

Robustly Assigning Unstable Items

Ananya Christman¹, Christine Chung², Nicholas Jaczko¹, Scott Westvold¹, and David S. Yuen³

¹ Department of Computer Science, Middlebury College, Middlebury VT
{achristman, njaczko, swestvold}@middlebury.edu

² Department of Computer Science, Connecticut College, New London, CT
cchung@conncoll.edu

³ Department of Mathematics, University of Hawaii, Honolulu, HI
yuen@math.hawaii.edu

Abstract. We study the Robust Assignment Problem where the goal is to assign items of various types to containers without exceeding container capacity. We seek an assignment that uses the fewest number of containers and is robust, that is, if any item of type t_i becomes corrupt causing the containers with type t_i to become unstable, every other item type $t_j \neq t_i$ is still assigned to a stable container. We begin by presenting an optimal polynomial-time algorithm that finds a robust assignment using the minimum number of containers for the case when the containers have infinite capacity. Then we consider the case where all containers have some fixed capacity and give an optimal polynomial-time algorithm for the special case where each type of item has the same size. When the sizes of the item types are nonuniform, we provide a polynomial-time 2-approximation for the problem. We also prove that the approximation ratio of our algorithm is no lower than 1.813. We conclude with an experimental evaluation of our algorithm.

1 Introduction

We study the Robust Assignment Problem (RAP) where we are given various types of items, each with a weight, where items of the same type have the same weight. We must assign the items to a set of containers with the constraint that if an item is found to be corrupt (we assume that there may be at most one such item), then every container containing an item of that type becomes unstable. Therefore, we would like at least one item of every other type to remain in at least one stable container. Such an assignment is considered *robust* and we would like a robust assignment that uses the fewest number of containers while satisfying their weight limit.

More formally, the input is n item types t_1, \dots, t_n with sizes (or weights) w_1, \dots, w_n , respectively, and container capacity C . The output is an assignment of types to subsets of containers, which uses the lowest number of containers and satisfies the following constraints: (1) Each type is assigned to k containers, for some $k \geq 1$ (2) Each container is assigned at most C total weight (3) The assignment is *robust*, that is, for any type t_i , if all containers having an item of type t_i become unstable, for all other types $t_j \neq t_i$, there is a stable container that contains t_j . Formally, let $S_i = \{s_{i,1}, s_{i,2}, \dots, s_{i,k}\}$ for $1 \leq i \leq n$ denote the set of k containers to which an item of type t_i was assigned. Then for every type $t_j \neq t_i$ such that an item of type t_j is also assigned to any container of S_i , an item of type t_j will exist on some container that is not in S_i . The goal is to find a robust assignment that uses the fewest containers.

RAP has many practical applications. For example, in distributed systems, multiple applications, including instances of the same app, are hosted on a cluster of servers. If a failure occurs in an app (and may therefore possibly occur in the other instances of the faulty app), then the app, all of its hosting servers, and hence all other app instances on those servers, are temporarily suspended. Therefore, the system would like an assignment of app instances to the minimal number of servers

such that if a failure occurs in an app and therefore all its hosting servers are temporarily suspended, there is still a running instance of every other app hosted on some unaffected server in the system. RAP can be used to find such an assignment - the apps correspond to the items and the servers correspond to the containers. The goals of our work were in fact motivated by a conversation with industry colleagues who encountered this problem in their company’s hosting platforms.

Ad placement on webpages is another application of the Robust Assignment Problem. Ad companies often have ads from multiple clients that must be displayed throughout various webpages of a website. If an ad crashes or slows down, it may affect the entire webpage and hence, the other ads displayed on that webpage as well. Other webpages displaying the faulty ad may need to be temporarily suspended to repair or check the faulty ad. Therefore ad companies would like an assignment of ads to webpages such that if a faulty ad temporarily suspends all of the webpages it is displayed on, there is still a running instance of every other ad on some webpage on the website. Here, the ads and webpages correspond to the items and containers, respectively.

RAP can also be presented as an application to gardening/agriculture. Avid gardeners often grow multiple plants of different varieties in several garden beds. Suppose that during the growing season, it becomes known that a particular plant variety has become disease prone. Therefore, all plants that are planted in the same bed as a disease-prone plant may become infected with the disease. Therefore, gardeners would like to find a way to plan their garden such that if a plant variety becomes prone to disease, then at least one plant of each variety still grows.

Our Results. For RAP we first give an optimal polynomial-time algorithm for finding the minimum number of containers needed to robustly assign the given set of item types, ignoring capacity constraints on the containers (Section 3.1). We then introduce the constraint of capacitated containers and give an optimal polynomial-time algorithm for the special case where each type of item has the same size (Section 4). For the general case of nonuniform sizes, we provide a polynomial-time 2-approximation for the problem (Section 4.2). I.e., our algorithm uses no more than twice the number of containers of the optimal robust assignment. We also prove that the approximation ratio of our algorithm is at least 1.813. We conclude with an experimental evaluation of our algorithm. (Section 5).

2 Related Work

To the best of our knowledge, our specific model for a robust assignment has not been previously studied. However, our solution ideas draw on those used for the bin-packing problem and some assignment problems, so we first discuss literature related to both problems. As mentioned above, in the context of distributed computing, our work applies to the problem of assigning replicas of applications to servers on a hosting platform, so we also discuss some literature on variations of this problem.

Our problem model has similarities to the problem of bin-packing with conflicts (or constraints) [2, 5, 6]. In the most general form of this problem, there are conflicts among the items to be packed and these conflicts are captured by a *conflict graph*, where the nodes represent the items and an edge exists between two items that are in a conflict [5]. The goal is to pack the items in the fewest number of bins while satisfying the capacity constraints on the bins and ensuring that no two items in a conflict are packed in the same bin. Jansen proposed an asymptotic approximation scheme for this problem for d -inductive graphs (i.e. where the vertices can be assigned distinct numbers $1 \dots n$ in such a way that each vertex is adjacent to at most d lower numbered vertices) including trees, grid graphs, planar graphs and graphs with constant treewidth [5]. For all $\epsilon > 0$, Jansen and Öhring [6] presented a $(2+\epsilon)$ -approximation algorithm for the problem on cographs and partial K -trees, and a 2-approximation algorithm for bipartite graphs. Epstein and Levin [2] improved on the 2.7-approximation of [6] on perfect graphs by presenting a 2.5-approximation. They

also presented a $7/3$ -approximation for a sub-class of perfect graphs and a 1.75-approximation for bipartite graphs.

