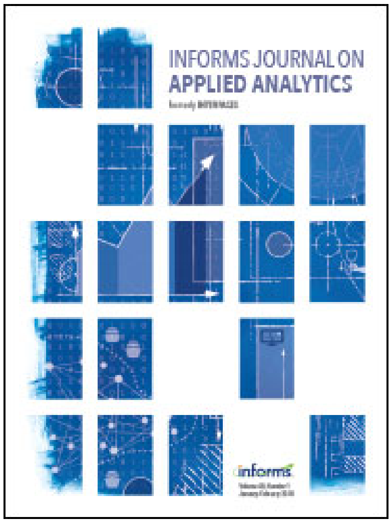

INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Smart Parcel Consolidation at Cainiao

Yujie ChenBiao Yuan, Yinzhi Zhou, Yuwei Chen, Haoyuan Hu

INFORMS JOURNAL ON APPLIED ANALYTICS
Vol. 54, No. 5, September–October 2024, pp. 417–430
ISSN 2644-0865 (print), ISSN 2644-0873 (online)
https://pubsonline.informs.org/journal/ijaa

Smart Parcel Consolidation at Cainiao

Yujie Chen,[a] Biao Yuan,[b,c,*] Yinzhi Zhou,[a] Yuwei Chen,[a] Haoyuan Hu[a]

[a] Cainiao Network, Hangzhou, Zhejiang 311100, China; [b] Data-Driven Management Decision-Making Laboratory, Shanghai Jiao Tong University, Shanghai 200030, China; [c] Sino-US Global Logistics Institute, Shanghai Jiao Tong University, Shanghai 200030, China
*Corresponding author

Contact: aisling.cyj@alibaba-inc.com (YujC); biaoyuan.ie@sjtu.edu.cn, https://orcid.org/0000-0002-3100-5993 (BY); yinzhi.zyz@cainiao.com (YZ); chenyuwei.chenyuwe@cainiao.com (YuwC); haoyuan.huhy@cainiao.com (HH)

https://doi.org/10.1287/inte.2024.0124

Abstract. Cainiao proposes a novel business model that consolidates parcels ordered by the same consumer from one or more merchants during the fulfillment process. The objective is to increase delivery speed without incurring additional costs for merchants and consumers. To support this business model, we develop three analytics methods: (1) a two-phase online optimization algorithm to determine which of a consumer's parcels constitute consolidated parcels and to select the shipping methods for the consolidated parcels that maximize the gains while satisfying the constraints (e.g., the 10-day on-time delivery rate of all consumer parcels created within a specified time should reach a target value), (2) a statistical method to calculate delivery time distributions to obtain on-time delivery rates within different days, and (3) a simulation-based optimization method to guide managers in setting appropriate target values for the constraints. In addition, we prove that the expected optimality gap and constraint violation of the online optimization algorithm have sublinear bounds, and we validate its effectiveness and robustness by testing instances generated from real-world data. Since 2020, Cainiao has utilized the system to consolidate numerous parcels shipped from China to more than 50 countries and regions, thus saving tens of millions of dollars annually and reducing delivery time by at least 50%.

History: This paper has been accepted for the *INFORMS Journal on Applied Analytics* Special Issue—2023 Daniel H. Wagner Prize for Excellence in the Practice of Advanced Analytics and Operations Research.

Funding: B. Yuan was supported by the National Natural Science Foundation of China [Grant 72301170] and the Startup Fund for Young Faculty at SJTU (Shanghai Jiao Tong University) [Grant 23X010502006].

Keywords: cross-border logistics • shipment consolidation • online optimization

Introduction

Cainiao Network, which we refer to as Cainiao in the remainder of the paper, was founded in 2013 and is a global logistics industrial Internet company and the logistics arm of Alibaba Group. Cainiao; FedEx and UPS both based in the United States; and Germany-based DHL constitute the four major global players in cross-border logistics. Cainiao's cross-border parcel network covers more than 200 cities in China, and the company has established more than 50 collection warehouses in major cities, serving more than 220 countries and regions. To date, its average daily cross-border parcel volume has reached 4.5 million. Because of its fast and reliable logistics services, Cainiao, which offers various shipping methods with varying delivery costs and times, has been selected as one of the logistics providers for AliExpress, Alibaba's cross-border e-commerce platform.

Shipping and delivery are important aspects of the shopping experience and may impact the willingness of consumers to purchase goods from merchants in other countries. The results of a global survey by solution provider ESW confirm this point, with 27% of the participants selecting long shipping times and high shipping costs as the main reasons they avoid cross-border purchases (Gillai and Lee 2023). To achieve the goal of fulfilling consumer orders in a cost-effective manner, Cainiao is continuously improving its parcel fulfillment process. In 2020, it proposed a novel business model, which during transportation consolidates the parcels ordered by the same consumer from one or more merchants. The objective is to increase the delivery speed without incurring extra costs for merchants and consumers.

To understand the benefits of parcel consolidation during fulfillment, we present a hypothetical example in Tables 1 and 2. The values in the tables are for illustrative purposes and do not reflect actual data. Table 1 shows three parcels ordered by a consumer from three merchants with their creation times, shipping methods,

Table 1. Details of Three Parcels Ordered by One Consumer

Parcel code	Merchant code	Creation time	Shipping method	Weight (g)
LP0001	M1	2023-03-20 13:40:20	Economy	100
LP0002	M2	2023-03-21 12:20:19	Economy	200
LP0003	M3	2023-03-22 10:30:19	Economy	150

and weights. Table 2 lists the delivery fees and times of the shipping methods, where the first weight fee is the price for the first 100 grams, and the additional weight fee is the price for each additional gram. For example, the delivery fee of Parcel LP0003 using the economy shipping method is $20 + (150 - 100) * 0.12 = 26.0$ CNY (CNY represents Chinese Yuan where 7.14 CNY equal approximately 1 USD). Without consolidation, the consumer receives all three parcels with a delivery time of 23 days and a fee of 78.0 CNY. However, if the first parcel (LP0001) waits for the last two parcels (LP0002 and LP0003) and is shipped with them using the standard shipping method, the consumer receives all three parcels with a delivery time of 13 days and a fee of 77.5 CNY. Therefore, with consolidation, Cainiao earns 0.5 CNY (the consumer or the merchant pays Cainiao 78.0 CNY, whereas Cainiao should only charge 77.5 CNY after consolidation), and the consumer receives all parcels 10 ($23 - 13$) days earlier than expected. In addition, the increased delivery speed without additional costs helps merchants on the e-commerce platform attract more consumers (Salari et al. 2022) and enhances the competitiveness of the platform. Not all parcel consolidations can earn money and increase delivery speed. For example, if the previous three parcels are combined and shipped using the economy shipping method, then that earns 16.0 CNY for Cainiao but delays the delivery time of the first two parcels. Therefore, if consolidation gains and losses are well balanced, Cainiao, the platform, merchants, and consumers can all benefit from this new business model.

Considering the operational costs and times, Cainiao built warehouses in China for parcel consolidation. After receiving the order, the company delivers the corresponding parcel from the merchant to the consumer, as Figure 1 shows. The figure shows that the parcel is collected and transported to the collection warehouse, the predefined sorting center, the consolidation warehouse if necessary, the origin Customs, and the origin port, thus completing the delivery in China. Once the parcel arrives at the destination port, it is transported to Customs at the destination, the sorting center, the delivery station, and the consumer.

After the consumer pays for the order, the merchant prepares the goods and creates a package. Meanwhile, the decision support system monitoring the process searches for all other parcels to be shipped to the same consumer. According to the consolidation rules set by Cainiao's managers, the system then determines whether these parcels can be combined into one or more consolidated delivery groups, each of which is determined by the consolidation decision algorithm. The parcels in one consolidated decision group are referred to as associated parcels. A consolidated delivery group always refers to a group of packages (two or more) to be shipped to one consumer and that satisfies the business delivery rules. A consolidation decision of one original parcel is marked as final (i.e., the decision cannot be changed) by the system if the parcel or one of its associated parcels (if one exists) has reached the latest predefined time to be shipped; otherwise, the parcel decision may be updated by the system when a new parcel for the same consumer arrives. In addition, once the decision has been marked as final, all associated parcels are updated as final (i.e., they cannot be considered for consolidation any longer even if a new parcel for that customer arrives). Thus, one parcel may be involved in multiple decisions until it is marked as final. Figure 2 shows that when a parcel arrives at the collection warehouse, it is sent to the consolidation warehouse if it was included in a consolidated delivery group in the last involved decision; otherwise, it follows the subsequent fulfillment process without consolidation. If all parcels in one consolidated delivery group have arrived at the consolidation warehouse, they are packaged into a large parcel and sent from the warehouse for fulfillment; otherwise, the parcels are stored in a cell at the consolidation warehouse, as Figure 3 shows.

The objective of the optimization problem in the consolidation decision process is to determine which of

Table 2. Delivery Fees and Times of Standard and Economy Shipping Methods

Shipping method	First weight fee (CNY)	Additional weight fee (CNY/g)	Delivery time (days)
Standard	25	0.15	10
Economy	20	0.12	20

Figure 1. (Color online) Parcel Fulfillment Process

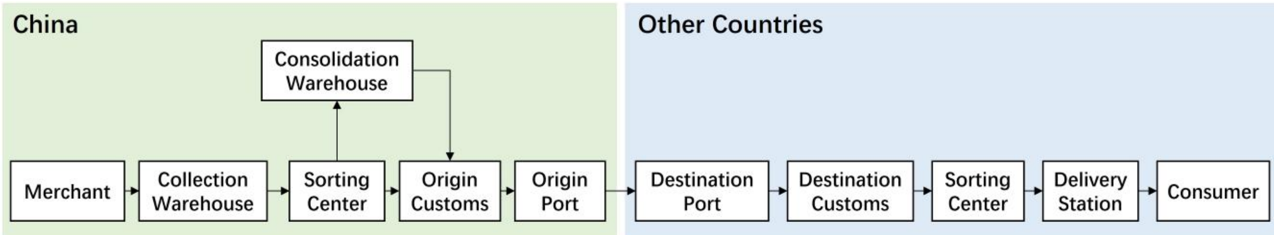

each consumer's parcels can be included in a consolidated delivery group and the shipping method that will provide the maximum benefit while meeting the delivery time constraints. The constraints in the problem focus not on each parcel but on a set of parcels. For example, the on-time delivery (OTD) rate within a given number of days (e.g., 10) should exceed a target value (e.g., 0.5), where the OTD rate refers to the ratio of the number of parcels delivered to consumers within the predefined days to the total number of parcels created in the given time. Obviously, when we decide on whether a parcel should be included in a consolidated delivery group, we only know the information of the parcels created before the current parcel instead of the global information of all parcels created within the considered horizon (e.g., 30 days), which makes building a mathematical model to solve the problem impossible. Therefore, solving such an online optimization problem with multiple decisions for a parcel, a problem that has not been investigated in the literature, is our main challenge in implementing our business model.

