

Evaluation of Dynamic Vector Bin Packing for Virtual Machine Placement

Zong Yu Lee and Xueyan Tang
Nanyang Technological University, Singapore
Email: zongyu001@e.ntu.edu.sg, asxytang@ntu.edu.sg

Abstract—Virtual machine placement is a crucial challenge in cloud computing for efficiently utilizing physical machine resources in data centers. Virtual machine placement can be formulated as a *MinUsageTime* Dynamic Vector Bin Packing (DVBP) problem, aiming to minimize the total usage time of the physical machines. This paper evaluates state-of-the-art *MinUsageTime* DVBP algorithms in non-clairvoyant, clairvoyant and learning-augmented online settings, where item durations (virtual machine lifetimes) are unknown, known and predicted, respectively. Besides the algorithms taken from the literature, we also develop several new algorithms or enhancements. Empirical experimentation is carried out with real-world datasets of Microsoft Azure. The insights from the experimental results are discussed to explore the structures of algorithms and promising design elements that work well in practice.

I. INTRODUCTION

Cloud computing implements the delivery of computing resources to clients through the Internet by adopting a pay-as-you-go model [1]. When a client makes a request, the cloud system typically creates a virtual machine, which uses a specified amount of computing resources such as CPU, memory, and storage in a physical machine. There are usually many physical machines in a data center. The cloud system needs to select a physical machine with sufficient resources available to host a requested virtual machine [1]. This is known as *virtual machine placement* [2]. To serve more clients, multiple virtual machines often share the resources in a physical machine through virtualization. However, if the packing of virtual machines into the physical machines is inefficient, it can cause a waste of resources, known as fragmentation. Fragmentation may lead to resource over-provisioning. According to [3], even a 1% reduction in fragmentation could save a cost in the order of \$100M per year for large cloud providers.

The virtual machine placement problem can be modelled as a *MinUsageTime* Dynamic Vector Bin Packing (DVBP) problem. *MinUsageTime* DVBP is a variation of the classical bin packing problem in which there are items (virtual machines) of vector sizes arriving and departing over time, and they are to be packed into bins (physical machines) such that the capacity constraints of bins are observed at any time. Unlike classical bin packing [4], [5], the objective of *MinUsageTime* DVBP is to minimize the total usage time of all bins, where the usage time of a bin is defined as the duration in which at least one item is placed in the bin (which indicates the time that the physical machine is used). Extensive research has been done to introduce online packing algorithms and theoret-

ically analyze their competitive ratios (worst-case ratio of the algorithm performance over the optimal solution). However, existing work suffers from two major deficiencies. First, most studies considered the single-dimensional case only, where both item sizes and bin capacities are represented by scalar values [6]–[15]. In contrast, virtual machines normally demand multiple types of resources (CPU, memory, storage, etc.) necessitating the consideration of the multi-dimensional case. Second, very little research evaluated the proposed algorithms empirically. The competitive ratio typically characterizes the algorithm performance in some corner cases. The practical performance of the packing algorithms for virtual machine placement remains unclear. To the best of our knowledge, only two studies investigated *MinUsageTime* DVBP in the multi-dimensional case and conducted empirical comparisons [16], [17]. Nevertheless, these evaluations used synthetic data only and not real-life datasets. Moreover, they did not explore the learning-augmented setting (where some item characteristics are predicted and they are prone to errors) and algorithms.

In this paper, we aim to bring the theory to practice by evaluating a comprehensive collection of packing algorithms using real-world datasets of Microsoft Azure. Besides the state-of-the-art algorithms taken from the literature, we also propose several new algorithms or enhancements and include them in the evaluation. We carry out extensive empirical evaluations of *MinUsageTime* DVBP algorithms in a variety of different settings. The evaluation results provide insights into the structures of algorithms that work well in practice. Our contributions are summarized as follows.

- We develop several new algorithms or enhancements, including Round Robin Next Fit, Nearest Remaining Time, modified Hybrid, modified RCP and PPE. We also conduct competitive analysis for the new algorithms and improve the analysis of some existing algorithms.
- Our evaluation shows that Any Fit (never open a new bin for an incoming item if it can fit in any current open bin) is a promising feature that can improve empirical packing efficiency if incorporated into the algorithm design.
- We show that categorizing items by their departure times outperforms that by their durations in the clairvoyant setting, but the former is less robust against prediction errors than the latter in the learning-augmented setting.
- Our evaluation also shows that dynamic categorization of bins performs better than static categorization.

The rest of this paper is organized as follows. Section II presents a formal definition of the *MinUsageTime* DVBP problem. Section III discusses the experimental setup. Sections IV, V and VI introduce and evaluate algorithms in the non-clairvoyant, clairvoyant and learning-augmented online settings respectively. Finally, Section VII concludes the paper.

II. PROBLEM DEFINITION AND PRELIMINARIES

The inputs to the *MinUsageTime* Dynamic Vector Bin Packing (DVBP) problem include a set of items \mathcal{R} , where each item $r \in \mathcal{R}$ is characterized by a d -dimensional vector $s(r)$ describing the item size and a time interval $I(r)$. The d -dimensional vector $s(r)$ represents the sizes of the item along d different dimensions. The time interval $I(r) = [I(r)^-, I(r)^+)$ represents the active interval of the item, where $I(r)^-$ denotes the arrival time and $I(r)^+$ denotes the departure time of the item. The interval length $I(r)^+ - I(r)^-$ is called the duration of the item. The item is said *active* during $I(r)$.

There is a sufficiently large pool \mathcal{P} of bins available to pack items. All the bins have the same capacity. The capacity of a bin is also represented by a d -dimensional vector. Without loss of generality, we assume that the capacity of a bin is given by $\langle 1, 1, \dots, 1 \rangle$, that is, all dimensions are normalized. Accordingly, it can be assumed that the size of an item along each dimension is no larger than 1.

In the context of virtual machine placement, each item models a virtual machine (VM). The d -dimensional size vector models the resource demands of the VM for d different types of resources (CPU, memory, storage, etc.). The arrival and departure of the item model the creation and termination of the VM respectively. Each bin corresponds to a physical machine (PM). The d -dimensional capacity vector models the resource capacities of the PM for the d types of resources.

A packing algorithm places each item into a bin while ensuring that at any time, the total size of the items placed in a bin does not exceed its capacity along all dimensions. The *usage time* of a bin is the length of time in which there is at least one item placed in the bin. The accumulated bin usage time is the sum of the usage times of all the bins. From another perspective, the accumulated bin usage time can be considered as the integration of the number of bins used at all time instants over the time horizon.

The objective of the *MinUsageTime* DVBP problem is to design a packing algorithm that minimizes the accumulated bin usage time. The key to optimizing the bin usage time is to increase the packing efficiency and hence the capacity utilization of bins. Inefficient placement of items would result in fragmentation and unnecessary use of additional bins. The total bin usage time of an optimal packing solution for a set of items \mathcal{R} has the following lower bound [16]:

$$\int_{-\infty}^{\infty} \left[\left\| \sum_{r \in \mathcal{R} \text{ and } t \in I(r)} s(r) \right\|_{\infty} \right] dt, \quad (1)$$

since the number of bins used at any time t must be at least the ℓ_{∞} -norm (the maximum value across all dimensions) of the aggregate size vector of all the active items at t .

