Biased random walks, lyapunov functions, and stochastic analysis of best fit bin packing.

Claire Kenyon, Yuval Rabani, Alistair Sinclair

► To cite this version:

Claire Kenyon, Yuval Rabani, Alistair Sinclair. Biased random walks, lyapunov functions, and stochastic analysis of best fit bin packing.. [Research Report] LIP RR-1995-14, Laboratoire de l'informatique du parallélisme. 1995, 2+19p. hal-02102057

HAL Id: hal-02102057

<https://hal-lara.archives-ouvertes.fr/hal-02102057>

Submitted on 17 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Laboratoire de l'Informatique du Parallélisme

Ecole Normale Supérieure de Lyon
Unité de recherche associée au CNRS n°1398

Biased Random Walks, Lyapunov Functions, and Stochastic Analysis of Best Fit Bin Packing

Claire Kenyon
Yuval Rabani
Alistair Sinclair

September 1995

Research Report N° 95-14

Ecole Normale Supérieure de Lyon

46 Allée d'Italie, 69364 Lyon Cedex 07, France

Téléphone : (+33) 72.72.80.00   Télécopieur : (+33) 72.72.80.80

Adresse électronique : lip@lip.ens-lyon.fr

Biased Random Walks, Lyapunov Functions, and Stochastic Analysis of Best Fit Bin Packing

Claire Kenyon
Yuval Rabani
Alistair Sinclair

September 1995

Abstract

We study the Best Fit algorithm for on-line bin packing under the distribution in which the item sizes are uniformly distributed in the discrete range $\{1/k, 2/k, \dots, j/k\}$. Our main result is that, in the case $j = k - 2$, the asymptotic expected waste remains bounded. This settles an open problem of Coffman *et al* [3], and involves a detailed analysis of the infinite multi-dimensional Markov chain underlying the algorithm.

Keywords: Best Fit, Bin-Packing, On-line Algorithms, Average-Case Analysis, Markov Chains

Résumé

Nous étudions l'algorithme "Meilleur choix" pour le problème de la mise en boîtes en-ligne lorsque les données sont tirées de façon aléatoire uniforme dans l'ensemble $\{1/k, 2/k, \dots, j/k\}$. Notre résultat principal est que, dans le cas $j = k - 2$, l'espace moyen perdu reste asymptotiquement borné. Ceci résout un problème de Coffman et al. [3], et utilise une analyse détaillée de la chaîne de Markov infinie multi-dimensionnelle sous-jacente à l'algorithme.

Mots-clés: Mise en Boîtes, Algorithmes En-Ligne, Analyse en Moyenne, Chaînes de Markov.

Biased Random Walks, Lyapunov Functions, and Stochastic Analysis of Best Fit Bin Packing

Claire Kenyon*   Yuval Rabani†   Alistair Sinclair‡

1 Introduction

In the one-dimensional bin packing problem, one is given a sequence $a_1, \dots, a_n \in (0, 1]$ of items and asked to pack them into bins of unit capacity in such a way as to minimize the number of bins used. This problem is well known to be NP-hard, and a vast literature has developed around the design and analysis of efficient approximation algorithms for it. The most widely studied among these is the Best Fit algorithm, in which the items are packed on-line, with each successive item going into a partially filled bin with the smallest residual capacity large enough to accommodate it; if no such bin exists, a new bin is started. The performance of this and other approximation algorithms on a given sequence is typically measured by the *waste*, which is the difference between the number of bins used by the algorithm and the sum of the sizes of all the items.

Best Fit was first analyzed in the worst case by Johnson *et al* [8], who proved that the waste is never more than a factor 0.7 of the sum of the item

---

*LIP, URA CNRS 1398, Ecole Normale Supérieure de Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France. E-mail: kenyon@lip.ens-lyon.fr. This work was done while this author was visiting UC Berkeley, partially supported by a NATO Fellowship.

†Department of Computer Science, University of Toronto. Email: rabani@cs.toronto.edu.

‡Computer Science Division, UC Berkeley. Email: sinclair@cs.berkeley.edu. Supported by NSF grant CCR-9505448 and a UC Berkeley Faculty Research Grant.

sizes. When items are drawn from the uniform distribution on $(0, 1]$, the expected waste of Best Fit was shown by Shor [11] to be $\Theta(n^{1/2} \log^{3/4} n)$. Thus among on-line algorithms Best Fit is currently the best available, in the sense that no algorithm is known which beats it both in the worst case and in the uniform average case. This, together with its intuitive appeal and ease of implementation, make it the algorithm of choice in most applications.

With the goal of achieving a better understanding of the Best Fit algorithm, researchers have recently considered its behavior under various other input distributions, notably the class of discrete distributions $U\{j, k\}$ for integers $j < k$. Here the item sizes, instead of being chosen from the continuous real interval $(0, 1]$, are selected uniformly from the finite set of equally spaced values i/k , for $1 \leq i \leq j$. Equivalently, we may think of the bins as having capacity k and the item sizes being uniformly distributed on the integers $\{1, \dots, j\}$. This family of distributions is of interest for two reasons. Firstly, it is an important step towards exploring the robustness of Best Fit under *non-uniform* distributions (because the distribution is biased towards smaller items); and secondly it applies to the more realistic case of discrete rather than continuous item sizes. (For more extensive background, the reader is referred to [3] and the upcoming survey by Johnson.)

