

STABILIZING AN UNCERTAIN PRODUCTION SYSTEM

Costas COURCOUBETIS¹, Panagiotis KONSTANTOPOULOS²,
Jean WALRAND² and Richard R. WEBER³

¹ AT&T Bell Laboratories, Murray Hill, NJ07974, U.S.A.

² EECS, Univ. of California at Berkeley, Berkeley CA 94720, U.S.A.

³ Queens' College, University of Cambridge, Cambridge, U.K.

Received 22 July 1988; revised 3 January 1989

Abstract

Consider a production system that consists of m machines each of which can produce parts of n types. When machine k is used, it produces a part of type i with probability p_{ki} . Requests arrive for parts, one at a time. With probability λ_i an arriving request is for a part of type i . The requests must be served without waiting. Thus, if a requested part is not available, it must be produced. We find necessary and sufficient conditions for the existence of a strategy (a choice of the machines to be used) which makes the inventory of parts stable and we provide such a strategy.

Two variations of this model are also considered: the case of batch arrivals of requests, and that of a system where the requests can be queued.

Keywords: Stability, control strategies, inventory processes

1. Introduction and basic model

It is not unusual for machines to yield uncertain products. For instance, when cutting a crystal for electronics applications, it may happen that the result is not a crystal of the desired frequency. Thus if a crystal at frequency f_i is desired, the machine may produce one at frequency f_j with some probability. If $f_j \neq f_i$, then this crystal will not meet the request and a new attempt for producing the desired frequency has to be made. It may happen that crystals at the frequency f_j can be requested later. In that case, the products that do not meet the original requests should be stored. This paper investigates the stability of the inventory of such a system. The model we use is the following.

There are m machines ($k = 1, 2, \dots, m$) that can produce parts of n types ($i, j = 1, 2, \dots, n$). When machine k is used, it produces a part of type i with probability p_{ki} . Thus $p_{ki} \geq 0$ and $\sum_{i=1}^n p_{ki} = 1$ for $k = 1, \dots, m$. This production is assumed to be instantaneous.

Requests arrive one at a time. With probability $\lambda_i \geq 0$ an arriving request is for a part of type i (for $i = 1, \dots, kn$). Thus $\sum_{i=1}^n \lambda_i = 1$. Requests must be served without waiting. Thus a request of type i is served by removing a part of type i from the inventory, if there is one. If no part of the requested type is in the inventory, then one must choose one of the machines to produce one. Because of the uncertainty of the production system, one may have to run the machine a random number of times, and, as a consequence, produce a random number of parts of types that are not requested. A *strategy* is a method for choosing the machines to produce the requested parts that are not available. The method must be non-anticipative and can be randomized. That is, the decision taken at time t should only depend on the past (including the current request) and, possibly, on some additional random variables which are independent of everything else. We are interested in the effect of different strategies on the growth of the inventory of such a system.

Denote by $x_i(t)$ the number of parts of type i in the inventory at time t for $t = 0, 1, \dots$ and $i = 1, \dots, n$. We will say that a strategy is *stable* if there is some $B < \infty$ such that

$$E[x_i(t)] < B \quad \text{for all } t \geq 0 \quad \text{and } i = 1, \dots, n. \quad (1.1)$$

In this paper we investigate the conditions for the existence of stable strategies for the above system. A “dual” problem has been investigated in [2]. That paper considers a system in which objects arrive to be placed in bins where objects and bins have integer size. Once an object is placed in an empty bin it generates some wasted space (the remaining unused space in the bin). In this system the “inventory” of wasted space (sum of empty space of the partially filled bins) is reduced every time an object that arrive is placed in some bin which is partially filled, and is increased every time a new bin starts being filled. The results of the paper are the necessary and sufficient conditions for the existence of stable packing strategies, i.e., strategies that produce expected wasted space uniformly bounded for all times. The stability conditions obtained are similar to the ones in this paper and are given in terms of the vector of the arrival rates of the different types of objects λ being in the interior of some convex cone. The stabilizing strategy proposed is also similar to the one in this paper and uses randomization. Although there is some connection between the two problems, the stabilizing strategy proposed in this paper is fundamentally different than the one in [2] in that it requires substantially less information about the system.

In order to simplify the discussion let us observe that one can eliminate from consideration the machines that produce parts that are never requested. That is, if $\lambda_i = 0$ and $p_{ki} > 0$, then machine k should not be used. Thus one can assume without loss of generality that $\lambda_i > 0$ for all i . This assumption will be made throughout the paper. Furthermore, let p_k denote the probability vector (p_{k1}, \dots, p_{kn}) associated with machine k .

The basic stability results are introduced in section 2 and are extended in the subsequent sections. Section 3 deals with the case where requests arrive in batches and cannot be queued. In section 4 we analyze the case of queued requests.

2. Stable strategies for the basic model

COMPLETE INVENTORY INFORMATION

This corresponds to the controller of the system having complete information about the total inventory x at all times. The following result provides the necessary and sufficient conditions for the existence of stable strategies for the above system. It also describes such a stable strategy.

