Batching and In-Building Delivery Routing with Capacitated Residential Parcel Lockers

Dipayan Banerjee[a,*] and Ignacio Erazo[b]

[a] Quinlan School of Business, Loyola University Chicago; dbanerjee2@luc.edu
[b] Amazon Fulfillment Technologies & Robotics; erazoignacio@gmail.com
[*] Corresponding author

Extended abstract submitted for presentation at the 12th Triennial Symposium on Transportation Analysis conference (TRISTAN XII)
June 22-27, 2025, Okinawa, Japan

Keywords: urban logistics, last-mile delivery, routing, branch-and-price

1 INTRODUCTION

Parcel delivery to single-family dwellings is generally straightforward: the driver parks the delivery vehicle, walks to the building, leaves the parcel(s) by the door, confirms the delivery (e.g., by taking a picture), and returns to the vehicle. On the other hand, parcel delivery to modern high-rise urban residential buildings is often facilitated by *residential parcel lockers* (RPLs). Unlike a 'public' (or 'shared') parcel locker, which can be used to make deliveries to any customer, an RPL is dedicated to privately serving the residents of one particular building. In the presence of an RPL, a driver places a resident's parcel into an available compartment, where it is stored securely until the resident retrieves the parcel with a private code. This allows drivers to quickly deliver several parcels in succession without having to move from the vicinity of the locker.

Empirical research has demonstrated the utility of RPLs to recipients and delivery firms (Mohri *et al.*, 2024, Ranjbari *et al.*, 2023a). However, the benefits of RPLs — security for parcel recipients and streamlined operations for delivery drivers — are only realized if a locker compartment is available (i.e., empty) when a driver attempts to deliver a parcel. When all RPL compartments are occupied, the driver may need to spend significant time walking and traveling on elevators to deliver parcels to recipients' doorsteps inside the building. This is a particular concern when the number of parcels to be delivered is atypically high due to retail events such as Black Friday or Amazon Prime Day, motivating the need to carefully optimize in-building delivery operations in these scenarios. Nevertheless, optimization-based research incorporating RPLs is scarce in contrast to the numerous studies optimizing aspects of public parcel lockers. To our knowledge, only Ranjbari *et al.* (2023b) optimize RPL-related decisions, albeit at a strategic level: the authors use simulation to choose an optimal RPL size and configuration.

In this work, we optimize the operations of a parcel delivery firm that has access to an RPL in the lobby of a high-rise building. We assume that the number of parcels to be delivered on a particular day exceeds the number of available RPL compartments; parcels that cannot be delivered to the RPL must be delivered to residents' doorsteps on upper floors of the building. Given certain building features, such as floor layouts and elevator waiting time distributions, we seek to determine the optimal subset of parcels to deliver to the RPL. For the remaining parcels, we seek to optimize the driver's travel route within the building; when the number of non-RPL packages exceeds the driver's carrying capacity, we additionally seek to determine an optimal separation of the non-RPL packages into batches. Our objective is to minimize the expected time to deliver all of the building's parcels. Under mild assumptions, we first derive a polynomial-time algorithm for determining the optimal sequence of floor visits given a set of parcels to deliver at residents' doorsteps. We then leverage this algorithm to develop a branch-and-price solution approach capable of solving realistically sized instances of the full problem to optimality.

2 PROBLEM DESCRIPTION

A parcel carrier sends a driver from a local depot to a high-rise residential building to deliver a set of parcels P. The carrier has access to all or part of an RPL located just inside the main entrance to the building; the RPL is made up of several compartments, each of which can hold exactly one parcel. The RPL is known by the carrier to have c available compartments to be used by the driver for delivering some of the parcels in P. Motivated by real-life examples (e.g., Kusisto, 2015), the parcels that are not placed in the RPL cannot be left in the lobby; rather, they must be delivered by the driver to recipients' doorsteps by traveling through the building.

Prior to the driver's departure from the depot, the carrier physically partitions the parcels P into batches $B_1, B_2, \ldots, B_m$, each with a capacity of q parcels. Within each batch B_ℓ, some parcels $R_\ell \subseteq B_\ell$ are designated to be delivered directly to the RPL, and the remaining parcels $D_\ell = B_\ell \setminus R_\ell$ are designated to be delivered directly to recipients' doorsteps. Upon parking at the building, the driver brings a tote containing batch B_1 to the RPL. The driver delivers the parcels R_1 to the RPL; the setup time at the RPL (e.g., for logging into the system) is α minutes, and the per-parcel service time at the RPL (e.g., for confirming delivery) is β minutes. Upon completing the RPL deliveries in $\alpha + \beta|R_1|$ minutes, the driver travels through the building to deliver the tote's remaining packages to recipients' doorsteps. The per-parcel service time at a doorstep is γ minutes; further details on in-building travel are given in the following subsection. The driver then returns to the vehicle and repeats this process until no totes remain.

