

BIN-PACKING WITH FRAGILE OBJECTS

Nikhil Bansal

*Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA
nikhil@cs.cmu.edu*

Zhen Liu

*IBM T. J. Watson Research Center
PO Box 704
Yorktown Heights, NY 10598, USA
zhenl@us.ibm.com*

Arvind Sankar

*Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139-4307, USA
arvinds@mit.edu*

Abstract We consider an extension of the classical bin packing problem, motivated by a frequency allocation problem arising in cellular networks. The problem is as follows: Each object has two attributes, weight and fragility. The goal is to pack objects into bins such that, for every bin, the sum of weights of objects in that bin is no more than the fragility of the most fragile object in that bin.

We look for approximation algorithms for this problem. We provide a 2-approximation to the problem of minimizing the number of bins. We also show a lower bound of $3/2$. Unlike in traditional bin packing, this bound holds in the asymptotic case. We then consider the approximation with respect to fragility and provide a 2-approximation algorithm. Our algorithm uses the same number of bins as the optimum but the weight of objects in a bin can exceed the fragility by a factor of 2.

1. Introduction

We consider a generalization of the classical bin-packing problem. In the traditional bin-packing problem, we are given a collection of n objects, where each object has an arbitrary non-negative weight no more than 1. These objects are to be placed in bins, such that the total weight of the objects in each bin is at most 1, and the total number of bins used is minimized. In our problem the objects are fragile, that is, in addition to its weight, each object has a fragility associated with it. The more fragile an object, the lower is its fragility value. In this model, an object breaks down if the total weight of objects in the bin in which it is placed exceeds the fragility of this object. Thus, we seek to place the objects in bins such that for every bin, the sum of the weights of objects in that bin is no more than the fragility of the most fragile object in that bin.

Clearly, if we set the fragility of each object to 1, this places an upper bound of 1 on the total weight in each bin, and hence our problem reduces to the traditional bin-packing problem. As classical bin packing is known to be NP-hard, we look for efficient approximation algorithms for our fragile bin packing problem. There are two natural notions of approximation for the problem. The first (as in the classical bin packing problem) is the number of bins used. We give a factor 2 approximation with respect to the number of bins. We also show that, unlike the traditional bin packing problem which admits an asymptotic PTAS [de la Vega and Lueker, 1981, Karmarkar and Karp, 1982], our problem cannot be approximated by a polynomial time algorithm to a factor better than $3/2$, unless $P=NP$. Second, we can also consider an approximation with respect to fragility. Specifically, we can imagine an algorithm which produces an assignment of objects to bins that uses the same number of bins as the optimal algorithm, however we might violate the fragility upto some factor greater than 1. That is, the sum of weights of objects in a bin can be upto some factor times more than the fragility of the most fragile object in that bin. We will give an algorithm which achieves an approximation ratio of 2 with respect to this measure.

The bin-packing problem and its various variants have been extensively studied, in the context of approximation algorithms [de la Vega and Lueker, 1981, Karmarkar and Karp, 1982], online algorithms [Garey et al., 1972, Yao, 1980, van Vliet, 1996] and average case analysis [Shor, 1984, Coffman and Lueker, 1991]. A survey article by Coffman et. al. [Coffman et al., 1997] and numerous references therein provide a comprehensive list of the results in this area. While our variant of the problem is fairly natural, we do not know of any previous work on this problem. Our variant is motivated by the following problem arising in frequency allocation in cellular networks.

Consider a base station in a cellular network which has various users communicating with it on various frequency channels. In CDMA (a commonly used technology in wireless systems), a wide channel with capacity much larger than an individual user's information rate is allocated and multiple users share a single channel. To utilize the bandwidth efficiently it is necessary to assign as many users as possible to a single channel. However, the tradeoff in assigning

many users to a particular channel is that there is loss of quality (given by signal to noise ratio (SNR), usually denoted by β), due to interference among users sharing the channel. The goal of the channel assignment scheme is to minimize the number of channels used while guaranteeing that each user achieves a minimum SNR of β .

To be more specific, consider n users communicating with a central base station. If user i transmits with power p_i , then the signal received by the base station is $s_i = p_i g_i$, where g_i is the channel gain for user i . Let $U(t)$ be the set of users transmitting to the base station at time t , then transmission from user $i \in U(t)$ is successful if and only if

$$\frac{s_i}{N_0 + \sum_{j \in U(t), j \neq i} s_j} \geq \beta$$

where β is the signal-to-noise ratio (SNR) requirement for successful communication and N_0 is the background noise power. While we consider only single cell networks in this paper, N_0 can be used to model the interference due to thermal noise and the interference due to the transmission of the neighboring channels in a multicellular network.