Very little is known in rigorous terms about the performance of Best Fit under this distribution, with the exception of a few extreme cases: when $j = k - 1$, the behavior can be readily related to that for the continuous distribution on $(0, 1]$, yielding expected waste $\Theta(n^{1/2})$ [2]; and if j is very small compared to k (specifically, if $j \leq \sqrt{2k} + 2.25 - 1.5$) then the expected waste is known to be bounded by a constant as $n \rightarrow \infty$ [2]. The expected waste is also easily seen to be bounded when $j \leq 2$ for all $k > j + 1$.

Nonetheless, there is much experimental evidence to suggest that the behavior of Best Fit for various pairs (j, k) is complex and interesting. For example, it appears that the waste remains bounded when j is sufficiently close to k or to 1, but that it grows linearly when the ratio j/k is close to a critical value around 0.80. Moreover, in all cases (except $j = k - 1$) where the waste is unbounded it appears to grow linearly with n . Some large scale simulation results, together with some conjectures, are described in [3].

In an attempt to explain this behavior, Coffman *et al* introduced an interesting approach based on a view of the algorithm as a multi-dimensional Markov chain [3]. The states of the chain are positive integer vectors $(s_1, \dots, s_{k-1})$, where s_i represents the current number of open bins of residual

capacity i . Note that such a vector contains all relevant information about the state of the algorithm: in Best Fit, the ordering on the open bins is insignificant, and since we are measuring waste we need not concern ourselves with bins that have already been filled. It is a simple matter to write down the new vector s' that results from the arrival of any item $i \in \{1, \dots, j\}$; since each item arrives with probability $1/j$, this immediately gives the transition probabilities of the chain. (See Section 2.1 below for a more formal definition.) Thus we have a Markov chain on the infinite $(k-1)$ -dimensional space Z_+^{k-1} . The expected waste of Best Fit is intimately related to the asymptotic behavior of this chain.

We note in passing that similar Markov chains have been an object of study in queueing theory for over four decades; in computer science, they have also received attention in the stochastic analysis of packet routing [9]. Despite this extensive body of research few general analytical tools exist, and even the simplest questions, such as showing ergodicity, seem hard to answer. The most notable exception is the method of constructive use of Lyapunov functions, developed in recent years mainly by Malyshev, Menshikov and Fayolle (see [4] for a comprehensive account). The range of situations in which they are able to apply their method appears to be quite limited, however; the highlights are a complete classification of two- and three-dimensional *jump-bounded* Markov chains (i.e., the transitions are limited to geometrically close states). Obviously, the Markov chains that arise in the analysis of Best Fit are jump-bounded, but of much higher dimension.

This was the starting point for Coffman *et al*, who proceeded to analyze the Best Fit Markov chain for small values of j and k , using a novel approach. In the absence of analytical tools for high dimensional Markov chains, they used a computer program to search in an appropriate class of functions for a Lyapunov function (i.e., a potential function obeying certain properties, notably a systematic expected drift over some bounded number of steps). The existence of a suitable Lyapunov function for a given pair (j, k) implies bounded or linear waste. Coffman *et al* were able to classify the waste as bounded or linear for values of k up to 14 and most corresponding values $j < k - 1$.

This approach, while interesting, suffers from several obvious drawbacks, as observed by the authors themselves. Evidently, there is no prospect that this method can lead to proofs for infinite sequences of (j, k) pairs; in fact, the time and space resources consumed by the search make it infeasible to extend

the study beyond a very small finite range of values for j and k . Perhaps most importantly, the technique seems to yield almost no useful insight into *why* the algorithm performs as it does: for example, the Lyapunov function that proves bounded waste for $j = 5, k = 7$ is a linear function based on 23 steps of the Markov chain, while that for $j = 7, k = 10$ is a 15-step quadratic function, neither of which has any intuitive basis.

In this paper, we aim at *analytical* results on the behavior of Best Fit for an infinite sequence of values (j, k) . Specifically, we explore the line $j = k - 2$, the “smallest” interesting case beyond $j = k - 1$, which is the discrete analog of the continuous uniform distribution. Coffman *et al* exhibit computer proofs that the waste is bounded in this case for $k \leq 10$, and also conjecture on the basis of simulations that the waste is bounded for larger values of k . Our main result proves this conjecture for all k :

Theorem 1 *The expected waste of the Best Fit algorithm under the discrete uniform distribution $U\{k - 2, k\}$ is bounded for all k .*

Note the dramatic contrast with the apparently very similar case $j = k - 1$, in which the waste grows unboundedly with n .

Of at least as much interest as this result itself, in our view, are the techniques we use to prove it. Our starting point is again the multi-dimensional Markov chain of Coffman *et al*. However, we develop an alternative view of the chain that seems rather easier to visualize: in this view, the state s of the chain at any time is represented by $k - 1$ *tokens* placed on the non-negative integers, with token i at position s_i . The tokens move around as a dynamical system under the influence of item insertions. With the aid of this view, and the intuition that comes with it, we are able to design an explicit Lyapunov function that proves bounded waste for all pairs (j, k) with $j = k - 2$.

