

Random-order bin packing

Edward G. Coffman, Jr.* János Csirik† Lajos Rónyai‡
Ambrus Zsbán§

November 20, 2007

1 Abstract

The average case analysis of algorithms usually assumes independent, identical distributions for the inputs. In [?], Kenyon introduced the *random-order ratio*, a new average case performance metric for bin packing heuristics, and gave upper and lower bounds for it for the Best Fit heuristics. We introduce an alternative definition of the random-order ratio and show that the two definitions give the same result for Next Fit. We also show that the random-order ratio of Next Fit equals to its asymptotic worst case, i.e., it is 2.

2 Introduction

An instance of the classical bin packing problem consists of a positive real C and a list $L = (a_1, a_2, \dots, a_n)$ of items with sizes $0 < s(a_i) \leq C$, $1 \leq i \leq n$; a solution to the problem is a partition of L into a minimum number of blocks, called *bins*, such that the sum of the sizes of the items in each bin is at most the *capacity* C . The capacity is just a scaling parameter; as is customary, we put $C = 1$, and restrict item sizes to the unit interval.

Research on the bin packing problem started over 30 years ago [?], [?]. As the problem is NP-complete [?], many approximation algorithms have been proposed and analyzed. Next Fit (NF) is arguably the most elementary, as it packs items bin by bin, not starting a new bin until an item is encountered that does not fit into the current, *open* bin; in this event the open bin is closed, the new bin becomes the open bin, and no further attempt is made to pack items

*Department of Electrical Engineering, Columbia University, 1312 S.W. Mudd, 500 West 120th Street, New York, NY 10027, USA. egc@ee.columbia.edu

†Department of Computer Science, University of Szeged, Árpád tér 2, H-6720 Szeged, Hungary. csirik@inf.u-szeged.hu

‡MTA SZTAKI, Kende u. 13-17, Budapest, Hungary, Department of Algebra, Budapest University of Technology and Economics, Budapest. ronyai@sztaki.hu

§Department of Algebra, Budapest University of Technology and Economics, ambrus@math.bme.hu

in the bin just closed. A natural generalization of NF is the First Fit algorithm (FF), which never closes bins; it packs each successive item from L in the first (lowest indexed) bin which has enough space for it. Another improvement on NF is the Best Fit algorithm (BF), which packs the next item into the bin which can accommodate it with the smallest capacity left over (with ties resolved in favor of the lower indexed bin).

The most common ways of appraising an approximation algorithm are performance ratios, which give the performance of an approximation algorithm relative to an optimal algorithm. We use the term *competitive ratio* for online algorithms and *approximation ratio* for offline algorithms. Informally, *asymptotic* bounds for algorithm A typically take the form: For given constants $\alpha \geq 1$, $\beta \geq 0$, $A(L) \leq \alpha OPT(L) + \beta$ holds for all lists L ; the bound α is called an asymptotic *worst-case ratio*, or *performance guarantee*. If $\beta = 0$ is a constraint, then the corresponding α is said to be *absolute* rather than asymptotic.

In probabilistic, or average-case, analysis the item sizes are usually assumed to be independent, identically distributed random variables. For a given algorithm A , $A(L)$ is a random variable, whose distribution is the object of the analysis, along with the expected ratio $\mathbf{E}(A(L)/OPT(L))$ or simply the expected performance $\mathbf{E}A(L)$, usually in terms of $\mathbf{E}OPT(L)$. In most cases, computing the distribution of $A(L)$ presents a very difficult problem, so weaker results, such as asymptotic expected values and perhaps higher moments are computed.

Kenyon [?] introduced a new performance metric for an online algorithm A , which compares optimal performance with the performance of A when the ordering of its input is randomized. Specifically, let π denote a permutation of $(1, \dots, n)$ and let L_π denote L reordered by the permutation π of the item indices. Then the *Random-order* performance of A on list L is defined as

$$RR_A(L) = \frac{\mathbf{E}_\pi A(L_\pi)}{OPT(L)}$$

where, for given $|L| = n$, the expectation is taken over all $n!$ equally likely permutations π of the item indices. Let

$$RR_A(n) = \sup_{(L:|L|=n)} RR_A(L).$$

The asymptotic *random-order ratio* is then defined as

$$RR_A := \limsup_{n \rightarrow \infty} RR_A(n).$$

Again, one seeks bounds of the form $\mathbf{E}_\pi A(L_\pi) \leq \alpha OPT(L) + \beta$ for constants α, β with α as small as possible. This new measure gives another perspective on the pessimism of classical worst-case analysis.

