

STABILITY OF FLEXIBLE MANUFACTURING SYSTEMS

COSTAS COURCOUBETIS

University of Crete, and Institute of Computer Science, Forth, Heraklion, Greece

RICHARD WEBER

University of Cambridge, Cambridge, United Kingdom

(Received October 1989; revisions received August 1990, June 1991, July 1992; accepted September 1992)

We consider a flexible manufacturing facility that can be operated in any of m different modes. While running in mode k certain intermediate products are consumed and other intermediate or finished products are created. There may be variability in the manufacturing process, as well as random arrivals of raw materials and orders for finished products. We establish conditions that ensure demands can be satisfied while maintaining bounded levels of inventories. These results that may be viewed as generalizing to the flexible manufacturing context the notion of stability for a queueing system.

The aim of a manufacturing system is to transform raw materials, parts, and subassemblies into other parts, subassemblies, and finished products. In this paper, we consider a model of a flexible manufacturing facility (FMS) in which production can take place in any of m different modes or configurations. There is no cost or delay in changing modes. When production takes place in mode k certain intermediate products are consumed and other intermediate or finished products are created. We suppose that in total the system is concerned with n different products. Here we use the word *product* to refer to any unit of order, raw material, component, subassembly, work in progress, or finished product. Suppose that running production for one period in mode k changes the inventory level of product i by $S_{k,i}$ units. It is a key idea in this paper that $S_{k,i}$ may be a random variable. This reflects the fact that unpredictable variations in both the demand and production processes are usually major sources of difficulty in managing a manufacturing facility. It is often the case that inventories of intermediate products or raw materials are required to cope with these variations. However, there is usually a cost to carrying inventories and it is desirable that inventories remain small, even zero. Other researchers have considered such questions. For example, Kimemia and Gershwin (1983) and Gershwin, Akella and Choong (1985) considered the problem of how parts should be dispatched into an FMS to meet production requirements when machines are

unreliable and subject to breakdown. They imagine that production surplus is penalized by a quadratic cost function and propose a hierarchical policy for determining a near-optimal part-loading scheme. Bielecki and Kumar (1988) also considered a model with unreliable manufacture and found that even with uncertainty it may sometimes be optimal to aim for zero inventory.

In this paper, we do not try to model a cost function or determine an optimal policy. Rather, we address a more basic question: Given certain demands, production capacity and flexibility, is it possible to operate the system so that inventory levels are kept under control? Can we ensure that the inventory level is not arbitrarily large for long periods of time? The reader might view this paper as generalizing to the context of an FMS the notion that is familiar in queueing theory: That the mean queue length in a single-server queue is finite if the traffic intensity is less than 1. The corresponding notion is less obvious and more complex in the manufacturing context. This is because there are many ways in which an FMS can be operated; how this is done will affect the rate at which existing inventories are depleted and the rate at which finished products are produced in response to a random demand.

By allowing $S_{k,i}$ to be a random variable we model the random arrival of raw materials and orders for finished products. We also model the variability in manufacture that can arise because of variability in

Subject classifications: Probability, stochastic model applications: stabilizing a flexible manufacturing system. Production/scheduling, flexible manufacturing: stabilizing control of a manufacturing model.

Area of review: STOCHASTIC PROCESSES AND THEIR APPLICATIONS.

processing times or in the output itself. For example, in manufacturing crystal wafers that are used in mobile radio sets, one can set up the system for manufacture of wafers of certain frequency, but because of variability in the process wafers with a range of frequencies around the target value will be produced. Those that are not within tolerance of the target value may be held in stock and used to meet future demands for other frequencies. A similar example arises in fabricating microprocessors; each chip is tested and different chips are certified for use at different clock speed. It can also happen that some output must be scrapped.

Let the integer-valued random variable $x_i(t)$ denote the number of units of product i that are in the inventory at time t . This 'surplus production' satisfies

$$x_i(t+1) = x_i(t) + S_{k,i}. \quad (1)$$

If $S_{k,i}$ is negative, then manufacturing in mode k consumes units of product i . This means that if $x_i(t)$ is negative, then the backlog of product i is increased. A model in which $x_i(t)$ may be negative is appropriate if it is possible to borrow units of product i from another factory and then manufacture replacement units at some later time. In this case, we think of $x_i(t)$ as the surplus of cumulative production of product i over cumulative demand at time t and allow a negative surplus. Note that randomly arriving orders for a finished product of type i can be modeled by letting $S_{k,i}$ equal the number of units of product i created by manufacturing in mode k for one period, minus the number of new orders for that item which are received (a number that is independent of k). In Section 3 we also consider a model in which backlogging is not allowed.

It is convenient to imagine that at each discrete time t ($t = 0, 1, \dots$) one mode of manufacture must be selected and production take place in that mode for the next time period. We will discuss two requirements that are necessary if costs of carrying inventory are to be kept finite. The first requirement is that $P(\sum_{i=1}^n |x_i(t)| \rightarrow \infty) = 0$. If this can be achieved by operating some policy, we say the system is *weakly-stabilizable*. If the inventory process converges to a stationary regime, then this implies that $x(t)$ has a well-defined distribution. A stronger requirement is that the expected value of $\sum_{i=1}^n |x_i(t)|$ be bounded by a constant, uniformly in t . If this is possible we say that the system is *stabilizable*. Since, as observed above, some negative components of $x(t)$ may denote the backlogs of unfilled orders for finished products of various types, these requirements are also certainly necessary if the manufacturing capacity is to be

sufficient to meet all the demand placed upon it. Our aim is to give easily verifiable conditions under which the above model of a flexible manufacturing system is stabilizable or weakly-stabilizable.

A similar stability question for a deterministic FMS model was studied by Perkins and Kumar (1989). They call a scheduling policy stable if the cumulative production under that policy lags the cumulative demand by no more than a constant, and they analyze the stability of a number of policies. This sort of stability is interesting because it guarantees that finite buffers are sufficient to operate the system. For stochastic models, finite buffers will not suffice in general but the desire is that the average work in progress should be small. There are results for on-line bin packing problems that are similar to ones in this paper. One can think of n types of items that randomly arrive to a bin packing system and can be packed into bins in various ways. For example, a bin might be fully packed either by one item of each of types 1, 2 and 4, or by five items of type 3. The idea is to pack items in bins as they arrive, and ensure that the time-average number of partially full bins remains bounded. This is similar to the flexible manufacturing problem if we think of the spaces for items as arising from a production process and the demand for these spaces as being generated by randomly arriving orders. In Courcoubetis and Weber (1986a) we gave a necessary and sufficient condition for a bin packing system to be stabilizable when items arrive according to independent Bernoulli processes. In Courcoubetis and Weber (1986b) this was more fully developed for a variety of packing configurations. Courcoubetis and Rothblum (1991) considered models in which rewards are obtained when different packing configurations used. Rhee and Talagrand (1988) considered a model in which sizes of items are uniformly distributed over a subinterval $[a,b]$ of $[0,1]$ and a bin can contain any number of items whose sizes add to no more than 1. They characterized the choices of $[a,b]$ for which the wasted space in partially full bins can be held finite in time average by using an off-line packing algorithm. Courcoubetis and Weber (1990) considered stabilizing systems when there are side constraints on the frequencies with which bins may be packed in different ways. A production model with batch arrivals was considered by Courcoubetis et al. (1989).