The carrier seeks to minimize the expected time spent by the driver delivering the $|P|$ parcels. To this end, the firm must determine which parcels are placed into each capacity-constrained batch (i.e., tote), including determining which parcels are to be delivered to the RPL. It is clear that these decisions depend heavily on the layout of the building, including the behavior of the building's elevator(s), which we formally describe next.

2.1 Building layout and elevator behavior

We consider a high-rise residential building with 20+ floors, all connected by one or more adjacent elevators. The building includes a lobby (floor 0) that contains the RPL but no residences. Directly above the lobby, there may exist some 'non-residential' floors containing amenities such as a gym and parking garage. Directly above these floors are a set of m 'residential' floors. Each residential floor is assumed to be arranged as a tree graph with the elevators as the root node; beyond this assumption, the exact layout of residences' doorsteps on each floor may differ. Finally, there may exist one or more additional non-residential floors above the residential floors.

The driver walks at a constant speed within the building. A significant amount of the time spent delivering packages, however, is likely to be spent waiting and traveling on elevators. Concrete, publicly available data on residential building elevator operations is virtually nonexistent. Because elevator algorithms are generally proprietary, it is unrealistic to expect reliable knowledge of any given building's elevators' programmed behavior. Therefore, we proceed with the following mild assumptions on elevator behavior that do not require such algorithmic knowledge.

First, we assume that the time spent waiting on an elevator to arrive after pressing the 'up' (resp. 'down') button on floor i is given by an iid random variable $W_i^{\uparrow}$ (resp. $W_i^{\downarrow}$). Second, we assume that the probability of a building resident using an elevator to travel between any two *residential* floors is negligible; residents may take trips beginning or ending at non-residential floors, which delay the driver's travel. Third, we assume that it is never optimal for the driver to 'double-back' between floors; that is, for arbitrary floors $\{i, j, k\}$ with $i < j$ and $i < k$, an optimal in-building tour will never include travel from j to i followed by travel from i to k.

3 ANALYTICAL RESULTS AND SOLUTION METHOD

This setting is closely related to computationally intensive optimization problems such as the VRP and prize-collecting TSP. We show, however, that the structure of the residential building can be leveraged within column generation to facilitate optimization. In this section, we briefly summarize some key results and give a high-level overview of our solution method.

3.1 Inter-floor routing

Within the larger problem, a fundamental component is the subproblem of optimally routing the driver within the building to deliver all parcels in a given tote. We define the *inter-floor TSP* (IFTSP) as the formal optimization problem of determining the quickest route through the building (route over the floors to be visited) to deliver a given set of parcels to recipients' doorsteps. Because of potential asymmetries between each $\mathbb{E}[W_i^{\uparrow}]$ and $\mathbb{E}[W_i^{\downarrow}]$, it is not immediately clear whether the IFTSP is $\mathcal{NP}$-hard like the asymmetric TSP. We prove the opposite:

Proposition 1. *The IFTSP can be solved to optimality in polynomial time.*

The argument broadly proceeds as follows. First, we show that a 'backbone' of the optimal tour is an upward journey from the lobby to the highest floor represented in the given set of parcels, followed by a downward return journey to the lobby. Then, we show that every other floor in the given set must be visited on either the upward or downward journey (but not both), and the decision for each floor can be made in polynomial time independent of other floors. Of particular computational and managerial usefulness, an intermediate step within the proof of this result allows for the expression of the IFTSP objective as a linear function of the elevators' waiting time parameters and the topmost floor represented within a given set of parcels.

3.2 Branch-and-price

A standard integer programming model of the full problem becomes intractable even for fairly small instances, necessitating the development of a specialized solution method. We develop a branch-and-price approach to solve the full problem. Variables $x_{\ell,S}$ in the restricted master problem represent an assignment of a subset $S \subseteq P$ to the 'doorstep' sub-batch of B_ℓ; i.e., $x_{\ell,S} = 1$ if and only if $D_\ell = S \subseteq B_\ell$. At a very high level, we add columns to the restricted master problem's LP relaxation by seeking the solution to pricing problems of the form

$$\max_{\ell,S} \quad \sum_{j\in S} \left[e_j + \theta|S| - \texttt{TSP}(S)\right] \tag{1a}$$

$$\text{s.t.} \quad S \text{ is feasible for bin } \ell \ , \tag{1b}$$

where (1b) encodes several specific mathematical constraints omitted for brevity.

