Probability in the Engineering and Informational Sciences, 4, 1990, 447–460. Printed in the U.S.A.

STABILITY OF ON-LINE BIN PACKING WITH RANDOM ARRIVALS AND LONG-RUN-AVERAGE CONSTRAINTS

Costas Courcoubetis
AT&T Bell Laboratories
Murray Hill, NJ 07974

Richard Weber
Engineering Department
Cambridge University
Cambridge CB3 9ET, England

Items of various types arrive at a bin-packing facility according to random processes and are to be combined with other readily available items of different types and packed into bins using one of a number of possible packings. One might think of a manufacturing context in which randomly arriving subassemblies are to be combined with subassemblies from an existing inventory to assemble a variety of finished products. Packing must be done on-line; that is, as each item arrives, it must be allocated to a bin whose configuration of packing is fixed. Moreover, it is required that the packing be managed in such a way that the readily available items are consumed at predescribed rates, corresponding perhaps to optimal rates for manufacturing these items. At any moment, some number of bins will be partially full. In practice, it is important that the packing be managed so that the expected number of partially full bins remains uniformly bounded in time. We present a necessary and sufficient condition for this goal to be realized and describe an algorithm to achieve it.

1. ON-LINE BIN PACKING WITH RANDOM ARRIVALS AND INVENTORY

Items of type $a_1, \ldots, a_m$ arrive at a bin-packing facility according to random processes with rates $\lambda = (\lambda_1, \ldots, \lambda_m)$. These items are to be combined with

© 1990 Cambridge University Press 0269-9648/90 $5.00 + .00

other readily available items of type $b_1, \ldots, b_n$ and packed into bins using one of K possible packings. There are many interesting instances of systems that we can model by applying a suitable interpretation to type b items. One might think, for example, of a manufacturing context in which randomly arriving subassemblies, or orders, are to be combined with subassemblies from an existing inventory to produce K types of finished products. An alternative interpretation is to consider a finished product as one consisting of a number of parts that arrive according to some random process and a number of operations that must take place together with those parts to complete the product. In this case, the parts correspond to the items of type $a_1, \ldots, a_m$ and the operations to the items of type $b_1, \ldots, b_n$. Packing must be done on-line. That is, as each item arrives, it must be allocated to a bin whose configuration of packing is fixed. In a manufacturing environment, this corresponds to the case in which arriving parts are assigned to products immediately when they arrive. When a bin is completely filled with the type a items it can store, a mechanism fills it with the remaining type b items and the bin leaves the facility.

At each time, there is some number of type a items that are placed in bins that are not completely full. A key concept in this article is that the packing should be managed in such a way that type b items are used at some prede-scribed rate $f = (f_1, \ldots, f_n)$ and the number of partially full bins remains bounded in time average. If this is the case, we say the system is *f-stabilizable* and note that this is required if the facility is to be managed without running out of storage while meeting the above requirements. We present a necessary and sufficient condition for the system to be f-stabilizable over a large class of arrival processes and describe an algorithm that achieves this.

An interesting application of the above is to control the rates at which different configurations are used in a bin-packing environment. If one associates a unique type b item with each possible packing configuration, then controlling the rate at which type b items are used is equivalent to controlling the rate at which filled bins with different packing configurations are produced. In a manufacturing context (see, for example, Eaves and Rothblum [6]), the desire to use type b items at rate f might reflect a constraint on the rates at which these items can be manufactured or a desire to operate the machines that produce the items at optimal rates. For example, if we associate $b_1^1, \ldots, b_K^1$ with each of the K finished products, $b_1^2, \ldots, b_n^2$ with the distinct operations needed to fabricate these products, and the corresponding frequencies are g_i for type b_i^1 and f_i for type b_i^2, then f, g could be the solution of the linear program

$$\text{maximize} \sum_{i=1}^{K} g_i R_i - \sum_{i=1}^{n} f_i C_i^0$$

$$\sum_{j=1}^{n} A_{ij} f_j \leq d_i, \qquad i = 1, \ldots, l,$$

$$\sum_{j=1}^{K} g_j c_i^j = f_i, \qquad i = 1, \ldots, n.$$

Here, R_i is the reward for producing a type i, product C_i^0 is the cost for an operation of type i, the matrix A and vector d express constraints on available resources, and c_k^j is the number of type k operations needed to manufacture a type j product. In this paper, we will not be concerned with the details of the derivation of a particular f; we will consider it as given.

The class $L(\lambda)$ of arrival processes for type a items we consider is the following. In our model, time proceeds in discrete time steps, $t = 0,1,\ldots$. Some integer constant T exists such that if we divide time in blocks of constant length T, then the following holds. During successive blocks, the number of arrivals of type a_i items are independent and identically distributed with the same distribution as the random variable A_i, $i = 1,\ldots,m$. Each A_i has mean $\lambda_i T$ and finite variance. A further assumption is that the A_i's have distributions that have exponentially bounded tails, that is, $P(A_i > k) < e^{-\omega k}$, for some $\omega > 0$, $i = 1,\ldots,m$. A special case is the Bernoulli model in which, at each time t, there is a probability λ_i of a type a_i arrival. Another interesting case is when the variance of A_i is zero, which corresponds to deterministic arrivals of item a_i.

