

COMPUTATION OF PRODUCTION CONTROL POLICIES
BY A DYNAMIC PROGRAMMING TECHNIQUE

by

Joseph Kimemia
Bell Telephone Laboratories
Holmdel, NJ 07733

Stanley B. Gershwin

Dimitri Bertsekas

ABSTRACT

The problem of production management for an automated manufacturing system is described. The system consists of machines that can perform a variety of tasks on a family of parts. The machines are unreliable, and the main difficulty the control system faces is to meet production requirements while machines fail and are repaired at random times. A multi-level hierarchical control algorithm is proposed which involves a stochastic optimal control problem at the first level. Optimal production policies are characterized and a computational scheme is described.

COMPUTATION OF PRODUCTION CONTROL POLICIES
BY A DYNAMIC PROGRAMMING TECHNIQUE

by

Joseph Kimemia
Bell Telephone Laboratories, NP2D108, Holmdel, NJ 07733

Stanley B. Gershwin
Dimitri Bertsekas

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology, Cambridge, MA 02139

ABSTRACT

The problem of production management for an automated manufacturing system is described. The system consists of machines that can perform a variety of tasks on a family of parts. The machines are unreliable, and the main difficulty the control system faces is to meet production requirements while machines fail and are repaired at random times. A multi-level hierarchical control algorithm is proposed which involves a stochastic optimal control problem at the first level. Optimal production policies are characterized and a computational scheme is described.

This research was carried out in the M.I.T. Laboratory for Information and Decision Systems with support extended by the National Science Foundation Grant DAR78-17826 and ECS 79-20834.

1. Introduction

Flexible Manufacturing Systems (FMS) are being introduced in an effort to increase productivity in the manufacture of small and medium sized batches of related parts [Cook, 1975]. An FMS is a set of one or more workcenters each consisting of workstations at which operations are carried out on workpieces. A material handling system transports parts to and from the workstations. Overall control is exercised by one or more computers which control the transportation mechanism, the scheduling of operations and the downloading of appropriate control programs to the workstations [Lerner, 1981].

The FMS produces a family of parts that are related by similar operational requirements. The members of the part family are manufactured simultaneously. The flexibility of the system allows parts the choice of one or more stations for each operation. This allows production to continue when a workstation is out of service because of a failure or maintenance.

The ability of an FMS to produce different part types simultaneously results in increased productivity because of reduced part inventories and increased utilization of available time at the workstations. However, to reap the full benefit of flexible automation, careful planning and control of production is necessary [Hutchinson, 1977]. This is made difficult by the fact that the workstations in a flexible workcenter are prone to failures. Production planning and control algorithms must take into account the reliability of the workstations. Otherwise, the advantage of reduced inventories offered by flexible automation may be lost.

In most implementations, flexible workcenters are part of a multi-stage manufacturing system. The parts coming into the workcenter have undergone one or more processing stages. The output is a family of parts that are assembled into final products or sub-assemblies.

The management of a manufacturing firm makes production plans for finished products.

From the resulting master production schedule, the requirements for all the components that go into the final product can be made [Orlicky, 1979]. The various departments responsible for the manufacture of the components schedule their activities so as to meet the demands dictated by the master production and the requirements plans [Hitomi, 1979].

In an automated flexible workcenter, most operational decisions are made by one or more control computers. It is important therefore that control algorithms should generate production schedules which satisfy the demand requirements placed on the workcenter and exercise control over the system so that the output conforms to the schedule.

In a workcenter of reasonable size, the material flow process is complex and does not lend itself to direct centralized control. A multi-level control algorithm is proposed. The hierarchy is illustrated in Figure 1, in which the workcenter controller is embedded in the larger hierarchy of production management. The objective of the controller is to satisfy a known, possibly time varying demand for a family of parts that is dictated by the master production plan.