We focus on the online setting, where items are released at their respective arrival times. Each item must be placed into a bin immediately when it arrives, without any information of the items that will arrive in the future. The online setting can further be divided into non-clairvoyant and clairvoyant settings depending on whether the departure time of an item is known at its arrival. In addition, a new setting known as the learning-augmented setting [18], [19] has emerged recently, which will be explained later. Theoretically, the competitive ratio of an online algorithm refers to the worst-case ratio between a solution constructed by the algorithm and an optimal solution [20].

III. EXPERIMENT SETUP

We use two publicly available datasets from Microsoft Azure and Huawei Cloud to empirically evaluate *MinUsageTime* DVBP.

Microsoft Azure [3] releases part of the workload on Microsoft's Azure Compute over a 14-day period (<https://github.com/Azure/AzurePublicDataset/blob/master/AzureTracesForPacking2020.md>).

The dataset consists of a VMType table and a VMRequest table. There are different virtual machine (VM) types with distinct resource demands and different physical machine (PM) types with distinct resource capacities. Each record in the VMType table specifies the resource usage of a virtual machine type on a physical machine type. The fields of the VMType schema include:

- virtual machine (VM) type,
- physical machine (PM) type,
- core, memory, HDD, SSD, NIC (fractional usage of core, memory, hard disk drive, solid-state drive, network bandwidth by the VM on the PM),


Each record in the VMRequest table specifies a virtual machine request. The fields of the VMRequest schema include:

- virtual machine (VM) type requested,
- starttime (the starting time of the VM),
- endtime (the termination time of the VM).


We clean and convert the dataset into *MinUsageTime* DVBP instances as follows.

Records. The VMType table has 488 different VM types, 35 different PM types, and a total of 4619 records (not all records of distinct VM and PM type pairs are available because some pairs are not compatible with each other [3]). The VMRequest table has 5,559,800 requests.

Time Information. The VMRequest table contains complete workload in the 14-day period of data collection. In the VMRequest table, there are records where the *endtime* field has null values, which means that the VM is not terminated by the end of data collection. We removed them because their VM lifetimes are unknown. Meanwhile, there are records where the *starttime* field has negative values, which means that the VM is created before the start of data collection; and there are also records where the *endtime* field has values greater than 14 days, which indicates that the data collector continues



(a) Number of VM requests with lifetime in $[2^{i-1}, 2^i]$ seconds



(b) Q-Q plot of the log lifetime of all VM requests

Fig. 1: Exploratory analysis of VM lifetimes

to monitor the created VM (at most until 90 days) and the VM is terminated within the monitoring period. The above requests with VM lifetimes partially falling outside the 14-day period are not included in the evaluation because we do not have the complete workload before and after the 14-day period. After the aforesaid records are removed, there are a total of 4,551,058 requests (about 82%) left whose VM lifetimes completely fall in the 14-day period. Figure 1a shows the distribution of these VM lifetimes and Figure 1b presents the normal quantile-quantile plot for the logarithm of VM lifetimes, which shows that the VM lifetimes almost follow a log-normal distribution.

Instances. Our evaluation assumes that there is a sufficiently large number of physical machines available and all the physical machines are of the same type with the same resource capacity. Thus, 35 possible *MinUsageTime* DVBP instances can be created from the dataset, one for each PM type. Given a PM type, the VM requests for VM types compatible with the PM type are extracted to form the input to the instance. Among 35 PM types, 2 of them do not have any VM requests because there is no request for compatible VM types. In addition, some VM types are exact duplications of other VM types in terms of compatibility and resource usage with respect to PM types. Therefore, some instances created are exactly the same. After

removing the two empty instances and the redundant instances (there are 4 of them), we end up having 29 distinct instances. We further exclude one particular instance in our evaluation because it is trivial (due to the low number and resource demands of VM requests in this instance, all the active VMs can fit into one PM at any time).

Resource Dimensions. In the VMType records, there are five resource dimensions ($d = 5$): core, memory, HDD, SSD and NIC. Some VMType records have null values or zero values in the HDD field. We assume that null values mean zero (i.e., the VM does not use any resource of the corresponding dimension). For some PM types, all the VM types compatible with the PM type have zero/null values in the HDD field. As a result, the resource dimension of HDD does not affect VM placement and thus can be ignored (the PM type probably does not have HDD resources at all). Hence, in such cases, only four resource dimensions ($d = 4$) are considered in our evaluation: core, memory, SSD and NIC.

We run different packing algorithms for each of the 28 instances to obtain the total usage time of all bins. We compute the ratio of the total usage time produced by an algorithm over the best-known lower bound of the optimal solution as defined in Equation (1). The ratio is intended to capture the normalized performance of an algorithm compared with the optimal solution for a particular instance. We refer to it as the *performance ratio*. The lower bound is used instead of the exact optimal solution because *MinUsageTime* DVBP is NP-hard, and it is computationally expensive to derive the exact optimal solution. Obviously, the actual ratio between the total bin usage times of an algorithm and the optimal solution is never higher than our performance ratio. We present the box plot of the performance ratios to examine the performance of an algorithm across all instances. We use the common box plot showing the first quartile (25th percentile) and the third quartile (75th percentile) as a box, and the boundaries of the whiskers based on 1.5 times the interquartile range (which is defined as the distance between the third and first quartiles).

Huawei Cloud also releases a VM placement dataset [21] (<https://github.com/huaweicloud/VM-placement-dataset>) that has a schema largely similar to the Microsoft Azure dataset but has a much smaller size (116,313 VM requests after cleaning and only two resource dimensions: CPU and memory). The evaluation results using the Microsoft Azure and Huawei Cloud datasets show similar trends. Due to space limitations, we shall focus on presenting the results of the Microsoft Azure dataset in this paper. The results of the Huawei Cloud dataset are provided in an extended version of this paper [22].

IV. NON-CLAIRVOYANT ALGORITHMS

This section introduces and evaluates algorithms in the non-clairvoyant setting – the departure time of each item is not known at its arrival and hence cannot be used in making packing decisions. Since the packing algorithm is blind to item departure times, it cannot avoid mixing items with short and long durations in the same bin. Indeed, it has been proved that

no deterministic packing algorithm can achieve a competitive ratio lower than μ (the max/min item duration ratio) in the single-dimensional case [11].

In the packing process, a bin is said *opened* when it receives the first item. When all the items in a bin depart, the bin is said *closed*. Any Fit is an important feature of many online packing algorithms [11]. It means that the algorithm never opens a new bin for an incoming item if the item can fit into any current open bin. Any Fit is generally good for *MinUsageTime* DVBP because item placement will always try to make use of open bins that do not incur additional bin usage time immediately. Theoretically, it has been shown that no Any Fit packing algorithm can achieve a competitive ratio lower than $(\mu + 1)d$ in the d -dimensional case [16].

A. First Fit and Most Recently Used

First Fit [11], [12] and Most Recently Used (MRU) [7], [16] are Any Fit algorithms. These algorithms maintain different orderings of the open bins. Whenever an item arrives, the algorithm goes through the particular order of the open bins and places the item into the first bin that has enough capacity available for it in all resource dimensions.