Our problem differs from these previous problems in at least two important ways. First, the conflicts among our items cannot be easily captured by a conflict graph as they do not pertain to specific pairs of items, but rather to *all* pairs of items. Second, for our problem, the total number of items that are packed into bins is not predefined, so an algorithm may create more or less if doing so yields fewer bins.

The wide variety of problems that address the task of assigning items to containers while satisfying constraints and minimizing or maximizing some optimization objective are typically classified as Generalized Assignment Problems [1, 11]. While (to our knowledge) no previous works have considered the requirement of a robust assignment as in our model, a few works have had some similarities to ours. Fleischer et al. [3] studied a general class of maximizing assignment problems with packing constraints. In particular, they studied the Separable Assignment Problems (SAP), where the input is a set of n bins, a set of m items, values $f_{i,j}$ for assigning item j to bin i ; and a separate packing constraint for each bin – i.e. for bin i , a family of subsets of items that fit in bin i . The goal is to find an assignment of items to bins with the maximum aggregate value. For all examples of SAP that admit an approximation scheme for the single-bin problem, they present an LP-based algorithm with approximation ratio $(1 - \frac{1}{e} - \epsilon)$ and a local search algorithm with ratio $(\frac{1}{2} - \epsilon)$. Korupolu et al. [8] studied the Coupled Placement problem, in which jobs must be assigned to computation and storage nodes with capacity constraints. Each job may prefer some computation-storage node pairs more than others, and may also consume different resources at different nodes. The goal is to find an assignment of jobs to computation nodes and storage nodes that minimizes placement cost and incurs a minimum blowup in the capacity of the individual nodes. The authors present a 3-approximation algorithm for the problem.

One application of our work is the problem of assigning replicas of applications to servers on a hosting platform so that the system is fault-tolerant to a single application failure. There have been a wide variety of studies on related problems and here we discuss a few. Rahman et al. [10] considered the related Replica Placement Problem where copies of data are stored in different locations on the grid such that if one instance at one location becomes unavailable due to failure, the data can be quickly recovered. They present extensive experimental results for this problem. Mills et al. [9] also studied a variation of this problem in the setting where dependencies exist among the failures and the general goal is to find a placement of instances that does not induce a large number of failures. They give two exact algorithms for dependency models represented by trees. Uргаonkar et al. [15] also studied the problem of placing apps on servers, but do not consider fault tolerance and focus instead on satisfying each application’s resource requirement. The authors study the usefulness of traditional bin-packing heuristics such as First-Fit and present several approximation algorithms for variations of the problem.

More recently, Korupolu and Rajaraman [7] studied the problem of placing tasks of a parallel job on servers with the goal of increasing availability under two models of failures: adversarial and probabilistic. In the adversarial model, each server has a weight and the adversary can remove any subset of servers of total weight at most a given bound; the goal is to find a placement that incurs the least disruption against such an adversary. For this problem they present a PTAS. In the probabilistic model, each node has a probability of failure and the goal is to find a placement that maximizes the probability that at least a certain minimum number of tasks survive at any time. For the most basic version of the problem they study they give an algorithm that achieves an additive ϵ -approximation. Stein and Zhong [13] studied a related problem of processing jobs on machines to minimize makespan. The jobs must be grouped into sets before the number of machines is known and these sets must then be scheduled on machines without being separated. They present an algorithm that is guaranteed to return a schedule on any number of machines

that is within a factor of $(\frac{5}{3} + \epsilon)$ of the optimal schedule, where the optimum is not subject to the restriction that the sets cannot be separated.

3 Preliminaries

As a concrete example, Figures 1 and 2 show two assignments and the corresponding states of the containers for $n = 6$ item types. In Figure 1, the assignment is not robust – if type 2 is found to be corrupt, then no items of type 5 will exist in any other containers since all items of type 5 are in the same set of containers as items of type 2 (a similar problem occurs with types 3 and 4). Figure 2 depicts one robust assignment using the optimal number of containers: 4. Note that if an item of type 2 fails, then all item types contained in B and D exist in some other container. Further note that in this assignment if *any* of the item types are found to be corrupt, this robustness property holds.

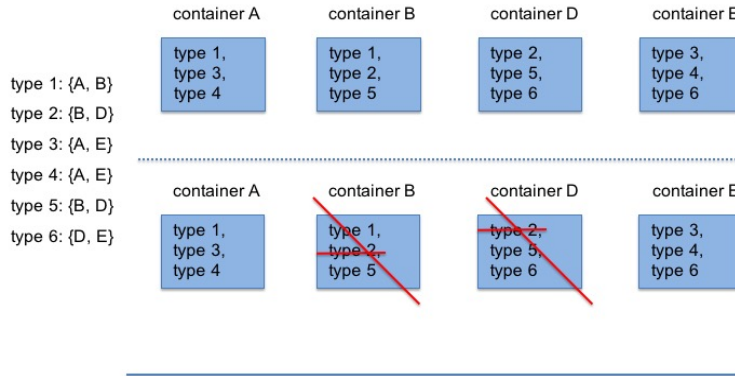


Fig. 1. Example of a non-robust assignment. If type 2 is corrupt, no items of type 5 exist.

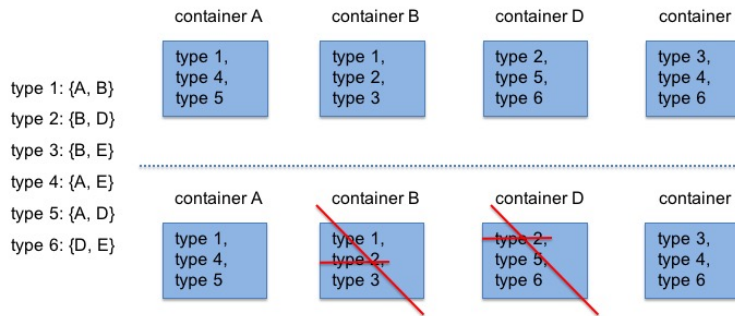


Fig. 2. An optimal robust assignment. If any type is corrupt, all other types still exist.