Thus, we can model the frequency channels as bins, the users are objects with weights equal to the amount of power received at the base station. Since the users can tolerate a power of up to $1/\beta$ times their power from other users on the same frequency channel, we set their fragilities to be $1 + 1/\beta$ times their weight. Thus the frequency allocation problem is a special case of our general bin-packing problem with fragile objects.

We note that a widely studied solution (assuming certain hardware capabilities) to the frequency allocation problem is *power control* [Hanly and Tse, 1999, Yates, 1995, Yates and Huang, 1995, Zander, 1993, Viterbi, 1995]. Here, the users control their transmission power levels such that the power received at the base station (the weight in our bin packing problem) is almost equal for all users. However, power control involves expensive and relatively sophisticated hardware. For example, most of the wireless Ethernet cards available presently cannot adjust their transmission power [Feeney and Nilsson, 2001]. Thus, our work applies to scenarios where power control is not available.

2. Problem Formulation

We are given n objects $1, \dots, n$ with weights w_1, \dots, w_n and fragility f_1, \dots, f_n respectively. We have bins (of potentially infinite capacity), in which we want to place the objects. Suppose objects j_1, j_2, \dots, j_{n_i} are assigned to bin B_i . Then the assignment to bin B_i is feasible for bin B_i if $w_{j_1} + \dots + w_{j_{n_i}} \leq \min\{f_{j_1}, f_{j_2}, \dots, f_{j_{n_i}}\}$. In words, the total weight of the objects assigned to bin B_i does not exceed the fragility of any object in the bin.

Let A be an assignment of objects to bins. We say that A is feasible if each object is assigned to some bin, and each bin is feasible. The cost of the assignment A is the number of bins used by the assignment. Let OPT be the assignment which uses the minimum number of bins. We also abuse the

notation to let OPT also denote the number of bins used the the optimum solution.

We consider two measures of approximation. In the first, we measure the approximation ratio by the number of bins used by our algorithm against that used by the optimum. We will denote this as “approximation with respect to the number of bins”.

For the second measure, we define an assignment of objects to bin B_i to be c -feasible if $w_{j_1} + \dots + w_{j_{n_i}} \leq c \min\{f_{j_1}, f_{j_2}, \dots, f_{j_{n_i}}\}$. That is, the bin is allowed to deviate from its fragility constraint by a factor of up to c . An assignment is c -feasible if each bin is c -feasible. According to this measure, our algorithm has an approximation ratio of c if it finds an assignment which uses at most OPT bins, but only guarantees that each bin is c -feasible (as opposed to being 1-feasible as in the optimum solution). We will denote this as “approximation with respect to fragility”.

3. Approximation with Respect to the Number of Bins

3.1. Algorithm

Before we describe the algorithm, we first fix some notation and labeling conventions.

We will label the users according to the non-decreasing order of their fragility values, thus $f_n \geq f_{n-1} \geq \dots \geq f_2 \geq f_1$. We will say that i is to the left of j if $f_i > f_j$, in the above ordering. Thus 1 is the rightmost object. Consider an assignment A of objects into bins. For a bin B , let $r(B)$ denote the index of the object with the least fragility assigned to bin B . Suppose A uses m bins, we will label the bins B_1, \dots, B_m such that $r(B_1) < r(B_2) < \dots < r(B_m)$. Thus, the bin containing user 1 is bin 1, the bin containing the rightmost object other than the ones in bin 1 will be denoted by bin 2 and so on.

To avoid trivialities, we will assume that $w_i \leq f_i$ for each object i . Observe that if $w_i > f_i$ for some i , then the object cannot be placed in any bin, hence there is no feasible solution.

Consider the following greedy algorithm:

- 1 Sort and label all the objects according to non-decreasing f_i , i.e. $f_n \geq \dots \geq f_1$.
- 2 Initialize $i \leftarrow 1$, $j \leftarrow 1$ and $w \leftarrow 0$, $f \leftarrow f_i$.
- 3 While $i \leq n$
 - If $B_j \cup \{i\}$ is feasible ($w + w_i \leq \min\{f, f_i\}$)
 - $B_j \leftarrow B_j \cup \{i\}$, $w \leftarrow w + w_i$, $f \leftarrow \min\{f, f_i\}$
 - Else
 - $j \leftarrow j + 1$
 - $B_j \leftarrow \{i\}$, $w \leftarrow w_i$, $f \leftarrow f_i$
 - $i \leftarrow i + 1$

Observe that the algorithm above fills the bins greedily starting from object 1. It adds objects to a bin, until adding the next object makes the bin infeasible. At this point, the algorithm creates a new bin and continues to add objects into it.