First Fit arranges the bins in ascending order of open time, so an incoming item is placed in the earliest-opened bin. Most Recently Used arranges the bins in descending order of last access time, where a bin is said accessed when a new item is placed in the bin. In this way, receiving a new item moves the bin to the front of the ordering – (some literature named this algorithm Move to Front [7], [16]). First Fit can achieve a competitive ratio of $(\mu + 2)d + 1$ [12], [16], which in the single-dimensional case is pretty close to the lower bound of μ on the competitiveness of any deterministic algorithm [11]. MRU can achieve a competitive ratio of $(2\mu + 1)d + 1$ [16].

B. Next Fit and Round Robin Next Fit

In Next Fit [7], only one open bin accepts new item placement at any time. Once an incoming item cannot fit into this bin, this bin will no longer accept any further item, and a new bin will be opened to receive the incoming item and subsequently arriving items.

Next Fit originates in a static setting where items never depart. The main feature of Next Fit is to prevent mixing items arriving far apart in the same bin. In a dynamic setting, Next Fit is not likely to perform very well because all the open bins except the most recently opened one are not allowed to receive new items. An earlier-opened bin may be able to accommodate new items after some items placed in it depart. Next Fit is not an Any Fit algorithm as a new bin may be opened for an incoming item even if the item can fit into an earlier-opened bin. We propose Round-Robin Next Fit as a new algorithm that converts the concept of Next Fit into an Any Fit algorithm, which is more suitable for a dynamic setting. When an incoming item cannot fit into the open bin currently under consideration, instead of opening a new bin, we can loop around to search for any possible open bin that can accommodate the item. That is, we maintain the open bins

in ascending order of open time (same as First Fit) and keep a flag indicating the open bin currently under consideration. For each incoming item, the flag initially points to the bin receiving the last arrived item and the flag will be moved around in a round-robin manner to search for an open bin for placing the incoming item. Only when no open bin can accommodate the incoming item will a new bin be opened. Next Fit achieves a competitive ratio of $2\mu d + 1$ [7], [16]. We prove that Round-Robin Next Fit achieves a competitive ratio of $(2\mu + 1)d + 1$. The main idea is as follows. The usage time of each bin can be partitioned into alternating periods of receiving items (the flag points to the bin) and not receiving items (the flag does not point to the bin). At any time, there is only one bin in the receiving period. The non-receiving period of any other open bin b is initiated by an item r placed into another bin (i.e., r causes the flag to move from b to the latter bin). By the algorithm definition, the total size of items in the non-receiving period of b and the item r must exceed the bin capacity. If we duplicate the item r and extend its duration by μ , and extend the durations of the items in b by $\mu + 1$, all these items would be active throughout the non-receiving period of b . Hence, the number of open bins is bounded by the total size of active items in the original and duplicate sets, which in turn is bounded by the optimal packing solution. The full proof is deferred to an extended version of this paper [22]. We also construct an adversary example to show that the competitive ratio of Round Robin Next Fit is at least $2\mu d$ in [22].

C. Best Fit

The guiding principle of Best Fit [11] is to reduce the wastage of bin capacity. Among all open bins, an incoming item will be placed into the open bin whose available capacity “best” fits the size of the item, i.e., the open bin that will leave the least remaining capacity after accepting the item. In the single-dimensional case, this simply means to place the item into the most-loaded open bin. In the multi-dimensional case, however, the item size and bin capacity are described by vectors due to multiple resource types. Hence, there are many different ways to define what is the “least” vector. We adopt the definition of ℓ_p -norm to characterize vectors. Let $\vec{a} = \langle a_1, a_2, \dots, a_d \rangle$ be the size of an incoming item, and $\vec{b} = \langle b_1, b_2, \dots, b_d \rangle$ be the available capacity of an open bin. Then, the remaining capacity of the bin after accepting the item is given by $\vec{b} - \vec{a} = \langle b_1 - a_1, b_2 - a_2, \dots, b_d - a_d \rangle$. We explore three common ℓ_p -norms to compute the fit score: ℓ_1 -norm (sum of available capacity over all dimensions: $\sum_{1 \leq i \leq d} (b_i - a_i)$), ℓ_2 -norm (square root of the sum of squares: $\sqrt{\sum_{1 \leq i \leq d} (b_i - a_i)^2}$), and ℓ_∞ -norm (maximum available capacity across all dimensions: $\max_{1 \leq i \leq d} (b_i - a_i)$). The incoming item will be placed in the open bin with the lowest fit score that can accommodate it. It has been proved that Best Fit cannot achieve any bounded competitive ratio, even in the single-dimensional case [11].

Experimental Comparison: Figure 2 compares the performance ratios of different norms for Best Fit. In general, ℓ_∞ -




Fig. 2: Performance of Best Fit

norm works the best for Best Fit. As seen from Figure 2, l_∞ -norm produces the lowest mean performance ratio. A possible reason is that l_∞ -norm captures the dimension with the maximum available capacity and acts as an upper bound of the available capacity in other dimensions. As such, l_∞ -norm is a good representation of the available capacity of a bin. In what follows, we shall focus on using the l_∞ -norm in our evaluation.

D. Experimental Comparison in Non-clairvoyant Setting

Figure 3 presents the box plot for various algorithms in the non-clairvoyant setting. To facilitate the comparison among better-performing algorithms, we limit the cap of the vertical axis (performance ratio) to 2.5, which cuts off the box/outliers for some algorithms.¹ We could draw the following insights:

Best Performer: The First Fit algorithm produces the lowest mean performance ratio among all the algorithms tested, which concurs with that it has the lowest competitive ratio.

Any Fit Feature: It is beneficial to enhance a non-clairvoyant algorithm with the Any Fit feature. Round Robin Next Fit considerably improves the performance ratio over Next Fit, though they have very similar competitive ratios. This shows that by utilizing open bins as much as possible, the Any Fit feature can help improve the packing efficiency dramatically.

Best Fit and MRU: Best Fit performs close to First Fit, because minimizing the leftover capacity in the open bins makes item placement concentrate in fewer bins. Most Recently Used (MRU) performs close to Round Robin Next Fit. We note that earlier studies [16], [17] showed that MRU is empirically the best in their evaluations with synthetic data, but this is not the case in our evaluation with real-world data of VM placement. This could be due to different distributions of item sizes and durations (they were generated from uniform distributions in [16] and [17]).

V. CLAIRVOYANT ALGORITHMS

This section presents and evaluates algorithms in the clairvoyant setting – the departure time (or equivalently the duration) of each item is known upon its arrival. Algorithms in the clairvoyant setting usually utilize the departure time or item duration information to categorize the items and




Fig. 3: Performance of non-clairvoyant algorithms

bins. The items of each category are placed in the bins of the same category in a First Fit manner unless stated otherwise. Note, however, that algorithms of this paradigm are not Any Fit algorithms if there are multiple categories. This is nevertheless necessary for packing algorithms to achieve competitive ratios of $o(\mu d)$, since the lower bound of $(\mu + 1)d$ on the competitiveness of Any Fit algorithms still applies in the clairvoyant setting. In what follows, we introduce the details of each algorithm and its intuition. We start with algorithms categorizing items based on departure times, followed by algorithms conducting categorization based on item durations.

