The Stochastic Container Relocation Problem *

V. Galle † V. H. Manshadi ‡ S. Borjian Boroujeni § C. Barnhart ¶ P. Jaillet ‖

January 25, 2018

Abstract

The container relocation problem (CRP) is concerned with finding a sequence of moves of containers that minimizes the number of relocations needed to retrieve all containers, while respecting a given order of retrieval. However, the assumption of knowing the full retrieval order of containers is particularly unrealistic in real operations. This paper studies the stochastic CRP (SCRP), which relaxes this assumption. A new multistage stochastic model, called the *batch model*, is introduced, motivated, and compared with an existing model (the *online model*). The two main contributions are an optimal algorithm called *Pruning-Best-First-Search (PBFS)* and a randomized approximate algorithm called *PBFS-Approximate* with a bounded average error. Both algorithms, applicable in the batch and online models, are based on a new family of lower bounds for which we show some theoretical properties. Moreover, we introduce two new heuristics outperforming the best existing heuristics. Algorithms, bounds, and heuristics are tested in an extensive computational section. Finally, based on strong computational evidence, we conjecture the optimality of the "leveling" heuristic in a special "no information" case, where, at any retrieval stage, any of the remaining containers is equally likely to be retrieved next.

Introduction

With the growth in international container shipping in maritime ports, there has been an increasing interest in improving operations in container terminals, both on the sea side and land side. The operations on the sea side involve the assignment of quay cranes to ships, the loading of export containers on vessels, and the discharging of import containers from vessels onto internal trucks. Import containers are then transferred to the land side and are stacked in the storage area. Operations on the land side (also called yard operations) include the routing of internal trucks within the yard, the stacking of containers for storage, and the delivery of import containers to external trucks for delivery to another location. This work focuses on the latter problem.

Due to limited space in the storage area, containers are stacked on top of each other. The resulting stacks create rows of containers as shown in Figure 1. If a container that needs to be retrieved (*target*

*Accepted for publication in Transportation Science.

†MIT, Operations Research Center. Email: vgalle@mit.edu.

‡Yale School of Management. Email: vahideh.manshadi@yale.edu

§MIT, Operations Research Center. Email: setareh@alum.mit.edu

¶MIT, Operations Research Center and Civil & Environmental Engineering. Email: cbarnhar@mit.edu

‖MIT, Operations Research Center and Electrical Engineering & Computer Science. Email: jaillet@mit.edu

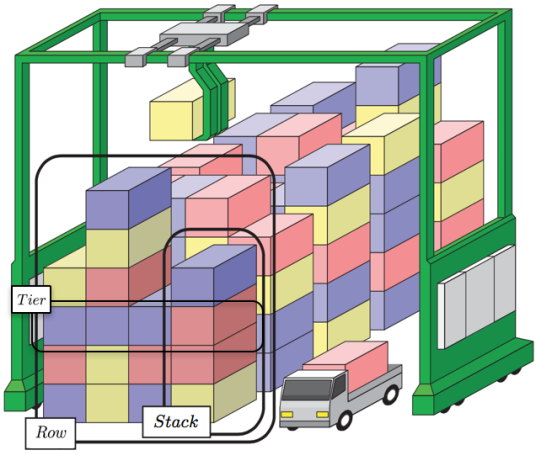

Figure 1: Illustration of stacks of containers in a storage yard (figure from Tanaka and Takii (2016)).

container) is not located at on a topmost tier and is covered by other containers, the blocking containers must be relocated to another stack. As a result, during the retrieval process, the yard cranes perform one or more *relocation* moves. Such relocations (also called reshuffles) are costly for the port operators and result in delays in the retrieval process. Thus, reducing the number of relocations is one of the main goals of port operators. The *container relocation problem (CRP)* (also known as the block relocation problem) addresses this challenge by minimizing the number of relocations. As this problem is the main subject of this paper, we provide a formal definition and an extensive literature review.

The container relocation problem. First, it is commonly known that the time to relocate a container within a row is insignificant compared to the time to relocate a container between two distinct rows. Therefore, in most cases, port operators tend to avoid relocations between rows. The CRP assumes that only relocations within rows are allowed, and problems for different rows should be considered independently. Furthermore, a row usually stores containers of the same type for the sake of stability and simplicity.

Using these facts, the CRP models one row using a two-dimensional array of size (T, S), where S is the number of stacks and T is the maximum height, i.e., the maximum number of containers in a stack as limited by the height of the crane. Each element of this array represents a potential slot for a container, and contains a number only if a container is currently stored in this slot. Stacks are numbered from left (1) to right (S) and tiers from bottom (1) to top (T). We refer to this array as a *configuration*. The common assumptions of the CRP are the following:

> **$\mathbf{A_1}$**: The initial configuration has T tiers, S stacks, and C containers. In order for the problem to always be feasible, we suppose that the triplet (T, S, C) satisfies $0 \leqslant C \leqslant ST - (T - 1)$ (see Caserta et al. (2012)).
>
> **$\mathbf{A_2}$**: A container can be retrieved/relocated only if it is in the topmost tier of its stack, i.e., no other container is blocking it.

A_3: A container can be relocated only if it is blocking the target container. This assumption was suggested by Caserta et al. (2012), and the problem under this assumption is commonly referred to as the *restricted* CRP. Most studies focus on this restricted version, because it is the current practice in many yards, and it helps decrease the dimensionality of the problem, while not losing much optimality (see Petering and Hussein (2013)). As is common practice, we will not mention the term "restricted" in the rest of the paper even though we always assume A_3.

A_4: The cost of relocating a container from a stack does not depend on the stack to which the container is relocated. This allows us to consider the stacks of a configuration as interchangeable. In addition, it motivates the objective of minimizing the number of relocations, since the cost of each relocation can be normalized to 1. Note that this assumption is not required for all the results stated below, hence our approaches could be easily extended to the case when Assumption A_4 does not hold.

A_5: The retrieval order of containers is known, so that each container can be labeled from 1 to C, representing the departure order, i.e., Container 1 is the first one to be retrieved, and C the last one.

The CRP involves finding a sequence of moves to retrieve Containers $1, 2, \ldots, C$ (respecting the order) with a minimum number of relocations. Figure 2 provides a simple example of the CRP. The CRP with the above classical assumptions is referred to as *static and full information*: "static" because no new containers arrive during the retrieval process (see Assumption A_1) and "full information" because we know the full retrieval order at the beginning of the retrieval process (see Assumption A_5). This problem was first formulated by Kim and Hong (2006) in a dynamic programming model.

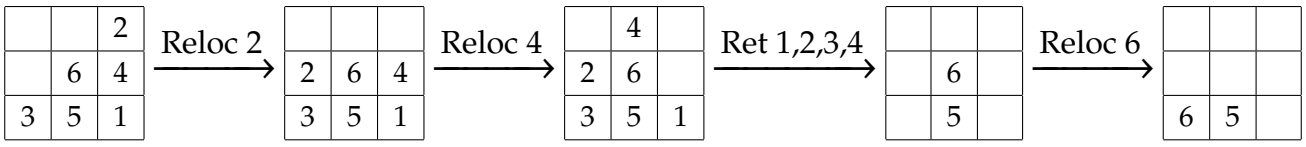

Figure 2: Configuration for the CRP with 3 tiers, 3 stacks, and 6 containers. The optimal solution performs 3 relocations.

Researchers have tackled the static CRP with full information from two points of view. The first approach aims to find the optimal solution. Primarily, researchers have used integer programming (IP) to address this problem. For example, Caserta et al. (2012) propose an intuitive formulation of the problem. Petering and Hussein (2013) develop a more tractable formulation that is, however, unable to solve real-sized instances efficiently. Zehendner et al. (2015) fix the formulation from Caserta et al. (2012) and improve it by removing some variables, tightening some constraints, introducing a new upper bound, and applying a preprocessing step to fix several variables. In all these IP formulations, due to the combinatorial nature of the problem, the number of variables and constraints dramatically increases as the size of the problem grows, and the IP cannot be solved for large instances. To bypass this problem, a recent trend has been to look at more efficient ways to explore the branch-and-bound tree, or even decrease its size using the structural properties of the problem. Ünlüyurt and Aydın (2012) and Expósito-Izquierdo et al. (2015) suggest two branch-and-bound approaches with several heuristics based on this idea. Another solution using the A^* algorithm is explored by Zhu et al. (2012), and built

upon by Tanaka and Takii (2016) and Borjian et al. (2015a). Another solution using branch-and-price is presented by Zehendner and Feillet (2014b).

As the problem is $\mathcal{NP}$-hard (Caserta et al. (2012)), an alternative approach is to use quick and efficient heuristics providing suboptimal solutions. For the sake of conciseness, we mention only some of them that are relevant to this paper. Caserta et al. (2012) introduce a "MinMax" policy that is defined and generalized later in this paper. Wu and Ting (2010) propose a beam search heuristic, and Wu and Ting (2012) develop the group assignment heuristic (GAH) also generalized in Section 2. Finally, in Kim and Hong (2006) and Zhu et al. (2012), lower bounds for the CRP are introduced: Kim and Hong (2006) count the number of blocking containers as a straightforward lower bound, and Zhu et al. (2012) refine this idea by taking into account additional unavoidable relocations using a family of lower bounds.

Finally, there are many problems related to the CRP. The stacking problem is concerned with how to store incoming containers in a configuration given an arrival order of containers. The pre-marshalling problem deals with rearranging the containers prior to the retrieval process in order to minimize future relocations, but no container is removed in this process. For both problems, general review and classification surveys of the existing literature on the CRP can be found in Stahlbock and Voß (2008), Steenken et al. (2004), and Lehnfeld and Knust (2014). In addition, Assumption A_1 can be relaxed, which leads to the dynamic CRP where stacking and retrieving are done simultaneously as new containers arrive. For this problem, see Borjian et al. (2015b) and Hakan Akyüz and Lee (2014).

Finally, the main focus of this paper is an extension of the CRP where the full information assumption (A_5) is relaxed. Indeed, Assumption A_5 is unrealistic given that arrival times of external trucks at the terminal are generally unpredictable due to uncertain conditions. Nevertheless, new technology advancements such as truck appointment systems (TASs) and GPS tracking can help predict relative truck arrival times. Thus, although the exact retrieval order might not be known, some information on trucks' arrival times might be available. This leads us to introduce a stochastic version of the CRP.

The stochastic CRP (SCRP). A common assumption is that, for each container, there is a time window during which a truck driver will arrive to retrieve it. We refer to a *batch of containers* as the set of containers stacked in the same row and with the same arrival time window. This information can be either inferred using machine learning algorithms, not yet much discussed in the literature, or obtained by using the appointment time windows in a TAS, which has gained attention over the last decade. The first TAS was implemented by Hong Kong International Terminals (HIT) in 1997. It uses 30-minute time slots, where trucks can register (Murty et al. (2005)). Another TAS was introduced in New Zealand in 2007. Two other studies, Giuliano and O'Brien (2007) and Morais and Lord (2006), evaluate the benefits of TAS, in reducing truck idling time by increasing on-time ratio. More recent information can be found in Phillips (2015) and Bonney (2015).

On the modeling side, Zehendner and Feillet (2014a) formulate an IP to get the optimal number of slots a TAS should offer for each batch. Very few studies have tackled the SCRP, also referred to as CRP with Time Windows. This problem was first modeled by Zhao and Goodchild (2010). In their original model, each container is assigned to a batch. Batches of containers are ordered such that all containers in a batch must be retrieved before any containers from a later batch are retrieved. Furthermore, the relative retrieval order of containers within a given batch is assumed to be a random permutation. From now on, we will refer to the model of Zhao and Goodchild (2010) as the *online model*. In Section 1, we discuss in

more detail how this model assumes information is revealed. For the online model, Zhao and Goodchild (2010) develop a myopic heuristic (called RDH) and study, in different settings with two or more groups, the value of information using RDH. They conclude that a small improvement in the information system reduces the number of relocations significantly. Van Asperen et al. (2013) use a simulation tool to evaluate the effect of a TAS on many statistics including the ratio of relocations to retrievals. Their decision rules are based on several heuristics including the "leveling," Random, or "traveling distance" heuristics. More recently, Ku and Arthanari (2016) also use the online model. They formulate the SCRP under the online model as a finite horizon dynamic programming problem and suggest a decision tree scheme to solve it optimally. They also introduce a new heuristic called expected reshuffling index (ERI), which outperforms RDH, and they perform computational experiments based on available test instances. We will refer frequently to this work and use some of their techniques as well as their available test instances to evaluate our algorithms.

In another recent study related to the SCRP, Zehendner et al. (2017) study the *online* container relocation problem, which corresponds to an adversarial model. They prove that the number of relocations performed by the leveling policy can be upper-bounded by a linear function of the number of blocking containers and provide a tight competitive ratio for this policy. Moreover, Galle et al. (2016) show that the ratio of the expected number of relocations to the expected blocking lower bound converges to one. Finally, Tierney and Voß (2016) study the robust pre-marshalling problem which also considers uncertainty in the retrieval times of containers.

For a general review of techniques on finite horizon dynamic programming, we refer the reader to Bertsekas (2005) and Sennott (2009).

Contributions of the paper. The contributions of this paper are:

1. A new stochastic model, referred to as the *batch model*. This new model uses the same probability distribution as the online model. However, the two models differ in the way each reveals new information on the retrieval order. The batch model is motivated, described, and compared with the online model.

2. Lower and upper bounds for the SCRP. We derive a **new family of lower bounds** for which we show theoretical properties. Furthermore, we develop **two new fast and efficient heuristics**.

3. **A novel optimal algorithm scheme based on decision trees and pruning strategies referred to as *Pruning-Best-First-Search* (*PBFS*)**, which takes advantage of the properties of the aforementioned lower bounds. The algorithm is explained with pseudocode in Algorithm 2.

4. **A second novel algorithm tuned for the case of larger batches referred to as *PBFS-Approximate* (*PBFSA*)**. We build upon *PBFS* and derive a sampling strategy resulting in an approximate algorithm with an expected error that we bound theoretically. The pseudocode of the second algorithm is presented in Algorithm 3.

5. **Extensive computational experiments using an existing set of instances**. The first experiment exhibits the efficiency of *PBFS*, our lower bounds and two new heuristics for the batch model based on existing instances, presented by Ku and Arthanari (2016), where batches of containers

are small (2 containers per batch on average). The second experiment illustrates the advantage of using *PBFSA* when batches of containers are larger, based on instances obtained by modifying the existing set. In addition, most of our techniques including lower bounds, heuristics, and the *PBFS* algorithm also apply to the online model. The third experiment shows that, in this model, *PBFS* outperforms the algorithm introduced in Ku and Arthanari (2016) in the sense that it is faster for instances that Ku and Arthanari (2016) could solve, and it can solve problems of larger size. Furthermore, our two new heuristics also outperform the best existing heuristic (ERI) for the online model. Finally, the last experiment is used to test the conjecture about the optimality of the leveling heuristic in the special case of the online model with a unique batch of containers.

The rest of the paper is structured as follows: Section 1 thoroughly describes the batch model, its assumptions and objective, the difference with the online model, and the general theory of decision trees applied to the SCRP. Section 2 gives a good intuition into the problem and defines heuristics and a class of lower bounds for the SCRP used in subsequent sections. Then Sections 3 and 4 introduce respectively the *PBFS* and *PBFSA* algorithms. Computational experiments for both batch and online models are carried through Section 5. We conclude the paper in Section 6 with a discussion of future directions for the SCRP.

1 SCRP and Decision Trees

1.1 Motivation

Before stating the general assumptions of the batch model, let us motivate our problem using a typical example. We consider a port with a TAS offering 30-minute time windows during which truck drivers who want to retrieve a container can register to arrive at the port. For the sake of the example, we consider the time window between 9:00 a.m. and 9:30 a.m. Multiple trucks can be registered in this time window; in this example, presented in Figure 3, 3 trucks (designated i_1, i_4 and i_6) are registered for this time window. We assume that **all 3 trucks arrive on time** (between 9:00 a.m. and 9:30 a.m.) and that their containers (similarly designated i_1, i_4, and i_6) form a batch to be retrieved. We display the configuration of interest in Figure 4 (3 tiers, 3 stacks, and 6 containers).

We assume that trucks arrive randomly within the time window, so **each truck arrival order is equally likely to happen**. In this example, there are 6 potential arrival orders, each with 1/6 chance of occurring.

At 9:00 a.m., none of the 3 trucks has arrived, and **their relative retrieval order is unknown**. Consequently, these 3 containers are all labeled 1 in Figure 4a. In Figure 4b, the IDs of all containers and their locations are depicted.

Between 9:00 a.m. and 9:30 a.m., trucks arrive in a particular order (*e.g.*, Truck i_4 first, then i_6, and i_1 last). In busy terminals, trucks typically queue up as they wait to be served. Their place in line is based on their arrival order, so the port operator generally retrieves containers based on the arrival order. Processing in this way, **on a first-come, first-served basis**, avoids issues with truck unions and maintains fairness among drivers. Consequently, we take the retrieval order to be exogenously determined and we do not consider it a potential decision for port operators.

To provide a specified level of service to the truck drivers, the terminal operator often sets a target average waiting time. If the appointment time window is about the same as or shorter than the target average waiting time, the operator has information about the retrieval order of containers in the batch

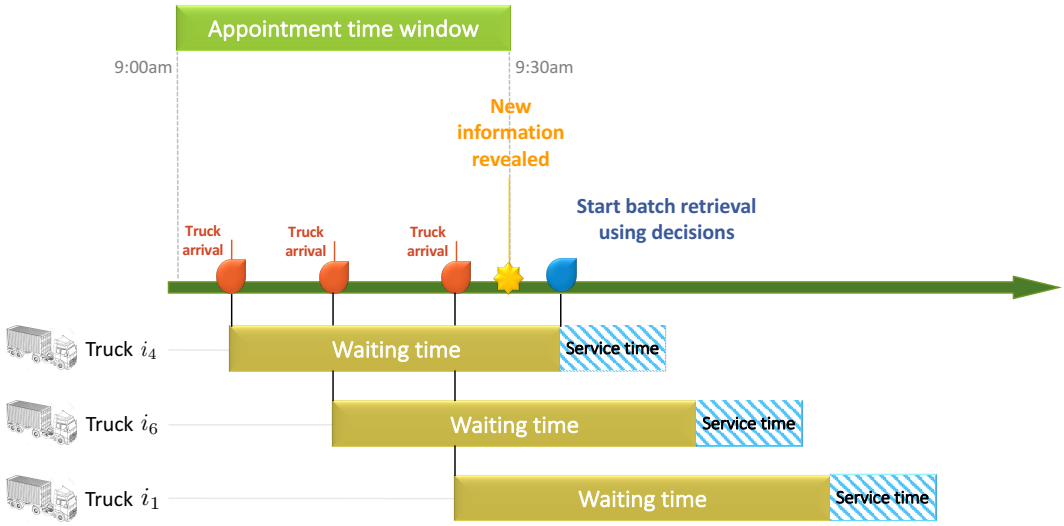

Figure 3: Timeline of events for the batch model with three trucks.

before the retrieval of those containers must begin in order to meet the target waiting time. Given these conditions, we make the simplifying assumption in this work that the retrieval of a batch begins at the end of the appointment time window associated with the batch, and **the retrieval order of all containers in the batch is known at the moment the retrieval of the batch commences**. In our example, the target average waiting time is 30 minutes. At 9:30 a.m., the retrieval order of the batch (i_4, i_6, i_1) is known and the retrieval of the batch commences soon after. The updated information is depicted in the configuration of Figure 4c.

The assumption that containers to be retrieved are revealed on a batch basis models the reality that port operators typically know information about all the containers in the same batch before starting to retrieve them. This is especially true for busy ports that have a TAS. Moreover, we assume that no information about future batches is available when making decisions for the current batch. Similar modeling assumptions have been made in previous work (see Zhao and Goodchild (2010) and Ku and Arthanari (2016)).

		1
	5	4
1	5	1

4a Before any truck has arrived (9:00 a.m.)

		i_6
	i_3	i_5
i_1	i_2	i_4

4b IDs to match containers with trucks

		2
	5	4
3	5	1

4c Before the first container gets retrieved (9:30 a.m.)

Figure 4: SCRP example. The left configuration is the input to our problem. The configuration in the middle denotes each container with an ID i_l where $l = 1, \ldots, 6$. The configuration on the right denotes the order of the first batch after it is revealed.

The general assumptions we apply to our model are formally stated in Section 1.2 (A_5^* and A_6^*) and result in the *batch model, the main focus of this paper*. The goal of the SCRP is to find a sequence of moves minimizing the expected number of relocations needed to empty the initial configuration.

Labels in Figures 4a and 4c are defined such that two containers have the same label only if they are in the same batch and their relative order is yet to be revealed. In our example, since Container i_5 is the only container in the second batch and is retrieved after the first 3 containers, it is necessarily the fourth container to be retrieved (thus labeled 4). Containers i_2 and i_3 are labeled 5, and when their relative order is revealed, one will be labeled 5 and the other 6.

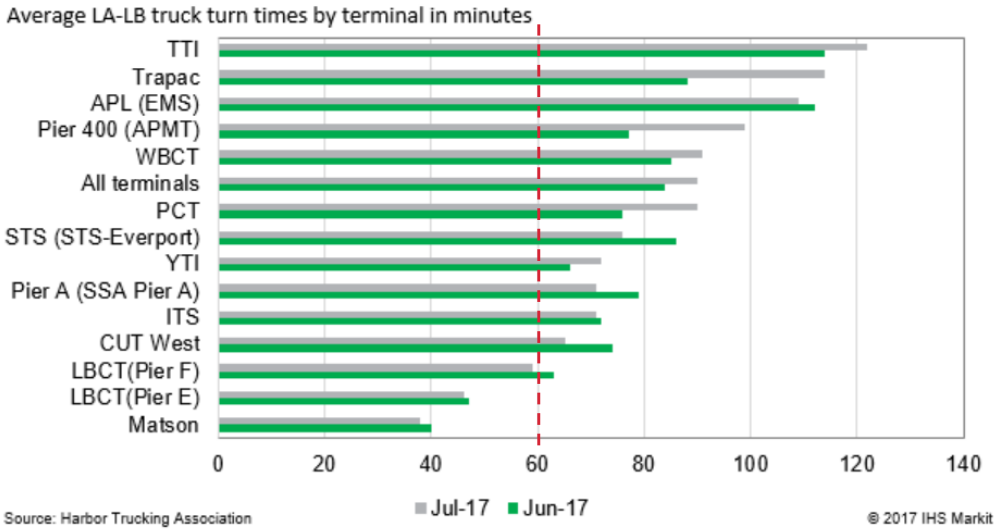

Figure 5: Average truck turn times in minutes by terminal at Los Angeles-Long Beach port in June and July 2017 (source: JOC.com). The dashed line shows the length of a time window (60 minutes) in the truck appointment system.

