

From Computers and Intractability

by Garey & Johnson

6.1 PERFORMANCE GUARANTEES FOR APPROXIMATION ALGORITHMS

guarantee for a problem. Finally, in Section 6.3, we discuss some practical considerations concerning the applicability of this type of result, and we mention some recent theoretical work directed toward analyzing the "average performance" of heuristic algorithms.

6.1 Performance Guarantees for Approximation Algorithms

Let us begin by presenting a formal description of what we will mean by an "optimization problem."

A *combinatorial optimization problem* Π is either a *minimization problem* or a *maximization problem* and consists of the following three parts:

- (1) a set D_Π of *instances*;
- (2) for each instance $I \in D_\Pi$, a finite set $S_\Pi(I)$ of *candidate solutions* for I ; and
- (3) a function m_Π that assigns to each instance $I \in D_\Pi$ and each candidate solution $\sigma \in S_\Pi(I)$ a positive rational number $m_\Pi(I, \sigma)$, called the *solution value* for σ .

If Π is a minimization [maximization] problem, then an *optimal solution* for an instance $I \in D_\Pi$ is a candidate solution $\sigma^* \in S_\Pi(I)$ such that, for all $\sigma \in S_\Pi(I)$, $m_\Pi(I, \sigma^*) \leq m_\Pi(I, \sigma)$ [$m_\Pi(I, \sigma^*) \geq m_\Pi(I, \sigma)$]. We will use $\text{OPT}_\Pi(I)$ to denote the value $m_\Pi(I, \sigma^*)$ of an optimal solution for I (usually dropping the subscript Π when the problem is clear from context).

An algorithm A is an *approximation algorithm* for Π if, given any instance $I \in D_\Pi$, it finds a candidate solution $\sigma \in S_\Pi(I)$. The value $m_\Pi(I, \sigma)$ of the candidate solution σ found by A when applied to I will be denoted by $A(I)$. If $A(I) = \text{OPT}(I)$ for all $I \in D_\Pi$, then A is called an *optimization algorithm* for Π .

These definitions can be illustrated by considering our old friend the traveling salesman problem. It is a minimization problem, and the set of instances consists of all finite sets of cities together with their intercity distances. The candidate solutions for a particular instance are all the permutations of the given cities. The solution value for such a permutation is the length of the corresponding tour. Thus an approximation algorithm for this problem need only find some permutation of the given set of cities, whereas an optimization algorithm must always find a permutation that corresponds to a minimum length tour.

If, as in this case, the optimization problem is NP-hard, then we know that a polynomial time optimization algorithm cannot be found unless $P = \text{NP}$. A more reasonable goal is that of finding an approximation algorithm A that runs in low-order polynomial time and that has the property that, for all instances I , $A(I)$ is "close" to $\text{OPT}(I)$. The following example illustrates the type of results we will be interested in.

Consider the “bin packing” problem. Given a finite set $U = \{u_1, u_2, \dots, u_n\}$ of “items” and a rational “size” $s(u) \in [0, 1]$ for each item $u \in U$, find a partition of U into disjoint subsets U_1, U_2, \dots, U_k such that the sum of the sizes of the items in each U_i is no more than 1 and such that k is as small as possible. We can view each subset U_i as specifying a set of items to be placed in a single unit-capacity “bin,” with our objective being to pack the items from U in as few such bins as possible.

This problem is NP-hard in the strong sense (it contains 3-PARTITION as a special case), so there is little hope of finding even a pseudo-polynomial time optimization algorithm for it. However, there are a number of simple approximation algorithms for it that are worth considering.

One of these is known as the “First Fit” algorithm. Imagine that we start with an infinite sequence B_1, B_2, \dots of unit-capacity bins, all of which are empty. The algorithm then places the items into the bins, one at a time, in order of increasing index. It does so according to the following simple rule: always place the next item u_i into the lowest-indexed bin for which the sum of the sizes of the items already in that bin does not exceed $1 - s(u_i)$. In other words, u_i is always placed into the first bin in which it will fit (without exceeding the bin capacity). Figure 6.1 shows an example, where each item is represented by a rectangle having height proportional to its size.




Figure 6.1 An example of a First Fit placement, where u_6 is placed in bin B_2 since that is the lowest indexed bin in which it fits.

Intuitively this seems to be a very natural and reasonable algorithm. It never starts a new bin until all the nonempty bins are too full. What can be proved about its performance?

A first observation relates the number of bins used by First Fit to a natural function of the problem parameters. Let us use “FF” as an abbreviation for First Fit. Then we have that, so long as $FF(I) > 1$,

This is because the packing whose content is higher indexed such could not have been essentially the best form $U = \{u_1, u_2, \dots\}$ items will fit in the item sizes is $(n/2) +$ choosing $\epsilon > 0$ suitable.