In this article, we adopt the following convention. We will use superscripts to denote different row vectors, subscripts to denote components of such vectors, and if v is an $m + n$ vector, then v_a, v_b will denote the vectors corresponding to the projection of v onto its first m and last n components, respectively.

We describe the K possible packing configurations of the a's and b's by row vectors c^j, $j = 1,\ldots,K$. Here, c^j is a vector of $m + n$ components. Components c_i^j and c_{m+i}^j are the number of occurrences in the packing configuration c^j of type a_i and b_i items, respectively.

Let $x_i(t)$ denote the empty space for type a_i items in the partially filled bins and $n(t,j)$ the number of completely filled bins with configuration j at the end of the tth time step. We define a system to be stabilizable if, for any arrival process in $L(\lambda)$, a policy exists under which $E[x_i(t)] < B < \infty$ for all times t. A system is f-stabilizable if for any arrival process in $L(\lambda)$ there exists a policy that stabilizes it and also under this policy

$$\lim_{t\to\infty} \frac{1}{t} \sum_{j=1}^{K} n(t,j) c_b^j = f. \tag{1.1}$$

The model here is the same as that in Refs. [2] and [3] if all items are type a only and arrivals are Bernoulli. The basic result in Courcoubetis and Weber [2,3] is that the system is stabilizable if and only if λ lies in the interior of the cone generated by the vectors c^j, $j = 1,\ldots,K$, where c_i^j is the number of a_i type items that fit in a bin with configuration c^j. To indicate that it is nontrivial to decide if a system is f-stabilizable, we provide the following example. Consider the type of items a_1,a_2,a_3,b_1,b_2, and the packing configurations $c^1 = (1,0,1,1,0)$, $c^2 = (0,1,1,0,1)$, $c^3 = (1,0,0,1,0)$, $c^4 = (0,1,0,0,1)$, where c_i^j, $1 \le i \le 3$, is the number of a_i type items and c_{3+i}^j, $1 \le i \le 2$, is the number of b_i type items in configuration c^j. The arrival process for the type a items is Bernoulli with probabilities $\lambda = (1/2,1/2,1/2)$. Let the desired rate for the type

b items be $f = (1/2, 1/2)$. First observe that (λ, f) is in the cone generated by the c^j's and that this cone is a halfspace in R^5. Assume first that the type b items are not readily available, but they arrive randomly as a Bernoulli process with parameter f. Then it follows by Courcoubetis and Weber [2,3] that the system cannot be stabilized since the cone has no interior. On the other hand, one might suspect that this system is f-stabilizable if the type b items are readily available since there will be fewer fluctuations in the arrivals to the system. Furthermore, if we ignore type b items, since λ is in the interior of $c_a^1 = (1,0,1)$, $c_a^2 = (0,1,1)$, $c_a^3 = (1,0,0)$, $c_a^4 = (0,1,0)$, this system can be stabilized. Intuitively, if there is enough "flexibility," one could devise a policy that besides being stabilizing will also consume the type b items at the predetermined rate. The results in this article show that indeed, this system is f-stabilizable, whereas a second system that differs from the first since $c^2 = (0,1,1,1,1)$ and for which the same remarks hold is not f-stabilizable.

Besides providing necessary and sufficient conditions for f-stabilizability, we improve on the results in Refs. 2 and 3 in two directions. First, we allow more general arrival processes of a renewal type, which includes deterministic arrivals. This is an important extension since many of the input processes in manufacturing environments are deterministic. Second, we provide an algorithm for stabilizing the system that is simpler than those in Courcoubetis and Weber [2,3] and Courcoubetis and Rothblum [5].

This article is organized as follows. In Section 2 we provide two different equivalent Conditions 1 and 2 for the system to be f-stabilizable. These conditions are necessary and sufficient. The first condition is a natural candidate as a sufficient condition for f-stabilizability. From the second condition, it is easy to derive a simple computational procedure for checking f-stabilizability. This procedure is described in Section 3. Finally, in Section 4 we discuss some related work.

2. A NECESSARY AND SUFFICIENT CONDITION FOR f-STABILIZABILITY

One can think of the system as consisting of the packing facility and an infinite inventory of empty unconfigured bins. In the packing facility, all bins are already configured (assigned to particular packing configurations), and as items arrive, they are placed in these bins. When a bin is filled with all the type a items its configuration allows, then it immediately is filled with the remaining type b items and leaves the facility. A partially filled bin is one currently residing in the packing facility, with space remaining for some items. Note that some partially full bins might actually be empty, although already assigned to a particular packing configuration.

