = \gamma(d_{i,t-1}^1 - \ell_{i,t-1}^1)$ that in period t will increase its initial SCG inventory $I_{i,t}^2$, where $\gamma > 1$ is a multiplier factor representing the ratio between the amount of SCG produced and that of coffee consumed. Clearly, we assume that for each coffee shop i , the stocking capacities for the two products are sufficiently large to stock at least what is demanded (produced) each time period, that is, $0 \leq \max_i\{d_{it}^1\} \leq \bar{U}_i^1 - \underline{U}_i^1$, and $0 \leq \gamma \max_i\{d_{it}^1\} \leq \bar{U}_i^2 - \underline{U}_i^2$.

Each coffee shop i can be visited by at most one vehicle per period $t \in T$, and let $q_{it}^1 \geq 0$ and $q_{it}^2 \geq 0$ be the amounts of coffee delivered and SCG picked-up, respectively. Without loss of generality, we assume that coffee delivery and SCG pick-up can also be done at the same time t without significant additional cost, with coffee delivery performed before SCG pick-up.

Denoting with $\bar{d}_{i,t}^1$ the forecast of coffee demand in period t of coffee shop i , for the two inventory levels of coffee shop i we have the following relations for $t \in T$, where I_{i0}^{+1} and I_{i0}^{+2} are the initial inventory on-hand levels for coffee and SCG, respectively, and $\gamma(\bar{d}_{i,0}^1 - \ell_{i,0}^1) = \gamma(d_{i,0}^1 - \ell_{i,0}^1)$ is the initial SCG produced and stockpiled:

$$I_{i,t-1}^{+1} + q_{it}^1 = I_{it}^{+1} + (\bar{d}_{i,t}^1 - \ell_{it}^1), \text{ with } 0 \leq \underline{U}_i^1 \leq I_{it}^{+1} \leq \bar{U}_i^1 \text{ and } 0 \leq \bar{d}_{i,t}^1 - \ell_{it}^1,$$

$$I_{i,t-1}^{+2} + \gamma(\bar{d}_{i,t-1}^1 - \ell_{i,t-1}^1) = I_{it}^{+2} + \ell_{it}^2 + q_{it}^2 \text{ with } 0 \leq \underline{U}_i^2 \leq I_{it}^{+2} \leq \bar{U}_i^2 \text{ and } \underline{U}_i^2 \leq (I_{i,t-1}^{+2} + \gamma(\bar{d}_{i,t-1}^1 - \ell_{i,t-1}^1) - \ell_{it}^2) \leq \bar{U}_i^2.$$

Inventory decisions define the amount q_{it}^1 of product 1 (coffee) to be delivered and the amount q_{it}^2 of product 2 (SCG) to be picked up for each day $t \in T$ in each coffee shop i by the fleet of vehicles, each day starting at the (start) depot 0 and ending at the (end) depot $n + 1$. Inventory and vehicle routing decisions have to be made to meet coffee demands, while minimizing the average inventory, the loss of coffee demand and of SCG to recycle, as well as the total routing cost.

3. Methodology

We solve the problem with a rolling horizon heuristic approach, where at each decision round we first decide the pick-up and delivery amounts for each coffee shop and for each time period and then we solve the pick-up and delivery vehicle routing problem for each frozen time period. In the following, we present the rolling horizon scheme, the adopted inventory policy for the first (inventory control) phase and provide some details for the second (routing) phase.

3.1. Rolling horizon scheme

The daily coffee demands faced by the coffee shops are not known in advance, and a rolling horizon scheme is adopted to deal with their stochasticity.

Consequently, decisions are assumed to be made every f time periods, that is, at times $\theta = 1, 1 + f, \dots, 1 + (r - 1)f$, with $r = \lceil \frac{|T|}{f} \rceil$. Each planning stage $j = 1, \dots, r$ considers the subset $L_j \subseteq T$ of $|L_j| = \min\{\lambda, |T| - \theta_j + 1\}$ time periods (with $\lambda \geq f$), starting from the current time period $\theta_j = 1 + (j - 1)f$, i.e., $L_j = \{\theta_j, \dots, \theta_j + |L_j| - 1\}$, and then assuming frozen the decisions on the first $|F_j| = \min\{f, |L_j|\}$ time periods of L_j , i.e., $F_j = \{\theta_j, \dots, \theta_j + |F_j| - 1\}$.