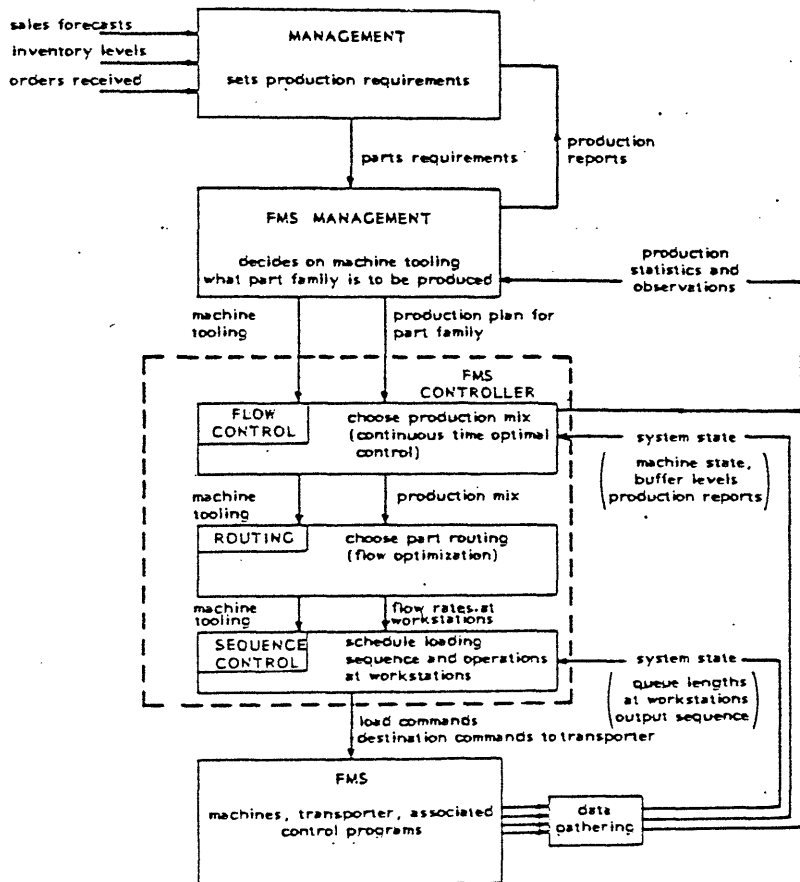


Figure 1.
The Hierarchical Control Algorithm

The flow control level of the algorithm adjusts the instantaneous production rate of the workcenter. The flow regulation is done continuously so as to respond to random failures and repairs of the workstations. The flow control model is shown in Figure 2. The part flow is modelled by a continuous process. The workcenter is modelled as a processing system whose state depends on the operational state of the workstations. The productive capacity of the workcenter therefore varies with time. The downstream buffers hold finished workpieces and serve to decouple the workcenter from downstream production stages.

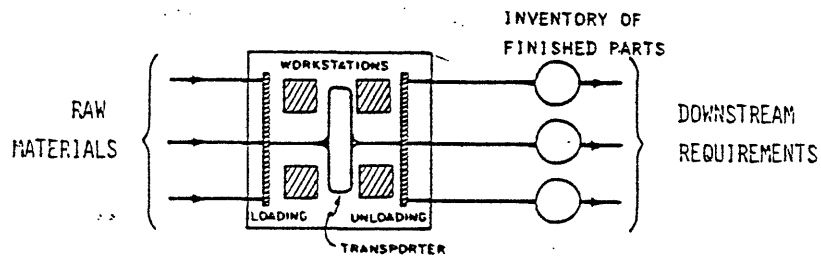


Figure 2. The Flow Control Model of the Workcenter

Within the workcenter, parts often have a choice of one or more stations for some of the required operations. The routing algorithm determines the proportion of parts that go to each

station whenever such a choice is available.

The sequence controller has the task of scheduling the introduction of pieces into the system and controlling their movement between workstations. The objective of the sequence control schedule is to maintain the throughput and the proportions determined by the flow and routing control algorithms.

The hierarchical controller is designed for application in systems where the mean time between failures and to repair is long compared to the time to produce a single part. This allows the controller to account for workstation reliability at the flow control level. Throughput rates determined by the flow controller are at all times feasible for the current system configuration. This guarantees feasible solutions to the routing and sequence control problems.

In this paper, we examine the flow control level of the hierarchy. Section 2 describes the flow control model of the workcenter and formulates the stochastic optimal flow control problem. We show that for each failure state of the workcenter, optimal flow control policies are piece-wise constant functions of the downstream buffer levels. In Section 3, we develop an estimate based (EB) sub-optimal control policy. The EB-controller uses estimates of the optimal value function to generate feedback control laws which like the optimal policies are piece-wise constant.

A hierarchical control scheme has been proposed by Hildebrandt [1980] for the problem of minimizing the time to produce a given quantity of parts. A static optimization level of the hierarchy gives part routing for all failure conditions. However, feedback information on the current state of production is not utilized. Olsder and Suri [1980] use a dynamic programming formulation for the minimum time production problem. In this case, a feedback policy results which depends on the current failure state and production levels.