Another random-order performance metric, which may be easier to analyze in some cases, focuses on random orderings of lists that give largest performance ratios. Formally, let $\sigma = (L^{(1)}, L^{(2)}, \dots)$ denote a sequence of worst-case lists of

n items under A, i.e., no list of n items produces a larger ratio $A(L)/OPT(L)$ than $L^{(n)}$ does. Define

$$RR_A^* := \limsup_{n \rightarrow \infty} RR_A(L^{(n)})$$

Clearly, $RR_A^* \leq RR_A$. $RR_{NF}^* = 2$ is proved in Section 3, but the evaluation of RR_{BF}^* remains an open problem.

By means of the following example for BF, Kenyon illustrates the dramatic differences one can expect in performance as measured by random-order ratios. For the list

$$L_{2n} = (\underbrace{1/2 - \epsilon, \dots, 1/2 - \epsilon}_n, \underbrace{1/2 + \epsilon, \dots, 1/2 + \epsilon}_n),$$

an optimal algorithm gives, by matching the smaller and larger items, the value $OPT(L_{2n}) = n$. FF and BF give for this list $3n/2 - 1 \leq FF(L_{2n}) = BF(L_{2n}) \leq 3n/2$, and hence an asymptotic ratio of $3/2$, which is not much less than the asymptotic worst-case ratio of $17/10$.

In the random-order scenario, Kenyon approximates random permutations of the input by taking each item independently to be $1/2 + \epsilon$ or $1/2 - \epsilon$ with equal probability, i.e., by a sequence of Bernoulli trials. The resulting sequences can be viewed as unbiased random walks where at each step we move one up or down depending on whether the arriving item is larger or smaller than $1/2$. As is easily verified,¹ the number of unpaired items is bounded by the vertical span of the walk associated with the input sequence. The expected value of the vertical span of an unbiased random walk is well known to be $O(\sqrt{n})$, and so in the random-order scenario, Best Fit is asymptotically optimal for these near worst-case examples, as the expected value of the optimum is $O(n)$.

In fact, the same conclusion holds in the precise model where we consider permutations of the list L_{2n} . Then the corresponding walk will always return to the origin. One can show that the expected vertical span of this random walk is $o(n)$. This can be obtained from the bound for the unbiased walk above by using the chopping technique of Section 2.2. There we exploit the fact that short segments of sufficiently long random permutations behave like Bernoulli sequences.

Kenyon proves that the random-order ratio of BF satisfies

$$1.08 \leq RR_{BF} \leq 1.5,$$

which clearly leaves considerable scope for improvement. Prospects are dimmed by Kenyon's observation that the exact result is thought to be near the lower bound, but the upper bound is by far the more difficult one to prove and hence, presumably, to tighten.

In this paper we will investigate the random-order performance of Next Fit. It is known that 2 is both the absolute and asymptotic worst-case performance of NF, and that the average case performance under the $U(0, 1)$ distribution is

¹This random-walk approach originated with Richard Karp.

4/3 [?]. The next section first applies Kenyon’s initial approach to NF, which is an approximate analysis of the random-order performance on lists that bring out NF’s worst-case behavior. It is then verified that, in contrast to the corresponding BF analysis, this estimate is in fact exact to within constants hidden by our asymptotic notation. In summary, for these lists we get a ratio of 10/7. Section 3 verifies that this analysis does **not** yield the random-order ratio for NF; it shows that, in fact, $RR_{NF}^* = RR_{NF} = 2$, which is the same as the combinatorial worst-case performance [?].

3 Next Fit

We start with an estimate of $RR_{NF}^*(L)$ for worst-case lists L . Section 2.2 then shows that these random-order ratios are in fact exact.