In the present paper we consider a flexible manufacturing model, but we assume far less than in other papers about the stochastic processes that describe production and demands. In previous papers, the policies that were used to stabilize systems relied on artificially complicated representations of the state of

the system and required randomized actions. The policies constructed for the proofs in this paper are simpler. Section 1 presents an algorithm for stabilizing a model of a flexible manufacturing system when the changes in inventory levels resulting from operating the system in various modes are independent from period to period. In Sections 2–4 we consider more general processes and present sufficient conditions for systems to be weakly-stabilizable or stabilizable. The necessity of these conditions is addressed in Section 5.

1. STABILITY OF A SIMPLE MANUFACTURING SYSTEM

Consider a model in which the changes in inventory due to running the system are independent from period to period. Specifically, consider a *Bernoulli model* in which $S_{k,i}$ is a random variable, with mean $\theta_{k,i}$, that in each period is distributed independently of all other random variables in the model. Let $S_k = (S_{k,1}, \dots, S_{k,n})$ and $\theta_k = (\theta_{k,1}, \dots, \theta_{k,n})$. We assume that the tails of the distribution of $S_{k,i}$ are exponentially bounded, which means that there is a $\omega > 0$ such that $P(|S_{k,i}| > a) \leq e^{-\omega a}$ for all a sufficiently large. This assumption will be met in any practical setting because there will be some absolute bound on how much inventories levels can change during one period of manufacture. The assumption is used in the proof of Theorem 1, in which the expectation of the Lyapunov function $f(x(t)) = \|x(t)\|$ is shown to be uniformly bounded for all t . The proof uses a criterion due to Hajek (1982), which states that if the increment $f(x(t + 1)) - f(x(t))$ is bounded independently of $x(t)$ by a random variable whose distribution has an exponentially bounded tail, and there exists $\delta > 0$ such that for all x outside a finite set the drift $E[f(x(t + 1)) - f(x(t)) | x(t) = x]$ is less than $-\delta$, then the stochastic process $f(x)$ is stable in the sense that there exists a $B < \infty$ such that $E[f(x(t))] < B$ for all t .

Theorem 1. *Suppose that the convex cone generated by the vectors $\theta_k, k = 1, \dots, m$, is the whole of \mathcal{R}^n . Then the system is stabilizable.*

One can interpret the theorem as saying that the system is stabilizable if given any inventory position, say x , it is possible to manufacture according to some mix of the possible modes such that after a time t , the expected level of all inventories will be reduced to zero, i.e., $x + t \sum_{k=1}^m \alpha_k \theta_k = 0$, where $\sum_{k=1}^m \alpha_k = 1$. It is easy to imagine circumstances in which the θ_k s do not span all of \mathcal{R}^n . Consider, for example, a chemical

plant that is used to make two final products. Each unit that is manufactured of product 1 takes one period and results in 2 units of a certain by-product 3. One unit of by-product 3 is used to make each unit of product 2, which also takes one period. Suppose that the demand arising during a period is 0.5 unit of each product. Here we would have $\theta_1 = (0.5, -0.5, 2)$ and $\theta_2 = (-0.5, 0.5, -1)$. Clearly, it is not possible to operate the facility without scrapping some by-product or failing to meet all the demand for product 1.

A small numerical example may further help the reader to understand the theorem. Suppose that $n = 2$ and

- S_1 is equally likely to be $(1, 1)$ or $(-1, -1)$;
- S_2 is equally likely to be $(-1, -1), (-1, 5), (5, -1)$ or $(5, 5)$;
- S_3 is identically distributed to $-S_2$;
- S_4 is equally likely to be $(-5, 5)$ or $(3, -3)$;
- S_5 is identically distributed to $-S_4$.

Then $\theta_1 = (0, 0), \theta_2 = -\theta_3 = (2, 2), \theta_4 = -\theta_5 = (-1, 1)$. For this example, we see that although production mode 1 has $\theta_1 = (0, 0)$, using this mode alone cannot stabilize the system. If only mode 1 is employed, $x_i(t)$ follows a symmetric random walk and $E[|x_i(t)|] \rightarrow \infty$.

There are sensible rules-of-thumb that might be applied in practice. One such rule is to find at each time t that i for which $|x_i(t)|$ is greatest and then run for one period the production mode that minimizes $E[|x_i(t + 1)|]$. However, as we will see in Section 5, this rule cannot stabilize the system, because it will only use modes 2 and 3, and the cone generated by θ_2 and θ_3 is not all of \mathcal{R}^2 . An alternative rule-of-thumb is to run at time t a production mode that minimizes $E[\sum_i |x_i(t + 1)|]$. However, it is difficult to see whether or not such a rule-of-thumb will stabilize the system. The value of Theorem 1 is that it tells us when there exists a policy that will stabilize the system and whether a certain set of manufacturing modes is sufficient. In the example above, a stabilizing policy can be constructed using modes 2–5. It must use all these modes, because the cone generated by any strict subset of $\theta_2, \dots, \theta_5$ would not be \mathcal{R}^2 .

We will prove Theorem 1 by constructing a policy that stabilizes the system. The simplest way to do this will be to use a randomized policy that chooses a mix of modes that will reduce the expected levels of the inventories to 0 in a certain time. In Section 2 we will show that a nonrandomized policy could also stabilize the system.

Proof of Theorem 1. Take as a Lyapunov function the L_2 norm $f(x) = \|x\| = (\sum x_i^2)^{1/2}$. We prove stability by verifying Hajek's criterion for $f(x)$.

Suppose that $x(t) = x$ and consider the drift in $[f(x)]^2 = \|x(t)\|^2$ over one time period when the system is operated in mode k for period t . Let

$$\begin{aligned} d_k &= E(\|x + S_k\|^2) - \|x\|^2 \\ &= 2 \sum_{i=1}^n \theta_{k,i} x_i + \sum_{i=1}^n E(S_{k,i}^2). \end{aligned}$$