A packing policy is used at each time t. It looks at all type a items that arrived during this time step and then decides, first, if new bins have to be brought in from the inventory and how these will be configured, and second, in which bins to place these items. The class of policies we consider consists of policies that use the past information about the history of the system and are also time-dependent. In what follows, we provide two conditions that are equivalent and also necessary and sufficient for the existence of policies that achieve f-stability. We begin by proving that the first condition is sufficient for f-stabilizability by constructing a policy that achieves f-stability. Then we prove that the second condition is a necessary condition for f-stabilizability, and finally we prove that the second condition implies the first.

Let S be the convex cone spanned by the c^j's in R^{m+n}. The conditions are the following.

Condition 1: There exist m vectors $C^1,\ldots,C^m$ in S and positive $\gamma_1,\ldots,\gamma_m$ such that $C_a^1,\ldots,C_a^m$ are linearly independent and $\sum_{j=1}^m \gamma_j C^j = (\lambda, f)$.

Condition 2: There exists a subset $c^{j_1},\ldots,c^{j_r}$ of the c^j's such that $c_a^{j_1},\ldots,c_a^{j_r}$ span R^m, and positive $\alpha_1,\ldots,\alpha_r$ such that $\sum_{l=1}^r \alpha_l c^{j_l} = (\lambda, f)$.

THEOREM 1: *Condition* 1 *implies that the system is f-stabilizable.*

PROOF: We will define a packing policy that stabilizes the system and that also consumes type b items with rate f. The following definitions are needed in our proof. Let

$$E = \Big\{\beta \Big| \beta \in R^m,\ \beta_i \in \{1+\epsilon,\ldots,1+m\epsilon\},\ \epsilon > 0, \quad \text{and} \quad \beta_i \neq \beta_j \quad \text{for} \quad i \neq j\Big\}. \tag{2.1}$$

Define $\lambda(\beta)$ to be the vector $(\lambda_1\beta_1,\ldots,\lambda_m\beta_m)$ and assume that ϵ is small enough so that $\lambda(\beta)$ is in the cone generated by the $C_a^1,\ldots,C_a^m$ for all $\beta \in E$. Let the $\gamma_j(\beta)$'s be defined by

$$\sum_{j=1}^m \gamma_j(\beta) C_a^j = \lambda(\beta), \tag{2.2}$$

and let the δ_{jk}'s be any nonnegative solution of

$$\sum_{k=1}^K \delta_{jk} c^k = C^j. \tag{2.3}$$

We define first a mechanism $\boldsymbol{M}$ for introducing and configuring new bins in the facility. This mechanism takes as input a value β from E and adds a new bin in the facility whose configuration has been randomly chosen to be of type $k \in I$ with probability $p_k(\beta)$. These probabilities need to satisfy the equation

$$\sum_{k=1}^K p_k(\beta) c^k = \eta \sum_{j=1}^m \gamma_j(\beta) C^j, \tag{2.4}$$

where η is some positive scalar depending on β. Clearly, such probabilities exist for all $\beta \in E$, and a possible candidate is

$$p_k(\beta) = \frac{\sum_{j=1}^{m} \gamma_j(\beta)\delta_{jk}}{\sum_{k=1}^{K} \sum_{j=1}^{m} \gamma_j(\beta)\delta_{jk}}. \tag{2.5}$$

Let $Z = (Z_1, \ldots, Z_m)$ be the empty space for the type a items generated by a single invocation of $\boldsymbol{M}$. Then

$$E[Z] = \sum_{k=1}^{K} p_k(\beta) c_a^k = \eta\lambda(\beta). \tag{2.6}$$

Consider the following policy π. At every step t, this policy assigns the type a items that arrived during t to partially filled bins if there is some empty space of the corresponding type available; otherwise, it uses the mechanism $\boldsymbol{M}$ to generate new space. Since $\boldsymbol{M}$ produces a single bin each time it is invoked whose configuration might not be the appropriate one to store all the above items, the policy repeats the mechanism until eventually the correct configurations are produced and all items are stored. When a bin is filled with all the type a items it can store, it is immediately filled with the remaining type b items and is removed from the system. Note that by the definition of $\boldsymbol{M}$, at all times t there is some a_i for which $x_i(t) \le \max_{i,k}\{c_i^k\} - 1 = M$, where M is a positive constant. Also, the number of times $\boldsymbol{M}$ will be invoked until it produces space for a particular item is bounded above by a geometric random variable G with parameter $\min_{k,\beta}\{p_k(\beta)\}$.

In implementing policy π, the value of β that is supplied to $\boldsymbol{M}$ is updated every MT steps as follows. Let $\beta(x)$ be a vector in E such that $x_k/\lambda_k < x_i/\lambda_i$ implies $\beta_i < \beta_k$. Then for $kMT < t \le (k+1)MT$, whenever $\boldsymbol{M}$ is invoked, it is supplied with the same value $\beta[x(kMT)]$ that was computed based on the available empty space after the kMTth step. We will show that π achieves f-stability.