A *robust* assignment is characterized by whether each type is assigned to a set of containers that is not a subset of the set of containers assigned to any other type. We present this characterization formally as our first Lemma

Lemma 1. Let $S_i = \{s_{i,1}, s_{i,2}, \dots, s_{i,k}\}$ for $1 \leq i \leq n$ denote the set of k containers to which an item of type t_i was assigned. An assignment of item types to containers is robust if and only if there is no pair of item types t_i, t_j such that $S_i \subseteq S_j^\dagger$.

Proof. First we show that a robust assignment implies no pair of types t_i, t_j will be such that $S_i \subseteq S_j$. Suppose by way of contradiction that there is some pair of types t_i, t_j where $S_i \subseteq S_j$. This means if type t_j is found to be corrupt, and all the containers in S_j become unstable, then all the containers in S_i also become unstable. In this case there are no items of type t_i in stable containers so the assignment was not robust. We now prove the other direction of the lemma. Suppose for contradiction we have no pair of types t_i, t_j such that $S_i \subseteq S_j$, but the assignment is not robust. If it is not robust, there is some type k such that removing the containers in S_k will leave another type k' in no remaining containers. But for this to be true, it must be that $S_{k'} \subseteq S_k$ which is a contradiction.

3.1 Uncapacitated Robust Assignment Problem

In this section we begin by considering the special case of the RAP where the containers have infinite capacity. To tackle the Uncapacitated Robust Assignment Problem we first consider the inverse problem: given m containers, what is the maximum number of item types we can assign robustly? Due to Lemma 1 this problem can be modeled as the combinatorics problem of finding the maximum cardinality *antichain* of a set. Specifically, let P denote the set of subsets of m elements $\{1, 2, \dots, m\}$. An antichain of P is a set $\bar{P} = \{s_1, s_2, \dots, s_k\} \subseteq P$ such that for any pair of subsets s_i, s_j in \bar{P} , $s_i \not\subseteq s_j$. For a table of all antichains for $m = 1, 2$, and 3 , please refer to the full version of the paper.

Sperner's Theorem [12] states that the maximum cardinality of an antichain \bar{P} of an m -sized set is $\binom{m}{\lfloor m/2 \rfloor}$ and each subset of \bar{P} has size $m/2$. (If m is odd then there will be two maximum cardinality antichains whose subsets will have size $\lfloor m/2 \rfloor$ and $\lceil m/2 \rceil$, respectively.) Therefore, Sperner's Theorem yields the maximum number of item types that can be assigned to m containers as well as the number of containers to which each type is assigned. The values in Table 1 were derived from Sperner's Theorem.

We can thus use this theorem in conjunction with Lemma 1 to solve our original assignment problem – that is, given n types, find the minimum number of containers required to assign these types. Specifically, given n types, we would like to find the smallest m such that $\binom{m}{\lfloor m/2 \rfloor} \geq n$. See Algorithm 1 for further details.

m (# of containers)	maximum number of item types that can be robustly assigned
1	1
2	2
3	3
4	6
5	10
6	20
7	35
8	70
9	126
10	252

Table 1. The maximum number of types that can be robustly assigned to $1 \leq m \leq 10$ containers.

[†] Note that $S_i \subseteq S_j$ is the general condition for nonuniform k ; for uniform k the condition is $S_i = S_j$.

Algorithm 1: Input is n item types.

- 1: Use Sperner's Theorem to find the minimum integer m such that $\binom{m}{\lfloor m/2 \rfloor} \geq n$.
 - 2: Set up m empty containers.
 - 3: Generate all $\binom{m}{\lfloor m/2 \rfloor}$ of the $\lfloor m/2 \rfloor$ -combinations of the m containers.
 - 4: Assign each item type one of the $\lfloor m/2 \rfloor$ -combinations, i.e. for each type, assign an item of that type to each of the $\lfloor m/2 \rfloor$ containers in the $\lfloor m/2 \rfloor$ -combination that this type was assigned to.
-

Theorem 1. *The Uncapacitated Robust Assignment Problem is solvable in time polynomial in n , the number of item types.*

Proof. Due to Sperner's Theorem and Lemma 1, Algorithm 1 correctly returns the minimum number of containers required. Steps 1 and 2 take no more than time linear in the number of types as n serves as a trivial upperbound on the value of m that satisfies the Sperner's Theorem condition. (Furthermore, we can potentially find the solution more quickly by computing upper and lower bounds on m using Stirling's approximation which states that $\binom{m}{\lfloor m/2 \rfloor} \approx m + \frac{1}{2} - \frac{1}{2} \log_2(m\pi)$ [14].) Steps 3 and 4 of the algorithm require enumeration of the $\binom{m}{\lfloor m/2 \rfloor}$ combinations; the number of combinations is exponential in m , but since the chosen m will be $O(\log(n))$, the composite run time is still polynomial in n .

4 The Robust Assignment Problem with Capacity Constraints

In Section 3.1, we implicitly assumed that any number of items can be assigned to any one container. However, in practical settings, constraints such as storage space, memory, or other demands will impose limits on the number of items a container may hold. We therefore consider a model where there is one such constraint. We will use the example of a storage constraint for expository purposes.

The problem now becomes: given n types t_1, t_2, \dots, t_n with integer-valued *sizes*, w_1, w_2, \dots, w_n , respectively, where items of type t_i have size w_i ; and an integer-valued container capacity C , where $1 \leq w_i \leq C$, find an assignment of items to the minimal number of containers such that the assignment is both (1) robust and (2) satisfies the following *capacity constraint*: if A_j is the set of items assigned to container s_j , then for all containers $j = 1 \dots m$, where m is the number of containers used in the assignment, $\sum_{a \in A_j} w(a) \leq C$, where $w(a)$ is the size of item a . We refer to this variant as the *Capacitated Robust Assignment Problem*.

As a small example, suppose in Fig. 2, types 1, 2, \dots , 6 have sizes 1, 2, \dots , 6, respectively. Then if $C = 12$, the assignment shown in the figure would not satisfy the capacity constraint since both containers D and E currently use 13 units of size. Figure 3 shows an assignment that satisfies both the robustness and capacity constraints.




Fig. 3. An assignment that satisfies both the robustness and capacity constraints for capacity equal to 12.