A. Classify By Departure Time

Wastage of bin capacity is likely to occur when some items in a bin depart while other items in the bin do not depart. With the knowledge of item departure times, an intuitive approach to reduce the bin usage time is to place items with similar departure times in the same bin [6]. In this way, the items in a bin would depart at around the same time, after which the bin can be closed timely. As a result, the bin usage time can be saved. The Classify-By-Departure-Time algorithm operates with a parameter ρ defined as the width of the departure time range. It partitions the time horizon into disjoint intervals of length ρ and puts all the items with departure times in each interval into one distinct category. The algorithm can achieve a competitive ratio of $O(\sqrt{\mu})$ in the single-dimensional case by setting ρ to the geometric mean of the maximum and minimum item durations if they are known upfront [6]. There is no competitive analysis for the multi-dimensional case in the existing literature. Therefore, we conduct experiments for different ρ values.

Experimental Comparison: As seen from the box plot in Figure 4, the median and mean of performance ratios exhibit a U-shaped trend. When ρ is small, the algorithm creates many categories, which potentially degrades the packing efficiency because items of different categories are packed into separate bins. On the other hand, when ρ is large, the algorithm creates fewer categories, but it degenerates towards the First Fit algorithm which does not exploit any departure time information. As such, there is a trade-off between overusing and underusing the departure time information of items in the packing. For the Microsoft Azure dataset, setting ρ to 0.25 day appears to produce the best performance ratios. Thus, we shall use this ρ setting in subsequent comparisons.

¹Similar caps on the vertical axis are also imposed in subsequent figures.




Fig. 4: Performance of Classify-By-Departure-Time

B. Nearest Remaining Time (NRT)

We propose an algorithm called Nearest Remaining Time that extends the idea of classification by departure time. The central intuition is again that items with similar departure times should be placed in the same bin. However, Classify-By-Departure-Time is not an Any Fit algorithm. To improve the algorithm design, we integrate the Any Fit concept by dynamically defining the departure time range. Note that the latest departure time of all items currently placed in a bin indicates the bin closing time (referred to as the *indicated closing time* hereafter). Among all the open bins with enough available capacity, we can place an incoming item into the bin that has the indicated closing time nearest to the departure time of the item. From an alternative perspective, if we arrange the bins in ascending order of indicated closing time, this strategy implies that each bin will receive items with departure times in the range from the midpoint between the previous bin and this bin to the midpoint between this bin and the next bin. Therefore, it dynamically defines the range of accepting departure times for each bin.

To enhance the algorithm design, we can further distinguish between case (a) where the bin's indicated closing time is no earlier than the item departure time and case (b) where the bin's indicated closing time is earlier than the item departure time. In case (a), placing the item into the bin does not increase the bin usage time, whereas in case (b), it extends the bin's indicated closing time and hence increases the bin usage time. Thus, to reduce the bin usage time, priorities can be given to the bins of case (a) over the bins of case (b) in item placement. Therefore, we evaluate two versions of the Nearest Remaining Time (NRT) strategy as follows:

- **Standard NRT** – Consider case (a) and case (b) together and choose the bin with indicated closing time nearest to the departure time of an item for placement.
- **Prioritized NRT** – An item is placed in a bin of case (b) only if the item cannot fit into any bin of case (a). For each case, choose the bin with indicated closing time nearest to the item departure time for placement.

Standard NRT does not have a bounded competitive ratio. In fact, an adversary can force Standard NRT to maintain many open bins indefinitely, whereas the optimal packing solution needs just one bin. First, the adversary releases n items each of size $\frac{1}{2} + \epsilon$ ($\epsilon > 0$) to open n bins b_1, b_2, \dots, b_n with increasing




Fig. 5: Performance of NRT

indicated closing times forming an arithmetic sequence. Then, the adversary releases items in rounds. In each round, the adversary releases n items each of size ϵ , where the departure time of the i -th item falls between the indicated closing times of b_{n-i+1} and b_{n-i} and is slightly closer to the former than the latter. Standard NRT will place the i -th item in bin b_{n-i+1} and extend b_{n-i+1} 's indicated closing time. After each round, the indicated closing times of the bins again form an arithmetic sequence. As $\epsilon \rightarrow 0$, all the active items can be placed into just one bin after the initial items depart. Hence, an optimal packing solution will only need one bin to accommodate all the items. Thus, the competitive ratio is not bounded as $n \rightarrow \infty$.

We prove that Prioritized NRT achieves a competitive ratio of $(\mu + 2)d + 1$, which is rather close to the lower bound of $(\mu + 1)d$ on the competitiveness of Any Fit algorithms. The main idea is as follows. At any time t , we sort the open bins by the earliest time instant when the bin has its indicated closing time greater than t due to a new item placed in it (call it the triggering item). By the algorithm definition, for any j , the total size of active items in the $(j - 1)$ -th bin and the triggering item of the j -th bin must exceed the bin capacity. If we duplicate the items and extend their durations by $\mu + 1$, all these items would be active at time t . Hence, the number of open bins at t is bounded by the total size of items active at t in the original and duplicate sets, which in turn is bounded by the optimal packing solution. The full proof is deferred to an extended version of this paper [22] due to space limitations.

Experimental Comparison: Figure 5 compares the performance ratios of Standard NRT and Prioritized NRT. It can be seen that Prioritized NRT works significantly better than Standard NRT. Though they both have the Any Fit feature, Prioritized NRT prevents the bin closing time from increasing as far as possible and thus avoids additional bin usage time. Therefore, we shall focus on Prioritized NRT in our evaluation.

C. Greedy

The Greedy algorithm [17] also makes placement decisions based on the relation between the bin's indicated closing time and the item's departure time. Different from NRT, it places an incoming item into the bin with the latest closing time, among all the open bins that can accommodate the item. If the item's departure time is later than the indicated closing time of the chosen bin, the bin will have its indicated closing time extended to the item's departure time after placement.




Fig. 6: Performance of Classify-By-Duration

Li *et al.* [17] proved that the Greedy algorithm can achieve a competitive ratio of $3\mu d + 1$ in the d -dimensional case. We further improve its competitive analysis in an extended version of this paper [22] and show that the Greedy algorithm achieves a competitive ratio of $(\mu + 2)d + 1$, following the same proof idea as Prioritized NRT.

D. Classify By Duration

Another strategy is to place items with similar durations in the same bin [6]. This is motivated by the observation that in the non-clairvoyant setting, the competitiveness of online packing is limited by the variation of item durations. By reducing the variation of item durations in the same bin, it is possible to improve the performance of the packing algorithm. Specifically, items are classified into categories such that the max/min item duration ratio for each category is a given constant β . Each category i includes all the items with durations in the range $[\beta^{i-1}, \beta^i]$ for an integer i (where i can be negative, zero or positive). If μ is not known in advance, the Classify-By-Duration algorithm achieves a competitive ratio of $O(\log \mu)$ in the single-dimensional case [6].