3.1 Approximate random-order performance on worst-case lists

The standard example giving asymptotic worst-case bounds for Next Fit is defined by

$$L_{2n} = \underbrace{(1/2, \epsilon, \dots, 1/2, \epsilon)}_{2n \text{ pairs}}.$$

Here $OPT(L_{2n}) = n + 1$ and $NF(L_{2n}) = 2n$ when $\epsilon < 1/(2n)$.

If we now take the approximate approach of Kenyon, then $4n$ items are drawn independently, each taken to be $1/2$ or ϵ with equal probability. Call the $1/2$ items (i.e., the items with sizes $1/2$) *big* items, and the ϵ items *small* items. The NF packing process is described by the following Markov chain with just four states:

- a: The open bin is empty or it is full with two big items. The open bin is empty only in the initial state.
- b: There is just one item in the open bin and it is big.
- c: There is at least one small item, but no big item in the open bin.
- d: There is one big item and at least one small item in the open bin.

Transitions are shown in Figure ?? and each has probability $1/2$. Note that, if no more than $2n$ items are packed, addition of items of size ϵ can never start a new open bin, since $\epsilon < 1/(2n)$. The chain is aperiodic and irreducible. The stationary probabilities can be computed from the following equations, in which p_x denotes the stationary probability of state x .




Figure 1:

A Markov chain describing Next Fit packing of items, each with size ϵ or $1/2$, and with each size equally likely. Each transition has probability $1/2$.

$$\begin{aligned}
 p_a &= p_b/2 \\
 p_b &= p_a/2 + p_d/2 \\
 p_c &= p_a/2 + p_c/2 \\
 p_d &= p_b/2 + p_c/2 + p_d/2
 \end{aligned}$$

The unique probability distribution solving these equations is $p_a = p_c = 1/7$, $p_b = 2/7$, $p_d = 3/7$. NF starts a new open bin from state d with probability $1/2$, and NF always starts a new open bin in transitions out of state a . Thus, for large n , each item packed starts a new open bin with (asymptotic) probability $1 \cdot \frac{1}{7} + \frac{1}{2} \cdot \frac{3}{7} = \frac{5}{14}$ and since there are $4n$ items, $\mathbf{ENF}(L_{2n}) \sim 20n/14$, $n \rightarrow \infty$. Since $\mathbf{EOPT}(L_{2n}) \sim n$, as $n \rightarrow \infty$, it follows that

$$\lim_{n \rightarrow \infty} \frac{\mathbf{ENF}(L_{2n})}{\mathbf{EOPT}(L_{2n})} = \frac{10}{7} \tag{1}$$

We note in passing that this is only slightly larger than $4/3$, the average-case performance of NF when item sizes are drawn independently from $U(0, 1)$.

We mention that if we start with the list $(1, \epsilon, 1, \epsilon, \dots, 1, \epsilon)$ studied in [?] we will get a ratio of $4/3$ from the same analysis.

3.2 Exact random-order performance on worst-case lists

The analysis below uses well-known monotonicity and subadditivity properties that NF shares with OPT (see, e.g., [?], pages 30, 146). We omit the routine proofs.

Proposition 1. Let $L = (a_1, a_2, \dots, a_n)$ be an arbitrary list. Delete any prefix (a_1, a_2, \dots, a_k) satisfying the condition $\sum_{i=1}^k a_i \leq 1$ from the beginning of the list, and let $L^* = (a_{k+1}, a_{k+2}, \dots, a_n)$ be the new list. Then

$$NF(L) - 1 \leq NF(L^*) \leq NF(L).$$

■

Proposition 2. Suppose L' and L'' are two arbitrary lists. Then

$$NF(L') + NF(L'') - 1 \leq NF(L'L'') \leq NF(L') + NF(L''),$$

where $L'L''$ denotes the concatenation of L' and L'' .