The analysis of the Lyapunov function is somewhat subtle, which perhaps explains why it had not been discovered before. In order to establish the drift in the Lyapunov function, we have to consider $T(j)$ steps of the Markov chain, where $T(j)$ is an exponential function of j ; the drift is proved by a detailed comparison of the Lyapunov function with symmetric random walk on the non-negative integers. More specifically, we are able to relate the projections of the multi-dimensional chain onto the individual coordinates to one-dimensional symmetric random walks that are biased by a limited adversary. This adversary model corresponds to a worst case assumption on the effect of other coordinates, and we believe it to be of independent interest.

It is similar in flavor to, but differs essentially from, the biased random walk model considered by Azar *et al* [1] in a different context. The model in [1] is allowed to bias the transition probabilities slightly on every step, whereas our adversary may intervene overwhelmingly but only on a limited number of steps. The techniques required to analyze the two models seem to be very different.

In addition to settling an open problem posed in [3], our result, more significantly, is the first proof that exploits the detailed structure of the multi-dimensional Markov chain, and thus the first that provides an understanding of its behavior. We are optimistic that our techniques can be extended to analyze the Best Fit Markov chain for other pairs of values (j, k) , and perhaps also to other situations in the analysis of algorithms in which homogeneous multi-dimensional Markov chains of this kind arise.

The remainder of this paper is structured as follows. In Section 2 we introduce the token model as a convenient representation of the Markov chain underlying the algorithm, and establish various fundamental properties of it. In Section 3 we construct our Lyapunov function and analyze its behavior using comparisons with symmetric random walks.

2 The token model

2.1 Definitions

As advertised in the Introduction, we describe the behavior of the Best Fit algorithm over time in terms of the evolution of a dynamical system. In this system, $k - 1$ *tokens* move among the non-negative integer points under the influence of item insertions, as follows. The tokens are labeled $1, 2, \dots, k - 1$. At any time instant t , the position of token i is the number of open bins at time t with residual capacity exactly i . We shall denote the state of the system at time t by $s(t) = (s_1(t), \dots, s_{k-1}(t))$, a vector random variable taking values in Z_+^{k-1} . Initially, the state of the system is $s(0) = (0, \dots, 0)$, reflecting the fact that there are no open bins.

Now suppose the state of the system at time t is $s(t)$ and the next item to be inserted is ℓ , where $1 \leq \ell \leq j$. Let i be the smallest index such that $i \geq \ell$ and $s_i(t) > 0$, if such exists: in this case, the algorithm inserts item ℓ into a bin with capacity i , so we have $s_i(t + 1) = s_i(t) - 1$ and, if $i > \ell$,

$s_{i-\ell}(t+1) = s_{i-\ell}(t) + 1$; all other components of $s(t)$ are unchanged. If no such i exists, then the algorithm inserts item ℓ into an empty bin, so we have $s_{k-\ell}(t+1) = s_{k-\ell}(t) + 1$ and all other components of s are unchanged. This completes the description of the dynamical system.

Note that the above system is nothing other than a convenient pictorial representation of a multi-dimensional Markov chain, with state space Z_+^{k-1} , in which token i executes a random walk in dimension i . The motions of individual tokens are, of course, not independent. However, the transition probabilities of any given token at any time depend only on which of the tokens are at zero at that time, i.e., on the set $\{i : s_i = 0\}$. This is an important property which makes analysis of the chain feasible.

Our goal will be to investigate the behavior of the *waste* in the algorithm after packing t items, which in this model is defined as the number of (non-full) open bins, i.e., $\sum_{i=1}^{k-1} s_i(t)$. In particular, we will be concerned with determining, for particular pairs (j, k) , whether or not the waste remains bounded for an infinite stream of items, i.e., as $t \rightarrow \infty$.

2.2 Classification of tokens

It will be convenient for us to partition the tokens into two classes, which we will call “large” and “small.” This idea is motivated by the fact that tokens behave in two distinct ways, as we shall see in a moment. The *small* tokens are tokens i with $1 \leq i \leq \lceil \frac{j}{2} \rceil$. The *large* tokens are tokens $k - i$ with $1 \leq i \leq \lceil \frac{j}{2} \rceil$. Note that the numbers of small and large tokens are equal. In the case that j is even there is an additional token, namely $\lceil \frac{j}{2} \rceil + 1$, which is neither small nor large: we call this the *middle* token.

We first establish a fundamental constraint on the states that are reachable from the initial state $s(0)$. This fact is implicit in [3]; the proof is a straightforward induction on time which we omit from this abstract. The reader may enjoy figuring it out.

Proposition 2 *State s is reachable from the initial state $s(0)$ only if*

1. $i + i' \geq k \implies s_i = 0$ or $s_{i'} = 0$. (I.e., no two tokens whose index sum is k or greater can simultaneously be at non-zero positions.)
2. $\sum_{i \text{ not small}} s_i \leq 1$. (I.e., the large and middle tokens cannot move beyond

position 1; moreover, at most one of them can be away from zero at any time.) $\square$

It is not hard to see that all states satisfying the conditions of Proposition 2 are in fact reachable from the initial state. From now on, we shall therefore assume that the state space of our Markov chain is precisely this set S of reachable states.

