

NECESSARY AND SUFFICIENT CONDITIONS FOR STABILITY OF A BIN-PACKING SYSTEM

C. COURCOUBETIS,* *AT&T Bell Laboratories*
R. R. WEBER,** *University of Cambridge*

Abstract

Objects of various integer sizes, o_1, \dots, o_n , are to be packed together into bins of size N as they arrive at a service facility. The number of objects of size o_i which arrive by time t is A_i^t , where the components of $A^t = (A_1^t, \dots, A_n^t)$ are independent renewal processes, with $A^t/t \rightarrow \lambda$ as $t \rightarrow \infty$. The empty space in those bins which are neither empty nor full at time t is called the *wasted space* and the system is declared *stabilizable* if for some finite B there exists a bin-packing algorithm whose use guarantees the expected wasted space is less than B for all t . We show that the system is stabilizable if the arrival processes are Poisson and λ lies in the interior of a certain convex polyhedral cone Λ . In this case there exists a bin-packing algorithm which stabilizes the system without needing to know λ . However, if λ lies on the boundary of Λ the wasted space grows as $O(\sqrt{t})$ and if λ is exterior to Λ it grows as $O(t)$; these conclusions hold even if objects may be repacked as often as desired.

ONLINE ALGORITHMS; RANDOM WALKS; STOCHASTIC ALGORITHMS

1. Problem description and results

In a one-dimensional bin-packing model we suppose that objects of various integer sizes in the set $O = \{o_1, \dots, o_n\}$ are to be packed together into bins of integer size N as they arrive at a service facility. The number of objects of size o_i which arrive by time t is A_i^t , where the components of $A^t = (A_1^t, \dots, A_n^t)$ are independent renewal processes, with $A^t/t \rightarrow \lambda = (\lambda_1, \dots, \lambda_n)$ as $t \rightarrow \infty$. The empty space in those bins which are neither empty nor full at time t is called the *wasted space* and this is denoted by W^t . The system is declared *stabilizable* if for some finite B there exists a bin-packing algorithm whose use guarantees that the expected wasted space is less than B for all t . Of course this is equivalent to the notion that the expected number of partially full bins remains uniformly bounded as time increases.

Received 26 July 1985; revision received 29 October 1985.

* Postal address: AT&T Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974, USA.

** Postal address: Queens' College, Cambridge CB3 9ET, UK.

We are interested in characterizing the parameter values N , O and A' for which the system is stabilizable. It is assumed that the bin-packing must be online: that is, the objects must be placed in bins as they arrive and may not be repacked later. The wasted space can certainly be less if objects may be repacked as often as desired, but even if repacking is allowed the system may not be stabilizable (for example, if $O = \{2\}$ and $N = 3$).

Coffman et al. (1984) survey over 100 articles on bin-packing models. Almost all of these concern the worst-case analysis of suboptimal algorithms for deterministic scenarios. Some authors have studied stochastic bin-packing models, with both online and offline algorithms. Their principal interest has been the average performance of classical algorithms such as 'first fit' or 'best fit' when the object sizes are drawn from a known distribution (see Bentley et al. (1984)). Hoffmann (1982) has studied the design of online algorithms which employ knowledge about the distribution from which object sizes are drawn. He has shown that there is an online algorithm making the expected value of the ratio of the number of used bins to the sum of the object sizes as close to 1 as desired as $t \rightarrow \infty$.

We ask whether it is possible for the expected wasted space to remain uniformly bounded as time increases. This would be desirable in a cutting-stock context, where as orders arrive they are cut from rolls of stock which are manufactured in equal lengths N , and operating costs are proportional to the amount of stock which has been manufactured but which is as yet only partially cut. We begin our discussion by defining the set of *configurations* $C = \{c_1, \dots, c_m\}$ as the set of all distinct ways in which a bin of size N can be packed completely with objects from the set O . Configuration c_j is given by the vector (c_{j1}, \dots, c_{jm}) , where $c_{jk} \in O$ denotes the size of object to occupy the k th slot of the configuration, and where clearly $\sum_k c_{jk} = N$. In order that the c_j 's have a unique designation, without duplication of equivalent packings, we always order the slots so that $c_{j1} \leq \dots \leq c_{jm}$.