In this setting, when programming coffee deliveries and SCG pick-ups at time θ_j for the next $|L_j|$ periods starting from θ_j , we assume to have an updated forecast of coffee demands $\tilde{d}_{i,t}^1$ of each coffee shop i , for $t = \theta_j, \dots, \theta_j + |L_j| - 1$, affected by an error-prone forecast. The farther ahead in time, the higher the forecast error. Demand forecasts are obtained following the ideas of Clark [2005], assuming that the accuracy of the forecast improves as we approach the actual demand event, because new information becomes available to support the forecast method.

Therefore, for the j^{th} planning round involving time periods $t = \theta_j, \dots, \theta_j + |L_j| - 1$, given the updated forecast demand for the demand and the actual inventory on-hand at the end of time period $\theta_j - 1$, we determine for each coffee shop i the amount $q_{i,t}^1$ of coffee delivery and the amount $q_{i,t}^2$ of SCG pick-up, according to the proposed inventory policy, freezing the decisions for the first $|F_j| \leq |L_j|$ time periods. These decisions provide the pick-up and delivery transportation request for each period $t = \theta_j, \dots, \theta_j + |F_j| - 1$, for each coffee shop i , as input for the routing decision phase.

3.2. Inventory control phase

For product 1 (coffee) and for each coffee shop i , we apply a dynamic inventory control policy $(r_{i,t}^1, S_{i,t}^1)$ where the coffee inventory on-hand is reviewed every period t and if not greater than the reorder point $r_{i,t}^1$, a coffee delivery is performed to raise the inventory on-hand to the order-up-to level $S_{i,t}^1 > r_{i,t}^1$ (see, e.g., [Babai and Dallery, 2009] and [Coelho et al., 2014]).

The values of $r_{i,t}^1$ and $S_{i,t}^1$ should be evaluated to minimize the total expected inventory cost (including the loss demand cost) and the number of deliveries. In particular, assuming that inventory control decisions are taken in the first period (i.e., at time period 1), the reorder point evaluated for period t is assumed to be $r_{i,t}^1 = \tilde{d}_{i,t}^1(1 + pt)$, that is, equal to the forecast demand of period t multiplied by the factor $(1 + pt)$ to consider forecast uncertainty, where $p = 0.02$ is the uncertainty degree of the forecast demand per period. The order-up-to level is assumed to be $S_{i,t}^1 = \sum_{\theta=1}^{\tau_i} \tilde{d}_{i,t+\theta-1}^1(1 + p(t + \theta - 1)) > r_{i,t}^1$, that is, equal to the forecast demand of the next τ_i periods multiplied by a factor to take into account the uncertain degree, where τ_i is the expected delivery cycle time (it may depend on the type of the coffee shop: e.g., if it has mostly weekdays or weekly holiday demands). Finally, the inventory on-hand lower limit is assumed $\underline{U}_i^1 = 0$ and the upper limit \bar{U}_i^1 sufficiently large so as $0 = \underline{U}_i^1 < r_{i,t}^1 < S_{i,t}^1 \leq \bar{U}_i^1$.

Therefore, the delivery of coffee to shop i in time period t is carried out if $I_{i,t-1}^1 \leq r_{i,t}^1$ (i.e., if the initial level of the inventory on-hand is low), with delivery amount $q_{i,t}^1 = S_{i,t}^1 - I_{i,t-1}^1$.

