Mathematical Methods of Operations Research (1997) 45:455–481

Using Information about Machine Failures to Control Flowlines

C. Roger Glassey[1]

Dept. of Industrial Engineering and Operations Research, University of California, Berkeley, CA 94720, USA

Sridhar Seshadri[1]

Department of Statistics and Operations Research & Operations Management Area, Leonard N. Stern School of Business, New York University, NY 10012, USA

J. G. Shanthikumar[2]

Walter A. Haas School of Business and Dept. of Industrial Engineering and Operations Research, University of California, Berkeley, CA 94720, USA

Abstract: We study the problem of using information about the type of machine failures to control production. When each station in a two stage flowline has a single machine, the processing time distributions are Erlang, and the time to fail as well repair a machine are exponentially distributed; we show that the optimal policy for controlling the flowline is completely specified by selecting a produce up to level called threshold level corresponding to each type of machine failure. These threshold levels are shown to be ordered as per the rate at which the failed downstream station can be repaired. These results are partially extended to the case when there are multiple machines at the two stations of the flowline. We then summarize numerical results that illustrate when it is worth while to use information about failures in controlling inputs to the flowline.

Key Words: Production/Scheduling, Markov Decision Theory, Semiconductor Industry.

1 Introduction

Machine failures are usually a major source of uncertainty in production planning and control. This is especially true in the context of semiconductor fabrication where the operations are machine dependent, critical machines are very

1 The work of C. Roger Glassey and Sridhar Seshadri was partly supported by grants to the University of California from the Semiconductor Research Corporation and the California State MICRO program.
2 The work of J. G. Shanthikumar was partly supported by NSF Grant DDM-9113008 and the Sloan Foundation grants for the study on "Competitive Semiconductor Manufacturing."

1432-2994/97/45:3/455–481 $2.50 © 1997 Physica-Verlag, Heidelberg

expensive, and work can not be subcontracted out. It is well known that the effect of machine breakdowns can be mitigated by carrying out properly planned preventative maintenance. However relatively little effort has been made to understand when and how information about the type of machine breakdowns can be used in production control. In this paper we study how information about the *type* of machine failures can be used for solving the production control problem of minimizing the average inventory in two stage flow lines subject to achieving a given rate of production (also called throughput rate). In our models the flowline consists of two stations, see Figure 1. There may be several identical machines at each station that are subject to multiple types of failures. There is unlimited supply of raw materials at the first station and unlimited storage space (also called the buffer space) between the stations. The decision required is when to release material for processing to the first station.

Using this model, we show that when there are single machines at the two stations, machine failure as well as repair times are exponentially distributed, and processing time distributions are Erlang, the optimal control policy can be described by specifying a threshold (or produce up to) level for each type of failure of the downstream station. Thus the information about the type of failure translates into the optimal control of releasing material into the flowline till the work for the second station in the line reaches the threshold level for that type of failure. We then show that the threshold levels are ordered naturally (nested), in the sense, the longer it will take to repair a given type of machine failure the smaller should be the threshold level for that type of failure. Because we use Erlang distributions for the processing times, the results extend to the case when processing times are deterministic. These results are partially extended to the case when there are multiple machines at the two stations of the flowline. We then summarize numerical results that illustrate *when* it is worth while to use information about the type of failure in controlling inputs to the flowline. Our proof techniques comprise studying an equivalent problem of controlling a Markov chain; and setting up a dynamic program to find the optimal control. Most of the methods used by us are standard, except a coupling argument used for establishing the nestedness of the threshold levels.

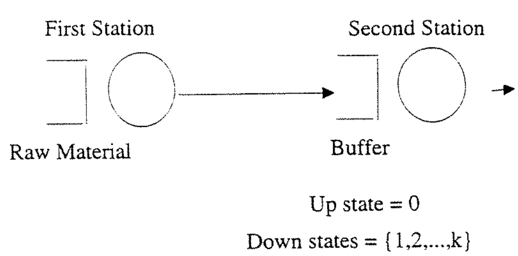

Fig. 1. A two stage flowline with multiple types of machine failures

The paper is organized as follows, in section 2 we discuss the modeling issues and review related literature, in section 3 we describe the model with single machines at each station and present the analysis of the optimal control policy, section 4 is on the analysis of the model with multiple machines at each station, and in section 5 we summarize numerical results, and suggest extensions of this work.

2 Modeling Issues and Review of Literature

The work reported in this paper is part of a larger study, described in Seshadri (1993) and Glassey, Shanthikumar and Seshadri (1996), whose objective was to determine efficient scheduling and input control rules for use in semiconductor wafer fabrication. The facility for producing semiconductor wafers, each wafer containing from tens to hundreds of chips, is called a *fab*. In this section, we briefly describe a typical fab, how some of the modeling assumptions stated in the introduction arise in this context, and review related literature.

The Flowline Model for Fabs: There are three major stages to making an IC, namely Wafer Preparation, Wafer Fabrication and Assembly. We study the wafer fab because most of the complicated production work and investment in facilities are concentrated at the clean room fabrication stage – in which layers of intricate circuitry gets built upon the wafers. Because of the large investment in fabs, the stated goal of most fabs is to simultaneously achieve high levels of production and high yield of good wafers. While this goal has to be achieved to recoup the investment, it is often at odds with the goal of maintaining low inventories. Moreover, high inventories that result as a consequence of high levels of production can conceal problems, leading to lowered yields, which means more pressure for higher production and higher inventories and so forth – a well understood downward spiral (Thomas (1993)). This explains the choice of the trade-off in our objective function, namely, "to minimize the average inventory subject to achieving a given throughput."

In the entire wafer fabrication there are *three* critical processes: oxidation, photolithography and diffusion, and the wafer visits these processes a number of times before leaving the fab. Kumar (1993, 94) proposed a reentrant flow line model to capture this repeated sequence of processing which is necessary to build up the layers of circuitry. While his model captures the actual dynamics, it is also rather complex for applying dynamic programming (DP) – the technique used in this paper. Therefore, we have chosen to study two stage flowlines in this paper, and focus more on understanding when information about the type of machine failures is important, and how this information should be used by a production controller. We do suggest guidelines for extending our work to real life fabs in section 6.

There are two major sources of uncertainty in a fab, namely, yield variations and machine failures. We attempt to model only uncertainties arising from machine failures in this paper, and some of the yield management aspects have been reported in Nurani (1995), Seshadri and Shanthikumar (1994) and Nurani, Seshadri, and Shanthikumar (1995). The decision to use the exponential distribution in our model for the time to fail and repair machines was based on informal discussions with fab managers. In addition we assume that the sequence of durations when a given machine is operational (called *up* times), and the sequence of the repair times (called *down* times), are independent of one another, and that the successive up (down) times are independent and identically distributed random variables. The statistical work of Hanifin discussed in Buzacott and Shanthikumar (1993) shows that these assumptions may not always hold good in practice and that subsequent times to fail could well be correlated.

Related Literature: We review literature related to incorporating machine failures in production control. For work done on production planning of semiconductor fabs, the reader is referred to the excellent survey provided by Uzsoy, Lee and Martin-Vega (1992, 94). There has been ample work done in the area of scheduling single stage failure prone facilities, see for example Glazebrook (1987), Birge and Glazebrook (1988), Pinedo and Ramouz (1988), Posner and Berg (1988), Buzacott and Shanthikumar (1993), and also Pinedo (1995) which contains a description of the research on single machine stochastic models that deal with a large class of objective functions. In most of these studies, the objective of controlling a single stage facility is usually given in terms of meeting due dates or avoiding backlog while minimizing the cost of holding finished products. Thus the single stage models do not deal with release control. Because of this the results for single stage failure prone facilities are not directly applicable in our context.

Several studies have been carried out to evaluate the performance of failure prone flowlines in which the control is exercised solely by fixing the sizes of inter-stage buffers. Literature survey, examples, and analysis of these models are given in Buzacott and Shanthikumar (1993). There is however no guarantee that the control of a flowline using fixed buffer sizes will give results close to the optimal performance. Flowlines have also been analyzed under the assumption that the machines are completely reliable. Examples of such work include the analysis of the *PAC* system in Buzacott and Shanthikumar (1992), the work reported in Veatch and Wein (1993), and the unifying notion of Linear Control Rules for release control developed in Glassey, Shanthikumar, and Seshadri (1996).