Hahne [1981] and Tsitsiklis [1982] study the problem of maximizing throughput in a system in which parts can be routed from an upstream machine to one of two unreliable downstream

machines. They show that optimal policies are piece-wise constant functions of intermediate buffer levels. Calculation of exact optimal policies for the three machine system has large computational requirements.

2. The Flow Control Model

The workcenter consists of M workstations on which N parts are produced. Let $u(t) \in R^N$ be the production rate vector of the system. The downstream demand is $d(t)$ and is known over an interval of time $[0, T]$. Define $x(t) \in R^N$ by the following differential equation

$$\frac{dx(t)}{dt} = u(t) - d(t) \quad (1)$$

The vector $x(t)$, termed the buffer state, measures the cumulative difference between production and demand for the part family.

The state of the workstations is described as the machine state and is denoted by an M -tuple of binary variables $\alpha(t)$ with the m th component defined by

$$\alpha_m(t) = \begin{cases} 1 & \text{if station } m \text{ is operational} \\ 0 & \text{otherwise} \end{cases}$$

The times between failures and the times to repair are modelled by independent exponentially distributed random variables with means $1/p_m$ and $1/r_m$ respectively. The machine state can thus be modelled by an irreducible Markov chain with 2^M states.

Let S be an index set corresponding to the machine states. Then for $i, j \in S$ and $i \neq j$,

$$P(\alpha(t + \delta t) = j \mid \alpha(t) = i) = \lambda_{ij} \delta t \quad (2)$$

By definition, $\lambda_{ii} = -\sum_{j \neq i} \lambda_{ij}$. The transition rates λ_{ij} are functions of the failure and repair rates of the workstations.

To define the capacity set $\Omega(\alpha(t))$ of the workcenter, consider the machine state $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_M)$. Let $w_{nm}^k \geq 0$ be the rate at which type n parts are sent to station m for operation k . Let τ_{nm}^k be the time required to complete the operation. Since it is assumed that no material is accumulated within the system, the number of type n parts undergoing operation k per unit time is equal to the throughput u^n for the part. This is expressed as

$$\sum_m w_{nm}^k = u^n \quad \text{for all } n, k \quad (3)$$

The limited capacity of station m is expressed as [Kimemia, 1982]

$$\sum_n \sum_k w_{nm}^k \tau_{nm}^k \leq \alpha_m \quad (4)$$

The control constraint set $\Omega(\alpha)$ is the projection of the set of feasible w_{nm}^k into \mathcal{R}^N via (3). From (3) and (4), $\Omega(\alpha)$ is a convex polyhedral set lying entirely in the positive orthant and containing the origin.

The flow control problem can now be stated. Given a workcenter as described above, an initial buffer state $x(0)$ and machine state $\alpha(0)$, we wish to specify a production plan $u(t) \in \Omega(\alpha(t))$ for $0 \leq t \leq T$ that minimizes the performance index

$$J_0^i(x) = E \left\{ \int_0^T g(x(t)) dt \mid x(0) = x, \alpha(0) = i \right\} \quad (5)$$

subject to (1) and (2). The function $g(x)$ penalizes the controller for failing to meet demand and for keeping an inventory of parts in the downstream buffer. It is given by

$$g(x) = \sum_{n=1}^N g_n(x^n) \quad (6)$$

where $g_n(x^n)$ are convex functions satisfying

$$\lim_{|z| \rightarrow \infty} g_n(z) = \infty \quad (7)$$

The performance index (5) is thus the expected total cost over the period (0,T) given the initial conditions x and i .

We define the admissible control set U_A to be the set of all functions $u : R^N \times S \times R \rightarrow \Omega(\alpha)$ such that for all $i \in S$, discontinuities of $u(x,i,t)$ divide $R^N \times R$ into a finite set of regions within which $u(x,i,t)$ is a continuous function of x and t . Furthermore, the boundary separating two regions should be continuous and piece-wise smooth.

Intuitively, whenever the buffer state crosses a discontinuity in the control policy, the control vector $u(x,i,t)$ experiences a jump. Physically, there is a limit to the number of such jumps that can occur in a finite time interval. The limitation of possible discontinuities in admissible control policies eliminates control laws with an infinite number of discontinuities in the buffer state trajectory. Piece-wise smooth boundaries again ensure that any trajectory which moves along a boundary also experiences a control vector with a finite number of discontinuities in any finite time interval.