The above proposition expresses general constraints on the motions of the tokens. In the following three subsections, we establish further properties of the behavior of tokens under certain assumptions about the distribution of other tokens. These properties will be used in our analysis in the next section.

2.3 Behavior of large and middle tokens

We have already seen that the large and middle tokens behave in an extremely restricted fashion. Their behavior becomes even more restricted under the condition that $s_{\lceil j/2 \rceil} > 0$. This condition will arise naturally in our analysis in the next section.

Proposition 3 *Suppose that $s_{\lceil j/2 \rceil}$ remains strictly positive throughout some time interval. Then during this interval:*

- *all large tokens remain at zero;*
- *the middle token (if it exists) oscillates between 0 and 1 independently of the positions of all other tokens.*

Proof. The first claim is immediate from condition 1 of Proposition 2. To see the second claim, note that, because all larger tokens are at zero, insertions of item $\lceil \frac{j}{2} \rceil + 1$ are placed alternately in an empty bin (thus creating a bin with capacity $\lceil \frac{j}{2} \rceil + 1$) and in this newly opened bin. No other insertions can affect token $\lceil \frac{j}{2} \rceil + 1$, which therefore oscillates between 0 and 1 as claimed. $\square$

In view of the second claim of Proposition 3, we will assume from now on that j is odd, so that there is no middle token to worry about. This assumption is justified because our analysis will hinge on the behavior of the system when $s_{\lceil j/2 \rceil} > 0$; but Proposition 3 then tells us that the behavior of

the middle token under this condition is degenerate. With this observation, the argument we will give for j odd trivially extends to the case when j is even.

2.4 Behavior of small tokens

Most of this paper is concerned with the detailed behavior of the small tokens: since the other tokens remain very severely bounded, it is really only the small tokens that are interesting from the point of view of the asymptotic behavior of the algorithm. In the next proposition, we isolate an essential feature of the motion of the small tokens under a certain condition that will again arise naturally from our analysis in the next section.

Proposition 4 *Let i be a small token. If s_{i-1} is strictly positive throughout some time interval, then during this interval the motion of token i has the following properties:*

Property A *When s_i is not at zero, it executes random walk on the positive integers with non-negative drift and holding probability at most $1 - \frac{2}{j}$.*

Property B *The time spent by s_i on each visit to zero is stochastically dominated by a random variable D with constant expectation (that depends only on j).*

Proof. Consider first the case when $s_i > 0$. Since $s_{i-1} > 0$, the only way in which s_i can decrease is through the insertion of item i . On the other hand, s_i will certainly increase on insertion of item $k - i$; to see this, note from condition 1 of Proposition 2 that $s_{i'} = 0$ for all $i' \geq k - i$, so the algorithm must insert item $k - i$ into an empty bin. Hence s_i decreases with probability $\frac{1}{j}$ and increases with probability at least $\frac{1}{j}$, which is exactly equivalent to Property A.

Now consider what happens when $s_i = 0$. If $s_{i'} = 0$ for all $i' \geq k - i$, then as above we can conclude that s_i moves to 1 with probability at least $\frac{1}{j}$. However, now we cannot exclude the possibility that $s_{k-i} = 1$, in which case item $k - i$ will be inserted into the bin with capacity $k - i$ so s_i cannot leave 0. On the other hand, in this situation we see that two consecutive insertions of item $k - i$ will certainly have the effect of moving s_i to 1. This crude argument indicates that the time spent by token i at 0 is stochastically

dominated by the random variable D defined as follows over the sequence of item insertions immediately following the arrival of i at 0:

$$D = \begin{cases} 1 & \text{if first insertion is } k - i; \\ N & \text{otherwise,} \end{cases}$$

where N is the numbers of insertions until the first pair of consecutive insertions of $k - i$ has occurred. Notice that the events that item $k - i$ is inserted at time t are mutually independent for all t , and all have probability $\frac{1}{j}$. Hence it is easy to see that the tail of D has the form $\Pr[D > n] \leq \beta^n$ for some constant $\beta > 1$ that depends only on j . This in turn implies that the expectation of D is bounded above by a constant that depends only on j .

□

2.5 Behavior of the waste

In this subsection, we investigate what happens to the waste in the system, again under the assumption that $s_{\lceil j/2 \rceil} > 0$. Define $f(t) = \sum_{i=1}^{\lceil j/2 \rceil} i s_i(t)$, which is essentially just the waste due to the small tokens at time t (weighted by coefficients in a bounded range). By Proposition 2, the total waste is bounded above by $f(t) + 1$. The following proposition shows that $f(t)$ has negative drift under our assumption about $s_{\lceil j/2 \rceil}$.