There have been several studies devoted to the design of release control systems for flexible manufacturing systems (FMS). Machine failures are the major source of uncertainty in FMS and their control systems have to use feedback information on failures to perform correctly. The earliest work on failure prone FMS is that of Hildebrandt (1980). More recent examples include the work of Akella and Kumar (1986), Sharifnia (1988) and Boukas and Hourie (1990). An excellent and accessible summary of this work is given in Gershwin (1994). In the models of FMS, the flow of work is approximated by *fluid flow* and production is assumed to be instantaneous (i.e., the stocks and flows in the model are connected by differential equations). The assumption of fluid flow is appropriate in the context of FMS because there is little or no storage space between stations in FMS, and operations take place on a time scale which is much smaller compared to the times to failure and the repair times. The production control problem in FMS is to *match* the output from the FMS with a given demand pattern. Thus the time scales in the models of FMS as well as tradeoffs in the production control problem for FMS are entirely different and we should expect little of the results for FMS to carry over to the case of flowlines. It is interesting then, that the structural results in our study correspond to those given in Sharifnia (1988). Sharifnia using the fluid flow approximation, studied the problem of controlling a single stage facility subject to different types of failures in order to minimize the combined cost of backlogged demand and the cost of holding finished product in the output store. He too found that there were threshold levels corresponding to each type of failure, and given a type of failure it was optimal to produce whenever the inventory in the output store was below the corresponding threshold level. This similarity of optimal policies is due to the ubiquity of threshold type optimal control policies for control of Markovian systems; and interesting because we are able to obtain a similar result for a two stage flow line with discrete job flow, and deterministic processing times.

There are a number of studies related to control of queueing networks that incorporate machine failures, such as the work of Harrison (1988), Harrison and Wein (1990), Wein (1990), and Kushner (1990). In these studies, the time to repair a failed machine is considered to be a part of the job processing times. This assumption is valid when the duration of failures and job processing times are comparable. On the other hand there are situations in which the times to fail and repair are very much larger than the processing times. In such circumstances the occurrence of failures and repairs *can not be* ignored and the state of machines *can not be* taken as fixed for the purpose of controlling the release and scheduling of jobs. Buzacott and Shanthikumar (1993, p. 234) discuss these two views.

3 Single Machines in Each Station – Model Description and Analysis

In this section, we consider a two stage flow line with single machines at each station and infinite inter-stage buffer. The assumptions are listed below.

Station #1: There is infinite supply of raw materials, the service time to process a job at the first station has an Erlang distribution with r phases and mean r/μ_1, and the station is completely reliable.
Station #2: Processing times have an Erlang distribution with s phases and mean s/μ_2, the failure rate is exponentially distributed with rate ζ, it is assumed that the machine can fail even when it is idle (called the time dependent failure model, see Buzacott and Shanthikumar (1993)), and that the job under process at the time of failure will "resume" its processing after the machine has been repaired. There are k types of failures and the probability that the next failure is of type j is given by p_j, $j = 1, 2, \ldots, k$, $\sum_{j=1}^{k} p_j = 1$.
Demand and Costs: Demand for the product is unlimited, every completed unit (job) yields a profit of p' per unit produced, there is a per unit holding cost of Λc per unit time spent in the system after the first phase of processing has been completed at station #1 (discussed below), and all profits and costs are continuously discounted using a factor of α.
Control: The controller is allowed to process a job at the first station using a rate, $u \in [0, 1]$. The control space is $[0, 1]^N$, where $N = \{1, 2, \ldots\}$.

The Dynamic Programming Recursions: The objective is to maximize the average discounted net profit over the infinite time horizon. This problem is equivalent to that of minimizing the average discounted inventory holding cost subject to achieving a given throughput rate. Instead of analyzing the problem in continuous time as stated, we use a technique due to Lippman (1977) and work with the equivalent discrete time Markov chain to analyze the problem. The Markov chain is obtained by uniformizing the continuous time process of part production, machine failures and repairs, using a fast Poisson process of sufficiently large rate Λ. The state space for this chain is given by,

$$S = \{(u, v, n, j); u = 0, 1, 2, \ldots, r-1; v = 0, 1, 2, \ldots, s-1;$$
$$n = 0, 1, 2, \ldots, \ldots; j = 0, 1, 2, \ldots, k\} \;;$$

where the first two values, i.e. u and v, denote the phase of service that is underway at the first and second machines. The third value, n, is the number of jobs in the buffer plus the job in service if any at the second machine. The last value denotes the type of breakdown, with 0 denoting that the second machine is operational (up). We have assumed that the information on the phase under service is available to the controller, for analytical convenience. Define n' to be the number of jobs in the system that have completed the first phase of processing at the first station. Redefine the profit per unit to be, $p = p' - \frac{c\Lambda}{\mu_1}$. The logic of redefining the profit is that as stated we shall be charging inventory holding costs on a job from the instance one phase of service gets completed at the first station. This is not a shortcoming because we are primarily interested in the average cost case and use discounting to facilitate inductive proofs of structural properties. For the average cost case the correction to the profit is correct when the first station is never idled whenever there are no jobs downstream of the first station. Also note that in extending the results to the deterministic case, we will have to let the number of phases get very large and increase the processing rates proportionately, in which case the correction term goes to zero. Denote the indicator function of set A as $I\{A\}$. The expected single stage reward function, $R(u, v, n, j; u_1)$, received from using the control u_1 in the state (u, v, n, j), can be written as,

$$R(u, v, n, j; u_1)$$

$$= \frac{\mu_2 I\{v = s-1; j = 0; n \neq 0\}}{\Lambda + \alpha}(p - (n' - 1)c) \quad \text{(Profit – Holding Cost of } (n'-1) \text{ jobs if a job leaves the system)}$$

$$+ \frac{u_1 \mu_1}{\Lambda + \alpha}(-n'c - cI\{u = 0\}) \quad \text{(Holding cost of } n' \text{ jobs } \textit{plus} \text{ the cost of holding a new job that progresses beyond the 1st stage of processing at the first machine)}$$

$$- \frac{(\Lambda - u_1\mu_1 - \mu_2 I(v = s-1; j = 0; n \neq 0))}{\Lambda + \alpha} n'c \quad \text{(Holding cost of } n' \text{ jobs if "nothing", i.e., a transition in the fast Poisson process, happens)}$$

$$= \frac{\mu_2 I(v = s-1; j = 0; n \neq 0)}{\Lambda + \alpha}(p + c) - \frac{\Lambda}{\Lambda + \alpha} n'c - \frac{u_1\mu_1}{\Lambda + \alpha} cI(u = 0) \; . \quad (1)$$

The three components of the expected reward correspond to the three possibilities: the second station completes one stage of processing, the first station completes a stage of processing, or a transition of the fast Poisson process occurs. The denominator of $\Lambda + \alpha$ in these expressions arises due to the discounting of the profit and costs by the factor, $\Lambda/(\Lambda + \alpha)$. Let $V(u, v, n, j; u_1)$, be the expected net discounted profits over the infinite horizon, obtained by using the control u_1 immediately in the state (u, v, n, j), and the optimal control thereafter. Let $V(u, v, n, j) = \max_{u_1 \in [0,1]} V(u, v, n, j; u_1)$, be the optimal value function in the state (u, v, n, j). Then, using equation (1), and a similar logic, we obtain,

$$V(u, v, n, j; u_1) = R(u, v, n, j; u_1)$$

$$+ \frac{u_1\mu_1}{\Lambda + \alpha} V(0, v, n+1, j) I\{u = r - 1\}$$

$$+ \frac{u_1 \mu_1}{\Lambda + \alpha} V(u+1, v, n, j) I(u \neq r-1)$$

$$+ \frac{\mu_2 I(j=0)}{\Lambda + \alpha} (V(u, 0, n-1, j) I\{v = s-1; n \neq 0\}$$

$$+ V(u, v+1, n, j) I(v \neq s-1; n \neq 0))$$

$$+ \frac{\lambda(j) I\{j \neq 0\}}{\Lambda + \alpha} V(u, v, n, 0) + \frac{\sum_l \zeta p_l}{\Lambda + \alpha} V(u, v, n, l) I\{j = 0\}$$