4.1 Uniform Sizes

We first consider the special case where each type, and therefore each item, has the same size w . Given n such types and containers of capacity $C \geq w$, the problem is to find an assignment of types to the minimal number of containers such that the assignment is robust and also satisfies the capacity constraints. In this case the capacity constraint is that if $|A_j|$ denotes the number of items assigned to container s_j , and m is the number of containers in the assignment, then for all $j = 1 \dots m$, $|A_j|w \leq C$.

Theorem 2. *The Capacitated Robust Assignment Problem with uniform sizes is solvable in time polynomial in n , the number of item types.*

Proof. Algorithm 2 solves this problem optimally. Recall that k denotes the number of items of each type. The algorithm effectively performs an exhaustive search to find the minimum m over all possible k for which the robustness and capacity constraints are satisfied. Specifically, the algorithm starts with the lower bound for m (given by Sperner's Theorem) and searches every possible integral value of k given this m (i.e. starting from $k = \lfloor m/2 \rfloor$ down to $k = 1$) that will satisfy both the robustness and capacity constraints. Robustness is satisfied if $\binom{m}{k} \geq n$ and the capacity constraint is satisfied if $\lfloor \frac{C}{w} \rfloor \geq \frac{nk}{m}$. If no value for k for the given m satisfies both constraints, the algorithm increments m and repeats the search for k .

Note that the algorithm will eventually terminate: if eventually m is incremented to n and k is decremented to 1 both conditions of the while loop will be true. There will be $O(n^2)$ iterations of the while loop. Each iteration takes constant time so the runtime of the loop is $O(n^2)$. The polynomial run-time and correctness of step 11 is addressed in the full version of the paper. So the overall run time of Algorithm 2 is polynomial in n .

The procedure for assigning the n types robustly to the m containers computed by Algorithm 2 is rather technical, so we refer the reader to the full version of the paper for the details on this procedure. We also note that while there are ways to optimize the run-time of our algorithm, the exhaustive-search version we present here is for the sake of simplicity and clarity.

Algorithm 2: Input is the container capacity C , n item types, and item size $w \leq C$.

- 1: Use Sperner's Theorem to find minimum m such that $\binom{m}{\lfloor \frac{m}{2} \rfloor} \geq n$. Note that m is a lower bound on the number of containers required to assign the types.
 - 2: $k = \lfloor \frac{m}{2} \rfloor$
 - 3: **while** not $(\binom{m}{k} \geq n$ and $\lfloor \frac{C}{w} \rfloor \geq \frac{nk}{m})$ **do**
 - 4: **if** $k > 1$ **then**
 - 5: $k --$ //decrease the number of items per type
 - 6: **else**
 - 7: $m ++$ //add another container
 - 8: $k = \lfloor \frac{m}{2} \rfloor$ //re-initialize k for the new m
 - 9: **end if**
 - 10: **end while**
 - 11: For details on how to assign k items of each type to a distinct subset of the m containers, refer to the full version of the paper.
-

We note that if the problem is simply to find the minimum number of containers needed for the robust assignment, without also requiring the robust assignment itself, one can do so in $\text{polylog}(n)$ time by formulating the problem as a fixed-dimension integer program. Namely, given inputs (n, C)

where for simplicity we assume $w = 1$, then we want to solve the system $1 \leq k \leq m/2$, $kn \leq mC$, $\binom{m}{k} \geq n$ for k and m with m minimal. The key observation is that for any fixed m , the k satisfying the first two equations that yield the largest $\binom{m}{k}$ is $k = \lfloor \min(m/2, mC/n) \rfloor$. Thus whether or not an m has a corresponding k that satisfies the three equations is equivalent to whether or not $\binom{m}{\lfloor \min(m/2, mC/n) \rfloor} \geq n$. Because we can show $\binom{m}{\lfloor \min(m/2, mC/n) \rfloor}$ is an increasing function in m , then we can do a binary search for the minimal m that satisfies $\binom{m}{\lfloor \min(m/2, mC/n) \rfloor} \geq n$ in the interval $[1, n]$. This would be far more efficient than the brute force while-loop that we give for simplicity in Algorithm 2.

4.2 Nonuniform Sizes

In this section we consider the variant of the problem where there may be a different size $1 \leq w_i \leq C$ for each type t_i . In this case, if we ignore the robustness constraint, the problem would be NP-hard due to its equivalence to bin-packing [4]. We first present our algorithm, Robust First Fit (RFF), for this problem and prove that its approximation ratio is at most 2. We then show that the approximation ratio of RFF is not lower than 1.813.

4.2.1 The Robust First Fit Algorithm

RFF begins by sorting the types in descending order by size. When finding an assignment for type t_i the algorithm first finds the set S of all the containers that have enough empty space to fit an item of type t_i . It then assigns an item of type t_i to the (lexicographically) first container assignment that can be created from the containers in S that has not already been used by a previous type.

Algorithm 3: ROBUST FIRST FIT (RFF). Input is the container capacity C and a set T of n types where all items of type t_i have size $w_i \leq C$ for $1 \leq i \leq n$.

- 1: Sort the types in descending size order.
 - 2: Use Sperner's Theorem to find minimum m such that $\binom{m}{\lfloor \frac{m}{2} \rfloor} \geq n$. Note that m is a lower bound on the number of containers required to hold the items.
 - 3: **for** $k = \lfloor \frac{m}{2} \rfloor$ to 1 **do**
 - 4: Set up $m' = m$ empty containers.
 - 5: **for** each type t_i in T **do**
 - 6: Let S denote the subset of the m' containers that still have sufficient space to fit an item of type t_i
 - 7: Assign k items of type t_i to the lexicographically first k -combination of containers in S that is still available (i.e. no type has already been assigned to it)
 - 8: **if** t_i is still unassigned **then** {there were no available k -combinations in S }
 - 9: Add a new container to S , maintaining lexicographic order and increment m' . (Note that k does not change.)
 - 10: Go back to Step 7.
 - 11: **end if**
 - 12: **end for**
 - 13: Store m' along with the corresponding assignment.
 - 14: **end for**
 - 15: Return the assignment from the iteration of the outer-most **for**-loop (Step 3) that used the fewest containers.
-

Whenever no suitable assignment can be created for type t_i with the existing containers, a new empty container is created. An assignment with a storage constraint will never require fewer

containers than an assignment without a storage constraint, so the number of containers (m) and number of items of each type (k) are initialized to the values given by Sperner's Theorem. The **for**-loop in step 3 accounts for the fact that decreasing the number of items of each type might decrease the number of containers required. Therefore, we start with $k = \lfloor m/2 \rfloor$, as given by Sperner's Theorem, and try decreasing the number of items from there.