Crucial to our analysis is the $n \times m$ -dimensional *packing matrix* Q , of which the element q_{ij} is the number of objects of size o_i used in packing configuration c_j . To avoid triviality we suppose that each row of Q is non-zero, so that each object in O can be packed with some other objects in O to completely fill a bin of size N . Another important entity is the set, $\Lambda = \{\lambda : \lambda = Qf, f \geq 0\}$, spanned by non-negative linear combinations of the columns of Q . Let $\text{bdy}(\Lambda)$, $\text{ext}(\Lambda)$ and $\text{int}(\Lambda)$ denote the boundary, exterior and interior of Λ respectively. Notice that Λ is a closed convex polyhedral cone in n -dimensional space and $\text{int}(\Lambda)$ is non-empty if and only if the null space $N(Q) = \{x : Q'x = 0\} = \{0\}$. To illustrate the importance of Λ we describe a simple example.

Example. Suppose objects of sizes 1 and 2 are to be packed in bins of size 3: $O = \{1, 2\}$, $N = 3$. There are two configurations, with

$$Q = \begin{pmatrix} 3 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \Lambda = \{\lambda : \lambda_1 \geq \lambda_2\}.$$

For reasons which we shall explain following the statement of Theorem 4, we restrict attention to $\lambda > 0$. If $\lambda \in \text{ext}(\Lambda)$ then $\lambda_2 > \lambda_1$ and it is clear that even if repacking is allowed the number of bins which are only partially filled by a single 2 will grow as $O(t)$. If $\lambda \in \text{bdy}(\Lambda)$ then $\lambda_1 = \lambda_2$ and each 1 must be matched with a 2 if the wasted space is not to grow as fast as $O(t)$. If we try to do this, either all partially full bins at time t will contain a single 2, or all will contain a single 1. The number of partially full bins will perform a symmetric random walk on the non-negative integers with reflection at 0, and hence the expected number of such bins will grow as $O(\sqrt{t})$. If $\lambda \in \text{int}(\Lambda)$ then $\lambda_1 > \lambda_2$. Suppose we use an online algorithm which, if there are partially full bins, always puts a 1 in a partially full bin with a 2, or if there is no such bin puts it in a partially full bin with other 1's. The number of partially full bins containing a 2 performs a random walk on the non-negative integers, with drift towards 0 and reflection at 0. The number of partially full bins containing 1's is never more than 1. Thus the expected wasted space is bounded by some B for all t .

The observations made for the above example can be generalized. We show that a system is stabilizable only if $\lambda \in \text{int}(\Lambda)$. When the objects arrive according to independent Poisson processes this condition is sufficient for stability. We describe a probabilistic bin-packing algorithm which stabilizes the system in this case and which does not need to know λ . When λ is not in the interior of Λ the system is not stabilizable even if repacking is allowed. If $\lambda \in \text{bdy}(\Lambda)$ the wasted space grows as $O(\sqrt{t})$, and if $\lambda \in \text{ext}(\Lambda)$ it grows as $O(t)$. These results are stated in the following theorems.

Theorem 1. If $\lambda > 0$ is exterior to Λ , then $W' \sim O(t)$, even when repacking is allowed.

Theorem 2. If $\lambda > 0$ is on the boundary of Λ and the interarrival times of the renewal processes A_i^j have non-zero variance, then for every algorithm there is a $B > 0$ such that $E[W'] > B\sqrt{t}$, even when repacking is allowed. If the interarrival times for each object size are constant, then there exists an online algorithm and $B < \infty$ such that $W' < B$ for all t .

Theorem 3. If $\lambda > 0$ is in the interior of Λ , and objects may be repacked as desired, then there exists an algorithm, a bound $B < \infty$ and an almost surely finite random time τ such that $W' < B$ for all $t > \tau$.

Theorem 4. If $\lambda > 0$ is in the interior of Λ and the arrival processes are Poisson then the system is stabilizable by an online algorithm which does not need to know λ .