The first part of the proof consists in showing that for any pair of a_i, a_j, the process $\{Y(k), \boldsymbol{F}(k)\}$, $k = 0, 1, \ldots,$ where $Y(k) = |[x_i(kMT)/\lambda_i - x_j(kMT)/\lambda_j]|$, and $\boldsymbol{F}(k)$ is the σ-field generated by $\{x(sMT), s = 0, 1, \ldots, k\}$, behaves as a supermartingale with (1) negative drift uniformly bounded away from zero outside some finite set, and (2) has increments a.s. bounded above by a random variable with a distribution having an exponentially bounded tail. This implies by using the results of Hajek [7] that $E[|x_i(kMT) - x_j(kMT)|] = E[Y(k)] < B < \infty$ for all $k > 0$, and since at all times k there is an a_l for which $x_l(kMT) \le M$, it follows that $E[x_i(kMT)] < B'$, for some $B' < \infty$, and for all $i = 1, \ldots, m$ and $k > 0$. From this, it trivially follows by the definition of the arrival process and the policy π that $E[x_i(t)]$ is also bounded uniformly in t for all $i = 1, \ldots, m$. Finally, one can easily observe that the above together with the definition of π imply that $E[x_i(t)]$ is also bounded uniformly in t for all $i = m+1, \ldots, m+n$. The second part of the proof consists of showing that the foregoing policy besides being stable also consumes type b items with rate f.

We start by proving that $Y(k)$ is a supermartingale with the above properties for $Y(k)$ sufficiently large. The proof consists of the following two lemmas.

LEMMA 1: *If* $E[Y^2(k+1)|\boldsymbol{F}(k)] \leq Y^2(k) - \psi Y(k) + \psi^2/4$ *for* $Y(k) > K_1$, $K_1 < \infty$ *and,* $\psi > 0$, *then* $E[Y(k+1)|\boldsymbol{F}(k)] \leq Y(k) - \psi/2$ *for* $Y(k) > K_1$.

PROOF: Since $(E[Y(k+1)|\boldsymbol{F}(k)])^2 \leq E[Y^2(k+1)|\boldsymbol{F}(k)]$, the proof trivially follows by taking the square roots on both sides of the inequality. ∎

LEMMA 2: $E[Y^2(k+1)|\boldsymbol{F}(k)] \leq Y^2(k) - \psi Y(k) + \psi^2/4$ *for* $Y(k) > K_1$, *where* $K_1 < \infty$ *and* $\psi > 0$, *independent of* $Y(k)$.

PROOF: Let $W = (W_1, \ldots, W_m)$ denote the total amount of empty space created for the type a items by $\boldsymbol{M}$ during $kMT < t \leq (k+1)MT$, and let $N = (N_1, \ldots, N_m)$ denote the total number of arrivals of type $a_1, \ldots, a_m$ items during the same period of time. To simplify notation, let x denote $x(kMT)$ and β denote $\beta[x(kMT)]$, and assume that $x_i/\lambda_i > x_j/\lambda_j$. Observe first that

$$\begin{aligned}\left(\frac{x_i + W_i - N_i}{\lambda_i} - \frac{x_j + W_j - N_j}{\lambda_j}\right)^2 &= \left(\frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j}\right)^2 + 2\left(\frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j}\right)\left(\frac{W_i - N_i}{\lambda_i} - \frac{W_j - N_j}{\lambda_j}\right) \\ &\quad + \left(\frac{W_i - N_i}{\lambda_i} - \frac{W_j - N_j}{\lambda_j}\right)^2. \end{aligned} \tag{2.7}$$

Note also that since the arrival process regenerates itself every T steps and π does not depend on the past, $E[Y^2(k+1)|\boldsymbol{F}(k)] = E[Y^2(k+1)|x]$. From this, it follows that

$$\begin{aligned} E[Y^2(k+1) - Y^2(k)|\boldsymbol{F}(k)] &= E[Y^2(k+1) - Y^2(k)|x] \\ &= 2\left(\frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j}\right)E\left[\frac{W_i - N_i}{\lambda_i} - \frac{W_j - N_j}{\lambda_j}\,\middle|\, x\right] \\ &\quad + E\left[\left(\frac{W_i - N_i}{\lambda_i} - \frac{W_j - N_j}{\lambda_j}\right)^2 \,\middle|\, x\right]. \end{aligned} \tag{2.8}$$

Consider first the last term in (2.8). Clearly, N_i, N_j do not depend on x and have finite expectation by the assumptions of our model. Also, the W_i, W_j are bounded above by some random variable with finite expectation since there are, at most, $\sum_{i=1}^{m} N_i$ items to be packed in new bins, and for each item, the number of times we need to run $\boldsymbol{M}$ is bounded above by a geometric random variable. Clearly, this bound is also independent of the state x. This implies that this term is bounded above by some constant D independent of x.