THEOREM 2.1

There exists a stable strategy if and only if $\lambda = (\lambda_1, \dots, \lambda_n)$ belongs to the interior of the cone

$$\Lambda = \left\{ \sum_{k=1}^m \alpha_k p_k : \alpha_k \geq 0, k = 1, \dots, m \right\}. \tag{2.1}$$

Proof

a) Sufficiency

We will exhibit a stable strategy. Notice first that Λ is a closed convex polyhedral cone in n -dimensional space and the interior of the cone is non-empty if and only if the null space $\{x: Px = 0\} = \{0\}$, where $P = \{p_{kj}\}$. Suppose that the vector λ lies in the interior of the cone Λ . Given the state $x = (x_1, x_2, \dots, x_n)$ of the inventory, it can be easily seen that there is a vector $\beta = \beta(x) = (\beta_1, \dots, \beta_n)$ that satisfies the following conditions:

- (i) $\beta_i > 1$, for all i .
- (ii) The vector $\beta(x)$ depends only on the ordering of the vector

$$\left(\frac{x_1}{\lambda_1}, \dots, \frac{x_n}{\lambda_n} \right)$$

in such a way that

$$\beta_i > \beta_j \quad \text{if} \quad \frac{x_i}{\lambda_i} < \frac{x_j}{\lambda_j}.$$

In other words, if

$$\left(\frac{x_1}{\lambda_1}, \dots, \frac{x_n}{\lambda_n} \right)$$

is ordered the same way as

$$\left(\frac{x'_1}{\lambda_1}, \dots, \frac{x'_n}{\lambda_n} \right)$$

then $\beta(x) = \beta(x')$.

(iii) If $\tilde{p}_{kj} = \beta_j p_{kj}$ then the vector λ lies in the interior of the cone formed by the vectors $\tilde{p}_k = (\tilde{p}_{k1}, \dots, \tilde{p}_{kn})$, $k = 1, \dots, m$. In other words, λ can be written as

$$\lambda = \sum_{k=1}^m \alpha_k \tilde{p}_k, \quad \text{where } \alpha_k > 0, \quad \text{for all } k.$$

For example, we can define β as a permutation of the numbers $1 + \mu, 1 + 2\mu, \dots, 1 + n\mu$ for some $\mu > 0$. This choice will satisfy (i), (ii) and furthermore, if μ is small enough, it will also satisfy (iii). Note that it is important that λ is in the interior of Λ , else we may not be able to fulfill condition (iii). Note also that if

$$\frac{x_i}{\lambda_i} < \frac{x_j}{\lambda_j} \quad \text{then} \quad \frac{1}{\beta_i} - \frac{1}{\beta_j}$$

is bounded above by a negative number $-\varepsilon$ which, because there are only finitely many choices for β , does not depend on x . Suppose that the demand is l when the state is x . If $x_l > 0$ then the demand is immediately fulfilled and no machine needs to be chosen. If $x_l = 0$ we choose machine k with probability

$$\gamma_{kl}(x) = \frac{\alpha_k \tilde{p}_{kl}}{\lambda_l}. \tag{2.2}$$

(Note that $\sum_{k=1}^m \gamma_{kl}(x) = 1$.) This machine is run until one part of type l is produced which is then used to fulfill the request. Observe that this strategy is a Markov and stationary since it only depends on the current state and the current request and is independent of time.

It remains to show that this strategy is stable. This is done next. Fix two distinct indices i, j from the set $\{1, \dots, n\}$ and let

$$f(x) = f_{ij}(x) = \left| \frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j} \right|. \tag{2.3}$$

Define the drift of the function f at the point x as

$$E[\Delta f(x) | x] := E[f(x(t+1)) - f(x(t)) | x(t) = x]. \tag{2.4}$$

We show that

LEMMA 2.1

$$\sup_{x: f(x) > K} E[\Delta f(x) | x] < 0 \tag{2.5}$$

where K is some positive number.

Proof

An important observation is that if we assume that the system starts with empty inventory, the above strategy is such that at all times there is some part of type l in the inventory for which $x_l = 0$.

First assume that

$$\frac{x_i}{\lambda_i} > \frac{x_j}{\lambda_j} > 0, \quad \text{and that } \frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j} > K$$

for some $K > 0$. Then

$$E[\Delta f | x] = \sum_l E[\Delta f | x, l] \lambda_l \tag{2.6}$$

where l is the type of request. For $l = i$ we have $E[\Delta f | x, i] = -1/\lambda_i$, and for $l = j$ we have $E[\Delta f | x, j] = 1/\lambda_j$. Hence the expression (2.6) is

$$E[\Delta f | x] = \sum_{l \neq i, j} E[\Delta f | x, l] \lambda_l = \sum_{l: x_l = 0} E[\Delta f | x, l] \lambda_l. \tag{2.7}$$

since if $x_l > 0$ there is no change of f . Note also that the sum is never over an empty set since there is always an l such that $x_l = 0$. Now for such an l ,

$$E[\Delta f | x, l] = \sum_k E[\Delta f | x, l, k] \gamma_{kl}(x)$$

that is, we condition on the chosen machine k . Given that the demand is l (with $x_l = 0$) and the chosen production process is k , f changes to

$$\left| \frac{x_i + w_i}{\lambda_i} - \frac{x_j + w_j}{\lambda_j} \right|,$$

where w_i has mean p_{ki}/p_{kl} and similarly for w_j . Consider the event

$$U = \left\{ \frac{x_i + w_i}{\lambda_i} > \frac{x_j + w_j}{\lambda_j} \right\}.$$

Observe that on U ,

$$\Delta f = (\Delta f)' = \frac{w_i}{\lambda_i} - \frac{w_j}{\lambda_j}$$

while on U^c ,

$$\Delta f = (\Delta f)'' = \frac{w_j}{\lambda_j} - \frac{w_i}{\lambda_i} - 2 \left(\frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j} \right).$$