In this high-level representation, θ and each e_j are variables in the LP relaxation's dual problem, while $\texttt{TSP}(S)$ denotes the expected length of time required to deliver the parcels S to recipients' doorsteps. We note that $\texttt{TSP}(S)$ can be calculated efficiently; Proposition 1 shows that optimal travel paths between floors can be computed efficiently, and the tree layout of each residential floor allows for efficient computation of optimal intra-floor walking paths. As is typical, branching occurs when the restricted master problem's LP is solved to optimality with a fractional solution. The resulting branch-and-price algorithm is more efficient with respect to both computational time and memory requirements.

4 COMPUTATIONAL RESULTS AND DISCUSSION

We present a brief computational illustration to demonstrate the benefits associated with optimizing the composition of each batch. Consider the delivery of 24 parcels to a building with 640 total apartments distributed across 20 residential floors.

We compare the following solutions: (i) the optimal solution (`OPT`) with optimal batching and assignment of parcels to the RPL; (ii) a top-down heuristic (`TDH`) solution, in which parcels are grouped by the floor of their recipient, and parcels designated for the topmost floors are delivered to the RPL; (iii) a randomized heuristic (`RH`) solution, in which the best solution is chosen from one million random assignments of parcels to batches; and (iv) a doorstep-only (`DO`) solution, in which batches are optimized without using the RPL.

In-building travel and elevator routing for each batch is optimized for all four approaches. We test tote and RPL capacities of $q = c = 6$, $q = c = 8$, and $q = c = 12$ in this abstract. Table 1 displays the mean expected delivery time (in seconds, including both RPL and doorstep deliveries) across 50 randomly generated instances. For the latter three approaches, the table also records the percentage gap compared to the `OPT` solution's mean expected delivery time.

		Mean expected delivery time (sec) and gap vs. `OPT` (%)			
q	c	`OPT`	`TDH`	`RH`	`DO`
6	6	1930.7	2091.1 (8.3%)	2029.7 (5.1%)	2376.1 (23.1%)
8	8	1631.7	1787.9 (9.6%)	1692.2 (3.7%)	2189.7 (34.2%)
12	12	1251.9	1391.8 (11.2%)	1259.3 (0.6%)	2002.9 (60%)

Table 1 – *Comparison of parcel batching policies*

We highlight key insights from these results. First, as expected, simultaneously increasing q and c entails improved performance for each approach. Increasing q allows for greater economies of scale on in-building travel routes, and increasing c allows the driver to avoid in-building travel for more parcels. Second, the mere presence of the RPL improves performance: in each of the three cases, `DO` (which does not utilize an RPL) is clearly dominated by `RH`, `TDH`, and `OPT` (all of which utilize the RPL). The relative underperformance of `DO` is significant even for $q = c = 6$, and the disparity only becomes more apparent as $q = c$ increase. Finally, we observe that the benefits of the RPL are not fully realized without optimized batch creation *and* optimized in-building routing, as `OPT` outperforms both the `RH` and `TDH` heuristics. In particular, we note that `RH` is able to achieve near-optimality for $q = c = 12$ because the space of feasible batches is relatively small (and thus more easily 'searched' by randomization). However, the search space grows as $q = c$ decrease, thereby increasing the `RH`-`OPT` gap. Further experiments with larger instances and sensitivity analyses will be available in the full manuscript.

4.1 Ongoing and future work

Some firms may require a driver to deliver parcels to a number of high-rise buildings in succession before returning to the depot; the problem of determining the optimal visit sequence of buildings; unlike the single-building problem, involves significant stochasticity because the number of available RPL compartments at each building may change over the course of the day.

References

Kusisto, Laura. 2015. *Web-Shopping Deluge Boxes In Landlords: Camden Property Trust takes a tough stand and stops accepting package deliveries.* The Wall Street Journal, URL `https://t.ly/0NLsy`.

Mohri, Seyed Sina, Nassir, Neema, Thompson, Russell, & Ghaderi, Hadi. 2024. Last-mile logistics with on-premises parcel lockers: who are the real beneficiaries? *Transportation Research Part E*, **183**, 103458.

Ranjbari, Andisheh, Diehl, Caleb, Dalla Chiara, Giacomo, & Goodchild, Anne. 2023a. Do parcel lockers reduce delivery times? Evidence from the field. *Transportation Research Part E*, **172**, 103070.

Ranjbari, Andisheh, Diehl, Caleb, Dalla Chiara, Giacomo, & Goodchild, Anne. 2023b. What is the right size for a residential building parcel locker? *Transportation Research Record*, **2677**(3), 1397–1407.