Experimental Comparison: We test different β values. As shown by the box plot in Figure 6, the median and mean of performance ratios generally have a U-shaped trend. Similar to Classify-By-Departure-Time, when β is small, the algorithm creates many categories, which potentially downgrades the packing efficiency; when β is large, the algorithm creates fewer categories and approaches the First Fit algorithm. For the Microsoft Azure dataset, it appears that when $\beta = 2$, the median performance ratio is the lowest. Hence, in what follows, we shall set $\beta = 2$ in our evaluation.

E. Hybrid Algorithm

The Hybrid algorithm (HA) [8] combines the First Fit and Classify By Duration strategies. It categorizes items based on their durations and arrival times. Each category (i, c) includes all the items with durations in the range $[2^{i-1}, 2^i)$ for an integer i and arrival times falling in the interval $[(c-1) \cdot 2^i, c \cdot 2^i)$ for an integer c . The Hybrid algorithm defines two types of bins – general bins and category bins. General bins can receive items of any categories, while each category bin can receive items of a particular category only. Items of each category (i, c) are placed into the general bins when their total size in

the general bins is no larger than a threshold $\frac{1}{2\sqrt{i}}$, and placed into the category bins when their total size in the general bins exceeds the threshold. The First Fit strategy is adopted for item placement among the general bins as well as among the category bins of the same category. The intuition behind the Hybrid algorithm is to implement the Classify By Duration strategy only when there is sufficient load of each category and prevent items from being unnecessarily distributed to many separate bins. Azar and Vainstein [8] proved that the Hybrid algorithm achieves a competitive ratio of $O(\sqrt{\log \mu})$, which is asymptotically optimal (best achievable) in the clairvoyant setting. The original Hybrid algorithm is for the single-dimensional case. To adapt it to the multi-dimensional case, we define the total size of active category- (i, c) items in the general bins as the ℓ_∞ -norm of their aggregate size vector, i.e., $\|\sum_{r \in \mathcal{R}^*} s(r)\|_\infty$, where \mathcal{R}^* is the set of active category- (i, c) items placed in the general bins. We prove that such adaptation allows the Hybrid algorithm to achieve a competitive ratio of $O(d\sqrt{\log \mu})$. Due to space limitations, we defer the full proof to an extended version of this paper [22]. We also note that the original Hybrid algorithm assumes the minimum item duration is at least 1, so that i in the category index is always a positive integer (hence \sqrt{i} is well-defined). In our implementation, for each instance created from the dataset, we find the minimum item duration and the power-of-2 range $[2^{z-1}, 2^z)$ including this minimum item duration. We set $i = 1$ for this range $[2^{z-1}, 2^z)$ and set $i = j - z + 1$ for each subsequent range $[2^{j-1}, 2^j)$ where $j > z$ accordingly.

Recently, Liu and Tang [13] found that item categorization can be simplified while maintaining the asymptotic optimal $O(\sqrt{\log \mu})$ competitiveness of packing. It is sufficient to use a more coarse-grained item categorization: define categories based on item durations only. That is, each category i includes all the items with durations in the range $[2^{i-1}, 2^i)$ for an integer i . This can significantly reduce the number of categories. The concepts of general bins and category-specific bins remain unchanged in the Hybrid framework. We also include this Reduced Hybrid algorithm in our evaluation. Again, we adapt it to the multi-dimensional case by using the ℓ_∞ -norm to characterize the total size of active category- i items in the general bins. We prove that the Reduced Hybrid algorithm also achieves a competitive ratio of $O(d\sqrt{\log \mu})$ in [22].

Very recently, Li *et al.* [17] proposed a technique called *direct-sum* to extend a packing algorithm in the single-dimensional case to the multi-dimensional case. The idea is to divide all items into d classes based on the dimension in which the item has the largest size. Then, the packing algorithm is applied to each class of items separately, since the packing needs to look at only the dimension of largest item sizes. Li *et al.* [17] proved that such adaptation introduces a multiplicative factor of d (the number of dimensions) to the competitive ratio. We also evaluate the Hybrid and Reduced Hybrid algorithms adapted using the *direct-sum* technique to the d -dimensional case, which also have a competitive ratio of $O(d\sqrt{\log \mu})$.

Experimental Comparison: Figure 7 shows the experimental results. Comparing the left two algorithms, Reduced




Fig. 7: Performance of Hybrid algorithms

Hybrid considerably improves the packing efficiency over Hybrid. This is because eliminating the factor of item arrival times cuts down the number of categories which in turn leads to better utilization of bin capacity. Comparing our ℓ_∞ -norm characterization (left two algorithms) with the direct-sum adaptation (right two algorithms), our adaptation is much more effective than direct-sum. This is because direct-sum essentially further increases the number of item categories by a factor of d , which downgrades the packing efficiency. In what follows, we shall focus on the Reduced Hybrid algorithm using our ℓ_∞ -norm characterization in the evaluation.

F. Experimental Comparison in Clairvoyant Setting

Figure 8 presents the box plot for various algorithms in the clairvoyant setting. We could draw the following insights:

Best Performer: Prioritized NRT performs very well and better than all the other algorithms. This demonstrates the effectiveness of matching item departure times with bin closing times and prioritizing bins that do not need to extend their closing times upon accepting items.

Departure Time vs. Item Duration Information: Classify-By-Departure-Time, Prioritized NRT and Greedy utilize the departure time information of items in the placement, whereas Classify-By-Duration and Reduced Hybrid utilize the item duration information in the placement. From Figure 8, algorithms using departure time information tend to perform better than algorithms using item duration information. While categorizing items based on duration can optimize the asymptotic competitive ratios from the theoretical perspective, it is practically less successful than leveraging the departure time information in reducing the total bin usage time. Aligning the departure times of items in the same bin allows the bin to be closed timely, thereby reducing the wastage of bin capacity effectively.

Any Fit Feature: Prioritized NRT and Greedy are Any Fit algorithms and they perform better than others which are not Any Fit algorithms. Even though Any Fit algorithms are never better than $O(\mu d)$ in competitiveness, empirically Any Fit is still promising. A good algorithm design is to integrate Any Fit with the departure time information of items.

VI. LEARNING-AUGMENTED ALGORITHMS

In practical applications, it is usually difficult to know the exact item durations (such as VM lifetimes or job lengths) at item arrivals. Item durations often need to be predicted




Fig. 8: Performance of clairvoyant algorithms

based on history by using machine learning techniques. The prediction accuracy, however, varies widely with techniques and environments. Hence, it will be important to evaluate whether the packing efficiency is resilient to prediction errors. Making use of predictions to enhance online algorithm design has led to an arising research area called learning-augmented algorithms [23]. This section presents and evaluates algorithms in the learning-augmented setting – a predicted duration of each item is known upon its arrival, but it may not always be the same as the real duration of the item and can be arbitrarily different from the latter in the worst case. The goal is to make wise use of the predictions to improve the algorithm’s performance.