In the first term of (2.8), we have

$$E\left[\frac{W_i - N_i}{\lambda_i} - \frac{W_j - N_j}{\lambda_j}\bigg| x\right] = \frac{E[W_i|x] - \lambda_i MT}{\lambda_i} - \frac{E[W_j|x] - \lambda_j MT}{\lambda_j}$$
$$= E\left[\frac{W_i}{\lambda_i} - \frac{W_j}{\lambda_j}\bigg| x\right]. \tag{2.9}$$

Let $u(l)$ be the random variable such that $u(l) = 1$ iff $\boldsymbol{M}$ has been invoked for the lth time during $kMTt \leq (k+1)MT$, and let $Z^{(l)} = [Z_1^{(l)}, \ldots, Z_m^{(l)}]$ denote the empty space generated by $\boldsymbol{M}$ when it was invoked for the lth time. Recall also that by the definition of $\boldsymbol{M}$ we have for all $l = 1, 2, \ldots, E[Z^{(l)}] = \eta\lambda(\beta)$ for some $\eta > 0$ that depends on β. Then,

$$\frac{W_i}{\lambda_i} - \frac{W_j}{\lambda_j} = \sum_{l=1}^{\infty} u(l)\left[\frac{Z_i^{(l)}}{\lambda_i} - \frac{Z_j^{(l)}}{\lambda_j}\right], \tag{2.10}$$

and we use monotone convergence to obtain

$$E\left[\frac{W_i}{\lambda_i} - \frac{W_j}{\lambda_j}\bigg| x\right] = \sum_{l=1}^{\infty} E\left\{u(l)\left[\frac{Z_i^{(l)}}{\lambda_i} - \frac{Z_j^{(l)}}{\lambda_j}\right]\bigg| x\right\}$$
$$= \sum_{l=1}^{\infty} E[u(l)|x]\, E\left[\frac{Z_i^{(l)}}{\lambda_i} - \frac{Z_j^{(l)}}{\lambda_j}\bigg| x\right] \tag{2.11}$$

since the $u(l)$'s and the $Z^{(l)}$'s are mutually independent. By using the fact that the $Z^{(l)}$ are identically distributed and depend on x only through β, we obtain through (2.6)

$$E\left[\frac{W_i}{\lambda_i} - \frac{W_j}{\lambda_j}\bigg| x\right] = \sum_{l=1}^{\infty} E[u(l)|x]\left(\frac{\eta\lambda_i\beta_i}{\lambda_i} - \frac{\eta\lambda_j\beta_j}{\lambda_j}\right)$$
$$\leq \eta E[u(1)|x]\,(\beta_i - \beta_j) < 0 \tag{2.12}$$

since our assumption about $x_i/\lambda_i > x_j/\lambda_j$ implies that $\beta_i < \beta_j$. Observe now that $E[u(1)]$ is the probability that during $kMT < t \leq (k+1)MT$, the mechanism $\boldsymbol{M}$ is invoked at least once, which is strictly positive since, at time kMT, there is some type a_l item for which $x_l \leq M$, and every T steps, the probability of an arrival of such an item is positive by the assumptions of our model. From this, it follows that there is some positive constant δ for which $E[u(1)|x] \geq \delta$ for all states x consistent with our policy π. Similarly, since η in the previous inequality depends on β, let η' be the minimum η achieved over all values $\beta \in E$. Now, since $\beta_i - \beta_j \leq -\epsilon$, it follows that

$$E\left[\frac{W_i}{\lambda_i} - \frac{W_j}{\lambda_j}\bigg| x\right] \leq -\epsilon\eta'\delta = -\zeta < 0, \tag{2.13}$$

where $\zeta > 0$ does not depend on x.

We return now to the proof of Lemma 2. We just proved that if $x_i/\lambda_i > x_j/\lambda_j$, then

$$E[Y^2(k+1)|x] \le Y^2(k) - 2\zeta Y(k) + D. \tag{2.14}$$

Choose any positive $\psi < 2\zeta$ and let $K_1 = (D - \psi^2/4)/(2\zeta - \psi)$. Then one can easily check that by using the above ψ and K_1, and the conditions of the lemma are satisfied. The case of $x_i/\lambda_i < x_j/\lambda_j$ can be handled in a similar way, which completes the proof of the lemma. ■

We now return to the proof of Theorem 1. The foregoing two lemmas imply that $Y(k)$ has a drift of, at most, $-\psi/2 < 0$ when $Y(k) > K_1$. To complete the first part of the proof, we must show that the increments of $Y(k)$ are a.s. bounded above by some random variable having a distribution with an exponentially bounded tail. It is easy to see that the increments of the $Y(k)$ are bounded above by $W_i/\lambda_i + W_j/\lambda_j$, which is also bounded above by the random variable

$$\Delta = \left(\frac{1}{\lambda_i} + \frac{1}{\lambda_j}\right) \max_j \left\{\sum_{i=1}^{m} c_i^j\right\} \sum_{k=1}^{A(kMT)} G_k,$$

where $A(kMT)$ is the total number of arrivals for all objects that arrived during $kMT < t \le (k+1)MT$, and the G_k's are i.i.d. random variables distributed as the random variable G introduced before. One can easily check that the assumptions about the arrival processes imply that $A(kMT)$ has a distribution with an exponentially bounded tail, which implies the same for Δ as well. This completes the first part of the proof of the theorem.