Furthermore note that each r.v. w_i is bounded above by a geometrically distributed r.v. T representing the number of times that machine k is run until it produces one part of type l . Hence

$$\left| \frac{w_i}{\lambda_i} - \frac{w_j}{\lambda_j} \right| \leq \frac{T}{\min(\lambda_i, \lambda_j)}$$

and so

$$\begin{aligned}
 P(U^c) &= P\left(\frac{x_i + w_i}{\lambda_i} \leq \frac{x_j + w_j}{\lambda_j}\right) = P\left(\frac{w_j}{\lambda_j} - \frac{w_i}{\lambda_i} \geq \frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j}\right) \\
 &\leq P\left(\frac{T}{\min(\lambda_i, \lambda_j)} \geq K\right),
 \end{aligned}$$

which goes to zero geometrically fast as $K \rightarrow \infty$. From this it follows that we can choose a suitably large K so that for

$$\frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j} > K,$$

$E[\Delta f | x, l, k]$ can be arbitrarily close to $E[(\Delta f)' | x, l, k]$... This is true since

$$\begin{aligned}
 &|E[\Delta f | x, l, k] - E[(\Delta f)' | x, l, k]| \\
 &\leq |E[(\Delta f)' 1_{U^c} | x, l, k] - E[(\Delta f)' | x, l, k]| + E[(\Delta f)'' 1_{U^c} | x, l, k],
 \end{aligned}$$

and as K increases, U^c tends to an event of measure zero and T has all its moments. Since the above observation is true for all l, k , by choosing a large enough K the expression (2.7) becomes arbitrarily close to

$$\begin{aligned}
 \sum_{l: x_l=0} \lambda_l \sum_k E[(\Delta f)' | x, l, k] \gamma_{kl}(x) &= \sum_{l: x_l=0} \lambda_l \sum_k \left(\frac{1}{\lambda_i} \frac{p_{ki}}{p_{kl}} - \frac{1}{\lambda_j} \frac{p_{kj}}{p_{kl}}\right) \frac{\alpha_k \beta_l p_{kl}}{\lambda_l} \\
 &= \sum_{l: x_l=0} \beta_l \left(\frac{1}{\beta_i} - \frac{1}{\beta_j}\right) < -\epsilon < 0.
 \end{aligned}$$

This last inequality follows from the comments following the definition of the vector β . We also used the equation $\sum_k p_{ki} \alpha_k = \lambda_i / \beta_i$. We have thus so far proved that there is a suitably large K for which

$$\sup\left\{E[\Delta f | x] : \frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j} > K, x_j > 0\right\} < 0. \tag{2.8}$$

Consider now the case where $x_j = 0$ and $x_i / \lambda_i > K$. Then

$$\begin{aligned}
 E[\Delta f | x] &= \sum_{\substack{l: x_l=0, \\ l \neq i, j}} E[\Delta f | x, l] \lambda_l + \sum_{l=i, j} \lambda_l E[\Delta f | x, l] \\
 &= \sum_{\substack{l: x_l=0 \\ l \neq i, j}} \lambda_l \sum_k E[\Delta f | x, l, k] \gamma_{kl}(x) + \sum_{l=i, j} \lambda_l E[\Delta f | x, l].
 \end{aligned}$$

By conditioning on the event U as before, we see that the above expression is

approximately equal (for suitably large K) to:

$$\begin{aligned} & \sum_{\substack{l: x_l=0, \\ l \neq i, j}} \lambda_l \sum_k \left(\frac{1}{\lambda_i} \frac{p_{ki}}{p_{kl}} - \frac{1}{\lambda_j} \frac{p_{kj}}{p_{kl}} \right) \frac{\alpha_k \beta_l p_{kl}}{\lambda_l} + \lambda_i \left(-\frac{1}{\lambda_i} \right) + \lambda_j \sum_k \frac{1}{\lambda_i} \frac{p_{ki}}{p_{kj}} \frac{\alpha_k \beta_j p_{kj}}{\lambda_j} \\ &= \sum_{\substack{l: x_l=0, \\ l \neq i, j}} \beta_l \left(\frac{1}{\beta_i} - \frac{1}{\beta_j} \right) - 1 + \frac{\beta_j}{\beta_i} \\ &\leq -1 + \frac{\beta_j}{\beta_i} = \beta_j \left(\frac{1}{\beta_i} - \frac{1}{\beta_j} \right) < -\varepsilon. \end{aligned}$$

This shows that

$$\sup \left\{ E[\Delta f | x] : \frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j} > K, x_j = 0 \right\} < 0. \tag{2.9}$$

The inequalities (2.8) and (2.9) (together with their corresponding ones when i and j are reversed) show that (2.5) holds and this proves the lemma. \square

Since (2.5) holds one could try to apply Foster’s criterion (see [1]). However this fails for the reason that the set $\{x: f(x) > K\}$ is not finite. Alternatively, if we observe that Δf is bounded above by some geometrically distributed r.v., we can use the results in [3] to get that $E[f_{ij}(x(t))] \leq C$ for all times t , where C is a positive constant. Fix now some i . Since it is always the case that $x_l = 0$ for some l we have

$$E \left| \frac{x_i(t)}{\lambda_i} \right| \leq \sum_{j: j \neq i} |f_{ij}(x(t))| \leq C \cdot n < \infty \quad \text{for all } t$$

which shows (1.1) and consequently that the Markov chain $x(t)$ defined by means of the above strategy is positive recurrent.

b) Necessity

Suppose that there is a strategy (as defined in the introduction) that makes the process $x(t)$, $t = 0, 1, 2, \dots$ stable. This means that (1.1) holds. Let $N_j(t)$, $A_j(t)$ be the total number of parts of type j produced and requested, respectively, by time t . Their difference is certainly equal to the net increase of the stock from time 0 up to time t :

$$N_j(t) - A_j(t) = x_j(t) - x_j(0). \tag{2.10}$$

From (1.1) it follows that

$$\lim_{t \rightarrow \infty} \frac{E x_j(t)}{t} = 0.$$

Taking expectations in (2.10) we see that

$$\lim_{t \rightarrow \infty} \frac{E N_j(t)}{t} = \lambda_j. \tag{2.11}$$

Write now $N_j(t) = \sum_{k=1}^m N_{kj}(t)$ where $N_{kj}(t)$ is the total number of parts of type j produced by time t by machine k . Since $EN_{kj}(t) = p_{kj}\alpha_k(t)$, where $\alpha_k(t)$ is the number of times that machine k has been used by time t , we see, from (2.11), that

$$\lim_{t \rightarrow \infty} \sum_{k=1}^m \frac{\alpha_k(t)}{t} p_k = \lambda.$$

Since

$$\sum_{k=1}^m \frac{\alpha_k(t)}{t} p_k$$

is in Λ for all t and Λ is a closed set, it follows that λ is also in Λ .