We mention that these assumptions apply to an increasing number of port terminals. For example, Figure 5 presents the average truck turn time in minutes for all 15 terminals of Los Angeles-Long Beach port (the largest U.S. port for containers) in June and July 2017. In many of these terminals, a truck appointment system with 60-minutes time windows (see Davies (2009)) has been implemented, and Figure 5 shows that most terminals have an average truck turn time longer than 60 minutes.

The online model. The main distinguishing difference between the batch model and the online model introduced by Zhao and Goodchild (2010) is the information revelation process. The online model disregards any within-batch information available when it plans the moves to retrieve a container. Hence **new information is revealed one container at a time.** The online model applies especially to less busy ports where the waiting time is significantly shorter than the appointment time windows. In this case, because there is a limited number of trucks waiting, limited information about the future is known. We show through Lemma 1 that ignoring information (if available) results in a potential loss of operational efficiency as measured by the expected number of relocations. In addition, most of our batch-based approaches also apply to the online model, and we provide numerical results based on it in Section 5.

1.2 Assumptions, Notations, and Formulation

To define our problem as a multistage stochastic optimization problem (the number of "stages" is the number of batches), we need to define *a probabilistic model of the container retrieval order, a scheme for revealing new information about this order, and an objective function.*

Batch model. Let us state the assumptions and objective of the stochastic CRP under the batch model. Assumptions A_1^*, A_2^*, A_3^*, and A_4^* are respectively identical to Assumptions A_1, A_2, A_3, and A_4.

> $\mathbf{A_5^*}$: (probabilistic model) Given an ordering of batches, the probability distribution of the retrieval order is such that: 1) **all the containers in a given batch are retrieved before any containers in a later batch**, and 2) **within each batch of containers, the order of the containers is drawn from a uniform random permutation**. This paper focuses mainly on this latter assumption, but our model can be extended to the more general case of any probability distribution on permutations (not necessarily uniform) that respects the order of batches.
>
> $\mathbf{A_6^*}$: (revelation of new information) For each batch w, **the full relative order of containers from the w^{th} batch (i.e., the specific random permutation drawn) is revealed after all containers in batch** 1 **through** $w-1$ **have been retrieved**.

Given these assumptions, we want to find the minimum expected number of relocations to retrieve all containers from a given initial configuration. We refer to this problem as the "*stochastic CRP*." Let us introduce some notations:

- The problem size is given by (T, S, C), respectively the number of tiers, stacks, and containers in the initial configuration.
- The number of batches of containers in the initial configuration is denoted by W. We consider that the batches are ordered from 1 to W.
- For each batch $w \in \{1, \dots, W\}$, let C_w be the number of containers in w. By definition $\sum_{w=1}^{W} C_w = C$.
- Each container has two attributes:
  - The first attribute, denoted by $(c_l)_{l\in\{1,\dots,C\}}$, is the label and is defined as follows: initially, containers in batch w are labeled K_w such that $K_w = 1 + \sum_{u=1}^{w-1} C_u$, where the sum is empty for $w = 1$. Then, for $k \in \{1, \dots, C\}$, if a container is revealed to be the k^{th} container to be retrieved, its label changes to k. Using this labeling, at any point in the retrieval process, two containers have the same label only if they are in the same batch and their relative order is yet to be revealed.
  - The second attribute is a unique ID denoted by $(i_l)_{l\in\{1,\dots,C\}}$. This ID is used only to identify uniquely containers in the initial configuration (see Figure 4b) and for the sake of clarity of the following probabilistic model. Note that for $l \in \{1, \dots, C\}$, if Container i_l is in batch w, then $c_l = K_w$ (until the actual retrieval order of i_l is revealed).

- For $k \in \{1, \dots, C\}$, let ζ_k be a random variable taking values in $(i_l)_{l\in\{1,\dots,C\}}$, such that $\{\zeta_k = i_l\}$ is the event that Container i_l is the k^{th} container to be retrieved. According to Assumption A_5^*, the distribution of $(\zeta_k)_{k\in\{1,\dots,C\}}$ is given by

$$\mathbb{P}[\zeta_k = i_l] = \begin{cases} \frac{1}{C_w} \text{ , if } w = \min\{u \in \{1, \dots, W\} \mid K_u \geqslant k\} \text{ and } c_l = K_w, \\ 0 \text{ , otherwise.} \end{cases}$$

  In this case, there are a total of $\prod_{w=1}^{W}(C_w!)$ orders with equal probabilities. More generally, we consider the case where the probability of each order within each batch is not necessarily equally likely. However, we still assume that the batches are ordered, thus $\mathbb{P}[\zeta_k = i_l]$ can be positive only if $w = \min\{u | K_u \geqslant k\}$ and $c_l = K_w$. In the practical case where probabilities are not considered to be uniform, a list of potential retrieval orders and their associated probability is given for each batch of containers (based on historical data), thus $\mathbb{P}[\zeta_k = i_l]$ can easily be inferred from these probabilities.

- An action corresponds to a sequence of relocations to retrieve one container. For $k \in \{1, \dots, C\}$, we denote the action for the k^{th} retrieval by a_k, and the feasible set of actions is defined according to Assumptions A_2^* and A_3^*.

- For a given batch $w \in \{1, \dots, W\}$,

  1. let y_w denote the configuration before the retrieval order of containers in batch w is revealed, i.e., before $\zeta_{K_w}, \dots, \zeta_{K_w+C_w-1}$ are revealed. Note that y_1 corresponds to the initial configuration. We denote x_{K_w} the configuration after the retrieval order of containers in batch w is revealed, and before action a_{K_w} is taken. If $\xrightarrow{\zeta_{K_w},\dots,\zeta_{K_w+C_w-1}}$ represents the revelation of the random variables $\zeta_{K_w}, \dots, \zeta_{K_w+C_w-1}$, we can write $y_w \xrightarrow{\zeta_{K_w},\dots,\zeta_{K_w+C_w-1}} x_{K_w}$.
  2. After the retrieval order of batch w has been revealed, actions to retrieve the revealed containers must be made. If $C_w > 1$, then $\{K_w, \dots, K_w + C_w - 2\} \neq \emptyset$. In this case, for $k \in \{K_w, \dots, K_w + C_w - 2\}$, let x_{k+1} be the configuration after applying action a_k to state x_k and before action a_{k+1}. Therefore, if $\xrightarrow{a_k}$ represents the application of action a_k, we have $x_k \xrightarrow{a_k} x_{k+1}$.
  3. The last container to be retrieved in batch w is the $(K_w + C_w - 1)^{th}$ container, thus, according to the previous point, $x_{K_w+C_w-1}$ corresponds to the configuration before $a_{K_w+C_w-1}$ is taken. After this retrieval, the order of the next batch (batch $w + 1$) has to be revealed, and according to the first point, the configuration is y_{w+1}. The configuration after retrieving batch W will be empty, thus we define y_{W+1} to be the empty configuration. So if $\xrightarrow{a_{K_w+C_w-1}}$ represents the application of action $a_{K_w+C_w-1}$, then we have $x_{K_w+C_w-1} \xrightarrow{a_{K_w+C_w-1}} y_{w+1}$.

  In summary, we have

$$\forall\, w \in \{1, \dots, W\}, \begin{cases} y_w \xrightarrow{\zeta_{K_w},\dots,\zeta_{K_w+C_w-1}} x_{K_w}, \\ x_k \xrightarrow{a_k} x_{k+1}, \text{ if } C_w > 1\ ,\ \forall\, k \in \{K_w, \dots, K_w + C_w - 2\}, \\ x_{K_w+C_w-1} \xrightarrow{a_{K_w+C_w-1}} y_{w+1}. \end{cases}$$

- Let the function $r(.)$ be such that $r(x)$ is number of relocations to retrieve the target container in configuration x. It is also equal to the number of containers blocking the target container. This

function is defined only for configurations in which the target container is identified. Specifically, it is defined for $(x_k)_{k=1,\dots,C}$ (but not for $(y_w)_{w=1,\dots,W}$). For $k \in \{1,\dots,C\}$, we refer to $r(x_k)$ as the *immediate cost* for the k^{th} retrieval.

- Let the function $f(.)$ be such that $f(x)$ is the minimum expected number of relocations required to retrieve all containers from configuration x. In addition, $f(x)$ is commonly referred to as the *cost-to-go* function of configuration x. Note that it is well defined for both $(x_k)_{k=1,\dots,C}$ and $(y_w)_{w=1,\dots,W}$.

By definition, we have:

$$\forall\, w \in \{1,\dots,W\},\ \begin{cases} f(y_w) = \underset{\zeta_k,\dots,\zeta_{k+C_w-1}}{\mathbb{E}} [f(x_k)] & \text{, where } k = K_w, \\ f(x_k) = r(x_k) + \min\limits_{a_k} \{f(x_{k+1})\} & \text{, if } C_w > 1 \text{ and } \forall\, k \in \{K_w,\dots,K_w + C_w - 2\}, \\ f(x_k) = r(x_k) + \min\limits_{a_k} \{f(y_{w+1})\} & \text{, where } k = K_w + C_w - 1, \end{cases}$$

which can be written as

$$\forall\, w \in \{1,\dots,W\},\ \begin{cases} f(y_w) = \underset{\zeta_{K_w},\dots,\zeta_{K_w+C_w-1}}{\mathbb{E}} [f(x_{K_w})], \\ f(x_{K_w}) = \min\limits_{a_{K_w},\dots,a_{K_w+C_w-1}} \left\{ \left(\sum\limits_{k=K_w}^{K_w+C_w-1} r(x_k) \right) + f(y_{w+1}) \right\}, \end{cases} \tag{1}$$

and $f(y_{W+1}) = 0$. Therefore, the SCRP is concerned with finding $f(y_1)$, where by induction:

$$f(y_1) = \underset{\zeta_{K_1},\dots,\zeta_{K_1+C_1-1}}{\mathbb{E}} \left[\min_{a_{K_1},\dots,a_{K_1+C_1-1}} \left\{ \underset{\zeta_{K_2},\dots,\zeta_{K_2+C_2-1}}{\mathbb{E}} \left[\dots \underset{\zeta_{K_W},\dots,\zeta_{K_W+C_W-1}}{\mathbb{E}} \left[\min_{a_{K_W},\dots,a_{K_W+C_W-1}} \left\{ \sum_{k=1}^{C} r(x_k) \right\} \right] \dots \right] \right\} \right]. \tag{2}$$

As a final remark, we note that batches should be as small as possible if only information is at stake. Indeed, smaller batches correspond to an efficient information system since more information is known about the retrieval order. But the size of batches is restricted by two intrinsic constraints:

1. **Batches should be at least larger than a certain size.** Indeed, a terminal offers time slots for trucks to register, and these slots cannot be too brief, as trucks would most certainly not arrive during their appointed slot due to traffic or other uncertain factors. Therefore, given the minimum time of a slot, the terminal will allow at least a certain number of trucks to register for each slot, i.e., the minimum batch size.

2. **Batches cannot be too large** for the batch model to be applicable, since in this model, the appointment time windows are supposed to be the same as or shorter than the target waiting time. As the target waiting time is limited, only a limited number of containers can be retrieved in a certain batch.

The online model. Using our notations, we briefly present the SCRP under the online model. Instead of Assumption A_6^*, the online model assumes that for each retrieval $k \in \{1,\dots,C\}$, only the next target container is revealed (i.e., ζ_k). Therefore, we consider the states $\left(y_k^o, x_k^o\right)_{k=1,\dots,C}$ defined such that $k \in \{1,\dots,C\}$, $y_k^o \xrightarrow{\zeta_k} x_k^o \xrightarrow{a_k} y_{k+1}^o$, where y_{C+1}^o is the empty configuration. In this case, if f^o denotes the cost-to-go function, then by definition we have $f^o\left(y_k^o\right) = \underset{\zeta_k}{\mathbb{E}}\left[f^o\left(x_k^o\right)\right]$ (with $f^o\left(y_{C+1}^o\right) = 0$), and $\forall\, k \in$

$\{1,\dots,C\}$, $f^o\left(x_k^o\right) = \min_{a_k}\left\{r\left(x_k^o\right) + f^o\left(y_{k+1}^o\right)\right\}$. By induction, the SCRP under the online model is hence concerned with finding:

$$f^o\left(y_1^o\right) = \mathop{\mathbb{E}}_{\zeta_1}\left[\min_{a_1}\left\{\mathop{\mathbb{E}}_{\zeta_2}\left[\dots\mathop{\mathbb{E}}_{\zeta_C}\left[\min_{a_C}\left\{\sum_{k=1}^{C} r\left(x_k^o\right)\right\}\right]\dots\right]\right\}\right].$$

The next lemma compares the batch and the online models theoretically (the proof can be found in Appendix A). It states that it is beneficial in terms of the expected number of relocations to use the batch model rather than the online model, if the first one is applicable. Practically, this suggests that the operator should always use available information.

Lemma 1. *Let y be a given initial configuration, then we have*

$$f(y) \leqslant f^o(y).$$

1.3 Decision Trees

Multistage stochastic optimization problems can be solved using decision trees in which *chance nodes* and *decision nodes* typically alternate. A *chance node* embodies the stochasticity of the model, whereas a *decision node* models the possible actions of the algorithm. In a decision tree for the SCRP, a node represents a configuration. The root node (denoted by 0) is the initial configuration, and the leaf nodes are configurations for which we can compute the cost-to-go function.

In our case, we slightly modify the structure of a typical decision tree, in the sense that chance nodes and decision nodes do not necessarily alternate. A chance node is a configuration for which the target container is not known yet, and information needs to be revealed (note that this occurs only at the beginning of each batch). A decision node is a configuration for which the target container is known. Using our notations, chance nodes correspond to $(y_w)_{w=1,\dots,W}$ and decision nodes correspond to $(x_k)_{k=1,\dots,C}$.

Let n be a node corresponding to a configuration in the decision tree. Thus $f(n)$, the cost-to-go function of n, is defined for all nodes n, and $r(n)$, the immediate cost function of n, is well defined when n is a decision node. We denote by λ_n the level of n in the decision tree, and define it as the number of containers remaining in the configuration. We denote the lowest level of the tree by $\lambda^* = \min_{n\in Tree}\{\lambda_n\}$. It corresponds to the level such that if $\lambda_n = \lambda^*$, $f(n)$ can be computed efficiently (the empty configuration being an obvious candidate with a cost-to-go of 0). Moreover,

- If n is a chance node, then there exists a unique $w \in \{1,\dots,W\}$ such that $\lambda_n = C - K_w + 1$. We denote by Ω_n the set of offspring of n, each offspring being a decision node corresponding to a realization of the random variables $\zeta_{K_w},\dots,\zeta_{K_w+C_w-1}$, i.e., the full retrieval order of containers in batch w. Note that n has a priori $|\Omega_n| = C_w!$ offspring.

- If n is a decision node, then $r(n)$ is well defined and is equal to the number of containers blocking the target container in configuration n (i.e., the $(C-\lambda_n+1)^{th}$ container to be retrieved). We denote by Δ_n the set of offspring of n, which can either be chance nodes if there exists $w \in \{1,\dots,W\}$ such that $\lambda_n + 1 = C - K_w + 1$, or decision nodes otherwise. For the sake of simplicity, we compute Δ_n greedily by considering all feasible combinations of relocations of the $r(n)$ containers blocking the target container in n. Note that $|\Delta_n|$ is of the order of $(S-1)^{r(n)}$, where S is the number of stacks.

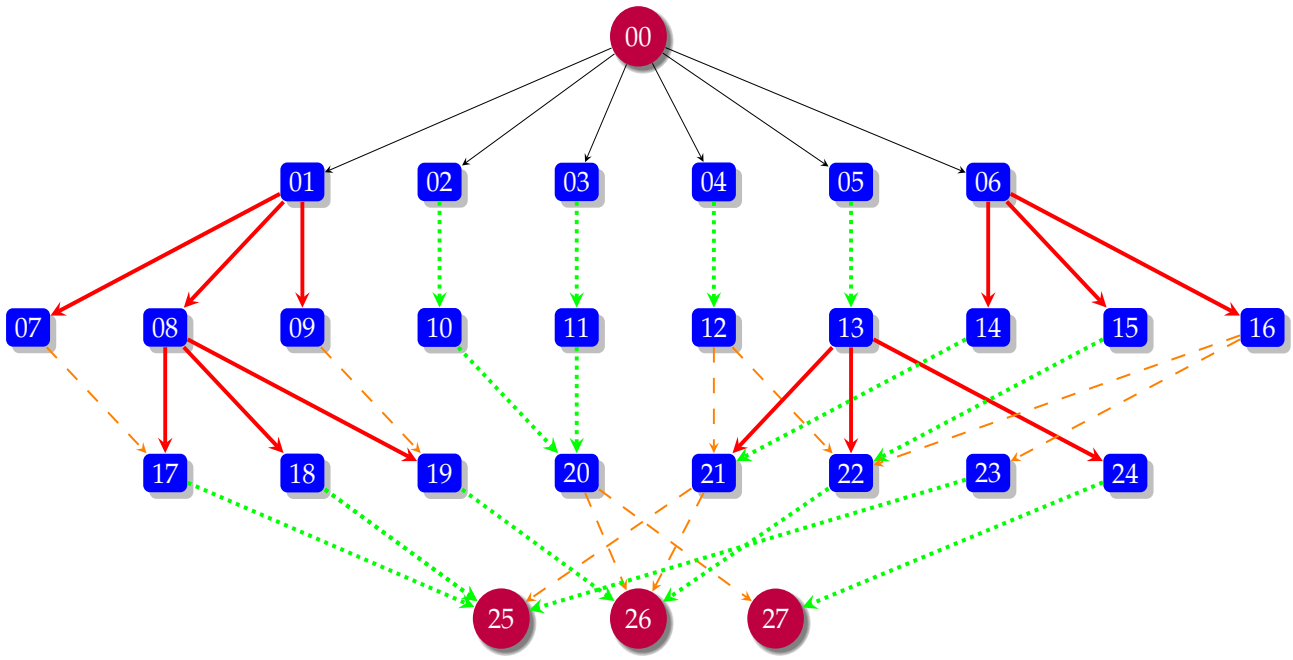

Figure 6: Decision tree represented with nodes. The colored arrows represent different values of immediate cost, i.e., the number of containers blocking the target container (dotted green: 0, dashed orange: 1, thick solid red: 2).

Equation (1) provides the relation to compute the cost-to-go by backtracking. For all n in the decision tree, we have

$$f(n) = \begin{cases} \dfrac{1}{|\Omega_n|} \displaystyle\sum_{n_i \in \Omega_n} f(n_i) \text{ , if } n \text{ is a chance node,} \\ r(n) + \min\limits_{n_i \in \Delta_n} \{f(n_i)\} \text{ , if } n \text{ is a decision node.} \end{cases} \tag{3}$$

In the case in which the probability of each permutation (in each batch) is not uniform, we mentioned that in practice, operators provide the probability of potential orders for each batch. Given a chance node n and one of its offspring $n_i \in \Omega_n$, this input probability is exactly the probability that the actual retrieval order is the one revealed in n_i. For a given node n, we denote these probabilities by $(p_{n_i})_{n_i \in \Omega_n}$. In this case, Equation (3) is replaced by:

$$f(n) = \begin{cases} \displaystyle\sum_{n_i \in \Omega_n} p_{n_i} f(n_i) \text{ , if } n \text{ is a chance node,} \\ r(n) + \min\limits_{n_i \in \Delta_n} \{f(n_i)\} \text{ , if } n \text{ is a decision node,} \end{cases}$$

for all n in the decision tree.

Figures 6 and 7 provide the description of the decision tree corresponding to the example in Figure 4, using chance/decision nodes and configurations, respectively. A chance node is depicted with a circle, and a decision node with a square.