$$+ \frac{(\Lambda - u_1 \mu_1 - (\mu_2 I\{n \neq 0\} + \zeta) I\{j=0\} - \lambda(j) I\{j \neq 0\})}{\Lambda + \alpha}$$

$$\times V(u, v, n, j) \ . \tag{2}$$

The existence of an optimal policy for this discounted version of the DP follows from standard arguments, for example see Ross (1970). For the average cost case, when the holding cost is positive it can be shown that there is a finite upper bound for the value of n' for the states that are not transient in the optimally controlled Markov chain. This implies that the state (0, 0, 0, 0) will recur infinitely often and thus there exists an optimal control policy for the average cost too. To formalize this argument we would only need to show that the original problem is equivalent to solving a dynamic program which terminates when the state of the chain hits (0, 0, 0, 0), and the proof is omitted. In equation (2), collecting the terms involving control u_1 in the expression for $V(u, v, n, j; u_1)$ gives,

$$u_1(-cI\{r-1=0\} + V(0, v, n+1, j) - V(r-1, v, n, j))$$
$$\text{for } u = r-1 \ ; \quad n > 0 \ , \tag{3a}$$

$$u_1(-cI\{u=0\} + V(u+1, v, n, j) - V(u, v, n, j))$$
$$\text{for } u \neq r-1 \ ; \quad n > 0 \ , \tag{3b}$$

$$u_1(-cI(r-1=0) + V(0, 0, 1, j) - V(r-1, 0, 0, j))$$
$$\text{for } u = r-1 \ ; \quad n = 0 \ , \tag{3c}$$

$$u_1(-cI(u=0) + V(u+1, 0, 0, j) - V(u, 0, 0, j))$$
$$\text{for } u \neq r-1 \ ; \quad n = 0 \ . \tag{3d}$$

From these expressions, it is evident that the optimal control is of the bang-bang type. The decision to produce in state (u, v, n, j) will be either 0 or 1, depending on whether the terms in the parentheses in (3) are negative or non-negative. In other words, the optimal control u_1^* is a function of the state (u, v, n, j), with $u_1^*(u, v, n, j)$ taking values in the set $\{0, 1\}$.

The Main Result: The main result of this section is given next. It is based on four lemmas, with Lemma 3 being the key to the theorem.

Theorem I: For a two stage flowline, with single servers at both stations, Erlang processing time distributions, reliable first station; when the time to fail and repair the second station are exponentially distributed, there are multiple types of failures possible, profit p per unit produced, constant holding cost per unit per unit time spent in the system, and under continuous discounting; the optimal policy to maximize the expected net value of discounted profits over the infinite time horizon has the following structure,

(i) For every state j, $j = 0, 1, 2, \ldots, k$ of the second machine, there exists a threshold level $Z(j)$, such that when the second machine is in state j, it is optimal to continue production at the first machine (and produce at the maximum possible rate) only if the inventory downstream of the first machine measured in terms of processing phases for the second machine is less than or equal to $Z(j)$.
If all threshold values are finite then:
(ii) If $Z(0)$ is not zero, then $Z(0) \geq Z(j)$, $j = 1, 2, \ldots, k$.
(iii) If the repair rate (of the second machine) in failed state j is smaller than that in failed state l, then $Z(l) \geq Z(j)$.
(iv) It is never optimal to idle the first station except when the first phase of processing has to be undertaken. It is never optimal to idle the second station.

Proof:

(i) In Lemma 2, we show that $V(u, v, n, j) - V(u-1, v, n, j)$ is non-decreasing in v. Part (i) of the theorem now follows from equations (3).
(ii) In Lemma 3, we show that $V(a, v, n, j) - V(b, v, n, j) \leq V(a, v, n, 0) - V(b, v, n, 0)$, for $a > b$, $n > 0$. Part (ii) follows from this and equations (3).
(iii) In Lemma 4, we show that if $\lambda(l) > \lambda(j)$, then $V(a, v, n, j) - V(b, v, n, j) \leq V(a, v, n, l) - V(b, v, n, l)$, for $a > b$, $n > 0$. Part (iii) follows from this and equations (3).
(iv) This part follows immediately from Lemma 1, where we show that (a) $V(u, v, n, j) - V(u-1, v, n, j) \geq 0$, for $(u-1) > 0$, and (b) $V(u, v, n, j) - V(u, v-1, n, j) \geq 0$. ■

There are two practical implications of this theorem. The first implication is that using thresholds that are stated in terms of jobs (instead of phases) downstream of the first machine will result in a policy that is at most "one job away" from the optimal policy. The second implication is that all the results of the theorem extend to the case when processing times are deterministic. The four lemmas used in this theorem are proved below.

Lemma 1:

(a) $V(u, v, n, j) - V(u - 1, v, n, j) \geq 0$, for $(u - 1) > 0$ and
(b) $V(u, v, n, j) - V(u, v - 1, n, j) \geq 0$

Proof: For proving (a), consider two systems I and II. Let the initial state in system I be $(u - 1, v, n, j)$ and be (u, v, n, j) in system II, with $u - 1 > 0$. Couple the two systems by letting the sequence of up and down times and processing times at the second station be the same in both systems. Assume that system I is operated under the optimal control. Consider the (sub-optimal) policy of *idling* the first station in system II till the first station in system I has completed the processing of the $(u - 1)st$ phase, an event that is bound to happen eventually as the systems are assumed to be profitable to operate. At that instant either the two systems hit the same state or system II may have processed jobs for a longer time at the second station (if $u \geq r$). Part (a) now follows by use of part (b).

For proving part (b), consider two systems I and II. Let the initial state in system I be $(u, v - 1, n, j)$ and in system II, let the initial state be (u, v, n, j). Assume that the sequence of up and down times at the second station, as well as the processing times at the first station are the same in the two systems. Use the optimal policy to operate system I. In system II, follow the same input control policy as used in system I (which is a sub-optimal policy for system II). It is then easy to see that both systems will have the same stream of new jobs joining the buffer between the stations – but the second station in II will always either be ahead in processing jobs, or in step with the second station in system I. This will lead to net profits in II being at least as high as that in system I. This lemma may be summarized by saying that inserted idleness is not beneficial. ■

Lemma 2:

(i) $V(u, v, n, j) - V(u - 1, v, n, j) \geq V(u, v - 1, n, j) - V(u - 1, v - 1, n, j)$ for $0 \leq u \leq r - 1; 0 \leq v \leq s - 1$.
(ii) $V(u, 0, 0, j) - V(u - 1, 0, 0, j) \geq V(u, v, n, j) - V(u - 1, v, n, j)$

Proof: We shall prove only part (i) of the claim. The proof of part (ii) is similar and is omitted. To prove part (i) we use a standard induction technique using the finite time horizon framework. The proof is tedious and involves enumerating several cases, and analyzing each separately. One unit of time is one transition in the uniformizing Poisson process. Define, $V^m(u, v, n, j)$ to be the optimal net discounted profit, when starting at state (u, v, n, j) and there are m time units left to go. The time units (periods) left to go will also be termed the *stage* of the DP, and $V^m(u, v, n, j)$ in this terminology is called the m stage value function. The recursion equations for the $(m + 1)$ stage value function are, compare with equation (1),

$$
\begin{aligned}
V^{m+1}(u, v, n, j) = \max_{u_1 \in [0,1]} \Big(& R(u, v, n, j; u_1) \\
& + \frac{u_1 \mu_1}{\Lambda + \alpha} V^m(0, v, n+1, j) I\{u = r - 1\} \\
& + \frac{u_1 \mu_1}{\Lambda + \alpha} V^m(u+1, v, n, j) I\{u \neq r - 1\} \\
& + \frac{\mu_2 I(j = 0)}{\Lambda + \alpha} (V^m(u, 0, n-1, j) I\{v = s - 1; n \neq 0\} \\
& + V^m(u, v+1, n, j) I\{v \neq s - 1; n \neq 0\}) \\
& + \frac{\lambda(j) I\{j \neq 0\}}{\Lambda + \alpha} V^m(u, v, n, 0) + \frac{\sum_l \zeta p_l}{\Lambda + \alpha} V^m(u, v, n, l) I\{j = 0\} \\
& + \frac{(\Lambda - u_1\mu_1 - (\mu_2 I\{n \neq 0\} + \zeta) I\{j = 0\} - \lambda(j) I\{j \neq 0\})}{\Lambda + \alpha} \\
& \times V^m(u, v, n, j) \Big) \,.
\end{aligned} \tag{4}
$$