Throughout the above we restrict attention to $\lambda > 0$. If the i th component of λ is 0 and other components are positive then the dimension of the system should be reduced and the theorems should be read for a Λ^* embedded in $n - 1$ -dimensional space, where $O^* = O - \{o_i\}$, $\lambda^* > 0$, and Q^* is derived from Q by deleting the i th row of Q and all columns of Q which have a positive entry in the i th row (thereby deleting all configurations which cannot be used). Thus, in the above example, we have $\lambda = (1, 0)' \in \text{bdy}(\Lambda)$ if Λ is viewed as being embedded in R^2 . However, for the set $\Lambda^* = \{\lambda_1: \lambda_1 > 0\}$, which is embedded in R^1 , we do have $(1) \in \text{int}(\Lambda^*)$. It is therefore Theorems 3 and 4 which apply when $\lambda = (1, 0)'$, rather than Theorem 2. For $\lambda = (0, 1)'$, the reduction of Q to Q^* leaves nothing and Λ^* is empty.

A linear programming algorithm can be used to determine whether or not a particular λ lies in the interior or exterior of Λ . The computation complexity of this calculation is therefore polynomial in the dimensions of the matrix Q . However, m may be very large. If $O = \{1, 2, \dots, n\}$ and the m columns of Q express all the ways in which a bin of size n can be filled, then the calculation will grow exponentially, since $m(n)$, the number of partitions of n , is asymptotically $(1/4n\sqrt{3})\exp(\pi\sqrt{2n/3})$.

In the proof of Theorem 4 we introduce a probabilistic bin-packing algorithm which stabilizes the system if and only if the system can be stabilized. The algorithm does not need to know λ . The key idea is to view the placing of an object in a bin as filling a slot in a particular configuration. The state of the system is given by a list of partially full configurations. This is a more useful description of the state of the system than one which simply states the various amounts of empty space in partially full bins. The wasted space can be studied in terms of the random walks on the non-negative integers, and the theorems are proved using well-known results for these walks.

2. Characterization of stabilizable systems

It is helpful to elaborate the algebraic description of the closed convex polyhedral cone $\Lambda = \{Qf: f \geq 0\}$. The cone is formed from non-negative combinations of a finite number of generators. These generators, which correspond to edges of the cone, are a subset of the m columns of Q . The cone is bounded by hyperplanes. Suppose h is a unit-length normal to the boundary hyperplane given by $\{\lambda: h'\lambda = 0\}$, and suppose the direction of h is away from Λ , in the sense that $h'\lambda \leq 0$ for all $\lambda \in \Lambda$. If H denotes the matrix whose columns are all such vectors h then $\Lambda = \{\lambda: H'\lambda \leq 0\}$.

Remark. If we know H then it is easy to determine whether or not a particular λ lies in Λ . However, the calculation of H from Q is equivalent to

finding the set of extreme points of a linear program and it is therefore a calculation whose computational complexity grows exponentially in m and n .

We now describe the packing process. Throughout the following we use capital letters for processes which take only integer values and small letters for processes which may take real values. We define the filled bin process N^t by $A^t = QN^t + B^t$, where N^t, B^t are non-negative integer-valued processes with N^t minimizing $\|B^t\|_2 = \sqrt{\sum(B_i^t)^2}$ over all integer-valued processes. N_j^t is the number of fully packed bins at time t which use configuration j . Note that N^t may not be unique. B_i^t is the number of objects of size o_i packed in bins which are only partially full. The number of objects in partially full bins is $\|B^t\|_1 = \sum |B_i^t|$. Since the size of each object is at least 1 there are between $\|B^t\|_1/N$ and $\|B^t\|_1$ partially full bins. Thus the wasted space satisfies $\|B^t\|_1/N \leq W^t \leq \|B^t\|_1$ for all possible packings. The following lemma bounds the wasted space in terms of the normals to boundary hyperplanes of Λ .

Lemma 1. Suppose h is a unit-length vector normal to a boundary hyperplane of Λ and h points away from Λ . Then $\|B^t\|_2 \geq (h^t A^t)^+$.