■

Let L_n denote a list having n big (i.e., $1/2$) and n small (i.e., ϵ) items in some order. We compute below the asymptotic performance of NF averaged over all permutations of L_n . Let ξ_n be a permutation of L_n drawn uniformly at random from the set \mathcal{U}_n of $\binom{2n}{n}$ such permutations. Let η_n be a random length- $2n$ list containing only big and small items; η_n has a uniform distribution on the set of 2^{2n} such lists. It is easy to see that η_n can be analyzed by the unconstrained random-walk method. Indeed, we have already proved that $\lim_{n \rightarrow \infty} \mathbf{ENF}(\eta_n) / \mathbf{EOPT}(L_n) = 10/7$; we will now show the following

Theorem 3.

$$\lim_{n \rightarrow \infty} \frac{\mathbf{ENF}(\xi_n)}{\mathbf{EOPT}(L_n)} = \frac{10}{7}.$$

Proof. Let L'_n be a random sequence drawn uniformly from ξ_n . Let us divide L'_n into sublists each of length m , except possibly for the last sublist which has length $2n \bmod m$, where m is an integer to be defined later. Let us denote the sublists by $L'_{n,1}, L'_{n,2}, \dots, L'_{n, \lceil \frac{2n}{m} \rceil - 1}, L'_{n, \lceil \frac{2n}{m} \rceil}$, so that

$$L'_{n,i} = (a_{(i-1) \cdot m + 1}, a_{(i-1) \cdot m + 2}, \dots, a_{i \cdot m})$$

for $1 \leq i \leq \lceil \frac{2n}{m} \rceil - 1$ and

$$L'_{n, \lceil \frac{2n}{m} \rceil} = (a_{(\lceil \frac{2n}{m} \rceil - 1) \cdot m + 1}, \dots, a_{2n}).$$

By repeated application of Proposition 2 we get

$$\mathbf{ENF}(L'_{n,1}) + \mathbf{ENF}(L'_{n,2}) + \dots + \mathbf{ENF}(L'_{n, \lceil \frac{2n}{m} \rceil - 1}) - \left(\frac{2n}{m} - 1 \right) \leq \mathbf{ENF}(L'_n) \leq$$

$$\mathbf{ENF}(L'_{n,1}) + \mathbf{ENF}(L'_{n,2}) + \dots + \mathbf{ENF}(L'_{n, \lceil \frac{2n}{m} \rceil - 1}) + m,$$

where we made use of

$$\mathbf{ENF}(L'_{n, \lceil \frac{2n}{m} \rceil}) \leq m.$$

As $NF(L'_{n,i}), 1 \leq i \leq \lceil \frac{2n}{m} \rceil - 1$ are identically distributed random variables, we have that

$$\left(\left\lceil \frac{2n}{m} \right\rceil - 1 \right) \mathbf{E}NF(L'_{n,1}) - \frac{2n}{m} + 1 \leq \mathbf{E}NF(L'_n) \leq \left(\left\lceil \frac{2n}{m} \right\rceil - 1 \right) \mathbf{E}NF(L'_{n,1}) + m.$$

Now, if $n \rightarrow \infty$ and m is chosen in such a way that $m \rightarrow \infty$ and $n/m \rightarrow \infty$ then $m/n \rightarrow 0$ and so $m = o(n)$. We get

$$\mathbf{E}NF(L'_n) = \frac{2n}{m} \mathbf{E}NF(L'_{n,1}) + o(n).$$

So it is sufficient just to prove that

$$\lim_{n \rightarrow \infty} \frac{\mathbf{E}NF(L'_{n,1})}{\mathbf{E}OPT(L'_{n,1})} = \frac{10}{7}. \quad (2)$$

To this end, we show that for any $L'_{n,1}$

$$e^{-\frac{1}{n^{1/4}}} \leq P(L'_{n,1} = s) 2^m \leq e^{\frac{1}{n^{1/4}}}, \quad (3)$$

where $P(L'_{n,1} = s)$ is the probability that $L'_{n,1}$ contains s large items and $(m-s)$ small items.