Given the business benefits of the parcel consolidation model and the challenges in building a solution approach to support it, our contributions are as follows:

- From a methodological point of view, we proposed a two-phase online optimization algorithm to effectively solve Cainiao's parcel consolidation problem. Unlike other approaches in the literature that address shipment

Figure 2. Simplified Consolidation Process

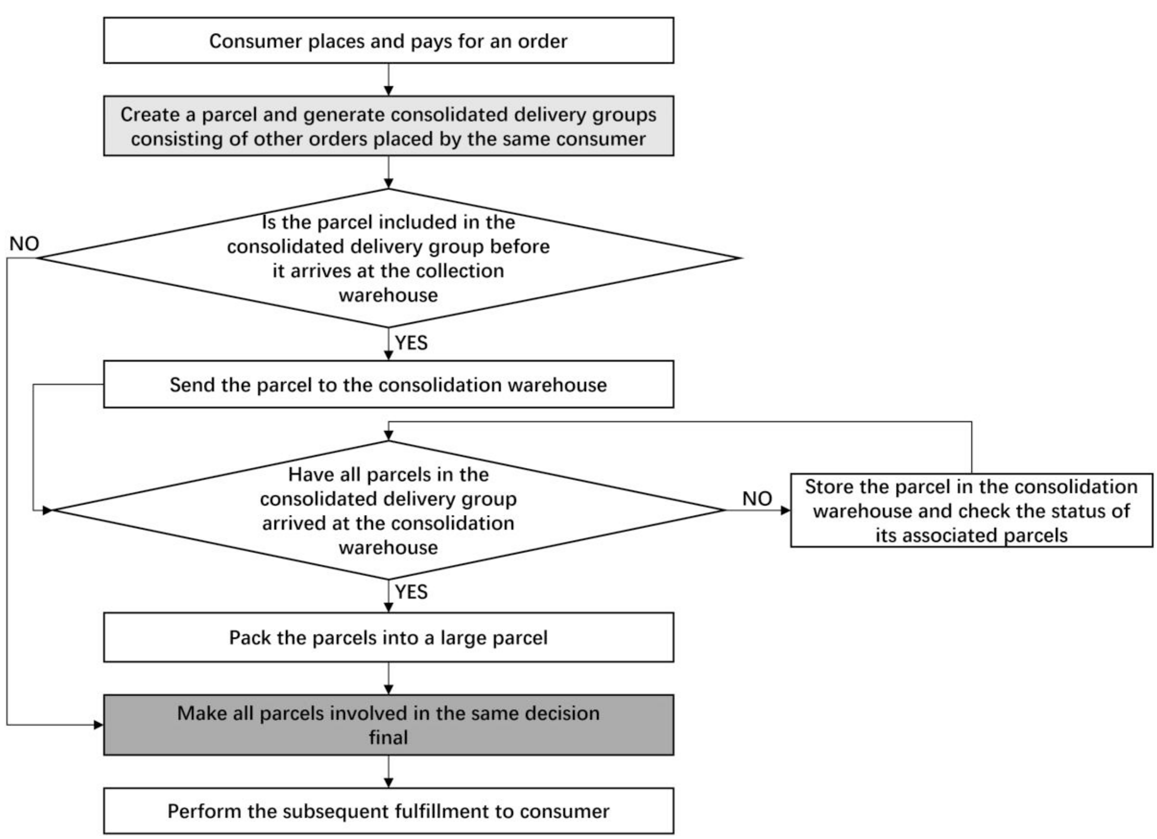

Note. Light gray rectangle indicates the consolidation decision, and the dark gray rectangle indicates that the decision is final.

Figure 3. (Color online) Cells Used to Store the Waiting Packages in Consolidation Warehouses

consolidation, the proposed algorithm requires no knowledge of future parcels. Moreover, we proved that the expected optimality gap and constraint violation of the algorithm have sublinear bounds and verified its effectiveness and robustness by testing instances generated from real-world data. In addition, we developed a statistical method and a simulation-based optimization method to accurately define the left- and right-hand-side (LHS and RHS) coefficients of the constraints for maximizing the benefits of the business model. The previous three methods work as a unit to address the challenges that Cainiao faced.

- From an application point of view, we introduced a new shipment consolidation model used in cross-border e-commerce logistics, which provides many benefits to Cainiao, e-commerce platforms, merchants, and consumers. We also demonstrated how to develop and implement the three analytics methods to support this model. After implementing the proposed methods in the decision support system, we consolidated numerous parcels shipped from China to more than 50 countries and regions, saving tens of millions of dollars annually and reducing delivery time by at least 50%. Such a successful implementation can also guide other companies in solving similar real-world problems.

We organized the rest of the paper as follows. In the *Related Literature* section, we review the relevant literature. In the *Solution Approach* section, we introduce the solution framework and its three analytical modules. In the *Implementation and Results* section, we present the implementation of the solution approach and the validation of the proposed online algorithm. The *Conclusion* section concludes the paper.

Related Literature

In this section, we review two streams of literature that are most closely related to our paper: shipment consolidation and online linear programming.

Shipment Consolidation

Shipment consolidation is a widely used logistics strategy to achieve economies of scale in transportation by combining two or more shipments (or orders) into one larger shipment, thereby reducing costs (Higginson and Bookbinder 1995). It can also result in reducing harmful emissions that affect air quality, such as CO_2, emitted from the exhaust of delivery vehicles (Ülkü 2012). Shipment consolidation, which has been extensively investigated in different scenarios in the literature (Çetinkaya 2005), can be implemented on its own (Zhang et al. 2019, Wei et al. 2021, Xu et al. 2023), or in coordination or integration with some other decisions, such as lot-sizing (Chan et al. 2002a, b), inventory replenishment (Chan et al. 2002a, b; Lee et al. 2003), order fulfillment (Wei et al. 2021), and production planning (Li et al. 2020). Our problem belongs to the former implementation. Hence, we review several relevant papers without involving other decisions to demonstrate various considerations in shipment consolidation. Zhang et al. (2019) study the order consolidation problem in the last-mile delivery, deciding whether to consolidate multiple shipments for the same customer and how many periods to postpone each shipment for a set of orders with known information while minimizing the sum of vehicle-dispatching cost, shipping cost, and inventory cost. The authors develop a three-phase algorithm to solve the problem and compare it with the first-in-first-out rule. Wei et al. (2021) study the shipping consolidation across two warehouses for e-commerce and omni-channel retailers. In the scenario of one warehouse with fixed and variable costs, the authors consider whether it is more economical to delay the shipments of some orders to consolidate them with future orders that will arrive with a given probability and use dynamic programming to study the optimal policy and its structure for balancing the fixed cost and the additional cost of expedited shipments. Xu et al. (2023) investigate the shipment consolidation problem in long-haul and short-haul shipping, determining how to consolidate orders into shipments, when to send a shipment, and which shipping method to use for a shipment to deliver a set of orders that were already placed while minimizing the total shipping cost and inventory cost. The authors develop the first-due-first-delivered and no-wait policies to solve the problem efficiently.

Our problem differs from most previous problems discussed in the consolidation literature in three ways. (1) We are the first to consider the delivery service performance of a set of parcels in a dynamic environment using statistical indicators (i.e., the OTD rate). Unlike the constraints considered in the existing literature, where the committed delivery time of each order (or package) should be met during transportation, in our constraints, we do not consider whether one package meets the delivery time requirement but rather the ratio of the number of parcels delivered within a predefined time. (2) Our problem has no knowledge of future parcels, whereas studies in the existing literature assume that the parcels to be merged are known or described with an arrival distribution. (3) Unlike most transportation operations in the existing literature, where all orders are shipped to a warehouse before consolidation decisions are made, we allow parcels not likely to be consolidated to skip the step of being shipped to the consolidation warehouse, which is also one of the decisions involved in our problem.

Online Linear Programming

Our parcel consolidation problem can be formulated as an underlying online optimization problem with set partitioning constraints, which is related to the online linear programming (LP) problem. There is a large body of literature on the design and analysis of algorithms for the online LP problem. Some of them focus on the particular forms of the LP problem and design the algorithms based on special constraint structures, for example, the secretary problem with the all-one constraint matrix (Arlotto and Gurvich 2019), the network revenue management problem with the finite support of random coefficients (Jasin 2015), and the resource allocation problem with the binary constraint matrix (Asadpour et al. 2020). Other studies consider the LP problem in a general form. Kesselheim et al. (2014) develop a primal-based algorithm to solve online LP problems. In each step, their algorithm solves a primal LP problem with revealed requests and a scaled capacity vector and randomly rounds the optimal LP solution to obtain an integer solution, which may be discarded if the constraints are not satisfied. Agrawal et al. (2014) propose a primal-dual algorithm for the problem that dynamically updates a threshold price vector at geometric time intervals and uses the price vector (i.e., the dual values corresponding to the constraints obtained by solving an LP problem formed by the revealed columns in the previous periods) to determine the sequential decisions in the current period. Li and Ye (2021) develop an action-history-dependent learning algorithm to solve the online LP problem, which considers the past input data and decisions/actions to improve the performance of the previous algorithm. The above algorithms may not be computationally efficient because they need to solve an LP problem, which can be prohibitive when the problem size is large. To deal with this issue, Li et al. (2023) and Balseiro et al. (2020) propose primal-dual algorithms for the problem, which perform the projected stochastic subgradient descent or dual mirror descent in the dual space and determine the primal solution based on the dual solution in an online manner.

Our problem has the following characteristics that make existing algorithms no longer applicable to it. (1) The set partitioning constraints incurred by assigning parcels into consolidated delivery groups make the problem more complex than those in the previous studies. (2) The operation of executing the last decision of multiple decisions involved in the fulfillment process of one package invalidates the current dual update procedure, like Li et al. (2023).

Solution Approach

To support the parcel consolidation model application, Cainiao needs both an optimization algorithm to make decisions and the methods to obtain the LHS and RHS coefficients of the constraints on the delivery performance. To meet these requirements, we develop three analytics modules as follows:

- A two-phase online optimization module makes consolidation decisions in real time.
- A statistical module computes delivery time distributions to periodically update OTD rates within a given number of days.
- A simulation-based optimization module identifies the appropriate target values of the constraints.