Proof. If $h^t A^t \leq 0$ the conclusion of the lemma is trivial. If $h^t A^t > 0$ then A^t is exterior to Λ and $h^t A^t$ is the distance of A^t to the boundary hyperplane of Λ . Let n^t minimize $\|A^t - Qn^t\|_2$, where n^t is a real-valued process. Then the lemma follows from

$$\|B^t\|_2 = \|A^t - QN^t\|_2 \geq \|A^t - Qn^t\|_2 \geq h^t A^t.$$

Proof of Theorem 1. Suppose $\alpha = \min\{\|\lambda - x\|_1 : x \in \Lambda\} > 0$ and that this minimum is achieved by $x = \hat{\lambda}$. Let \hat{f} be such that $\hat{\lambda} = Q\hat{f}$ and let n^t be any filled bin process. Then

$$\begin{aligned} \|A^t - Qn^t\|_1 &= \|A^t - \lambda t + \lambda t - Qn^t\|_1 \geq \|\lambda t - Qn^t\|_1 - \|A^t - \lambda t\|_1 \\ &\geq \|\lambda - \hat{\lambda}\|_1 t - \|A^t - \lambda t\|_1. \end{aligned}$$

Hence $\|A^t - Qn^t\|_1 \geq \alpha t - \|A^t - \lambda t\|_1$ for all t and

$$\|B^t\|_1 = \|A^t - QN^t\|_1 \geq \|A^t - Qn^t\|_1 \sim O(t) - o(t) \quad \text{almost surely.}$$

Proof of Theorem 2. Consider first the deterministic case, in which interarrival times for each object size are constant. Suppose $\lambda = Qf$ for $f \geq 0$. Then

$$A^t \geq \lfloor \lambda t \rfloor = \lfloor Qft \rfloor \geq Q \lfloor ft \rfloor,$$

and

$$A^t \leq \lfloor \lambda t \rfloor + \mathbf{1} \leq \lfloor Qft \rfloor + \mathbf{1} \leq Q \lfloor ft \rfloor + Q\mathbf{1} + \mathbf{1} \leq Q \lfloor ft \rfloor + (mN + 1)\mathbf{1},$$

where $\lfloor ft \rfloor$ denotes the integer part of ft taken componentwise. One can see from

the above that it is possible to implement an online algorithm which at time t has completely filled at least $N'_j = \lfloor ft \rfloor_j$ bins with configuration j and leaves no more than $mN + 1$ objects of size o_i in partially full bins. Thus the wasted space is always bounded by $nN(mN + 1)$.

When the interarrival times have non-zero variance and $\lambda \in \text{bdy}(\Lambda)$, then λ lies in some boundary hyperplane to the cone Λ . Let h be the unit-length vector which is normal to this hyperplane and which points away from Λ , and consider the process $y_t = (h'A')^+$. The central limit theorem implies that for some σ the distribution of $(h'A' - h'\lambda t)/\sigma\sqrt{t}$ converges to that of the standard normal random variable as $t \rightarrow \infty$. Since $h'\lambda = 0$ we have $E[(y_t)^+] \approx \sigma\mu\sqrt{t}$, where μ is the mean of the absolute value of a standard normal random variable. This fact combined with Lemma 1 completes the proof of the theorem.

Proof of Theorem 3. Consider an open ball $B(\lambda, \delta) \subseteq \text{int}(\Lambda)$ centred on λ . Since $A'/t \rightarrow \lambda$ almost surely, there exists a random time τ which is almost surely finite such that $A'/t \in B(\lambda, \delta)$ for all $t > \tau$. Therefore $A' \in B(\lambda t, \delta t) \subseteq \Lambda$, since Λ is a cone, and we can solve $A' = Qn'$ for all $t > \tau$. We have

$$Q[n'] \leq A' \leq Q[n'] + Q1 \leq Q[n'] + (mN)1,$$

and so it is possible to restrict the wasted space to no more than nmN^2 for $t > \tau$ by repacking so that there are at least $\Sigma[n']_i$ full bins at time t .