Coffee shop i produces product 2 (SGC) stocked into dedicated bins of total capacity \bar{U}_i^2 , that should be collected from time to time to maintain the related on-hand inventory level within the given range $[\underline{U}_i^2, \bar{U}_i^2]$, where the lower limit is assumed to be $\underline{U}_i^2 = 0$. Therefore, unlike coffee delivery, SGC pick-up should be done from a coffee shop i at time t to avoid the loss of recyclable SGC which will be $\ell_{i,t}^2 = \max\{I_{i,t-1}^2 + \gamma(\tilde{d}_{i,t-1}^2 - \ell_{i,t-1}^2) - \bar{U}_i^2, 0\}$, and hence it should be done if $I_{i,t-1}^2 + \gamma(\tilde{d}_{i,t-1}^2 - \ell_{i,t-1}^2) - \ell_{i,t}^2 \geq r_{i,t}^2$, where $r_{i,t}^2$ is the SGC recollecting point with $0 = \underline{U}_i^2 < r_{i,t}^2 < \bar{U}_i^2$. The value of $r_{i,t}^2$ should be evaluated to minimize the number of SCG pick-ups but at the same time to ensure that

SCG are collected with a sufficient frequency since SCG is a perishable product; moreover, the value of $r_{i,t}^2$ should be sufficiently less than the upper limit \bar{U}_i^2 to minimize the amount of loss of SCG. In particular, we assume that $r_{i,t}^2 = \max\{\bar{U}_i^2 - \bar{d}_i^2(1 + pt), 0\}$, where $\bar{d}_{i,t}^2$ is the forecast amount of SCG produced in period t , equal to $\bar{d}_{i,t}^1 - \ell_{i,t}^1$. Finally, due to the perishable nature of SCG, we assume to empty the recycle bins where SCG are stored (assuming $\underline{U}_i^2 = 0$) whenever a pick-up is performed; hence, the pick-up SCG amount is $q_{i,t}^2 = I_{i,t-1}^{+2} + \gamma(\bar{d}_{i,t-1}^1 - \ell_{i,t-1}^1) - \ell_{i,t}^2$.

In addition to these inventory-based policies, a synchronization mechanism between coffee deliveries and SCG pick-ups is introduced to reduce the number of visits, which indirectly will lead to a reduction in the routing costs. This mechanism is opportunistic in nature, referred here as the *Consolidated Inventory-Routing Policy (CIRP)*. Synchronization is considered at time t if at least one of the two inventory conditions (for coffee delivery or SCG pick-up) is satisfied. In such cases, the system evaluates whether the second action can be performed as well without violating inventory logic.

More precisely, if a coffee delivery is required (i.e., $I_{i,t-1}^{+1} \leq r_{i,t}^1$), an SCG pick-up is also performed, provided that the amount to be collected exceeds half of the total bin capacity. Conversely, if a SCG pick-up is required (based on $r_{i,t}^2$), a coffee delivery is made in the same visit if the related inventory on-hand level is not too high, i.e., $I_{i,t-1}^{+1} \leq S_{i,t}^1 - r_{i,t}^1$. These additional rules do not override the original reorder policies but extend them to enable consolidation whenever possible. For comparison purposes, the standard unsynchronized policy is denoted as the *Baseline Inventory Policy (BIP)*.

According to the rolling horizon approach, the dynamic lot sizing model for the integrated inventory control of the two products is applied for each decision round at time period θ_j , for each coffee shop i . We assume that before planning pick-ups and deliveries at time θ_j , the initial levels of inventory on-hand I_{i,θ_j-1}^{+1} of coffee and $I_{i,\theta_j-1}^{+2} + \gamma(d_{i,\theta_j-1}^1 - \ell_{i,\theta_j-1}^1)$ of SCG are known and available in advance, thanks to the use of specific IoT technology.

3.3. Routing phase

The routing phase at the j^{th} rolling horizon round takes as input the amount $q_{i,t}^1$ of coffee delivery and the amount $q_{i,t}^2$ of SCG pick-up, for each coffee shop i , determined by the inventory control phase, for the frozen time period $t = \theta_j, \dots, \theta_j + |F_j| - 1$. However, since we do not know in advance the actual amount of SCG stocks in the coffee shop i and, accordingly with the pick-up inventory policy, we must recollect all SCG stocks when a pick-up is done, we overestimate the pick-up amount in next period $t = \theta_j + 1, \dots, \theta_j + |F_j| - 1$, that is, if $q_{i,t}^2 > 0$, then we assume $q_{i,t}^2 = \bar{U}_i^2$. For each time period (day) $t = \theta_j, \dots, \theta_j + |F_j| - 1$, we heuristically solve the related VRP with simultaneous pick-up & delivery and route time limitation H , by solving its mathematical model with a commercial solver within a given time limit.

