



Order Release and Product Mix Coordination in a Complex PCB Manufacturing Line with Batch Processors

ROHIT BHATNAGAR

Nanyang Technological University, Singapore

PANKAJ CHANDRA

Indian Institute of Management, Ahmedabad, India

RICHARD LOULOU

Faculty of Management, McGill University, Montreal, Canada

JIN QUI

Merrill Lynch, Hong Kong

Abstract. In this paper, we study the role of order releases and product mix coordination in a complex manufacturing line with batch processors. We develop a planning methodology for synchronizing production in such manufacturing lines and discuss the decision-making process in the context of a PCB production environment at Northern Telecom's Fiberworld Division. The planning methodology includes developing mathematical programming models for determining a configuration of batch processors, order releases to the shop floor, and daily loading decisions at the batch processors. The optimization models are linked to a simulation model of the shop, which provides key statistics like lead-time, work in process, and utilization rates. The objective is to reduce lead-time for manufacturing different products in this environment while meeting the demand. We analyze the performance of such a line, study the efficacy of various types of shop floor synchronization policies, and establish the role of batch processors in managing such complex lines effectively. We exhibit how batch processors (which are bottleneck operations) could be scheduled effectively to incorporate the logical constraints that govern their operations and react to variability in the manufacturing line.

Key Words: order release, batch processor, shop-floor synchronization

I. Introduction

In this paper we address a problem encountered in various industries where a batch processor is a part of a complex manufacturing line. A batch processor can process more than one part of any type at a time (unlike a discrete processor, which processes one part at a time), and once a batch process begins, it cannot be interrupted until the entire batch is processed. Such batch processors include burn-in ovens in a semiconductor packaging or a printed circuit board (PCB) manufacturing line, diffusion and oxidation stages in wafer fabrication, and heat treatment facilities found in steel and ceramic industries as well as various bath treatments in the chemical and metallurgical industries. The processing time on a batch processor, usually, is the longest of any station in such a manufacturing line. A batch processor usually is subject to constraints that limit the batch size or composition. In the simplest case, a maximum number of parts is allowed in the batch processor at any time, whereas in other situations, constraints may be placed on the total space available or on the mix of items allowed in the batch.

The generic system under study comprises two stages: The first stage consists of several (serial or parallel) discrete processors, while the second stage consists of one or more parallel (although not necessarily identical) batch processing machines. The batch processor has certain restrictions on the mix of products that can be processed in a batch. Monthly demand is provided for each product, and the objective is to reduce the lead-time for manufacturing these products. Parts



released into the line are first processed by the first stage but have to wait at the batch processor until a "logical" batch (i.e., one that satisfies the mix constraints of the processor) has accumulated and may be processed. It can be appreciated that large release batches or stringent mix constraints can cause work-in-process (WIP) inventory to build up at the batch processor. As a result, any order release policy should aim at coordinating the product mix to meet the mix requirements of the batch processor. In other words, the batch processor should "pull" its requirement from the preceding stage in the manufacturing line.

In this article, we develop a planning methodology that can be used for the coordination of production in a manufacturing line that includes batch processors. We describe the objectives of the problem and the decision-making process in the context of a complex PCB manufacturing line at Northern Telecom's Fiberworld Division in St. Laurent, Quebec. The two stages in this PCB manufacturing operation are as follows: The first stage consists of several operations such as component mounting, insertion, and testing of individual PCBs, whereas the second stage is a burn-in operation that performs a final test on collections of PCBs forming fully functional systems. The goal of the project was to reduce the lead times on the shop floor. The current lead times, as observed in the plant, were about 15 days per PCB.

Much of the previous work related to multistage manufacturing is based on the study of a static shop; that is, one where the effective processing time at a station for any part is known and fixed. Karmarkar (1987a,b) has shown that the effective processing time at any station is a function of the lot size. As a result, the effect on congestion of the shop needs to be incorporated in choosing an order release policy, if we are to estimate lead times correctly. Most analytical results are derived for lines with very few stations and rarely in conjunction with batching type decisions on lines with setups (e.g., Awate and Sastry, 1987; Dallery and Gershwin, 1991).