It remains to prove that, under the assumption of existence of a stable strategy, λ is in the interior of Λ . Suppose that λ is in the boundary of Λ . Then there exists a hyperplane H that contains λ and separates the n -dimensional space into two parts such that only one of them contains Λ . Consider a vector ν , orthogonal to H , such that

$$\langle \nu, p_k \rangle \geq 0 \quad \text{for all } k = 1, \dots, m,$$

where $\langle \cdot, \cdot \rangle$ denotes inner product. Consider the process

$$y(t) = \langle \nu, x(t) \rangle. \tag{2.12}$$

Let $\mathbf{F}(t)$ be the σ -field generated by $x(s)$, $s = 0, 1, \dots, t$. Let also $a(t)$ denote the arrival at time t . We then have:

$$\begin{aligned} E[y(t+1) - y(t) | \mathbf{F}(t)] &= \sum_{i=1}^n \nu_i E[x_i(t+1) - x_i(t) | \mathbf{F}(t)] \\ &= \sum_{i=1}^n \nu_i \sum_{l=1}^n E[x_i(t+1) - x_i(t) | \mathbf{F}(t), a(t) = l] \lambda_l \\ &= - \sum_{i=1}^n \nu_i \lambda_i 1(x_i(t) > 0) \\ &\quad + \sum_{i=1}^n \nu_i \sum_{l \neq i} E[x_i(t+1) - x_i(t) | \mathbf{F}(t), a(t) = l] \lambda_l. \end{aligned} \tag{2.13}$$

The second term of (2.13) is non-zero only on the set $\{x_i(t) = 0\}$. If we let $\delta_{kl}(t)$ denote the expected number of times machine k is used between t and $t+1$ given the request l at time t and the past $\mathbf{F}(t)$, we see that, on the set $x_i(t) = 0$,

$$E[x_i(t+1) - x_i(t) | \mathbf{F}(t), a(t) = l] = \sum_{k=1}^m p_{ki} \delta_{kl}(t), \quad \text{if } i \neq l, \tag{2.14}$$

and

$$\sum_{k=1}^m p_{ki} \delta_{kl}(t) = 1, \tag{2.15}$$

since the total number of parts of type l produced between t and $t + 1$ is, by definition, equal to 1, if there is a request for a part of type l at time t . Substituting (2.14) into the second term of (2.13) and rearranging terms we find that this term equals

$$\begin{aligned} & \sum_{l=1}^n \sum_{k=1}^m \delta_{kl}(t) 1(x_l(t) = 0) \lambda_l \sum_{i=1}^n \nu_i p_{ki} 1(i \neq l) \\ &= \sum_{l=1}^n \sum_{k=1}^m \delta_{kl}(t) 1(x_l(t) = 0) \lambda_l \{ \langle \nu, p_k \rangle - \nu_l p_{kl} \} \\ &= - \sum_{l=1}^n \nu_l \lambda_l 1(x_l(t) = 0) + \sum_{l=1}^n \sum_{k=1}^m \langle \nu, p_k \rangle \delta_{kl}(t) 1(x_l(t) = 0) \lambda_l, \end{aligned} \quad (2.16)$$

where we used (2.15). Substituting (2.16) into (2.13) we get

$$\begin{aligned} & E[y(t + 1) - y(t) | \mathbf{F}(t)] \\ &= - \langle \nu, \lambda \rangle + \sum_{l=1}^n \sum_{k=1}^m \langle \nu, p_k \rangle \delta_{kl}(t) 1(x_l(t) = 0) \lambda_l. \end{aligned} \quad (2.17)$$

The first term in (2.17) is zero and the second one non-negative, from the definition of the vector ν . It follows that the process y defined in (2.12) is a submartingale. Furthermore, since the strategy is stable, $E | y(t) | \leq \text{const.}$ for all t . Hence $\lim_{t \rightarrow \infty} y(t)$ exists path-wise. But this cannot happen, simply from the way that the process x evolves in time. We therefore conclude that λ cannot be in the boundary of Λ . \square

COROLLARY 2.1

The system is not stabilizable if the number of available machines is strictly smaller than the number of types.

Proof

Indeed, if $m < n$ then the cone Λ will have no interior in \mathbf{R}^n and theorem 2.1 allows us to conclude that the system is not stabilizable.

COROLLARY 2.2

If there is a stable strategy (i.e., if λ is in the interior of Λ) then there is also a stable *stationary* and *Markovian* stable strategy, being a strategy that depends on the past only through the current state and is independent of time.

LOCAL INVENTORIES AND PARTIAL INFORMATION

Assume that the m machines are in different locations. Denote by $x_{ki}(t)$ the number of parts in the inventory at time t that are of type i and that were produced by machine k . The vector $(x_{k1}(t), \dots, x_{kn}(t))$ will be called the local

inventory of machine k at time t . Clearly, if the controller of the system knows the exact states of the local inventories of the machines it can construct a “virtual” global inventory based on the above information and use the strategy described in the previous section in order to control in a stable way the system. It turns out that, under the condition of theorem 2.1, there exists a stable strategy under which the controller does not require complete information about the local inventories of the machines. Instead, the information needed about the local inventory of machine k consists of the relative ordering of the x_{ki} 's, $i = 1, \dots, n$. This strategy has the strange feature that it may produce a requested part even when it can be found in some local inventory.