Figure 4 shows the three-layer relationship between the three modules and other input data, where the arrow indicates the information/data flow. The bottom layer contains consolidation rules (i.e., checking the feasibility of consolidated delivery groups) and delivery cost evaluation rules (i.e., calculating the delivery cost of packages using a given shipping method). These rules are defined by Cainiao's managers. If one rule is changed, managers must use the simulation-based optimization module in the middle layer to find appropriate target values of the constraints. The statistical module is used to periodically update the distributions associated with the OTD rates. Because the simulation-based optimization module and the statistical module are not used in the online decision environment, we call them offline modules. The top layer is the online optimization algorithm, which is used to make real-time decisions. The three methods are described in detail in the following three sections.

Online Optimization Module

The optimization problem we seek to address is tacitly described as follows:

We seek to maximize the gains from fulfilling the parcels arriving in the system over a finite horizon by determining which parcels destined for the same consumer constitute consolidated delivery groups and selecting the shipping methods for the consolidated delivery groups while satisfying delivery performance constraints.

Given the difficulty of solving such an online decision problem, we adopt the online primal-dual algorithm

Figure 4. (Color online) Framework of Our Solution Approach

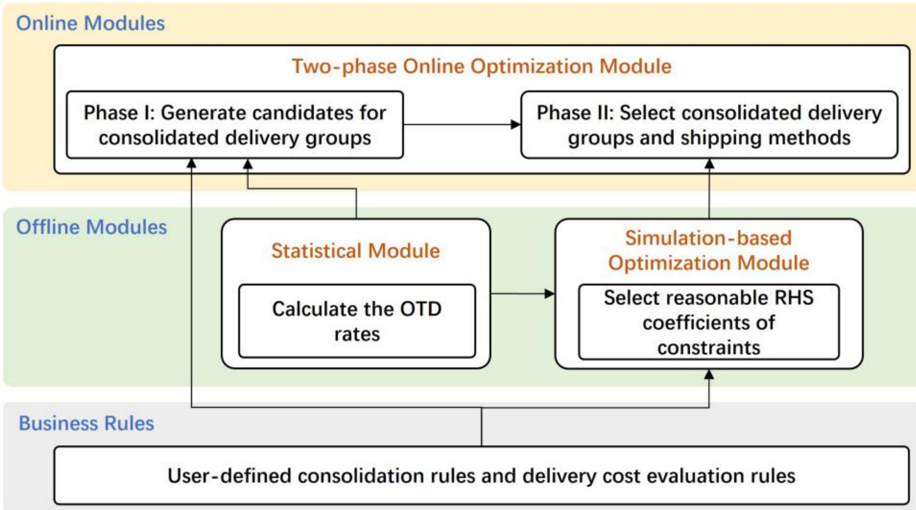

proposed by Li et al. (2023) as the framework of our algorithm. Li et al. (2023) devise the algorithm to solve a class of binary integer LP problems with constraint form $\mathbf{1}^\top \boldsymbol{x} \leq 1$, which arise in general resource allocation problems. Because we need to determine which parcels for the same consumer constitute consolidated delivery groups (i.e., to find the optimal partition of the set of parcels), our problem cannot be modeled in the form of Li et al. (2023). Moreover, in the dual update procedure of Li et al. (2023), the dual value is updated immediately after the decision is made. However, in our problem, the decision is made when a new parcel is created but can be changed before it is marked as final (Figure 2). To ensure the theoretical sublinear bound of the algorithm, the dual value should be updated only after the decision is marked as final. That is, the feedback of decisions has arbitrary delays in our problem. These factors motivate us to adapt their algorithm to the characteristics of our problem.

The online optimization module includes two phases: (1) generate consolidated delivery group candidates to formulate set partitioning constraints and (2) select consolidated delivery groups and their shipping methods. Cainiao's logistics partners always have varying delivery requirements or rules. For example, the total price of goods in the parcel cannot exceed a certain value to avoid the extra tariff, and the sum of the length, width, and height of the parcel cannot exceed 90 cm. The former can be easily checked by summing the prices of the parcels in a consolidated group; for the latter, we have to invoke the three-dimensional bin packing routine (Fontaine and Minner 2023) to check the feasibility. Moreover, for most consumers, the number of packages to be combined does not exceed five, that is, there are at most 31 ($2^5 - 1$) candidates if the individual parcel is also viewed as a candidate. Considering the previous two practical factors and the limited computational time, we develop an enumeration and filtering heuristic to generate candidates. Specifically, if the number of packages is not greater than five, we first enumerate all possible packages and then filter them with user-defined consolidation rules; otherwise, we sort the orders according to some criteria (e.g., the creation time) and then sequentially generate the candidates for every five packages with a given step size (e.g., two). At the end of the enumeration and filtering process, we associate the candidates with all possible shipping methods and calculate the cost difference (equal to the cost of the individual packages minus that of the consolidated package, that is, the gain from consolidation) and the OTD rates within predefined days (e.g., 10, 20, and 30). Therefore, once a candidate is selected, its shipping method is also determined. In the following description, we also consider the individual parcels as candidates.

The candidates are regenerated when a new parcel arrives but become final when the associate decision is marked as final. Before introducing the second phase, we construct an offline parcel consolidation problem by using the candidates generated in the final decision of each consumer:

$$
\begin{aligned}
\max_{x_{i,h}} \ & \sum_{i\in\mathcal{I}} \sum_{h\in\mathcal{H}_i} c_{i,h} x_{i,h} \\
\text{s.t.} \ & \sum_{i\in\mathcal{I}} \sum_{h\in\mathcal{H}_i} (r_{i,l,h} - \beta_l) n_{i,h} x_{i,h} \geq 0, \quad \forall l \in \mathcal{L} \\
& \sum_{h\in\mathcal{H}_i} \alpha_{i,o,h} x_{i,h} = 1, \quad \forall i \in \mathcal{I}, o \in \mathcal{O}_i \\
& x_{i,h} \in \{0,1\}, \quad \forall i \in \mathcal{I}, h \in \mathcal{H}_i,
\end{aligned}
\tag{1}
$$

where i and $\mathcal{I}$ are the index and the set of consumers, respectively; $\mathcal{O}_i$ and $\mathcal{H}_i$ are the sets of parcels and consolidated delivery group candidates of consumer i, respectively. When a decision related to a consumer is fixed, the consumer should be considered as another consumer if they create some packages at the same time, which we do not model explicitly for simplicity of description. Because a consumer corresponds to a final decision, i can also be the index of final decisions. The decision variable $x_{i,h}$ is one if the candidate h is selected and zero otherwise. Each candidate $h \in \mathcal{H}_i$ has a gain of $c_{i,h}$ and the OTD rates within l days, $r_{i,l,h}$, $l \in \mathcal{L}$ (e.g., {10, 20, 30}). Moreover, parameter $\alpha_{i,o,h}$ is one if parcel o is included in the candidate h and zero otherwise, β_l is the target value to be reached, at a minimum, for the average OTD rate in l days for delivering all parcels, and $n_{i,h}$ denotes the number of parcels in the consolidated delivery group h. The objective of Problem (1) is to maximize the total gain from consolidation. The first constraint guarantees that the average OTD rate within l days exceeds the predefined rate. The second constraint is the set partitioning constraint, which ensures that each parcel is assigned to exactly one candidate.

Algorithm 1 describes the overall process of solving Problem (1) in an online manner. In practice, each time a new parcel is created, the parcel with its associated parcels will incur a decision. That is, there are repeated decisions in the online setting, but the offline Problem (1) is constructed only by the coefficients of final decisions. The set of final decisions in the offline problem is only a subset of the decisions made in the online manner. Therefore, Algorithm 1 denotes the computation/decision round with i', not i in Problem (1). The key idea of Algorithm 1 is to maintain the values of dual variables, called *dual prices*. When a computation is triggered, it uses the dual prices to select the consolidated delivery groups with the highest profits. Specifically, we define the dual variable, μ_l, $\forall l \in L$, corresponding to the first constraint of Problem (1). At the decision i', we first find the relevant parcels and generate the candidate set, which is the first phase, and in the second phase, we select the proper consolidated delivery groups by solving Problem (2) in Line 5 of Algorithm 1. In parallel with Lines 3–5, whenever a decision is fixed, we update the dual prices μ_l by Line 7 with the step size η, the only parameter of the algorithm. We can prove that under some mild assumptions the expected optimality gap and constraint violation of Algorithm 1 have a sublinear upper bound. Interested readers can refer to Appendix A.

Algorithm 1 (Two-Phase Online Parcel Consolidation Algorithm)

1: Input step size η and initialize dual prices $\mu_l = 0$, $\forall l \in \mathcal{L}$
2: **for** $i' = 1, \ldots, N'$ **do**
  Phase I:
3:   Find the parcel set $\mathcal{O}_{i'}$ in the status of being able to be consolidated ordered by the same consumer as the parcel that triggers the decision i'
4:   Generate the candidate set $\mathcal{H}_{i'}$ from $\mathcal{O}_{i'}$ and compute $c_{i',h}, \alpha_{i',o,h}$ and $r_{i',l,h}$
  Phase II:
5:   Choose the consolidated delivery groups by solving

$$\begin{aligned} \max_{x_{i',h}} \sum_{h \in \mathcal{H}_{i'}} & \left(c_{i',h} + \sum_{l \in \mathcal{L}} \mu_l (r_{i',l,h} - \beta_l) n_{i',h} \right) x_{i',h} \\ \text{s.t.} \sum_{h \in \mathcal{H}_{i'}} & \alpha_{i',o,h} x_{i',h} = 1, \quad \forall o \in \mathcal{O}_{i'}, \\ & x_{i',h} \in \{0,1\}, \quad \forall h \in \mathcal{H}_{i'}. \end{aligned} \tag{2}$$

6:   **for** each decision k fixed during the decisions i' and $i'+1$ **do**
7:     Update dual prices $\mu_l \leftarrow \max\{\mu_l + \eta \sum_{h \in \mathcal{H}_k} (\beta_l - r_{k,l,h}) n_{k,h}, 0\}$, $\forall l \in \mathcal{L}$
8:   **end for**
9: **end for**

Statistical Module

The OTD rates $r_{i,l,h}$ are critical inputs of the proposed problem. To meet the manager's requirements for setting the OTD rates within different days, we first compute the delivery time distributions between different pairs of fulfillment states of parcels (e.g., between the creation and signed states, as Figure 5 shows) using historical data. According to the simplified consolidation process shown in Figure 1, we define the fulfillment states of packages and their transition relations in Figure 5, where each rectangle contains the fulfillment

Figure 5. Package Fulfillment States and Their Transition Relationships

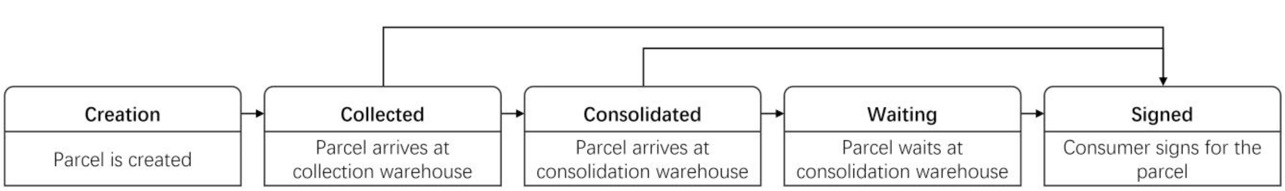

status and its meaning, and the arrows indicate the possible next states from the current one. For each online decision, we update the OTD rates of a parcel based on spatial and temporal changes in its fulfillment process.