Let the lemma be true for the m stage value function. As a check, the induction can be started in the last period by setting the zeroth period value function to be identically zero. Let u_1, u_2, u_3 and u_4 be the optimal controls used in the states (u, v, n, j), $(u - 1, v, n, j)$, $(u, v - 1, n, j)$, and $(u - 1, v - 1, n, j)$ when there are $(m + 1)$ periods left to go. Denote the difference, $(V^m(u, v, n, j) - V^m(u - 1, v, n, j)) = \Delta^m(u, v, n, j)$. Define $(-1, v, n, j) = (r - 1, v, n - 1, j)$ and $(u, -1, n, j) = (u, s - 1, n + 1, j)$. Consider the difference $(V^{m+1}(u, v, n, j) - V^{m+1}(u - 1, v, n, j)) - (V^{m+1}(u, v - 1, n, j) - V^{m+1}(u - 1, v - 1, n, j))$ given by,

$$\Delta^{m+1}(u, v, n, j) - \Delta^{m+1}(u, v-1, n, j)$$

$$= -\left(\frac{u_1\mu_1}{\Lambda+\alpha} cI\{u=0\} - \frac{u_2\mu_1}{\Lambda+\alpha} cI\{u-1=0\}\right) \tag{5-1}$$

$$+ \frac{u_1\mu_1}{\Lambda+\alpha}(V^m(0, v, n+1, j)I\{u=r-1\}$$

$$+ V^m(u+1, v, n, j)I\{u \neq r-1\}) \tag{5-2}$$

$$- \frac{u_2\mu_1}{\Lambda+\alpha} V^m(u, v, n, j) \tag{5-3}$$

$$+ \frac{(\Lambda - u_1\mu_1 - \mu_2 I\{j=0\} - \lambda(j)I\{j\neq 0\})}{\Lambda+\alpha} V^m(u, v, n, j)$$

$$+ \frac{\lambda(j)I\{j\neq 0\}}{\Lambda+\alpha} V^m(u, v, n, 0) \tag{5-4}$$

$$- \frac{(\Lambda - u_2\mu_1 - \mu_2 I\{j=0\} - \lambda(j)I\{j\neq 0\})}{\Lambda+\alpha} V^m(u-1, v, n, j)$$

$$- \frac{\lambda(j)I\{j\neq 0\}}{\Lambda+\alpha} V^m(u-1, v, n, 0) \tag{5-5}$$

$$+ \frac{\mu_2 I(j=0)}{\Lambda+\alpha}(\Delta^m(u, v+1, n, j)I\{v < s-1\}$$

$$+ \Delta^m(u, 0, n-1, j)I\{v = s-1\}) \tag{5-6}$$

$$+ \left(c\frac{u_3\mu_1}{\Lambda+\alpha} I\{u=0\} - c\frac{u_4\mu_1}{\Lambda+\alpha} I\{u-1=0\}\right) \tag{6-1}$$

$$- \frac{u_3\mu_1}{\Lambda+\alpha}(V^m(0, v-1, n+1, j)I\{u=r-1\}$$

$$+ V^m(u+1, v-1, n, j)I\{u \neq r-1\}) \tag{6-2}$$

$$+ \frac{u_4\mu_1}{\Lambda+\alpha} V^m(u, v-1, n, j) \tag{6-3}$$

$$-\frac{(\Lambda - u_3\mu_1 - \mu_2 I\{j=0\} - \lambda(j)I\{j \neq 0\})}{\Lambda+\alpha} V^m(u, v-1, n, j)$$

$$-\frac{\lambda(j)I\{j \neq 0\}}{\Lambda+\alpha} V^m(u, v-1, n, 0) \tag{6-4}$$

$$+\frac{(\Lambda - u_4\mu_1 - \mu_2 I\{j=0\} - \lambda(j)I\{j \neq 0\})}{\Lambda+\alpha} V^m(u-1, v-1, n, j)$$

$$+\frac{\lambda(j)I\{j \neq 0\}}{\Lambda+\alpha} V^m(u-1, v-1, n, 0) \tag{6-5}$$

$$-\frac{\mu_2 I(j=0)}{\Lambda+\alpha}(\Delta^m(u, 0, n-1, j)I\{v-1 = s-1\}$$

$$+\Delta^m(u, v, n, j)I\{v-1 < s-1\}) \;. \tag{6-6}$$

The proof will be based on manipulating these expressions so that the induction hypothesis can be used. We shall tackle the case $(u, v, n, j) = (0, 0, 1, 0)$ separately. In all other cases, by using the induction hypothesis, only the terms involving u_i's, $i = 1, 2, 3, 4$, have to be analyzed to see whether the difference $\Delta^{m+1}(u, v, n, j) - \Delta^{m+1}(u, v-1, n, j)$, i.e., the sum of terms in (5) and (6), is non-negative. In other words, if u_1, u_2, u_3 and u_4 were all equal to 1 then grouping the terms in (5) and (6), and applying the induction hypothesis will be sufficient proof of the lemma (see case (i) below), but analysis is necessary when some of the controls do not equal 1. Using Lemma 1, as well as the induction hypothesis, there are only five cases (out of a possible 16 cases) to consider. A complete analysis is given for case (i), and proof showing $\Delta^{m+1}(u, v, n, j) - \Delta^{m+1}(u, v-1, n, j)$ is non-negative for the remaining cases is given in Glassey, Seshadri, and Shanthikumar [1996a].

Case (i) $(u_1, u_2, u_3, u_4) = (1, 1, 1, 1)$: The terms in (5-1) cancel with those in (6-1). Combining the terms in (5-2) & (5-3) with those in (6-2) & (6-3), and use of the induction hypothesis yields,

$$\frac{\mu_1}{\Lambda+\alpha}\Delta^m(u, v, n, j) - \frac{\mu_1}{\Lambda+\alpha}\Delta^m(u, v-1, n, j)$$

$$= \frac{\mu_1}{\Lambda+\alpha}(\Delta^m(u, v, n, j) - \Delta^m(u, v-1, n, j)) \geq 0 \;.$$

Using the induction hypothesis and combining the terms in (5-6) & (6-6) gives,

$$+\frac{\mu_2 I(j=0)}{\Lambda+\alpha}(\Delta^m(u, v+1, n, j) - \Delta^m(u, v, n, j)) \geq 0$$

Similarly, using the induction hypothesis and adding the terms in (5-4) & (5-5) and (6-4) & (6-5) gives,

$$+\frac{(\Lambda - \mu_1 - \mu_2 I\{j=0\} - \lambda(j)I\{j \neq 0\})}{\Lambda+\alpha}(\Delta^m(u, v, n, j) - \Delta^m(u, v-1, n, j))$$

$$+\frac{\lambda(j)I\{j \neq 0\}}{\Lambda+\alpha}(\Delta^m(u, v, n, 0) - \Delta^m(u, v-1, n, 0)) \geq 0 \ .$$

Analysis of the Boundary Case: The problem with the boundary state $(u, v, n, j) = (0, 0, 1, 0)$ arises due to transitions at the second station, i.e. *terms involving* μ_2. This is called a boundary state because there are no states below this with v positive. When (u, v, n, j) is $(0, 0, 1, 0)$ the four states in the lemma are $(0, 0, 1, 0)$, & $(r-1, 0, 0, 0)$ and the pair $(0, s-1, 2, 0)$, & $(r-1, s-1, 1, 0)$. The terms involving μ_2 in (5) and (6) can be grouped into,

$$(V^m(0, 1, 1, 0) - V^m(r-1, 0, 0, 0)) - (V^m(0, 0, 1, 0) - V^m(r-1, 0, 0, 0))$$

$$= V^m(0, 1, 1, 0) - V^m(0, 0, 1, 0)$$

Appealing to proposition 1, that inserted idleness is not beneficial, we have $(V^m(0, 1, 1, 0) - V^m(0, 0, 1, 0) \geq 0$. ■

Lemma 3: $V(a, v, n, j) - V(b, v, n, j) \leq V(a, v, n, 0) - V(b, v, n, 0)$, for $a \geq b$, $n > 0$.