To illustrate how to use Equation (3), we derive the calculations using the example in Figure 6. Suppose that $f(25) = f(26) = f(27) = 0.5$ are known, then we get $f(17) = f(18) = f(19) = 0.5$, $f(07) = f(09) = 1.5$,

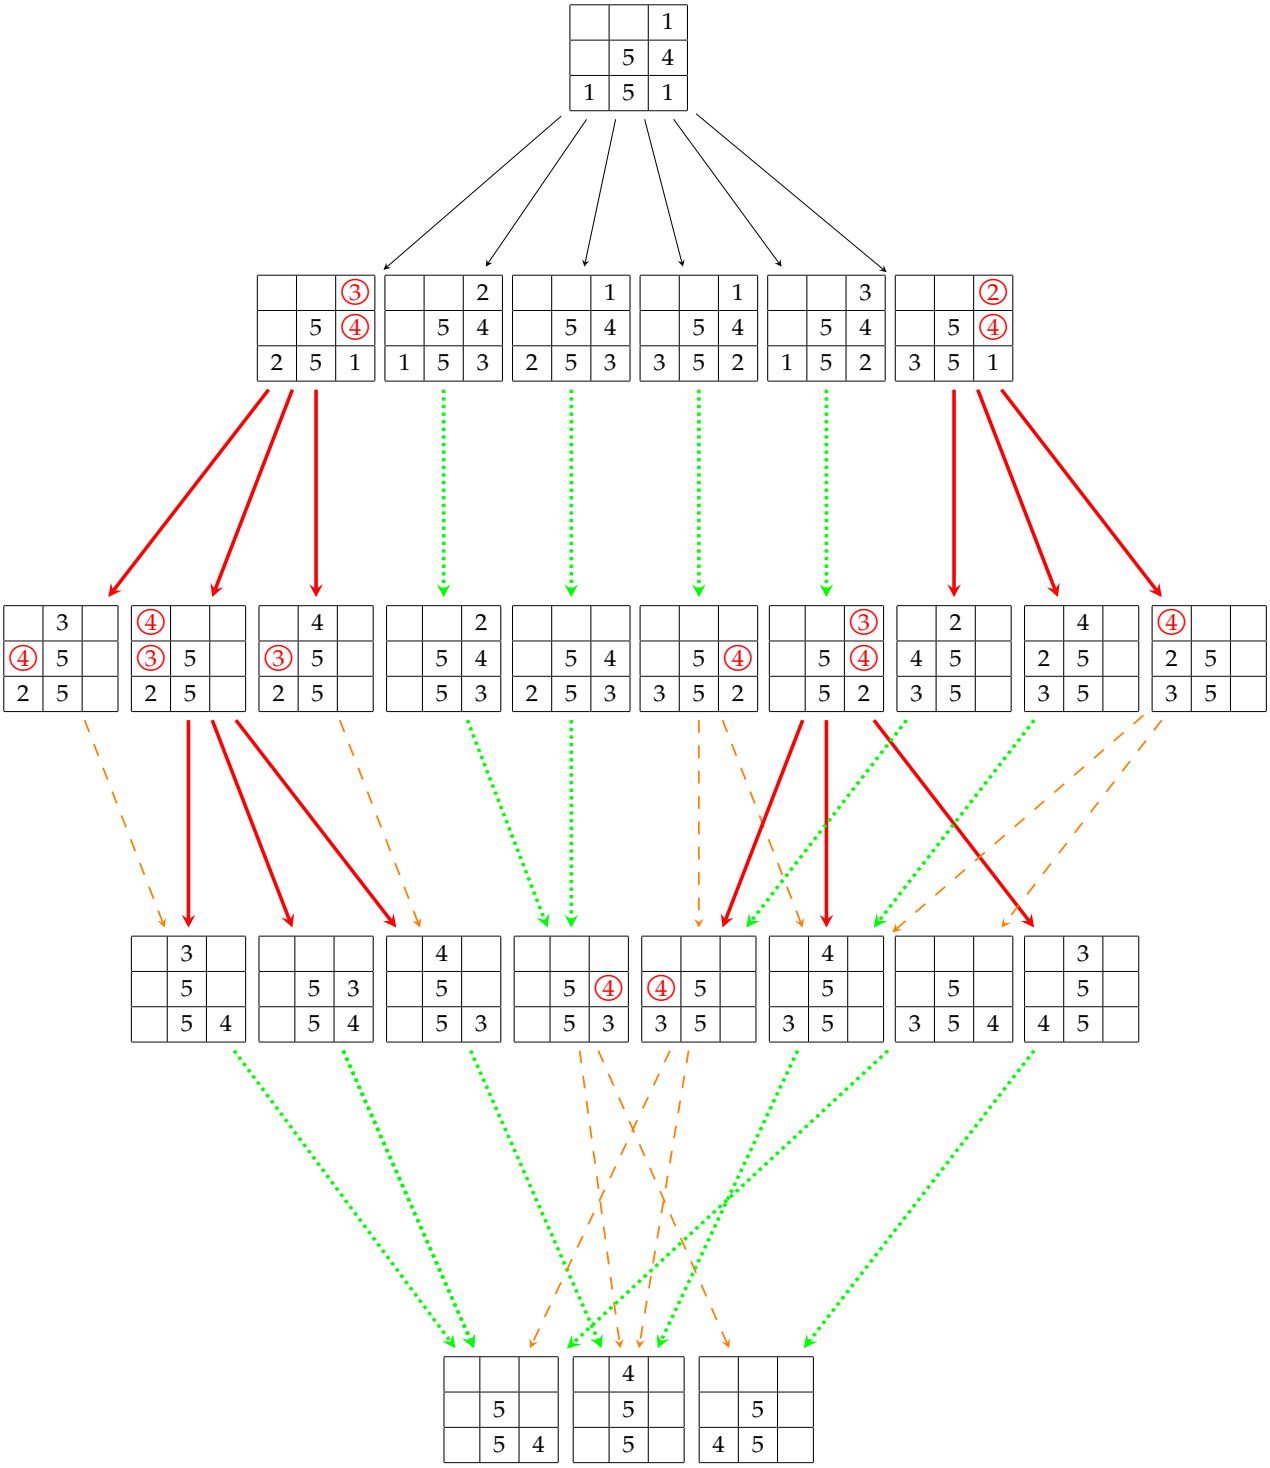

Figure 7: Decision tree represented with configurations. The colored arrows represent different values of immediate cost, i.e., the number of containers blocking the target container (dotted green: 0, dashed orange: 1, thick solid red: 2). Red circled numbers highlight containers blocking the target container.

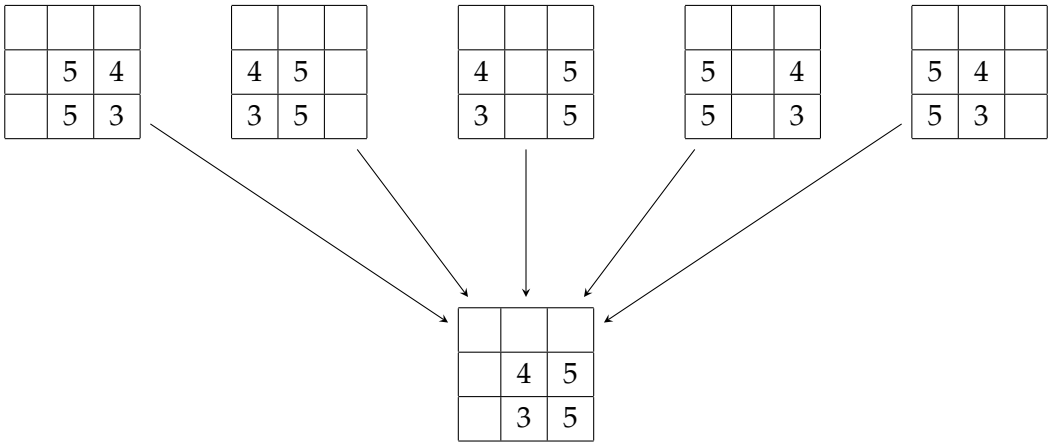

Figure 8: "Abstraction" procedure.

and $f(08) = 2.5$, leading to $f(01) = 3.5$. Similarly, by backtracking, we can compute $f(02) = f(03) = f(04) = 1.5$ and $f(05) = f(06) = 2.5$, giving us $f(00) = 13/6$.

As the example shows, considering the full decision tree can become intractable even for small examples, hence quickly impossible for larger problems. As previously mentioned, the number of decision offspring of a chance node scales with $C_w!$, and the number of offspring of a decision node is of the order of $(S-1)^{r(n)}$. So the size of the tree grows exponentially with the size of the problem. However, there exist general and specific techniques to reduce substantially the size of this tree, as we discuss now.

First, there are suboptimal approaches. One way to approximate $f(n)$, when n is a chance node, is to sample from its offspring. When Ω_n is large (which can happen with large batches), one might sample a certain number of realizations, resulting in a set of offspring $\Psi_n \subset \Omega_n$. By sampling "enough," we show in Section 4, we can ensure guarantees on expectation for such an algorithm. Another popular suboptimal approach is to use techniques from Approximate Dynamic Programming. These techniques can provide good empirical results, but no guarantee on how far from optimality can be obtained. This direction is not discussed in this paper but can be a direction for future work. Finally, another approach is to consider heuristics such as the ones described in the next section, which select a subset of the offspring of decision nodes, and lead to upper bounds on the optimal value $f(0)$.

Second, there exist ways to decrease the size of the tree while ensuring optimality. One is to reduce the number of nodes using the problem structure of the SCRP. In the online setting, Ku and Arthanari (2016) propose an "abstraction" technique that significantly shrinks the size of the tree. Thanks to Assumption A_4^*, we can consider that the stacks of a configuration are interchangeable, making many configurations equivalent in terms of number of relocations. For instance, in Figure 6, nodes 20 and 21 are identical in terms of number of relocations.

We describe the abstraction step with an example in Figure 8. The five configurations at the top are all equivalent to the configuration at the bottom. The abstraction transformation first ranks the stacks by increasing height. For stacks with equal height, it breaks ties by ranking them lexicographically going from bottom to top. Stacks are rearranged in order to have the first ranked on the left and the last on the right. Ku and Arthanari (2016) use a slightly different rule and add the extra step of removing empty

stacks. Using the abstraction procedure, the proposed algorithm should not generate a node twice with identical abstracted versions. Even though Ku and Arthanari (2016) introduce this rule for the online model, this abstraction step also applies in the batch setting. Throughout the rest of the paper, we will refer in pseudocode to the function ABSTRACT(n), when applying this method to a given node n. We mention that Ku and Arthanari (2016) also suggest caching strategies that could be added on the top of this simplification step, including caching part of the tree or using a transportation table.

Finally, the performance of a decision-tree-based algorithm depends on the exploration strategy of the tree. For the online model, Ku and Arthanari (2016) use depth-first-search (DFS), and we propose to explore the best-first-search (BFS) approach. Note that BFS requires some kind of measure that we define in Section 3.

In further sections, we explore two other ways to decrease the size of the tree while still ensuring optimality. The first one is specific to the SCRP, and uses properties of the problem to increase λ^*. Recall that λ^* is the minimum level of the tree at which we can find the optimal expected number of relocations, without further branching. We show that we can set λ^* to $\max\{S, C_W\}$, where S is the number of stacks, and C_W is the number of containers in the last batch. Comparatively, Ku and Arthanari (2016) branch until $\lambda^* = 0$. The second optimal pruning strategy uses lower bounds in a BFS scheme.

After we introduce the batch model for the SCRP (as well as the online model) and some preliminary concepts about decision trees, the next three sections develop approaches to solve the SCRP.

2 Heuristics and Lower Bounds

Before introducing the two main algorithms, we describe in this section heuristics and lower bounds for the SCRP. Indeed, *PBFS* and *PBFSA* build upon some of these bounds. In addition, these algorithms provide good intuition on how to solve the SCRP.

Let n be a given configuration with S stacks and T tiers. We say that a Container c is a *blocking container* in n if it is stacked above at least one container which is to be removed before c. Note that all the bounds mentioned below apply both in the batch and online models.

2.1 Heuristics

For the sake of completeness of this paper, we first describe three existing heuristics used in proofs and/or our computational experiments before describing two novel heuristics.

2.1.1 Existing heuristics

Random. For every relocation of a blocking Container c from Stack s, the Random heuristic picks any Stack $s' \neq s$ uniformly at random among stacks that are not "full," i.e., stacks containing strictly less than T containers.

Leveling (L). For every relocation of a blocking Container c from Stack s, L chooses the Stack $s' \neq s$ currently containing the fewest containers, breaking ties arbitrarily by selecting such leftmost stack.

Heuristic L is interesting for several reasons. Most important, it is an intuitive and commonly used heuristic in real operations in that it uses no more than the height of each stack. In addition, it does not

require any information about batches or departure times, which means it is robust with respect to the inaccuracy of information.

Expected reshuffling index (ERI). This index-based heuristic was introduced by Ku and Arthanari (2016) for the online model. For every relocation of a blocking Container c from Stack s, ERI computes a score called the expected reshuffling index for each Stack $s' \neq s$ that is not full, denoted by $ERI(s', c)$. ERI chooses the Stack $s' \neq s$ with the lowest $ERI(s', c)$. In the case of a tie, the policy breaks it by selecting the highest stack among the ones minimizing $ERI(s', c)$. Further ties are arbitrarily broken by selecting the leftmost stack verifying the two previous conditions. $ERI(s', c)$ corresponds to the expected number of containers in Stack s' that depart earlier than c. Let $H_{s'}$ be the current number of containers in s'. If $H_{s'} = 0$, then $ERI(s', c) = 0$. Otherwise, let $(c_1, \ldots, c_{H_{s'}})$ be the containers in s', then $ERI(s', c) = \sum_{i=1}^{H_{s'}} \mathbb{1}\{c_i < c\} + \frac{\mathbb{1}\{c_i = c\}}{2}$.

2.1.2 First new heuristic: Expected Minmax (EM)

EM considers an idea similar to that of Caserta et al. (2012) for the CRP. Let $min(s)$ be the smallest label of a container in s ($min(s) = C + 1$, if Stack s is empty). For every relocation of a blocking Container c from Stack s, we select the stack to which we relocate c using the following rules:

[Rule 1] If there exists $s' \neq s$ such that $min(s') > c$, let $M = \min_{s' \in \{1,\ldots,S\} \setminus s} \{\min(s') : \min(s') > c\}$. Select a stack such that $min(s') = M$, breaking ties by choosing from the highest ones, finally taking the leftmost stack if any ties remain.

[Rule 2] If for all Stacks $s' \neq s$, $min(s') \leqslant c$, let $M = \max_{s' \in \{1,\ldots,S\} \setminus s} \{\min(s')\}$. Select a stack such that $min(s') = M$. If there are several such stacks, select those with the minimum number of containers labeled M. Further ties are again broken by taking the highest ones, and finally choosing the leftmost one arbitrarily.

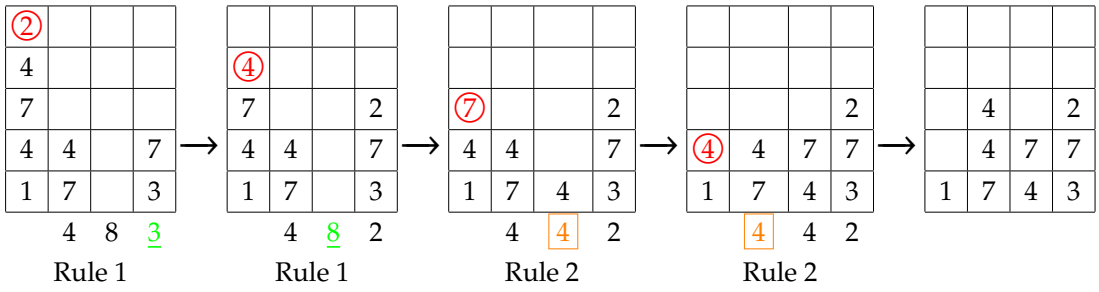

Figure 9: Decisions of the EM heuristic on an example with 5 tiers, 4 stacks, and 9 containers (3 per batch). Under the batch model, the first batch has been revealed, and we present the decisions to retrieve the first container made by EM. The container with the circled red label is the current blocking container. Numbers under the configuration correspond to the stack indices $min(s)$. The underlined green (respectively squared orange) indices correspond to the selected stack with the corresponding M when Rule 1 (respectively Rule 2) applies.

Rule 1 says: if there is a stack where min(s) is greater than c (c can almost surely avoid being relocated again), then choose such a stack where min(s) is minimized, since stacks with larger minimums can be useful for larger blocking containers.

If there is no stack satisfying min(s) $> c$ (Rule 2), then we have two cases following the same rule. On one hand, if $M = c$, then M is the maximum of the minimum labels of each stack, and c can potentially avoid being relocated again. If there are several stacks that maximize the minimum label, then by selecting the one with the least number of containers labeled M, EM minimizes the probability of c being relocated again. On the other hand, if $M < c$, c will almost surely be relocated again, and then EM chooses the stack where min(s) is maximized in order to delay the next unavoidable relocation of c as much as possible. We show how EM makes decision on a simple example in Figure 9.

2.1.3 Second new heuristic: Expected Group assignment (EG)

EM is quite intuitive because it tries to minimize the number of blocking containers after each retrieval. EG aims for the same goal, but uses some more sophisticated rules. EG is inspired by a heuristic designed by Wu and Ting (2012) for the complete information case, and we generalize this idea to the SCRP. It is different from ERI and EM because it considers a group of blocking containers together, whereas ERI and EM consider them one at a time. EG can be decomposed in two main phases for each retrieval. The decisions made by EG on the same example as shown in Figure 9 are given in Figure 10.

The first phase assigns the blocking containers for which there exists $s' \neq s$ such that $min(s') > c$. If this is not the case, the assignments of these containers will be ignored during the first phase. The acceptable containers are assigned in descending order of labels, i.e., the container with the highest label is assigned first (breaking ties for the highest one first). In order to assign these acceptable containers, the first phase applies the first of the EM rules. Finally, an acceptable container cannot be assigned to a stack if there is a container below it that was previously assigned to this stack.

The assignment in the second phase for the blocking containers not yet assigned might lead to additional relocations. These containers are assigned to other stacks in ascending order of labels. The second phase first computes a modified $min(s')$ index for each stack denoted by $Gmin(s')$, which is defined as follows: Let $H_{s'}$ be the height of Stack s' and $B(s')$ be the subset of containers assigned in the first phase to Stack s',

$$Gmin(s') = \begin{cases} -1, \ \textit{if } |B(s')| + H_{s'} = T, \\ min(s'), \ \textit{if } |B(s')| = 0, \\ B(s'), \ \textit{if } |B(s')| = 1, \\ 0, \ \textit{otherwise}. \end{cases}$$

If a stack is full after we assign the containers during the first phase, then it cannot be selected. If no container was assigned, the index remains as the *min*. If one container was assigned, it is "artificially" the new minimum of the stack. Finally, if more than one container was assigned, the index becomes very unattractive by being as low as possible (0). The second phase is similar to the EM heuristic, but it considers *Gmin* instead of the *min* index, breaking ties identically. Note that, after each assignment in the second phase, we update *Gmin* accordingly for the remaining containers to be assigned. For more details in the complete information case, we refer the reader to Wu and Ting (2012).

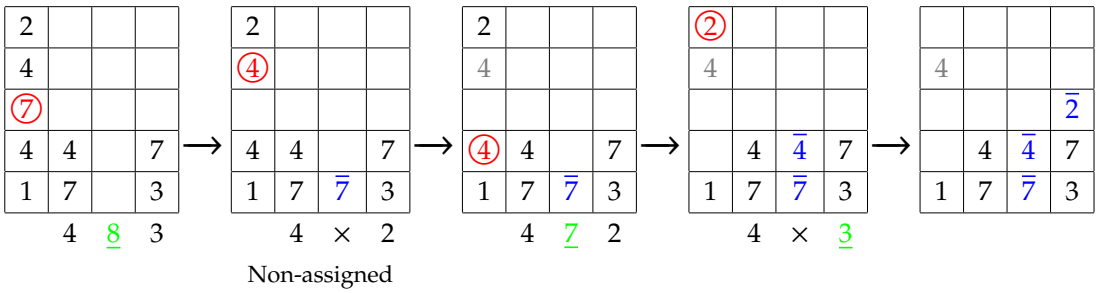

10a. First phase: EG assigns acceptable containers in descending order. The container with the circled red label is the next acceptable container that EG tries to assign to a stack. Containers with overlined blue labels are assigned, and with gray labels are unassigned. Below, we show the indices $min(s)$ to apply the first rule of EM ($\times$ means that a container below the considered container has already been assigned to a stack).

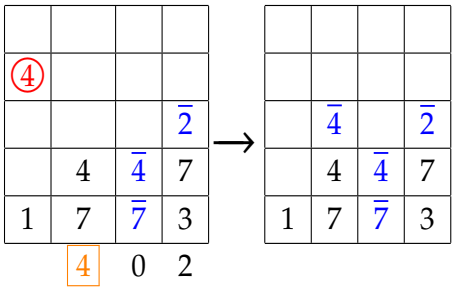

10b. Second Phase, EG assigns all unassigned containers using the index $Gmin(s)$.

Figure 10: Decisions of the EG heuristic in an example with 5 tiers, 4 stacks, and 9 containers (3 in each batch). Under the batch model, the first batch has been revealed, and we present the two-phases decisions to retrieve the first container made by EG.

2.2 Lower Bounds

After defining heuristics (upper bounds), we are now concerned with defining valid lower bounds for the SCRP. More specifically, we care about computing lower bounds for decision nodes in the decision tree defined before. Note that the computation of lower bounds easily extends to chance nodes.

2.2.1 Blocking lower bound

Suppose that the departure order is known, as in the CRP. The following lower bound was introduced by Kim and Hong (2006) and is based on the following simple observation. If a container is blocking in n, then it must be relocated at least once. Thus, the optimal number of relocations is lower bounded by the number of blocking containers.

In the SCRP, the retrieval order is a random variable, so the fact that a container is blocking is also random. Let us denote the expected number of blocking containers in n by $b(n)$. Therefore, by taking the expectation on the retrieval order of the previous fact, which holds for every retrieval order, we have the following observation.

Observation 1. *For all configurations n, $f(n)$ is the minimum expected number of relocations to empty n, and*

$b(n)$ is the expected number of blocking containers, then

$$f(n) \geqslant b(n).$$

Lemma 2 shows one way to compute the expected number of blocking containers for one stack, and $b(n)$ is the sum of the expected number of blocking containers of each stack of n. Mathematically, let $b_s(n)$ be the expected number of blocking containers in Stack s of n, we have $b(n) = \sum_{s=1}^{S} b_s(n)$.

Lemma 2. *Let n be a single stack configuration with T tiers, and $H \geqslant 0$ containers ($H \leqslant T$). If $H = 0$, we have*

$$b(n) = 0.$$

If $H \geqslant 1$, we denote the label of containers by $(c_i)_{i=1,\dots,H}$, where c_1 is the container at the bottom and c_H is the container at the top (see Figure 11), then we have:

$$b(n) = H - \sum_{h=1}^{H} \frac{\mathbb{1}\left\{c_h = \min\limits_{i=1,\dots,h} \{c_i\}\right\}}{\sum\limits_{i=1}^{h} \mathbb{1}\{c_h = c_i\}},$$

where $\mathbb{1}\{A\}$ is the indicator function of A.

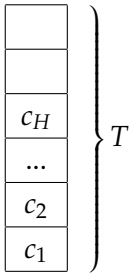

Figure 11: Example of a single stack configuration.

Proof. Clearly, if $H = 0$, then $b(n) = 0$. If $H \geqslant 1$, then by definition we have

$$b(n) = \mathbb{E}\left[\sum_{h=1}^{H} \mathbb{1}\{c_h \textit{ is a blocking container}\}\right] = \sum_{h=1}^{H} \mathbb{P}\left[c_h \textit{ is a blocking container}\right].$$

Let us fix $h \in \{1, \dots, H\}$, and compute the probability that c_h is blocking. We consider two cases:

- If $c_h > \min\limits_{i=1,\dots,h} \{c_i\}$, then c_h is almost surely blocking.
- Otherwise $c_h = \min\limits_{i=1,\dots,h} \{c_i\}$, and there are $\sum_{i=1}^{h} \mathbb{1}\{c_h = c_i\} - 1$ containers below c_h with the same label (or batch). Since each departure sequence between containers of the same batch is equally likely, the probability that c_h is blocking is equal to $\frac{\sum_{i=1}^{h} \mathbb{1}\{c_h = c_i\} - 1}{\sum_{i=1}^{h} \mathbb{1}\{c_h = c_i\}} = 1 - \frac{1}{\sum_{i=1}^{h} \mathbb{1}\{c_h = c_i\}}$.