We compute the distribution of delivery times as follows. First, we sample the parcels for which consumers signed in the past 30 days. Then, we classify the parcels according to the attributes of the parcels; examples include the province and country of the merchants and consumers, the shipping method, the category, and the consolidated flag. Let $\mathcal{G}$ denote the set of classes of parcels and $j \in \{1, 2, \dots\}$ denote the jth time interval. For each class of parcels $g \in \mathcal{G}$, we use the hour as the time unit and calculate the ratio of parcels whose delivery time between two fulfillment states is in the corresponding interval $(j-1, j]$ hour to the total number of parcels. The time intervals and the corresponding ratios form the discrete distribution $P(X = j|s, c)$, where $s \in \mathcal{S}$ is the fulfillment status and $c \in \mathcal{G}$ is the class of parcels. The computational process for each status and parcel class pair is similar to drawing a histogram.

Given the delivery time distribution, we calculate the OTD rate within l days of a newly created parcel o belonging to the parcel class g as $\sum_{j=1}^{j=24 \cdot l} P(X = j|s = \text{creation}, c = g)$. When the parcel consumes t_n days and reaches another state s', we update its OTD rate within l days as $\sum_{j=1}^{j=24 \cdot (l-t_n)} P(X = j|s = s', c = g)$. For a consolidated delivery group h containing a set of parcels $\mathcal{O}_h$ with different spatial and temporal states, its OTD rate within l days is defined as the average of the rates of the parcels.

Simulation-Based Optimization Module

Because the parcel consolidation process involves a stream of decisions, setting reasonable target value β_l for the first constraint of Problem (1) in such a dynamic environment is difficult for managers, which then affects the tradeoff between the gain and delivery service quality. Therefore, we develop a discrete event simulation-based optimization module to automate this task, which first evaluates all combinations of predefined possible target values using the simulation method, and then generates the best combination according to the objectives and constraint violations. Prior to the simulation, we prepare the data by first sampling the consumers who have signed for at least one package in the past 30 days and obtaining their corresponding packages with critical historical fulfillment information. The elapsed time between two fulfillment states is sampled from the corresponding delivery time distribution computed in the statistical module.

In the simulation, we use the fixed-time-increment principle instead of the event-paced principle to simulate the passage of time (Cassandras and Lafortune 2007). We do this because a computer cannot load the events of millions of parcels in its memory. Let s_o and t_o be the current status of package o and the time when reaching the current status, respectively. Let t_s and Δ be the start time and the time step of the simulation, respectively. The simulation details are as follows. For all packages, the values of s_o and t_o are initialized to the status of creation and the corresponding actual creation time, respectively. At the simulation step k, we first find the packages with $s_o \neq \text{signed}$ and $t_s + (k-1)\Delta \leq t_o < t_s + k\Delta$. Then, for each package sorted by the ascending order of t_o, we update the values of s_o and t_o to the next status and the corresponding occurrence time, as shown in Figure 5. For example, for a parcel o with $s_o = \text{creation}$, the value of s_o is set to the status of collected and t_o is set to the collected time. If the package incurs a decision, we invoke the two-phase online algorithm to generate consolidated delivery groups using the combination of input target values and then update s_o and t_o of associated parcels in the decision. The above iterations are repeated until all parcels are in the signed status.

Because the simulations for each combination of target values are independent, we run the simulation models in parallel with different combinations of target values to speed up the optimization process. Once all combinations are simulated, we calculate various performance metrics and provide them to the managers.

Implementation and Results

In this section, we present the details of the implementation, the performance of the proposed online optimization algorithm, and the practical benefits that Cainiao obtained.

Implementation

The decision support system, including the three analytics modules, is coded mainly in the Java programming language. Packages with detailed historical fulfillment information are stored in MaxCompute, a large-scale data warehousing and processing platform of Alibaba Cloud. In addition, the delivery time distributions are periodically computed by MaxCompute and then stored as key-value pairs in ApsaraDB for Redis, an in-memory database of Alibaba Cloud, for high-speed reading by the other two modules. The simulation-based optimization module does not use the existing commercial simulation software (e.g., AnyLogic and FlexSim) but uses Java to implement the fulfillment simulation logic described in the *Simulation-based Optimization Module* section to take full advantage of cloud computing. The online optimization module is designed as a real-time distributed system running on Alibaba Cloud to meet reliable, low-latency, and high-throughput service-level agreements. The optimization module uses the high-speed service framework (a microservice framework used internally by Alibaba Group) as an application

Table 3. Results of Five Instances Using the Online Algorithm

	$\beta_{10} = 0.45$				$\beta_{10} = 0.50$			
	Online		Offline		Online		Offline	
Instance	Gain	OTD rate	Gain	OTD rate	Gain	OTD rate	Gain	OTD rate
I1	2.9798	0.4391	3.0221	0.45	2.2463	0.4841	2.3128	0.50
I2	2.2158	0.4394	2.2995	0.45	1.7483	0.4874	1.8131	0.50
I3	1.7591	0.4464	1.8210	0.45	1.4041	0.4919	1.4712	0.50
I4	1.6859	0.4430	1.7429	0.45	1.5080	0.4877	1.5750	0.50
I5	2.2989	0.4423	2.3598	0.45	2.2211	0.4891	2.3079	0.50

Note. The scale of the data in the Gain column, which shows the average gain over all parcels, is obfuscated for confidentiality reasons.

programming interface that receives decision-making requests from the parcel fulfillment system and responds within tens of milliseconds with consolidated delivery groups for the requests to the fulfillment system. Meanwhile, the states of the algorithm (i.e., dual prices) are recorded in ApsaraDB for Redis.

Computational Results

To understand the performance of the two-phase online optimization algorithm, we perform a set of experiments by using our simulation module. First, we generate five instances using the parcels for which consumers signed in March 2022 and destined for five of the countries using the consolidation business model. Each instance contains millions of parcels. Then, we test each instance with two target values for the OTD rates within 10 days (i.e., β_{10}). For the step size η of the algorithm, in Appendix A, we prove the orders of the expected optimality gap and constraint violation are both $\sqrt{\bar{d}N}$ when the step size is $\eta = 1/\sqrt{\bar{d}N}$, where N is the total number of consumers and $\bar{d}$ is the average delay between decisions being made and fixed. Thus, in the experiments, we adopt $\eta^* = 1/\sqrt{\bar{d}N}$ as the step size. Table 3 presents the results of the five instances solved by Algorithm 1 when the value of β_{10} (i.e., the coefficient of the first constraint in Problem (1)) is set to 0.45 and 0.50, respectively. Column "Gain" reports the average gain over all parcels, which is obfuscated for confidentiality reasons. Column "OTD rate" shows the achieved OTD rate within 10 days of fulfilling all parcels. Columns "Online" and "Offline" denote the result of our online parcel consolidation algorithm and the optimal solution of the offline Problem (1) in the *Online Optimization Module* section, respectively. The offline Problem (1) is solved with optimization software. Noting that Problem (1) may be too large to solve, we split and solve the offline model containing a maximum of 500,000 consumers at each run.

Computational results in Table 3 show that Algorithm 1 tries to satisfy the predetermined OTD rate; the algorithm obtains the maximum constraint violation of 0.0159 for the instance I1 with $\beta_{10} = 0.50$. Moreover, it obtains a slightly lower OTD rate than the target. In practice, managers may set β_l slightly higher than the real target, making the result closer to their preferred outcome. The metrics of the gain and the achieved OTD rate are contradictory. Taking the instance I5 as an example, when $\beta_{10} = 0.45$ the achieved average gain is 2.2989; however, when $\beta_{10} = 0.50$, the achieved average gain is smaller: 2.2211. In conclusion, considering the unknown information, we believe that the algorithm can provide competitive solutions to the online parcel consolidation problem and adopt it as the algorithmic approach for Cainiao's decision support system.

The step size η is the only parameter of the online optimization algorithm (i.e., Algorithm 1). To observe its impact on the performance of the algorithm, we test the instance I4 by changing the step size from $0.1\eta^*$ to $5.0\eta^*$. Table 4 demonstrates the results of the instance solved by Algorithm 1 with different values of η. Columns "Gain" and "OTD rate" have the same meaning as in Table 3. Computational results in Table 4 show that the step size has little effect on the algorithm performance; that is, the gain and the OTD rate vary less at different step sizes, which proves its suitability for solving our problem.

Real-World Benefits

Cainiao implemented and adopted the parcel consolidation system and has used it since 2020. The system makes 3.8 million consolidation decisions daily for parcels delivered from China to more than 50 countries

Table 4. Results of Instance I4 Using Different Step Sizes in the Online Optimization

	$\beta_{10} = 0.45$		$\beta_{10} = 0.50$	
Step size	Gain	OTD rate	Gain	OTD rate
$0.1\eta^*$	1.7028	0.4390	1.5200	0.4864
$0.5\eta^*$	1.6939	0.4407	1.5155	0.4875
$1.0\eta^*$	1.6859	0.4430	1.5080	0.4877
$5.0\eta^*$	1.6760	0.4466	1.4991	0.4891

Note. As in Table 3, the scale of the data in the Gain column is obfuscated for confidentiality reasons.

and regions, saving tens of millions of dollars annually and shortening the delivery time by at least 50%. The substantial cost savings occur because the weight of most packages does not exceed the maximum allowable weight that incurs the additional weight fee and combining such packages can save the first weight fee, while the significant reduction in delivery time is due to the wide variation in promised delivery times for different shipping methods. As a result, upgrading a package's shipping method results in a twofold or greater increase in delivery speed. Taking logistics services from China to the United States as an example, the fastest and slowest shipping methods promised by Cainiao take 5 to 9 and 12 to 17 days, respectively. In addition, with the possibility of reducing costs and increasing delivery efficiency, Cainiao is introducing more cost-effective logistics services, such as 10-day delivery at a cost of 5 USD and 20-day delivery at a cost of 2 USD for selected countries. Also, the e-commerce platform can cooperate with Cainiao to provide promotional plans, whereby when the number or total price of goods purchased by a consumer exceeds a given threshold, the consumer can enjoy faster logistics services at no additional cost. Furthermore, combining parcels also reduces sorting and last-mile delivery operations in the subsequent fulfillment, which benefits energy conservation and emission reduction.