The rest of this section leads up to the proof of Theorem 4 and throughout what remains we suppose that each arrival process A'_j is a Poisson process of rate λ'_j . Let $X'_j = (X'_{j1}, \dots, X'_{jn_j})'$ be a vector describing the state at time t of bins which have been filled or are being filled according to packing configuration c_j . Here X'_{jk} is the number of objects (of size c_{jk}) which have been placed in bins for which the packing configuration has been chosen to be c_j and each of which has been placed in the the k th slot of that configuration (see Figure 1). Let $d'_{jkl} = |X'_{jk} - X'_{jl}|$. Consider a vector, $e'_j = (e'_{j1}, \dots, e'_{jn_j})'$, $e'_{jk} \in \{1, 2, \dots, n_j\}$. We say that e'_j is the *stabilizing effort vector* for X'_j if $e'_{jk} = s$ when X'_{jk} is the s th largest distinct integer in the set $\{X'_{j1}, \dots, X'_{jn_j}\}$. This means that $e'_{jk} > e'_{jl}$ whenever $X'_{jk} < X'_{jl}$ and $e'_{jk} = e'_{jl}$ whenever $X'_{jk} = X'_{jl}$. Because the arrival processes are Poisson, X'_j , $j = 1, \dots, m$, provides complete information about the state.




Figure 1. Bins filled using c_j . $X'_{j1} = 4$, $X'_{j2} = 1$, $X'_{j3} = 3$

We also let W_j^t denote the total empty space in those bins which are partially packed using configuration j .

Lemma 2. Suppose that all bins are to be packed using the one configuration c_j and that at time t the state is X_j^t . Let e_j^t be the stabilizing effort vector and suppose the arrival process for those objects of size c_{jk} which are to be packed in the k th slot of c_j is a non-homogeneous Poisson process with rate

$$(1) \quad \mu'_{jk} = g_j^t + \varepsilon e_{jk}^t, \quad k = 1, \dots, n_j,$$

with $\varepsilon > 0$, and $0 \leq g_j^t < G$ for some $G > 0$. Then there exists $B < \infty$ such that $E[W_j^t] < B$.

Proof. The process d'_{jkl} is an asymmetric random walk on the non-negative integers with unit jumps, a reflecting boundary at 0 and a negative drift towards 0 of at least ε . The state at time t of such a random walk has a finite expected value which is bounded uniformly in t . Hence $E[d'_{jkl}]$ can be uniformly bounded in t for each k and l . The lemma follows by taking the expected value of $W_j^t \leq N \sum_{k,l} d'_{jkl}$.

The following lemma describes the method by which we can construct a stabilizing bin-packing algorithm when λ is known.

Lemma 3. If the arrival processes are Poisson, and $\lambda > 0$ is in the interior of Λ and known, then there exists a $B < \infty$ and is a packing algorithm depending on λ for which $E[W^t] < B$.

Proof. The proof of the lemma is based upon the notion of decomposing the system into m subsystems, the j th of which is concerned with filling bins according to configuration c_j . Input processes to these m subsystems are constructed by splitting the one arrival process A^t into m arrival processes, one for each subsystem. The splitting is done so that the rates at which objects enter each subsystem satisfy (1) and so that Lemma 2 may be applied to each subsystem. The splitting is achieved by probabilistically routing the arrivals to the subsystems using the time-varying routing probabilities μ'_{jk} .

To achieve an appropriate splitting of A^t we show that there exists a vector $g^t = (g_1^t, \dots, g_m^t)$, $g \geq 0$, and some $\varepsilon > 0$ not depending on t , such that

$$(2) \quad \mu'_{jk} = g_j^t + \varepsilon e_{jk}^t, \quad j = 1, \dots, m \quad \text{and} \quad k = 1, \dots, n_j,$$

$$(3) \quad \lambda_i = \sum_{j=1}^m \sum_{k=1}^{n_j} 1(c_{jk} = o_i) \mu'_{jk},$$