This observation can be relative to an

We thus conclude that

However, First Fit is a following theorem [Graham, Johnson, & McGeoch, 1974]:

Theorem 6.1 For all

Furthermore, there exists

Thus Theorem 6.1 shows that the First Fit algorithm is significantly more than bad. (A slight improvement is given by (17/10)OPT, given by Graham, Johnson, & McGeoch, 1974, of Theorem 6.1, which is given by the theorem in Figure 6.2, for which a bound in the theorem is given.)

a finite set $U_i \in [0,1]$ for each $i = 1, 2, \dots, k$ such that more than 1 and set U_i as specified, with our observations as possible.

THE 3-PARTITION PROBLEM. Even a pseudo-polynomial time algorithm, there are a few worth considering.

Imagine that we have n bins, all of which are of size 1. We pack items into the bins, one at a time, to the following first-indexed bin for which the remaining space does not exceed the size of the item. The first bin in which it does not fit shows an example, which is roughly proportional to

$$FF(I) < \lceil 2 \sum_{i=1}^n s(u_i) \rceil$$

This is because there can be at most one nonempty bin in the First Fit packing whose contents total $1/2$ or less. (If not, the first item to go in the higher indexed such bin would have fit in the lower indexed such bin and could not have been placed elsewhere by First Fit.) That this bound is essentially the best possible is apparent when we consider instances of the form $U = \{u_1, u_2, \dots, u_n\}$ where $s(u_i) = 1/2 + \epsilon, 1 \leq i \leq n$. Here no two items will fit in the same bin, so $FF(I) = n$, even though the sum of the item sizes is $(n/2) + n\epsilon$, which can be made as close to $n/2$ as desired by choosing $\epsilon > 0$ suitably small.

This observation also gives us a bound on how bad a First Fit packing can be relative to an optimal packing, since we clearly have

$$OPT(I) \geq \lceil \sum_{i=1}^n s(u_i) \rceil$$

We thus conclude that, for all instances I ,

$$FF(I) < 2 \cdot OPT(I)$$

However, First Fit actually obeys a better bound of this form, given by the following theorem from [Johnson, Demers, Ullman, Garey, and Graham, 1974]:

Theorem 6.1 For all instances I of the bin packing problem,

$$FF(I) \leq \frac{17}{10} OPT(I) + 2$$

Furthermore, there exist instances I with $OPT(I)$ arbitrarily large such that

$$FF(I) \geq \frac{17}{10} (OPT(I) - 1)$$

Thus Theorem 6.1 characterizes the asymptotic worst-case performance of the First Fit algorithm. First Fit never differs from optimal by significantly more than 70 percent and it can on occasion be essentially this bad. (A slight improvement on the constant term in the upper bound, replacing $(17/10)OPT(I) + 2$ by $\lceil (17/10)OPT(I) \rceil$, is obtained in [Garey, Graham, Johnson, and Yao, 1976]). Although we omit the lengthy proof of Theorem 6.1, we note that worst-case behavior almost as bad as that given by the theorem can be seen from the class of examples described in Figure 6.2, for which $FF(I) = (5/3)OPT(I)$. (The proof of the lower bound in the theorem is a complicated extension of these examples.)

placed in bin B_2 since

able algorithm. It is full. What can be

by First Fit to a naive algorithm. 'FF' as an abbreviation for First Fit.

$$\text{Instance } I: U = \{u_1, u_2, \dots, u_{18m}\}, s(u_i) = \begin{cases} 1/7 + \epsilon & 1 \leq i \leq 6m \\ 1/3 + \epsilon & 6m < i \leq 12m \\ 1/2 + \epsilon & 12m < i \leq 18m \end{cases}$$

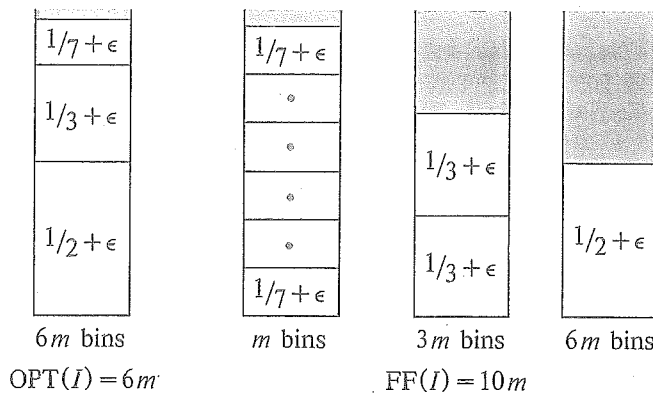


Figure 6.2 Instances I with $\text{OPT}(I)$ arbitrarily large such that $\text{FF}(I)$ equals $(5/3)\text{OPT}(I)$.