Consequently, we get

$$\mathbb{P}\left[c_h \text{ is a blocking container}\right] = 1 \times \mathbb{1}\left\{c_h > \min_{i=1,\dots,h}\{c_i\}\right\} + \left(1 - \frac{1}{\sum_{i=1}^{h}\mathbb{1}\{c_h = c_i\}}\right) \times \mathbb{1}\left\{c_h = \min_{i=1,\dots,h}\{c_i\}\right\}$$
$$= 1 - \frac{\mathbb{1}\left\{c_h = \min_{i=1,\dots,h}\{c_i\}\right\}}{\sum_{i=1}^{h}\mathbb{1}\{c_h = c_i\}}.$$

We sum the above expression for $h = 1, \dots, H$ to conclude the proof. □

Therefore, one can compute the blocking lower bound as follows: let H^s be the number of containers in Stack s, and $\left(c_1^s, \dots, c_{H^s}^s\right)$ be the containers in Stack s listed from bottom to top, then

$$b(n) = \sum_{\substack{s=1,\dots,S \\ H^s \geqslant 1}} \left(H^s - \sum_{h=1}^{H^s} \frac{\mathbb{1}\left\{c_h^s = \min_{i=1,\dots,h}\left\{c_i^s\right\}\right\}}{\sum_{i=1}^{h}\mathbb{1}\left\{c_h^s = c_i^s\right\}} \right). \tag{4}$$

Non-uniform case. In the case where probabilities are not uniform across retrieval orders, we still consider a similar lower bound. For each Container c_h^s, let $q_{c_h^s}$ be the probability that c_h^s is the first container to be retrieved among the ones with the same batch, and positioned below in its stack. Equation (4) extends to give:

$$b(n) = \sum_{\substack{s=1,\dots,S \\ H^s \geqslant 1}} \left(H^s - \sum_{h=1}^{H^s} q_{c_h^s}\mathbb{1}\left\{c_h^s = \min_{i=1,\dots,h}\left\{c_i^s\right\}\right\} \right).$$

2.2.2 Look-ahead lower bounds

Note that the blocking lower bound b is only taking into account the current configuration. However, some relocations lead necessarily to an additional relocation. We refer to such relocations as "bad." Formally, let s be a stack of a configuration, and $\min(s)$ be the smallest label of a container in s. Recall that, if s is empty, we set $\min(s) = C + 1$. We say that the relocation of Container c from Stack s is "bad" if $c > \max_{s'=1,\dots,S\ ,\ s'\neq s}\{\min(s')\}$. We propose to construct a lower bound that anticipates "bad" relocations.

The basic idea is based on a similar one used by Zhu et al. (2012) for the CRP. We consider the first look-ahead lower bound denoted by $b_1(n)$. By definition, we take $b_1(n) = b(n) + d_1(n)$, where $b(n)$ is the blocking lower bound, and $d_1(n)$ is the expected number of unavoidable "bad" relocations while performing the first removal. We compute the term $d_1(n)$ by considering all realizations of the first target container. For each realization, we compute the number of unavoidable "bad" relocations and average them. Formally, for a given configuration n, consider U_n the set of potential next target container in n (which can be a singleton if it is known already), i.e., $U_n = \left\{c \,\middle|\, c = \min_{s=1,\dots,S}(min(s))\right\}$. Based on the definition of a bad relocation, we compute the number of unavoidable "bad" moves for each $u \in U_n$ denoted by $\beta(n, u)$, and we take:

$$d_1(n) = \frac{1}{|U_n|}\sum_{u\in U_n}\beta(n, u),$$

or $d_1(n) = \sum_{u \in U_n} p_{n,u}\beta(n,u)$, where $p_{n,u}$ is the probability that u is the next target container in n if the probabilities considered are not uniform (which can be computed using $(p_{n_i})_{n_i \in \Omega_n}$ if n is a chance node).

		4
		3
1	3	1

Figure 12: Example for look-ahead lower bounds.

For example, in Figure 12, the presented configuration denoted by n is such that $b(n) = 2$. Now consider a container $u \in U_n$: if u is the container labeled 1 in Stack 1, then there is no blocking container, so $\beta(n,u) = 0$; if u is the other container labeled 1, the relocation of the container labeled 4 from Stack 3 is necessarily a bad relocation, since $min(1) = 1 < 4$ and $min(2) = 3 < 4$, but it is not the case for the blocking container labeled 3, hence $\beta(n,u) = 1$. Therefore, $d_1(1) = 0.5(0+1) = 0.5$, and $b_1(n) = 2 + 0.5 = 2.5$, hence giving a lower bound closer to the optimal solution than $b(n)$. Note that, if n has an empty stack, then $\beta(n,u) = 0$ for all $u \in U_n$, and hence $d_1(n) = 0$.

We can refine this idea by trying to find unavoidable "bad" relocations for the second removal. In this case, the configuration depends on the first removal and the decisions that have been made accordingly. For the sake of clarity, consider that the first target container has been revealed, and denote it u_1. After retrieving u_1, only containers blocking u_1 have changed from their initial position. It can be very challenging to detect future unavoidable "bad" moves for these containers. To bypass this issue, we consider that all containers blocking u_1 are also removed, resulting in a configuration without u_1 and its blocking containers. Given this new configuration denoted by $n(u_1)$, we can compute the expected number of unavoidable bad moves $d_1(n(u_1))$. Since u_1 is actually random, we have to consider each scenario with its associated probability and compute a new configuration where blocking containers are retrieved with the target container. We denote the result $d_2(n)$, and it is a lower bound on the expected number of unavoidable bad relocations for the first two removals starting at n. Finally, our second look-ahead lower bound is given by $b_2(n) = b(n) + d_2(n)$.

Algorithm 1 Lower bound on the number of unavoidable bad relocations for the k first removals

1: **procedure** $[d_k(n)]$ = UnavoidableBadReloc $(n,\ k)$
2: **if** $k = 0$ or n has an empty stack or n is empty **then** $d_k(n) = 0$
3: **else** let $U_n = \{$*containers with minimum label in n*$\}$
4: **if** $k = 1$ **then** $d_k(n) = \frac{1}{|U_n|}\sum_{u \in U_n}\beta(n,u)$
5: **else**
6: **for** $u \in U_n$ **do**
7: Let $n(u)$ be the configuration n without u and all containers blocking u
8: Compute recursively $d_{k-1}(n(u))$ = UnavoidableBadReloc $(n(u),\ k-1)$
9: Compute $d_k(n) = \frac{1}{|U_n|}\sum_{u \in U_n}\beta(n,u) + d_{k-1}(n(u))$

This idea can easily be generalized for $k \geqslant 2$ by induction with $b_k(n) = b(n) + d_k(n)$. Here k is the number of removals that the lower bound considers to compute the expected number of unavoidable bad relocations (see pseudocode of Algorithm 1).

3 *PBFS*, a New Optimal Algorithm for the SCRP

Building upon lower bounds introduced in the previous section, this section introduces, studies, and proves the optimality of one of the main contributions of this paper, the *PBFS* Algorithm.

3.1 PBFS Algorithm

We start by giving the pseudocode of our algorithm, and we derive its optimality in Lemmas 3 and 4. *PBFS* takes two inputs, the configuration n for which we aim to compute $f(n)$, and a valid lower bound l. This algorithm uses a combination of four features to return $f(n)$. The first one is the BFS exploration of the tree based on a given lower bound l. We first compute f for the "most promising nodes," because nodes with small lower bounds are more likely to result in small f. The second technique is stopping to compute f recursively after level $\lambda^* = \max\{S, C_W\}$, by calculating it either using b or the A^* algorithm defined later. In Algorithm 2, $A^*(n)$ denotes the optimal number of relocations for node n obtained using the A^* algorithm. The third feature is pruning with a lower bound, revealing the suboptimality of some nodes without actually computing f. As its fourth feature, the algorithm uses the abstraction technique described previously.

Algorithm 2 *PBFS* Algorithm

procedure $[f(n)] = PBFS(n,\ l)$
  if $\lambda_n \leqslant S$ (n has less than S containers) **then** $f(n) = b(n)$
  else
    if n is a *chance node* **then** start with $\Psi_n^{PBFS} = \{\}$
      for $n_i \in \Omega_n$ **do** $n_i \leftarrow$ Abstract(n_i)
        if there exists $m = n_i$ already in Ψ_n^{PBFS} **then** $p_m^n \leftarrow p_m^n + \frac{1}{|\Omega_n|}$
        else if there exists $m = n_i$ already in the decision tree **then** add m to Ψ_n^{PBFS} and $p_m^n = \frac{1}{|\Omega_n|}$
        else add n_i to Ψ_n^{PBFS}, $p_{n_i}^n = \frac{1}{|\Omega_n|}$ and compute $f(n_i) = PBFS(n_i,\ l)$
      Compute $f(n) = \sum_{n_i \in \Psi_n^{PBFS}} p_{n_i}^n f(n_i)$
    else n is a *decision node*
      if $\lambda_n = C_W$ (the full retrieval order is known) **then** $f(n) = A^*(n)$
      else construct Δ_n by considering all feasible sets of decisions to deliver the target container
        Compute $l(n_i)$ for each $n_i \in \Delta_n$
        Sort $\left(n_{(1)}, n_{(2)}, \ldots, n_{(|\Delta_n|)}\right)$ in nondecreasing order of $l(.)$
        Compute $f(n_{(1)}) = PBFS\left(n_{(1)},\ l\right)$
        Start with $\Gamma_n^{PBFS} = \{n_{(1)}\}$ and $k = 2$
        while $k \leqslant |\Delta_n|$ and $l(n_{(k)}) < \min_{j=1,\ldots,k-1} \left\{f(n_{(j)})\right\}$ **do** $n_{(k)} \leftarrow$ Abstract$(n_{(k)})$
          if there exists $m = n_{(k)}$ already in the decision tree **then** add m to Γ_n^{PBFS}
          else add $n_{(k)}$ to Γ_n^{PBFS} and compute $f(n_{(k)}) = PBFS\left(n_{(k)},\ l\right)$
          Update $k = k + 1$
        $f(n) = r(n) + \min_{n_i \in \Gamma_n^{PBFS}} \{f(n_i)\}$

3.1.1 Decreasing the size of decision tree by increasing λ^* to max$\{S, C_W\}$

If $C_W \leqslant S$, then compute f(n) using b(n). Recall that, for every relocation, heuristic L chooses the stack with the fewest containers, breaking ties arbitrarily by choosing the leftmost one. Note that L always provides a valid upper bound for the SCRP. So if we denote the resulting expected number of relocations to empty configuration n using L by $f_L(n)$, then we have $f_L(n) \geqslant f(n)$.

Lemma 3. *Let n be a configuration with S stacks, T tiers, and C containers such that $C \leqslant S$, then we have*

$$f_L(n) = b(n) = f(n)$$

Proof. Consider a retrieval order of containers from n that has a nonzero probability of occurring. If there are no blocking containers, then the lemma clearly holds. Otherwise, let c be one of the blocking containers for this retrieval order, and consider the first removal for which c has to be relocated. For this removal, there are at most S containers in the configuration, hence there exists at least one empty stack to relocate c. Since heuristic L chooses always empty stacks if one exists, L would move c to one of the existing empty stacks. Note that in this case, c will never be blocking again, and hence will never be relocated again. This observation holds for any blocking containers, thus L relocates each blocking container at most once.

Since this fact holds for any retrieval order with nonzero probability, by taking expectation on the retrieval order, we have $f_L(n) \leqslant b(n)$, thus $f_L(n) \leqslant b(n) \leqslant f(n) \leqslant f_L(n)$, which concludes the proof. □

Lemma 3 states that for configurations with S containers or fewer, heuristic L is optimal for the SCRP. This implies that, for nodes at level S, we have access to the cost-to-go function using $b(n)$, as well as an optimal solution (provided by heuristic L). Hence *PBFS* can stop branching at $\lambda^* = S$ (line 2 of Algorithm 2).

If $C_W > S$, then compute f(n) using the A^* algorithm. If n is a decision node at level C_W, the full order of retrieval is known, and computing $f(n)$ reduces to solving a classical CRP, so we can leverage the existence of efficient solutions to the classical CRP such as the A^* algorithm, and take $\lambda^* = C_W$. Throughout the rest of the paper, A^* refers to the improved version of this algorithm presented in Borjian et al. (2015a).

Combining with the two previous observations, we can take $\lambda^* = \max\{S, C_W\}$.

3.1.2 Decreasing the size of decision tree by pruning using lower bounds

We would also like to reduce the size of the tree before level λ^*. For a decision node n, *PBFS* considers only a subset Γ_n^{PBFS} of all the offspring Δ_n (line 21 of Algorithm 2). Our goal is to set Γ_n^{PBFS} in order to still guarantee optimality.

First, *PBFS* generates all nodes $n_i \in \Delta_n$ by considering all feasible sets of decisions to deliver the target container in n (line 12 of Algorithm 2), and for each of them, compute a lower bound $l(n_i)$, where l is the input lower bound (line 13 of Algorithm 2). Let $\left(n_{(1)}, n_{(2)}, \ldots, n_{(|\Omega_n|)}\right)$ be the list of offspring of n sorted by nondecreasing lower bound (line 14 of Algorithm 2). The algorithm considers first $n_{(1)}$, adds it to Γ_n^{PBFS} and computes $f(n_{(1)})$ recursively (lines 15-16 of Algorithm 2). Then for $k \geqslant 2$, we consider $n_{(k)}$'s

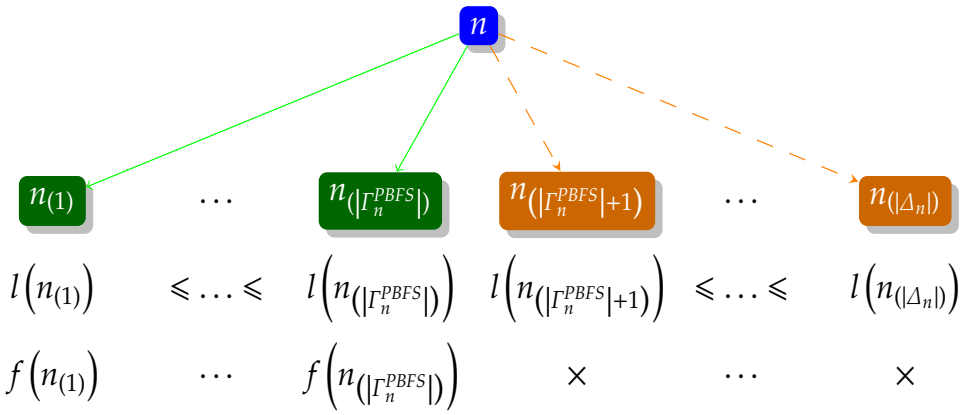

Figure 13: Illustration of the pruning rule. First, offspring are ordered by nondecreasing lower bounds. Then we start computing the objective function starting at $n_{(1)}$. We stop computing the objective functions once the pruning rule is reached. In the figure above, green nodes linked with full green arrows are nodes in Γ_n^{PBFS}, i.e., $f(.)$ has been computed. Orange nodes linked with dashed orange arrows are nodes in $\Delta_n \setminus \Gamma_n^{PBFS}$, i.e., $f(.)$ does not need to be computed, which is represented here by $\times$.

sequentially and check if $l(n_{(k)}) < \min_{j=1,\ldots,k-1} \left\{f(n_{(j)})\right\}$ (line 17 of Algorithm 2). If so, add $n_{(k)}$ to Γ_n^{PBFS} and compute $f(n_{(k)})$ recursively. If not, we stop branching on all nodes $n_{(k)}, \ldots, n_{(|\Omega_n|)}$. An illustration of the pruning rule is shown in Figure 13 and the next lemma shows the optimality of this rule.

Lemma 4. *Let n be a decision node in the decision tree, and Γ_n^{PBFS} be the subset of nodes considered for this node in Algorithm 2, and constructed as aforementioned, then we have*

$$\min_{m_i \in \Gamma_n^{PBFS}} \{f(m_i)\} = \min_{n_i \in \Delta_n} \{f(n_i)\}.$$

Proof. Let $\left(n_{(1)}, n_{(2)}, \ldots, n_{(|\Omega_n|)}\right)$ be the list of offspring of n, sorted by nondecreasing lower bounds. We consider two cases:

- If $\Gamma_n^{PBFS} = \Delta_n$, the statement clearly holds.
- Otherwise, there exists $k \leqslant |\Delta_n|$ such that $l(n_{(k)}) \geqslant \min_{j=1,\ldots,k-1} \left\{f(n_{(j)})\right\}$, and $\Gamma_n^{PBFS} = \left\{n_{(1)}, \ldots, n_{(k-1)}\right\}$. Note that we have $\forall k' \geqslant k$, $f(n_{(k')}) \geqslant l(n_{(k')}) \geqslant l(n_{(k)}) \geqslant \min_{j=1,\ldots,k-1} \left\{f(n_{(j)})\right\}$. Hence $\min_{n_i \in \Delta_n} \{f(n_i)\} = \min_{j=1,\ldots,k-1} \left\{f(n_{(j)})\right\} = \min_{m_i \in \Gamma_n^{PBFS}} \{f(m_i)\}$.

□

We claim that increasing λ^* to $\max\{S, C_W\}$ together with pruning in a Best-First-Search scheme, dramatically help in the efficiency of *PBFS* while guaranteeing optimality. However, this algorithm faces the issue that $|\Omega_n| = C_w!$ if n is a chance node. So if $\mathbf{C_w}$ **is large, typically** $\mathbf{C_w \geqslant 4}$, the number of nodes to consider gets too large. We tackle this issue by considering a near-optimal algorithm in the next section. This leads us to consider an alternative to *PBFS* in the case of larger batches.

4 *PBFSA*, Near-Optimal Algorithm with Guarantees for Large Batches

Building upon *PBFS* introduced in the previous section, this section describes the randomized algorithm *PBFSA* and shows some theoretical guarantees on expectation. This new algorithm is identical to *PBFS* except when computing the value function of a chance node (lines 4 to 17 of Algorithm 3). In order to decrease the number of decision offspring to consider for each chance node, we sample a certain number of *i.i.d.* permutations and consider only the decision nodes associated with these permutations as illustrated in Figure 14. In particular, *PBFSA* uses $f_{min}(.)$ and $f_{max}(.)$, lower and upper bound functions on $f(.)$ for offspring of chance nodes. In this paper, we use certain $f_{min}(.)$ and $f_{max}(.)$ defined in Equations (9)-(10), although others could be used. Combined with the fact that our problem has a finite number of sampling stages, this allows us to independently sample nodes in order to approximate the objective function. Using concentration inequalities, we can choose the number of samples needed to control the approximation error.

Since we perform a sampling at each chance node, *PBFSA* incurs an approximation error at each chance node. Lemma 5 proves that the total approximation error at the root node is on average the sum of all approximation errors from the root node to any leaf node. Therefore, consider a node n, then δ_n is the number of chance nodes between n and any leaf node in the decision tree. If the target error is ϵ at node n, *PBFSA* "allocates" evenly the remaining error to the next chance nodes, giving $\epsilon_n = \frac{\epsilon}{\delta_n}$ error at each remaining chance node where sampling occurs.

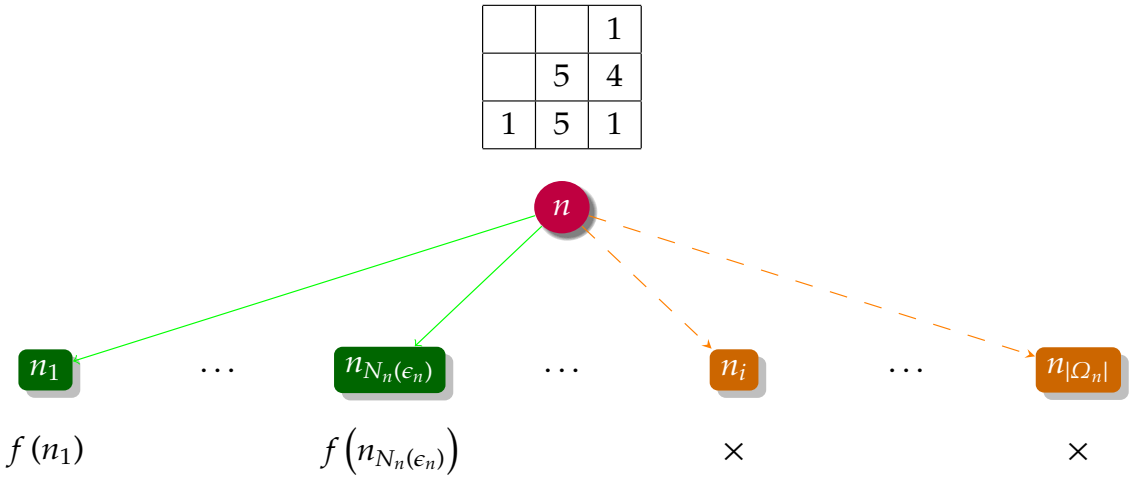

Figure 14: Illustration of the sampling rule. In this figure, the smallest batch is batch 1; therefore $w_{min} = 1$, and there are 6 containers, thus $\lambda_n = 6$. We have $\lambda^* = 3$ so $\delta_n = 1$ and thus $\epsilon_n = \epsilon$. These values allow us to compute the number of samples required $N_n(\epsilon_n)$. If $N_n(\epsilon_n)$ in less than the total number of offspring $|\Omega_n| = C_{w_{min}}! = 3!$, then we compute only $f(.)$ for sampled nodes. Ψ_n^{PBFSA} represents the subset of sampled nodes colored green and linked with full green arrows, and for which $f(.)$ needs to be computed. Note that $\left|\Psi_n^{PBFSA}\right| = N_n(\epsilon_n)$. Orange nodes linked with dashed orange arrows are nodes in $\Omega_n \setminus \Psi_n^{PBFSA}$, i.e., there were not sampled and $f(.)$ does not need to be computed, which is represented here by ×. Finally, the approximate value of $f(n)$ is the average of the objective values over all sampled nodes.