Now by the hypothesis of the theorem there exists a scalar $\delta(x) > 0$ and probabilities $\alpha_k, \sum_k \alpha_k = 1$, such that

$$\sum_{k=1}^m \alpha_k \theta_{k,i} = -2\delta(x)\text{sign}(x_i), \quad i = 1, \dots, n,$$

where $\text{sign}(x)$ is defined to be -1 or 1 as x is or is not negative. Since $\delta(x)$ depends on x only through the values of $\text{sign}(x_i)$ we may assume that there are only finitely many $\delta(x)$. Let δ be the minimum such $\delta(x)$. If we adopt a policy that chooses mode k with probability α_k , then the drift in $f(x)^2$ is

$$\begin{aligned} \sum_{k=1}^m \alpha_k d_k &= -4\delta(x) \sum_{i=1}^n |x_i| + \sum_{k=1}^m \sum_{i=1}^n \alpha_k E(S_{k,i}^2) \\ &\leq -4\delta(x)\|x\| + \sum_{k=1}^m \sum_{i=1}^n \alpha_k E(S_{k,i}^2) \\ &\leq -2\delta(x)\|x\| \leq -2\delta\|x\|, \end{aligned}$$

where the second inequality holds for all $\|x\|$ sufficiently large. It follows that for such x

$$\begin{aligned} [E(\|x(t+1)\|)]^2 &\leq E(\|x(t+1)\|^2) \\ &\leq \|x\|^2 - 2\delta\|x\| + \delta^2, \end{aligned}$$

where the expectations above are over the distribution of $x(t+1)$ conditional on $x(t) = x$. Taking square roots we have $E(\|x(t+1)\|) \leq \|x\| - \delta$. This verifies condition b of Hajek's criterion. To verify condition a note that the increment satisfies

$$\|x(t+1)\| - \|x\| = \|x + \Delta\| - \|x\| \leq \|\Delta\|,$$

where Δ is a random variable and $\|\Delta\|$ is stochastically smaller than $\sum_{k=1}^m \|S_k\|^2$. Thus

$$\begin{aligned} P(\|x(t+1)\| - \|x(t)\| > t|x(t) = x) &\leq P\left(\sum_{k=1}^m \|S_k\| > t\right) \\ &\leq \sum_{k=1}^m P(\|S_k\| > t/m) \\ &\leq \sum_{k=1}^m \sum_{i=1}^n P(|S_{k,i}| > t/mn), \end{aligned}$$

and because we have assumed that the distribution of $S_{k,i}$ has exponentially bounded tails, it follows that the same is true for the increment in $f(x(t))$. Hence, both a and b of Hajek's criteria are satisfied and $E[f(x(t))]$ is therefore bounded uniformly in t .

Notice that the above proof does not use the fact that $x(t)$ is Markov. Hence, the assumptions about the independence of the $S_{k,i}$ s could be relaxed and the theorem would remain valid.

2. STABILITY OF A SYSTEM WITH GENERAL ORDERS AND MANUFACTURING STATISTICS

In this section, we consider a model in which the demand and manufacturing processes are more general than the Bernoulli processes considered above. Thus, as might occur in practice because of variation in raw materials or machine operators, there may be correlation between successive periods of operation in a given mode. Suppose that on the t th occasion that the system is operated in mode k the change in the inventory of product i is $S_{k,i}(t)$. Let $N_{k,i}(t) = \sum_{s=1}^t S_{k,i}(s)$ denote the cumulative change in the inventory of product i due to the first t periods of operation in mode k . Assume that the processes $N_{k,i}(t)$ are independent. Let $N_k(t)$ be the vector $(N_{k,1}(t), \dots, N_{k,n}(t))$. Let $\mathcal{F}_k(s)$ be the σ -field generated by the entire history of production in mode k up to the time mode k has been used for the s th time. We will prove the system is weakly-stable under the following assumption about the processes $N_{k,i}(t)$.

Assumption A. *There is a function $g(t)$ such that for all k, i, s and $t > 0$,*

$$E(|N_{k,i}(s+t) - N_{k,i}(s) - t\theta_{k,i}| | \mathcal{F}_k(s)) \leq g(t),$$

where $g(t)/t \rightarrow 0$ as $t \rightarrow \infty$.

This weak assumption simply states that the change to the inventory occasioned by t successive periods of operation in mode k will have an absolute deviation from its mean whose expected value grows more slowly than linearly in t . Note that assumption A implies that there exists a $B < \infty$ such that $E(|N_{k,i}(s+t) - N_{k,i}(s) - t\theta_{k,i}|) \leq Bt$ for all k, i, s and t . One special case in which assumption A is satisfied is when $S_{k,i}(t)$ is an ergodic Markov process for each k, i . Another is when $S_{k,i}(t)$ is a second-order stationary process. Various kinds of renewal process also satisfy the assumption. The following theorem states a general condition on the demand and production processes under which the system can be

weakly-stabilized. The intuitive interpretation is the same as we gave following the statement of Theorem 1.

Theorem 2. *Consider the model described in the above paragraph. Suppose that orders can be backlogged, assumption A holds, and the convex cone generated by $\theta_1, \dots, \theta_m$ is the whole of \mathbb{R}^n . Then the system is weakly-stabilizable.*

Proof. Suppose that at time s the system has been operated in mode k for s_k periods, $s_1 + \dots + s_m = s$. The inventory level is $x(s) = \sum_k N_k(s_k)$. As before, let $\|x(s)\|$ be the L_2 norm. We will prove that the system is weakly-stabilizable by showing that there exists an increasing sequence of random times $\{\sigma_i\}$, tending to infinity, such that both $E(\sigma_{i+1} - \sigma_i)$ and $E(\|x(\sigma_i)\|)$ are uniformly bounded for all i . Clearly, the second of these ensures that $P(\|x(t)\| \rightarrow \infty) = 0$. Let H be the convex hull in \mathbb{R}^n of the points $\theta_1, \dots, \theta_m$. The origin must be in the interior of H , because as $\theta_1, \dots, \theta_m$ are generators for \mathbb{R}^n there must be points in H in every direction from the origin. Let δ_1 be chosen so that the ball of radius δ_1 about the origin lies within H . A sufficiently small δ_1 can be calculated by finding a nonnegative $m \times 2n$ matrix α such that $(\theta_1 | \dots | \theta_m) \alpha = (I_n | -I_n)$ and taking $\delta_1 = 1/m \max_j \{\sum_{k=1}^{2n} \alpha_{jk}\}$. If x is within a ball of radius δ_1 , then $x = (\theta_1 | \dots | \theta_m) \alpha'$, where $\alpha' = \alpha(x^+ | -x^-)^T$ and $0 \leq \alpha'_j \leq 1/m$ for all $j = 1, \dots, m$. Also, define $\delta_2 = \sum_{k=1}^m \|\theta_k\|$. Let $D = \{x : \|x\| \leq L\}$, where $L > m\delta_1$ and suppose that $x(s)$ is not in D . Choose nonnegative μ_1, \dots, μ_m to minimize $\sum_{k=1}^m \mu_k$ subject to $\sum_{k=1}^m \mu_k \theta_k = -x(s)$. Then $-x(s)/(\sum_{k=1}^m \mu_k)$ lies on the boundary of H , and hence $\sum_{k=1}^m \mu_k \leq \|x(s)\|/\delta_1$. Now defining $t_k = \lceil \mu_k \rceil$, we have

$$\left\| x(s) + \sum_{k=1}^m t_k \theta_k \right\| = \left\| \sum_{k=1}^m (t_k - \mu_k) \theta_k \right\| \leq \delta_2. \quad (2)$$