From a different but related angle, queuing network theory has allowed the computation of WIP and lead times for a very limited set of release policies; for example, open queuing networks tackle the case where a release is directly triggered by the occurrence of a demand, whereas closed networks can mimic a few other policies where a part completion triggers a new release. Karmarkar, Kekre, and Kekre (1985a) model job shops as an open network of queues and develop heuristics for determining lot sizes in a multimachine environment. Also see Buzacott and Yao (1986) and Hsu, Tapiero, and Lin (1993) for surveys of manufacturing applications of queuing models.

It generally is difficult to compute lead times analytically for complex, multiproduct production lines, especially when logical constraints (such as the ones just alluded to) govern the operation of some of the stations in the line. This inability to use tractable analytical models for solving problems with complex structures has led to simulation-oriented analyses. Lin and Cochran (1987) and McDowell and Randhawa (1989) have developed simulation-based decision support systems for analyzing PCB fabrication lines. Karmarkar, Kekre, Kekre, and Freeman (1985b) have combined simulation with analytical models to examine the characteristics of a manufacturing system at Eastman Kodak's Apparatus Division. Another recent paper that studies the role of WIP in serial lines using simulation is Conway Maxell, McClain, and Thomas (1988). Some others who have written on order release policies in the context of semiconductor manufacturing are Glassey and Resende (1988), Wein (1988), and Akella, Rajagopalan, and Singh (1992). None of these, however, models the more complex burn-in-like operations.

The presence of the burn-in operation and its related constraints pose a new set of optimization problems in trying to manage the larger manufacturing line. The loading pattern of the burn-in



ovens has a significant impact on the WIP and consequently on lead times. Any attempt to reduce lead times will have to synchronize the order release policies on the shop floor with the loading plans for the burn-in ovens. The interest in modeling batch processors is quite recent. Most papers on batch processing operations address the loading issues in isolation of the order release issue (e.g., Ikura and Gimple, 1988; Glassey and Weng, 1992; Fowler, Hogg, and Phillips, 1992; Lee, Uzsoy, and Martin-Vega, 1992; Dobson and Nambimadom, 1992; Chandru, Lee, and Uzsoy, 1993; and Uzsoy, 1994). These researchers have considered a single-batch operation with different types of constraints. Lee, Uzsoy, and Martin-Vega (1992) and Dobson and Nambimadom (1992) acknowledge the necessity of integrating these models with those at other stations in the shop, as it is not obvious that decisions from such isolated models will be valid in the context of a multistage manufacturing line. Ahmadi, Ahmadi, Dasu, and Tang (1992), Chandra and Gupta (1992) and Gurnani, Anupindi, and Akella (1992) apparently are the only researchers who model batch and discrete processors together in the context of a multistage manufacturing line. Ahmadi et al. (1992) models a two-machine system, where one machine is a batch processor and the other is a discrete processor with no setups. They devise order release policies to minimize the makespan in the shop. Chandra and Gupta (1992) model batch processors in a semiconductor packaging line where processing times are constant for all parts and there is no variability in the production process. They establish the relationship between burn-in oven loading and lead times and illustrate the need to manage the WIP at the batch processor to minimize lead times. Gurnani et al. (1992) use bulk queuing procedures for modeling a generic serial-batch system, where the machines are subject to failures and design a control limit policy (i.e., wait until a defined length of the queue is built up) for batch loading. Only Chandra and Gupta (1992) address the requirements of multiple-part production in this context.

In the following section, we describe the problem in detail and discuss the product and the process structures. In Section 3, we present the decision framework and describe a planning methodology and its models. Our approach incorporates optimization models within a simulation framework to evaluate various order release and loading policies. In Section 4, we describe how our planning methodology is used to substantially reduce average lead-time on the shop floor. We illustrate the role of synchronization in reducing WIP and show how shop floor performance deteriorates when the capacity of the batch processors is highly utilized. Section 5 provides the conclusion and points toward SI issues that remain to be addressed.