Proposition 5 *Suppose that $s_{\lceil j/2 \rceil}(t) > 0$. Then $E[f(t+1) - f(t) \mid f(t)] = -1/j$.*

Proof. For all $i = 1, 2, \dots, \lceil \frac{j}{2} \rceil$, define $f_i(t) = i s_i(t) + (i-1) s_{i-1}(t) + \dots + s_1(t)$. Thus, $f(t) = f_{\lceil j/2 \rceil}(t)$. For all $i = 1, 2, \dots, \lceil \frac{j}{2} \rceil$, define the set of i -requests to be $\{1, 2, \dots, i\} \cup \{k-i, k-i+1, \dots, k-2\}$. This is the set of all items that can potentially be absorbed by a bin with capacity i' , $1 \leq i' \leq i$, or can create a new bin with such a capacity from an empty bin. We will prove:

Claim *Let $i \in \{1, 2, \dots, \lceil \frac{j}{2} \rceil\}$. If $s_i(t) > 0$ then the expected change in f_i due to i -requests is $-1/j$.*

The Proposition follows from this Claim with $i = \lceil \frac{j}{2} \rceil$, since the set of $\lceil \frac{j}{2} \rceil$ -requests is the entire set of items. To prove the Claim, we proceed by induction on i :

Basis: $i = 1$. $f_1(t) = s_1(t)$. The only 1-request is 1. Its contribution to the expected change in f_1 is $-1/j$, since it arrives with probability $1/j$ and causes s_1 to decrease by 1.

Inductive step: Assume the Claim holds for i . We show that it holds for $i + 1$. Let $i' < i + 1$ be the largest index such that $s_{i'}(t) > 0$; set $i' = 0$ if all these tokens are at 0. We consider two cases:

Case 1: $i' > 0$. By the inductive hypothesis, the i' -requests cause an expected change of $-1/j$ in $f_{i'}$. Consider the remaining items in the set of $(i + 1)$ -requests in pairs $i'', k - i'', i' < i'' \leq i + 1$. For $i'' = i + 1$, both items affect only s_{i+1} and they cancel each other's contribution. For $i'' < i + 1$, item i'' reduces s_{i+1} by 1 and increases $s_{i-i''+1}$ by 1, so its contribution is $(-(i + 1) + (i - i'' + 1))/j = -i''/j$. On the other hand, item $k - i''$ causes a new bin of capacity i'' to be created (since $s_{i+1}(t) > 0$, there are no open bins of capacity $\geq k - i - 1$), so its contribution is $+i''/j$, and the total contribution of the pair is again 0.

Case 2: $i' = 0$. Pair the requests as before for $1 < i'' \leq i + 1$: again, each pair's contribution is 0. The remaining item, of size 1, causes a decrease of 1 in s_{i+1} and an increase of 1 in s_i . Its contribution to f_{i+1} is therefore $(-(i + 1) + i)/j = -1/j$.

This completes the proof of the Claim, and hence of the Proposition.

□

3 Analysis of the Markov chain

This section is devoted to proving our main result, Theorem 1 stated in the Introduction. Our proof makes use of the following result of [4], which establishes a general condition, in terms of the existence of a suitable Lyapunov function, for a multi-dimensional Markov chain to be ergodic. For more specialized variations on this theme, see [6, 10, 7, 3].

Lemma 6 [4, Corollary 7.1.3] *Let $\mathcal{M}$ be a Markov chain with state space $S \subseteq \mathbb{Z}^k$, and b a positive integer. Denote by $p_{ss'}^b$ the transition probability from s to s' in $\mathcal{M}^b$, the b -step version of $\mathcal{M}$. Let $\Phi : S \rightarrow \mathbb{R}_+$ be a non-negative real-valued function on S , which satisfies the following conditions:*

1. *There are $C_1, \mu > 0$ such that $\Phi(s) > C_1 \|s\|^\mu$ for all $s \in S$.*

2. There is $C_2 > 0$ such that $p_{ss'}^b = 0$ whenever $|\Phi(s) - \Phi(s')| > C_2$, for all $s, s' \in S$.
3. There is a finite $B \subset S$ and $\epsilon > 0$ such that $\sum_{s' \in S} p_{ss'}^b (\Phi(s') - \Phi(s)) < -\epsilon$ for all $s \in S \setminus B$.

Then $\mathcal{M}$ is ergodic with stationary distribution π satisfying $\pi(s) < Ce^{-\delta\Phi(s)}$ for all $s \in S$, where C and δ are positive constants. $\square$

To interpret this lemma, view Φ as a potential function that maps the state space to the non-negative reals, so that the image of the Markov chain under Φ becomes a dynamical system on the real line. Condition 2 requires this process to be well-behaved, in the sense that it has bounded variation. The key is condition 3, which says that, except for a finite set of states, Φ has negative drift over an interval of some constant length b . This implies that $\mathcal{M}$ is ergodic with a stationary distribution that decays exponentially with Φ .

In our application, $\mathcal{M}$ will be the Markov chain that governs the movements of the tokens, whose state space is the subset S of Z_+^{k-1} defined by Proposition 2, and Φ will be the function $\Phi(s) = 1 + \sum_{i=1}^{\lfloor j/2 \rfloor} is_i$. (Note that Φ is essentially just the function f of Proposition 5.) It is clear that conditions 1 and 2 hold for this Φ , with any choice of constant b . All our work will be devoted to proving the negative drift condition 3, for suitably chosen b , B and ϵ . Note that Theorem 1 will then follow immediately, since the asymptotic waste is bounded above by $\sum_{s \in S} \pi(s)\Phi(s)$, which by Lemma 6 is bounded.