We will say that a solution is *banded* if each bin consists only of consecutive objects. Formally, if $i > j$, then for any k, l such that $k \in B_i$ and $l \in B_j$, we have that $k > l$. It is easy to see that the greedy algorithm above produces the optimum banded solution, and requires time $O(n \log n)$.

We will show that the greedy algorithm above requires at most twice the number of bins required by the optimum algorithm. The idea of the proof is that, given an optimum solution using OPT bins, we first produce a *fractional banded solution* (defined below) which consists of OPT bins. We then round this fractional solution to produce a banded solution with at most $2 \cdot OPT$ bins.

3.2. Worst-Case Analysis

Consider an optimum assignment. We will abuse notation and let B_1, \dots, B_{OPT} denote the sets of objects assigned to bins B_1, \dots, B_{OPT} respectively. Define W_i to be the sum of all the weights of objects in B_i , thus $W_i = \sum_{j \in B_i} w_j$. From the optimum solution construct a fractional banded solution as follows:

Greedily assign objects, possibly fractionally, to bins B'_i such that sum of the weights of objects in bin B'_i is exactly W_i . That is, start from the rightmost object, put it in bin B'_1 . Let k be such that $w_1 + \dots + w_k \leq W_1$ and $w_1 + \dots + w_{k+1} > W_1$. In this case, bin B'_1 will comprise of objects $1, 2, \dots, k$ and a fraction x of the object $k+1$ such that $w_1 + \dots + w_k + xw_{k+1} = W_1$. Continue packing the remaining objects in bin B'_2 and so on.

Informally, if we think of the objects laid out (in the order of their fragility) along the line as segments of length equal to their respective weight, then what we are doing is marking intervals of length equal to the weight of each bin of the optimal solution. In the fractional banded solution, these intervals are interpreted as the bins, and each object is put into the bin corresponding to the interval in which it falls, fractionally if it overlaps two of the intervals.

We now define some notation to make this precise. Let l'_k and r'_k be the index of the leftmost and the rightmost object having a non-zero fraction in B'_k . Let $x_{l'_k}$ and $x_{r'_k}$ denote the fraction of l'_k and r'_k in B'_k , hence $0 \leq x_{l'_k}, x_{r'_k} \leq 1$. Then, $\forall k, r'_k \leq l'_{k-1}$ and $x_{l'_k} w_{l'_k} + \sum_{i, l'_k < i < r'_k} w_i + x_{r'_k} w_{r'_k} = W_k$.

Note that B'_k completely contains all objects i such that $l'_k < i < r'_k$. In the case when $l'_k = r'_k$, then B'_k contains at most one object, moreover, since $w_i \leq f_i$ for all i , we know that B'_k contains exactly one object and $x_{l'_k} = x_{r'_k} = 1$.

We now observe a few properties of this assignment.

Property 1 *The number of bins B' formed thus will be exactly OPT .*

Proof: This follows directly from the procedure to construct the fractional bins. ■

Property 2 Let r'_i and r_i be the index of the object with the least fragility (hence rightmost) in bin B'_i and B_i respectively. Then $r_i \leq r'_i$.

Proof: In OPT, all objects to the right of r_i will lie in bins B_j such that $j < i$. Hence the sum of the weights of the objects to the right of r_i will be at most $\sum_{j=1}^{i-1} W(j)$. Since the sum of the objects to the right of r'_i in the fractional solution is exactly $\sum_{j=1}^{i-1} W(j)$, it follows that the rightmost object of B'_i is to the left of the rightmost object of bin B_i in OPT. ■

We will now convert this fractional banded solution into a banded solution which uses at most $2 \cdot OPT$ bins.

Rounding Procedure:

- 1 If an object i has a fraction strictly less than 1 in some bin B'_j , then remove i from all bins B'_j , which contain some fraction of i . Put object i into a new bin by itself. Call these bins of type H .
- 2 Repeat Step 1 until all the bins B'_k contain a 0 or 1 fraction of every object. Now, each B'_k consists of objects in the range $[l''_k, \dots, r''_k]$. Denote these new bins by B''_k .

Clearly the number of bins B''_k will be at most OPT . Also, each bin of type H is feasible since it contains exactly one object. The next lemma shows that the number of bins of type H is also bounded, and all the bins B'' produced are feasible.

Lemma 3 *The above rounding procedure produces a banded assignment of objects to bins, the number of bins used is at most $2 \cdot OPT$ and each bin is feasible.*

Proof: We first bound the number of bins of type H . For any bin B'_j at most one object has the property that it lies in both B'_j and B'_{j+1} . Since the number of such bins B' is at most OPT , it follows that the number of objects covered fractionally is at most OPT .