We can also bound $t = \sum_{k=1}^m t_k$ in terms of $\|x(s)\|$ as

$$t = \sum_{k=1}^m t_k \leq m + \sum_{k=1}^m \mu_k \leq 2\|x(s)\|/\delta_1. \quad (3)$$

The second inequality follows from the fact that $x(s)$ is not in D and the definition of the μ_k s. We also have

$$\begin{aligned} \|x(s)\| &= \left\| \sum_{k=1}^m \mu_k \theta_k \right\| \leq \sum_{k=1}^m \mu_k \|\theta_k\| \\ &\leq \sum_{k=1}^m t \|\theta_k\| = t \delta_2 \end{aligned} \quad (4)$$

so

$$\|x(s)\|/\delta_2 \leq t \leq 2\|x(s)\|/\delta_1. \quad (5)$$

Next we define the sequence of observation times $\{\sigma_0, \sigma_1, \dots\}$ with $\sigma_0 = 0$. If $x(\sigma_i)$ is in D run the system for one period in any mode and let $\sigma_{i+1} = \sigma_i + 1$. If $x(\sigma_i)$ is not in D run the system for $\tau_i = t$ periods with t_k of these periods of manufacture in mode k , where the t_k s and t are calculated as above for $s = \sigma_i$. In this case, let $\sigma_{i+1} = \sigma_i + \tau_i$ and note that for $s = \sigma_i$ and $t = \tau_i$ we have that at time $s + t$, conditional on all the history up to time s ,

$$\begin{aligned} E(\|x(t+s)\|) &= E\left(\left\| x(s) + \sum_{k=1}^m (N_k(s_k+t_k) - N_k(s_k)) \right\|\right) \\ &= E\left(\left\| \left(x(s) + \sum_{k=1}^m \theta_k t_k \right) + \sum_{k=1}^m (N_k(s_k+t_k) - N_k(s_k) - \theta_k t_k) \right\|\right) \\ &\leq \left\| x(s) + \sum_{k=1}^m \theta_k t_k \right\| + \sum_{k=1}^m E\left(\left\| N_k(s_k+t_k) - N_k(s_k) - \theta_k t_k \right\|\right) \\ &= \left\| x(s) + \sum_{k=1}^m \theta_k t_k \right\| + \zeta(t) \leq \delta_2 + \zeta(t), \end{aligned}$$

where because each t_k is less than t it follows that $\zeta(t)$ is $o(t)$.

It follows from the left-hand inequality in (5) that when $\|x(s)\|$ is sufficiently large t will be large enough to imply $\zeta(t)/t \leq \delta_1/2$; suppose that L is large enough to ensure this. Using the right-hand inequality in (5) we have $\|x(s)\|/(t\delta_1) \geq 1$ and thus

$$\begin{aligned} E(\|x(t+s)\|) &\leq \delta_2 + \zeta(t)\|x(s)\|/(t\delta_1) \\ &\leq \delta_2 + (1/2)\|x(s)\| \leq \rho\|x(s)\| \end{aligned}$$

for some $\rho < 1$ when $\|x(s)\|$ is sufficiently large; again, suppose that L has been chosen large enough.

From the above discussion we have that if $x(\sigma_i)$ is not in D , then $E(\|x(\sigma_{i+1})\|) < \rho\|x(\sigma_i)\|$. On the other hand if $x(\sigma_i)$ is in D , then $E(\|x(\sigma_{i+1})\|)$ is uniformly bounded by $L + nB + \delta_2$, where B is as defined in the paragraph following the statement of assumption A. So for all observation times

$$E(\|x(\sigma_{i+1})\|) \leq \rho E(\|x(\sigma_i)\|) + L + nB + \delta_2. \quad (6)$$

Thus, if $\|x(\sigma_0)\| \leq (L + nB + \delta_2)/(1 - \rho) = M$, then (6) implies $E(\|x(\sigma_i)\|) \leq M$ for all i . This fact is sufficient to ensure that the system is weakly-stabilizable.

A concluding remark is that $\delta_1 t \leq \|x(\sigma_i)\|$ implies $(\sigma_{i+1} - \sigma_i) \leq \|x(\sigma_i)\|/\delta_1$. One might think that by

an application of assumption **A** it would be possible to bound $\|x(u)\|$ in the range $\sigma_i < u < \sigma_{i+1}$ by $\|x(\sigma_i)\| + (nB + \delta_2)(\sigma_{i+1} - \sigma_i)$. If this were valid it would imply $E(\|x(u)\|) \leq M(1 + (nB + \delta_2)/\delta_1)$ and the system would be stabilizable. However, this argument is not valid because in applying assumption **A** it is not correct to replace t by the random variable $\sigma_{i+1} - \sigma_i$.

To guarantee that a system is stabilizable an assumption is needed to ensure that there is little correlation between periods of operation that are not close together in time; if it holds then the average inventory level can be kept finite. Let $U_{k,i}(t) = S_{k,i}(t) - E[S_{k,i}(t)]$.

Assumption B. *There exists a $C > 0$ such that for all k, l, s_1, s_2, t_1, t_2 ,*

$$E\left(\left\|\sum_{u_1=s_1+1}^{s_1+t_1} U_k(u_1)\right\| \left\|\sum_{u_2=s_2+1}^{s_2+t_2} U_l(u_2)\right\|\right) \leq C(t_1 + t_2).$$

Assumption **B** holds if the variance of $S_{k,i}(t)$ is uniformly bounded and the covariance (or autocovariance when $k = l$) between $S_{k,i}(s)$ and $S_{l,j}(t)$ tends to zero sufficiently fast as $|t - s| \rightarrow \infty$. For example, it is not hard to see that the assumption holds when the $S_{k,i}(t)$ s are independent, irreducible, aperiodic Markov chains.

Theorem 3. *Suppose that assumption **B** and the conditions of Theorem 2 hold. Assume also that the production processes are such that by using the policy described in the proof of Theorem 2 the process $x(t)$ is stationary. Then the system is stabilizable.*

Proof. The stationarity assumption implies that both $\lim_{t \rightarrow \infty} E\|x(t)\|$ and $\lim_{t \rightarrow \infty} (1/t)E[\sum_{s=0}^{t-1} \|x(s)\|]$ exist and have the same value. Using the observations times $\{\sigma_i\}$ constructed in the proof of Theorem 2, we have

$$\begin{aligned} E\left[\sum_{s=0}^{t-1} \|x(s)\|\right] &\leq E\left[\sum_{i:\sigma_i < t} \sum_{s=\sigma_i}^{\sigma_{i+1}} \|x(s)\|\right] \\ &\leq \sum_{i:\sigma_i < t} E\left[E\left(\sum_{s=\sigma_i}^{\sigma_{i+1}} \|x(s)\| \mid \sigma_i, x(\sigma_i)\right)\right] \\ &\leq \sum_{i=0}^{t-1} E\left[E\left(\sum_{s=\sigma_i}^{\sigma_{i+1}} \|x(s)\| \mid \sigma_i, x(\sigma_i)\right)\right]. \end{aligned}$$