2. Problem description

Northern Telecom's (NT) Fiberworld Division produces various types of transmission vices in NT's family of SONET-based products (a proprietary technology), sold to telecom companies and other industrial users. Transmission devices are digital multiplexing systems that essentially transmit digital signals over short and long distances.

The objective of this research is to determine a production plan that meets the monthly demand at a minimum lead-time or equivalently to minimize the WIP while meeting demand requirements. This would allow the firm to (1) reduce inventory build-up on shop floor and (2) produce items closer to due dates. This is in congruence with the firm's desire to reduce customer dissatisfaction due to delivery delays. In the following paragraph we describe the manufacturing environment and discuss the issues under consideration

2.1. Production planning

Each transmission device constitutes a set of printed circuit boards and other accessories as specified by a customer. Production on the shop floor is driven by firm orders for the transmission devices and individual printed circuit boards as well as forecast quantities for

individual printed circuit boards. All the demands are translated into requirements for specific PCBs. The monthly production plans are determined in terms of these PCB requirements. Requirements are updated each month. For our study, we used the actual demand faced by the firm over an 11-month period. The total annual demand for all PCBs was 44,462. The mean monthly total demand was 4042 and the standard deviation was 847.

2.2. Product structure

The product, called a *transmission device*, is a collection of PCBs in a large bay. The average cost of a bay is \$150,000. The firm produces three types of transmission devices depending on the type of task required in the transmission and receipt of signals: OC-48, OC-12, and REGEN. Each device comprises a set of subassemblies (e.g., 7E02XX, 7E04XX, 8E03XX). There are 38 different subassemblies. Each subassembly consists of one or more *printed circuit boards* (e.g., 7E0210, 7E0208, 7E0204). There are 60 different types of PCBs. Table I shows the product structure; for example, one OC-12 comprises subassemblies 7E02XX, 7E08XX, and so forth. Subassembly 7E02XX is made up of PCBs 7E0210, 7E0208, and forth. Note that some subassemblies, such as 7E08XX, comprise a single PCB, 7E08AA. Also a PCB may be a part of more than one subassembly and a subassembly can be used