The following is an informal sketch of our strategy for proving condition 3:

- (i) We consider an interval of length b , and show that Φ has negative drift over this interval provided it is large enough at the start of the interval: i.e., we will take B to be the finite set of points on which Φ is “small.” Thus for $s \in S \setminus B$, we can be sure that, for some small token i , s_i is large at the start of the interval, and hence positive throughout the interval.
- (ii) Since $s_i > 0$ throughout the interval, by Proposition 4 the motion of s_{i+1} is a symmetric random walk: hence the time that s_{i+1} spends at 0 during the interval is small (about $\text{const} \times \sqrt{b}$).

- (iii) Iterating this argument, appealing to Proposition 4 each time, we can conclude that each of the tokens $s_{i+2}, \dots, s_{\lceil j/2 \rceil}$ spends little time at 0.
- (iv) Finally, since we have established that $s_{\lceil j/2 \rceil} > 0$ during most of the interval, Proposition 5 tells us that f (and hence Φ) has negative drift on most steps, and hence an overall negative drift over the entire interval.

The tricky part of the above argument is step (iii): at each stage we need to use the fact that $s_{i'} > 0$ to deduce from Proposition 4 that $s_{i'+1}$ behaves like a symmetric random walk. However, occasionally $s_{i'}$ will be at 0, and at these times we have no control over the motion of $s_{i'+1}$. We therefore assume that $s_{i'+1}$ behaves like a symmetric random walk most of the time, but that an adversary is able to control its motion on a small number of steps. Accordingly, we need to prove a lemma that quantifies the effect that such an adversary can have on the amount of time $s_{i'+1}$ spends at 0. This we now do.

Consider a symmetric random walk on $[0, \infty)$ of a given length, started at some specified position, and an adversary whose goal is to maximize the number of times the walk hits 0.¹ The adversary is allowed to intervene at some specified number of steps, selected according to any strategy: on these steps, the adversary may specify any desired probability distribution on the legal moves of the process from the current state; on all other steps, the process behaves as a symmetric random walk with a perfectly reflecting barrier at 0. It is perhaps not surprising that the optimal strategy for the adversary is always to intervene by driving the process deterministically towards the origin, and to use up all these interventions as early as possible. However, this claim requires some justification, which we now provide.

Lemma 7 *Let $p(i, n, y, m)$ be the probability that a symmetric random walk of n steps, starting at i and with y adversary steps, hits the origin at least m times. Let $q(i, n, y, m)$ be the same quantity for the particular adversary strategy in which downward steps are used as early as possible. Then $p(i, n, y, m) \leq q(i, n, y, m)$ for all i, n, y, m .*

To prove Lemma 7, we need a simple technical observation about symmetric random walk.

---

¹Throughout, for convenience, we shall take “hits 0” to mean “makes the transition $0 \rightarrow 1$.”

Proposition 8 *Let $p_0(i, n, m) = p(i, n, 0, m)$ denote the probability that an unbiased random walk of length n started at i hits 0 at least m times. Then $p_0(2, n, m - 1) \geq p_0(1, n + 1, m)$.*

Proof. Let W_2 be the random walk started from 2 and W_1 the random walk started from 1. Consider the first time T when W_2 reaches 1, and the first time T' when W_1 reaches 1 by a $0 \rightarrow 1$ transition. Then it is easy to see that T' is equal to $T + 1$ in distribution. The remainder of W_2 after T is a random walk started at 1 which must have at least $m - 1$ hits. The remainder of W_1 after T' is also a random walk started at 1 which must have at least $m - 1$ hits. Thus if $m > 1$, the probabilities are the same for both walks (conditioning on the event $T = T' - 1 \leq n$), and if $m = 1$, we trivially have $p_0(2, n, 0) = 1 \geq p_0(1, n + 1, 1)$. $\square$

We now prove Lemma 7. We will consider only deterministic strategies: the randomized case will follow by averaging.

Proof of Lemma 7. We use induction on n . Let $D^y R$ denote the strategy which uses the y forced down steps as soon as possible, and then follows the unbiased random walk. Let RD^y denote the strategy which starts with a truly random step as soon as possible, and then uses the y forced down steps as soon as possible. Notice that a transition from 0 to 1 is neither a forced down step (obviously) nor a truly random step, since it has probability 1. Let $q(i, n, y, m)$ denote the probability that the walk of length n defined by the strategy $D^y R$, started at i and with y forced down steps has at least m hits; let $r(i, n, y, m)$ be defined similarly for the strategy RD^y . We claim that

$$q(i, n, y, m) \geq r(i, n, y, m) \tag{1}$$

Note that the lemma will then follow by induction on n : consider the first time the adversary may intervene. Either way, after this step we are left to deal with fewer than n steps. If the adversary does force a down step, by induction the best strategy to continue is $D^{y-1} R$, so the strategy for the entire walk is $D^y R$. If the adversary does not intervene, using induction again, the best strategy to continue is $D^y R$, so the strategy for the entire walk is $RD^y R = RD^y$. Inequality (1) shows that the strategy $D^y R$ is better than RD^y .