Now conditional on knowing σ_i and $x(\sigma_i)$, the value of σ_{i+1} is determined by the construction described in the proof of Theorem 2. So from time σ_i to time σ_{i+1} the value of $E(\|x(s)\|)$ can increase by at most $(nB + \delta_2)(\sigma_{i+1} - \sigma_i)$ and the above is

$$\begin{aligned} &\leq \sum_{i=0}^{t-1} E[(\sigma_{i+1} - \sigma_i)(\|x(\sigma_i)\| \\ &\quad + (nB + \delta_2)(\sigma_{i+1} - \sigma_i))] \\ &\leq \sum_{i=0}^{t-1} E[\|x(\sigma_i)\|^2(1/\delta_1 + (nB + \delta_2)/\delta_1^2)], \end{aligned}$$

where the final inequality follows from $(\sigma_{i+1} - \sigma_i) \leq \|x(\sigma_i)\|/\delta_1$. Hence, the system will be stabilizable if $E\|x(\sigma_i)\|^2$ is uniformly bounded for all i . To see this is the case, consider

$$E(\|x(\sigma_i)\|^2) \leq (E\|x(\sigma_i)\|)^2 + E(\|x(\sigma_i) - E[x(\sigma_i)]\|^2). \tag{7}$$

The first term on the right-hand side is uniformly bounded by Theorem 2. Let s_k be the number of times production mode k has been run by time σ_{i-1} , and let t_k be the number of times mode k is run between times σ_{i-1} and σ_i . The second term on the right-hand side of (4) is

$$\begin{aligned} &E\left[E\left(\left\|\sum_{k=1}^m \sum_{t=s_k}^{s_k+t_k} U_k(t)\right\|^2 \mid x(\sigma_{i-1})\right)\right] \\ &\leq E\left[E\left(\sum_{k=1}^m \sum_{l=1}^m \left\|\sum_{t=s_k}^{s_k+t_k} U_k(t)\right\| \left\|\sum_{t=s_l}^{s_l+t_l} U_l(t)\right\| \mid x(\sigma_{i-1})\right)\right] \\ &\leq CE\left[\sum_{k=1}^m \sum_{l=1}^m (t_k + t_l)\right] \\ &\leq m^2 CE(\sigma_i - \sigma_{i-1}) \leq m^2 CE\|x(\sigma_{i-1})\|/\delta_1, \end{aligned}$$

which is uniformly bounded by Theorem 2.

It is an assumption of Theorem 3 that $x(t)$ is stationary under the policy described in the proof of Theorem 2. This holds, for example, if the $S_{k,i}(t)$ s are independent, irreducible, aperiodic Markov chains. In this case, Theorem 1 is covered by Theorem 3. Observe that the stabilizing policy in Theorem 3 is not randomized and is based on simple information about the state of the system.

3. STABILITY WITHOUT BACKLOGGING

In this section, we suppose that production is not allowed to fall behind orders for finished products. To do this we separate the notions of adding and subtracting from inventories. The system is said to be in a *producing mode* when any of modes $r + 1, \dots, m$ are used. Such modes only add to the inventories: That is, $S_{k,i}(t)$ is nonnegative for each of these modes. The system is said to be in a *request mode* when one of modes $1, \dots, r$, is used. In such a mode

$S_{k,i}(t)$ is nonpositive. If the requests arising from running the system in a request mode for one time period can be met from inventory, then this is done and the system is next run in a request mode without any further production. However, if requests arise for products which are not in the inventory then before allowing further requests the system is run in producing modes until all outstanding requests are satisfied. Only after satisfying all requests that arise from one period's operation in a request mode do we consider requests arising from the next period of operation in a request mode. For example, we might imagine there to be one period of a request mode at the beginning of each week. If the requests cannot be met from the inventory, then during the week the system is run in a producing mode, for as many periods as are needed, until all requests are satisfied. Notice that we do not allow orders to be rejected, but rather assume that arbitrarily many periods of production may be inserted between two periods' operation in a request mode. In practice, there will be an upper limit on the rate of production and this will be reflected in the lead time. However, for fairly regular production and request processes it may be possible to guarantee that one period of requests can be met by, at most, a certain number of periods of production, or that a finite amount of buffer inventory will be sufficient to ensure zero lead time. As before, there is a choice among several producing and request modes. The choice among request modes might be realized by a marketing strategy that takes account of present inventory levels. There are two reasons why it is more difficult to stabilize the model just described. First, the choice between modes of operation is more restricted: Request and producing modes may not be interleaved arbitrarily. Second, during periods that the system runs in a producing mode to meet requests many unwanted products may be added to the inventories.

Prior to stating Theorem 4 we observe the following consequence of assumption A. Suppose that mode k has been run s times. Then there exists a constant T' , that is independent of s and the history up to the s th time that mode k has been run, such that for all $k, i, \theta_{k,i} > 0$, the expected number of periods for which it is required to run mode k until one more product of type i is produced is less than T' . Again, this is a condition that is certainly fulfilled in practical circumstances because it simply says that if mode k can produce products of type k , then it only takes a finite time to do so on average. We see this because if $t\theta_{k,i} > 1$,

$$\begin{aligned} P(N_{k,i}(s+t) - N_{k,i}(s) < 1 | \mathcal{F}_k(s)) &= P(N_{k,i}(s+t) - N_{k,i}(s) - t\theta_{k,i} < 1 \\ &\quad - t\theta_{k,i}\mathcal{F}_k(s)) \\ &= P((N_{k,i}(s+t) - N_{k,i}(s) - t\theta_{k,i}) - \geq t\theta_{k,i} \\ &\quad - 1 | \mathcal{F}_k(s)) \\ &\leq E((N_{k,i}(s+t) - N_{k,i}(s) \\ &\quad - t\theta_{k,i})^- | \mathcal{F}_k(s)) / (t\theta_{k,i} - 1) \end{aligned}$$

By assumption A the final term is less than 1/2 for all t greater than some integer T sufficiently large. Thus, for any history up to the s th time that mode k has been run the probability of producing at least one product of type i during the next T periods of use of mode k is at least 1/2. That is,

$$P(N_{k,i}(s+T) - N_{k,i}(s) \geq 1 | \mathcal{F}_k(s)) \geq 1/2.$$

Since this holds for $s = u, u + T, u + 2T, \dots$, we see that the expected time until one product of type i is produced is less than $T' = 2T$. Clearly, we can take T large enough that this holds uniformly for all k and i .

The following theorem is the natural generalization of Theorems 1-3 to backlogging.