These results for First Fit provide a starting point for analyzing approximation algorithms for bin packing. One can now go on to analyze other algorithms that might have better guarantees. An obvious modification to the First Fit algorithm, for example, is that obtained by using the following more sophisticated placement rule: Always place the next item u_i in that bin which has current contents closest to, but not exceeding, $1 - s(u_i)$ (choosing the one with lowest index in case of ties). This is known as the "Best Fit" algorithm. Unfortunately, and perhaps surprisingly, Best Fit has essentially the same worst case performance as First Fit [Johnson et al., 1974].

A better approximation algorithm is obtained by observing that the worst performance for First Fit (and Best Fit) seems to occur when the smaller items appear before the larger items in the ordering used by the algorithm. Suppose that, instead of merely taking the items from U in the given order, we first sort them by size and reindex them so that $s(u_1) \geq s(u_2) \geq \dots \geq s(u_n)$. The algorithm that applies First Fit to such a reordered list is called the "First Fit Decreasing" algorithm (FFD). Its performance is characterized by the following theorem, due to Johnson [1973] (the proof is sketched in [Johnson et al., 1974]):

Theorem 6.2 For all instances I of the bin packing problem,

$$\text{FFD}(I) = \frac{11}{9} \text{OPT}(I) + 4$$

Furthermore, there exist instances I with $\text{OPT}(I)$ arbitrarily large such that

Thus First Fit Decreasing is at most 11 percent worse than optimal result holds for the class of problems illustrated in Figure 6.3. Figure 6.3 illustrates a class of problems for which both cases.

Instance $I: U =$




Figure 6.3 Instances I with $\text{FFD}(I) = (11/9)\text{OPT}(I) + 4$

The proof of the above theorem, whose recapitulation is given in the appendix, is based on the rule for problems of the form $(1/4 - 2\epsilon, 1/4 + \epsilon, 1/2 + \epsilon)$ for other problems has indeed even for bin packing problems. We are willing to settle for we


Further modifications to the First Fit algorithm (see [Johnson, 1973], [Yeh, 1973]) have been proposed, but no sub-constant approximation algorithm is known.

In summary, our analysis of the bin packing problem might be considered a step forward but apparently so far. It is worth noting that by proving bounds on the performance of approximation algorithms, we can use examples to verify the

$$\text{FFD}(I) = \frac{11}{9} \text{OPT}(I)$$

Thus First Fit Decreasing is guaranteed never to be more than about 22 percent worse than optimal, and it can on occasion be this bad. An identical result holds for the analogous "Best Fit Decreasing" algorithm. Figure 6.3 illustrates a class of examples that suffice to prove the lower bound in both cases.

$1 \leq i \leq 6m$
 $6m < i \leq 12m$
 $12m < i \leq 18m$



that $\text{FF}(I)$ equals

For analyzing approximation to analyze other algorithm modification to the using the following item u_i in that bin g , $1 - s(u_i)$ (chosen known as the "Best", Best Fit has essentially [Johnson et al., 1974]. observing that the to occur when the ring used by the algorithms from U in the index them so that s First Fit to such a thm (FFD). Its performance to Johnson [1973]

$$\text{Instance } I: U = \{u_1, u_2, \dots, u_{30m}\}, s(u_i) = \begin{cases} 1/2 + \epsilon & 1 \leq i \leq 6m \\ 1/4 + 2\epsilon & 6m < i \leq 12m \\ 1/4 + \epsilon & 12m < i \leq 18m \\ 1/4 - 2\epsilon & 18m < i \leq 30m \end{cases}$$

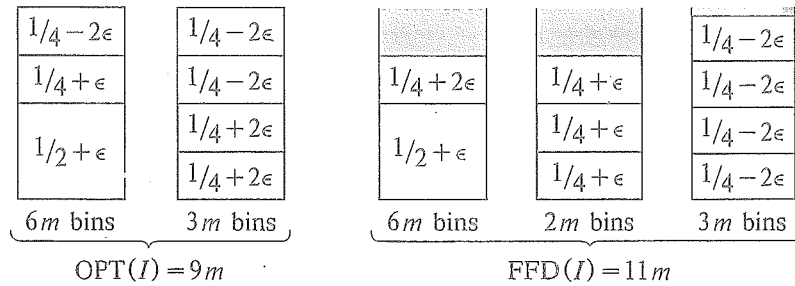


Figure 6.3 Instances I with $\text{OPT}(I)$ arbitrarily large such that $\text{FFD}(I)$ equals $(11/9)\text{OPT}(I)$.