To complete the second part of the proof of the theorem, we must show that π consumes type b items with rate f. Let $u(t,l,\beta)$ be the indicator function on the event that during the first t time steps the mechanism M was invoked for the lth time and this occurred with parameter β. Let $V(t)$ denote the total space (for item types a and b) produced up to time t. Then

$$V(t) = \sum_{l=1}^{\infty} \sum_{\beta \in E} u(t,l,\beta) Z^{(l)}(\beta), \tag{2.15}$$

where in order to simplify notation, this time $Z^{(l)}(\beta)$ denotes the $m+n$ vector of the empty space for type a, b items produced by the lth invocation of M with parameter β. Then monotone convergence implies that

$$\begin{aligned} E[V(t)] &= \sum_{l=1}^{\infty} E\left[\sum_{\beta \in E} u(t,l,\beta)\right] E[Z^{(l)}(\beta)] \\ &= \sum_{l=1}^{\infty} E\left[\sum_{\beta \in E} u(t,l,\beta)\right] \eta(\beta) \sum_{j=1}^{m} \gamma_j(\beta) C^j \\ &= \sum_{\beta \in E} \eta(\beta) \sum_{j=1}^{m} \gamma_j(\beta) C^j \sum_{l=1}^{\infty} E[u(t,l,\beta)], \end{aligned}$$

and by defining $\sigma(t,\beta) = \sum_{l=1}^{\infty} E[u(t,l,\beta)]$, we get

$$E[V(t)] = \sum_{\beta \in E} \eta(\beta)\sigma(t,\beta) \sum_{j=1}^{m} \gamma_j(\beta) C^j = \sum_{j=1}^{m} \rho(t,j) C^j, \tag{2.16}$$

where we define $\rho(t,j) = \sum_{\beta} \eta(\beta)\sigma(t,\beta)\gamma_j(\beta)$.

Let the $m + n$ vector $N(t)$ denote the total number of items of each type that are packed in bins up to time t. This corresponds to the total number of arrivals for each type a item and the total number of type b items that have been assigned to the completely filled bins up to time t. Now, since the system is stable, it follows that $E[V_a(t) - N_a(t)] < B$ for all t, which in turn implies that

$$\lim_{t \to \infty} \sum_{j=1}^{m} \frac{\rho(t,j)}{t} C_a^j - \lambda = 0. \tag{2.17}$$

Since $\lambda = \sum_{j=1}^{m} \gamma_j C_a^j$ by Condition 1, it follows that

$$\lim_{t \to \infty} \sum_{j=1}^{m} \left[\frac{\rho(t,j)}{t} - \gamma_j \right] C_a^j = 0. \tag{2.18}$$

This implies that $\lim_{t \to \infty} \rho(t,j)/t = \gamma_j$ since he C_a^j's are linearly independent.

Consider now the $\lim_{t \to \infty} E[N_b(t)]/t$. Since $E[V_b(t) - N_b(t)] < B$ for all t, it follows by using Condition 1 that

$$\lim_{t \to \infty} E \frac{[N_b(t)]}{t} = \lim_{t \to \infty} \sum_{j=1}^{m} \frac{\rho(t,j)}{t} C_b^j = \sum_{j=1}^{m} \gamma_j C_b^j = f. \qquad \blacksquare$$

THEOREM 2: *f-stabilizability implies Condition* 2.

PROOF: We show first that (λ, f) must be in the core generated by the c^j's. Assume that there exists some packing strategy that is f-stable. Let $g = (\lambda, f)$. Then if $V_i(t)$ denotes the total number of empty slots for type i items produced by time t, we have

$$\lim_{t \to \infty} \frac{E[V_i(t)]}{t} = g_i.$$

Let $a_k(t)$ denote the number of times that a type k configuration has been used by time t. Then since $E[V(t)]/t = \sum_k E[a_k(t)]/t \, c^k$ is in the cone generated by the c^k's for all t and this cone is a closed set, it follows that g is also in the cone.

Consider now the "reduced" system with bins that do not require any type s items, but have the same set of configuraitons c_a^j for type a packing items and the same arrival process as the original system. In Courcoubetis and Weber [2,3] it is proved that if the arrival processes are Bernoulli, then the system is stabilizable if and only if λ is in the interior of the cone generated by the c_a^j's. Clearly, if the original system is f-stabilizable for all arrival processes satisfying the conditions of the model, then the "reduced" system must also be stabilizable for the special case of Bernoulli arrivals with parameter λ. This together with the previous observations implies that the c_a^j's must span R^m so that the cone generated by the c_a^i's has an interior, and λ must be in the interior of this cone. This completes the proof of the theorem.