2.1 Characterization of Optimal Control Policies

We provide an informal presentation of the main results that characterize optimal policies. For a related mathematically rigorous analysis, we refer to Tsitsiklis [1982]. Define the expected cost when the policy u is applied in the interval (t,T) with initial buffer level x and machine state i as

$$J_u(x,i,t) = E \left\{ \int_t^T g(x_u(s)) ds \mid x(t) = x, \alpha(t) = i \right\} \quad (8)$$

The expected cost (8) satisfies, for $i \in S$ and $x \in R^N$, [Rishel, 1975]

$$J_u(x, i, t) = \int_t^T \left\{ g(x_u^i(s; x, t)) + \sum_{j \in S} \lambda_{ij} J_u(x_u^i(s; x, t), j, s) \right\} ds \quad (9)$$

with $J_u(x_u^i(T; x, t), i, T) = 0$ for all $i \in S$. The integration in (9) is performed along the trajectory $x_u^i(s; x, t)$ which is the solution for $s \in (t, T)$ to the deterministic differential equation

$$\frac{dx}{ds} = u(x, i, s) - d(s) \quad (10)$$

with initial conditions $x(t) = x$ and $\alpha(t) = i$. It can then be shown that an optimal policy u^* and associated cost $J_{u^*}(x, i, t)$ satisfies [Rishel, 1975]

$$\min_{u \in \Omega(i)} \left\{ g(x) + \frac{\partial}{\partial t} J_{u^*}(x, i, t) + \frac{\partial}{\partial x} J_{u^*}(x, i, t) [u - d(t)] + \sum_{j \in S} \lambda_{ij} J_{u^*}(x, j, t) \right\} = 0 \quad (11)$$

Furthermore, the optimal cost-to-go function $J_{u^*}(x, i, t)$ is continuous in x and t for all i and convex in x for all i and t [Kimemia, 1982].

The partial differential equation (11) is linear in u and the control constraint set $\Omega(i)$ is polyhedral. Consequently, whenever the derivative $\frac{\partial}{\partial x} J_{u^*}(x, i, t)$ exists, the optimal control $u^*(x, i, t)$ takes a value at an extreme point of $\Omega(i)$. For each machine state and time t , an optimal policy therefore divides the buffer state into a set of regions within which the control is constant. However, the regions do not cover the whole space. If the derivative $\frac{\partial}{\partial x} J_{u^*}(x, i, t)$ does not exist or is orthogonal to a face of $\Omega(i)$ in some subset of R^N , then there is no unique minimizing value to (11). The optimal policy then depends on the extreme point policies in the regions neighboring the subset. This behavior is similar to singular control in deterministic optimal control problems.

3. An Estimate Based (EB) Control Scheme

3.1 The Approach

The optimal policy in the flow control problem is determined from the optimal value function $J_u^*(x, i, t)$ by the linear program (11). An optimal policy is a feedback law which for every machine state divides the buffer state space into regions within which the control is constant at an extreme point of the control constraint set.

The gradient $\frac{\partial}{\partial x} J_u^*(x, i, t)$ can be regarded as a weighting on part production for the optimal control law. The parts that are most vulnerable to machine failures have the most weight and so when there is a backlog, an optimal policy is more inclined to produce them than less vulnerable parts. The calculation of the optimal value function takes into account the relative costs of backlogs and inventory storage determined by the functions $g_n(x)$. Thus a part that has a high value index and is at the same time sensitive to machine failures would have correspondingly a large component in the value function.

The hedging point $x_{hi}(t)$, which is the minimum of $J_u^*(x, i, t)$ with respect to x , is the optimal buffer level with which to hedge against future failures. When demand is close to the capacity of the system, the hedging points are at high buffer levels because failures quickly result in deficits and recovery from a deficit is slow.

Optimal policies cannot be computed in practice because of the large dimension of the flow control problem. We need a practical method for calculating sub-optimal control laws which produce good results when used in the flow control level of the hierarchy.

Given convex functions $\psi(x, i, t)$ which are estimates of the optimal value function, consider the control policy \hat{u} determined by

$$\hat{u}(x, i, t) = \underset{u \in \Omega(i)}{\operatorname{argmin}} \left\{ \frac{\partial}{\partial x} \psi(x, i, t) \right\} u \quad (12)$$

The sub-optimal policy \hat{u} like an optimal policy divides the buffer state space into a set of regions in each of which it takes values at an extreme point of $\Omega(i)$.