This means that the differences between the probabilities of a list from η_n (which is $1/2^m$) and from $L'_{n,1}$ can be made really arbitrarily small, independently of s . Formally, this suffices because $NF(L'_{n,1})$ and $NF(\eta_m)$ are both nonnegative, and therefore

$$e^{-\frac{1}{n^{1/4}}} \leq \frac{\mathbf{E}NF(L'_{n,1})}{\mathbf{E}NF(\eta_m)} \leq e^{\frac{1}{n^{1/4}}},$$

and likewise

$$e^{-\frac{1}{n^{1/4}}} \leq \frac{\mathbf{E}OPT(L'_{n,1})}{\mathbf{E}OPT(\eta_m)} \leq e^{\frac{1}{n^{1/4}}},$$

so (??) follows from (??).

We turn now to the proof of (??). Let S be a sequence drawn from \mathcal{U}_n ; it contains precisely n large and n small items. Suppose further, that $L'_{n,1}$ consists of s large and $m-s$ small items. Then the probability of this prefix is

$$p_{n,s} = \frac{\binom{2n-m}{n-s}}{\binom{2n}{n}}$$

for $0 \leq s \leq m$.

Assume now that $m = \lfloor n^{1/4} \rfloor$. Then clearly $m \rightarrow \infty$ and $n/m \rightarrow \infty$ as $n \rightarrow \infty$.

Let p_n and P_n be the minimal and maximal values of $p_{n,0}, \dots, p_{n,m}$, respectively. We have 2^m possible (short) sequences $L'_{n,1}$, hence

$$p_n \leq \frac{1}{2^m} \leq P_n. \quad (4)$$

From the monotonicity properties of binomial coefficients, we see that $p_n = p_{n,0}$ and $P_n = p_{n,k}$ with $k = \lfloor m/2 \rfloor$. We have

$$\begin{aligned} 1 &\leq \frac{P_n}{p_n} = \frac{\binom{2n-m}{n-k}}{\binom{2n-m}{n}} = \frac{n!(n-m)!}{(n-k)!(n-m+k)!} = \\ &= \frac{n(n-1)\cdots(n-k+1)}{(n-m+k)(n-m+k-1)\cdots(n-m+1)} \leq \frac{n^k}{(n-m+1)^k} = \\ &= \left(1 + \frac{m-1}{n-m+1}\right)^k. \end{aligned}$$

For sufficiently large n we have $n/2 < n-m+1$ and $2m \leq n^{1/2}$, hence

$$\begin{aligned} \left(1 + \frac{m-1}{n-m+1}\right)^k &< \left(1 + \frac{2m}{n}\right)^k \leq \left(1 + \frac{1}{n^{1/2}}\right)^k \leq \left(1 + \frac{1}{n^{1/2}}\right)^{n^{1/4}} = \\ &= \left(\left(1 + \frac{1}{n^{1/2}}\right)^{n^{1/2}}\right)^{n^{-1/4}} < e^{\frac{1}{n^{1/4}}} \rightarrow 1, \end{aligned}$$

as $n \rightarrow \infty$.

This together with (??) gives that

$$e^{-\frac{1}{n^{1/4}}} \leq p_{n,s} 2^m \leq e^{\frac{1}{n^{1/4}}},$$

which is exactly inequality (??), and hence completes the argument. ■

4 Random-order performance of NF

Theorem 4.

$$RR_{NF} = RR_{NF}^* = 2.$$

Proof. It is clear that $RR_{NF} \leq 2$ since for any list L , $NF(L) < 2 \cdot OPT(L)$. Next, let us define

$$L_{2n,k} = \underbrace{(1/2, \underbrace{\epsilon/k, \dots, \epsilon/k}_k, \dots, 1/2, \underbrace{\epsilon/k, \dots, \epsilon/k}_k)}_{2n}.$$

Thus, $L_{2n,k}$ consists of $2n(k+1)$ items; out of these $2n$ are large and $2nk$ are small. Here n is an arbitrarily fixed positive integer, and k will be selected to be sufficiently large (compared to n). Now we have $OPT(L_{2n,k}) = n+1$ when ϵ is small enough. For the random order performance we have to compute the average number of bins over all permutations, i.e., over $P_n := \binom{2nk+2n}{2n}$

permutations. For any permutation we will use at least the optimal number of bins, namely $n+1$ bins. On the other hand, we can characterize a subset of those permutations where we will use exactly $2n$ bins: these are those permutations where we do not have consecutive $1/2$ items. We will show that almost all permutations are of this type.