A. Robust and Consistent Packing (RCP) and Packing with Prediction Error (PPE)

Liu and Tang [13] developed a learning-augmented algorithm called Robust and Consistent Packing (RCP). It extends the Hybrid framework by defining four types of bins (large bins, general bins, base bin, category bins) and dynamically turning item categories on and off. The major extensions are as follows:

- Items are categorized according to their predicted durations rather than real durations;
- Each large item (of size larger than $1/2$) is placed in a separate large bin.²
- The threshold setting is simplified by applying a uniform and adjustable threshold value $1/\sqrt{x}$ for all the categories, where x is the number of distinct categories seen by the algorithm so far (which typically grows as time goes on).
- Small items (of size no larger than $1/2$) go to the general bins first. For each category, when the total size of active items in the general bins exceeds the threshold $1/\sqrt{x}$, incoming items do not go to the category bins immediately. Instead, such incoming items of all categories go to a base bin. The base bin is a logical concept. At any time, there is at most one base bin. When the total size of items in the base bin exceeds $1/2$, the base bin is converted to a category bin of the category of a selected item in the bin, turning on that category.

²A “large bin” means a bin for placing large items only. It does not mean that the bin has a larger capacity than others.

- When a category is on, incoming items of the category (exceeding the threshold) go to the category bins. When the total size of items in the category bins falls low, the category is turned off, after which incoming items of the category (exceeding the threshold) go to the base bin again.

The base bin is designed to guarantee the resilience of the algorithm against terrible predictions. In the worst case, all incoming items can be very small in size and belong to different categories due to mispredictions. Directly applying the Hybrid framework will open many bins concurrently. In this case, the base bin is used to receive all items and avoid opening unnecessary bins. Liu and Tang [13] proved that the RCP algorithm achieves $O(\mu)$ consistency (competitive ratio under perfect predictions) and $O(\sqrt{\log \mu})$ robustness (competitive ratio under terrible predictions), matching the asymptotic optimal competitive ratios in the clairvoyant and non-clairvoyant settings respectively.

More recently, Liu and Tang [14] further refined the RCP algorithm and named it Packing with Prediction Error (PPE). The main idea is to monitor the prediction errors of departed items (note that the real duration of an item is revealed when it departs) and adjust the threshold value accordingly. Specifically, the threshold value is set to α/\sqrt{x} , where α is an online estimate of the maximum prediction error using a guess-and-double approach. As a result, PPE achieves an asymptotic optimal competitive ratio of $O(\min\{\max\{\epsilon \cdot \sqrt{\log \mu}, \epsilon^2\}, \mu\})$ over the entire spectrum of prediction errors, where ϵ is the maximum multiplicative prediction error among all items (the multiplicative prediction error is the ratio between the real duration and the predicted duration). We include both RCP and PPE in our evaluation. Similar to the Hybrid algorithm, we adapt RCP and PPE to the multi-dimensional case by using the ℓ_∞ -norm to characterize the total size of a set of items.

In addition, we find that placing each large item (of size larger than $1/2$) in a separate large bin is not really necessary for RCP and PPE to achieve asymptotic optimal competitive ratios. Intuitively, these large bins could result in significant wastage of bin capacity, since each bin hosts one item only. In fact, small items may still be placed in large bins to improve the bin utilization. Therefore, we also propose variants of RCP and PPE which eliminate the concept of large bins. Specifically, we modify the algorithms as follows:

- All items (irrespective of large or small) go to the general bins first, when the total size of active items in the same category does not exceed the threshold $1/\sqrt{x}$ in the general bins.
- For each incoming large item (exceeding the threshold), if the category of the item is on, the item goes to the category bins.
- Otherwise, if the item can fit into the remaining capacity of the base bin, it goes to the base bin. After receiving the large item, the base bin is converted to a category bin, since the total size of items in it exceeds $1/2$.
- Finally, if the item cannot fit into the remaining capacity

of the base bin, a category bin is opened for the item, and as a result, the category of the item is turned on immediately.

B. Lifetime Alignment (LA)

Recently, Barbalho *et al.* [24] designed a Lifetime Alignment (LA) algorithm for VM placement with lifetime predictions. It is built upon the Classify-By-Duration strategy, i.e., categorizing items (VMs) according to their predicted durations (lifetimes). The Lifetime Alignment algorithm introduces a new idea to dynamically classify bins (PMs) based on their predicted remaining usage times. The predicted remaining usage time of a bin is given by the duration from the current time to the latest predicted departure time of all items currently placed in the bin (or equivalently the longest predicted remaining duration of all items in the bin). This contrasts the static bin classification used by the Classify-By-Duration, Hybrid, RCP and PPE algorithms, where each bin is always associated with a fixed item category. By the aforesaid dynamic classification, a bin can be associated with different item categories as time goes on. The intuition is to prevent placing an item of long duration into a bin that is about to be closed (all items in it will depart soon, though they had long durations). Lifetime Alignment is an Any Fit algorithm. Let X_0, X_1, X_2, \dots be disjoint ranges of durations for categorization, where X_0 contains the smallest values. For each category except X_0 , an incoming item will first consider the open bins of the same category for placement and if unsuccessful, then consider other open bins. For category X_0 (items with shortest durations), an incoming item will consider all open bins together for placement, since these items can help filling the leftover capacity of any bins and do not impact their usage times. Among the open bins considered with enough available capacity, the Best Fit strategy is used for item placement. In our evaluation, we use the ℓ_∞ -norm to compute the fit score (as discussed in Section IV-C). We evaluate two different ways of categorization:

- **Binary classification:** There are only two categories: short-lived and long-lived, by setting $X_0 = [0, 120)$ minutes and $X_1 = [120, \infty)$ minutes following [24].
- **Geometric ranges:** Similar to the RCP and PPE algorithms, we set $X_0 = [0, 1)$ second, and $X_i = [2^{i-1}, 2^i)$ seconds for each $i \geq 1$.

Neither version of Lifetime Alignment has a bounded competitive ratio. In fact, when all the items are short items belonging to category X_0 , Lifetime Alignment is reduced to Best Fit which has been proved not to have a bounded competitive ratio [10], [11].

C. Experimental Comparison in Learning-augmented Setting

Additional Experiment Setup: To evaluate learning-augmented algorithms, we further generate a predicted duration of each item (a predicted lifetime of each VM requested) in the experiment setup. Recall that in a *MinUsageTime* DVBP instance, each item r is characterized by an active interval $I(r) = [I(r)^-, I(r)^+)$, where $I(r)^-$ and $I(r)^+$ are the arrival




Fig. 9: Classify-By-Departure-Time and Classify-By-Duration




Fig. 10: RCP and PPE




Fig. 11: Lifetime Alignment, Classify-By-Duration and Reduced Hybrid

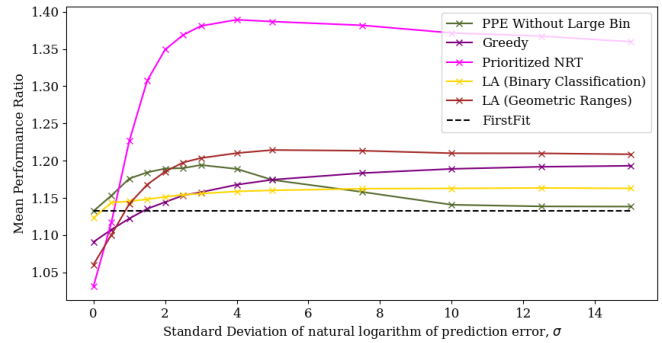


Fig. 12: Mean performance ratios of selected algorithms

and departure times of the item respectively. The real duration of the item is given by $Rdur(r) = I(r)^+ - I(r)^-$. We generate a predicted duration of the item by scaling the real duration $Rdur(r)$ with a multiplicative prediction error. For each item, we randomly generate a multiplicative prediction error δ according to a log-normal distribution,³ where the mean and standard deviation of δ 's natural logarithm are μ and σ respectively. The log-normal distribution simulates rare but significant prediction failures of the learning algorithm. It was used in the experimentation of other work on learning-augmented algorithms [19]. On generating δ , the predicted duration is set to $Pdur(r) = \delta \cdot Rdur(r)$, which means that the item is predicted to depart at time $I(r)^- + Pdur(r)$. We set the mean $\mu = 0$ and vary the standard deviation σ in the experiment. $\sigma = 0$ means the prediction error is always $e^\mu = e^0 = 1$, i.e., perfect predictions. For each setting, we repeat the experiment 5 times and plot the mean performance ratio.