RFF runs in polynomial time. Since the initial value of m from step 2 can be at most n , the **for**-loop in Step 3 will run for at most $O(n)$ iterations. Each iteration of the loop finds an assignment for n types. To find an assignment for each type t_i , the algorithm searches for an available k -combination whose containers have sufficient space for an item of type t_i . This can be done in poly-time as there will never be more than $O(n)$ k -combinations to check before finding one that is available. Since the number of containers will be no more than $n \cdot \frac{n}{2} = O(n^2)$ (i.e. if one item of each type was assigned to its own dedicated container), the algorithm may reach Step 10 (which causes a new iteration from Step 7) $O(n^2)$ times. Hence the overall run-time of RFF is polynomial.

We note that there are clearly ways to optimize the run-time of our algorithm if one wishes to implement it on a real-world system (for example, using binary search instead of linear search). The version we present here is for the sake of simplicity and clarity.

4.2.2 Upper Bound

We now show that RFF has an approximation ratio of no worse than 2.

Theorem 3. *RFF is a 2-approximation for the Capacitated Robust Assignment Problem with nonuniform sizes. I.e., RFF will use at most $2m^*$ containers to robustly assign all n types, where m^* is the number of containers that an optimal solution uses.*

Proof. Consider any input instance. Let n denote the number of item types to be assigned, let OPT be an optimal robust assignment for them, and let k be the number of items assigned per type in OPT. It suffices to show that for the optimal k RFF uses at most $2m^*$ number of containers since RFF tries each potential value of k and chooses the value of k that minimizes the number of containers required. Hence, if RFF uses no more than $2m^*$ containers when it assigns k items per type, it must ultimately not use more than $2m^*$ containers. We also assume $k \geq 2$, since in the case of $k = 1$ RFF and OPT will both use a dedicated assignment (one item of each type per container) so RFF will return an optimal assignment, using $m = m^*$ containers.

Suppose for contradiction that OPT uses m^* containers while RFF uses strictly more than $2m^*$ containers when assigning the n types. Consider the moment during the execution of the RFF algorithm that container number $2m^* + 1$ was opened and added to S . Let t_i be the type that was being assigned when RFF opened this $(2m^* + 1)$ th container. Let w_i be the size of type t_i . Let S_{-i} be the set of $2m^*$ containers already in use by the algorithm when it tried to assign t_i , but before it added container number $2m^* + 1$. Note that RFF has sorted and re-indexed the types in descending size order.

Case 1: $2 \leq k \leq \lfloor m^*/2 \rfloor$, $w_i > C/3$. Let B denote the set of all types t_j for whom $w_j > C/3$. Note that type $t_i \in B$. Due to their size, no more than two items of types in B can fit on a single container, so there can be no more than $2m^*$ such items in total, i.e., $k|B| \leq 2m^*$. Note however, that RFF must be able to assign the $k|B| \leq 2m^*$ items to at most $2m^*$ containers because $2m^*$ containers would indeed be sufficient for even a dedicated assignment: one item of each type per container. This contradicts the assumption that type t_i required RFF to open a $(2m^* + 1)$ th container.

Case 2: $2 \leq k \leq \lfloor m^*/2 \rfloor$, $w_i \leq C/3$. In this case, we consider two sub-cases. Subcase 1: there are at least m^* containers in S_{-i} with available space at least w_i (i.e. enough space for an item of type t_i). In this case, we would then have a robust assignment from the set S_{-i} for t_i because

OPT needed only m^* containers total to assign all n types, so having m^* containers must provide enough k -combinations to have at least one left for t_i .

Subcase 2: there are fewer than m^* containers in S_{-i} with available space at least w_i . So, in this case there must be $m^* + x$ containers $S_f \subseteq S_{-i}$, where $x > 0$, that have less than w_i available space. We can say that each of these containers in S_f already has filled capacity $C_f > C - w_i$. So if $w(S_f)$ is the total size of all of the items in the containers in S_f , then $w(S_f) > (m^* + x)(C - w_i)$. Since OPT used m^* containers of capacity C to assign all k items of each of the n types robustly, we have $(m^* + x)(C - w_i) < m^*C$. Recalling that we are in the case where $w_i \leq C/3$, we then have

$$(m^* + x) \left(C - \frac{C}{3} \right) = (m^* + x) \left(\frac{2C}{3} \right) < m^*C.$$

This implies $2xC/3 < m^*C/3$, which implies

$$x < m^*/2. \quad (1)$$

Let $S_g = S_{-i} - S_f$ be the set of containers in S_{-i} that still have enough remaining capacity to store an item of type t_i . For type t_i to be unable to be assigned to these $|S_g| = |S_{-i}| - |S_f| = m^* - x$ containers, it must be due to robustness: they must have no remaining available unique combinations of containers. We will show however, that if this were true, it would also lead to a contradiction.

If there are no unique combinations of containers remaining in S_g to assign t_i to, there must be at least $\binom{m^* - x}{k}$ distinct types that are already assigned to those containers. In other words, if

$$T_g = \{t_j \in T : \text{an item of type } t_j \text{ is assigned to some container in } S_g\},$$

then $|T_g| \geq \binom{m^* - x}{k}$. This is true because if $|T_g| < \binom{m^* - x}{k}$ then there would be at least one remaining available k -combination of the containers in S_g on which to assign t_i .

RFF considers types in descending order by size so each item of the $|T_g| \geq \binom{m^* - x}{k}$ types must take up at least as much space as w_i . Thus, $w(S_g) \geq \binom{m^* - x}{k}kw_i$, where $w(S_g)$ is the total size of all the items on the $m^* - x$ containers of S_g .