That strategy is defined as follows. Say that the local inventories are $\{x_{ki}, 1 \leq i \leq n\}$ for $k = 1, 2, \dots, m$. Assume, as before, that the vector λ lies in the interior of the cone formed by the vectors p_1, \dots, p_n . Then there are numbers $\beta_{ki} = \beta_{ki}(x)$, $k = 1, \dots, m$; $i = 1, \dots, n$ defined by the following conditions:

- (i) $\beta_{ki} > 1$, for all i .
- (ii) The vector $(\beta_{k1}(x), \dots, \beta_{kn}(x))$ depends only on the ordering of the vector

$$\left(\frac{x_{k1}}{p_{k1}}, \dots, \frac{x_{kn}}{p_{kn}} \right)$$

in such a way that $\beta_{ki} > \beta_{kj}$ if

$$\frac{x_{ki}}{p_{ki}} < \frac{x_{kj}}{\lambda_{kj}}.$$

In other words, if

$$\left(\frac{x_{k1}}{p_{k1}}, \dots, \frac{x_{kn}}{p_{kn}} \right)$$

is ordered the same way as

$$\left(\frac{x'_{k1}}{p_{k1}}, \dots, \frac{x'_{kn}}{p_{kn}} \right)$$

then $\beta_{ki}(x) = \beta_{ki}(x')$.

(iii) If $\tilde{p}_{kj} = \beta_j p_{kj}$ then the vector λ lies in the interior of the cone formed by the vectors $\tilde{p}_k = (\tilde{p}_{k1}, \dots, \tilde{p}_{kn})$, $k = 1, \dots, m$. In other words, λ can be written as

$$\lambda = \sum_{k=1}^m \alpha_k \tilde{p}_k, \quad \text{where } \alpha_k > 0, \quad \text{for all } k.$$

Note that since the controller knows the relative ordering of the x_{ki} 's it can compute the β_{ki} 's and then finally compute the α_k 's. Suppose now that the controller has a request for a part of type i . He then chooses machine k with probability

$$\gamma_{ki}(x) = \frac{\alpha_k \tilde{p}_{ki}}{\lambda_i}.$$

If $x_{ki} > 0$ the request can be immediately fulfilled and x_{ki} alone is reduced by one unit. If, on the other hand, $x_{ki} = 0$ we run this machine k (even if there may be a k' with $x_{k'i} > 0$) until one part of type i is produced which is used to fulfill the request. This process will simultaneously increase the sizes of the other components of local inventory k , while it will leave the other local inventories unaffected.

The proof of stability of the strategy defined above is similar to the proof of theorem 2.1. The test functions that we now choose are the following:

$$f_{k,ij}(x) = \left| \frac{x_{ki}}{p_{ki}} - \frac{x_{kj}}{p_{kj}} \right|,$$

where k ranges over $\{1, \dots, m\}$ and i, j are distinct indices ranging over $\{1, \dots, n\}$. It is then shown that the drift corresponding to each $f_{k,ij}$ is negative and bounded away from 0 for large values of $f_{k,ij}$

$$\sup_{x: f_{k,ij}(x) > K} E[\Delta f_{k,ij}(x) | x] < 0,$$

and hence by using a similar argument as in theorem 2.1 we get that $E[f_{k,ij}(t)] \leq C$ for all times t , where C is a positive constant. One can also easily see that if the system starts with all inventories being empty it will always be the case that for each k , $x_{kl} = 0$ for some l . This and $E[f_{k,ij}(t)] \leq C$ imply that each local inventory is stable and hence the whole system is stable.

Note that this strategy based on partial information on the sizes of the local inventories is, loosely speaking, less efficient than the one using complete information in that it may produce a part even if it already exists in some local inventory. It is surprising that the above information is sufficient for stabilizing the system since information about the relative availability of parts in the local inventories does not give us any information about the ordering of the numbers of different parts available in the system as a whole.

3. Batch arrivals

DEFINITIONS

This section considers the case where requests arrive in i.i.d. batches distributed as the random vector $N = (N_1, N_2, \dots, N_n)$. The interpretation is that N_i is the number of requests for parts of type i for $i = 1, 2, \dots, n$. As before, these requests have to be met without waiting (again assuming that production of missing parts is virtually instantaneous). The fabrication of parts by machines is modeled as before.

This model raises the possibility of using the information about the set of requests so as to optimize the sequencing of production of parts. However, in this paper, we will only be concerned with the existence of a stable strategy.

STABLE STRATEGIES

We saw, in the previous two sections, that a stable strategy for the model with single arrivals exists if and only if the vector of arrival rates λ is in the cone λ formed by the vectors p_k . Looking at the proof of the necessity of this condition (theorem 2.1) it is not difficult to see that it holds even for the model with batch arrivals if λ_i is defined to be the mean E_{N_i} of the number of requests for parts of type i . The question now is whether this condition is still sufficient. The answer is affirmative and is shown in the following theorem.

THEOREM 3.1

Theorem 2.1 holds for this model, with $\lambda_i := E[N_i]$.

Proof

a) **Sufficiency**

Start by defining the quantities $\beta_i, \tilde{p}_k, \alpha_k$ and $\gamma_{kl}(x)$ as in the proof of theorem 2.1. The strategy that we use can then be described as follows: Given the request vector N and the state of the system x we choose any d for which $x_d = 0$ and we satisfy first the request N_d by running a machine randomly chosen according to $\gamma_{kd}(x)$ until N_d parts of type d have been produced by that same machine. Now we decide on an arbitrary order for fulfilling the remaining $n - 1$ requests. Let \bar{x}_l be the number of parts of type l in the inventory just before we start fulfilling the request for this particular type (hence $\bar{x}_d = x_d = 0$ for the d defined before). Then we take out of the inventory $\min(N_l, \bar{x}_l)$ parts of type l and we run a machine chosen according to $\gamma_{kl}(x)$ until $(N_l - \bar{x}_l)^+$ parts of type l have been produced by that same machine. Observe that this strategy implies that always there is some l for which $x_l = 0$. Also this strategy has the particular feature that the decision taken concerning the machine to be used to fulfill if necessary the request (in the batch) for parts of type l is based on the state x at the time the batch of the requests arrived and once such a machine is selected, the same machine is used to produce all parts of type l needed to fulfill the request.