Besides the learning-augmented algorithms introduced above, we also include key algorithms in the clairvoyant setting in the evaluation. Clairvoyant algorithms can be considered as learning-augmented algorithms that blindly take predictions as truths. That is, they operate based on predicted

³We have also experimented with other distributions of prediction errors and observed similar performance trends. The results using a uniform distribution of prediction errors are provided in [22].

item durations or departure times directly. Note that the Nearest Remaining Time, Greedy and Lifetime Alignment algorithms exploit the bin's indicated closing time (the latest departure time of all items currently placed in a bin). If all the active items in a bin have their departure times underestimated, it may happen that the bin's indicated closing time based on predictions is earlier than the current time. In this case, we adapt the algorithms to simply take the current time as the bin's indicated closing time. This essentially means that if the prediction underestimates the item duration, after the predicted duration expires, the item is predicted to depart immediately, so its predicted departure time is set to the current time.

Experimental Comparison: For clarity of illustration, we divide algorithms into several groups and plot their performance separately. Figures 9, 10, 11, and 12 compare the mean performance ratios of various algorithms as the standard deviation σ increases (prediction errors spread over a wider range and the maximum prediction error increases). We also include the First Fit algorithm (which is the best-performing non-clairvoyant algorithm) in these figures. Since First Fit does not use any predicted information, its performance is a constant and independent of the prediction error.

Classify-By-Departure-Time vs. Classify-By-Duration: Figure 9 shows that the mean performance ratio of Classify-By-Departure-Time increases much faster than Classify-By-Duration as prediction errors spread more widely. This is

because the number of item categories created by Classify-By-Departure-Time increases linearly with the maximum prediction error, whereas that created by Classify-By-Duration increases logarithmically. Hence, the former is less robust against prediction errors than the latter, despite that the former outperforms the latter in the clairvoyant setting (perfect predictions) as was shown in Figure 8.

RCP and PPE: Figure 10 shows that if predictions are perfect, RCP and PPE behave exactly the same, and if prediction errors are small, they also perform similarly. When prediction errors further increase, PPE starts to outperform RCP, which demonstrates the benefit of adjusting the threshold value according to the prediction error. If prediction errors are very large, the threshold value α/\sqrt{x} set by PPE becomes very high (because α increases linearly with the maximum prediction error whereas x increases logarithmically). As a result, most items go to the general bins, reducing PPE towards First Fit. Thus, PPE’s performance ratio even decreases with increasing prediction errors. In contrast, the threshold value $1/\sqrt{x}$ set by RCP drops with increasing prediction errors. Hence, more items go to the category bins. Since there are also more item categories as prediction errors increase, the packing efficiency deteriorates. Thus, RCP’s performance ratio continues to increase with prediction errors. Comparing RCP/PPE with their modified versions, it can be seen that the modified versions perform better than the original versions. This is because eliminating large bins improves the bin utilization, thereby reducing the number of bins that need to be opened. Therefore, we shall choose the modified PPE algorithm without large bins for further comparison below.

Lifetime Alignment vs. Classify-By-Duration and Hybrid: Figure 11 compares the two variants of Lifetime Alignment (LA) as well as Classify-By-Duration and Reduced Hybrid. The performance ratio of LA(Binary Classification) increases slower than the other three algorithms as prediction errors increase. This is because LA(Binary Classification) uses the item duration information gently for categorization and hence is less vulnerable to prediction errors. In contrast, the other three algorithms use the item duration information more heavily for categorization and thus are more likely to suffer from prediction errors. Comparing these three algorithms, we can see that LA(Geometric Ranges) consistently outperforms the other two across different prediction errors. This demonstrates that dynamically categorizing the bins can achieve higher packing efficiency than fixing the categories of bins for item placement.

Prioritized NRT and Greedy: Figure 12 shows that Prioritized NRT is not resilient to prediction errors, despite that it has the best performance when predictions are perfect. In contrast, Greedy is more robust than Prioritized NRT against prediction errors. Interestingly, we find a similarity between Greedy and Prioritized NRT – both algorithms consider bins with indicated closing times after the item’s departure time for item placement before other bins, aiming to place the item in a bin without increasing its indicated closing time. The difference is that Greedy lists these bins in decreasing

order of indicated closing time, whereas Prioritized NRT lists these bins in ascending order of indicated closing time. In a sense, Greedy is *conservative* in that it prefers bins with indicated closing times far beyond the item’s departure time, while Prioritized NRT is *aggressive* since it prefers bins with indicated closing times near to the item’s departure time. In the presence of prediction errors, the bins preferred by Greedy are more likely to have real closing times after the item’s departure time than the bins preferred by Prioritized NRT. That is why Greedy is more robust against prediction errors than Prioritized NRT.

Overall Comparison: Figure 12 also includes the best-performing algorithms from previous comparisons. We also include First Fit – the best performer in the non-clairvoyant setting (it is a flat curve because First Fit is oblivious to prediction errors). As seen from Figure 12, at high prediction errors, the modified PPE and First Fit algorithms perform the best. At medium prediction errors, modified PPE is somewhat inferior to First Fit because modified PPE does not have the Any Fit feature. In this case, the Greedy and LA(Binary Classification) algorithms outperform modified PPE. When predictions are close to perfect, the Prioritized NRT algorithm performs the best. An interesting future direction is to develop an adaptive algorithm that switches between different algorithms based on the prediction errors observed so that it can combine their advantages to optimize performance over a wider range of prediction errors.

VII. CONCLUSION

We have conducted extensive experiments to evaluate DVBP algorithms for virtual machine placement. We have also proposed and analyzed new algorithms or enhancements including Round Robin Next Fit, Nearest Remaining Time, Reduced Hybrid, and modified RCP/PPE without large bins. Our evaluations show that algorithms conducting fixed classification of bins based on item departure times or durations, while necessary to break the lower bound of Any Fit algorithms in competitive ratios, do not perform as well as Any Fit algorithms empirically. Our proposed Prioritized NRT algorithm not only achieves a near-optimal competitive ratio among Any Fit algorithms, but also demonstrates the best empirical performance in the clairvoyant setting. The modified PPE, Greedy and LA(Binary Classification) algorithms show stronger resilience to prediction errors than other algorithms in the learning-augmented setting. Interesting future directions include (1) constructing adaptive algorithms to take advantage of the strengths while avoiding the drawbacks of different approaches empirically; (2) closing the gap between the lower bound on the competitiveness of DVBP and the competitive ratios achieved by DVBP algorithms in the multi-dimensional case theoretically; (3) analyzing the average-case performance of DVBP algorithms under practical input distributions.