We now show that the bins B''_k are feasible. The sum of the weights of the objects in bin B''_k is clearly no more than that of B'_k and hence less than W_k . Also, the index of the rightmost object of B''_k is at least as large as that of B'_k , which by Property 2 is at least as large as that of B_k . Thus, as the fragility of the least fragile object of B''_k is not less than that of B_k and the weight of objects in B''_k is not greater than that of B_k , it follows that B''_k is feasible. Thus the result follows. ■

Since the greedy algorithm produces the optimum banded solution, Lemma 3 gives,

Theorem 1 *The number of bins used by the greedy algorithm is at most twice the optimum, and it runs in time $O(n \log n)$.*

3.3. Lower Bound

In this section we show that no algorithm can achieve an approximation ratio of less than $\frac{3}{2}$ unless $P = NP$. This lower bound holds even in the asymptotic case, unlike the traditional bin packing which has a PTAS in the asymptotic case.

We reduce the partition problem to our packing problem. Given a set A with each object having a size $s(a) \in \mathbb{Z}^+$, the partition problem is to decide if $\exists A' \subset A$, such that $\sum_{a \in A'} s(a) = \sum_{a \in A - A'} s(a)$ [Garey and Johnson, 1979]. We can easily transform this into our problem. Let s be the sum of objects in A . We form an instance of our problem, where we have n objects with weights $s(a)$ and all fragilities set to $s/2$.

Clearly, there exists a solution to the partition problem iff the objects can be packed in 2 bins. To show a lower bound for the asymptotic case, we construct m copies, A_1, \dots, A_m of the set A . The set A_i is obtained from A by scaling both the weight and the fragility of the objects by s^{i-1} (i.e. in A_i the weight of object a is $s^{i-1}s(a)$ and the fragility is $s/2 \cdot s^{i-1}$).

We solve the bin packing problem on the instance $A_1 \cup \dots \cup A_m$. Observe that, if $i \neq j$ no two objects from A_i and A_j can lie in the same bin, because one of the objects will always have more weight than the fragility of the other object. Thus, if we obtain a solution which uses less than $3m$ bins we know that the partition problem has a solution. Hence the bin packing problem has an optimum solution which either uses only $2m$ bins or at least $3m$ depending on whether the partition problem has a solution. Thus we have shown that

Theorem 2 *The number of bins used the bin packing problem cannot be approximated to within a factor of $\frac{3}{2}$ by a polynomial time algorithm, unless $P = NP$.*

4. Approximation with Respect to Fragility

In this section we will construct a solution where the number of bins used is OPT , but each of the bins is only guaranteed to be 2-feasible rather than 1-feasible as in the optimum solution.

Given some optimum solution OPT , we consider the fractional solution using OPT bins as in Section 3. However we modify our rounding procedure.

Rounding Procedure:

- 1 Given a bin B'_i in the fractional solution, form a new bin B''_i by including the rightmost object of B'_i completely in B''_i .
- 2 If $x_{i'} < 1$, (that is if the leftmost object of B'_i occurs fractionally in B'_i) then remove it from B''_i .

Lemma 4 *The rounding procedure gives an assignment of objects to bins which is banded, has number of bins at most OPT , and each bin is 2-feasible.*

Proof: In the rounding procedure, observe that the rightmost object of the bin B''_i remains the same as the rightmost object of B'_i . Moreover, if B'_i contains an

object which is not contained by B_i'' , we know that the object is l_i' and $x_i' < 1$. Hence $x_{r_{i+1}'} < 1$ and l_i' will be included in B_{i+1}'' . Thus every object (except possibly n) lies in some B_i'' for some i . To see that n also lies in B_{OPT}'' , observe that $l_n' = 1$, since $w_n \leq f_n$. Hence the number of bins B_i'' is exactly OPT , and they contain all the objects.

We now show that each bin B_i'' is 2-feasible. Consider the bin B_i'' , we know that its rightmost object is r_i' , and by Property (2), $r_i' \geq r_i$. Moreover we know that the sum of weights of objects in B_i'' is at most $W_i + w_{r_i}$.

Now since bin B_i was feasible, $W_i \leq f_{r_i}$. As $w_{r_i} \leq f_{r_i}$ it follows that $W_{r_i} + w_{r_i} \leq 2f_{r_i} \leq 2f_{r_i'}$. ■

Thus there exists a banded solution where the bins are 2-feasible and the number of bins used is at most OPT . So our algorithm will be as follows.