Next, define the test functions f_{ij} as in (2.3). Our objective is to show that the drift (2.4) satisfies (2.5). First assume that x is such that

$$\frac{x_i}{\lambda_i} - \frac{x_j}{\lambda_j} > K > 0.$$

Let $w^l = (w_1^l, \dots, w_n^l)$, $l = 1, \dots, n$ denote vector of changes to the state x brought by the decision taken in order to satisfy the request for parts of type l . Thus, w_i^l has, conditional on N , mean equal to

$$(N_l - \bar{x}_l)^+ \sum_{k=1}^m \frac{p_{ki}}{p_{kl}} \gamma_{kl}(x) \quad \text{if } l \neq i$$

and is equal to $-\min(N_i, \bar{x}_i)$ if $l=i$. The total change brought to state x is clearly equal to the vector $W = \sum_{l=1}^n w^l$. For given i, j consider the event

$$U = \left\{ \frac{x_i + W_i}{\lambda_i} - \frac{x_j + W_j}{\lambda_j} > 0, x_i - N_i > 0 \right\}.$$

On this event the change of the function $f = f_{ij}$ is equal to

$$\frac{W_i}{\lambda_i} - \frac{W_j}{\lambda_j} \tag{3.1}$$

and the inventory has enough parts in order to fulfill the request for parts of type i . It is not difficult to see that $P(U^c) \rightarrow 0$ as $K \rightarrow \infty$ by observing that each W_i and N_i are dominated, in absolute value, by random variables with finite expectation. This shows that, for suitably large K , the drift of the function f is, conditionally on N , very close to the mean of (3.1) which is equal to:

$$\begin{aligned} & \frac{-\min(N_i, \bar{x}_i)}{\lambda_i} + \frac{1}{\lambda_i} \sum_{l \neq i} (N_l - \bar{x}_l)^+ \sum_{k=1}^m \frac{p_{ki}}{p_{kl}} \gamma_{kl}(x) \\ & + \frac{\min(N_j, \bar{x}_j)}{\lambda_j} - \frac{1}{\lambda_j} \sum_{l \neq j} (N_l - \bar{x}_l)^+ \sum_{k=1}^m \frac{p_{kj}}{p_{kl}} \gamma_{kl}(x). \end{aligned} \tag{3.2}$$

Substituting the value of $\gamma_{kl}(x)$ from (2.2) and taking expectation once more, (3.2) becomes:

$$\begin{aligned} & \frac{-E \min(N_i, \bar{x}_i)}{\lambda_i} + \frac{E(N_j - \bar{x}_j)^+}{\lambda_j} \frac{\beta_j}{\beta_i} + \sum_{l \neq i, j} \frac{E(N_l - \bar{x}_l)^+}{\lambda_l} \frac{\beta_l}{\beta_i} \\ & + \frac{E \min(N_j, \bar{x}_j)}{\lambda_j} - \frac{E(N_i - \bar{x}_i)^+}{\lambda_i} \frac{\beta_i}{\beta_j} - \sum_{l \neq i, j} \frac{E(N_l - \bar{x}_l)^+}{\lambda_l} \frac{\beta_l}{\beta_j}. \end{aligned} \tag{3.3}$$

We have to consider the following two cases. First assume that $x_j > 0$. From this it follows that there is some $d \neq i, j$ for which $x_d = 0$ and hence by the definition of the strategy $\bar{x}_d = x_d = 0$. We can now write (3.3) as follows

$$\begin{aligned} & \frac{-E \min(N_i, \bar{x}_i)}{\lambda_i} + \frac{E(N_j - \bar{x}_j)^+}{\lambda_j} \frac{\beta_j}{\beta_i} + \sum_{l \neq i, j, d} \frac{E(N_l - \bar{x}_l)^+}{\lambda_l} \frac{\beta_l}{\beta_i} + \frac{EN_d}{\lambda_d} \frac{\beta_d}{\beta_i} \\ & + \frac{E \min(N_j, \bar{x}_j)}{\lambda_j} - \frac{E(N_i - \bar{x}_i)^+}{\lambda_i} \frac{\beta_i}{\beta_j} - \sum_{l \neq i, j, d} \frac{E(N_l - \bar{x}_l)^+}{\lambda_l} \frac{\beta_l}{\beta_j} - \frac{EN_d}{\lambda_d} \frac{\beta_d}{\beta_j}. \end{aligned} \tag{3.4}$$

Noting the identity $\min(a, b) + (a - b)^+ = a$ and using the inequality $\beta_j/\beta_i < 1$

we find that (3.4) is

$$\begin{aligned} &\leq \sum_{l \neq i, j, d} \frac{E(N_l - \bar{x}_l)^+}{\lambda_l} \beta_l \left(\frac{1}{\beta_i} - \frac{1}{\beta_j} \right) + \beta_d \left(\frac{1}{\beta_i} - \frac{1}{\beta_j} \right) \\ &\leq \beta_d \left(\frac{1}{\beta_i} - \frac{1}{\beta_j} \right) \leq -\varepsilon < 0. \end{aligned}$$