Algorithm 3 *PBFSA* Algorithm

procedure $[\tilde{f}(n)] = PBFSA\,(n,\ l,\ \epsilon)$
  if $\lambda_n \leqslant S$ **then** $\tilde{f}(n) = b(n)$
  else
    if n is a *chance node* **then** start with $\Psi_n^{PBFSA} = \{\}$. Let w_{min} be such that $\lambda_n = C - K_{w_{min}} + 1$
      Compute $\delta_n = \min\left\{w \in \{w_{min}, \ldots, W\} \left| \sum_{u=w_{min}}^{w} C_u \geqslant \lambda_n - \lambda^* \right.\right\}$ to get $\epsilon_n = \dfrac{\epsilon}{\delta_n}$
      Compute $f_{max}(n)$ and $f_{min}(n)$ from Equations (9)-(10) to get $N_n(\epsilon_n) = \dfrac{\pi\left(f_{max}(n) - f_{min}(n)\right)^2}{2\epsilon_n^2}$
      if $N_n(\epsilon_n) \leqslant C_{w_{min}}!$ **then**
        for $i = 1, \ldots, N_n(\epsilon_n)$ **do**
          Sample a random permutation, get corresponding $n_i \in \Omega_n$ and $n_i \leftarrow$ Abstract(n_i)
          if there is $m = n_i$ already in Ψ_n^{PBFSA} **then** $p_m^n \leftarrow p_m^n + \frac{1}{N_n(\epsilon_n)}$
          else if there is $m = n_i$ already in the decision tree **then** add m to Ψ_n^{PBFSA}, $p_m^n = \frac{1}{N_n(\epsilon_n)}$
          else add n_i to Ψ_n^{PBFSA}, $p_{n_i}^n = \frac{1}{N_n(\epsilon_n)}$ and compute $\tilde{f}(n_i) = PBFSA\,(n_i,\ l,\ \epsilon - \epsilon_n)$
      else
        for $n_i \in \Omega_n$ **do** $n_i \leftarrow$ Abstract(n_i)
          if there exists $m = n_i$ already in Ψ_n^{PBFSA} **then** $p_m^n \leftarrow p_m^n + \frac{1}{|\Omega_n|}$
          else if there exists $m = n_i$ already in decision tree **then** add m to Ψ_n^{PBFSA} and $p_m^n = \frac{1}{|\Omega_n|}$
          else add n_i to Ψ_n^{PBFSA}, $p_{n_i}^n = \frac{1}{|\Omega_n|}$ and compute $\tilde{f}(n_i) = PBFSA\,(n_i,\ l,\ \epsilon - \epsilon_n)$
      Compute $\tilde{f}(n) = \sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \tilde{f}(n_i)$
    else n is a *decision node*
      if $\lambda_n \leqslant C_W$ **then** $\tilde{f}(n) = A^*(n)$
      else Construct Δ_n by considering all feasible sets of decisions to deliver the target container
        Compute $l(n_i)$ for each $n_i \in \Delta_n$
        Sort $\left(n_{(1)}, n_{(2)}, \ldots, n_{(|\Delta_n|)}\right)$ in nondecreasing order of $l(.)$
        Compute $\tilde{f}(n_{(1)}) = PBFSA\left(n_{(1)},\ l,\ \epsilon\right)$
        Start with $\Gamma_n^{PBFSA} = \{n_{(1)}\}$ and $k = 2$
        while $k \leqslant |\Delta_n|$ and $l(n_{(k)}) < \min_{j=1,\ldots,k-1}\left\{\tilde{f}(n_{(j)})\right\}$ **do** $n_{(k)} \leftarrow$ Abstract$(n_{(k)})$
          if there exists $m = n_{(k)}$ already in the decision tree **then** add m to Γ_n^{PBFSA}
          else add $n_{(k)}$ to Γ_n^{PBFSA} and compute $\tilde{f}(n_{(k)}) = PBFSA\left(n_{(k)},\ l,\ \epsilon\right)$
          Update $k = k + 1$
        $\tilde{f}(n) = r(n) + \min_{n_i \in \Gamma_n^{PBFSA}}\left\{\tilde{f}(n_i)\right\}$

Formally, when n is a given chance node, recall that *PBFS* computes $f(n) = \dfrac{1}{|\Omega_n|}\sum_{n_i \in \Omega_n} f(n_i)$, where each $n_i \in \Omega_n$ represents one retrieval order (a random permutation) of batch w (if $\lambda_n = C - K_w + 1$). Instead, *PBFSA* computes the number of chance nodes between n and a leaf node denoted by δ_n. If ϵ is the target error at node n, the algorithm allocates $\epsilon_n = \frac{\epsilon}{\delta_n}$ error for sampling at node n and the remaining

(i.e., $\epsilon - \epsilon_n$) to its offspring (lines 12 and 17 of Algorithm 3). Using ϵ_n, we compute $N_n(\epsilon_n)$ and obtain a subset $\Psi_n \subset \Omega_n$ from $N_n(\epsilon_n)$ offspring drawn *i.i.d.* uniformly from Ω_n. The important part is to define $N_n(\epsilon_n)$, such that $\tilde{f}(n) = \frac{1}{|\Psi_n|} \sum_{m \in \Psi_n} f(m)$ is a "good" approximation of $f(n)$, i.e., $\left|\tilde{f}(n) - f(n)\right|$ is bounded by ϵ on average. Note that if $N_n(\epsilon_n) > C_{w_{min}}!$, we would need to sample more elements than the total number of offspring of n, and thus we do not sample (lines 14-17 of Algorithm 3).

PBFSA takes three input arguments, the configuration n for which we want to evaluate f, a valid lower bound l and an upper bound $\epsilon > 0$ on the total expected "error" ensured by the algorithm. It outputs $\tilde{f}(n)$, which is a randomized approximation of $f(n)$. Because of the samplings performed in line 9 in Algorithm 3, the output of *PBFSA* is random. The average error incurred by the algorithm is $\mathbb{E}\left[\left|\tilde{f}(n) - f(n)\right|\right]$, where the expectation is taken over the aforementioned samplings. Our main result (Lemma 5) states that *PBFSA* ensures $\mathbb{E}\left[\left|\tilde{f}(n) - f(n)\right|\right] \leqslant \epsilon$, in other words, *PBFSA* guarantees an average error of at most ϵ. The proof of Lemma 5 can be found in Appendix B.

Lemma 5. *Let n be a configuration with $\lambda_n \geqslant 0$ containers, l be a valid lower bound function, and $\epsilon > 0$. If $\tilde{f}(n) = PBFSA(n, l, \epsilon)$, then*

$$\mathbb{E}\left[\left|\tilde{f}(n) - f(n)\right|\right] \leqslant \epsilon.$$

Sketch of the proof. Lemma 5 is proven by backtracking from leaf nodes to the root node, i.e., consider a node n at level λ_n, then the proof is done by induction on λ_n. To show that the expected absolute value of the error at node n (i.e., $\tilde{f}(n) - f(n)$) is bounded by ϵ, we actually show that the expected positive and negative parts of the error are both bounded by $\epsilon/2$, which implies our result.

First, we consider the case where n is a decision node and we show that there is no additional error incurred by *PBFSA* at such node, i.e., if all the offspring of node n (in Γ_n^{PBFSA} and at level $\lambda_n - 1$) have the expected positive and negative parts of their error bounded by $\epsilon/2$ (the induction hypothesis), then the expected positive and negative parts of the error at node n are also bounded by $\epsilon/2$.

Second, we consider the case where n is a chance node, and we show that an additional error is incurred due to sampling. Using the lemmas proven below, the positive and negative parts of this additional error are bounded by $\epsilon_n/2$. Since all the offspring of node n (in Ψ_n^{PBFSA}) are decision nodes at level λ_n and *PBFSA* sets a target error of $(\epsilon - \epsilon_n)/2$ for these nodes, then the first part of the proof shows that the positive and negative parts of the error of each offspring of node n are bounded by $(\epsilon - \epsilon_n)/2$. Combining both observations leads to n having the positive and negative parts of its error bounded by $\epsilon/2$, which proves the lemma. □

4.1 Hoeffding's Inequality Applied to the SCRP

To prove this result, we use Hoeffding's inequality to compute the number of samples to ensure probabilistic guarantees. We first state the well-known inequality and a direct corollary. Proofs of Corollary 1, Lemma 6, and Lemma 7 can be found in Appendix C.

Theorem 1 (Hoeffding's inequality)**.** *Let $X \in [x_{min}, x_{max}]$ be a real-valued bounded random variable with mean value $\mathbb{E}[X]$. Let $N \in \mathbb{N}$ and $(X_1, \ldots, X_N)$ be N i.i.d. samples of X. If $\overline{X} = \frac{1}{N} \sum_{i=1}^{N} X_i$, then we have*

$$\forall \delta > 0\ ,\ \mathbb{P}\left(\overline{X} - \mathbb{E}[X] > \delta\right) \leqslant exp\left(\frac{-2N\delta^2}{(x_{max} - x_{min})^2}\right), \tag{5}$$

and

$$\forall \delta > 0\ ,\ \mathbb{P}\left(\overline{X} - \mathbb{E}\left[X\right] < -\delta\right) \leqslant exp\left(\frac{-2N\delta^2}{\left(x_{max} - x_{min}\right)^2}\right). \tag{6}$$

Corollary 1. *Let $X \in [x_{min}, x_{max}]$ be a real-valued bounded random variable with mean value $\mathbb{E}[X]$. Let $N \in \mathbb{N}$ and $(X_1, \ldots, X_N)$ be N i.i.d. samples of X. If $\overline{X} = \frac{1}{N}\sum_{i=1}^{N} X_i$, then $\forall \epsilon > 0$ such that $N \geqslant \frac{\pi\left(x_{max} - x_{min}\right)^2}{2\epsilon^2}$, we have*

$$\mathbb{E}\left[\left(\overline{X} - \mathbb{E}\left[X\right]\right)^+\right] \leqslant \frac{\epsilon}{2}, \tag{7}$$

$$\mathbb{E}\left[\left(\overline{X} - \mathbb{E}\left[X\right]\right)^-\right] \leqslant \frac{\epsilon}{2}, \tag{8}$$

where $x^+ = \max\{x, 0\}$ (resp. $x^- = -\min\{x, 0\}$) is the positive (resp. negative) part of x.

To use Hoeffding's inequality, we need to define lower (f_{min}) and upper (f_{max}) bound functions, such that for each chance node n, $f_{min}(n) \leqslant \min\limits_{n_i \in \Omega_n}\{f(n_i)\}$ and $f_{max}(n) \geqslant \max\limits_{n_i \in \Omega_n}\{f(n_i)\}$.

Lemma 6. *Let n be a chance node, if*

$$f_{min}(n) = \min_{n_i \in \Omega_n}\{b(n_i)\}, \tag{9}$$

and

$$f_{max}(n) = \min\left\{\left(\left(\lambda_n - S\right)\left(T-1\right)\right)^+ + \left(\min\left\{S, \lambda_n\right\} - 1\right)\ ,\ \left(2\left\lceil\frac{\lambda_n}{S}\right\rceil - 1\right)\max_{n_i \in \Omega_n}\{b(n_i)\}\right\}, \tag{10}$$

then

$$f_{min}(n) \leqslant \min_{n_i \in \Omega_n}\{f(n_i)\} \text{ and } f_{max}(n) \geqslant \max_{n_i \in \Omega_n}\{f(n_i)\}.$$

Notice that the previous lemma involves computing $\min\limits_{n_i \in \Omega_n}\{b(n_i)\}$ and $\max\limits_{n_i \in \Omega_n}\{b(n_i)\}$. The following corollary provides an efficient way to compute these values.

Lemma 7. *Let n be a chance node, and $w_{min} \in \{1, \ldots, W\}$ be such that $\lambda_n = C - K_{w_{min}} + 1$ (i.e., the minimum batch in n). For each Stack s of n with $H^s \geqslant 1$ containers, let $(c_h^s)_{h=1,\ldots,H^s}$ be the containers in s, where c_1^s is the container at the bottom and $c_{H^s}^s$ is the container at the top (see Figure 11, for the case $H = H^s$). Finally, consider $C_{w_{min}}^s = \left|\left\{c_h^s = K_{w_{min}},\ h = 1, \ldots, H^s\right\}\right|$. Then we have*

$$\min_{n_i \in \Omega_n}\{b(n_i)\} = \sum_{\substack{s=1,\ldots,S \\ H^s \geqslant 1}}\left(H^s - C_{w_{min}}^s - \sum_{\substack{h=1,\ldots,H^s \\ c_h^s \neq K_{w_{min}}}} \frac{\mathbb{1}\left\{c_h^s = \min\limits_{i=1,\ldots,h}\left\{c_i^s\right\}\right\}}{\sum\limits_{i=1}^{h}\mathbb{1}\left\{c_h^s = c_i^s\right\}}\right), \tag{11}$$

and

$$\max_{n_i \in \Omega_n}\{b(n_i)\} = \sum_{\substack{s=1,\ldots,S \\ H^s \geqslant 1}}\left(H^s - \sum_{\substack{h=1,\ldots,H^s \\ c_h^s \neq K_{w_{min}}}} \frac{\mathbb{1}\left\{c_h^s = \min\limits_{i=1,\ldots,h}\left\{c_i^s\right\}\right\}}{\sum\limits_{i=1}^{h}\mathbb{1}\left\{c_h^s = c_i^s\right\}}\right). \tag{12}$$

Non-uniform case. Similar to the blocking lower bound, we can extend Lemma 7 to the case where probabilities are not uniform across retrieval orders. Recall that $q_{c_h^s}$ denotes the probability that c_h^s is the first one to be retrieved among the ones positioned below in its stack and with the same label. Then we have

$$\min_{n_i \in \Omega_n} \{b(n_i)\} = \sum_{\substack{s=1,\dots,S \\ H^s \geqslant 1}} \left(H^s - C^s_{w_{min}} - \sum_{\substack{h=1,\dots,H^s \\ c_h^s \neq K_{w_{min}}}} q_{c_h^s} \mathbb{1} \left\{ c_h^s = \min_{i=1,\dots,h} \left\{ c_i^s \right\} \right\} \right),$$

and

$$\max_{n_i \in \Omega_n} \{b(n_i)\} = \sum_{\substack{s=1,\dots,S \\ H^s \geqslant 1}} \left(H^s - \sum_{\substack{h=1,\dots,H^s \\ c_h^s \neq K_{w_{min}}}} q_{c_h^s} \mathbb{1} \left\{ c_h^s = \min_{i=1,\dots,h} \left\{ c_i^s \right\} \right\} \right).$$

5 Computational Experiments

Having introduced lower and upper bounds, *PBFS*, *PBFSA*, and theoretical guarantees in previous sections, we present several experimental results in this section to understand the effectiveness of our algorithms for the SCRP. For clarity, we refer to the set of instances from Ku and Arthanari (2016) as the *existing dataset*. We present four sets of experiments:

1. Based on instances from the existing dataset, which have relatively small batches, we test the *PBFS* algorithm, as well as the two new heuristics and our lower bounds.

2. We slightly modify the existing dataset to obtain the *modified dataset*, in order to obtain instances with relatively larger batches. We test the efficiency of *PBFSA* on this modified dataset.

3. Based on the existing dataset, we show that *PBFS* improves on the algorithm proposed in Ku and Arthanari (2016) for the online model. Moreover, the two new heuristics (EM and EG) outperform the ERI algorithm on expectation for the majority of the instances in the dataset.

4. We change the existing dataset by considering that all containers belong to a unique batch. We show strong computational evidence to support Conjecture 1, which states that the leveling policy is optimal for the SCRP under the online model with a unique batch.

All experiments are performed on a MacBook Pro with 2.2 GHz Intel Core i7 processor and 8.00 GB of RAM, and the programming language is MATLAB 2016a. Finally, all results and instances used in this section are available at https://github.com/vgalle/StochasticCRP.

Implementation of heuristics.

1. Computing the number of relocations using b when there are S containers or less: In the retrieval process, **when there are S containers or fewer remaining in the configuration, the expected number of relocations performed by ERI, EM, EG, and L is computed using b**. This is motivated by the following observation: ERI, EM, EG, and L are optimal when there are S containers or fewer remaining in the configuration, and Lemma 3 shows that the optimal expected number of

relocations in this case is equal to b. Therefore, for all heuristics (except Random), instead of running simulations until there are no containers left, we stop when there are S containers left and compute the expected number of relocations using b instead.

2. Estimating the expected number of relocations using sampling: To estimate the exact objective value for a given heuristic, one would have to consider all possible retrieval scenarios. Instead, for each heuristic **unless specified otherwise, we report the average over 5000 samples (of retrieval orders) for each instance**, where samples are uniformly drawn at random.

Existing dataset description. The full description of the dataset can be found in Ku and Arthanari (2016), and the original data set is available at http://crp-timewindow.blogspot.com. Note that:

- Configuration sizes vary from $\mathbf{T} = \mathbf{3}, \ldots, \mathbf{6}$ tiers, and $\mathbf{S} = \mathbf{5}, \ldots, \mathbf{10}$ stacks.
- Two occupancy rates are considered, **50 and 67 percent**. The occupancy rate ($\mu \in [0, 1]$) is defined such that the initial number of containers is $C = round\,(\mu \times S \times T)$, where $round(x)$ rounds x to the closest integer. Therefore, a given triplet (T, S, μ) is equivalent to a given triplet (T, S, C), and note that if $C = round\,(0.67 \times S \times T)$, the condition $0 \leqslant C \leqslant ST - (T - 1)$ is satisfied.
- Given a configuration size (T and S) and an occupancy rate (μ) resulting in a given initial number of containers (C), the dataset includes **30 different initial configurations**.
- For all 1,440 instances, the ratio between the number of batches and C is taken to be around half, i.e., there are on average two containers per batch, which is the smallest size for a batch.

In all our experiments, the time limit is set to an hour, and the first look-ahead lower bound b_1 is used as input for both *PBFS* **and** *PBFSA*. All instances are solved by heuristics and lower bounds within seconds or less.

5.1 Experiment 1: Batch Model with Small Batches

Table 1 gives a summary of the results as follows: ✔ indicates that all 30 instances are solved optimally by *PBFS*. In this case, the average solution time to solve these instances is given in seconds. Otherwise, the number of instances solved optimally is provided in red and in the form x/30. This table shows the efficiency of *PBFS* as it can solve all instances except two, for $T = 3$ and $T = 4$. Most important, the average time to solve these instances is under 10 seconds for these problem sizes. Since many ports today have a maximum tier requirement of 4 and need fast solutions, *PBFS* could be used in practice in the case of small batches. However, for $T = 5$ and $T = 6$, *PBFS* cannot solve all instances optimally in a timely manner. This suggests that, as the problem grows slightly, given instances become very hard to solve, which should not be a surprise, given the $\mathcal{NP}$-hardness of the problem. To avoid such situations in real operations, heuristics can be used to provide a "good" suboptimal solution (good in the sense of being not too far from optimality). Therefore, we want to evaluate the performance of these heuristics in order to know which one should be used in real operations.

We measure the performance of heuristics and the tightness of lower bounds in Tables 5 and 6. Concerning lower bounds, b encompasses a significant number of relocations. Adding unavoidable

	T	3		4		5		6	
S	Fill rate	50 percent	67 percent	50 percent	67 percent	50 percent	67 percent	50 percent	67 percent
5	C	8	10	10	13	13	17	15	20
	Solved	✔	✔	✔	✔	✔	28/30	✔	15/30
	Time (s)	0.01	0.02	0.03	0.12	0.17		5.17	
6	C	9	12	12	16	15	20	18	24
	Solved	✔	✔	✔	✔	✔	25/30	✔	14/30
	Time (s)	0.01	0.03	0.04	0.86	2.90		15.94	
7	C	11	14	14	19	18	23	21	28
	Solved	✔	✔	✔	✔	✔	24/30	23/30	5/30
	Time (s)	0.02	0.04	0.04	0.83	1.37			
8	C	12	16	16	21	20	27	24	32
	Solved	✔	✔	✔	✔	✔	20/30	22/30	5/30
	Time (s)	0.01	0.06	0.16	10.04	6.84			
9	C	14	18	18	24	23	30	27	36
	Solved	✔	✔	✔	✔	29/30	10/30	19/30	2/30
	Time (s)	0.03	0.10	0.37	8.84				
10	C	15	20	20	27	25	34	30	40
	Solved	✔	✔	✔	28/30	29/30	12/30	22/30	2/30
	Time (s)	0.03	0.10	0.54					

Table 1: Instances solved by *PBFS* in the batch model with small batches.

"bad" relocations in b_1 and b_2 improves the lower bound slightly. But experiments seem to confirm that $b_2(n) - b_1(n) \leqslant b_1(n) - b(n)$ holds, supporting our intuition that the relative increase of lower bounds $b_k(n) - b_{k-1}(n)$ decreases with k.

Concerning heuristics, EG and EM clearly outperform ERI as they result in lower expected numbers of relocations. When we have access to the optimal solutions, both heuristics are on average at most 2 percent more than the optimal solution. We expect this behavior to be similar for larger instances, however we only have access to lower bounds to evaluate their performances. In this case, heuristics are on average at most 11 percent more than b_2, hence at most 11 percent from the optimal solution (even though we believe that this number is very conservative, as our lower bounds are not "tight"). Therefore, both EG and EM appear to be good solutions for the SCRP under the batch model with small batches. In this case, we recommend using EM for its simplicity of implementation and its understandability.

5.2 Experiment 2: Batch Model with Larger Batches

5.2.1 Modifying existing instances

For the sake of reproducibility, we use the existing set of instances, but slightly modify it to consider larger batches. In order to create these instances, for each original instance n, consider n' with the same containers in the same configuration. But, if w is the batch of a container c in n, then we take the batch of c in n' to be $w' = \left\lceil \frac{w}{\gamma} \right\rceil$, where $\gamma > 1$, i.e., we merge γ batches together. In these experiments, we take $\gamma = 2$, which implies that batches have an average size of four.

	T	3		4		5		6	
S	Fill rate	50 percent	67 percent	50 percent	67 percent	50 percent	67 percent	50 percent	67 percent
5	C	8	10	10	13	13	17	15	20
	Solved	✔	✔	✔	✔	✔	21/30	✔	3/30
	Time (s)	0.08	0.29	0.14	4.55	3.20		72.70	
6	C	9	12	12	16	15	20	18	24
	Solved	✔	✔	✔	✔	✔	18/30	27/30	1/30
	Time (s)	0.08	0.47	0.25	126.37	14.74			
7	C	11	14	14	19	18	23	21	28
	Solved	✔	✔	✔	29/30	✔	9/30	14/30	0/30
	Time (s)	0.13	0.71	0.58		17.74			
8	C	12	16	16	21	20	27	24	32
	Solved	✔	✔	✔	28/30	29/30	6/30	17/30	1/30
	Time (s)	0.08	1.67	1.26					
9	C	14	18	18	24	23	30	27	36
	Solved	✔	✔	✔	26/30	25/30	5/30	15/30	0/30
	Time (s)	0.13	1.49	1.47					
10	C	15	20	20	27	25	34	30	40
	Solved	✔	✔	✔	22/30	29/30	7/30	14/30	0/30
	Time (s)	0.17	0.79	3.10					

Table 2: Instances solved by *PBFSA* in the batch model with larger batches.

5.2.2 Target error ϵ

To set our target error, we consider the following. Let n_0 be a given instance, and set $\epsilon = \frac{b(n_0)}{2}$. In this case, we know that $\epsilon \leqslant \frac{f(n_0)}{2}$, which implies that we are making an error of at most 50 percent. Note that this error is very conservative due to two major things: first, $b(n_0)$ is not necessarily representative of $f(n_0)$, especially if n_0 has many containers. Second, the number of samples given by Hoeffding's inequality is also very conservative, probably making our approximation substantially more accurate than what we can theoretically prove.