4. Case Study and Results

The proposed approach has been tested on a real-world case study involving 100 coffee shop locations in Milan. The system includes a starting depot for coffee deliveries and a separate end depot for the collection of spent coffee grounds (SCG). These elements are shown in Figure 1, where coffee shops, the starting depot, and the end depot are represented in blue, green, and red, respectively.

Customer demand at each coffee shop was synthetically generated by distinguishing between locations with typically higher weekend demand and those with greater weekday demand. High-demand shops were assigned a mean daily number of coffee cups between 350 and 450, while low-demand shops ranged from 200 to 300. To capture variability in customer flow, each shop was also assigned a coefficient of variation ranging from 5% to 10%. Based on the simulated number of coffee cups demanded, daily coffee consumption (in kilograms) and the corresponding amount of spent coffee grounds (SCG) produced were estimated for each shop.

4.1. Rolling Horizon Planning Strategy

To evaluate the proposed methodology, we simulate a rolling-horizon planning environment over a 30-day period. At each decision round, forecast demand is used to compute inventory control parameters, and specifically the reorder

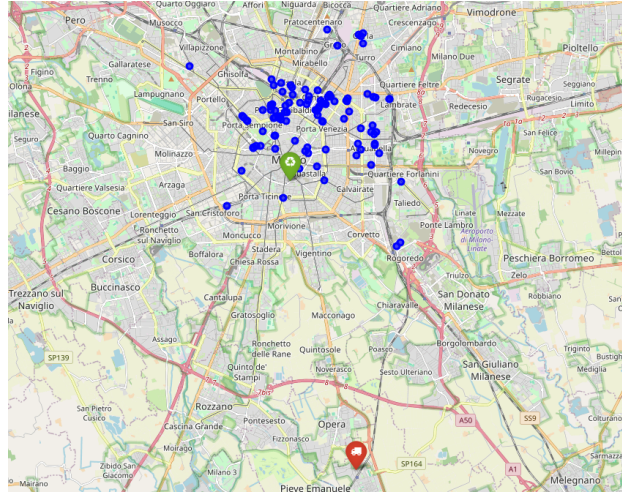


Fig. 1. Geographical distribution of the 100 coffee shops in Milan (blue), the starting depot for coffee delivery (green), and the ending depot for SCG collection (red).

point $r_{i,t}^1$ and the order-up-to level $S_{i,t}^1$ and the recollecting point $r_{i,t}^2$ for SCG, based on predicted coffee consumption. The RH heuristic is applied biweekly, every Monday and Thursday, using a 7-day planning window (i.e., $\lambda = 7$). However, only the first $f = 3$ or $f = 4$ days of the schedule, referred to as the frozen period, are fixed for execution. The remaining part of the horizon remains flexible and is re-evaluated at the next planning round with updated information.

Figure 2 illustrates the RH logic. The blue area represents the full 7-day planning window, while the green portion denotes the frozen segment. After the frozen window is executed, the process rolls forward and a new plan is generated for the next 7-day horizon.




Fig. 2. Rolling Horizon planning cycle: at each step, decisions are made for 7 days ahead, but only the pick-up and delivery operations of the first 3 (or 4) days are executed before the next re-planning occurs.

After the execution of the planned operations in each frozen window, the inventory levels are updated using the actual (realized) customer demand. These updated values serve as the new starting point for the next round of the planning process. This adaptive mechanism enables the model to respond to deviations from the forecast, improving the robustness and effectiveness of the generated plans.

In each round, delivery and pick-up schedules are determined while respecting the current inventory levels and minimizing the number of coffee shop visits. If synchronization is enabled, according to CIRP, the model combines deliveries and SCG pick-ups in the same visit whenever inventory constraints allow.

4.2. Performance analysis

The results here presented are based on simulated scenarios covering a 30-day planning horizon (i.e., $|T| = 30$). For each day, pick-up and delivery amounts and inventory levels were updated, and a daily vehicle routing problem (VRP) was solved accordingly. Table 1 reports the average per-bar performance metrics obtained under the Baseline Inventory Policy (BIP), where deliveries and pick-ups are handled independently, and the Consolidated Inventory-Routing Policy (CIRP), where coffee deliveries and SCG pick-ups are synchronized.