The second case is $x_j = 0$. Then if there is some other d such that $x_d = 0$ and the requests for parts of type d is served before the requests for type j we can repeat the previous steps to infer that the drift of f is $< -\varepsilon < 0$. If the requests for type j are served first then $\bar{x}_j = x_j = 0$ and since $UN_i < x_i < \bar{x}_i$, (3.3) becomes

$$\begin{aligned} &-\frac{EN_i}{\lambda_i} + \frac{EN_j}{\lambda_j} \frac{\beta_j}{\beta_i} + \sum_{l \neq i, j} \frac{E(N_l - \bar{x}_l)^+}{\lambda_l} \frac{\beta_l}{\beta_i} - \sum_{l \neq i, j} \frac{E(N_l - \bar{x}_l)^+}{\lambda_l} \frac{\beta_l}{\beta_j} \\ &= -1 + \frac{\beta_j}{\beta_i} + \sum_{l \neq i, j} \frac{E(N_l - \bar{x}_l)^+}{\lambda_l} \beta_l \left(\frac{1}{\beta_i} - \frac{1}{\beta_j} \right) \\ &\leq -1 + \frac{\beta_j}{\beta_i} \leq -\varepsilon < 0. \end{aligned}$$

The case

$$\frac{x_j}{\lambda_j} - \frac{x_i}{\lambda_i} > K$$

is symmetric and can be treated as above by interchanging i and j .

b) Necessity

As indicated in the introduction of this section, the proof of necessity is the same as that of theorem 2.1. \square

4. Queued requests

DEFINITIONS

Consider the following situation. There are m machines and n types of parts. When machine k is used, the number x_i of parts of type i changes by a random variable distributed as N_{ki} (and so it becomes $x_i + N_{ki}$), independently of the past evolution. This number is the difference between the number of parts of type i produced by machine k and the number of requests for parts of the same type.

The random variable N_{ki} can therefore take both positive and negative values. Hence the number x_i of parts of type i may, at some given, be positive (corresponding to some excess inventory) or negative (corresponding to a backlog).

A strategy for the system is a non-anticipative, possibly randomized method for choosing a machine at each time t . We want to investigate the existence of stable strategies, i.e., strategies for which $E[x_i(t)]$ is a bounded function of t for all i . A strategy can be *Markovian*, that is it depends on the past only through the current state, or *non-Markovian*. A Markovian strategy is *stationary* if the decision taken at time t is the same for all t . In the latter case only it can be defined by a probability vector

$$\gamma(x) = (\gamma_1(x), \dots, \gamma_m(x))$$

whose k -th component $\gamma_k(x)$ is the probability of choosing machine k when the system is at state $x = (x_1, \dots, x_n)$.

Let N_k denote the vector (N_{k1}, \dots, N_{kn}) , $k = 1, \dots, m$ and let e_k be the mean vector of N_k . It is assumed that the random variables N_{ki} have finite means and finite second moments.

STABLE STRATEGIES

The main result is as follows.

THEOREM 4.1

There exists a stable strategy if and only if the vectors $e_k = (EN_{k1}, \dots, EN_{kn})$, $k = 1, \dots, m$ span \mathbf{R}^n and are positively dependent, i.e., there exist nonnegative numbers α_k , $k = 1, \dots, m$ that are not all zero such that

$$\sum_{k=1}^m \alpha_k e_k = 0. \tag{4.1}$$

Proof

a) Sufficiency

The assumptions about the vectors e_1, \dots, e_m imply that $m \geq n + 1$. Without loss of generality assume $m = n + 1$. (If $m > n + 1$, keep only $n + 1$ out of m vectors that still satisfy the assumptions.) Then the coefficients $\alpha_1, \dots, \alpha_{n+1}$ must necessarily be all strictly positive.

Consider now the Lyapunov (test) function

$$f(x) = \|x\|^2 = \sum_{i=1}^n x_i^2.$$

For this function we first show the following:

Claim

There is a $\gamma(x)$ (that defines a Markov and stationary strategy) such that the drift $d(x) = E(\Delta f | x)$ corresponding to the function f is negative and bounded away from 0 outside a finite set of states S , i.e.,

$$\sup_{x \in S^c} d(x) < 0. \tag{4.2}$$

Proof

To show this, let us first evaluate the drift:

$$\begin{aligned}
 d(x) &= E[\Delta f | x] = \sum_{k=1}^m E(\|x + N_k\|^2 - \|x\|^2) \gamma_k(x) \\
 &= \sum_{k=1}^m \sum_{i=1}^n (2x_i EN_{ki} + EN_{ki}^2) \gamma_k(x) \\
 &= \langle \gamma(x), 2Ax + b \rangle,
 \end{aligned} \tag{4.3}$$

where A is an $m \times n$ matrix whose k -th row is e_k , x is the state (thought as a column vector) and b is a column vector whose k -th element equals $\sum_{i=1}^n EN_{ki}^2$. The symbol $\langle \cdot, \cdot \rangle$ denotes inner product.

Consider the mapping $g: \mathbf{R}^n \rightarrow \mathbf{R}^{n+1}$ defined by

$$g(x) = 2Ax + b.$$

By assumption, A has full rank and so g is one-to-one. Let Π be the range of g . Clearly then, Π is a hyperplane in \mathbf{R}^{n+1} and the mapping $g: \mathbf{R}^n \rightarrow \Pi$ is invertible. Note also that Π satisfies the equation

$$\langle \alpha, y \rangle = c, \quad y \in \mathbf{R}^{n+1},$$

where $\alpha = (\alpha_1, \dots, \alpha_{n+1})$ (the α_k 's are the constants appearing in the statement of this theorem) and $c = \alpha \cdot b = \sum_{i=1}^{n+1} \alpha_i b_i > 0$. Hence α is the normal vector of Π and since all its components are strictly positive, the intersection of Π with the positive orthant \mathbf{R}_+^{n+1} is a bounded set.