ACKNOWLEDGMENT

This research is supported by the Ministry of Education, Singapore, under its AcRF Tier 1 (Award RG15/25).

REFERENCES

- [1] AWS, “Overview of Amazon Web Services - Overview of Amazon Web Services — docs.aws.amazon.com,” <https://docs.aws.amazon.com/whitepapers/latest/aws-overview/introduction.html>, [Accessed 31-12-2023].
- [2] M. Masdari, S. S. Nabavi, and V. Ahmadi, “An overview of virtual machine placement schemes in cloud computing,” *Journal of Network and Computer Applications*, vol. 66, pp. 106–127, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1084804516000291>
- [3] O. Hadary, L. Marshall, I. Menache, A. Pan, E. E. Greeff, D. Dion, S. Dorminey, S. Joshi, Y. Chen, M. Russinovich, and T. Moscibroda, “Protean: VM allocation service at scale,” in *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. USENIX Association, Nov. 2020, pp. 845–861. [Online]. Available: <https://www.usenix.org/conference/osdi20/presentation/hadary>
- [4] Y. Azar, I. R. Cohen, S. Kamara, and B. Shepherd, “Tight bounds for online vector bin packing,” in *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, ser. STOC ’13. New York, NY, USA: Association for Computing Machinery, 2013, p. 961–970. [Online]. Available: <https://doi.org/10.1145/2488608.2488730>
- [5] E. G. C. Jr., M. R. Garey, and D. S. Johnson, “Dynamic bin packing,” *SIAM J. Comput.*, vol. 12, no. 2, pp. 227–258, 1983. [Online]. Available: <https://doi.org/10.1137/0212014>
- [6] R. Ren and X. Tang, “Clairvoyant dynamic bin packing for job scheduling with minimum server usage time,” in *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures*, ser. SPAA ’16. New York, NY, USA: Association for Computing Machinery, 2016, p. 227–237. [Online]. Available: <https://doi.org/10.1145/2935764.2935775>
- [7] S. Kamali and A. López-Ortiz, “Efficient online strategies for renting servers in the cloud,” in *SOFSEM 2015: Theory and Practice of Computer Science*, G. F. Italiano, T. Margaria-Steffen, J. Pokorný, J.-J. Quisquater, and R. Wattenhofer, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, pp. 277–288.
- [8] Y. Azar and D. Vainstein, “Tight bounds for clairvoyant dynamic bin packing,” *ACM Trans. Parallel Comput.*, vol. 6, no. 3, Oct. 2019. [Online]. Available: <https://doi.org/10.1145/3364214>
- [9] X. Tang, Y. Li, R. Ren, and W. Cai, “On first fit bin packing for online cloud server allocation,” in *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2016, pp. 323–332.
- [10] Y. Li, X. Tang, and W. Cai, “Dynamic bin packing for on-demand cloud resource allocation,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 27, no. 1, pp. 157–170, 2016.
- [11] —, “On dynamic bin packing for resource allocation in the cloud,” in *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures*, ser. SPAA ’14. New York, NY, USA: Association for Computing Machinery, 2014, p. 2–11. [Online]. Available: <https://doi.org/10.1145/2612669.2612675>
- [12] R. Ren, X. Tang, Y. Li, and W. Cai, “Competitiveness of dynamic bin packing for online cloud server allocation,” *IEEE/ACM Trans. Netw.*, vol. 25, no. 3, p. 1324–1331, Jun. 2017.
- [13] M. Liu and X. Tang, “Dynamic bin packing with predictions,” *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 6, no. 3, Dec. 2022. [Online]. Available: <https://doi.org/10.1145/3570605>
- [14] —, “Tight bounds for dynamic bin packing with predictions,” *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 8, no. 3, Dec. 2024. [Online]. Available: <https://doi.org/10.1145/3700437>
- [15] N. Buchbinder, Y. Fairstein, K. Mellou, I. Menache, and J. S. Naor, “Online virtual machine allocation with lifetime and load predictions,” in *Abstract Proceedings of the 2021 ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS ’21. New York, NY, USA: Association for Computing Machinery, 2021, p. 9–10. [Online]. Available: <https://doi.org/10.1145/3410220.3456278>, <https://www.microsoft.com/en-us/research/publication/online-virtual-machine-allocation-with-lifetime-and-load-predictions/>
- [16] A. Murhekar, D. Arbour, T. Mai, and A. B. Rao, “Brief announcement: Dynamic vector bin packing for online resource allocation in the cloud,” in *Proceedings of the 35th ACM Symposium on Parallelism in Algorithms and Architectures*, ser. SPAA ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 307–310. [Online]. Available: <https://doi.org/10.1145/3558481.3591314>
- [17] Y. Li, M. Masoori, L. Narayanan, and D. Pankratov, “Renting servers for multi-parameter jobs in the cloud,” in *Proceedings of the 26th International Conference on Distributed Computing and Networking*, ser. ICDCN ’25. New York, NY, USA: Association for Computing Machinery, 2025, pp. 36–45. [Online]. Available: <https://doi.org/10.1145/3700838.3700850>
- [18] R. Kumar, M. Purohit, and Z. Svitkina, “Improving online algorithms via ML predictions,” in *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS)*, 2018, pp. 9661–9670.
- [19] T. Lykouris and S. Vassilvitskii, “Competitive caching with machine learned advice,” in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018, pp. 3302–3311.
- [20] A. Borodin and R. El-Yaniv, “Online computation and competitive analysis,” in *Online computation and competitive analysis*, 1998. [Online]. Available: <https://api.semanticscholar.org/CorpusID:5431684>
- [21] J. Sheng, S. Cai, H. Cui, W. Li, Y. Hua, B. Jin, W. Zhou, Y. Hu, L. Zhu, Q. Peng *et al.*, “Vmagent: Scheduling simulator for reinforcement learning,” *arXiv preprint arXiv:2112.04785*, 2021.
- [22] Z. Y. Lee and X. Tang, “Evaluation of dynamic vector bin packing for virtual machine placement (extended version),” *arXiv preprint arXiv:2602.14704*, 2026.
- [23] M. Mitzenmacher and S. Vassilvitskii, “Algorithms with predictions,” *Commun. ACM*, vol. 65, no. 7, p. 33–35, Jun. 2022. [Online]. Available: <https://doi.org/10.1145/3528087>
- [24] H. Barbalho, P. Kovaleski, B. Li, L. Marshall, M. Molinaro, A. Pan, E. Cortez, M. Leao, H. Patwari, Z. Tang, T. Santos, L. R. Gonçalves, D. Dion, T. Moscibroda, and I. Menache, “Virtual machine allocation with lifetime predictions,” in *MLSys*, 6 2023. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/virtual-machine-allocation-with-lifetime-predictions/>