5.2.3 Results

Results are summarized in Table 2. Similarly to Table 1, ✔indicates that all 30 instances are solved approximately by *PBFSA* within the given expected error. In this case, the average solution time to solve these instances is given in seconds. Otherwise, the number of instances solved is provided in red and in the form x/30. This table shows that *PBFSA* presents several advantages. First, it solves most of instances with $T = 4$ and $S \leqslant 9$ approximately within two minutes, while we note that *PBFS* was not able to solve most of these. Moreover, as can be seen in Tables 7 and 8, *PBFSA* still outperforms the best heuristics despite the fact that we only set the theoretical guarantee to 50 percent of optimality. Together, these two advantages show the practicality of *PBFSA* for problem sizes typically encountered in real ports. Moreover, we note that increasing the batch size appears to make the problem significantly more complicated to solve as we can solve optimally larger instances in Experiment 1 (see Table 1). Finally, we remark that similar conclusions of Experiment 1 can be drawn for lower and upper bounds (see Tables 7

and 8).

5.3 Experiment 3: Online Model and Comparison with Ku and Arthanari (2016)

Table 3 gives a summary similar to the two previous experiments. In addition, we report the results of Ku and Arthanari (2016) who set a *time limit of eight hours* for each instance. In this table, ✓(✓) indicates that all 30 instances are solved optimally by both *PBFS* and Ku and Arthanari (2016). In this case, the average solution time in seconds to solve these instances is given for *PBFS* and for Ku and Arthanari (2016) in parentheses. ✔indicates that all 30 instances are solved optimally only by *PBFS* but not by Ku and Arthanari (2016). In this case, only the average solution time to solve these instances with *PBFS* is given in seconds. Otherwise, the number of instances solved by *PBFS* is provided in red and in the form x/30.

	T	3		4		5		6	
S	Fill rate	50 percent	67 percent	50 percent	67 percent	50 percent	67 percent	50 percent	67 percent
5	*C*	8	10	10	13	13	17	15	20
	Solved	✓(✓)	✔	✓(✓)	✔	✓(✓)	28/30	✔	18/30
	Time (s)	0.01 (0.02)	0.02	0.01 (2.51)	0.09	0.16 (2483.30)		3.74	
6	*C*	9	12	12	16	15	20	18	24
	Solved	✓(✓)	✔	✓(✓)	✔	✔	25/30	✔	15/30
	Time	0.01 (0.01)	0.04	0.06 (139.08)	0.81	1.85		14.92	
7	*C*	11	14	14	19	18	23	21	28
	Solved	✓(✓)	✔	✓(✓)	✔	✔	24/30	23/30	5/30
	Time (s)	0.02 (0.33)	0.04	0.04 (207.62)	0.67	1.38			
8	*C*	12	16	16	21	20	27	24	32
	Solved	✓(✓)	✔	✔	✔	✔	20/30	22/30	5/30
	Time (s)	0.01 (0.33)	0.05	0.10	8.29	5.85			
9	*C*	14	18	18	24	23	30	27	36
	Solved	✓(✓)	✔	✔	✔	29/30	10/30	19/30	2/30
	Time (s)	0.02 (32.24)	0.09	0.38	7.26				
10	*C*	15	20	20	27	25	34	30	40
	Solved	✓(✓)	✔	✔	28/30	29/30	12/30	16/30	2/30
	Time (s)	0.03 (58.85)	0.08	0.52					

Table 3: Instances solved by *PBFS* and Ku and Arthanari (2016) in the online model with small batch.

Results show strong evidence that our solution significantly improves the best existing results for the SCRP under the online model, given that we solve many larger instances optimally. Furthermore, it also outperforms the most recent algorithm in solution time for the problem sizes it can solve. It appears that, for problems for which we can solve all (or almost all) instances, most instances are "easy" to solve as the algorithm finds a solution within seconds. However, as in the batch model, there exist some instances for which the optimal solution still requires an exponential number of nodes, which makes our algorithm not tractable.

In Tables 9 and 10, we also report in parentheses the averages for ERI and Random found by Ku and Arthanari (2016). The results for Random are consistent. However, we find significantly better results for our implementation of ERI. This is unexpected since the only difference between the two implementations

is the use of lower bound b, when the configuration has fewer than S containers remaining. Nevertheless, ERI should also be optimal in this case, as it reduces to heuristic L. So this should not affect the expected number of relocations, and we cannot explain this difference. Finally, we point out that the results are quite similar to those in Experiment 1. Indeed, the existing data set has relatively small batches (on average 2 containers), which inherently makes the two models, batch and online, very close to each other.

5.4 Experiment 4: Online Model with a Unique Batch

	T	3		4		5		6	
S	Fill rate	50 percent	67 percent	50 percent	67 percent	50 percent	67 percent	50 percent	67 percent
5	*C*	8	10	10	13	13	17	15	20
	PBFS	2.08	3.33	3.54	6.53	6.56	12.05	9.13	17.28
	L	2.08	3.33	3.54	6.52	6.57	12.06	9.14	17.29
6	*C*	9	12	12	16	15	20	18	24
	PBFS	2.10	4.04	4.23	8.02	7.01	13.48	10.73	20.13
	L	2.10	4.04	4.23	8.02	7.01	13.48	10.73	20.13
7	*C*	11	14	14	19	18	23	21	28
	PBFS	2.69	4.60	4.82	9.55	8.58	14.75	12.22	23.14
	L	2.69	4.61	4.82	9.55	8.58	14.74	12.22	23.14
8	*C*	12	16	16	21	20	27	24	32
	PBFS	2.61	5.19	5.51	9.96	9.12	17.75	13.83	-
	L	2.62	5.19	5.51	9.95	9.12	17.75	13.83	26.04
9	*C*	14	18	18	24	23	30	27	36
	PBFS	3.31	5.72	6.10	11.58	10.89	19.15	-	-
	L	3.31	5.72	6.10	11.58	10.89	19.14	15.40	28.84
10	*C*	15	20	20	27	25	34	30	40
	PBFS	3.36	6.38	6.68	12.98	11.13	22.07	-	-
	L	3.36	6.38	6.67	12.99	11.13	22.06	16.87	31.77

Table 4: Instances solved with *PBFS* and heuristic L in the online model with a unique batch.

In this experiment, we consider the existing dataset, but assign all containers into a unique batch ($W = 1$). We consider the SCRP under the online model, where containers are revealed one at a time. Note that, in this case, each container is equally likely to be retrieved, and it is equivalent to know no-information about containers relative retrieval order. For each instance, we solve it twice: first using *PBFS*, and then using heuristic L, for which we sample *10,000 scenarios* (this is different from the 5000 samples considered in previous experiments). We report the results in Table 4. In this table, for each problem size, we report the expected optimal number of relocations averaged over 30 instances. "-" means that all 30 instances could not be solved optimally with *PBFS* within the given time limit of an hour. Note that the expected number of relocations using heuristic L reported in this experiment might be less than the one of *PBFS*; this is only due to the fact that we are sampling. Intuitively, L should be the optimal solution in this setting, and this experiment shows strong evidence that the next conjecture holds.

Conjecture 1. *Consider a configuration n with a unique batch. Let $f^o(n)$ be the minimum expected number of relocations to empty n under the online model, and let $f^{o,L}(n)$ be the expected number of relocations performed by the leveling heuristic under the online model, then*

$$f^o(n) = f^{o,L}(n). \tag{13}$$

This conjecture could also be made in the dynamic case, when containers arrive to be stacked. These results would have important ramifications for port operations, namely: *the optimal policy to minimize relocations, when no information is given in advance, is leveling configurations.*

6 Discussion

Managing relocation moves is one of the main challenges in the storage yard of container terminals, because it directly affects the costs and efficiency of yard operations. The container relocation problem, notorious for its computational intractability, addresses this issue. In this paper, we extend the CRP to the more practical case in which the retrieval order of containers is not known far in advance. First, we introduce a new stochastic model, called the batch model, show the applicability of this model, and compare it theoretically with the existing model of Zhao and Goodchild (2010). Then, we derive lower bounds and fast and efficient heuristics for the SCRP. Subsequently, we develop two novel algorithms (*PBFS* and *PBFSA*) to solve the stochastic CRP in different settings. Efficiencies of all algorithms are supported through computational experiments, for which all results are made available online at https://github.com/vgalle/StochasticCRP. Finally, using our solution methods and based on extensive experiments, we conjecture the optimality of the simple leveling heuristic in the online stochastic setting. More generally, the methods developed in this paper apply to multistage stochastic optimization problems, where the number of stages is finite, the set of feasible actions at each stage is finite, the objective function is bounded, and bounds on the objective function can be easily computed.

Future work could include the proof of Conjecture 1. Important future work can also be done on the optimal design of time windows for a TAS. On the one hand, small time windows imply more information on the retrieval sequence, hence higher operational efficiency of port operators. On the other hand, large time windows insure higher flexibility for truck drivers and a high rate of on-time arrivals. To balance this trade-off, one would need to quantify two important metrics with respect to the expected number of relocations: the "value of information" and the assignment of containers to "wrong" batches. Finally, in the grand scheme of port operations, the study of stacking and retrieving simultaneously, as well as the extension in the row dimension of blocks, is important for future studies of operations to take into account.

Acknowledgments

The authors would like to thank anonymous reviewers and the editor for their suggestions that led to significant improvements in this paper.

References

Bertsekas, D. P. (2005). In *Dynamic Programming and Optimal Control, Third Edition*, volume 1 of *Lecture Notes in Computer Science*. Athena Scientific.

Bonney, J. (2015). US ports move toward truck appointment model. *Information Handling Services IHS*. (Accessed 11.01.2017).

Borjian, S., Galle, V., Manshadi, V. H., Barnhart, C., and Jaillet, P. (2015a). Container Relocation Problem: Approximation, Asymptotic, and Incomplete Information. *CoRR*, abs/1505.04229. (Accessed 11.01.2017).

Borjian, S., Manshadi, V. H., Barnhart, C., and Jaillet, P. (2015b). Managing Relocation and Delay in Container Terminals with Flexible Service Policies. *CoRR*, abs/1503.01535. (Accessed 11.01.2017).

Caserta, M., Schwarze, S., and Voß, S. (2012). A mathematical formulation and complexity considerations for the blocks relocation problem. *European Journal of Operational Research*, 219(1):96–104.

Davies, P. (2009). Container terminal reservation systems. In *3rd Annual METRANS National Urban Freight Conference, Long Beach CA*.

Expósito-Izquierdo, C., Melián-Batista, B., and Moreno-Vega, J. M. (2015). An exact approach for the Blocks Relocation Problem. *Expert Systems with Applications*, 42(17):6408–6422.

Galle, V., Borjian Boroujeni, S., Manshadi, V., Barnhart, C., and Jaillet, P. (2016). An average-case asymptotic analysis of the Container Relocation Problem. *Operations Research Letters*, 44(6):723–728.

Giuliano, G. and O'Brien, T. (2007). Reducing port-related truck emissions: The terminal gate appointment system at the Ports of Los Angeles and Long Beach. *Transportation Research Part D: Transport and Environment*, 12(7):460–473.

Hakan Akyüz, M. and Lee, C.-Y. (2014). A mathematical formulation and efficient heuristics for the dynamic container relocation problem. *Naval Research Logistics (NRL)*, 61(2):101–118.

Kim, K. H. and Hong, G.-P. (2006). A heuristic rule for relocating blocks. *Computers & Operations Research*, 33(4):940–954.

Ku, D. and Arthanari, T. S. (2016). Container relocation problem with time windows for container departure. *European Journal of Operational Research*, 252(3):1031–1039.

Lehnfeld, J. and Knust, S. (2014). Loading, unloading and premarshalling of stacks in storage areas: Survey and classification. *European Journal of Operational Research*, 239(2):297–312.

Morais, P. and Lord, E. (2006). Terminal Appointment System Study. *Technical Report, Transportation Development Center of Transport Canada*. (Accessed 11.01.2017).

Murty, K. G., Wan, Y.-w., Liu, J., Tseng, M. M., Leung, E., Lai, K.-K., and Chiu, H. W. (2005). Hongkong International Terminals gains elastic capacity using a data-intensive decision-support system. *Interfaces*, 35(1):61–75.

Petering, M. E. and Hussein, M. I. (2013). A new mixed integer program and extended look-ahead heuristic algorithm for the block relocation problem. *European Journal of Operational Research*, 231(1):120–130.

Phillips, E. E. (2015). Southern California Ports to try Trucking Appointment System. *The Wall Street Journal*, 32(3):593–615. (Accessed 11.01.2017).

Sennott, L. I. (2009). *Stochastic dynamic programming and the control of queueing systems*, volume 504. John Wiley & Sons.

Stahlbock, R. and Voß, S. (2008). Operations research at container terminals: a literature update. *OR Spectrum*, 30(1):1–52.

Steenken, D., Voß, S., and Stahlbock, R. (2004). Container terminal operation and operations research - a classification and literature review. *OR Spectrum*, 26(1):3–49.

Tanaka, S. and Takii, K. (2016). A Faster Branch-and-Bound Algorithm for the Block Relocation Problem. *IEEE Transactions on Automation Science and Engineering*, 13(1):181–190.

Tierney, K. and Voß, S. (2016). Solving the Robust Container Pre-Marshalling Problem. In *Computational Logistics: 7th International Conference, ICCL 2016, Lisbon, Portugal, September 7-9, 2016, Proceedings*, pages 131–145.

Ünlüyurt, T. and Aydın, C. (2012). Improved rehandling strategies for the container retrieval process. *Journal of Advanced Transportation*, 46(4):378–393.

van Asperen, E., Borgman, B., and Dekker, R. (2013). Evaluating impact of truck announcements on container stacking efficiency. *Flexible Services and Manufacturing Journal*, 25(4):543–556.

Wu, K.-C. and Ting, C.-J. (2010). A beam search algorithm for minimizing reshuffle operations at container yards. In *Proceedings of the International Conference on Logistics and Maritime Systems*, pages 703–710.

Wu, K.-C. and Ting, C.-J. (2012). Heuristic approaches for minimizing reshuffle operations at container yard. In *Proceedings of the Asia Pacific industrial engineering & management systems conference*, pages 1407–1451.

Zehendner, E., Casserta, M., Feillet, D., Schwarze, S., and Voß, S. (2015). An improved mathematical formulation for the blocks relocation problem. *European Journal of Operational Research*, 245(2):415–422.

Zehendner, E. and Feillet, D. (2014a). Benefits of a truck appointment system on the service quality of inland transport modes at a multimodal container terminal. *European Journal of Operational Research*, 235(2):461–469.

Zehendner, E. and Feillet, D. (2014b). A branch and price approach for the container relocation problem. *International Journal of Production Research*, 52(24):7159–7176.

Zehendner, E., Feillet, D., and Jaillet, P. (2017). An algorithm with performance guarantee for the Online Container Relocation Problem. *European Journal of Operational Research*, 259(1):48–62.

Zhao, W. and Goodchild, A. V. (2010). The impact of truck arrival information on container terminal rehandling. *Transportation Research Part E: Logistics and Transportation Review*, 46(3):327–343.

Zhu, W., Qin, H., Lim, A., and Zhang, H. (2012). Iterative Deepening A* Algorithms for the Container Relocation Problem. *IEEE Transactions on Automation Science and Engineering*, 9(4):710–722.

A Theoretical and Computational Comparison of the Batch and the Online Models

A.1 Theoretical Comparison: Proof of Lemma 1

Lemma 1. *Let y be a given initial configuration, then we have*

$$f(y) \leqslant f^o(y).$$

Proof. We prove this lemma by induction on the number of batches W. The lemma clearly holds if y is empty (i.e., $W = 0$). Now consider $W \geqslant 1$ and $C_1 \geqslant 1$. For the sake of clarity of the proof, we define the following notation:

$$\forall\, d \in \{1,\dots,C_1\}\,, \begin{cases} y \xrightarrow{\zeta_1,\dots,\zeta_d} x_1^d \\ x_k^d \xrightarrow{a_k} x_{k+1}^d, \text{ if } d > 1\,,\ \forall\, k \in \{1,\dots,d-1\} \\ x_k^d \xrightarrow{a_k} y_{k-d+2}^d\,,\ \forall\, k \in \{d,\dots,C\} \\ y_{k-d+1}^d \xrightarrow{\zeta_k} x_k^d\,,\ \forall\, k \in \{d+1,\dots,C\}. \end{cases} \tag{14}$$

These notations correspond to the following process: the first d containers to be retrieved are all revealed at once. Then decisions to retrieve these d containers are made. Afterwards, each of the $C - d$ remaining containers is revealed one at a time (as in the online model). Under this revelation process, the minimum expected number of relocations is given by

$$f^d(y) = \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_d}\left[\min_{a_1,\dots,a_d}\left\{\sum_{k=1}^{d} r\left(x_k^d\right) + f^o\left(y_2^d\right)\right\}\right],\ \forall\, d \in \{1,\dots,C_1\}.$$

Moreover, using the recursion formula from the online model, we have

$$f^o\left(y_2^d\right) = \mathop{\mathbb{E}}_{\zeta_d}\left[\min_{a_d}\left\{r\left(x_{d+1}^d\right) + f^o\left(y_3^d\right)\right\}\right].$$

In particular, by definition of the online model, we have $f^o(y) = f^1(y)$.

Using these relations, let us prove that

$$f^d(y) \leqslant f^{d-1}(y)\,,\ \forall\, d \in \{2,\dots,C_1\}. \tag{15}$$

Let $d \in \{2, \dots, C_1\}$, we have

$$
\begin{aligned}
f^d(y) &= \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_d}\left[\min_{a_1,\dots,a_d}\left\{\sum_{k=1}^{d} r\left(x_k^d\right) + f^o\left(y_2^d\right)\right\}\right] \\
&= \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_{d-1}}\left[\mathop{\mathbb{E}}_{\zeta_d}\left[\min_{a_1,\dots,a_{d-1}}\left\{\min_{a_d}\left\{\sum_{k=1}^{d} r\left(x_k^d\right) + f^o\left(y_2^d\right)\right\}\right\}\right]\right] \qquad (16) \\
&\leqslant \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_{d-1}}\left[\min_{a_1,\dots,a_{d-1}}\left\{\mathop{\mathbb{E}}_{\zeta_d}\left[\min_{a_d}\left\{\sum_{k=1}^{d} r\left(x_k^{d-1}\right) + f^o\left(y_3^{d-1}\right)\right\}\right]\right\}\right] \qquad (17) \\
&= \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_{d-1}}\left[\min_{a_1,\dots,a_{d-1}}\left\{\sum_{k=1}^{d-1} r\left(x_k^{d-1}\right) + \mathop{\mathbb{E}}_{\zeta_d}\left[\min_{a_d}\left\{r\left(x_d^{d-1}\right) + f^o\left(y_3^{d-1}\right)\right\}\right]\right\}\right] \qquad (18) \\
&= \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_{d-1}}\left[\min_{a_1,\dots,a_{d-1}}\left\{\sum_{k=1}^{d-1} r\left(x_k^{d-1}\right) + f^o\left(y_2^{d-1}\right)\right\}\right] = f^{d-1}(y),
\end{aligned}
$$

where the equality (18) holds since x_k^{d-1} for $k \in \{1, \dots, d-1\}$ does not depend on a_d and ζ_d. Finally, the inequality holds because we have $\mathbb{E}[\min\{Z_1, \dots, Z_m\}] \leqslant \min\{\mathbb{E}[Z_1, \dots, Z_m]\}$ for any $Z_1, \dots, Z_m$ random variables. Note that we changed x_k^d in x_k^{d-1} and y_2^d in y_3^{d-1}. This change is necessary to stay consistent with the definition of Equation (14). Indeed, the order between the expectations and the minimums in Equation (16) implies that the process of the first d retrievals corresponds to

$$
y \xrightarrow{\zeta_1,\dots,\zeta_d} x_1^d \xrightarrow{a_1} x_2^d \xrightarrow{a_2} \dots \xrightarrow{a_{d-1}} x_d^d \xrightarrow{a_d} y_2^d,
$$

whereas the order between the expectations and the minimums in Equation (17) corresponds to the following process for the first d retrievals:

$$
y \xrightarrow{\zeta_1,\dots,\zeta_{d-1}} x_1^{d-1} \xrightarrow{a_1} x_2^{d-1} \xrightarrow{a_2} \dots \xrightarrow{a_{d-1}} y_2^{d-1} \xrightarrow{\zeta_d} x_d^{d-1} \xrightarrow{a_d} y_3^{d-1}.
$$

Recall Equation (1) and apply it with $w = 1$ (note that $K_1 = 1$ thus $K_1 + C_1 - 1 = C_1$) to get

$$
f(y) = \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_{C_1}}\left[\min_{a_1,\dots,a_{C_1}}\left\{\sum_{k=1}^{C_1} r(x_k) + f(y_2)\right\}\right].
$$

By induction, for all configuration y_2 with $W - 1$ batches we have $f(y_2) \leqslant f^o(y_2)$, thus

$$
f(y) = \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_{C_1}}\left[\min_{a_1,\dots,a_{C_1}}\left\{\sum_{k=1}^{C_1} r(x_k) + f(y_2)\right\}\right] \leqslant \mathop{\mathbb{E}}_{\zeta_1,\dots,\zeta_{C_1}}\left[\min_{a_1,\dots,a_{C_1}}\left\{\sum_{k=1}^{C_1} r\left(x_k^{C_1}\right) + f^o\left(y_2^{C_1}\right)\right\}\right] = f^{C_1}(y),
$$

where we replaced x_k by $x_k^{C_1}$ and y_2 by $y_2^{C_1}$ because, on the right-hand side of the inequality, the revelation process after the first C_1 containers is the online model. Finally, since $f^o(y) = f^1(y)$, by applying Equation (15) for each value of $d \in \{C_1, \dots, 2\}$, we complete the proof as

$$
f(y) \leqslant f^{C_1}(y) \leqslant f^{C_1-1}(y) \leqslant \dots \leqslant f^2(y) \leqslant f^1(y) = f^o(y).
$$

□

As a final remark, Lemma 1 is tight in the general setting. Indeed, there exists an initial configuration y for which $f(y) = f^o(y)$. For instance, consider the configuration in Figure 4a, then we have $f(y) = f^o(y) = 13/6$.

A.2 Computational Comparison

There also exist configurations for which $f(y) < f^o(y)$. The difference between these two values represents the value of taking into account available information (if possible).

To show a positive difference, we could have compared Experiments 1 and 3. However, since the average batch size is two, these experiments do not show a positive difference between both models. Another possibility would have been to use the instances from Experiment 2. However, as we previously mentioned, such instances are hard to solve optimally and not approximately.

Instead, we consider another set of simpler instances randomly generated: 100 instances with $T = 4$ tiers, $S = 4$ stacks, and $C = 12$ containers. Each instance has $W = 3$ batches and each batch has $C_w = 4$ containers (for $w = 1, 2, 3$). We solve each of these 100 instances under the batch and the online models. The code and detailed results are available at https://github.com/vgalle/StochasticCRP. We are especially interested about $\frac{f^o(.) - f(.)}{f(.)} \times 100$, which reveals the percentage difference of relocations between the batch and the online models.