The proof of the upper bound involves an extremely detailed case analysis, whose recapitulation here would require more pages than we have allotted to this entire chapter. (Although such lengthy proofs appear to be the rule for problems similar to bin packing, we note that results like these for other problems have been obtained without such Herculean effort, and indeed even for bin packing much shorter proofs are obtainable if we are willing to settle for weaker bounds.)

Further modifications of First Fit Decreasing have been suggested ([Johnson, 1973], [Yao, 1978a]) in hopes of obtaining a polynomial time approximation algorithm for bin packing with an even better performance guarantee, but no substantially better bound has yet been proved for any of them.

In summary, our analysis of approximation algorithms for the bin packing problem might be described as follows: We started with a straightforward but apparently sensible algorithm and analyzed its performance, both by proving bounds on what could happen in the worst case and by devising examples to verify that these bounds could not be improved. With this

blem,

arbitrarily large such that

analysis in mind, and especially the insight it provided as to the drawbacks of our initial algorithm, we could then seek alternative algorithms (perhaps just more complicated versions of the original one) and analyze them. We also settled on a general form for our guarantees, in terms of ratios, which was useful for comparison purposes and which seems to express nearness to optimality in a reasonable way. This general approach can serve as a model for our study of other NP-hard optimization problems and indeed has been widely applied (although, of course, on occasion other types of guarantees may be more appropriate or easier to prove, for example, see [Cornuejols, Fisher, and Nemhauser, 1977], [Nemhauser, Wolsey, and Fisher, 1978]).

To formalize this approach, let us make a few more definitions. If Π is a minimization [maximization] problem, and I is any instance in D_{Π} , we define the ratio $R_A(I)$ by

$$R_A(I) = \frac{A(I)}{\text{OPT}(I)} \quad \left[R_A(I) = \frac{\text{OPT}(I)}{A(I)} \right]$$

The *absolute performance ratio* R_A for an approximation algorithm A for Π is given by

$$R_A = \inf \{ r \geq 1 : R_A(I) \leq r \text{ for all instances } I \in D_{\Pi} \}$$

The *asymptotic performance ratio* R_A^{∞} for A is given by

$$R_A^{\infty} = \inf \left\{ r \geq 1 : \begin{array}{l} \text{for some } N \in \mathbb{Z}^+, R_A(I) \leq r \text{ for all} \\ I \in D_{\Pi} \text{ satisfying } \text{OPT}(I) \geq N \end{array} \right\}$$

Notice that we have defined these ratios in such a way that the ratio for a minimization problem is the reciprocal of that for a maximization problem. This has been done so that we will have a uniform scale on which to consider approximation algorithms for different types of problems, always having $1 \leq R_A \leq \infty$ and $1 \leq R_A^{\infty} \leq \infty$, with a ratio that is closer to 1 indicating better performance.

Notice also that R_A need not equal R_A^{∞} . Although Theorem 6.2 shows that $R_{\text{FFD}}^{\infty} = 11/9$, it is easy to give instances I for which $\text{OPT}(I) = 2$ and $\text{FFD}(I) = 3$, so that $R_{\text{FFD}} \geq 3/2$. The asymptotic ratios seem to be the more important ones for bin packing, although for other problems the absolute ratios may be more appropriate, or it may be the case that $R_A = R_A^{\infty}$ for all the approximation algorithms in which we are interested. At any rate, it will be convenient to have both types of ratios available, and the differences between them are worth keeping in mind when analyzing an approximation algorithm.

As a second example, let us return once more to the traveling salesman problem, only this time with an added restriction. An instance I is still a set C of cities and a specification of the distances between them, but we also require that these distances obey the "triangle inequality"; i.e., for every triple a, b, c of cities from C ,

This condition is met, actual shortest distance tours that visit some of the given distances d_{ij} path from c_i to c_j . It is hard under this restriction considering will have I

Consider the following "Nearest Neighbor" algorithm. In a distance, in [Gavett, 1977] ties. The first city in a partial tour built up so far algorithm chooses for $c_{\pi(i)}$ among all such cities, is as small as possible such city). The next [1977], shows that the behavior than any of the

Theorem 6.3 For all r with triangle inequality

N

Furthermore, for arbitrary for which

NN

The main importance $R_{\text{NN}} = \infty$, a not very recursive construction in Figure 6.4.

The performance much to be desired. Apparently sensible heuristic by arbitrarily large mu

Fortunately, other much better. In fact, described in [Rosenkrantz algorithm that has this being tree." Let us view terms of a complete g