where $1(c_{jk} = o_i)$ is 0 or 1 as $c_{jk} \neq o_i$ or $c_{jk} = o_i$ respectively. Now at each time t that an object of size o_i arrives we calculate the e_{jk}^t 's and solve (2) and (3) for the g_j^t 's and the μ'_{jk} 's. We then place the object in a slot of size o_i , by choosing with

probability μ'_{jk} to place the object in slot k of configuration c_j , where $c_{jk} = o_i$. To show that we can indeed solve (2) and (3) we combine them to give

$$(4) \quad \lambda_i = \sum_{j=1}^m q_{ij}g'_j + \varepsilon \sum_{j=1}^m \sum_{k=1}^{n_j} 1(c_{jk} = o_i) e'_{jk},$$

$$(5) \quad \lambda = Qg' + \varepsilon u', \quad \text{where } u'_i = \sum_{j=1}^m \sum_{k=1}^{n_j} 1(c_{jk} = o_i) e'_{jk}.$$

Now every u'_i is no more than mN^2 . Let $D = mN^2$ and let $d(\lambda)$ denote the minimum distance of λ from the boundary of Λ . Since λ is in the interior of Λ the open ball $B(\lambda, d(\lambda))$ is in Λ . Let $\varepsilon = \varepsilon(\lambda) = d(\lambda)/(D\sqrt{n})$. Then $\lambda - \varepsilon u' \in \Lambda$ for all t and the equation $\lambda - \varepsilon u' = Qg'$ has a solution $g' \geq 0$ for all t .

To summarize: the splitting of the arrival process is accomplished in the following way. When an object of size o_i arrives at time t we calculate the stabilizing effort vector e'_j for each X'_j . We solve (5) for g' and then make a randomized decision about where the object should be packed. Amongst all slots having $c_{jk} = o_i$, with probability μ'_{jk} we choose to place the object in slot k of configuration c_j . Now (2) states that the arrival rate of those objects which are to be put in slot k of a bin being packed with configuration c_j can be written as the sum of a process of rate g'_j , which is independent of k , and a term which depends on the stabilizing effort vector for X'_j at time t . Since each g'_j is less than $\max[\lambda_i]$ we can apply Lemma 2 to each subsystem, taking $G = \max[\lambda_i]$. From Lemma 2 we deduce that for each subsystem j there is a $B < \infty$ such that $E[W'_j] < B/m$. The sum over j gives $E[W'] < B$.

The above proof of Lemma 3 describes a bin-packing algorithm which requires knowledge of λ and ε in solving (5) and (2). Since $\varepsilon = d(\lambda)/D\sqrt{n}$ the algorithm will be denoted $ALG(\lambda)$. The following states that the system is also stabilizable using $ALG(\hat{\lambda})$ if $\hat{\lambda}$ is sufficiently close to λ .

Lemma 4. There exists an open ball $B(\lambda, \delta)$ centred on the actual arrival rate λ such that $ALG(\hat{\lambda})$ stabilizes the system if $\hat{\lambda}$ is in $B(\lambda, \delta)$. That is, if $\hat{\lambda}$ and $\varepsilon(\hat{\lambda}) = d(\hat{\lambda})/D$ are used instead of λ and ε to compute the μ'_{jk} 's from (5) and (2), then the system is still stabilizable when these μ'_{jk} 's are used to determine the routing probabilities.

Proof. First note that a sufficient condition for $E[W'] < B < \infty$ is

$$(6) \quad \text{for all } j, k, l: \quad \mu'_{jk} < \mu'_{jl} - \varepsilon \quad \text{if } X'_{jk} > X'_{jl}.$$

Suppose $ALG(\hat{\lambda})$ is used but the true arrival rate is λ . In this case, let $\mu'_{jk}(\lambda, \hat{\lambda})$ denote the arrival rate of those objects which will be placed in the k th slot of bins being packed with configuration c_j . We need to show that there exists an open ball around λ such that if $\hat{\lambda}$ is in this ball these $\mu'_{jk}(\lambda, \hat{\lambda})$'s satisfy (6).