THEOREM 3: *Condition* 2 *implies Condition* 1.

PROOF: Assume without loss of generality that $j_1,\ldots,j_r$ is $1,\ldots,r$. Then Condition 2 is equivalent to $\alpha_1 c^1 + \cdots + \alpha_r c^r = (\lambda,f)$, $\alpha_i > 0$, $i = 1,\ldots,r$, and $c_a^1,\ldots,c_a^r$ span R^m. Note that $r \geq m$. Suppose without loss of generality that the vectors $\alpha_1 c_a^1,\ldots,\alpha_m c_a^m$ are linearly independent and define $v^1 = \alpha_1 c_a^1,\ldots,$ $v^m = \alpha_m c_a^m$. By definition, these vectors satisfy

$$v^1 + \cdots + v^m + \alpha_{m+1} c^{m+1} + \cdots + \alpha_r c^r = (\lambda,f), \tag{2.19}$$

and the $v_a^1 + \cdots + v_a^m$ are linearly independent. Consider now the $\beta_1,\ldots,\beta_{m+1}$ satisfying

$$\beta_1 \alpha_1 c_a^1 + \cdots + \beta_m \alpha_m c_a^m + \beta_{m+1} \alpha_{m+1} c_a^{m+1} = 0, \tag{2.20}$$

and assume without loss of generality that β_1 and β_{m+1} have opposite sign. Define the m vectors $\tilde{v}^1 = \alpha_1 c^1 + \alpha_{m+1} c^{m+1}$, $\tilde{v}^2 = \alpha_2 c^2,\ldots,\tilde{v}^m = \alpha_m c^m$. These vectors satisfy

$$\tilde{v}^1 + \cdots + \tilde{v}^m + \alpha_{m+2} c^{m+2} + \cdots + \alpha_r c^r = (\lambda,f), \tag{2.21}$$

and the $\tilde{v}_a^1,\ldots,\tilde{v}_a^m$ are linearly independent since if this is not the case, there must exist some $\delta_1,\ldots,\delta_m$ such that $\delta_1 \tilde{v}_a^1 + \cdots + \delta_m \tilde{v}_a^m = 0$. But since this identity must be the same as (2.20), modulo a constant factor, this would imply that the coefficients of c^1 and c^{m+1} in (2.20) must be of equal sign, contradicting the assumption we made earlier. Observe now that (2.21) corresponds to (2.19), with c^m being eliminated while preserving the properties of the first m vectors on the left-hand side of the equation. By repeating $r - m - 1$ times the above construction with the v^j's being the $\tilde{v}^j$'s of the previous step, the last set of the $\tilde{v}^j$, $j = 1,\ldots,m$, will correspond to the desired set of $\gamma_1 C^1,\ldots,\gamma_m C^m$, which by induction will satisfy Condition 1. ∎

3. A COMPUTATIONAL PROCEDURE FOR DETERMINING f-STABILIZABILITY

As mentioned in the introduction, Condition 2 is easier to check. In what follows, we will describe a procedure for checking the validity of Condition 2.

Consider first the following problem: Given the set $c^1,\ldots,c^K$, find the maximal subset of indices $L \subseteq \{1,\ldots,K\}$, for which there exist positive numbers α_l, $l \in L$, such that $\sum_{l\in L} \alpha_l c^l = (\lambda,f)$. One can easily see that if such a subset exists, then a maximal one must also exist since if there exist two different subsets of the c^j's that satisfy the above condition, their union also satisfies the condition. In order to determine if $i \in L$, $i = 1,\ldots,K$, we solve the linear program

$$\text{maximize } \alpha_i$$
$$\sum_{j=1}^{K} \alpha_j c^j$$
$$\alpha_j \geq 0, \qquad j = 1, \ldots, K.$$

One can easily see that $i \in L$ if and only if the solution α_i is positive.

Checking Condition 2 is now straightforward. If L is empty, then clearly Condition 2 does not hold. If L is not empty, we use Gauss elimination to determine the rank of the matrix with rows c_a^j's, $j \ni L$. Now Condition 2 holds only if the rank is at least m.

4. DISCUSSION

It is important to note that if we consider a particular arrival process in $L(\lambda)$ and ask if for this process there exists an f-stable policy, then Conditions 1 and 2 are sufficient but not necessary. For example, if we restrict ourselves to deterministic arrival processes with mean λ, then the condition that (λ, f) is in the interior or on the boundary of S is necessary and sufficient for f-stabilizability. The above observations together with the first part of the proof for Theorem 2 imply that for a particular arrival process (1) if λ is in the exterior of S, the system is not f-stabilizable; (2) if λ satisfies Condition 2 (and hence, Condition 1), then the system is f-stabilizable; and (3) if λ is in S and does not satisfy Condition 1 or 2, then the answer depends on the particular properties of the arrival process. In case (3) we leave as an open problem the determination of the necessary and sufficient conditions on the distributions of the arrival processes for the existence of f-stabilizing policies.