To prove (1), we also use induction on n . If $i = 0$, both strategies start the same way and we are done by induction. If $i \geq y + 1$, both strategies

give the same distribution of positions after $y + 1$ steps, and neither has hit 0 yet, so the two quantities are equal. The interesting case is for $1 \leq i \leq y$. Then, let $n' = n - (2y - i + 2)$ and $m' = m - (y - i + 1)$. It is easy to see that

$$\begin{aligned} q(i, n, y, m) &= \frac{1}{2}p_0(0, n', m') + \frac{1}{2}p_0(2, n', m'); & \text{and} \\ r(i, n, y, m) &= \frac{1}{2}p_0(0, n', m') + \frac{1}{4}p_0(0, n', m' + 1) + \frac{1}{4}p_0(2, n', m' + 1). \end{aligned}$$

Since $p_0(1, n' + 1, m' + 1) = \frac{1}{2}(p_0(0, n', m' + 1) + p_0(2, n', m' + 1))$, the difference is

$$q - r = \frac{1}{2}(p_0(2, n', m') - p_0(1, n' + 1, m' + 1)),$$

which is non-negative by Proposition 8. $\square$

In order to relate the above adversary result to our process, we need an elementary technical fact. Consider an arbitrary stochastic process over the non-negative integers. Assume it has arbitrary holding probabilities except at 0, where the holding probability is 0, and non-negative drift everywhere. Let Z denote the number of hits on 0 during the first T steps of this process. Let U be the similar quantity for the symmetric random walk with perfectly reflecting barrier at 0, starting at the same point. The following fact can be proved by a simple coupling argument, which we omit from this abstract.

Proposition 9 *Z is stochastically dominated by U .* $\square$

We are now in a position to proceed with the proof of our main result, following the sketch given after Lemma 6. As observed there, the main difficulty lies in step (iii): assuming that $s_i > 0$ throughout some interval, we want to conclude that $s_{\lceil j/2 \rceil} > 0$ during most of that interval. This is the subject of our next lemma, which makes essential use of the above adversary lemma.

Lemma 10 *Let T and a be positive constants, and suppose that $s_i(0) > T$. With probability at least $1 - C/a$, where C is a constant that depends only on j , $s_{\lceil j/2 \rceil}(t)$ is strictly positive at all but $a\sqrt{T}$ time instants t within the interval $[0, T]$.*

Proof. We will prove the claim for $i = 1$. The proof for general i is exactly the same. So, assume that $s_1(0) > T$. For each i , let the random variable T_i

denote the time spent at zero by token i during the interval $[0, T]$. Clearly $T_1 = 0$ with probability 1.

Next let us consider the behavior of the sequence s_2 . Consider a modified process s'_2 which is defined as follows. First, run s_2 for T steps. Then, have the token s'_2 follow an unbiased random walk with a holding time at zero distributed according to D . Finally, delete from s'_2 all stationary steps at 0. Let T'_2 be the number of hits on 0 of s'_2 during the time interval $[0, T]$. Then we have

$$T_2 \leq \sum_{r=1}^{T'_2+1} D_r, \quad (2)$$

where the D_r are i.i.d. with the same distribution as D , and T'_2 is independent of all of the D_r .

To analyze T'_2 , we compare it with U , the number of hits on 0 of symmetric random walk with perfectly reflecting barrier at 0. By Proposition 9, T'_2 is stochastically dominated by U . Taking expectations in (2) and using this observation, we get

$$\begin{aligned} \mathbb{E}T_2 &\leq \mathbb{E} \left[\sum_{r=1}^{T'_2+1} D_r \right] \\ &= \sum_{t=0}^{\infty} \sum_{r=1}^{t+1} \mathbb{E}[D_r \mid T'_2 = t] \Pr[T'_2 = t] \leq (\mathbb{E}U + 1)\mathbb{E}D \\ &= d(\mathbb{E}U + 1), \end{aligned}$$

where the constant d is the expectation of D from Proposition 4.

Now consider token s_ℓ , where $3 \leq \ell \leq \lceil \frac{j}{2} \rceil$. Define a modified process s'_ℓ and a random variable T'_ℓ similar to the fashion that s'_2 and T'_2 were defined. By analogy with (2) we may write

$$T_\ell \leq \sum_{r=1}^{T'_\ell+1} D_r. \quad (3)$$

Now our adversary argument, Lemma 7, implies that T'_ℓ is stochastically dominated by $U + T_{\ell-1}$. To see this, note that

$$\Pr[T'_\ell \geq m \mid T_{\ell-1} = y] \leq p(s_\ell(0), T, y, m) \leq q(s_\ell(0), T, y, m) \leq \Pr[U \geq m - y].$$

Taking expectations in equation (3), and using this fact, we get

$$ET_\ell \leq (EU + ET_{\ell-1} + 1)ED = d(EU + ET_{\ell-1} + 1).$$

Iterating this bound, and using the base case (3), gives

$$ET_\ell \leq \left(\sum_{r=1}^{\ell-2} d^r\right)(EU + 1) + d^{\ell-2}ET_2 \leq \ell d^\ell (EU + 1).$$

But $EU \leq c\sqrt{T}$ for some universal constant c . Hence by Markov's inequality

$$\Pr[T_{\lceil j/2 \rceil} > a\sqrt{T}] < \frac{\lceil j/2 \rceil d^{\lceil \frac{j}{2} \rceil - 1} (c + 1/\sqrt{T})}{a},$$

which is bounded above by C/a for some constant C as required. $\square$

Remarks: (a) The above proof actually demonstrates the stronger conclusion that $s_{i'}(t) > 0$ for all i' in the range $i \leq i' \leq \lceil \frac{j}{2} \rceil$, for a similar majority of the interval.