Suppose $\hat{\lambda} \in B(\lambda, d(\lambda)/2)$. This implies $d(\hat{\lambda}) \geq d(\lambda)/2$ and hence $\varepsilon(\hat{\lambda}) = d(\hat{\lambda})/D \geq d(\lambda)/2D = \varepsilon(\lambda)/2$. By construction, the $\mu'_{jk}(\hat{\lambda}, \hat{\lambda})$'s satisfy (6) with ε replaced by $\varepsilon(\hat{\lambda})$. But since $\varepsilon(\hat{\lambda}) \geq \varepsilon(\lambda)/2$ they actually satisfy (6) with ε replaced by $\varepsilon(\lambda)/2$.

Now consider $\tilde{\lambda}$ near $\hat{\lambda}$. By the construction in (2) and (5) the rate $\mu'_{jk}(\tilde{\lambda}, \hat{\lambda})$ is continuous in $\tilde{\lambda}$, uniformly for all $\hat{\lambda}$'s within any bounded set within Λ . Therefore, for all $\hat{\lambda} \in B(\lambda, d(\lambda)/2)$ we can choose a δ , with $\delta < d(\lambda)/2$, which does not depend on $\hat{\lambda}$ and is such that if $\tilde{\lambda} \in B(\hat{\lambda}, \delta)$ then the $\mu'_{jk}(\tilde{\lambda}, \hat{\lambda})$'s satisfy (6) with ε replaced by $\varepsilon(\lambda)/3$ (here choosing $\varepsilon(\lambda)/3$ as something less than the quantity $\varepsilon(\lambda)/2$ found at the conclusion of the previous paragraph). This proves the lemma, since (6) now holds for the $\mu'_{jk}(\lambda, \hat{\lambda})$'s when ε is replaced by $\varepsilon(\lambda)/3$, and $\|\hat{\lambda} - \lambda\|_2 < \delta$.

We use the above lemma to construct a stabilizing algorithm based upon a converging estimate of λ which is independent of the arrival process.

Lemma 5. Suppose that $\hat{\lambda}^t$ is a stochastic process which is independent of A^t and whose time-varying value is used as an estimate of λ . Suppose that for all $\delta > 0$ there exists a τ such that $\hat{\lambda}^t$ lies inside the ball $B(\lambda, \delta)$ for all $t > \tau$, and $E[\tau] < \infty$. Then $\text{ALG}(\hat{\lambda}^t)$ stabilizes the system.

Proof. Let δ be chosen so that the conclusions of Lemma 4 hold. By hypothesis, there is a random time τ such that $\hat{\lambda}^t$ lies within $B(\lambda, \delta)$ for all $t > \tau$, and $E[\tau] < \infty$. Let $K = \sum A_i^\tau$ be the number of objects which have arrived by time τ , where $E[K] = E[\tau] \sum \lambda_i < \infty$. Then for slots k and l of c_j and all t ,

$$(7) \quad E[d'_{jkl}] \leq E[K] + B,$$

for some finite B . This is trivial when $t < \tau$ since there are no more than K arrivals by time t and so $d'_{jkl} \leq K$. Recall that d'_{jkl} performs an asymmetric random walk on the non-negative integers. It increases and decreases in unit steps; the rates of increase and decrease are uniformly bounded above and there is a drift towards 0 of at least $\varepsilon(\lambda)/3$ if $\hat{\lambda}^t \in B(\lambda, \delta)$. Such a random walk returns to 0 infinitely often and it can be shown that the expected value of the maximum achieved before returning to 0 when the process starts in state k is bounded above by $k + B$, for some finite B . (Here B is just the expected value of the maximum achieved by a similar random walk, which starts at 0, but which does not have the reflecting boundary at 0.) Since d'_{jkl} is at most K , we see that as d'_{jkl} returns to 0 for the first time following τ , it reaches a maximum on the way whose expected value is bounded above by $E[K] + B$. Once the first return to 0 has been made, the same B is a bound for $E[d'_{jkl}]$ thereafter. The proof of the lemma follows from summing (7) over all j, k and l (with of course different K and B) to show $E[W^t]$ is uniformly bounded for all t .