Theorem 4. *Suppose that for the model described above assumption A holds for all $N_{k,i}(T)$ and the convex cone generated by $\theta_1, \dots, \theta_m$ is the whole of \mathcal{R}^n . Then the system is weakly-stabilizable. If the assumptions of Theorem 3 hold, then weakly-stabilizable may be strengthened to stabilizable.*

Proof. The proof is a modification of the proof for Theorem 2. At every observation time σ_i the inventory $x(\sigma_i)$ is nonnegative. As in the proof of Theorem 2 we solve (1) and show that for $\|x(\sigma_i)\|$ sufficiently large t can be an integer with $\|x(\sigma_i)\| / (\delta_2 \sqrt{n}) < t < \|x(\sigma_i)\| / \delta_1$. As in the proof of Theorem 2 we choose t_k such that $|\alpha_k t - t_k| < 1$ and $t = t_1 + \dots + t_m$. If $x(\sigma_i)$ is in D we run any producing mode for one period and set $\sigma_{i+1} = \sigma_i + 1$. If $x(\sigma_i)$ is not in D we first run each producing mode k for t_k periods, $k = r + 1, \dots, m$ and then run each request mode for t_k periods $k = 1, \dots, r$. Ignore for the moment the fact that in doing this some components of $x(t)$ may go negative. In a similar manner as in the proof of Theorem 2, we have that for $s = \sigma_i$ and $t = \tau_i$,

$$\begin{aligned} E(\|x(t+s)\|) &= E\left(\left\|x(s) + \sum_{k=1}^m (N_k(s_k + t_k) \right. \right. \\ &\quad \left. \left. - N_k(s_k))\right\|\right) \leq \delta_2 + \zeta(t). \end{aligned} \tag{8}$$

However, inventories are not allowed to go negative. So suppose that the effect of these t periods is to change the level of product k by C_k . The policy we will use is augmented by the proviso that if following any period during which the system is run in a request mode a product j is requested that is not in stock, then a producing mode k is chosen for which $\theta_{k,j} > 0$ and the system is run in mode k until one product of type j is produced. As demonstrated prior to the statement of the theorem, assumption **A** implies that the expected number of periods of production required to produce one more product of type j is bounded for all j by some $T' < \infty$. The number of times we will have to initiate extra production periods because of shortfalls in numbers of product j is at most $(x_j(\sigma_i) + C_j)^-$. The result of all such additional production periods, say ν_i periods in total, due to shortages for all types of product, will be to produce some further number of type j products, say Z_j . The next observation time is taken as $\sigma_{i+1} = \sigma_i + \tau_i + \nu_i$. It follows from assumption **A** and (5) that

$$E(\nu_i) \leq T' \sum_{j=1}^n E(x_j(\sigma_i) + C_j)^- \\ \leq T'E\left(\left\| \left(x(\tau_i) + \sum_{k=1}^m (N_k(s+t) - N_k(s)) \right) \right\| \right)$$

and therefore $E(\nu_i)$ is $o(t)$. Thus, $E(Z_j)$ is also $o(t)$. Let C and Z be the vectors with components C_j and Z_j . The inventory at the next observation time will be

$$x(\sigma_{i+1}) = (x(\sigma_i) + C)^+ + Z - (x(\sigma_i) + C)^- \\ \leq (x(\sigma_i) + C)^+ + Z. \tag{9}$$

Thus, $\|x(\sigma_{i+1})\|$ has an $o(t)$ bound of $\delta_2 + \zeta(t) + \|E(Z)\|$. The theorem is completed along the same lines as Theorem 2. Under the assumptions of Theorem 3, a similar proof shows that the system is stabilizable.

Thus far, we have not considered optimization models, but only questions of stability. However, a first step in the direction of optimization questions is the problem of stability under a constraint. Suppose that each use of request mode k is rewarded by r_k and each use of producing mode k costs c_k . Consider the problem of stabilizing the system while maximizing the rate at which profit is earned. To solve the problem we need to know for what β_k 's it is possible to stabilize the system while using mode k with a frequency β_k relative to other modes. One approach to this question is to augment the system by m dummy products, labeled $n + 1, \dots, n + m$. Let the action of mode S_k be extended to the dummy products by

$S_{k,n+k} = 1$, and $S_{k,n+j} = 0, j \neq k$. Suppose that β_1, \dots, β_m are rational numbers. Let there be an additional mode 0 for which $S_{0,n+k} = -N\beta_k$, where N is an integer chosen so that $N\beta_k$ is an integer for all k . Then by consideration of the inventory of dummy products, it is clear that if the system can be stabilized, then mode $k > 0$ must be used with a frequency β_k relative to the other modes. Theorem 4 may be applied to the augmented system to determine whether the system can be stabilized.

Although this is an appealing use of the ideas in previous sections it suggests a condition for stability that is actually too strong. Consider, for example, a system with $\theta_1 = (1, 1), \theta_2 = (-1, 0),$ and $\theta_3 = (0, -1)$. It is clearly possible to stabilize this system and use each mode 1/3 of the time, but that would not follow from the argument of the above paragraph, because $(0, 0, -1, -1, -1), (1, 1, 1, 0, 0), (-1, 0, 0, 1, 0),$ and $(0, -1, 0, 0, 1)$ do not span \mathcal{R}^5 . It is because the inventories of dummy products are altered in deterministic fashions that the condition can be weakened. A different necessary and sufficient condition can be established along the lines in Courcoubetis and Weber (1990). Although we do not prove it here, the conclusion is that the system can be stabilized if there exist $\mu_{lk} \geq 0$ and $\gamma_l > 0$ such that $\{\theta_l^i: \theta_l^i = \sum_{k=1}^n \mu_{lk}\theta_k, l = 1, \dots, n\}$ is a set of n linearly independent vectors whose nonnegative linear combinations span \mathcal{R}^n , and $\sum_{l=1}^n \gamma_l \theta_l^i = 0, \sum_{l=1}^n \gamma_l \mu_{lk} = \beta_k$.

4. A GENERAL MODEL FOR A MANUFACTURING SYSTEM

In this section we present a general model for a manufacturing system and apply some of the ideas from previous sections. We think of a manufacturing system that may be modeled by n distinct queues. Some queues are queues of orders for finished products. Other queues are inventories of raw materials, parts, and subassemblies, or work-in-progress waiting for processing at a given machine. Items arrive to these queues from outside the system according to stochastic processes. The arriving items are supplies of raw materials and orders for finished products. As above, we let $x_i(t)$ be the number of items in queue i at time t . A mode of manufacture makes *deterministic* changes to the queues, subtracting items from some queues and adding items to others. Subtractions correspond to the consumption of raw materials or the use of intermediate subassemblies to build other subassemblies and finished products. A final assembly operation subtracts from certain queues the resources

required to build a finished product and then subtracts one unit from the queue of unfilled orders for that product. A limiting constraint in this model is that the queues are not allowed to assume negative values, corresponding to the case in which a manufacturing mode cannot be executed if the subassemblies or materials it needs are not already available.