On average over the 100 instances, the optimal expected number of relocations under the batch model is 6.526 and under the online model 6.61, hence giving a difference of 0.084. We observe here that this difference represents more than 1.287 percent of the optimal solution under the batch model, which is quite significant considering the fact that heuristic EM experimentally lies within 2 percent above the optimal solution. In addition, we noticed that for 25 of these instances, this difference was more than 2 percent with a maximum of about 4 percent (see Figure 15a).

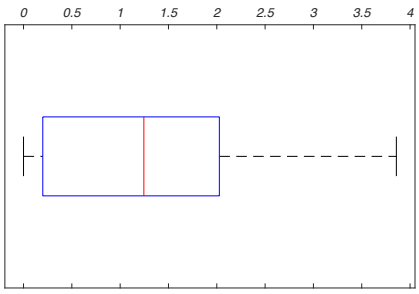

15a $T = 4$, $S = 4$, $C = 12$, $W = 3$ and $C_w = 4$ (for $w = 1, 2, 3$).

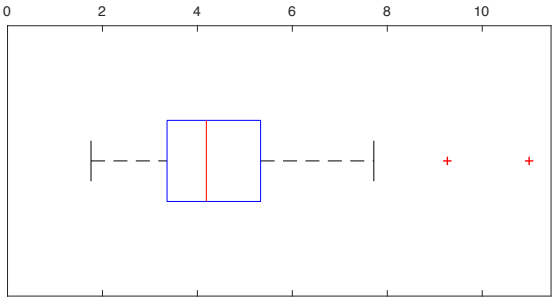

15b $T = 4$, $S = 4$, $C = 12$, $W = 2$ and $C_w = 6$ (for $w = 1, 2$).

Figure 15: Distributions of percentage difference between the batch and the online models from 100 randomly generated instances.

We also consider 100 instances for which $T = 4$, $S = 4$, and $C = 12$, but now $W = 2$ and each batch has $C_w = 6$ (for $w = 1, 2$). Figure 15b shows that this relative difference appears to increase when the batch size increases. Indeed, the average difference is about 4.251 percent (batch: 6.751, online: 7.038, difference: 0.287), with 25 instances having a difference of more than 5.3 percent.

B Proof of Lemma 5

Lemma 5. *Let n be a configuration with $\lambda_n \geqslant 0$ containers, l be a valid lower bound function, and $\epsilon > 0$. If $\tilde{f}(n) = PBFSA(n, l, \epsilon)$, then*

$$\mathbb{E}\left[\left|\tilde{f}(n) - f(n)\right|\right] \leqslant \epsilon.$$

Proof. The proof is by induction on λ_n. Throughout the proof, we use the same notations as the ones introduced in Algorithm 3. We say that $\tilde{f}$ verifies Conditions (A) and (B) at node n, if it verifies respectively the first and second inequalities below:

$$\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^+\right] \leqslant \frac{\epsilon}{2} \text{ and } \mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^-\right] \leqslant \frac{\epsilon}{2}.$$

Note that if $\tilde{f}$ verifies Conditions (A) and (B) at node n, then $\mathbb{E}\left[\left|\tilde{f}(n) - f(n)\right|\right] \leqslant \epsilon$, which would prove the lemma. Given $\epsilon > 0$, and l a valid lower bound, the induction hypothesis is:

$$\text{If } \tilde{f}(n) = PBFSA(n, l, \epsilon), \text{ then } \tilde{f} \text{ verifies Conditions } (A) \text{ and } (B) \text{ at node } n.$$

First, if $\lambda_n \leqslant S$, then $\tilde{f}(n) = b(n) = f(n)$, and therefore, $\tilde{f}$ verifies Conditions (A) and (B) at node n. In this case, $\tilde{f}(n)$ is actually deterministic since no sampling is performed by $PBFSA$. From now on, consider n such that $\lambda_n > S$.

Analysis of error at decision nodes

Let $\boldsymbol{n}$ be a decision node such that $\tilde{f}(n) = PBFSA(n, l, \epsilon)$ and $\lambda_n > S$. First, if $S < \lambda_n \leqslant C_W$, then $\tilde{f}(n) = A^*(n) = f(n)$, hence $\tilde{f}$ verifies Conditions (A) and (B) at n.

If $\lambda_n > \max\{S, C_W\}$, consider $\tilde{n} = \underset{n_i \in \Gamma_n^{PBFSA}}{argmin} \left\{\tilde{f}(n_i)\right\}$ and $n^* = \underset{n_i \in \Delta_n}{argmin} \{f(n_i)\}$. Note that $\tilde{f}(n) - f(n) = \tilde{f}(\tilde{n}) - f(n^*)$ almost surely $(a.s.)$, and by definition, $\tilde{n}$ and n^* are both such that $\lambda_{\tilde{n}} = \lambda_{n^*} = \lambda_n - 1 < \lambda_n$. Consider the following measurable event:

$$\mathcal{E} = \left\{\tilde{f}(n) - f(n) = \tilde{f}(\tilde{n}) - f(n^*) > 0\right\}. \tag{19}$$

- Conditioned on $\mathcal{E}$, we have $\left(\tilde{f}(n) - f(n)\right)^- = 0$ *a.s.*, thus

$$\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^- \;\middle|\; \mathcal{E}\right] = 0. \tag{20}$$

  Now let us show that conditioned on $\mathcal{E}$, $n^* \in \Gamma_n^{PBFSA}$ *a.s.*; we suppose by contradiction that *a.s.* $n^* \notin \Gamma_n^{PBFSA}$. If $k = \left|\Gamma_n^{PBFSA}\right|$, then $k < |\Delta_n|$ *a.s.*, and $\min_{j=1,\dots,k} \left\{\tilde{f}(n_{(j)})\right\} \leqslant l\left(n_{(k+1)}\right)$ *a.s.* By definition $\tilde{f}(\tilde{n}) = \min_{j=1,\dots,k} \left\{\tilde{f}(n_{(j)})\right\}$ so $\tilde{f}(\tilde{n}) \leqslant l\left(n_{(k+1)}\right)$ *a.s.* Since $n^* \notin \Gamma_n^{PBFSA}$, then there exists $k^* \in \{k+1, \dots, |\Delta_n|\}$ such that $n^* = n_{(k^*)}$. Since $\left(n_{(i)}\right)_{i \in \{1,\dots,|\Delta_n|\}}$ are ordered by nondecreasing $l(.)$, we have $l\left(n_{(k+1)}\right) \leqslant l\left(n_{(k^*)}\right) = l(n^*)$. Therefore $\tilde{f}(\tilde{n}) \leqslant l(n^*)$ *a.s.*; but, conditioned on $\mathcal{E}$, $\tilde{f}(\tilde{n}) > f(n^*) \geqslant l(n^*)$ *a.s.*, which leads to a contradiction. Thus conditioned on $\mathcal{E}$, $n^* \in \Gamma_n^{PBFSA}$

a.s.. Therefore, we have $\tilde{f}(n^*) = PBFSA(n^*, l, \epsilon)$. By induction, $\tilde{f}$ verifies Condition (A) at node n^*, thus

$$\mathbb{E}\left[\left(\tilde{f}(n^*) - f(n^*)\right)^+\right] \leqslant \frac{\epsilon}{2}. \tag{21}$$

Finally, since $\tilde{n} = \underset{n_i \in \Gamma_n^{PBFSA}}{argmin}\left\{\tilde{f}(n_i)\right\}$ and $n^* \in \Gamma_n^{PBFSA}$, then $\tilde{f}(\tilde{n}) \leqslant \tilde{f}(n^*)$ *a.s.*, so we have $\tilde{f}(\tilde{n}) - f(n^*) \leqslant \tilde{f}(n^*) - f(n^*)$ *a.s.* Consequently, we have $\left(\tilde{f}(n) - f(n)\right)^+ = \left(\tilde{f}(\tilde{n}) - f(n^*)\right)^+ \leqslant \left(\tilde{f}(n^*) - f(n^*)\right)^+$ *a.s.*, resulting in

$$\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^+ \,\middle|\, \mathcal{E}\right] \leqslant \mathbb{E}\left[\left(\tilde{f}(n^*) - f(n^*)\right)^+ \,\middle|\, \mathcal{E}\right]. \tag{22}$$

- Conditioned on $\overline{\mathcal{E}}$, we have $\left(\tilde{f}(n) - f(n)\right)^+ = 0$ *a.s.*, thus

$$\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^+ \,\middle|\, \overline{\mathcal{E}}\right] = 0. \tag{23}$$

Moreover, by definition $\tilde{n} \in \Gamma_n^{PBFSA}$ *a.s.*, and $\tilde{f}(\tilde{n}) = PBFSA(\tilde{n}, l, \epsilon)$, thus the induction hypothesis can be applied to $\tilde{n}$. In particular, we have

$$\mathbb{E}\left[\left(\tilde{f}(\tilde{n}) - f(\tilde{n})\right)^-\right] \leqslant \frac{\epsilon}{2}. \tag{24}$$

Finally, it is clear that $f(\tilde{n}) \geqslant f(n^*)$, then $\tilde{f}(\tilde{n}) - f(n^*) \geqslant \tilde{f}(\tilde{n}) - f(\tilde{n})$ *a.s.*, which is equivalent to $\left(\tilde{f}(n) - f(n)\right)^- = \left(\tilde{f}(\tilde{n}) - f(n^*)\right)^- \leqslant \left(\tilde{f}(\tilde{n}) - f(\tilde{n})\right)^-$ *a.s.*, resulting in

$$\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^- \,\middle|\, \overline{\mathcal{E}}\right] \leqslant \mathbb{E}\left[\left(\tilde{f}(\tilde{n}) - f(\tilde{n})\right)^- \,\middle|\, \overline{\mathcal{E}}\right]. \tag{25}$$

Finally, note the following observation: Let $Y \geqslant 0$ *a.s.*, and $\mathcal{F}$ be measurable, then we have

$$\mathbb{E}[Y \mid \mathcal{F}]\,\mathbb{P}(\mathcal{F}) \leqslant \mathbb{E}[Y] \text{ and } \mathbb{E}\left[Y \mid \overline{\mathcal{F}}\right]\mathbb{P}\left(\overline{\mathcal{F}}\right) \leqslant \mathbb{E}[Y].$$

Now we can derive

$$\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^+\right] = \mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^+ \,\middle|\, \mathcal{E}\right]\mathbb{P}(\mathcal{E}) \leqslant \mathbb{E}\left[\left(\tilde{f}(n^*) - f(n^*)\right)^+ \,\middle|\, \mathcal{E}\right]\mathbb{P}(\mathcal{E}) \leqslant \mathbb{E}\left[\left(\tilde{f}(n^*) - f(n^*)\right)^+\right] \leqslant \frac{\epsilon}{2},$$

where the first equality comes from Equation (23), the first inequality uses Equation (22), the second one holds thanks to $\left(\tilde{f}(n^*) - f(n^*)\right)^+ \geqslant 0$ *a.s.*, and the last one is Equation (21). Therefore, $\tilde{f}$ verifies Condition (A) at node n.

Similarly, we have

$$\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^-\right] = \mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^- \,\middle|\, \overline{\mathcal{E}}\right]\mathbb{P}\left(\overline{\mathcal{E}}\right) \leqslant \mathbb{E}\left[\left(\tilde{f}(\tilde{n}) - f(\tilde{n})\right)^- \,\middle|\, \overline{\mathcal{E}}\right]\mathbb{P}\left(\overline{\mathcal{E}}\right) \leqslant \mathbb{E}\left[\left(\tilde{f}(\tilde{n}) - f(\tilde{n})\right)^-\right] \leqslant \frac{\epsilon}{2},$$

where the first equality comes from Equation (20), the first inequality uses Equation (25), the second one holds thanks to $\left(\tilde{f}(\tilde{n}) - f(\tilde{n})\right)^- \geqslant 0$ *a.s.*, and the last one is Equation (24). Therefore, $\tilde{f}$ verifies Condition (B) at node n.

Therefore, we have proven that if n is a decision node with $\lambda_n > S$, $\tilde{f}$ verifies both Conditions (A) and (B) at node n, which proves the lemma for decision nodes.

Analysis of error at chance nodes

If $\boldsymbol{n}$ is a chance node such that $\tilde{f}(n) = PBFSA(n, l, \epsilon)$, and $\lambda_n > S$. Let us define $\overline{f}(n) = \sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n f(n_i)$, and show that

$$\mathbb{E}\left[\left(\tilde{f}(n) - \overline{f}(n)\right)^+\right] \leqslant \frac{\epsilon - \epsilon_n}{2} \text{ and } \mathbb{E}\left[\left(\tilde{f}(n) - \overline{f}(n)\right)^-\right] \leqslant \frac{\epsilon - \epsilon_n}{2}. \tag{26}$$

Recall that $\forall n_i \in \Psi_n^{PBFSA}$, $\lambda_{n_i} = \lambda_n$, and n_i are decision nodes such that $\tilde{f}(n_i) = PBFSA(n_i, l, \epsilon - \epsilon_n)$. Therefore, using the previous result, we know that $\mathbb{E}\left[\left(\tilde{f}(n_i) - f(n_i)\right)^+\right] \leqslant \frac{\epsilon - \epsilon_n}{2}$ and $\mathbb{E}\left[\left(\tilde{f}(n_i) - f(n_i)\right)^-\right] \leqslant \frac{\epsilon - \epsilon_n}{2}$. We derive the following calculations:

$$\begin{aligned}
\mathbb{E}\left[\left(\tilde{f}(n) - \overline{f}(n)\right)^+\right] &= \mathbb{E}\left[\left(\sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \left(\tilde{f}(n_i) - f(n_i)\right)\right)^+\right] \leqslant \mathbb{E}\left[\sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \left(\tilde{f}(n_i) - f(n_i)\right)^+\right] \\
&= \sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \mathbb{E}\left[\left(\tilde{f}(n_i) - f(n_i)\right)^+\right] \leqslant \sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \frac{\epsilon - \epsilon_n}{2} = \frac{\epsilon - \epsilon_n}{2}.
\end{aligned}$$

Similarly, we have

$$\begin{aligned}
\mathbb{E}\left[\left(\tilde{f}(n) - \overline{f}(n)\right)^-\right] &= \mathbb{E}\left[\left(\sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \left(\tilde{f}(n_i) - f(n_i)\right)\right)^-\right] \leqslant \mathbb{E}\left[\sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \left(\tilde{f}(n_i) - f(n_i)\right)^-\right] \\
&= \sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \mathbb{E}\left[\left(\tilde{f}(n_i) - f(n_i)\right)^-\right] \leqslant \sum_{n_i \in \Psi_n^{PBFSA}} p_{n_i}^n \frac{\epsilon - \epsilon_n}{2} = \frac{\epsilon - \epsilon_n}{2},
\end{aligned}$$

which proves Equation (26).

If $N_n(\epsilon_n) > C_{w_{min}}!$, then $f(n) = \overline{f}(n)$ so $\tilde{f}(n) - f(n) = \tilde{f}(n) - \overline{f}(n)$ *a.s.*, and since $\frac{\epsilon - \epsilon_n}{2} \leqslant \frac{\epsilon}{2}$, Equation (26) implies that $\tilde{f}$ verifies Conditions (A) and (B) at node n.

Otherwise, we have $N_n(\epsilon_n) \leqslant C_{w_{min}}!$. Since Ψ_n^{PBFSA} is constructed using $N_n(\epsilon_n) = \frac{\pi \left(f_{max}(n) - f_{min}(n)\right)^2}{2\epsilon_n^2}$ samples, thus by using Corollary 1, we have

$$\mathbb{E}\left[\left(\overline{f}(n) - f(n)\right)^+\right] \leqslant \frac{\epsilon_n}{2} \text{ and } \mathbb{E}\left[\left(\overline{f}(n) - f(n)\right)^-\right] \leqslant \frac{\epsilon_n}{2} \tag{27}$$

By combining Equations (26) and (27), we have

$$\begin{aligned}
\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^+\right] &\leqslant \mathbb{E}\left[\left(\tilde{f}(n) - \overline{f}(n)\right)^+\right] + \mathbb{E}\left[\left(\overline{f}(n) - f(n)\right)^+\right] \leqslant \frac{\epsilon - \epsilon_n}{2} + \frac{\epsilon_n}{2} = \frac{\epsilon}{2}, \\
\mathbb{E}\left[\left(\tilde{f}(n) - f(n)\right)^-\right] &\leqslant \mathbb{E}\left[\left(\tilde{f}(n) - \overline{f}(n)\right)^-\right] + \mathbb{E}\left[\left(\overline{f}(n) - f(n)\right)^-\right] \leqslant \frac{\epsilon - \epsilon_n}{2} + \frac{\epsilon_n}{2} = \frac{\epsilon}{2},
\end{aligned}$$

which shows that $\tilde{f}$ verifies Conditions (A) and (B) at node n and concludes the proof. □

C Technical Proofs of Section 4.1

Corollary 1. *Let $X \in [x_{min}, x_{max}]$ be a real-valued bounded random variable with mean value $\mathbb{E}[X]$. Let $N \in \mathbb{N}$ and $(X_1, \ldots, X_N)$ be N i.i.d. samples of X. If $\overline{X} = \frac{1}{N}\sum_{i=1}^{N} X_i$, then $\forall \epsilon > 0$ such that $N \geqslant \frac{\pi (x_{max} - x_{min})^2}{2\epsilon^2}$, we have*

$$\mathbb{E}\left[\left(\overline{X} - \mathbb{E}[X]\right)^+\right] \leqslant \frac{\epsilon}{2}, \tag{7}$$

$$\mathbb{E}\left[\left(\overline{X} - \mathbb{E}[X]\right)^-\right] \leqslant \frac{\epsilon}{2}, \tag{8}$$

where $x^+ = \max\{x, 0\}$ (resp. $x^- = -\min\{x, 0\}$) is the positive (resp. negative) part of x.

Proof of Corollary 1. For the first result, define $\Delta = \left(\overline{X} - \mathbb{E}[X]\right)^+ = \left(\overline{X} - \mathbb{E}[X]\right)\mathbb{1}\left\{\overline{X} - \mathbb{E}[X] > 0\right\}$. Note that Δ is a nonnegative random variable, and $\forall \delta > 0$, $\{\Delta > \delta\} = \left\{\overline{X} - \mathbb{E}[X] > \delta\right\}$. Let F_Δ denote the cumulative distribution function of Δ, thus, using Equation (5), $1 - F_\Delta(\delta) = \mathbb{P}(\Delta > \delta) = \mathbb{P}\left(\overline{X} - \mathbb{E}[X] > \delta\right) \leqslant exp\left(\frac{-2N\delta^2}{(x_{max} - x_{min})^2}\right)$, which gives

$$\begin{aligned}\mathbb{E}[\Delta] &= \int_{\delta=0}^{\infty} (1 - F_\Delta(\delta))\, d\delta \leqslant \int_{\delta=0}^{\infty} exp\left(\frac{-2N\delta^2}{(x_{max} - x_{min})^2}\right) d\delta \\ &= \frac{(x_{max} - x_{min})}{\sqrt{2N}} \int_{u=0}^{\infty} exp(-u^2) du = \frac{\sqrt{\pi}\,(x_{max} - x_{min})}{2\sqrt{2N}} \leqslant \frac{\epsilon}{2}.\end{aligned}$$

The proof of the second result is identical to the first one if we consider $\Delta' = \left(\overline{X} - \mathbb{E}[X]\right)^- = \left(\mathbb{E}[X] - \overline{X}\right)\mathbb{1}\left\{\mathbb{E}[X] - \overline{X} > 0\right\}$ and notice that $\forall \delta > 0$, $\{\Delta' > \delta\} = \left\{\overline{X} - \mathbb{E}[X] < -\delta\right\}$, hence $1 - F_{\Delta'}(\delta) \leqslant exp\left(\frac{-2N\delta^2}{(x_{max} - x_{min})^2}\right)$ using Equation (6). □

Lemma 6. *Let n be a chance node, if*

$$f_{min}(n) = \min_{n_i \in \Omega_n} \{b(n_i)\}, \tag{9}$$

and

$$f_{max}(n) = \min\left\{\left((\lambda_n - S)(T-1)\right)^+ + (\min\{S, \lambda_n\} - 1)\ ,\ \left(2\left\lceil\frac{\lambda_n}{S}\right\rceil - 1\right)\max_{n_i \in \Omega_n}\{b(n_i)\}\right\}, \tag{10}$$

then

$$f_{min}(n) \leqslant \min_{n_i \in \Omega_n}\{f(n_i)\} \text{ and } f_{max}(n) \geqslant \max_{n_i \in \Omega_n}\{f(n_i)\}.$$

Proof of Lemma 6. Since $b(n_i) \leqslant f(n_i)$, then we have $f_{min}(n) = \min_{n_i \in \Omega_n}\{b(n_i)\} \leqslant \min_{n_i \in \Omega_n}\{f(n_i)\}$.

By definition, $f_{max}(n)$ is the minimum of two valid upper bounds. The first one comes from a basic observation. If there are λ_n containers remaining to be retrieved in n, consider two cases:

- If $\lambda_n > S$, take the rth retrieval. If $S < r \leqslant \lambda_n$, then in order to perform this retrieval, there are at most $T-1$ containers blocking the target container so at most $T-1$ relocations are needed. When S or fewer containers remain, each container (except the lowest one) is at most relocated once, hence we need at most $S-1$ relocations. Combining these two facts, the maximum number of relocations is bounded by $(\lambda_n - S)(T-1) + (S-1) = ((\lambda_n - S)(T-1))^+ + (\min\{S, \lambda_n\} - 1)$.

- If $\lambda_n \leqslant S$, we know that $f(n_i) = b(n_i) \leqslant \lambda_n - 1 = ((\lambda_n - S)(T-1))^+ + (\min\{S, \lambda_n\} - 1)$.

This shows the validity of the first upper bound.