Proof: Consider the four systems:

System I: Starts with initial state (a, v, n, j) and is operated using the optimal control policy.
System II: Starts with initial state $(a, v, n, 0)$ and is operated using the control policy for system I till a random time, denoted by $\mathscr{J}$ and to be specified in the proof. Thereafter the system is controlled using the optimal control policy.
System III: Starts with initial state (b, v, n, j) and is operated using the control policy for system IV till the same random time $\mathscr{J}$ as in system II and thereafter is controlled using the optimal control policy.

System IV: Starts with initial state $(b, v, n, 0)$ and is operated using the optimal control policy.

Let $P(x)$, $x = I, II, III, IV$ denote the expected discounted value of profits less holding costs incurred in system x, over the infinite time horizon. First we note that $V(a, v, n, j) - V(b, v, n, j) \leq P(I) - P(III)$, as the control policy used for operating system III is sub-optimal. Similarly, the difference $P(II) - P(IV) \leq V(a, v, n, 0) - V(b, v, n, 0)$. The amount of processing done at the first station in systems I and II will be assumed to be identical in every finite time interval in $[0, \mathscr{I}]$. Similarly in systems III and IV, the amount of processing done at the first station will be assumed to be the same in every finite time interval in $[0, \mathscr{I}]$. Let T_j be the duration of the initial down period in systems I and III and be exponentially distributed with parameter λ_j, denoted as $T_j \sim \exp(\lambda_j)$. Let $Z \sim \exp(\zeta)$ be the duration of the initial up period in systems II and IV. Let l be the type of the first failure in systems II and IV. Let the duration of this failure be $T_l \sim \exp(\lambda_l)$ in both these systems. Let $P(x, [t_1, t_2], \varpi)$ be the discounted profit less holding cost in system x, $x = I, II, III, IV$ over the time interval $[t_1, t_2]$ *over a particular sample path.*

Let us first discard the possibility that during $[0, Z]$, the second station might have idled in either system II or IV. Let T_I be the first time before Z, that in either system II or IV the second station became idle. Equate the processing done at the second station in systems II and IV during the interval $[0, T_I]$. At time T_I let the states of the four systems be (x, v, n, j), $(x, 0, 0, 0)$, (z, v, n, j), and $(z, 0, 0, 0)$ respectively. We shall also assume that $x > z$. The reason is, if $x = z$ then systems I and III have coupled and so have II and IV, and $x < z$ will imply that $x = z$ was true at some previous instant but after time zero – so the systems would have coupled. As $x > z$, by Lemma 2, $V(x, 0, 0, 0) - V(z, 0, 0, 0)$ is not smaller than $V(x, v, n, 0) - V(z, v, n, 0)$. So we can resume comparing the four systems after time T_I, with initial states (x, v, n, j), $(x, v, n, 0)$, (z, v, n, j), $(z, v, n, 0)$ – start the argument all over from time T_I onwards. A formal argument for establishing this can be given, see case (i) below. From now onwards we assume that systems II and IV have *not* idled during $[0, Z]$. We shall therefore assume that processing done at the second station in both these systems are equal during time $[0, \min(Z, \mathscr{I})]$. We need to consider two cases now.

Case (i): $T_j \leq Z$: See Figure 2(a). In this case we shall terminate the coupling at time T_j, i.e. set $\mathscr{I} = T_j$. During this time period over a sample path where the second station has not idled,

$$P(I, [0, T_j], \varpi) - P(III, [0, T_j], \varpi) = P(II, [0, T_j], \varpi) - P(IV, [0, T_j], \varpi) \ .$$

Let the states of the four systems be (x, v, n_1, j), $(x, z, n_2, 0)$, (y, v, n_1, j), $(y, z, n_2, 0)$ at time T_j. Once again we argue (proof omitted) that we only need to consider $x > y$. Using Lemma 2, and the fact that $(s - 1 - z + n_2(s - 1))$ is not

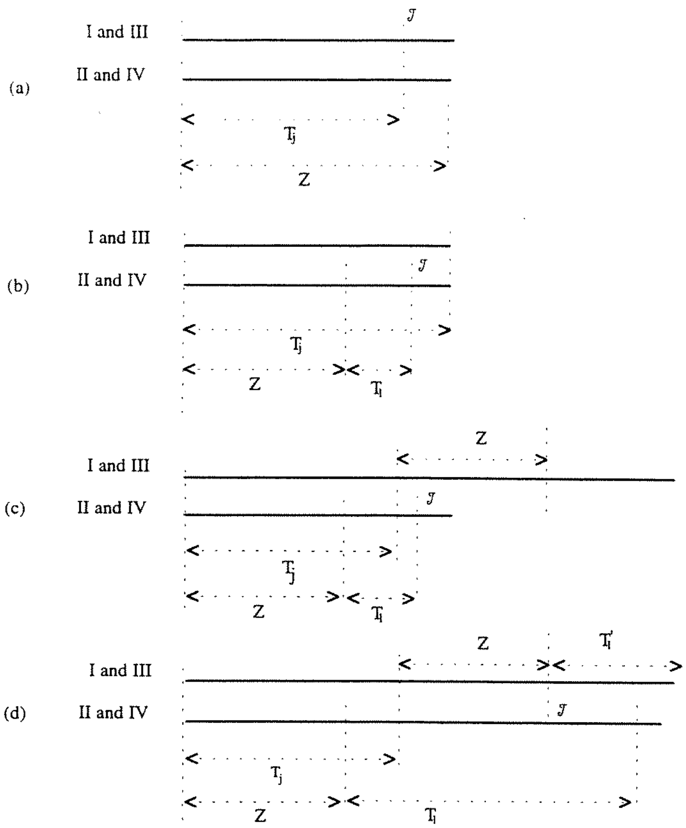

Fig. 2. Proof of proposition 3

larger than $(s - 1 - v + n_1(s - 1))$, the statement of the lemma is seen to be true on these sample paths.

We now formalize this argument. Let us denote A to be the event that $\{x > y$, *the second machine has not idled during* $[0, T_j]$, *and* $\{T_j \leq Z\}\}$. Let $V(a, v, n, j, A)$, $V(b, v, n, j, A)$, $V(a, v, n, 0, A)$, and $V(b, v, n, 0, A)$ be the expectations of the value received (till time infinity) using the optimal control policy and starting out from the four states (a, v, n, j), (b, v, n, j), $(a, v, n, 0)$, and $(b, v, n, 0)$ on the set A. Let $\Pr(A)$ be the probability of event A happening. Let $P(x, [0, T_j], A)$ be the expected value of discounted profits less holding cost incurred on the set in which event A happens in system x, $x = I, II, III, IV$. Let $E(\rho^{T_j})$ be the expected value of the discount factor till time T_j. Then,

$$
\begin{aligned}
&V(a, v, n, j, A) - V(b, v, n, j, A) \\
&\quad \leq P(I, [0, T_j], A) - P(III, [0, T_j], A) \\
&\qquad + E(\rho^{T_j})(V(x, v, n_1, 0) - V(y, v, n_1, 0)) \Pr(A)
\end{aligned}
$$

$$\le P(II, [0, T_j], A) - P(IV, [0, T_j], A)$$

$$+ E(\rho^{T_j})(V(x, z, n_2, 0) - V(y, z, n_2, 0)) \Pr(A)$$

$$\le V(a, v, n, 0, A) - V(b, v, n, 0, A) \ .$$

Case (ii): $T_j > Z$: We equate the processing done in the second station in systems *II* and *IV* during $[0, Z]$. There are two sub-cases to consider.