Table 1. Average per-coffee shop inventory and visit metrics under Baseline Inventory (BIP - no sync) and Consolidated Inventory-Routing (CIRP - sync) policies

Metric	BIP (no sync)	CIRP (sync)
Average frequency of coffee loss	0.014	0.009
Avg. coffee loss (kg), if present	0.162	0.148
Average frequency of SCG loss	0.014	0.004
Avg. SCG loss (kg), if present	0.282	0.061
Avg. coffee inventory on-hand (kg)	5.31	5.62
Avg. SCG inventory on-hand (kg)	7.24	6.18
Avg. daily visits with only delivery	6.83	0.51
Avg. daily visits with only pick-up	5.98	0.16
Avg. daily visits with pick-up & delivery	2.26	9.30
Avg. daily total number of visits	15.07	9.97

First, we observe that synchronization leads to lower product losses, both in terms of frequency and average amount (if loss occurs). The frequency of coffee loss drops, on average, from 0.014 to 0.009 per coffee shop, and the corresponding average loss decreases from 0.162 kg to 0.148 kg. Similarly, SCG losses show an even stronger reduction: from 0.014 to 0.004 in terms of frequency and from 0.282 kg to just 0.061 kg on average per coffee shop.

As expected, BIP produces a lower average inventory levels, since more frequent visits allow for tighter inventory control. CIRP leads to slightly higher inventory levels, but they remain comparable and within acceptable operational limits: 5.31 vs. 5.62 kg for coffee, and 7.24 vs. 6.18 kg for SCG.

The most significant difference emerges in the number and type of visits. Under CIRP, almost all interventions are joint visits combining coffee delivery and SCG collection (9.30 on average), while under BIP, visits are mostly disjoint (6.83 coffee-only and 5.98 SCG-only). The average total number of visits per coffee shop is also significantly reduced from 15.07 in BIP to 9.97 in CIRP, showing the efficiency gain in terms of number of visits for the latter policy.

To assess the routing performance over the full 30-day planning horizon, we averaged the results obtained from the daily Vehicle Routing Problems (VRPs). Table 2 summarizes the aggregated routing outcomes under the Baseline Inventory Policy (BIP) and the proposed Consolidated Inventory-Routing Policy (CIRP).

Table 2. Comparison between Baseline Inventory Policy (BIP) and Consolidated Inventory-Routing Policy (CIRP)

Policy	Average Cost	Average Distance	Average # of Vehicles	Average # of Visits
BIP (no sync)	548.73	27161.34	3.73	43.06
CIRP (sync)	395.41	19497.01	2.67	28.48

The consolidated strategy (CIRP) consistently outperforms the baseline in all evaluated metrics. The average daily routing cost is reduced from 548.73 to 395.41 (approximately 28%), primarily due to the significant drop in total distance traveled (from 27,161 to 19,497 meters). Moreover, the average number of vehicles used per day is reduced from 3.73 to 2.67.

These improvements are consistent with the reduction in the average number of daily coffee shop visits, which decreases from 43.06 under BIP to 28.48 under CIRP. This confirms that the reduction in routing cost and fleet size is largely due to better synchronization between coffee deliveries and SCG pick-ups. By consolidating visits, CIRP enables more efficient routing and minimizes redundant trips while maintaining service quality.

5. Conclusions

This study presents an integrated approach to manage coffee deliveries and spent coffee grounds (SCG) collection, combining dynamic inventory control with routing optimization. The proposed model reflects real operational constraints and leverages real-time data through a Vendor Managed Inventory framework supported by IoT infrastructure.

A Rolling Horizon (RH) heuristic was adopted to deal with the uncertainty and variability in daily customer demand. The RH approach allows regular re-planning, incorporating updated demand forecasts to enhance flexibility

and responsiveness. Additionally, a Consolidated Inventory-Routing Policy (CIRP) was introduced to synchronize deliveries and pick-ups whenever possible, with the aim of reducing the number of visits, and hence the routing costs.