(b) It is interesting to note that the only properties of the sequence $s(t)$ we have used in the above proof are properties A and B of Proposition 4. Thus Lemma 10 actually applies to any sequence of vector random variables satisfying these rather natural properties. We believe that this fact may be of independent interest.

We are finally in a position to complete the proof of Theorem 1, following our earlier sketch.

Proof of Theorem 1. Recall from the discussion immediately following the statement of the theorem that it suffices to show that condition 3 of Lemma 6 holds for the function $\Phi(s) = 1 + \sum_{i=1}^{\lceil j/2 \rceil} is_i$, with suitable choices of b , B and ϵ . The set $B \subseteq S$ will be defined as

$$B = \{s \in S : f(s) \leq T + 1\},$$

where T is some constant to be specified shortly.

Assume that $s = s(0) \in S \setminus B$, i.e., that $\Phi(s(0)) > T + 1$. Define $b = T/j^2$, and consider the time interval $[0, b]$. We will show that the expected drift of Φ over this interval is less than $-\epsilon$ for some $\epsilon > 0$, thus establishing condition 3.

Actually we will work with the function f of Proposition 5, which is just one less than Φ .

The assumption that $\Phi(s(0)) > T + 1$ implies that there is a token s_i , with $i \leq \lceil \frac{j}{2} \rceil$, such that $s_i(0) > T/j^2$. Then by Lemma 10, with probability at least $1 - C/a$ we have $s_{\lceil j/2 \rceil} > 0$ at all but at most $a\sqrt{T/j^2}$ time instants in the time interval $[0, T/j^2]$.

Now consider the change Δf in f after one step. By Proposition 5, if $s_{\lceil j/2 \rceil} > 0$ then $E\Delta f \leq -1/j$. In all other situations, there is the trivial bound $E\Delta f \leq j$. Putting these facts together, and conditioning on the event $\mathcal{A}$ that $s_{\lceil j/2 \rceil} > 0$ at all but at most $a\sqrt{T/j^2}$ time instants in the interval, we see that the drift of f over the entire interval $[0, T/j^2]$ is

$$\begin{aligned} E[f(b) - f(0) \mid f(0)] &\leq \\ &\leq E[f(b) - f(0) \mid f(0) \wedge \mathcal{A}] + (1 - \Pr[\mathcal{A}]) E[f(b) - f(0) \mid f(0) \wedge \neg \mathcal{A}] \\ &\leq \left\{ -\frac{1}{j} \left(\frac{T}{j^2} - a\sqrt{\frac{T}{j^2}} \right) + a\sqrt{T} \right\} + \frac{C}{a} \frac{T}{j}. \end{aligned}$$

By taking a large enough, and then T large enough, we can make this expression less than some negative constant ϵ .

This completes the verification of condition 3, and hence the proof of the theorem. $\square$

References

- [1] Y. Azar, A.Z. Broder, A.R. Karlin, N. Linial, and S. Phillips. Biased random walks. In *Proc. 24th Ann. ACM Symp. on Theory of Computing*, 1992, pp. 1–9.
- [2] E.G. Coffman, Jr, C. Courcoubetis, M.R. Garey, D.S. Johnson, L.A. McGeoch, P.W. Shor, R.R. Weber and M. Yannakakis. Fundamental discrepancies between average-case analyses under discrete and continuous distributions: A bin-packing case study. In *Proc. 23rd Ann. ACM Symp. on Theory of Computing*, 1991, pp. 230–240.
- [3] E.G. Coffman Jr., D.S. Johnson, P.W. Shor, and R.R. Weber. Markov chains, computer proofs, and average-case analysis of Best Fit bin packing. In *Proc. 25th Ann. ACM Symp. on Theory of Computing*, 1993, pp. 412–421.
- [4] G. Fayolle, V.A. Malyshev, and M.V. Menshikov. *Constructive Theory of Countable Markov Chains (Part I)*. Manuscript, May 1992.
- [5] W. Feller. *An Introduction to Probability Theory and its Applications, Volume I (3rd ed.)* John Wiley, New York, 1968.
- [6] F.G. Foster. On stochastic matrices associated with certain queueing processes. *Ann. Math. Stat.*, **24**(1953), pp. 355–360.
- [7] B. Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Adv. Appl. Prob.*, **14**(1982), pp. 502–525.
- [8] D.S. Johnson, A. Demers, J.D. Ullman, M.R. Garey and R.L. Graham. Worst case performance bounds for simple one-dimensional packing algorithms. *SIAM J. Comput.*, **3**(1974), pp. 299–325.
- [9] F.T. Leighton. Average case analysis of greedy routing algorithms on arrays. In *Proc. 2nd Ann. ACM Symp. on Parallel Algorithms and Architectures*, 1990, pp. 2–10.
- [10] V.A. Malyshev and M.V. Menshikov. Ergodicity, continuity, and analyticity of countable Markov chains. *Trans. Moscow Math. Soc.*, **39**(1979), pp. 3–48.

- [11] P.W. Shor. The average case analysis of some on-line algorithms for binpacking. *Combinatorica*, **6**(1986), pp. 179–200.