Because g is invertible on Π , the drift of f (see (4.3)) can be expressed as a function of $y = g(x)$,

$$d(y) = \langle \tilde{\gamma}(y), y \rangle, \quad y \in \Pi,$$

where $\tilde{\gamma} = \gamma \circ g^{-1}$. Let B be a ball in \mathbf{R}^{n+1} that contains $\Pi \cap \mathbf{R}_+^{n+1}$ in its interior. We claim that there exists $\tilde{\gamma}(y)$ such that

$$\sup_{y \in \Pi \cap B^c} d(y) < 0. \tag{4.4}$$

Write

$$\sup_{y \in \Pi \cap B^c} d(y) = \max_{\sigma} \sup_{y \in \Pi_{\sigma} \cap B^c} d(y)$$

where σ ranges over the $2^{n+1} - 2$ nonempty proper subsets of $\{1, 2, \dots, n + 1\}$ and $\Pi_{\sigma} = \{y \in \Pi: y_i < 0 \text{ if } i \in \sigma, y_i \geq 0 \text{ if } i \in \sigma^c\}$. Let σ_0 be a specific element of σ , say its minimum, and define $\tilde{\gamma}(y)$ as follows:

$$\begin{aligned}
 \tilde{\gamma}(y) &= 1, \quad i = \sigma_0 \\
 &= 0, \quad i \neq \sigma_0, \quad \text{if } y \in \Pi_{\sigma}.
 \end{aligned}$$

Then, on Π_{σ} , $d(y) = y_{\sigma_0}$ and, clearly,

$$\sup_{y \in \Pi_{\sigma} \cap B^c} d(y) < 0.$$

Hence, for this choice of $\tilde{\gamma}$ (4.4) is true. Now let $\gamma(x) = \tilde{\gamma}(g(x))$ and $S = g^{-1}(\Pi \cap B^c)$. Because $\Pi \cap B^c$ is a bounded set and g is one-to-one, the set S is also bounded and therefore the claim (4.2) is true. \square

It can now be seen that the results of [3] are valid here as well. This allows us to conclude that the stationary and Markovian strategy defined above is stable, i.e., (1.1) is true.

b) Necessity

The system can be described by the following recurrence relation:

$$x(t + 1) = x(t) + N_{\xi_t}(t), \quad t = 0, 1, 2, \dots \tag{4.5}$$

where ξ_t is a random index in $\{1, \dots, m\}$ denoting the machine used at time t . Let $F(t)$ be the σ -field generated by $x(s), s = 0, 1, \dots, t$. Let also $\gamma_k(t)$ denote the conditional probability of choosing machine k at time t given the past $F(t)$. Iterating (4.5) we get

$$x(t) = x(0) + \sum_{s=0}^{t-1} N_{\xi_s}(s). \tag{4.6}$$

Since we assume that there exists a stable strategy we have, from 92.1),

$$\lim_{t \rightarrow \infty} \frac{1}{t} Ex(t) = 0.$$

Taking expectations in (4.6) we get

$$\begin{aligned} 0 &= \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^{t-1} EN_{\xi_s}(s) \\ &= \frac{1}{t} \sum_{s=0}^{t-1} EE [N_{\xi_s}(s) | F(s), \xi_s] \\ &= \sum_{k=1}^m e_k \frac{1}{t} \sum_{s=0}^{t-1} \gamma_k(s). \end{aligned} \tag{4.7}$$

If C denotes the closed cone formed by all non-negative linear combinations of the vectors e_k , we have that

$$\sum_{k=1}^m e_k \frac{1}{t} \sum_{s=0}^{t-1} \gamma_k(s)$$

is in C for all t . Therefore so is its limit (4.7). In other words there are nonnegative numbers α_k such that $0 = \sum_{k=1}^m \alpha_k e_k$. From the fact that $\sum_{k=1}^m \gamma_k(s) = 1$ it follows that the α_k 's cannot be all zero. This shows that the e_k 's are positively dependent.

It remains to prove that the e_k 's span R^n . Suppose that they do not. Then there is a non-zero vector ν such that

$$\langle \nu, e_k \rangle = 0, \quad \text{for all } k = 1, 2, \dots, m.$$

Consider then the process

$$y(t) = \langle v, x(t) \rangle.$$

Then it is easy to see that

$$E[y(t+1) - y(t) | \mathbf{F}(t)] = \sum_{k=1}^m \langle v, e_k \rangle \gamma_k(t) = 0.$$

This shows that $y(t)$, $t \geq 0$ is a martingale. It has also bounded expectation, since (1.1) is true. Therefore it converges almost surely, which, as in the proof of theorem 2.1, is a contradiction. \square

COROLLARY 4.1

If the conditions of theorem 4.1 are satisfied then there exists a deterministic stable strategy.

Proof

It follows easily from the proof of theorem 4.1 that the state space can be partitioned into finitely many disjoint regions and with each such region there is associated an index k indicating the machine to be chosen when the state is in that region.

COROLLARY 4.2

The system is not stabilizable if the number of available machines is smaller than or equal to the number of types.

Proof

Indeed, if $m \leq n$ then it can not be simultaneously true that the vectors e_1, \dots, e_m span \mathbf{R}^n and $\sum_{k=1}^m \alpha_k e_k = 0$.

5. Conclusions

The stability of an uncertain production system has been analyzed. The issue is the existence of stabilizing strategies.

References

- [1] S. Asmussen, *Applied Probability and Queues* (J. Wiley, 1987).
- [2] C. Courcoubetis and R. Weber, Necessary and sufficient conditions for stability of a bin-packing system, *J. Appl. Prob.* 23 (1986) 989–999.
- [3] B. Hajek, Hitting-time and occupation-time bounds implied by drift analysis with applications, *Adv. Appl. Prob.* 14 (1986) 502–525.
- [4] R. Tweedie, Sufficient conditions for ergodicity and recurrence of Markov chains on a general state space, *Stoch. Proc. and Appl.* 3 (1975) 385–403.