As in Section 3, suppose that the t th period of arrivals adds to x the vector $S_0(t)$, which for this section is assumed to be nonnegative in every component. The effect of S_0 is to add to queues or orders and raw materials. Modes $1, \dots, m$ are the manufacturing modes and these make *deterministic* changes to the queues by addition to $x(t)$ of the vectors S_1, \dots, S_m ; these may have both negative and nonnegative components. Again, we emphasize that these changes are deterministic, and hence in this section $\theta_i = S_i$ for $k = 1, \dots, m$. The above model corresponds to a Petri net if the terms *items, modes of operations, and queues* are translated to *tokens, transitions, and places*, respectively. There is a randomized transition corresponding to mode 0. This transition does not consume any tokens, and when fired for the t th time it adds a random number $S_{0,i}(t)$ of tokens to place $i, i = 1, \dots, m$. Note that if there is a nonnegative linear combination of S_1, \dots, S_m equal to 0, it is possible, given a large enough initial inventory, to construct a policy that stabilizes the system but never uses mode 0. Since this does not model a practical alternative we will only consider policies that select mode 0 infinitely often. We desire to stabilize the system by ensuring that for all i and t we have $E(x_i(t)) < B$ for some B .

Assume without loss of generality that the first r components of θ_0 are the only ones that are nonzero. Thus, $S_0(t)$ only adds items to the first r queues. These queues are called *input queues*; the remaining queues are called *intermediate queues* because they are queues of intermediate products, not queues of orders or raw materials. Let $\Gamma = \{x \in \mathbb{R}^n : x \leq 0 \text{ and } x_i = 0 \text{ for } i > r\}$. Here Γ is the set of nonpositive vectors that are only nonzero in components corresponding to input queues. Let

$$\Lambda = \left\{ x \in \Gamma : x = \sum_{k=1}^m \alpha_k S_k, \alpha_k \geq 0 \right\},$$

where Λ is the set of vectors in Γ that can be represented as $E[\sum n_k S_k]$ for some nonnegative integer-valued random variables $\{n_k\}$. It contains all those random combinations of the manufacturing modes whose effect is not to increase the size of the input queues and to leave the sizes of intermediate queues

unchanged. We provide a lemma and theorem to characterize stability. The lemma says that in considering what changes can be made to existing inventories by various combinations of the manufacturing modes it is sufficient to restrict attention to a finite number of the possible combinations.

Lemma 1. *The vector is a finitely-generated cone whose generators may be taken to be nonpositive integer-valued vectors, each of which can be written as a nonnegative integral combination of the vectors S_1, \dots, S_m .*

Proof. Since the vectors S_1, \dots, S_m are integer-valued they generate a cone that can be written as $\{x : Ax \geq 0\}$ for some matrix A with integer coefficients. Similarly, $\Gamma = \{x : Bx \geq 0\}$ for some B with coefficients that are all $-1, 0,$ or 1 . Thus, $\Lambda = \{x : Ay \geq 0, Bx \geq 0\}$. Since the generators of Λ must be rational vectors and they lie in the cone generated by S_1, \dots, S_m , each can be expressed as a nonnegative linear combination of the S_k s. There must be such linear combinations in which the multipliers of the S_k s can be taken to be nonnegative rationals. It follows, by rescaling the generators, that there is a choice of generators for which the multipliers of the S_k s are nonnegative integers.

Let θ_k^1 and Λ^1 be the restrictions of θ_k and Λ to their first r components, i.e., the nonzero components of θ_0 . The following theorem, while similar to those of previous sections, exhibits the fact that the condition for stability is now expressed only in terms of the flexibility that is required in expressing $-\theta_0^1$ in terms of generators of Γ^1 . Since changes to intermediate queues are deterministic there is no danger of these queues becoming unstable provided that the mix of modes that is being used lies within Γ .

Theorem 5. *Suppose that the arrival process $N_0(t)$ satisfies assumption A and $-\theta_0^1$ lies in the interior of the convex cone Λ^1 . Then the system is weakly-stabilizable. Furthermore, there exists a weakly-stabilizing policy under which the intermediate queue sizes are uniformly bounded. If the assumptions of Theorem 3 hold, then weakly-stabilizable may be strengthened to stabilizable.*

Proof. Using Lemma 1, let S'_1, \dots, S'_d be the generators of Λ such that $S'_l = \sum_k \mu_{lk} S_k, l = 1, \dots, d$, where the μ_{lk} s are nonnegative integers. We can associate with each S'_l a new manufacturing mode that corresponds to running mode k for μ_{lk} times, $k = 1, \dots, m$. Consider the new system where the arrival

mode is the same as before and the production modes are the ones corresponding to the S'_l , $l = 1, \dots, d$. If we project the state of the system to one in which we watch only the input queues, then we can weakly-stabilize this system while keeping the input queues nonnegative. This follows from an application of Theorem 4 when there is a single producing mode corresponding to S_0 and d request modes corresponding to the S'_l 's. Let $K = \max_l \{ \sum_k \mu_{lk} \} \max_{k \geq 1, i} \{ |S_{k,i}| \}$. One can see that if before running any mode S'_l the size x_i of each intermediate queue is greater than K , then while running the S_k 's comprising S_l , this state will remain within the range $x_i \pm K$, and at the end reassume the value x_i . This implies that if the intermediate queues sizes are all at least K initially we can weakly-stabilize the projection of the system onto its input queues and thereby weakly-stabilize the complete system, while ensuring that the intermediate queue sizes are uniformly bounded. If we assume the conditions of Theorem 3 a similar proof shows the system is stabilizable.

5. NECESSARY CONDITIONS FOR STABILITY

In the previous sections we considered several variations of a model for a flexible manufacturing system and gave sufficient conditions for the system to be weakly-stabilizable or stabilizable. In this section, we show that the sufficient condition is also necessary for Theorem 3 in the sense that if the cone generated by $\theta_1, \dots, \theta_m$ is not the whole of \mathcal{R}^n then, apart from trivial cases, the system cannot be stabilized; the size of the inventory, or surplus production, will explode to infinity under any operating policy. Recall that in the model of Theorem 3 we run the system in one of k possible modes during each time period. Suppose that the $\{S_{k,i}(t)\}$ are mutually independent random variables and $\{S_{k,i}(t)\}_{t=1}^\infty$ are identically distributed. Assumption A will hold.

Without loss of generality assume that if any of the m modes are not available, then the process is no longer stabilizable. Assume also that the inventory level of each product is affected by at least one mode. Under the first assumption it is easy to see that to stabilize the system every mode must almost surely be used infinitely often. Suppose that there is such a policy, but the cone generated by $\theta_1, \dots, \theta_m$ is not the whole of \mathcal{R}^n . Then by Farkas' lemma there exists a $\neq q$, which may be chosen to have integer coefficients, such that $q^T \theta_k \geq 0$ for all k . Choose such a $q \neq 0$ and consider the integer-valued process $y(t) = q^T x(t)$. Note that if the system is run in mode k during period t we have $E[y(t+1) - y(t) | x(t)] = q^T \theta_k \geq$

0 and thus $y(t)$ is a sub-Martingale and by the Martingale convergence theorem it tends to a limit on every sample path. Now this must almost surely be a finite limit, because $|y(t)| \leq \|q\| \|x(t)\|$, and $E[x(t)]$ is uniformly bounded for all t because it is produced by a stabilizing policy. Since $y(t)$ is an integer, its limit on almost every sample path must be an integer and $y(t)$ must be equal to this limit for all t sufficiently large. But because we have assumed that each mode is almost surely used infinitely often this is impossible unless on almost every sample path the incremental change in y when mode k is used for the t th time satisfies $q^T \theta_k(t) = 0$ for all t sufficiently large. Now if q has only one nonzero component, say i , this would imply that on almost every sample path no mode affects the inventory of product i for t sufficiently large, contradicting our second assumption. If q is nonzero in more than one component, say in all components $i \in I$, then it is clearly impossible to have $q^T \theta_k(t) = 0$ unless the independent random variables $S_{k,i}(t)$, $i \in I$ are almost surely constants for t sufficiently large. Therefore, except in the special case that some modes change the inventories of some products by deterministic amounts, the system cannot be stabilizable.