For the second upper bound, Zehendner et al. (2017) prove that, in the online case with a unique batch, the number of relocations performed by the leveling heuristic (L) is at most $\left(2\left\lceil \frac{\lambda_n}{S} \right\rceil - 1\right) B$, where B is the number of blocking containers. Since L is not using any information about batches (only the height of stacks), this result holds for both batch and online models with any number of batches. Let $n_i \in \Omega_n$, using this result and taking expectation over the retrieval order of containers not unveiled in n_i yet, we have $f(n_i) \leqslant f_L(n_i) \leqslant \left(2\left\lceil \frac{\lambda_n}{S} \right\rceil - 1\right) b(n_i)$. By taking the maximum over all $n_i \in \Omega_n$, the latter inequality results in the second upper bound. □

Lemma 7. *Let n be a chance node, and $w_{min} \in \{1, \dots, W\}$ be such that $\lambda_n = C - K_{w_{min}} + 1$ (i.e., the minimum batch in n). For each Stack s of n with $H^s \geqslant 1$ containers, let $(c^s_h)_{h=1,\dots,H^s}$ be the containers in s, where c^s_1 is the container at the bottom and $c^s_{H^s}$ is the container at the top (see Figure 11, for the case $H = H^s$). Finally, consider $C^s_{w_{min}} = \left|\left\{c^s_h = K_{w_{min}},\ h = 1, \dots, H^s\right\}\right|$. Then we have*

$$\min_{n_i \in \Omega_n} \{b(n_i)\} = \sum_{\substack{s=1,\dots,S \\ H^s \geqslant 1}} \left(H^s - C^s_{w_{min}} - \sum_{\substack{h=1,\dots,H^s \\ c^s_h \neq K_{w_{min}}}} \frac{\mathbb{1}\left\{c^s_h = \min\limits_{i=1,\dots,h}\left\{c^s_i\right\}\right\}}{\sum\limits_{i=1}^{h} \mathbb{1}\left\{c^s_h = c^s_i\right\}} \right), \tag{11}$$

and

$$\max_{n_i \in \Omega_n} \{b(n_i)\} = \sum_{\substack{s=1,\dots,S \\ H^s \geqslant 1}} \left(H^s - \sum_{\substack{h=1,\dots,H^s \\ c^s_h \neq K_{w_{min}}}} \frac{\mathbb{1}\left\{c^s_h = \min\limits_{i=1,\dots,h}\left\{c^s_i\right\}\right\}}{\sum\limits_{i=1}^{h} \mathbb{1}\left\{c^s_h = c^s_i\right\}} \right). \tag{12}$$

Proof of Lemma 7. Let n be a chance node, and $n_i \in \Omega_n$ one of its decision offspring, such that all containers in batch w_{min} have been revealed. Recall that $b(n_i) = \sum_{s=1,\dots,S} b^s(n_i)$, where $b^s(n_i)$ is the expected number of blocking containers in Stack s. First, for each Stack s such that $H^s = 0$, $b^s(n_i) = 0$. Hence $b(n_i) = \sum_{\substack{s=1,\dots,S \\ H^s \geqslant 1}} b^s(n_i)$.

For each Stack s such that $H^s \geqslant 1$, consider the containers in this stack $(c^s_i)_{i=1,\dots,H^s}$. Since all containers labeled $K_{w_{min}}$, i.e., from batch w_{min}, are known in n_i, we can write

$$b^s(n_i) = \sum_{h=1,\dots,H^s} \mathbb{P}\left[c^s_h \text{ is blocking in } n_i\right] = \sum_{\substack{h=1,\dots,H^s \\ c^s_h = K_{w_{min}}}} \mathbb{1}\left\{c^s_h \text{ is blocking in } n_i\right\} + \sum_{\substack{h=1,\dots,H^s \\ c^s_h \neq K_{w_{min}}}} \mathbb{P}\left[c^s_h \text{ is blocking in } n_i\right].$$

Fix $h \in \{1, \dots, H^s\}$ and $c^s_h \neq K_{w_{min}}$, then the proof of Lemma 2 uses the fact that $\mathbb{P}\left[c^s_h \text{ is blocking in } n_i\right] = 1 - \frac{\mathbb{1}\left\{c^s_h = \min\limits_{i=1,\dots,h}\left\{c^s_i\right\}\right\}}{\sum_{i=1}^{h} \mathbb{1}\left\{c^s_h = c^s_i\right\}}$. Finally, it is clear that $0 \leqslant \sum_{\substack{h=1,\dots,H^s \\ c^s_h = K_{w_{min}}}} \mathbb{1}\left\{c^s_h \text{ is blocking in } n_i\right\} \leqslant C^s_{w_{min}}$. Therefore, we can get the corresponding formulas. As a final remark, note that each of these bounds is tight. Indeed,

consider the offspring of n, in which all containers in batch w_{min} are in the decreasing (resp. increasing) order of retrieval from top to bottom, then this offspring has no (resp. $C^{s}_{w_{min}}$) blocking container(s). □

D Computational Experiments Tables

			Lower bounds			PBFS	Heuristics				
S	T	C	b	b_1	b_2		EG	EM	ERI	L	Rand.
5	3	8	1.64	1.66	1.66	1.70	**1.70**	**1.70**	**1.70**	1.82	2.34
	4	10	2.88	2.96	2.99	3.11	**3.11**	3.13	3.14	3.51	4.62
	5	13	4.61	4.88	5.01	5.32	5.40	**5.38**	5.57	6.17	8.00
	6	15	6.28	6.64	7.06	7.59	7.85	**7.81**	8.09	9.41	12.35
6	3	9	1.68	1.69	1.69	1.74	1.76	**1.74**	**1.74**	1.84	2.43
	4	12	3.54	3.61	3.63	3.68	3.69	**3.68**	3.68	4.11	5.59
	5	15	5.37	5.57	5.68	5.91	5.97	**5.94**	6.01	7.21	9.55
	6	18	7.19	7.52	7.68	8.23	8.38	**8.29**	8.63	10.05	13.53
7	3	11	2.82	2.86	2.88	2.88	**2.88**	**2.88**	2.89	2.96	4.12
	4	14	3.97	4.06	4.10	4.16	**4.17**	**4.17**	4.20	4.65	6.47
	5	18	6.49	6.65	6.74	6.97	7.05	**7.00**	7.07	8.46	11.27
	6	21	8.82	9.21	9.51	-	10.40	**10.35**	10.76	12.47	17.69
8	3	12	2.29	2.30	2.30	2.31	**2.31**	**2.31**	**2.31**	2.43	3.23
	4	16	4.68	4.73	4.75	4.82	**4.83**	**4.83**	**4.83**	5.41	7.48
	5	20	7.20	7.42	7.54	7.85	7.96	**7.93**	8.06	9.32	13.44
	6	24	9.52	9.85	10.09	-	11.10	**10.99**	11.34	13.29	19.28
9	3	14	2.98	2.98	2.98	3.00	**3.00**	3.01	**3.00**	3.19	4.54
	4	18	5.63	5.71	5.71	5.73	**5.73**	**5.73**	**5.73**	6.52	9.29
	5	23	8.58	8.69	8.77	-	9.05	**9.02**	9.12	11.16	15.57
	6	27	10.38	10.78	10.98	-	11.59	**11.58**	11.76	14.62	20.93
10	3	15	3.18	3.18	3.18	3.19	**3.19**	3.20	3.20	3.27	4.75
	4	20	6.20	6.23	6.23	6.28	6.30	**6.28**	**6.28**	6.98	10.41
	5	25	9.10	9.37	9.39	-	**9.60**	**9.60**	9.73	11.38	16.64
	6	30	11.91	12.28	12.44	-	13.01	**12.92**	13.15	15.93	23.40

Table 5: Results of experiment 1: Performance of *PBFS*, heuristics, and tightness of lower bounds for a fill rate of 50 percent in the batch model, in the case of small batches. Bold numbers highlight the best heuristic for a given problem size.

			Lower bounds			PBFS	Heuristics				
S	T	C	b	b_1	b_2		EG	EM	ERI	L	Rand.
5	3	10	2.83	2.97	3.01	3.08	**3.08**	**3.08**	3.11	3.33	4.16
	4	13	4.69	4.97	5.09	5.58	5.69	**5.64**	5.75	6.50	8.22
	5	17	7.58	8.52	8.92	-	**10.44**	10.48	11.04	12.23	15.27
	6	20	9.69	10.79	11.46	-	**14.53**	14.65	15.77	18.47	22.93
6	3	12	3.60	3.70	3.72	3.89	**3.89**	3.90	3.90	4.32	5.54
	4	16	6.20	6.59	6.78	7.28	**7.41**	7.45	7.61	8.55	11.29
	5	20	8.28	8.85	9.16	-	10.39	**10.38**	10.80	12.85	16.43
	6	24	11.67	12.22	12.54	-	15.19	**15.17**	16.14	19.33	25.03
7	3	14	3.85	3.89	3.91	3.97	**3.98**	**3.98**	4.02	4.40	6.05
	4	19	6.25	6.60	6.86	7.29	7.37	**7.36**	7.54	8.89	11.60
	5	23	9.72	10.24	10.55	-	11.75	**11.71**	12.30	14.92	19.65
	6	28	13.52	14.41	14.93	-	17.72	**17.63**	18.81	23.10	30.78
8	3	12	4.47	4.57	4.61	4.66	**4.66**	**4.66**	4.68	5.14	6.94
	4	21	7.62	7.85	7.98	8.29	**8.33**	8.35	8.43	9.76	13.26
	5	27	11.61	12.08	12.52	-	13.56	**13.47**	14.10	17.15	23.11
	6	32	15.60	16.39	16.78	-	**19.28**	19.51	20.85	25.89	34.66
9	3	18	4.81	4.96	4.99	5.10	**5.10**	5.12	5.14	5.66	7.81
	4	24	8.98	9.18	9.30	9.58	9.63	**9.61**	9.76	11.66	16.00
	5	30	13.16	13.90	14.29	-	**15.65**	15.79	16.75	20.03	27.41
	6	36	16.77	17.36	17.83	-	**20.38**	20.40	21.86	28.12	38.12
10	3	20	5.21	5.21	5.21	5.27	**5.28**	**5.28**	**5.28**	5.79	7.86
	4	27	9.18	9.54	9.71	-	**10.27**	10.29	10.37	12.16	16.91
	5	34	14.46	14.88	15.16	-	**16.13**	16.19	16.69	21.06	29.03
	6	40	19.55	20.24	20.66	-	23.33	**23.20**	24.46	32.11	44.07

Table 6: Results of experiment 1: Performance of *PBFS*, heuristics, and tightness of lower bounds for a fill rate of 67 percent in the batch model, in the case of small batches. Bold numbers highlight the best heuristic for a given problem size.

			Lower bounds			PBFSA	Heuristics				
S	T	C	b	b_1	b_2		EG	EM	ERI	L	Rand.
5	3	8	1.68	1.68	1.68	1.76	1.77	**1.76**	**1.76**	1.87	2.38
	4	10	2.96	3.02	3.05	3.33	**3.36**	3.37	3.38	3.67	4.85
	5	13	4.58	4.82	4.93	5.50	5.69	**5.65**	5.74	6.33	8.22
	6	15	6.31	6.72	6.96	7.81	8.28	**8.03**	8.35	9.49	12.40
6	3	9	1.66	1.67	1.67	1.73	1.75	**1.74**	**1.74**	1.82	2.43
	4	12	3.61	3.71	3.73	3.93	**3.99**	**3.99**	4.00	4.34	5.82
	5	15	5.38	5.57	5.64	6.11	**6.23**	**6.23**	6.31	7.16	9.65
	6	18	7.01	7.26	7.44	-	8.49	**8.36**	8.61	9.92	13.45
7	3	11	2.76	2.79	2.79	2.85	2.84	2.84	**2.83**	2.95	4.06
	4	14	4.02	4.12	4.15	4.24	4.31	**4.29**	4.31	4.73	6.52
	5	18	6.29	6.39	6.43	6.77	7.00	**6.92**	7.01	8.20	11.08
	6	21	8.69	9.12	9.35	-	10.60	**10.52**	10.91	12.61	17.79
8	3	12	2.30	2.31	2.31	2.31	**2.31**	2.32	2.32	2.37	3.19
	4	16	4.61	4.62	4.63	4.71	**4.74**	**4.74**	4.75	5.25	7.40
	5	20	7.31	7.46	7.52	-	**8.01**	**8.01**	8.09	9.33	13.25
	6	24	9.65	9.95	10.12	-	**11.37**	**11.37**	11.67	13.44	19.51
9	3	14	2.93	2.93	2.93	2.95	**2.96**	**2.96**	**2.96**	3.15	4.48
	4	18	5.56	5.58	5.59	5.69	5.74	**5.70**	**5.70**	6.33	9.07
	5	23	8.49	8.64	8.73	-	9.16	**9.12**	9.16	10.99	15.39
	6	27	10.38	10.69	10.90	-	11.77	**11.75**	11.95	14.65	20.95
10	3	15	3.15	3.16	3.16	3.15	**3.17**	**3.17**	**3.17**	3.25	4.72
	4	20	6180	6.20	6.21	6.28	6.35	**6.34**	**6.34**	6.92	10.27
	5	25	9.13	9.31	9.36	-	9.66	**9.63**	9.68	11.44	16.73
	6	30	12.09	12.35	12.51	-	13.38	**13.23**	13.42	16.33	23.80

Table 7: Results of experiment 2: Performance of *PBFSA*, heuristics, and tightness of lower bounds for a fill rate of 50 percent in the batch model with larger batches. Bold numbers highlight the best heuristic for a given problem size.

			Lower bounds			PBFSA	Heuristics				
S	T	C	b	b_1	b_2		EG	EM	ERI	L	Rand.
5	3	10	2.78	2.87	2.90	3.07	**3.08**	**3.08**	3.09	3.36	4.18
	4	13	4.71	4.92	5.00	5.70	5.81	**5.80**	5.89	6.60	8.23
	5	17	7.53	8.17	8.44	-	10.38	**10.32**	10.68	11.96	14.9815
	6	20	9.69	10.56	11.12	-	15.00	**14.91**	16.01	18.22	22.79
6	3	12	3.56	3.63	3.64	3.90	**3.90**	3.91	3.92	4.28	5.55
	4	16	6.12	6.48	6.61	7.17	7.36	**7.41**	7.48	8.44	11.02
	5	20	8.33	8.72	8.90	-	10.36	**10.30**	10.57	12.54	16.15
	6	24	11.77	12.41	12.80	-	15.94	**15.91**	16.86	19.83	25.45
7	3	14	3.95	4.01	4.02	4.14	4.17	**4.15**	**4.15**	4.56	6.11
	4	19	6.27	6.56	6.77	-	7.36	**7.35**	7.52	8.66	11.46
	5	23	9.81	10.27	10.54	-	**12.08**	12.10	12.49	14.78	19.67
	6	28	13.64	14.47	14.91	-	18.36	**18.26**	19.32	23.10	31.03
8	3	12	4.65	4.74	4.76	4.85	4.86	**4.85**	4.88	5.34	7.14
	4	21	7.58	7.86	7.99	-	**8.38**	8.42	8.50	9.82	13.20
	5	27	11.46	11.98	12.28	-	13.73	**13.60**	14.11	17.00	22.84
	6	32	15.45	16.28	16.72	-	19.94	**19.83**	21.21	25.73	34.70
9	3	18	4.85	4.98	5.02	5.14	**5.17**	5.20	5.20	5.67	7.77
	4	24	8.82	9.00	9.11	-	9.66	**9.60**	9.72	11.52	15.70
	5	30	13.15	13.84	14.18	-	**15.91**	15.96	16.82	20.16	27.35
	6	36	16.85	17.39	17.80	-	20.99	**20.83**	21.97	28.05	38.04
10	3	20	5.19	5.21	5.22	5.31	**5.31**	**5.31**	**5.31**	5.79	7.92
	4	27	9.40	9.66	9.82	-	10.47	**10.46**	10.54	12.25	16.96
	5	34	14.44	14.83	15.07	-	**16.29**	**16.29**	16.62	21.12	28.86
	6	40	19.49	20.24	20.66	-	23.83	**23.65**	24.96	32.04	44.12

Table 8: Results of experiment 2: Performance of *PBFSA*, heuristics, and tightness of lower bounds for a fill rate of 67 percent in the batch model with larger batches. Bold numbers highlight the best heuristic for a given problem size.

			Lower bounds			PBFS	Heuristics				
S	T	C	b	b_1	b_2		EG	EM	ERI	L	Rand.
5	3	8	1.64	1.66	1.66	1.70	**1.70**	**1.70**	1.71 (1.71)	1.82	2.34 (2.34)
	4	10	2.88	2.96	2.99	3.11	**3.11**	3.13	3.14 (3.20)	3.51	4.62 (4.62)
	5	13	4.61	4.88	5.01	5.32	**5.38**	**5.38**	5.57 (5.58)	6.16	8.00 (8.00)
	6	15	6.28	6.64	7.06	7.59	7.85	**7.80**	8.08 (8.29)	9.41	12.36 (12.35)
6	3	9	1.68	1.69	1.69	1.74	1.76	**1.74**	**1.74** (1.75)	1.84	2.43 (2.43)
	4	12	3.54	3.61	3.63	3.68	3.69	**3.68**	**3.68** (3.75)	4.11	5.59 (5.59)
	5	15	5.37	5.57	5.68	5.91	5.96	**5.94**	6.00 (6.18)	7.21	9.54 (9.54)
	6	18	7.19	7.52	7.68	8.23	8.38	**8.29**	8.62 (8.77)	10.05	13.53 (13.53)
7	3	11	2.82	2.86	2.88	2.88	**2.88**	**2.88**	2.89 (2.88)	2.96	4.11 (4.11)
	4	14	3.97	4.06	4.1	4.16	**4.17**	**4.17**	4.21 (4.20*)	4.66	6.47 (6.03*)
	5	18	6.49	6.65	6.74	6.97	7.04	**7.00**	7.07 (7.18)	8.45	11.27 (11.27)
	6	21	8.82	9.21	9.51	-	10.40	**10.35**	10.76 (10.98)	12.46	17.69 (17.69)
8	3	12	2.29	2.3	2.3	2.31	**2.31**	**2.31**	**2.31** (2.32)	2.43	3.23 (3.23)
	4	16	4.68	4.73	4.75	4.82	**4.83**	**4.83**	**4.83** (4.88)	5.41	7.49 (7.49)
	5	20	7.20	7.42	7.54	7.85	7.97	**7.94**	8.07 (8.27)	9.32	13.44 (13.45)
	6	24	9.52	9.85	10.09	-	11.10	**10.98**	11.34 (11.61)	13.29	19.29 (19.29)
9	3	14	2.98	2.98	2.98	3.00	**3.00**	**3.00**	**3.00** (3.00)	3.19	4.54 (4.54)
	4	18	5.63	5.71	5.71	5.73	**5.73**	**5.73**	**5.73** (5.80)	6.52	9.29 (9.29)
	5	23	8.58	8.69	8.77	-	9.05	**9.02**	9.12 (9.36)	11.16	15.56 (15.57)
	6	27	10.38	10.78	10.98	-	11.59	**11.58**	11.76 (12.09)	14.62	20.94 (20.93)
10	3	15	3.18	3.18	3.18	3.19	**3.19**	3.20	3.20 (3.20)	3.27	4.75 (4.75)
	4	20	6.20	6.23	6.23	6.28	6.30	**6.27**	6.28 (6.33)	6.98	10.41 (10.41)
	5	25	9.10	9.37	9.39	-	9.61	**9.60**	9.73 (9.80)	11.38	16.64 (16.63)
	6	30	11.91	12.28	12.44	-	13.01	**12.92**	13.15 (13.51)	15.92	23.41 (23.41)

Table 9: Results of experiment 3: Performance of heuristics and tightness of lower bounds for a fill rate of 50 percent in the online model with small batches. Bold numbers highlight the best heuristic for a given problem size. Numbers in parentheses are taken from Ku and Arthanari (2016).

			Lower bounds			PBFS	Heuristics				
S	T	C	b	b_1	b_2		EG	EM	ERI	L	Rand.
5	3	10	2.83	2.97	3.01	3.08	**3.08**	**3.08**	3.12 (3.10)	3.33	4.16 (4.16)
	4	13	4.69	4.97	5.09	5.58	5.68	**5.64**	5.75 (5.80)	6.50	8.22 (8.22)
	5	17	7.58	8.52	8.92	-	**10.45**	10.48	11.04 (11.15)	12.24	15.28 (15.28)
	6	20	9.69	10.79	11.46	-	**14.53**	14.65	15.77 (16.14)	18.46	22.93 (22.93)
6	3	12	3.6	3.7	3.72	3.89	**3.89**	3.90	3.90 (3.92)	4.32	5.53 (5.53)
	4	16	6.2	6.59	6.78	7.28	**7.41**	7.45	7.61 (7.68)	8.54	11.29 (11.28)
	5	20	8.28	8.85	9.16	-	**10.38**	**10.38**	10.80 (10.97)	12.85	16.42 (16.42)
	6	24	11.67	12.22	12.54	-	**15.17**	**15.17**	16.14 (16.65)	19.33	25.04 (25.03)
7	3	14	3.85	3.89	3.91	3.97	**3.98**	**3.98**	4.02 (4.01)	4.40	6.05 (6.05)
	4	19	6.25	6.6	6.86	7.29	7.37	**7.36**	7.54 (7.68)	8.89	11.60 (11.61)
	5	23	9.72	10.24	10.55	-	11.76	**11.71**	12.30 (12.64)	14.92	19.66 (19.65)
	6	28	13.52	14.41	14.93	-	17.70	**17.64**	18.82 (19.49)	23.10	30.77 (30.79)
8	3	12	4.47	4.57	4.61	4.66	**4.65**	4.66	4.68 (4.7)	5.14	6.94 (6.94)
	4	21	7.62	7.85	7.98	8.29	**8.32**	8.35	8.43 (8.5)	9.75	13.26 (13.25)
	5	27	11.61	12.08	12.52	-	13.56	**13.47**	14.10 (14.44)	17.14	23.11 (23.12)
	6	32	15.6	16.39	16.78	-	**19.27**	19.51	20.85 (21.72)	25.89	34.64 (34.63)
9	3	18	4.81	4.96	4.99	5.10	**5.10**	5.12	5.14 (5.19)	5.66	7.80 (7.80)
	4	24	8.98	9.18	9.3	9.58	9.63	**9.61**	9.76 (9.92)	11.66	16.01 (16.00)
	5	30	13.16	13.9	14.29	-	**15.65**	15.79	16.75 (16.97)	20.03	27.38 (27.39)
	6	36	16.77	17.36	17.83	-	**20.38**	20.40	21.87 (22.73)	28.13	38.11 (38.14)
10	3	20	5.21	5.21	5.21	5.27	**5.28**	**5.28**	**5.28** (5.30)	5.79	7.85 (7.86)
	4	27	9.18	9.54	9.71	-	**10.27**	10.29	10.37 (10.50)	12.15	16.92 (16.91)
	5	34	14.46	14.88	15.16	-	**16.13**	16.19	16.69 (17.23)	21.07	29.03 (29.03)
	6	40	19.55	20.24	20.66	-	23.33	**23.20**	24.46 (25.58)	32.11	44.08 (44.07)

Table 10: Results of experiment 3: Performance of heuristics and tightness of lower bounds for a fill rate of 67 percent in the online model with small batches. Bold numbers highlight the best heuristic for a given problem size. Numbers in parentheses are taken from Ku and Arthanari (2016).