The size of all the items which are assigned to the $2m^*$ containers of S_{-i} is $w(S_{-i}) = w(S_f) + w(S_g) \geq (m^* + x)(C - w_i) + \binom{m^* - x}{k}kw_i$. Again, OPT used m^* containers of capacity C so we know the total size of all the items cannot be more than m^*C . Thus,

$$(m^* + x)(C - w_i) + \binom{m^* - x}{k}kw_i \leq m^*C \quad (2)$$

By expanding the left hand side of (2) we get

$$m^*C - m^*w_i + xC - xw_i + \binom{m^* - x}{k}kw_i \leq m^*C$$

and rearranging terms gives us:

$$xC + \binom{m^* - x}{k}kw_i \leq m^*w_i + xw_i \quad (3)$$

By combining Equations 3 and 1 with $w_i \leq C/3$ we get

$$3xw_i + \binom{m^* - x}{k}kw_i \leq m^*w_i + \frac{m^*}{2}w_i$$

which implies $\binom{m^* - x}{k} 2k + 6x \leq 3m^*$. Using the fact that $2 \leq k \leq \lfloor \frac{m^*}{2} \rfloor$ yields

$$\binom{m^* - x}{k} 4 + 6x \leq 3m^*. \quad (4)$$

It is a fact for any integers $a, b > 0$, where $b < a$, that $\binom{a}{b} \geq a$; and we know $m^* - x \geq k$ (since by Equation (1) we know $x < m^*/2$ and we are currently in the case where $k \leq m^*/2$). Hence we can say from Equation 4 that $4(m^* - x) + 6x \leq 3m^*$, which is a contradiction.

Both cases resulted in contradiction. So, RFF will never use more than $2m^*$ containers.

4.2.3 Lower Bound

We now provide a family of examples that give a lower bound on the approximation ratio of RFF. The family of examples is parameterized by a positive integer $d \geq 3$. We refer to the following instance as $I(d)$. There are $n = \binom{2d+3}{d}$ types, of which $\ell = 2d - 1$ are “large” types and $s = n - \ell$ are “small” types. The small types have size 1, while the large types have size $L = s$. Suppose the containers each have capacity $C = dL$. We first give an optimal assignment for this family.

Proposition 1. *For instance $I(d)$, an optimal assignment uses $m^* = 2d + 3$ containers.*

Proof. First we note that since $d \geq 2$, we have $\frac{1}{d+1} < \frac{2d+3}{(d+3)(d+2)}$. Then

$$\frac{(2d+2)!}{(d+1)d!(d+1)!} < \frac{(2d+3)(2d+2)!}{d!(d+1)!(d+3)(d+2)},$$

from which we get

$$\binom{2d+2}{d+1} < \binom{2d+3}{d} = n.$$

By Sperner’s Theorem, this says that instance $I(d)$ requires at least $m = 2d + 3$ containers and this number of containers is possible when $k = d$. Now, letting $k = d$, since

$$k\ell = d(2d - 1) = 2d^2 - d \leq 2d^2 + d - 3 = (d - 1)(2d + 3) = (d - 1)m,$$

we can store k items of each of the ℓ large types on the m containers with at most $d - 1$ items on each container (the full version of the paper describes how). Since the capacity of each container is $C = dL$, each container will have at least capacity L remaining. We will then use the remaining $\binom{2d+3}{d} - \ell$ combinations, which is exactly s , the number of small types, to assign the small items. By design, $Lm \geq ds$ and so there is enough remaining capacity to do this. Therefore $2d + 3$ is the optimal number of containers for the instance $I(d)$.

Given an instance $I(d)$, we now establish the number of containers returned by RFF.

Proposition 2. *For an instance $I(d)$, for each integer $1 \leq k \leq d + 1$, we define*

$$J(k) = \min\left\{j : \binom{j + 2d - 2k + 4}{j} \geq d - 1\right\}$$

$$z(k) = \min\left\{j : \binom{j}{k} \geq s\right\}$$

While using k items of each type, RFF will return the number of containers equal to:

$$m(k) = 2k - 1 - J(k) + z(k).$$

Then RFF will return the number of containers such that $m(k)$ is minimal over $1 \leq k \leq d + 1$.

Proof. (Please refer to Table 2 for example values of $J(k)$ and $z(k)$). RFF begins by calculating that at least $2d + 3$ containers are needed, and so RFF will loop from $k = d + 1$ down to $k = 1$ in search of the minimum number of containers needed. In what follows, we index both the containers and item types starting from 0. Consider a fixed k for $1 \leq k \leq d + 1$. RFF will assign each large item type t_j , for each $j = 0, \dots, d - 1$, to containers $\{0, \dots, k - 2, k - 1 + j\}$. Then the other remaining $d - 1$ large types are assigned to containers $k - 1, \dots, 2k - 2 - J$ and some order J subset of $\{2k - 1 - J, \dots, 2d + 2\}$, which has cardinality $J + 2d - 2k + 4$. We need $\binom{J + 2d - 2k + 4}{J} \geq d - 1$. For any such J , the containers numbered 0 through $2k - 2 - J$ would be filled to capacity with large types. Thus taking the minimum such J , calling it $J(k)$, exactly the first $2k - 1 - J(k)$ containers are filled; the other containers have at least capacity L remaining.

Let $z(k)$ be the smallest positive integer such that $\binom{z(k)}{k} \geq s$. To assign the s small types, it is clear we need at least $z(k)$ containers beyond the $2k - 1 - J(k)$. For $d \geq 3$, we can prove by induction that

$$s = \binom{2d + 3}{d} - (2d - 1) > \binom{2d + 2}{d + 1}. \quad (5)$$

This is true for $d = 3$, and assuming it is true for a particular d , then we multiply the lefthand side by $\frac{(2d + 4)(2d + 5)}{(d + 1)(d + 4)}$ and the righthand side by $\frac{(2d + 3)(2d + 4)}{(d + 2)^2}$, the latter of which we can prove is smaller by cross-multiplying. We then get

$$\binom{2d + 5}{d + 1} - (2d - 1) \frac{(2d + 4)(2d + 5)}{(d + 1)(d + 4)} > \binom{2d + 4}{d + 2}.$$

Now we can check by cross-multiplication that

$$(2d - 1) \frac{(2d + 4)(2d + 5)}{(d + 1)(d + 4)} > (2d + 1).$$

Then $\binom{2d + 5}{d + 1} - (2d + 1) > \binom{2d + 4}{d + 2}$, which is Equation 5 with d replaced by $d + 1$, completing the induction. Finally,

$$\binom{z(k)}{k} \geq \binom{2d + 3}{d} - (2d - 1) > \binom{2d + 2}{d + 1}$$

implies $z(k) \geq 2d + 3$.