Results on a real-world case involving 100 coffee shops in Milan show that synchronization significantly reduces the number of daily visits, vehicle usage, and total routing cost, without compromising inventory service levels. In particular, the CIRP achieves a better trade-off between service quality and logistics efficiency compared to the baseline policy without synchronization.

Future research will focus on several promising directions. First, the tuning of the rolling horizon parameters—such as the length of the planning window and the number of frozen days—could be optimized to balance responsiveness and stability across different demand scenarios. Finally, an anticipatory mechanism could be explored, where routing actions are scheduled ahead of time in the non-frozen portion of the horizon if significant inventory shortages are forecasted. These anticipations would still be re-evaluated in future planning iterations, but such proactive logic could further reduce loss and improve service continuity in highly dynamic environments.

Acknowledgements

This research has been partially supported by PRIN 2022 Project “PRISMA Platform - Physical internet Regenerative Sustainable MAterials”, n. 2022BE7XKZ PE8 PRIN 2022. Funded by European Union - Next Generation EU – Mission 4 Component 2 - Italian Minister of University and Research, Prin Project – PE08 - Physical Sciences and Engineering; CUP B53D23005830006.

References

- D. Kannan, S. Khademolqorani, N. Janatyan, and S. Alavi. Smart waste management 4.0: The transition from a systematic review to an integrated framework. *Waste management*, 174:1–14, 2023.
- L. Massaro Sousa and M.C. Ferreira. Spent coffee grounds as a renewable source of energy: An analysis of bulk powder flowability. *Particulology*, 43:92–100, 2019.
- A. Forcina, A. Petrillo, M. Travaglioni, S. Di Chiara, and F. De Felice. A comparative life cycle assessment of different spent coffee ground reuse strategies and a sensitivity analysis for verifying the environmental convenience based on the location of sites. *Journal of Cleaner Production*, 385:135727, 2023.
- N. Bostel, P. Dejax, P. Guez, and F. Tricoire. Multiperiod planning and routing on a rolling horizon for field force optimization logistics. In *Wasil, E. (eds), The Vehicle Routing Problem: Latest Advances and New Challenges. Operations Research/Computer Science Interfaces, vol 43*, pages 5151–5159. Springer, Boston, Ma. <https://doi.org/10.1007/978-0-387-77778-8-23>, 2008.
- J.F. Cordeau, M. Dell’Amico, S. Falavigna, and M. Iori. A rolling horizon algorithm for auto-carrier transportation. *Transportation Research Part B: Methodological*, 76:68–80, 2015.
- Y. Kim, D. Edirimanna, M. Wilbur, P. Pugliese, A. Laszka, A. Dubey, and S. Samaranyake. Rolling horizon based temporal decomposition for the offline pickup and delivery problem with time windows. In *AAAI’23/IAAI’23/EAAI’23: Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence, Article n. 575*, pages 5151–5159. <https://doi.org/10.1609/aaai.v37i4.25644>, 2023.
- A. Spinelli, F. Maggioni, T. R. P. Ramos, A. P. Barbosa-Póvoa, and D. Vigo. A rolling horizon heuristic approach for a multi-stage stochastic waste collection problem. *European Journal of Operational Research*, 323:1:276–296, 2024. doi: <https://doi.org/10.1016/j.ejor.2024.11.041>.
- M. Waller, M.E. Johnson, and T. Davis. Vendor-managed inventory in the retail supply chain. *Journal of business logistics*, 20:1:183, 1999.
- D. Aksen, K. Altinkemer, and S. Chand. The single-item lot-sizing problem with immediate lost sales. *European Journal of Operational Research*, 147:558–566, 2003.
- A.R. Clark. Rolling horizon heuristics for production planning and set-up scheduling with backlogs and error-prone demand forecasts. *Production Planning & Control*, 16:1:81–97, 2005.
- M.Z. Babai and Y. Dallery. Dynamic versus static control policies in single stage production-inventory systems. *International Journal of Production Research*, 47:2:415–433, 2009.
- L.C. Coelho, J.-F. Cordeau, and G. Laporte. Heuristics for dynamic and stochastic inventory-routing. *Computers & Operations Research*, 52: 55–67, 2014.