Conclusion

We present the three analytics methods for the support of a new business model of parcel consolidation proposed by Cainiao. The three methods have different objectives in the decision-making process. The two-phase online optimization module makes decisions in near-real time while respecting the delivery performance constraints, the statistical module provides the input data regarding the delivery time distributions for the other modules, and the simulation-based optimization module helps managers to select appropriate target OTD rates. These three modules are integrated to achieve a successful real-world application that helps Cainiao accomplish its business model purposes. In addition, the proposed framework is being extended to other organizations and similar business problems within Alibaba.

Appendix A. Proof for the Bounds of Algorithm 1

To obtain the bound on the expected optimality gap and constraint violation of Algorithm 1, we first transform the offline parcel consolidation problem and the online parcel consolidation algorithm into an allocation problem and an online allocation algorithm with arbitrary delays, respectively, which we show in Proposition A.1. We then prove in Theorem A.1 that the expected optimality gap (also called regret) and constraint violation of the equivalent online allocation algorithm are both on the order of $\sqrt{\bar{d}N}$, where N is the total number of requests (consumers in our setting) and $\bar{d}$ is the average of delays. Finally, by combining Proposition A.1 with Theorem A.1, we obtain the theoretical bounds of Algorithm 1 in Theorem A.2.

A.1. Problem Reformulation

First, we rewrite Problem (1) into the following compact form:

$$
\begin{aligned}
&\max_{\boldsymbol{x}_i} \sum_{i=1}^{N} \boldsymbol{c}_i^\top \boldsymbol{x}_i \\
&\text{s.t.} \sum_{i=1}^{N} (\boldsymbol{R}_i \boldsymbol{x}_i - \boldsymbol{\rho}) \le \boldsymbol{0} \\
&\qquad \boldsymbol{B}_i \boldsymbol{x}_i = \boldsymbol{1},\ \boldsymbol{x}_i \in \{0,1\}^{|\mathcal{H}_i|}, \quad \forall i = 1, \ldots, N,
\end{aligned}
\tag{A.1}
$$

where $\boldsymbol{c}_i = (c_{i,1}, \ldots, c_{i,|\mathcal{H}_i|})^\top$ is the vector of the gain of all consolidated delivery groups of consumer i, $\boldsymbol{x}_i = (x_{i,1}, \ldots, x_{i,|\mathcal{H}_i|})^\top$ is the decision vector, and $N = |\mathcal{I}|$ is the total number of consumers. $\sum_{i=1}^{N}(\boldsymbol{R}_i\boldsymbol{x}_i - \boldsymbol{\rho}) \le \boldsymbol{0}$ is the compact matrix form of $\sum_{i\in\mathcal{I}}\sum_{h\in\mathcal{H}_i}(r_{i,l,h} - \beta_l)n_{i,h}x_{i,h} \ge 0,\ \forall l \in \mathcal{L}$, where the entry in the lth row and hth column of $\boldsymbol{R}_i$ is $(\beta_l - r_{i,l,h})n_{i,h}$. $\boldsymbol{\rho}$ is a constant vector, which is equal to $\boldsymbol{0}$ in our case, but for more generality we keep $\boldsymbol{\rho}$ and all theoretical analyzes are valid for nonzero $\boldsymbol{\rho}$. $\boldsymbol{B}_i\boldsymbol{x}_i = \boldsymbol{1}$ is the compact matrix form of constraints $\sum_{h\in\mathcal{H}_i}\alpha_{i,o,h}x_{i,h} = 1,\ \forall o \in \mathcal{O}_i$.

Algorithm A.1 (Another Form of Online Parcel Consolidation Algorithm 1)

```
Input η and initialize z_1 = 0 and τ = 1
for i = 1, ..., N do
    Set μ_i = z_τ
    Compute x_i = arg sup_{x_i} {c_i^T x_i − μ_i^T R_i x_i | B_i x_i = 1, x_i ∈ {0,1}^{|H_i|}}
    Receive {R_k x_k | k ∈ F_i}
    for k ∈ F_i do
        Compute z_{τ+1} = max{z_τ − η(ρ − R_k x_k), 0}
                ▷ max{} is the elementwise maximum operator
        Set τ ← τ + 1
    end for
end for
```

Next, by using the previous compact form, Algorithm 1 can be rewritten as Algorithm A.1. In Algorithm A.1, Phase I is omitted because it can be seen as part of the input of the algorithm, $\boldsymbol{\mu}_i$ is the dual vector used at round i, and $\boldsymbol{z}_\tau$ is used to record the update process of the dual vector. Algorithm A.1 shows that at each round i we first compute the solution $\boldsymbol{x}_i$ of the ith request $(\boldsymbol{c}_i, \boldsymbol{R}_i, \boldsymbol{B}_i)$. Then we find a set of delayed gradients $\mathcal{F}_i$ (which can be empty), where $\mathcal{F}_i = \{k \in \mathbb{Z}^+ | k + d_k = i\}$ is the set of indices of gradients received during the rounds i and $i+1$ and d_k is the arbitrary delay of the request k. After receiving $\mathcal{F}_i$, for each request $k \in \mathcal{F}_i$, we update the dual vector $\boldsymbol{z}_\tau$ and set $\tau \leftarrow \tau + 1$. The repeated decisions are omitted in the previous algorithm because they do not affect the final decision or the dual value $\boldsymbol{z}_\tau$. The indices $i = 1, \ldots, N$ represent the requests that generate final decisions. Moreover, in practice, the order of requests in $\mathcal{F}_i$ is the order in which the consolidation decisions are fixed, whereas in the theoretical analysis, the order does not matter and can be arbitrary. Note that $\mathcal{F}_i$ can be empty in some rounds but contain several elements in other rounds. Furthermore, if $\mathcal{F}_i$ not received by round N are taken into account, the range of τ in Line 7 coincides with that of i, which spans from 1 to N.

In the following, we will convert Problem (A.1) into an allocation problem and Algorithm A.1 into an online allocation algorithm, which can be found in Algorithm A.2. Assume that the feasible points of the set partitioning constraints $\boldsymbol{B}_i\boldsymbol{x}_i = \mathbf{1}$, $\boldsymbol{x}_i \in \{0,1\}^{|\mathcal{H}_i|}$ are $\boldsymbol{x}_i^{(1)}, \boldsymbol{x}_i^{(2)}, \dots, \boldsymbol{x}_i^{(J_i)}$; then we have

$$\left\{ \boldsymbol{x}_i \middle| \begin{array}{l} \boldsymbol{B}_i\boldsymbol{x}_i = \mathbf{1} \\ \boldsymbol{x}_i \in \{0,1\}^{|\mathcal{H}_i|} \end{array} \right\} \Longleftrightarrow \left\{ \boldsymbol{x}_i \middle| \begin{array}{l} \boldsymbol{x}_i = \left(\boldsymbol{x}_i^{(1)}, \boldsymbol{x}_i^{(2)}, \dots, \boldsymbol{x}_i^{(J_i)}\right)\boldsymbol{y}_i \\ \mathbf{1}^\top \boldsymbol{y}_i = 1,\ \boldsymbol{y}_i \in \{0,1\}^{J_i} \end{array} \right\}, \tag{A.2}$$

where $\boldsymbol{y}_i$ is an J_i-dimensional one-hot vector and its bth dimension equal to one means that $\boldsymbol{x}_i^{(b)}$ is chosen as the solution $\boldsymbol{x}_i$. By substituting Formula (A.2) into Problem (A.1) and subtracting $\boldsymbol{x}_i$, Problem (A.1) is equivalent to

$$\begin{aligned} &\max_{\boldsymbol{y}_i} \sum_{i=1}^{N} \boldsymbol{v}_i^\top \boldsymbol{y}_i \\ &\text{s.t.} \sum_{i=1}^{N} (\boldsymbol{A}_i\boldsymbol{y}_i - \boldsymbol{\rho}) \leq \mathbf{0} \\ &\qquad \mathbf{1}^\top \boldsymbol{y}_i = 1,\ \boldsymbol{y}_i \in \{0,1\}^{J_i},\quad \forall i = 1, \dots, N, \end{aligned} \tag{A.3}$$

where $\boldsymbol{v}_i = \boldsymbol{X}_i^\top \boldsymbol{c}_i$, $\boldsymbol{A}_i = \boldsymbol{R}_i\boldsymbol{X}_i$, and $\boldsymbol{X}_i = (\boldsymbol{x}_i^{(1)}, \boldsymbol{x}_i^{(2)}, \dots, \boldsymbol{x}_i^{(J_i)})$. Problem (A.3) is an allocation problem and Algorithm A.1 is converted into Algorithm A.2, which is called the delayed online allocation algorithm in this paper. For the convenience of theoretical analysis, we record $w_\tau = k$ in Algorithm A.2.

Algorithm A.2 (Delayed Online Allocation Algorithm)

1: Input η and initialize $\boldsymbol{z}_1 = \mathbf{0}$ and $\tau = 1$
2: **for** $i = 1, \dots, N$ **do**
3: Set $\boldsymbol{\mu}_i = \boldsymbol{z}_\tau$
4: Compute $\boldsymbol{A}_i\boldsymbol{y}_i$ and $\boldsymbol{v}_i^\top\boldsymbol{y}_i$ where $\boldsymbol{y}_i = \arg\sup_{\mathbf{1}^\top \boldsymbol{y}=1, \boldsymbol{y}\geq \mathbf{0}} \{-\boldsymbol{\mu}_i^\top \boldsymbol{A}_i\boldsymbol{y} + \boldsymbol{v}_i^\top \boldsymbol{y}\}$
5: Receive $\{\boldsymbol{A}_k\boldsymbol{y}_k \mid k \in \mathcal{F}_i\}$
6: **for** $k \in \mathcal{F}_i$ **do**
7: Compute $\boldsymbol{z}_{\tau+1} = \max\{\boldsymbol{z}_\tau - \eta(\boldsymbol{\rho} - \boldsymbol{A}_k\boldsymbol{y}_k), \mathbf{0}\}$ ▷ max{} is the elementwise maximum operator
8: Set $\tau \leftarrow \tau + 1$ and record $w_\tau = k$
9: **end for**
10: **end for**

We note that $\boldsymbol{X}_i$ is unknown, but this does not matter. We only need to generate $\boldsymbol{A}_i\boldsymbol{y}_i$ and $\boldsymbol{v}_i^\top\boldsymbol{y}_i$ for theoretical analysis, both of which are available by using $\boldsymbol{A}_i\boldsymbol{y}_i = \boldsymbol{R}_i\boldsymbol{x}_i$ and $\boldsymbol{v}_i^\top\boldsymbol{y}_i = \boldsymbol{c}_i^\top\boldsymbol{x}_i$.