In fact, consider the orderings of the input $L_{2n,k}$ which have the following pattern:

$$s \dots sls \dots sls \dots sls \dots s,$$

where $s \dots s$ stands for a nonempty block of small items. The number S_n of these sequences is the same as the number of distributions of $2nk - 2n - 1$ indistinguishable balls into $2n + 1$ distinct boxes. The latter number is easily seen to be

$$\binom{2nk - 1}{2n},$$

see for example Section 1.7 in [?] for a discussion, and related counting problems (involving compositions and combinations with repetitions).

We have now

$$\begin{aligned} \frac{S_n}{P_n} &= \frac{\binom{2nk-1}{2n}}{\binom{2nk+2n}{2n}} = \frac{(2nk-1)(2nk-2)\cdots(2nk-2n)}{(2nk+2n)(2nk+2n-1)\cdots(2nk+1)} \geq \\ &\geq \left(\frac{2nk-2n}{2nk+2n}\right)^{2n} = \left(1 - \frac{2}{k+1}\right)^{2n} \geq 1 - \delta, \end{aligned}$$

for any $\delta > 0$ and any n , whenever k is sufficiently large. This means that at almost all permutations we do not have consecutive $1/2$ items and at these permutations we will pack $2n$ bins by the Next Fit. This will ensure that the average number of bins over all permutations can be made arbitrarily close to $2n$.

■

5 Open problems

It was shown here that for a worst case list of Next Fit the random-order performance is asymptotically the same as the average case performance. Is this true for all input lists? In more detail: let (b_1, b_2, \dots, b_m) denote the different sizes of $L = (a_1, a_2, \dots, a_n)$ and let c_i be the multiplicity of b_i in L . Clearly $\sum_{i=1}^m c_i = n$. Let \mathcal{L}_t be a list of t items, where the items are from the set $\{b_1, \dots, b_m\}$ and drawn independently with probabilities $c_1/n, \dots, c_m/n$. On the other hand let L^k denote a concatenation of k copies of L . Is now

$$\lim_{t \rightarrow \infty} \mathbf{ENF}(\mathcal{L}_t) = \lim_{k \rightarrow \infty} \frac{\mathbf{E}_\sigma \mathbf{NF}(L_\sigma^k)}{\mathbf{OPT}(L^k)}$$

true for all lists? We suspect that the answer is yes. We do not know whether the two performance measures are the same for the bin covering problem under the Next Fit algorithm.

More interestingly, is this true for more advanced algorithms like First Fit or Best Fit? This is well worth investigating.

References

- [BF03] J. Boyar and L. M. Favrholdt. The relative worst order ratio for on-line algorithms. CIAC 2003, LNCS 2653: 58–69.
- [CL91] E. G. Coffman, Jr. and G. S. Lueker. *Probabilistic Analysis of Packing and Partitioning Algorithms*. Wiley & Sons, New York, 1991.
- [Com74] L. Comtet. *Advanced Combinatorics*. D. Reidel Publishing Company, Dordrecht, 1974.
- [CSHY80] E. G. Coffman, Jr., K. So, M. Hofri, and A. C. Yao. A stochastic model of bin-packing. *Inf. and Cont.*, 44:105–115, 1980.
- [GGU72] M. R. Garey, R. L. Graham, and J. D. Ullman. Worst-case analysis of memory allocation algorithms. In *Proc. of the 4th Annual ACM Symposium on Theory of Computing*, pages 143–150, New York, 1972. ACM.
- [GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman, New York, New York, 1979.
- [Joh73] D. S. Johnson. *Near-Optimal Bin Packing Algorithms*. PhD thesis, Massachusetts Institute of Technology, Department of Mathematics, Cambridge, 1973.
- [Ken96] C. Kenyon. Best-fit bin-packing with random order. In *Proc. of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 359–364. SIAM, 1996.