- 1 Sort and label all the objects according to non-decreasing f_i , i.e. $f_n \geq \dots \geq f_1$.
- 2 Initialize $i \leftarrow 1$, $j \leftarrow 1$ and $w \leftarrow 0$, $f \leftarrow f_i$.
- 3 While $i \leq n$
 - If $B_j \cup \{i\}$ is 2-feasible (i.e. $w + w_i \leq 2 \min\{f, f_i\}$)
 - $B_j \leftarrow B_j \cup \{i\}$, $w \leftarrow w + w_i$, $f \leftarrow \min\{f, f_i\}$
 - Else
 - $j \leftarrow j + 1$
 - $B_j \leftarrow \{i\}$, $w \leftarrow w_i$, $f \leftarrow f_i$
 - $i \leftarrow i + 1$

The above algorithm produces an optimal banded solution which is 2-feasible. It follows from Lemma 4 that

Theorem 3 *The algorithm produces a solution where each bin is 2-feasible and the number of bins used is at most that used by the optimum 1-feasible solution. Moreover the algorithm runs in time $O(n \log n)$.*

5. Conclusion

In this paper we considered a variant of the bin packing problem arising in frequency allocation of cellular networks. We showed that two simple heuristics for assigning objects to bins based on their fragility are provably good. We believe that the results can be extended in various directions along the lines of other results for the traditional bin-packing problem. For example, an average case analysis for the case when the weights and fragilities come from some distribution would be interesting. It would also be interesting to study the online case, and the case when the objects are present for a temporary duration (this would correspond to users entering and leaving the cellular network). Also the problem of d -dimensional bin packing with fragile objects would be interesting.

Acknowledgments

The authors are grateful to Professors Avrim Blum and Volker Diekert for useful comments which helped improving the presentation of the paper.

References

- [Coffman et al., 1997] Coffman, E., Garey, M., and Johnson, D. (1997). Approximation algorithms for bin packing: a survey. In *D. S. Hochbaum, ed. Approximation Algorithms for NP-hard Problems*, pages 46–93. PWS Publishing, Boston.
- [Coffman and Lueker, 1991] Coffman, E. G. and Lueker, G. S. (1991). *Probabilistic Analysis of Packing and Partitioning Algorithms*. Wiley-Interscience.
- [de la Vega and Lueker, 1981] de la Vega, W. F. and Lueker, V. (1981). Bin packing can be solved within $1 + \epsilon$ in linear time. *Combinatorica*, pages 349–355.
- [Feeney and Nilsson, 2001] Feeney, L. and Nilsson, M. (2001). Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In *Proceedings of IEEE INFOCOM*.
- [Garey and Johnson, 1979] Garey, M. and Johnson, D. (1979). *Computers and Intractability: A Guide to the Theory of NP Completeness*. W.H. Freeman and Company, New York.
- [Garey et al., 1972] Garey, M. R., Graham, R. L., and Ullman, J. D. (1972). Worst-case analysis of memory allocation algorithms. In *ACM Symposium on Theory of Computing*, pages 143–150.
- [Hanly and Tse, 1999] Hanly, S. and Tse, D. (1999). Power control and capacity of spread-spectrum wireless networks. *Automatica*, 35(12):1987–2012.
- [Karmarkar and Karp, 1982] Karmarkar, N. and Karp, R. M. (1982). An efficient approximation scheme for the one-dimensional bin-packing problem. In *Proc. 23rd Ann. Symp. on Foundations of Computer Science*.
- [Shor, 1984] Shor, P. W. (1984). The average-case analysis of some on-line algorithms for bin packing. In *IEEE Symposium on Foundations of Computer Science*, pages 193–200.
- [van Vliet, 1996] van Vliet, A. (1996). On the asymptotic worst case behavior of harmonic fit. *J. Algorithms*, 20:113–136.
- [Viterbi, 1995] Viterbi, A. (1995). *CDMA – Principles of Spread Spectrum Communication*. Addison-Wesley.
- [Yao, 1980] Yao, A. C.-C. (1980). New algorithms for bin packing. *Journal of the ACM*, 27:207–227.
- [Yates, 1995] Yates, R. (1995). A framework for uplink power control in cellular radio systems. *IEEE Jour. Sel. Areas Communication*, 13(7):1341–1348.
- [Yates and Huang, 1995] Yates, R. and Huang, C. (1995). Integrated power control and base station assignment. *IEEE Transactions on Vehicular Technology*, 44(3):638–644.
- [Zander, 1993] Zander, J. (1993). Transmitter power control for co-channel interference management in cellular radio systems. In *Proc. WINLAB Workshop*.