The estimates $\psi(x, i, t)$ should exhibit the properties of the optimal value function described above. The value of the estimate should be largest for machine states with the smallest production capacity. The relative magnitudes of the components of the gradient $\frac{\partial}{\partial x} \psi$ should reflect both the relative value of parts and their vulnerability to machine failures. The minimum with respect to x of $\psi(x, i, t)$, which determines the hedging point for the sub-optimal policy, should be of a magnitude comparable to the optimal hedging buffer levels. If the estimates satisfy these criteria, we expect the sub-optimal policy to perform well and to meet demand requirements when they are close to system capacity. If the optimal value of the cost index is not sensitive to the location of the region boundaries, the cost $J_{\hat{u}}(x, i, t)$ corresponding to the estimate based control policy should be close to the optimal cost.

3.2 Calculation of the Estimates

The control constraint set $\Omega(i)$ is polyhedral, lies in the positive orthant and contains the origin. Define $\bar{H}(i)$ and $\underline{H}(i)$ to be sets such that

$$\bar{H}(i) = \left\{ u \in R^N \mid 0 \leq u^n \leq \bar{L}_{in} \right\} \quad n = 1, 2, \dots, N \quad (13)$$

$$\underline{H}(i) = \left\{ u \in R^N \mid 0 \leq u^n \leq \underline{L}_{in} \right\} \quad n = 1, 2, \dots, N \quad (14)$$

and

$$H(i) \subseteq \Omega(i) \subseteq \bar{H}(i) \quad (15)$$

$H(i)$ and $\bar{H}(i)$ are hypercubes, the former contained in $\Omega(i)$ and the latter containing the control constraint set. For example, Figure 3 shows the hypercubes for a sample control constraint set.




Figure 3. An Example of the Hypercubes for a Control Constraint Set

Define $\underline{\psi}(x, i, t)$ and $\bar{\psi}(x, i, t)$ by the following optimization problems.

$$\underline{\psi}(x, i, t) = \min_{u(t) \in \bar{H}(\alpha(s))} E \left\{ \int_t^T g(x(s)) ds \right\} \quad (16)$$

$$\bar{\psi}(x, i, t) = \min_{u(t) \in \bar{H}(\alpha(s))} E \left\{ \int_t^T g(x(s)) ds \right\} \quad (17)$$

both subject to (1) and (2) and with initial conditions $x(t)=x$ and $\alpha(t)=i$.

From (15), the following holds,

$$\underline{\psi}(x, i, t) \leq J_u^*(x, i, t) \leq \bar{\psi}(x, i, t) \quad (18)$$

Thus $\bar{\psi}$ and $\underline{\psi}$ are upper and lower bounds on the optimal value function. An estimate of J_u^* can be obtained by taking a convex combination of the lower and upper bounds.

The cost function $g(x)$ is given by (6) and is therefore separable. The constraints (13) and (14) affect each part separately. The optimization problems (16) and (17) are therefore decoupled and can be solved as a set of scalar problems one for each part.

The hypercubes $\underline{H}(i)$ and $\bar{H}(i)$ approximate the control constraint set. If the capacity for the production of part n is small in machine state i , the corresponding limits \underline{L}_{in} and \bar{L}_{in} of the hypercubes are small. Likewise, if the capacity is large, the limits are also large. The calculation of the upper and lower bounds thus takes into account the relative productive capacity in all machine states, demand rates and the value of the parts as given by the cost function $g(x)$. We expect therefore that the cost estimates satisfy the criteria above necessary for good performance by the estimate based controller.

3.3 Implementation of the Estimate Based Controller

There are two steps in the implementation of the EB-controller. Off-line, upper and lower bounds to the optimal value function are computed by solving (16) and (17). In practice this is done by discretizing the problems over discrete points in time and the buffer state. The estimates $\psi(x, i, t)$ are computed from the bounds and stored. On-line, whenever the system enters machine state i , the control $u(t)$ is determined by the linear program (12) using the

stored values of $\psi(x, i, t)$.

Computational and storage costs of the EB controller grow exponentially with the number of workstations M and linearly with the number of parts N . However, the computation is done off-line and the estimates of the optimal value function can be stored in peripheral devices. The on-line computation consists of the linear program (12) and has N variables and M constraints. Typically, N is between 5 and 10, and M between 10 and 20. The program can thus be easily solved on a small mini-computer.

The off-line computational cost can be reduced by pruning the machine state to exclude states with low probability. With the failure and repair rates typically found in manufacturing systems, a small number of states account for over 95% of the probability. A large number of states can therefore be eliminated without substantially altering the regions and hedging points corresponding to the estimate based control policy.