(a) $T_j > T_l$: We stop the coupling at time $Z + T_l$, i.e. set $\mathscr{J} = Z + T_l$. If the second machine is still down in systems *I* and *III* at this instant (see Figure 2(b)), let the states of the four systems be (x, v, n_1, j), $(x, z, n_2, 0)$, (y, v, n_1, j), $(y, z, n_2, 0)$ at time $Z + T_l$. Arguing as done earlier, we may start the argument afresh from time $Z + T_l$, with the four states: (x, v, n, j), $(x, v, n, 0)$, (y, v, n, j), $(y, v, n, 0)$. Else if the second machine is operational at time $\mathscr{J}$, then let the states of the four systems be (x, v, n_1, j), $(x, z, n_2, 0)$, (y, v, n_1, j), $(y, z, n_2, 0)$ at time $Z + T_l$, see Figure 2(c). In this case we equate the processing done in the second stations in systems *I* and *III* during $[T_j, Z + T_l]$ with that done at the second station in systems *II* and *IV* during $[0, Z + T_l - T_j]$. Set $\mathscr{J} = Z + T_l$. We observe that on every sample path on which the event $\{T_j > T_l; T_j \le Z + T_l\}$ occurs and also on which the second machine has not idled in any of the four systems,

$$P(I, [0, Z + T_l], \varpi) - P(III, [0, Z + T_l], \varpi)$$

$$= P(II, [0, Z + T_l], \varpi) - P(IV, [0, Z + T_l], \varpi)$$

By Lemma 2, and using the fact that $(s - 1 - v + n_1(s - 1))$ is not smaller than $(s - 1 - z + n_2(s - 1))$, the claim follows on these sample paths. The formal proof is omitted.

(b) $T_j \le T_l$: Let the type of failure that occurs at time $T_j + Z$ in systems *I* and *III* be of type l, and of duration $T_l' \sim \exp(\lambda_l)$. We shall assume that T_l is independent of T_l'. At time $Z + T_j$ we stop the coupling, see Figure 2(d). Set $\mathscr{J} = Z + T_j$. Equate the processing done in the second stations in systems *I* and *III* during $[T_j, Z + T_j]$ with the amount of processing done by the second station in systems *II* and *IV* during $[0, Z]$. At time $\mathscr{J}$ let the states of the four systems be: (x, v, n_1, l), (x, z, n_2, l), (y, v, n_1, l), (y, z, n_2, l). As argued previously the difference,

$$P(I, [0, Z + T_j], \varpi) - P(III, [0, Z + T_j], \varpi)$$

$$= P(II, [0, Z + T_j], \varpi) - P(IV, [0, Z + T_j], \varpi) \ .$$

and $(s - 1 - v + n_1(s - 1))$ is not smaller than $(s - 1 - z + n_2(s - 1))$. The claim now follows by use of Lemma 2. The formal proof is similar to the one given for case (i) and is omitted. ■

Lemma 4: If $\lambda(j) < \lambda(l)$, then $V(a, v, n, j) - V(b, v, n, j) \leq V(a, v, n, l) - V(b, v, n, l)$, for $a \geq b$, and $n > 0$.

Proof: Consider the four systems:

System I: Starts with initial state (a, v, n, j) and is operated using the optimal control policy.
System II: Starts with initial state (a, v, n, l) and is operated using the control policy for system *I* till time T_l (defined below) and thereafter is operated using the optimal control policy.
System III: Starts with initial state (b, v, n, j) and is operated using the control policy for system *IV* till time T_l (defined below) and thereafter is operated using the optimal control policy.
System IV: Starts with initial state (b, v, n, l) and is operated using the optimal control policy.

Let the duration of downtime in systems *I* and *III* be T_j and in systems *II* and *IV* be T_l. Equate the processing done at the first station in systems *I* and *II* during the time period $[0, T_l]$ and do similarly for systems *III* and *IV*. This coupling will stop at time T_l. Because, $\lambda(l) > \lambda(j)$, we can generate random variables $T_j \sim \exp(\lambda(j))$ and $T_l \sim \exp(\lambda(l))$ such that $T_j > T_l$ almost surely using a standard construction. Let the states of the four systems at time T_l be: (x, v, n, j), $(x, v, n, 0)$, (y, v, n, j), $(y, v, n, 0)$. As argued in Lemma 3, we can discard the possibility that x is smaller than y. Lemma 4 now follows by use of Lemma 3. ■

Accurately Predictable Repair: Consider now the case when the time to repair the jth type of failure has an Erlang distribution with $t(j)$ phases and mean $t(j)/\lambda(j)$. This model is appropriate when the repair times can be predicted very accurately. In that case, we can vary just the number of phases of the Erlang distribution associated with each failure type. The state space of the controlled Markov chain is expanded to $(u, v, w, n, j, u = 0, 1, \ldots, r - 1; v = 0, 1, \ldots, s - 1, j = 0, 1, \ldots, k; w = 0, 1, \ldots, t(j) - 1)$ where u, v, n, j are as before and w is the number of stages of repair completed. If in addition we assume that $\lambda(j) = \lambda(k)$, $\forall j, k \neq 0$, then a theorem very similar to Theorem I may be proved. The result now is that the control is purely a function of the number of stages of repair to go, say γ, and the threshold levels are given by a non-increasing function $Z(\gamma)$.

4 The Multiple Machines Model

We can extend the model to include a two stage flow line, in which there are N_1 machines in the first stage feeding N_2 machines in the second stage. All machines in a stage are assumed to be identical. Both stages are subject to machine failure, with failure rates ζ_1 and ζ_2. We assume that there are k_1 types of failures possible for the first stage machines and k_2 types of failures for the second stage machines, and the type of the next failure is modeled as in section 3. The processing times are assumed to be distributed exponential with parameters μ_1 and μ_2 at the two stages. There is a per unit holding cost of Λc per unit time, from the time the processing at the first station has been completed till the job leaves the system (see discussion in section 3). Each job on leaving the two stage flow line gives a profit of p. All profits and costs are continuously discounted using a factor α. It is assumed that the processing rate of any first stage machine, if it is operational, can be controlled continuously in the range $[0, \mu_1]$. The control space is $[0, 1]^{N^{N_1}}$. To simplify notation we represent a feasible control as the N_1 dimensional vector $\boldsymbol{u}$, with components $u(i)$, $i = 1, 2, \ldots, N$, and that the ith machine's processing rate is, $u(i)\mu_1$, where

$$u(i) \begin{array}{l} \in [0, 1], \text{ if machine } i \text{ is operational} \\ = 0, \text{ if machine } i \text{ is under repair} \end{array}$$

Assume that if a machine breaks down, the processing of the part if any on that machine can be taken up by an available (or next available) machine. The objective is to maximize the mean discounted net profit over the infinite time horizon. As in section 3, we work with the equivalent Markov chain. The state space for the Markov chain is $S = \{n_1, n_2, n, \lambda_1, \lambda_2; n \geq 0\}$; where the first three components are integers and stand for the number of machines working in the first and second stations and the number of jobs downstream of the first station. The last two components give the vector of repair rates for the failed machines in the two stations. The components of the repair rates are denoted as $\{\lambda_1(j), j = 1, \ldots, N_1\}$ and $\{\lambda_2(j), j = 1, \ldots, N_2\}$. If a machine is operational then its repair rate is set to zero. Let e_j be the unit vector of appropriate dimension (depending on the context) with a one in the jth position. Then as in section 3 the single stage reward obtained from using the control $\boldsymbol{u}$ in the state $(n_1, n_2, n, \lambda_1, \lambda_2)$, and the optimal value function can be written as,

$$R(n_1, n_2, n, \lambda_1, \lambda_2; \boldsymbol{u}\} = -\frac{\sum_i u(i)\mu_1}{\Lambda + \alpha}(n + 1)c + \frac{(n_2 \wedge n)\mu_2}{\Lambda + \alpha}(p - (n - 1)c)$$

$$- \frac{(\Lambda - \sum_i u(i)\mu_1 - (n_2 \wedge n)\mu_2)}{\Lambda + \alpha} nc \ ,$$