Now note that by definition of $z(k)$ that

$$\binom{z(k) - 1}{k} < s.$$

Because $z(k) \geq 2d + 3$ and $k \leq d + 1 < z(k)/2$, then

$$\binom{z(k) - 1}{k - 1} < \binom{z(k) - 1}{k} < s = L.$$

Thus

$$L z(k) \geq k \binom{z(k)}{k}.$$

We can robustly assign each of $\binom{z(k)}{k}$ small types to k out of $z(k)$ containers each with capacity at least L (the full version describes how) and in particular RFF would naturally do this because every combination of k out of $z(k)$ containers is used. Since $s \leq \binom{z(k)}{k}$, then RFF would successfully use $z(k)$ containers to robustly assign the s small types. Thus we have shown that RFF with $k \leq d + 1$ items of each type uses $2k - 1 - J(k) + z(k)$ containers. Thus RFF uses the number of containers equal to the minimum of $2k - 1 - J(k) + z(k)$ for $1 \leq k \leq d + 1$.

d	k	$J(k)$	$z(k)$	RFF output	OPT output	RFF/OPT
5	5	1	13	21	13	1.615
8	8	2	19	32	19	1.684
9	8	2	22	35	21	1.666
15000	10611	2	33185	54404	30003	1.81328
25000	17663	2	55348	90671	50003	1.81331
35000	24710	2	77521	126938	70003	1.81332

Table 2. The number of containers output by RFF and OPT for different values of d . RFF outputs the minimal $m(k)$ over $1 \leq k \leq d+1$ while OPT outputs $m^* = 2d+3$

Theorem 4. *The approximation ratio of RFF is no better (lower) than 1.813.*

Proof. Let $d = 15000$, and consider the instance $I(d)$ as defined above. By Proposition 2, RFF ends up using $k = 10611$ and $J = 2$, $z = 33185$ and $m = 54404$ for this instance, while (by Proposition 1) an optimal assignment requires only $m^* = 2d+3 = 30003$. (Please see Table 2.)

5 Experimental Results

As described in Section 1, RAP can be applied to assigning app instances to the minimal number of servers on a hosting platform while ensuring that if a failure occurs in an app and therefore all its hosting servers are temporarily suspended, there is still a running instance of every other app hosted on some unaffected server. Formally, we are given n apps where app i has size d_i and server capacity C . We would like to find an assignment of app instances to the minimal number of servers m , such that the assignment is robust and satisfies the capacity constraint.

To evaluate the performance of the RFF algorithm, we simulated a hosting platform and measured the number of servers used by the algorithm. Specifically, we tested four values for server capacity C (64GB, 128GB, 256GB, and 512GB), varied the number of apps from $n = 25$ to $n = 250$ apps (at increments of 25) and set app sizes d_i to be normally distributed between 4GB and 16GB. We compared RFF to a dedicated system (i.e. where the number of servers is simply the number of apps) and an “ideal” assignment, which does not correspond to any feasible robust assignment, but serves as a lower bound on the minimally required number of servers. (Recall that it even without the robustness constraint, it is NP-hard to compute OPT so we did not compute it for the experiments.) We computed the “ideal” assignment by determining the minimum number of servers needed to satisfy robustness alone and the minimum number of servers to satisfy the storage constraints alone and taking the maximum of these two values. I.e., the “ideal” number of servers is defined as: $\min_k \max\{m_r, m_c\}$ where $m_c = \min\{m : mC \geq k \sum_{i=1}^n d_i\}$ and $m_r = \min\{m : \binom{m}{k} \geq n\}$.

We tested each setting for 10 iterations and took the average of the results. The graphs in Table 5 show the results. The graphs show that for all settings, RFF performs significantly better than the dedicated system and almost as well as the ideal assignment. Specifically, the worst (minimum) ratio (over all values of n) of servers used by the dedicated system and RFF is 2.40, 3.25, 3.57, and 3.57 for 64GB, 128GB, 256GB, and 512GB, so RFF always assigned apps more than twice as efficiently as a dedicated system. Note that as the server capacity increases, these ratios either increase or stay the same. The average ratio of servers used by the dedicated system and RFF always increase: 2.64, 4.72, 7.34, and 9.89 for 64GB, 128GB, 256GB, and 512GB, respectively, so RFF on average performed as much as 9 times as efficiently as a dedicated hosting.

Comparing RFF with the lower bound on optimal, we find that the worst (maximum) ratio (over all values of n) of servers used by RFF and the ideal assignment is 1.17, 1.21, 1.13, and 1.20 for 64GB, 128GB, 256GB, and 512GB, respectively. So RFF always used close to the same number of servers as an optimal solution. (The average ratios are similar to these values.) The results indicate