Based on the previous discussion, we develop the following proposition.

Proposition A.1. *For any $\{\boldsymbol{c}_i, \boldsymbol{R}_i, \boldsymbol{B}_i\}_{i=1}^N$, Algorithms A.1 and A.2 have the same optimality gap and constraint violation, that is,*

$$Q^* - \sum_{i=1}^{N} \boldsymbol{c}_i^\top \boldsymbol{x}_i = Q^* - \sum_{i=1}^{N} \boldsymbol{v}_i^\top \boldsymbol{y}_i, \tag{A.4}$$

$$\left(\sum_{i=1}^{N} (\boldsymbol{R}_i\boldsymbol{x}_i - \boldsymbol{\rho})\right)^+ = \left(\sum_{i=1}^{N} (\boldsymbol{A}_i\boldsymbol{y}_i - \boldsymbol{\rho})\right)^+. \tag{A.5}$$

where $\{\boldsymbol{x}_i\}_{i=1}^N$ and $\{\boldsymbol{y}_i\}_{i=1}^N$ are computed by Algorithms A.1 and A.2, respectively, Q^* is the optimal objective value of Problems (A.1) and (A.3), and $\{\boldsymbol{v}_i, \boldsymbol{A}_i\}_{i=1}^N$ is generated from $\{\boldsymbol{c}_i, \boldsymbol{R}_i, \boldsymbol{B}_i\}_{i=1}^N$ by using $\boldsymbol{v}_i = \boldsymbol{X}_i^\top \boldsymbol{c}_i$, $\boldsymbol{A}_i = \boldsymbol{R}_i\boldsymbol{X}_i$ and $\boldsymbol{X}_i = (\boldsymbol{x}_i^{(1)}, \boldsymbol{x}_i^{(2)}, \dots, \boldsymbol{x}_i^{(J_i)})$.

A.2. Delayed Online Allocation

Consider the following linear relaxation of Problem (A.3):

$$\begin{aligned} &\max_{\boldsymbol{y}_i} \sum_{i=1}^{N} \boldsymbol{v}_i^\top \boldsymbol{y}_i \\ &\text{s.t.} \sum_{i=1}^{N} (\boldsymbol{A}_i\boldsymbol{y}_i - \boldsymbol{\rho}) \leq \mathbf{0} \\ &\qquad \mathbf{1}^\top \boldsymbol{y}_i = 1,\ \boldsymbol{y}_i \geq \mathbf{0},\quad \forall i = 1, \dots, N. \end{aligned} \tag{A.6}$$

Similar to the previous work (Balseiro et al. 2020, Li et al. 2023), we consider the following boundedness and i.i.d. assumptions. For convenience, we make the assumptions directly on Problem (A.6) instead of the original Problem (A.1).

Assumption A.1. *We assume that*

(a) *There exist upper bounds of the coefficients such that $\|\boldsymbol{v}_i\|_\infty \leq \bar{r}$ and $\|\boldsymbol{A}_i\|_\infty \leq \bar{a}$, $\forall i = 1, \dots, N$;*
(b) *The average of delays $\bar{d} = \frac{1}{N}\sum_{i=1}^{N} d_i$ is finite;*
(c) *There exists an upper bound of $\boldsymbol{\rho}$ such that $\|\boldsymbol{\rho}\|_\infty \leq \bar{\rho}$;*
(d) *$(\boldsymbol{v}_i, \boldsymbol{A}_i)$s are i.i.d. sampled from an unknown distribution $\mathcal{P}$;*
(e) *Problem (A.3) is feasible. (Thus, Problem (A.6) is also feasible.)*

In the sections that follow, we will present the theoretical analysis of the delayed online allocation algorithm, which shows that its expected optimality gap and constraint violation are both on the order of $\sqrt{\bar{d}}N$.

A.2.1. Preliminaries for the Dual Problem. Let $\boldsymbol{\mu}_i$ be an approximation of the dual vector associated with Constraint $\sum_{i=1}^{N}(\boldsymbol{A}_i\boldsymbol{y}_i - \boldsymbol{\rho}) \leq \mathbf{0}$ in Problem (A.6). Considering the dual counterpart of Problem (A.6):

$$\min_{\boldsymbol{\mu}\geq \mathbf{0}} \sum_{i=1}^{N} \phi_i(\boldsymbol{\mu}), \tag{A.7}$$

where $\phi_i(\boldsymbol{\mu}) = \sup_{\mathbf{1}^\top \boldsymbol{y}=1, \boldsymbol{y}\geq \mathbf{0}} \{-\boldsymbol{\mu}^\top \boldsymbol{A}_i\boldsymbol{y} + \boldsymbol{v}_i^\top \boldsymbol{y}\} + \boldsymbol{\rho}^\top \boldsymbol{\mu}$ and its gradient $\nabla\phi_i(\boldsymbol{\mu})$ is $\boldsymbol{\rho} - \boldsymbol{A}_i\boldsymbol{y}_i$, where $\boldsymbol{y}_i = \arg\sup_{\mathbf{1}^\top \boldsymbol{y}=1, \boldsymbol{y}\geq \mathbf{0}} \{-\boldsymbol{\mu}^\top \boldsymbol{A}_i\boldsymbol{y} + \boldsymbol{v}_i^\top \boldsymbol{y}\}$.

Moreover, we define the expectation form of Problem (A.7):

$$\min_{\boldsymbol{\mu}\geq \mathbf{0}} \mathbb{E}_{\mathcal{P}}[\phi_i(\boldsymbol{\mu})]. \tag{A.8}$$

The optimal solutions of Problems (A.7) and (A.8) are denoted by $\hat{\boldsymbol{\mu}}$ and $\boldsymbol{\mu}^*$, respectively. From these definitions, obviously, the following two inequalities hold.

$$\mathbb{E}_{\mathcal{P}}[\phi_i(\boldsymbol{\mu}^*)] \leq \mathbb{E}_{\mathcal{P}}[\phi_i(\boldsymbol{\mu})],\ \forall \boldsymbol{\mu} \geq \mathbf{0} \text{ independent of } (\boldsymbol{v}_i, \boldsymbol{A}_i). \tag{A.9}$$

$$\sum_{i=1}^{N} \phi_i(\hat{\boldsymbol{\mu}}) \leq \sum_{i=1}^{N} \phi_i(\boldsymbol{\mu}^*),\ \forall \{\boldsymbol{v}_i, \boldsymbol{A}_i\},\ i = 1, \dots, N. \tag{A.10}$$

A.2.2. Boundedness of the Dual Vector. Before deriving the bounds of optimality gap and constraint violation, we need to prove that the dual vector z_{N+1} is bounded in expectation, that is, Proposition A.2.

Remark 1. In Algorithm A.2. there could exist some gradients that arrive after round N. Although these gradients are not actually used, they are useful to facilitate our analysis. In the analysis, we virtually let Algorithm A.2 continue to perform Lines 5–9 at additional rounds $i = N+1, \ldots, N+\overline{d}-1$ without considering new requests $(\boldsymbol{c}_i, \boldsymbol{R}_i, \boldsymbol{B}_i)$ that occurs at these rounds. That is, all gradients are used and the sequence $\{z_\tau\}_{\tau=2}^{N+1}$ is generated.

Proposition A.2. *Under Assumption* A.1, z_{N+1} *generated by Algorithm* A.2 *is bounded in expectation*:

$$\mathbb{E}_{\mathcal{P}}[\|z_{N+1} - \boldsymbol{\mu}^*\|_2^2] \le \|\boldsymbol{\mu}^*\|_2^2 + \eta^2 m(\overline{\rho} + \overline{a})^2 N(2\overline{d} + 1).$$

Here $\boldsymbol{\mu}^*$ *is the optimal solution of Problem* (A.8), m *is the dimension of* $\sum_{i=1}^N (A_i \boldsymbol{y}_i - \boldsymbol{\rho}) \le \boldsymbol{0}$ *and* η *is the step size.* $\overline{\rho}$, $\overline{a}$, *and* $\overline{d}$ *are defined in* Assumption A.1.

Proof. Our proof starts with Equation (A.9). Recalling that $\boldsymbol{\mu}_i$ is computed before the ith request $(\boldsymbol{v}_i, \boldsymbol{A}_i)$ arrives, $(\boldsymbol{v}_i, \boldsymbol{A}_i)$ is independent of $\boldsymbol{\mu}_i$. Thus, it holds that

$$\mathbb{E}_{\mathcal{P}}[\phi_i(\boldsymbol{\mu}^*)] - \mathbb{E}_{\mathcal{P}}[\phi_i(\boldsymbol{\mu}_i)] \le 0, \quad \forall i = 1, \ldots, N. \tag{A.11}$$

Next, we have

$$\begin{aligned}
0 &\le \mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^N (\phi_i(\boldsymbol{\mu}_i) - \phi_i(\boldsymbol{\mu}^*))\right] \\
&\le \mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^N \nabla\phi_i(\boldsymbol{\mu}_i)^\top (\boldsymbol{\mu}_i - \boldsymbol{\mu}^*)\right] \\
&= \mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^N \nabla\phi_{w_\tau}(\boldsymbol{\mu}_{w_\tau})^\top (\boldsymbol{\mu}_{w_\tau} - \boldsymbol{\mu}^*)\right] \\
&= \mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^N \nabla\phi_{w_\tau}(\boldsymbol{\mu}_{w_\tau})^\top (z_\tau - \boldsymbol{\mu}^*)\right] \\
&\quad + \mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^N \nabla\phi_{w_\tau}(\boldsymbol{\mu}_{w_\tau})^\top (\boldsymbol{\mu}_{w_\tau} - z_\tau)\right] \\
&\le \mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^N \nabla\phi_{w_\tau}(\boldsymbol{\mu}_{w_\tau})^\top (z_\tau - \boldsymbol{\mu}^*)\right] \\
&\quad + \mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^N \sqrt{m}(\overline{\rho} + \overline{a})\|\boldsymbol{\mu}_{w_\tau} - z_\tau\|_2\right],
\end{aligned} \tag{A.12}$$

where the first inequality is due to the convexity of $\phi_i(\cdot)$, the first equality is because $w_1, \ldots, w_N$ is a permutation of $1, \ldots, N$, the last inequality comes from $\nabla\phi_{w_\tau}(\boldsymbol{\mu}_{w_\tau}) = \boldsymbol{\rho} - A_{w_\tau}\boldsymbol{y}_{w_\tau}$ and Assumption A.1, and m is the dimension of $\boldsymbol{\rho}$.