4. Simulation Results

The hierarchical controller has been implemented on a simulation model of a planned automated assembly system [Kimemia, 1982]. Simulation experiments with six and sixteen machine versions show that the performance of the algorithm conforms to theoretical predictions.

Estimate based flow control policies meet production targets that are within system capacity accurately and with low inprocess inventory. The time required to compute the cost estimates are fairly long. However since the estimates are calculated off-line and stored, the computational cost is not excessive. The on-line computation of the production rate $u(t)$ is easily done and the buffer state is responsive to changes in the production rate made by the flow control level of the of the hierarchical controller.

A preliminary experiment was performed to compare the hierarchical controller described here with a simpler, more intuitively comprehensible, strategy. Under this heuristic approach, a part is loaded into the system if its production lags demand, if the number of parts of that type inside the workcenter does not exceed a threshold, and if there is space available in a machine for that part. Thus, no explicit use is made of the repair rate, the failure rate, or other system parameters, and while feedback information on the buffer state (x) is considered, the current machine state (α) is not taken into account in making decisions.

The total production by the hierarchical controller was slightly larger than that of the heuristic controller. More significantly, however, was the fact that the production ratios were satisfied more exactly by the hierarchical controller. For example, two part types that were demanded in a ratio of 1:1 were produced in a ratio of 1.03:1 by the hierarchical method over a 16 hour day. They were produced in a ratio of 1.22:1 by the heuristic method. This difference can lead to important differences in total production of final assemblies and in inventory requirements.

Results of the simulation experiments show that feedback policies have definite advantages. In particular, dispatching parts at a rate that is within current system capacity reduces the proportion of time that parts spend waiting in internal buffers and hence the inprocess inventory is also reduced. Using buffer state information keeps production close to the desired demand rates.

5. Conclusions

The problem of controlling production in an automated flexible workcenter has been considered. The objective is to satisfy a master production plan while at the same time, utilizing fully the workstations in the system and minimizing the size of the in-process inventory.

A hierarchical production control algorithm depicted in Figure 1 is suggested. The algorithm is designed to fit in existing production management structures.

The flow control level of the hierarchy regulates the production rate of the workcenter ensuring at all times that the throughput is within the capacity of the operational workstations. The flow control problem is formulated as a stochastic control problem. The optimal production policy is a feedback law which, for each machine state, gives the production rate as a piece-wise constant function of the cumulative difference between actual and desired production.

The estimate based (EB) controller is a sub-optimal feedback control law which retains the piece-wise constant characteristics of the optimal policy. The EB controller computes, off-line, estimates of the optimal value function. On-line, the estimates are used to determine the throughput rate by means of a simple linear-program.

The hierarchical controller has been implemented on a simulation model of a planned automated assembly system with 16 workstations. Results of the simulation experiments show that the EB controller is capable of accurately meeting production targets that are within the capacity of the workcenter. Feedback control policies can thus help in achieving an important advantage of flexible automation, namely the ability to produce parts as they are required. This eliminates the need to keep large inventories of parts.

6. References

Cook, N. "Computer Managed Parts Manufacture," Scientific American, Vol.232, Feb 1975.

Hahne, E. "Dynamic Routing in an Unreliable Manufacturing Network With Limited Resources," M.I.T. Laboratory for Information and Decision Systems, Report No. LIDS-TH-1063, 1981.

Hildebrandt, R. "Scheduling Flexible Machining Centers When Machines are Prone to Failure," PhD Thesis, M.I.T Dept. of Aeronautics and Astronautics, May, 1980.

Hitomi, K Manufacturing Systems Engineering, Taylor and Francis, London, 1979

Hutchinson, G. "The Control of Flexible Manufacturing Systems: Required Information Structures," IFAC Symposium on Information-Control Problems in Manufacturing Technology, Japan, October, 1977

Kimemia, J. " Hierarchical Control of Production in Flexible Manufacturing Systems " PhD Thesis, M.I.T. Dept of Electrical Engineering and Computer Science, April,1982.

Lerner,E. " Computer Aided Manufacturing," IEEE Spectrum, Vol. 8, No. 11,

November, 1981.

Orlicky, J. Material Requirements Planning, McGraw Hill, 1975.

Rishel, R. "Dynamic Programming and Minimum Principles for Systems with Jump Markov Disturbances," SIAM Journal on Control, Vol. 13, No. 2, February 1975.

Tsitsiklis, J. "Characterization of Optimal Policies in a Dynamic Routing Problem," M.I.T. Laboratory for Information and Decision Systems Report, No. LIDS-R-1178, Feb, 1982.