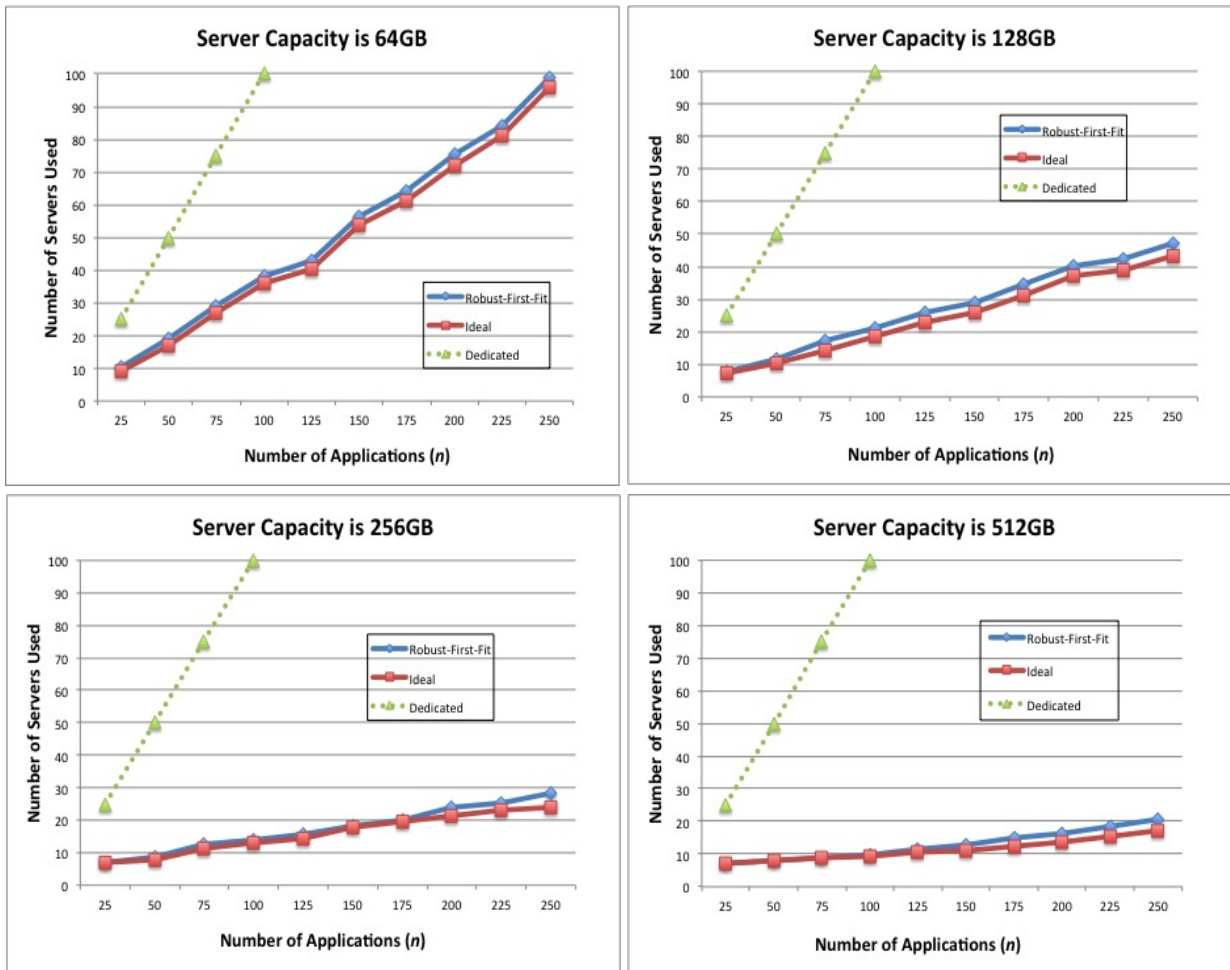


Table 3. Number of servers used when server capacity is 64GB, 128GB, 256GB, and 512 GB.

that when apps sizes are more realistic than those described in Theorem 4 of Section 4.2.3, RFF performs close to optimally.

6 Discussion and Conclusions

We proposed a new model for assigning items of various types to containers such that the system is *robust*. We presented an optimal poly-time algorithm in the setting without capacity constraints on the containers. We also presented an optimal poly-time algorithm when item sizes are uniform. Our main algorithm RFF is a poly-time 2-approximation algorithm for the setting where item sizes are nonuniform. Our experimental results suggest that when run on a simulated hosting platform, RFF performs well not only in the worst-case, but even more so on average.

In the lower bound instance, as d increases, it is not clear whether the corresponding ratio is converging (very slowly) to 2 or to a number less than 2, or whether the ratio converges at all; if the ratio does not converge, one can still ask for the limit supremum of the sequence of ratios. If the limit supremum is 2, then the upper bound of 2 is tight.

One direction for future work is to determine whether there is an algorithm with an approximation ratio better than 2. Also, our problem model assumes that the number of items is uniform over all types. A natural extension of this work would be to consider the case where this number is not required to be uniform.

References

1. C. Chekuri and S. Khanna. “A PTAS for the multiple knapsack problem”, in *Symposium on Discrete Algorithms (SODA)*, 2000.
2. L. Epstein, A. Levin. “On Bin Packing with conflicts,” in *Proceedings of the Workshop on Approximation and Online Algorithms (WAOA)*, 2006.
3. L. Fleischer, M. X. Goemans, V. S. Mirrokni, and M. Sviridenko. “Tight approximation algorithms for maximum general assignment problems”, in *Proceedings of the Symposium on Discrete Algorithms*, 2006.
4. M.R Garey and D.S Johnson. “Computers and Intractability: A Guide to the Theory of NP-Completeness Freeman,” 1979.
5. K. Jansen. “An approximation scheme for bin packing with conflicts,” in *Journal of Combinatorial Optimization*, vol. 3, issue 4, 1999.
6. K. Jansen and S. Öhring. “Approximation algorithms for time constrained scheduling,” in *Information and Computation*, vol. 132, issue2, 1997.
7. M. Korupolu and R. Rajaraman. “Robust and probabilistic failure-aware placement,” in *Proceedings of the Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pp. 213-224, 2016.
8. M. Korupolu A. Meyerson, R. Rajaraman, and B. Tagiku. “Robust and probabilistic failure-aware placement,” in *Mathematical Programming*, vol. 154, issue 1-2, pp. 493-514, 2015.
9. K. Mills, R. Chandrasekaran, and N. Mittal, “Algorithms for optimal replica placement under correlated failure in hierarchical failure domains,” in *Theoretical Computer Science* (pre-print), 2017.
10. R. Rahman, K. Barker, and R. Alhajj, “Replica placement strategies in data grid,” in *Journal of Grid Computing*, vol. 6, issue 1, pp. 103-123, 2008.
11. D. Shmoys and E. Tardos. “An approximation algorithm for the generalized assignment problem”, in *Mathematical Programming*, vol. 62, issue (3), pp. 461-474, 1993.
12. E. Sperner, “Ein Satz ber Untermengen einer endlichen Menge,” in *Mathematische Zeitschrift*, vol. 27, issue 1, pp. 544 - 548, 1928.
13. C. Stein and M. Zhong. “Scheduling When You Don’t Know the Number of Machines”, in *Proceedings of the Symposium on Discrete Algorithms (SODA)*, 2018.
14. J. Stirling, “Methodus differentialis, sive tractatus de summation et interpolation serierum infinitarum” *London*, 1730.
15. B. Urgaonkar, A. Rosenberg, and P. Shenoy, “Application placement on a cluster of servers,” in *International Journal of Foundations of Computer Science*, vol. 18, issue 5, pp. 1023-1041, 2007.