and

$$\begin{aligned}
V(n_1, n_2, n, \lambda_1, \lambda_2) = \max_u \, & (R(n_1, n_2, n, \lambda_1, \lambda_2; \boldsymbol{u}) \\
& + \frac{1}{\Lambda + \alpha}(\textstyle\sum_i u(i)\mu_1 V(n_1, n_2, n+1, \lambda_1, \lambda_2) \\
& + (n_2 \wedge n)\mu_2 V(n_1, n_2, n-1, \lambda_1, \lambda_2) \\
& + n_1\zeta_1 \textstyle\sum_{\lambda_1'} p(\lambda_1') V(n_1 - 1, n_2, n, \lambda_1', \lambda_2) \\
& + n_2\zeta_2 \textstyle\sum_{\lambda_2'} p(\lambda_2') V(n_1, n_2 - 1, n, \lambda_1, \lambda_2') \\
& + \textstyle\sum_j \lambda_1(j) V(n_1 + 1, n_2, n, \lambda_1 - \lambda_1(j)e_j, \lambda_2) I\{n_1 \neq N_1\} \\
& + \textstyle\sum_j \lambda_2(j) V(n_1, n_2 + 1, n, \lambda_1, \lambda_2 - \lambda_2(j)e_j) I\{n_2 \neq N_2\} \\
& + (\Lambda - \textstyle\sum_i u(i)\mu_1 - (n_2 \wedge n)\mu_2 - n_1\zeta_1 - n_2\zeta_2 \\
& - \textstyle\sum_j \lambda_1(j) - \sum_j \lambda_2(j)) V(n_1, n_2, n, \lambda_1, \lambda_2))) \; ,
\end{aligned}$$

where we have used λ_i' and $p(\lambda_i')$ to denote a new vector of repair rates and the probability of its occurrence given the current vector is λ_i, $i = 1, 2$, and $X \wedge Y = \min(X, Y)$. The terms involving the control, $u(i)$'s, when collected together, give,

$$+\sum_i u(i)(-c + V(n_1, n_2, n+1, \lambda_1, \lambda_2) - V(n_1, n_2, n, \lambda_1, \lambda_2)) \; . \tag{7}$$

It is clear from equation (7) that either all machines at the first station will produce at the maximum rate or not at all. We shall call $(n_1, n_2, \lambda_1, \lambda_2)$ as the *state of the flow line*. The number of possible states of the flowline are finite, and their collection will be denoted as S_F.

Theorem II: For the two stage flow line when there are multiple but identical machines at each stage, exponential processing time distributions, multiple types of failure at both stages, failure and repair times distributed exponential, constant per unit holding costs once a job has completed its processing at the first stage, constant profit per job on completion, under continuous discounting; and when the objective is to maximize the expected discounted value of net profit over the infinite time horizon, the structure of the optimal policy is such that: For each possible state of the flow line, $s \in S_F$ there exists a threshold level

$Z(s)$, such that if the number of jobs downstream of the first stage is less than or equal to $Z(s)$ it is optimal to process jobs at the highest possible rate at all machines at the first stage.

Proof: The proof follows from equation (7) and Lemma 5 in Glassey, Seshadri, and Shanthikumar [1996a], where it is shown that for each value of $(n_1, n_2, \lambda_1, \lambda_2)$, the value function $V(n_1, n_2, n, \lambda_1, \lambda_2)$ is concave in n. ■

Remark: Erlang repair times pose no problem in the multiple machine case too. However ordering the threshold levels is a difficult problem. This is because the vector of repair rates introduces only a partial ordering of the states of the flowline.

5 Numerical Results

The structural properties described in the previous sections by themselves will not help in answering questions such as: How does the distribution of processing times affect the optimal threshold levels? How will threshold levels change when failures occur? When is a policy that uses feedback on the machine status and type of failure clearly superior to one that does not use any feedback? What is the effect of multiple machines on threshold levels? And how to combine the theoretical results into a practical algorithm for release control? Several experiments were conducted to answer these questions. In this section we summarize the key observations from these experiments, describing a phenomenon called the separation effect and the value of using information about machine failures. The reader is referred to Glassey, Seshadri, and Shanthikumar [1996a] for the details of other experiments. The stations in the two stage lines contain multiple but always identical machines. The job processing time distributions are exponential and like in sections 3 and 4, the time to fail and repair are assumed to be exponentially distributed. Both stages are failure prone. The unit profit, holding cost, and the continuous discount factor are denoted as p, c, and r in the tables, typically these parameters were set to, $p = 10$, $c = 0.3$ and $r = 0.1$. We have used three combinations of processing rates, namely (1.5, 0.5), (1, 1), and (0.75, 1.25); where the first number denotes the processing rate of a typical machine at the first station and the second number that of a machine in the second station. This helps us compare situations in which (i) a fast station feeds a slower one, (ii) a balanced setup and (iii) one where a slow station feeds a faster station. These combinations are labeled as *processing rates* in the tables. The *machine state* is used to denote how many machines are working in each stage and also include where appropriate the repair rates of the failed machines. The machine state in the tables, indicated as $i \to j$, denotes i machines working in station 1 and feeding j operational machines in station 2.

Separation Effect: Consider just the results of Experiments I and VI shown in Table I. In Experiment I the machines are operational about 90% of the time and the repair rate is about a fifth of the processing rate. The optimal threshold levels for this set up differ quite a bit depending on the machine state. This is a phenomenon we call the *separation effect*. Contrast this with the results of Experiment VI, where the repair rate is comparable to the processing rate of the second station. In this experiment, there is almost no difference in the threshold levels across machine states. The results from Experiment VII show that even when the machines are nearly 100% reliable, there is no significant change in threshold levels when the repair rate is comparable to the processing rates of the second station. The gradual impact of the repair rate on threshold levels can be traced by examining the results in the following sequence: Experiment IV → I → V → VI → VII. Moreover, on comparing the threshold levels in Experiments I, II and III, we find that the parameters of failure times are apparently less important compared to the repair rate of the second station in determining the threshold levels.

Value of Information About Machine Failures: A static policy is one that does not react to machine failures, but keeps the threshold level unchanged. A dynamic policy is one that reacts to machine failures. In Table II, we show the results for a two stage flowline with two machines at each stage. For different combinations of processing rates, machine failure and repair rates we investigate the difference between the best static policy and the optimal policy. There are two measures used in making comparisons, for the first we simply compute the difference between the average profit rates per transition, and for the second we compute the largest deviation in the value functions over all recurrent states. In Table II we also show the optimal threshold level when *all* machines are operational, along with the value of the best static threshold level. It turns out that excepting in the case when repair rates are very small these two values are the same. This is not surprising as a static policy should be correct most of the time, and most of the time all machines in the examples in Table II are working. We see from the table that a dynamic policy is beneficial when, the processing rate of the first station is comparable to or is greater than that of the second stage, the fraction of down time of the second stage is not too small, and when the repair rate is small compared to the second station's processing rate.

Table I. Sensitivity of threshold level to repair rate
DATA: (2 machines per station, $c = 0.3$, $p = 10$, $r = 0.1$)

MACHINE STATE
(# working m/cs 1st stage → # working m/cs 2nd stage)

	Processing Rates	Failure Rate	Repair Rate	1 → 1	1 → 2	2 → 1	2 → 2	1 → 0	2 → 0
Experiment I	(1, 1)	0.01	0.1	9	14	5	9	4	2
Experiment II		0.033	0.1	8	13	6	10	4	2
Experiment III		0.044	0.1	7	13	6	10	4	1
Experiment IV		0.001	0.01	6	17	4	9	0	0
Experiment V		0.02	0.2	9	13	7	10	6	4
Experiment VI		0.08	0.8	9	10	8	9	7	7
Experiment VII		0.001	0.8	9	10	8	9	8	7
Experiment I	(1.5, 0.5)	0.01	0.1	3	5	2	4	0	0
Experiment II		0.033	0.1	3	5	2	4	1	1
Experiment III		0.044	0.1	3	5	2	4	1	1
Experiment IV		0.001	0.01	2	5	2	4	0	0
Experiment V		0.02	0.2	3	5	2	4	1	1
Experiment VI		0.08	0.8	3	4	3	4	3	2
Experiment VII		0.001	0.8	3	4	3	4	3	2
Experiment I	(0.75, 1.25)	0.01	0.1	15	23	11	18	8	5
Experiment II		0.033	0.1	14	21	11	17	8	5
Experiment III		0.044	0.1	14	20	12	17	8	6
Experiment IV		0.001	0.01	10	26	6	18	0	0
Experiment V		0.02	0.2	16	21	14	18	12	10
Experiment VI		0.08	0.8	15	17	15	16	14	13
Experiment VII		0.001	0.8	16	18	16	17	15	14