Note that it is essential that $\hat{\lambda}^t$ be independent of A^t . If $\hat{\lambda}^t$ were computed from $\{A^s: 0 \leq s \leq t\}$, say by $\hat{\lambda}^t = A^t/t$, then it would be incorrect to claim that A^t is Poisson after τ , since if this were so then $\hat{\lambda}^t$ could certainly exit $B(\lambda, \delta)$ after time τ . It is for this reason that we have approached the proof of Theorem 4 indirectly.

Proof of Theorem 4. The proof is almost complete. The key idea is to split the arrival process into two similar streams and let these feed arrivals to two identical bin-packing facilities, say F^1 and F^2 . It is convenient to suppose that $A^t/t \rightarrow 2\lambda$. As each object arrives, it is randomly sent to F^1 or F^2 with equal probabilities. Thus each of the two facilities is fed by a Poisson process of rate λ . Moreover, these arrival processes, say $A^{1,t}$ and $A^{2,t}$, are statistically independent. At F^1 we use algorithm $\text{ALG}(\hat{\lambda}^{2,t})$ with $\hat{\lambda}^{2,t} = A^{2,t}/[t]$ and at F^2 we use $\text{ALG}(\hat{\lambda}^{1,t})$ with $\hat{\lambda}^{1,t} = A^{1,t}/[t]$ (where $[t]$ denotes the integer part of t , these estimators are employed for $t \geq 1$ and we use any estimators for $t < 1$). By this device the estimate of λ used to determine the bin-packing at each facility is independent of the arrival process at that facility. Now $\hat{\lambda}^{1,t}$ and $\hat{\lambda}^{2,t}$ satisfy the conditions of Lemma 5. That is, they tend to λ almost surely (by the strong law of large numbers). They are also within δ of λ for all t greater than some random times τ_1, τ_2 which have finite means. To see this, we use a Chebyshev inequality to give

$$\begin{aligned} P(\|\hat{\lambda}^{1,t} - \lambda\|_2 > \delta) &= P(\|A^{2,[t]} - \lambda[t]\|_2 > \delta[t]) \\ &\leq E\left(\sum_i |A_i^{2,[t]} - \lambda_i[t]|\right)^6 / \delta^6 [t]^6. \end{aligned}$$

By calculations of moments of Poisson random variables we can show that the final term above is $O([t]^{-3})$. Then summing on $t = s, s+1, \dots$, gives $P(\tau_1 \geq s) \sim O([s]^{-2})$. Further summation on s gives $E[\tau_1] < \infty$. Lemma 5 applies to both facilities and the proof of Theorem 4 is complete.

3. Discussion

The algorithm constructed to prove Theorem 4 is a complicated one. We must, as each object arrives, calculate the e'_{jk} 's and then solve the n linear equations in m variables which appear in (5). This is a polynomial-time calculation in n and m , but its interest is theoretical rather than practical. We expect there to be some simpler algorithm which stabilizes the system if and only if it can be stabilized, and we expect it to do so if objects arrive according to renewal processes, not just Poisson processes. Perhaps some modifications of the best fit or first fit algorithms might work. We also hope to characterize stability for the continuous model, in which the bins are of size 1 and successive objects have sizes which are

independent, identically distributed samples from a distribution on $(0, 1]$. Bentley et al. (1984) and others have a number of results. For example, when the distribution is uniform on $(0, a]$, use of the first fit algorithm leads to $E[W'] \sim O(t^{4/5})$. For more results see Coffman et al. (1984).

References

BENTLEY, J. L., JOHNSON, D. S., LEIGHTON, F. T., MCGEOCH, C. C. AND MCGEOCH, L. A. (1984) Some unexpected results for bin-packing. Bell Laboratories Technical Memorandum.

COFFMAN, E. G., JR., GAREY, M. R. AND JOHNSON, D. S. (1984) Approximation algorithms for bin-packing — an updated survey, in *Algorithm Design for Computer System Design*, ed. G. Ausiello, M. Lucertini and P. Serafini, CISM Courses and Lectures No. 284, Springer-Verlag, New York.

HOFFMANN, U. (1982) A class of simple stochastic online bin-packing algorithms. *Computing* **29**, 227–239.