A problem related to the one discussed here has been considered in Courcoubetis and Rothblum [5]. In that paper, there was a simple cost structure associated with the bin-packing system in which a reward r_j was obtained for each bin packed according to configuration c^j. The system was also on-line and the items arrived as independent Poisson processes with rate λ. The goal was to find the conditions for the existence of policies that are optimal or nearly optimal while keeping the expected number of partially filled bins bounded or relatively slowly increasing in time. Necessary and sufficient conditions were given for a system to be stabilizable while operated in an optimal way. This article includes the above since we provide a characterization of all the possible f's that can be achieved while stabilizing the system, one among these being the optimal one corresponding to the cost structure in Courcoubetis and Rothblum [5]. We do not investigate the case of systems that are "slightly" unstable and achieve $f - \epsilon$ for arbitrarily small ϵ (quasistable and ϵ-optimal in Courcoubetis and Rothblum [5]).

The problem of stability of on-line bin-packing systems with item types from a finite set and Bernoulli arrivals has been investigated in Courcoubetis and Weber [2,3]. According to the terminology of Coffman et al. [1], these results provide necessary and sufficient conditions that must hold for a distribution over a finite set of the item sizes if perfect packing is to be achieved in an on-line fashion. The related problem of off-line packing has been investigated by Lueker [8] and Rhee and Talagrand [9].

The policy π that we introduced in section 2 to stabilize the system is an improvement over the policies used in Courcoubetis and Weber [2,3] and Courcoubetis and Rothblum [5] for the following two reasons. First, π uses significantly less information since it needs to know, for each a_i, only the number $x_i(t)$. In Courcoubetis and Weber [2,3], the policies used to achieve stability need to know the numbers $x_i^j(t)$, $i = 1, \ldots, m$, $j = 1, \ldots, K$, where $x_i^j(t)$ denotes the empty space for type a_i items in bins configured with configuration c^j. Second, unlike π, these policies may introduce empty bins to the system when new items arrive, even if these items can fit in existing partially filled bins.

There is an interesting connection between the policy introduced by Courcoubetis et al. [4] for stabilizing a production system and policy π in this article. Our bin-packing system corresponds to the production process setup of Courcoubetis et al. [4] where the production processes have a particular deterministic structure and arrivals are of the type considered in this article. One can view the arrivals of items as demands for empty space of a particular type, and introducing a new bin with a certain configuration into the facility as producing a deterministic amount of empty space of various types. Unfortunately, the proofs of Courcoubetis et al. [4] require that the production and arrival processes be independent and Bernoulli, and so they do not imply the results for stabilizing the systems considered here. On the other hand, our results extend those of Courcoubetis et al. [4] for the cases discussed above.

Acknowledgments

We would like to thank Mihalis Yannakakis for many helpful discussions concerning the formulation of Condition 2 in Section 2.

References

1. Coffman, E.G. Jr., Lueker, G. S., & Rinnooy Kan, A.H. (1988). Asymptotic methods in the probabilistic analysis of sequencing and packing heuristics. *Management Science* 34(3): 266-290.
2. Courcoubetis, C. & Weber, R.R. (1986). Necessary and sufficient conditions for stability of a bin-packing system. *Journal of Applied Probability*, 23: 989–999.
3. Courcoubetis, C. & Weber, R.R. (1986). A bin-packing system for objects from a finite set: necessary and sufficient conditions for stability and some applications. *Proceedings of 25th IEEE Conference on Decision and Control, Athens, Greece.*
4. Courcoubetis, C., Konstantopoulos, P., Walrand, J., & Weber, R.R. (1989). Stabilizing an uncertain production system. *Queueing Systems: Theory and Applications* 5(1–3): 37–54.
5. Courcoubetis, C. & Rothblum, U.G. (1990). On optimal packing of randomly arriving objects. *Mathematical Operations Research* (to appear).
6. Eaves, B.C. & Rothblum, U.G. (1988). A discrete time average cost flexible manufacturing and

operator scheduling model solved by deconvexification over time. *Operations Research* 36: 242–257.

7. Hajek, B. (1982). Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advanced Applied Probability* 14: 502–525.
8. Lueker, G.S. (1983). Bin-packing with items uniformly distributed over intervals $[a,b]$. *Proceedings of 24th Symposium on Foundation of Computer Science.* New York: IEEE Computer Society Press, pp. 289–297.
9. Rhee, W.T. & Talagrand, M. (1988). Some distributions that allow perfect packing. *Journal of the Association for Computing Machinery* 35(3): 564–578.