Next, we proceed to give the upper bounds of the two terms in the right-hand side of Formula (A.12). The following derivation is inspired by Wan et al. (2023).

For the first term in the right-hand side of Formula (A.12), we have

$$\begin{aligned}
\sum_{\tau=1}^N \nabla\phi_{w_\tau}(\boldsymbol{\mu}_{w_\tau})^\top (z_\tau - \boldsymbol{\mu}^*) &= \sum_{\tau=1}^N (\boldsymbol{\rho} - A_{w_\tau}\boldsymbol{y}_{w_\tau})^\top (z_\tau - \boldsymbol{\mu}^*) \\
&= \sum_{\tau=1}^N \frac{1}{\eta}(\eta(\boldsymbol{\rho} - A_{w_\tau}\boldsymbol{y}_{w_\tau}) + \boldsymbol{\mu}^* - \boldsymbol{\mu}^*)^\top (z_\tau - \boldsymbol{\mu}^*) \\
&= \sum_{\tau=1}^N \frac{1}{2\eta}(\eta^2\|\boldsymbol{\rho} - A_{w_\tau}\boldsymbol{y}_{w_\tau}\|_2^2 + \|z_\tau - \boldsymbol{\mu}^*\|_2^2 \\
&\quad - \|z_\tau - \eta(\boldsymbol{\rho} - A_{w_\tau}\boldsymbol{y}_{w_\tau}) - \boldsymbol{\mu}^*\|_2^2) \\
&\le \sum_{\tau=1}^N \frac{1}{2\eta}(\eta^2\|\boldsymbol{\rho} - A_{w_\tau}\boldsymbol{y}_{w_\tau}\|_2^2 + \|z_\tau - \boldsymbol{\mu}^*\|_2^2 - \|z_{\tau+1} - \boldsymbol{\mu}^*\|_2^2) \\
&\le \frac{\eta N m(\overline{\rho} + \overline{a})^2}{2} + \frac{1}{2\eta}(\|\boldsymbol{\mu}^*\|_2^2 - \|z_{N+1} - \boldsymbol{\mu}^*\|_2^2)
\end{aligned} \tag{A.13}$$

where the second equality is because $(a-b)(c-b) = \frac{1}{2}((a-b)^2 + (c-b)^2 - (c-a)^2)$, the first inequality is due to $\boldsymbol{\mu}^* \ge 0$ and Line 7 in Algorithm A.2 $z_{\tau+1} = \max\{z_\tau - \eta(\boldsymbol{\rho} - A_k \boldsymbol{y}_k), \boldsymbol{0}\}$ and $k = w_\tau$, and the last inequality is due to Assumption A.1 and $z_1 = \boldsymbol{0}$.

For the second term in the right-hand side of Formula (A.12), we derive

$$\begin{aligned}
\sum_{\tau=1}^N \|\boldsymbol{\mu}_{w_\tau} - z_\tau\|_2 &= \sum_{\tau=1}^N \|z_{1+\sum_{s=1}^{w_\tau - 1}|\mathcal{F}_s|} - z_\tau\|_2 \\
&\le \sum_{\tau=1}^N \sum_{\tau'=1+\sum_{s=1}^{w_\tau-1}|\mathcal{F}_s|}^{\tau-1} \|z_{\tau'} - z_{\tau'+1}\|_2 \\
&\le \sum_{\tau=1}^N \sum_{\tau'=1+\sum_{s=1}^{w_\tau-1}|\mathcal{F}_s|}^{\tau-1} \|\eta(\boldsymbol{\rho} - A_{w_{\tau'}}\boldsymbol{y}_{w_{\tau'}})\|_2 \\
&\le \eta\sqrt{m}(\overline{\rho} + \overline{a})\sum_{\tau=1}^N \left(\tau - 1 - \sum_{s=1}^{w_\tau - 1}|\mathcal{F}_s|\right) \\
&= \eta\sqrt{m}(\overline{\rho} + \overline{a})\sum_{\tau=1}^N \left(\tau - 1 - \sum_{s=1}^{\tau-1}|\mathcal{F}_s|\right) \\
&\le \eta\sqrt{m}(\overline{\rho} + \overline{a})\sum_{i=1}^N d_i = \eta\sqrt{m}(\overline{\rho} + \overline{a})N\overline{d},
\end{aligned} \tag{A.14}$$

where the first equality is because z has been updated $\sum_{s=1}^{w_\tau-1}|\mathcal{F}_s|$ times at the beginning of period w_τ, the second inequality is from the Line 7 in Algorithm A.2, the third inequality comes from Assumption A.1, and the second equality is because $w_1, \ldots, w_N$ is a permutation of $1, \ldots, N$. The last inequality is because the definition $\mathcal{F}_s = \{k \in \mathbb{Z}^+ \mid k + d_k = s\}$. More clearly, $\tau - 1 - \sum_{s=1}^{\tau-1}|\mathcal{F}_s|$ counts the number of $A_k\boldsymbol{y}_k$ that have not been received at the start of round τ and each $A_i\boldsymbol{y}_i$ will only be counted as not received at most d_i

times; therefore, $\sum_{\tau=1}^{N}(\tau-1-\sum_{s=1}^{\tau-1}|\mathcal{F}_s|)\le\sum_{i=1}^{N}d_i$. The last equality is from the definition of the average $\bar{d}$.

By substituting Formulas (A.13) and (A.14) into (A.12), we derive

$$0\le\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}(\phi_i(\boldsymbol{\mu}_i)-\phi_i(\boldsymbol{\mu}^*))\right]\le\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\nabla\phi_i(\boldsymbol{\mu}_i)^\top(\boldsymbol{\mu}_i-\boldsymbol{\mu}^*)\right]$$
$$\le\frac{1}{2\eta}(\|\boldsymbol{\mu}^*\|_2^2-\mathbb{E}[\|\boldsymbol{z}_{N+1}-\boldsymbol{\mu}^*\|_2^2])+\eta m(\bar{\rho}+\bar{a})^2N\left(\bar{d}+\frac{1}{2}\right).\tag{A.15}$$

Finally, by moving $\|\boldsymbol{z}_{N+1}-\boldsymbol{\mu}^*\|$ term to the left side, we obtain

$$\mathbb{E}_{\mathcal{P}}[\|\boldsymbol{z}_{N+1}-\boldsymbol{\mu}^*\|_2^2]\le\|\boldsymbol{\mu}^*\|_2^2+\eta^2m(\bar{\rho}+\bar{a})^2N(2\bar{d}+1).\tag{A.16}$$

This concludes the proof. □

A.2.3. Regret and Constraint Violation. Benefiting from the analysis of the boundedness of the dual vector $\boldsymbol{z}_{N+1}$, we can easily obtain the bounds for the expected optimality gap and constraint violation.

Theorem A.1. *Under Assumption A.1, if $\eta=1/\sqrt{\bar{d}N}$, the expected optimality gap and constraint violation of Algorithm A.2 are $O(\sqrt{\bar{d}N})$, that is,*

$$\mathbb{E}_{\mathcal{P}}\left[\boldsymbol{Q}^*-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i\right]\le\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i^*-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i\right]\le O\left(\sqrt{\bar{d}N}\right),$$

$$\mathbb{E}_{\mathcal{P}}\left[\left\|\left(\sum_{i=1}^{N}(A_i\boldsymbol{y}_i-\boldsymbol{\rho})\right)^+\right\|_2\right]\le O\left(\sqrt{\bar{d}N}\right),$$

where $\boldsymbol{Q}^*$ is the optimal objective value of Problem (A.3), $\{\boldsymbol{y}_i\}_{i=1}^N$ is the output of Algorithm A.2, $(\cdot)^+$ denotes the positive part function, and $\{\boldsymbol{y}_i^*\}_{i=1}^N$ is the optimal solution of Problem (A.6).

Proof. For the expected constraint violation, we have

$$\mathbb{E}_{\mathcal{P}}\left[\left\|\left(\sum_{i=1}^{N}(A_i\boldsymbol{y}_i-\boldsymbol{\rho})\right)^+\right\|_2\right]=\mathbb{E}_{\mathcal{P}}\left[\left\|\left(\sum_{\tau=1}^{N}(A_{w_\tau}\boldsymbol{y}_{w_\tau}-\boldsymbol{\rho})\right)^+\right\|_2\right]$$
$$\le\frac{1}{\eta}\mathbb{E}_{\mathcal{P}}\left[\left\|\left(\sum_{\tau=1}^{N}(\boldsymbol{z}_{\tau+1}-\boldsymbol{z}_\tau)\right)^+\right\|_2\right]=\frac{1}{\eta}\mathbb{E}_{\mathcal{P}}[\|\boldsymbol{z}_{N+1}\|_2]$$
$$\le\left(\sqrt{\|\boldsymbol{\mu}^*\|_2^2+\frac{1}{\bar{d}N}m(\bar{\rho}+\bar{a})^2N(2\bar{d}+1)}+\|\boldsymbol{\mu}^*\|_2\right)\sqrt{\bar{d}N}$$
$$=O(\sqrt{\bar{d}N}),\tag{A.17}$$

where the first inequality comes from Line 7 in Algorithm A.2, the first equality is because $w_1,\dots,w_N$ is a permutation of $1,\dots,N$, the second equality is due to $\boldsymbol{z}_1=\boldsymbol{0}$, and the last inequality is from Proposition A.2 and $\eta=1/\sqrt{\bar{d}N}$.

For the optimality gap, obviously, $\mathbb{E}_{\mathcal{P}}[\boldsymbol{Q}^*-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i]\le\mathbb{E}_{\mathcal{P}}[\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i^*-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i]$ holds. Then, for $\mathbb{E}_{\mathcal{P}}[\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i^*-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i]$ we have

$$\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i^*-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i\right]=\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\phi_i(\hat{\boldsymbol{\mu}})-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i\right]$$
$$\le\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\phi_i(\boldsymbol{\mu}^*)-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i\right]$$
$$=\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\phi_i(\boldsymbol{\mu}^*)-\sum_{i=1}^{N}\phi_i(\boldsymbol{\mu}_i)\right]+\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\phi_i(\boldsymbol{\mu}_i)-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i\right]$$
$$\le\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\phi_i(\boldsymbol{\mu}_i)-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i\right],\tag{A.18}$$

where the first equality is due to the strong duality of Problem (A.6), the first inequality is from Formula (A.10), and the second inequality comes from Formula (A.9).