Table II. Dynamic versus static policies
Data: (Profit = 10, cost rate = 0.3, All failures same type; Two machines at each station; # of iterations = 300, Average Cost Criterion)

Processing Rates	Repair Rate	Failure Rate	Threshold		Profit Rate		% Diff.	MAX Diff. Value Fn.
			Optim All Up	Static Level	Optimal Policy	Static Policy		
(1, 1)	0.05	0.005	11	11	3.00064	2.97538	−0.84	26.35
		0.01	12	11	2.59045	2.55493	−1.37	22.36
		0.05	12	11	1.07578	1.03820	−3.49	9.85
	0.1	0.001	12	12	2.65007	2.63653	−0.51	6.67
		0.05	12	12	1.58421	1.56450	−1.24	3.75
		0.01	12	12	0.98526	0.96722	−1.83	2.10
	0.2	0.02	12	12	2.16535	2.16071	−0.21	1.22
		0.1	12	12	1.31780	1.31144	−0.48	0.48
		0.2	12	12	0.83372	0.82633	−0.89	0.36
(1.5, 0.5)	0.05	0.005	4	4	1.63517	1.60742	−1.70	6.76
		0.01	4	4	1.45064	1.41573	−2.41	5.60
		0.05	4	4	0.62243	0.59594	−4.26	2.77
	0.1	0.001	4	4	1.43685	1.41434	−1.57	2.24
		0.05	4	4	0.90170	0.87443	−3.02	0.99
		0.01	4	4	0.55632	0.53871	−3.17	0.41
	0.2	0.02	4	4	1.16163	1.14637	−1.31	0.65
		0.1	4	4	0.72768	0.71586	−1.62	0.11
		0.2	4	4	0.45384	0.44563	−1.81	0.02

Table II. (*Cont.*)

Processing Rates	Repair Rate	Failure Rate	Threshold		Profit Rate		% Diff.	MAX Diff. Value Fn.
			Optim All Up	Static Level	Optimal Policy	Static Policy		
(0.75, 1.25)	0.05	0.005	32	14	2.67903	2.67467	−0.16	9.55
		0.01	31	14	2.35333	2.34647	−0.29	8.31
		0.05	24	14	1.02716	1.01456	−1.23	3.55
	0.1	0.001	31	31	2.37971	2.37624	−0.15	4.13
		0.05	26	26	1.50521	1.49030	−0.99	2.84
		0.01	26	26	0.96292	0.95141	−1.20	1.47
	0.2	0.02	31	31	1.94304	1.94246	−0.03	0.28
		0.1	25	25	1.25144	1.24757	−0.31	0.42
		0.2	20	20	0.81560	0.81025	−0.66	0.35

6 Summary

We have shown that threshold type policies, with threshold levels nested according to the rate of repair are optimal for controlling production in a class of failure prone two stage flowlines having a single machine at each stage. Threshold type policies are also optimal under our assumptions for models with multiple machines and types of failures. Numerical investigations reveal that the threshold level corresponding to a given machine state is dependent mostly upon, the ratio of the repair rate to the processing rate of the second station. Dynamic policies that react to machine failures are most beneficial when, the second machine is a bottleneck, the fraction of downtime is not too small, and when the ratio of the repair rate to the processing rate of the second station is small. We plan to extend and apply these results to controlling failure prone reentrant flowlines in future work.

References

Akella A; Kumar PR (1986) Optimal control of production rate in a failure prone manufacturing system. IEEE Trans. Automatic Control 31:116–126

Birge JR, Glazebrook KD (1988) Assessing the effects of machine breakdowns in stochastic scheduling. OR Letters 7, 6:267–271

Boukas E, Haurie A (1990) Manufacturing flow control and preventive maintenance: A stochastic control approach. IEEE Trans. Automatic Control 36, 5:1024–1031

Buzacott JA, Shanthikumar JG (1992) PAC: A general approach for coordinating production in multiple-cell manufacturing systems. Prod. Oper. Manag.: 34–52

Buzacott JA, Shanthikumar JG (1993) Stochastic models of manufacturing systems. Prentice-Hall, New York

Gershwin SB (1988) Hierarchical flow control: A framework for scheduling and planning discrete events in manufacturing systems. IEEE Proc. Special Issue on Discrete Event Systems

Glassey CR, Seshadri S, Shanthikumar JG (1996) Linear control rules for production control of semiconductor fabs. IEEE Trans. Semiconductor Manufacturing 9, 4:536–549

Glassey CR, Seshadri S, Shanthikumar JG (1996a) Using information about machine failures to control flowlines. Working Paper. Department of Statistics & Operations Research, New York University

Glazebrook KD (1987) Evaluating the effect of machine breakdowns in stochastic scheduling problems. Naval Research Logistics 34, 3:319–335

Gershwin SB (1994) Manufacturing systems engineering. Prentice Hall, NJ

Harrison JM (1988) Brownian models of queueing networks with heterogeneous customer populations. In: Fleming W, Lions PL (eds) Stochastic differential systems, stochastic control theory and applications, IMA Volume 10, Springer-Verlag, Berlin: 147–186

Harrison JM, Wein LM (1990) Scheduling networks of queues: Heavy traffic analysis of a two-station closed network. Operations Research 38, 6:1052–1064

Hildebrandt RR (1980) Scheduling machining systems when machines are prone to failures. Ph.D. Thesis, MIT, Dept. of Astronautics and Aeronautics

Kumar PR (1993) Re-entrant lines. Queueing Systems 13:87–110

Kumar PR (1994) Scheduling semiconductor manufacturing plants. IEEE Control Systems: 33–40

Kushner HJ (1990) Numerical methods for stochastic control problems in continuous time. SIAM J. Control and Optimization 28, 5:999–1048

Lippman SA (1975) Applying a new device in the optimization of exponential queuing systems. Operations Research 23, 4:687–710

Nurani RK (1995) Role of in-line process monitoring and control in wafer fab management. Ph.D. Dissertation, GSIA, Carnegie Mellon University, Pittsburgh, PA 15213

Nurani RK, Seshadri S, Shanthikumar JG (1995) Optimal control of a single stage production system subject to random process shifts. To appear in Operations Research

Pinedo M, Ramouz E (1988) A note on stochastic scheduling on a single machine subject to breakdown and repair. Prob. Eng. Inform. Sciences 2:41–49

Pinedo M (1995) Scheduling: Theory, algorithms and systems, Prentice Hall, NJ

Posner MJM, Berg M (1988) Analysis of a production-inventory system with unreliable production facility. Operations Research Letters 8:339–345

Ross SM (1970) Applied probability models with optimization applications. Holden-Day, San Francisco

Seshadri S (1993) Stochastic models of semiconductor fabs. Ph.D. Thesis, Haas School of Business Administration, University of California, Berkeley, CA 94720

Seshadri S, Shanthikumar JG (1994) Allocation of chips to wafers in a production problem of semiconductor chips. To appear in Operations Research

Sharifnia A (1988) Production control of a manufacturing system with multiple machine states. IEEE Trans. AC 33, 7:620–625

Thomas LJ (1993) Stockless and fast production: Review and research agenda. In: Sarin RK (ed.), Perspectives in operations management: Essays in honor of Elwood S. Buffa, Kluwer Academic Publishers, Boston, MA

Uzsoy R, Lee CY, Martin-Vega LA (1992) A review of production planning and scheduling models in the semiconductor industry. 1. System Characteristics, Performance Evaluation and Production Planning. IIE Trans. 24, 4:47–60

Uzsoy R, Lee C, Martin-Vega LA (1994) A review of production planning and scheduling models in the semiconductor industry, Part II: Shop-floor control. IIE Transactions 26:44–55
Veatch MH, Wein LM (1993) Optimal control of a two-station tandem production/inventory system. Manuscript, Operations Research Center and Sloan School of Management, M.I.T.
Wein LM (1990) Optimal control of a two-station Brownian network, Math. Operations Research 15, 2:215–242