The above argument demonstrates that the sufficient condition in Theorem 3 is necessary if the system is to be stabilizable for any $S_k(t)$'s. The exception is when some of the S_k 's are deterministic and there is a positive linear combination of these modes whose sum is zero. The necessity argument for the other stability theorems in the paper is similar.

6. DISCUSSION

What implications do the above results have for the management of real-life flexible manufacturing systems? We have seen that if the models in this paper are reasonable for a particular application and we can estimate the vectors θ_k , then Theorems 1–5 provide criteria by which to test whether the system can operated in a manner that keeps expected inventory levels uniformly bounded through time. For example, in Theorem 2 we have said how one could compute the time t_k for which the system is to be run in mode k before the policy is next reviewed. Choosing δ_1 in the proof of Theorem 2 requires an estimate of the speed with which $g(t)/t$ tends to 0 in assumption A. This could be estimated from observations of the $N_{k,i}(t)$'s. The criteria of the theorems can be tested using simple algorithms. For example, in Theorems 1–4 we required that the cone generated by nonnegative linear combinations of $\theta_1, \dots, \theta_m$ be the

whole of \mathcal{R}^n . This holds if and only if $\theta_1, \dots, \theta_m$ span \mathcal{R}^n and there exist $\gamma_1, \dots, \gamma_m$ such that $\gamma_k \geq 1$ for all k and $\sum_k \gamma_k \theta_k = 0$. This can be checked by a linear program in m variables. Theorem 5 required that $-\theta_0^1$ be in the interior of Λ^1 . An algorithm to check this can be based on the following ideas. Let $\Omega = \{x \in \mathcal{R}^n : x_i = 0 \text{ for } i > r\}$ and define I as the set of those indices $i > 0$ for which there exists $\{\alpha_k \geq 0, k \neq i\}$ such that $\theta_i + \sum_{k \neq i} \alpha_k \theta_k \in \Omega$. The condition of Theorem 5 is equivalent to requiring that the vectors $\{\theta_i, i \in I\}$ span Ω and there exist $\{\gamma_k \geq 1, k \in I \cup \{0\}\}$ such that $\sum_{k \in I \cup \{0\}} \gamma_k \theta_k = 0$.

There are at least two important questions outstanding. First, one would like to identify *practical* policies that achieve stability. The policies described in the proofs of Theorems 1–5 can be implemented by relatively simple calculations, but they are wasteful in the use of inventory space and certainly not rules-of-thumb, or methods one would recommend in practice. Clearly, it is important to *minimize* average inventory levels, not just to ensure that their expected values remain finite. One would really like to find policies that stabilize the system and require minimum average inventory levels to do so. Although the work-in-progress inventory is likely to be an important component of the total cost, one might also like to understand stability for a cost function depending on other factors, such as lost orders, or the frequency of switching production mode. Clearly, there are many such issues that might be investigated.

Second, the *maximum* inventory size is likely to be constrained. Interestingly, we can say something about this sort of constraint when assumption **A** is replaced by

Assumption C. *There is a function $g(t)$ such that for all k, i, s and $t > 0$*

$$|N_{k,i}(s+t) - N_{k,i}(s) - t\theta_{k,i}| \leq g(t),$$

where $g(t)/t \rightarrow 0$ as $t \rightarrow \infty$.

Under this assumption it is straightforward to modify the proofs of Theorems 2–5 to show stability in the sense that $|x_i(t)|$ is bounded by a constant uniformly in t .

As we have explained at the end of Section 3, there may be constraints on the rates at which some of the manufacturing modes are used. For example, we might pose the problem of stabilizing the system subject to a constraint that a particular mode k be used at rate β_k . This problem was considered in Courcoubetis and

Weber (1990) in the bin packing context, and similar results can be obtained for the models in this paper.

ACKNOWLEDGMENT

We are grateful to Alan Weiss for his help in formulating assumption **B** and the proof of Theorem 3. The referees also provided a number of helpful comments and references. We are grateful to one of the referees for the construction of the t_k s in the proof of Theorem 2.

REFERENCES

- BIELECKI, T., AND P. R. KUMAR. 1988. Optimality of Zero-Inventory Policies for Unreliable Manufacturing Systems. *Opns. Res.* **36**, 532–541.
- COURCOUBETIS, C., P. KONSTANTOPOULOS, J. WALRAND AND R. R. WEBER. 1989. Stabilizing an Uncertain Production System. *Queue. Syst. Theory and Applic.* **5**, 39–54.
- COURCOUBETIS, C., AND R. R. WEBER. 1986a. Necessary and Sufficient Conditions for Stability of a Bin Packing System. *J. Appl. Prob.* **23**, 989–999.
- COURCOUBETIS, C., AND R. R. WEBER. 1986b. A Bin Packing System for Objects From a Finite Set: Necessary and Sufficient Conditions for Stability and Some Applications. In *Proceedings of 25th IEEE CDC*, Athens, Greece.
- COURCOUBETIS, C., AND R. R. WEBER. 1990. Stability of On-Line Bin Packing With Random Arrivals and Long-Run-Average Constraints. *Prob. Engin. and Inf. Sci.* **4**, 447–460.
- COURCOUBETIS, C., AND U. G. ROTHBLUM. 1991. On Optimal Packing of Randomly Arriving Objects. *Math. O. R.* **16**, 176–194.
- KIMEMIA, J. G., AND S. B. GERSHWIN. 1983. An Algorithm for Computer Control of Production in Flexible Manufacturing Systems. *IIE Trans.* **15**, 353–362.
- GERSHWIN, S. B., R. AKELLA AND Y. F. CHOONG. 1985. Short-Term Production Scheduling of an Automated Manufacturing Facility. *IBM J. Res. Dev.* **29**, 392–400.
- HAJEK, B. 1982. Hitting-Time and Occupation-Time Bounds Implied by Drift Analysis With Applications. *Adv. Appl. Prob.* **14**, 502–525.
- PERKINS, J. R., AND P. R. KUMAR. 1989. Stable, Distributed, Real-Time Scheduling of Flexible Manufacturing/Assembly/Disassembly Systems. *IEEE Trans. Auto. Control* **34**, 139–148.
- RHEE, W. T., AND M. TALAGRAND. 1988. Some Distributions That Allow Perfect Packing. *J. Assoc. Comput. Mach.* **35**, 564–578.