Next, by substituting $\phi_i(\boldsymbol{\mu}_i)=-\boldsymbol{\mu}_i^\top A_i\boldsymbol{y}_i+\boldsymbol{v}_i^\top\boldsymbol{y}_i+\boldsymbol{\rho}^\top\boldsymbol{\mu}_i$, the right-hand side of Formula (A.18) becomes

$$\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}\phi_i(\boldsymbol{\mu}_i)-\sum_{i=1}^{N}\boldsymbol{v}_i^\top\boldsymbol{y}_i\right]=\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}(\boldsymbol{\rho}-A_i\boldsymbol{y}_i)^\top\boldsymbol{\mu}_i\right].\tag{A.19}$$

Then, by performing a similar reformulation to (A.12), we derive

$$\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N}(\boldsymbol{\rho}-A_i\boldsymbol{y}_i)^\top\boldsymbol{\mu}_i\right]=\mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^{N}(\boldsymbol{\rho}-A_{w_\tau}\boldsymbol{y}_{w_\tau})^\top\boldsymbol{\mu}_{w_\tau}\right]$$
$$=\mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^{N}(\boldsymbol{\rho}-A_{w_\tau}\boldsymbol{y}_{w_\tau})^\top\boldsymbol{z}_\tau\right]$$
$$+\mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^{N}(\boldsymbol{\rho}-A_{w_\tau}\boldsymbol{y}_{w_\tau})^\top(\boldsymbol{\mu}_{w_\tau}-\boldsymbol{z}_\tau)\right]$$
$$\le\mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^{N}(\boldsymbol{\rho}-A_{w_\tau}\boldsymbol{y}_{w_\tau})^\top\boldsymbol{z}_\tau\right]$$
$$+\sqrt{m}(\bar{\rho}+\bar{a})\mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^{N}\|\boldsymbol{\mu}_{w_\tau}-\boldsymbol{z}_\tau\|_2\right].\tag{A.20}$$

To bound the first term in the right-hand side of Formula (A.20), we use the dual update formula in Line 7 in Algorithm A.2 again. From the update formula, we know

$$\|\boldsymbol{z}_{\tau+1}\|_2^2\le\|\boldsymbol{z}_\tau\|_2^2+\eta^2\|\boldsymbol{\rho}-A_{w_\tau}\boldsymbol{y}_{w_\tau}\|_2^2-2\eta(\boldsymbol{\rho}-A_{w_\tau}\boldsymbol{y}_{w_\tau})^\top\boldsymbol{z}_\tau$$
$$\le\|\boldsymbol{z}_\tau\|_2^2+\eta^2m(\bar{\rho}+\bar{a})^2-2\eta(\boldsymbol{\rho}-A_{w_\tau}\boldsymbol{y}_{w_\tau})^\top\boldsymbol{z}_\tau.$$

By moving the cross-term to the right-hand side, we obtain

$$\mathbb{E}_{\mathcal{P}}\left[\sum_{\tau=1}^{N}(\boldsymbol{\rho}-A_{w_\tau}\boldsymbol{y}_{w_\tau})^\top\boldsymbol{z}_\tau\right]\le\frac{1}{2\eta}\sum_{\tau=1}^{N}(\|\boldsymbol{z}_\tau\|_2-\|\boldsymbol{z}_{\tau+1}\|_2)$$
$$+\frac{\eta mN}{2}(\bar{\rho}+\bar{a})^2\le\frac{\eta mN}{2}(\bar{\rho}+\bar{a})^2.\tag{A.21}$$

Finally, by substituting Formulas (A.14), (A.20), (A.21), and the step size $\eta=1/\sqrt{\bar{d}N}$ into (A.19), and substituting (A.19) into Formula (A.18), we obtain the desired result:

$$\mathbb{E}_{\mathcal{P}}\left[\sum_{i=1}^{N} \boldsymbol{v}_i^\top \boldsymbol{y}_i^* - \sum_{i=1}^{N} \boldsymbol{v}_i^\top \boldsymbol{y}_i\right] \le \frac{mN}{2\sqrt{dN}}(\overline{\rho} + \overline{a})^2 + m(\overline{\rho} + \overline{a})^2 \sqrt{dN} = O(\sqrt{dN}). \tag{A.22}$$

This completes the proof. □

A.3. Theoretical Bounds of Algorithm 1

Proposition A.1 state that Algorithms A.1 and A.2 have the same optimality gap and constraint violation, and Theorem A.1 prove that the regret and the expected constraint violation of Algorithm A.2 are $O(\sqrt{dN})$. By combining Proposition A.1 with Theorem A.1, we get the following Theorem A.2. For ease of description, we still use the symbols in the compact form (A.1) instead of the original Problem (1).

Theorem A.2. *Under Assumption A.1, if $\eta = 1/\sqrt{dN}$, the expected optimality gap and constraint violation of Algorithm 1 are $O(\sqrt{dN})$, that is,*

$$\mathbb{E}_{\mathcal{P}}\left[\boldsymbol{Q}^* - \sum_{i=1}^{N} \boldsymbol{c}_i^\top \boldsymbol{x}_i\right] \le O(\sqrt{dN}),$$

$$\mathbb{E}_{\mathcal{P}}\left[\left\| \left(\sum_{i=1}^{N} (\boldsymbol{R}_i \boldsymbol{x}_i - \boldsymbol{\rho})\right)^+ \right\|_2\right] \le O(\sqrt{dN}),$$

where $\boldsymbol{Q}^*$ is the optimal objective value of the offline Problem (A.1) (the compact form of Problem (1)), and $\{\boldsymbol{x}_i\}_{i=1}^N$ is the output of Algorithm A.1 (i.e., the final decisions in Algorithm 1), $(\cdot)^+$ denotes the positive part function.

References

Agrawal S, Wang Z, Ye Y (2014) A dynamic near-optimal algorithm for online linear programming. *Oper. Res.* 62(4):876–890.

Arlotto A, Gurvich I (2019) Uniformly bounded regret in the multisecretary problem. *Stochastic Systems* 9(3):231–260.

Asadpour A, Wang X, Zhang J (2020) Online resource allocation with limited flexibility. *Management Sci.* 66(2):642–666.

Balseiro S, Lu H, Mirrokni V (2020) Dual mirror descent for online allocation problems. Daumé H, Singh A, eds. *Proc. 37th Internat. Conf. Machine Learning* (JMLR.org), 613–628.

Cassandras CG, Lafortune S, eds. (2007) *Introduction to Discrete-Event Simulation* (Springer, Boston), 557–615.

Çetinkaya S (2005) *Coordination of Inventory and Shipment Consolidation Decisions: A Review of Premises, Models, and Justification* (Springer, Boston).

Chan LMA, Muriel A, Shen ZJM, Simchi-Levi D (2002a) On the effectiveness of zero-inventory-ordering policies for the economic lot-sizing model with a class of piecewise linear cost structures. *Oper. Res.* 50(6):1058–1067.

Chan LMA, Muriel A, Shen ZJM, Simchi-Levi D, Teo CP (2002b) Effective zero-inventory-ordering policies for the single-warehouse multiretailer problem with piecewise linear cost structures. *Management Sci.* 48(11):1446–1460.

Fontaine P, Minner S (2023) A branch-and-repair method for three-dimensional bin selection and packing in e-commerce. *Oper. Res.* 71(1):273–288.

Gillai B, Lee HL (2023) Exploring the right logistics strategy for China-U.S. B2C cross-border e-commerce. Technical report, Stanford Graduate School of Business, Stanford, CA.

Higginson JK, Bookbinder JH (1995) Markovian decision processes in shipment consolidation. *Transportation Sci.* 29(3):242–255.

Jasin S (2015) Performance of an LP-based control for revenue management with unknown demand parameters. *Oper. Res.* 63(4):909–915.

Kesselheim T, Tönnis A, Radke K, Vöcking B (2014) Primal beats dual on online packing LPs in the random-order model. *Proc. 46th Ann. ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 303–312.

Lee CY, Çetinkaya S, Jaruphongsa W (2003) A dynamic model for inventory lot sizing and outbound shipment scheduling at a third-party warehouse. *Oper. Res.* 51(5):735–747.

Li X, Ye Y (2021) Online linear programming: Dual convergence, new algorithms, and regret bounds. *Oper. Res.* 70(5):2948–2966.

Li X, Sun C, Ye Y (2023) Simple and fast algorithm for binary integer and online linear programming. *Math. Programming* 200(2):831–875.

Li F, Xu Z, Chen ZL (2020) Production and transportation integration for commit-to-delivery mode with general shipping costs. *INFORMS J. Comput.* 32(4):1012–1029.

Salari N, Liu S, Shen ZJM (2022) Real-time delivery time forecasting and promising in online retailing: When will your package arrive? *Manufacturing Service Oper. Management* 24(3):1421–1436.

Ülkü MA (2012) Dare to care: Shipment consolidation reduces not only costs, but also environmental damage. *Internat. J. Production Econom.* 139(2):438–446.

Wan Y, Yao C, Song M, Zhang L (2023) Non-stationary online convex optimization with arbitrary delays. Accessed January 29, 2024, https://arxiv.org/pdf/2305.12131v1.pdf.

Wei L, Kapuscinski R, Jasin S (2021) Shipping consolidation across two warehouses with delivery deadline and expedited options for e-commerce and omni-channel retailers. *Manufacturing Service Oper. Management* 23(6):1634–1650.

Xu Z, Li F, Chen ZL (2023) Shipment consolidation with multiple shipping methods under nonlinear cost structures. *Management Sci.*, ePub ahead of print June 12, https://doi.org/10.1287/mnsc.2023.4835.

Zhang Y, Sun L, Hu X, Zhao C (2019) Order consolidation for the last-mile split delivery in online retailing. *Transportation Res. Part E Logist. Transportation Rev.* 122(February):309–327.

Yujie Chen is a senior algorithm expert at the Artificial Intelligence Department of Cainiao Network. Her research interests include supply chain management and logistics optimization.

Biao Yuan is an assistant professor at Sino-US Global Logistics Institute, Shanghai Jiao Tong University. His research interests lie in developing solution approaches for optimization problems in logistics.

Yinzhi Zhou is an algorithm expert at the Artificial Intelligence Department of Cainiao Network. Her research interests include supply chain management and logistics optimization.

Yuwei Chen is a senior algorithm engineer at the Artificial Intelligence Department of Cainiao Network. His research interests include the development of online learning algorithms in logistics.

Haoyuan Hu is a director of the Artificial Intelligence Department of Cainiao Network. He was a finalist for the 2021 INFORMS Franz Edelman Award. His research interests include the integration of operations research and machine learning, large-scale optimization in logistics, financial engineering, and server scheduling.