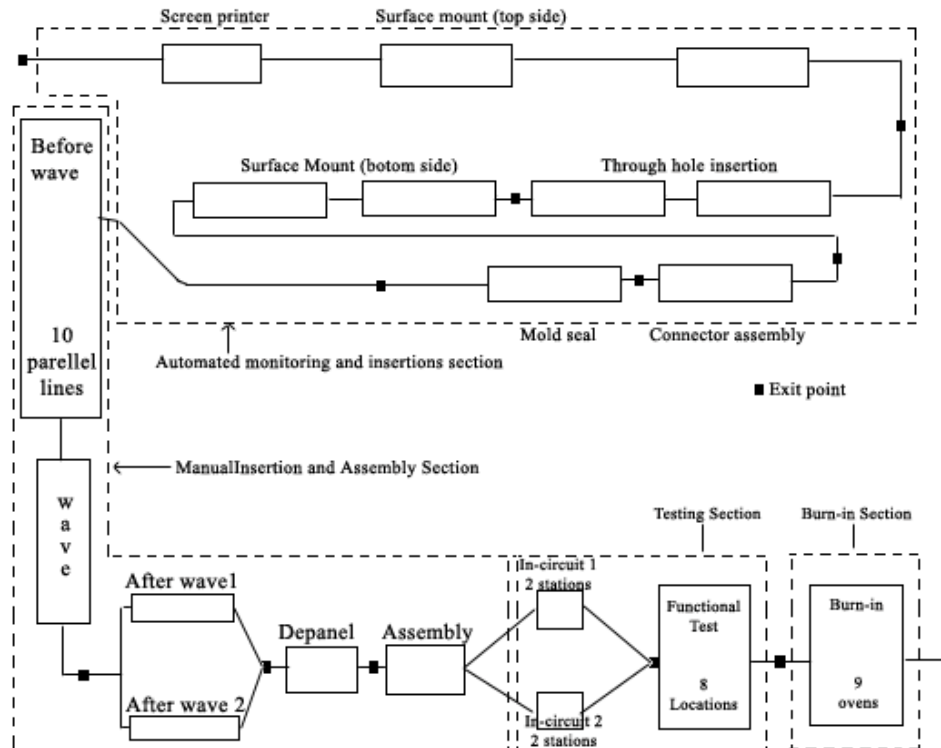
Table 1. The Product structure

Transmission devices	Subassemblies	PCB's
OC-12	7E02XX	7E0210
		7E0208
		⋮
	7E08XX	7E08AA
	⋮	
OC-48	8E01XX	8E0142
		8E0134
		⋮
	7E08XX	7E08AA
	⋮	
REGEN	8E03XX	8E0302
		8E0134
		⋮

in more than one product. Each printed circuit board is released into the line in the form of a *panel* of standard size. A panel may contain a single PCB or several identical PCBs, depending on the physical size of the PCB (e.g., a certain panel contains 27 copies of the 7E0208 PCB). The standard size of panels facilitates the automation of material handling by rail conveyors.

2.3. A guided tour of the production line

The manufacturing line comprises four sections: the automated mounting and insertion section, the manual insertion and assembly section, the testing section, and the burn-in section (figure I). The first three sections comprise discrete processors. In the automated mounting and insertion section, circuitry is printed and solder paste is spread selectively on the PCB by a screen printer,



and then the board is populated with components using surface mount and through-hole technologies. All these can withstand wave soldering. Components of different types and shapes are either mounted on both sides of the PCB, using top side or bottom side machines, or inserted using through hole insertion machines. Some PCBs go through an additional set of operations (connector assembly, mold seal, and modification before wave) before passing on to the next section. The manual insertion and assembly section consists of two segments: In the first segment, the manual insertion, additional components that can withstand wave soldering are inserted into the PCB, which then passes through an automated wave soldering machine. Next, the remaining components, which cannot withstand wave soldering, are inserted on the PCB. In the second segment, individual PCBs are depaneled (broken apart into individual PCBs) and a set of different PCBs (the assembly set) is assembled into a subassembly at the assembly stations. These operations are done manually. Up to the depaneling station, PCBs flow in a standard panel size. From this stage onward, subassemblies are tested.

Two types of tests are performed in the testing section. The first test, called the *in-circuit test*, checks the integrity of the circuits. This test is performed at two locations: one dedicated to the 7E product line and the other to 8E product line. Each location comprises two parallel testing stations. Any "reject" is sent for troubleshooting and repair before being retested. The first-pass yield at the in-circuit stations is 30-60%. In the second type of testing in this section, called the *functional test*, the subassembly is tested as to whether it is performing its collective function. The functional test is performed at eight locations, and each location may have more than one test station. Each PCB type is allocated a unique location and more than one type of PCB can be allocated to each location. Rejects, once again, go through troubleshooting, repair, and retesting.



The first-pass yield, here, ranges 35-98%. The second-pass yield at both the tests normally is 100%. The last section in the line, the burn-in section, comprises a number of batch processors (burn-in ovens). Each PCB is supposed to contribute to the functionality of the product (as a whole) and that too in the field. This is tested (i.e., simulated) in the burn-in ovens. Hence, the purpose of these ovens is twofold. First, it is to test whether the product (i.e., a transmission device consisting of a set of assemblies) functions as a unit. The second purpose of the oven is to subject the product to various temperature profiles while continuously testing its functions. The burn-in oven is the most challenging process in the entire line. Ovens are available in two sizes: small and large. Currently, the manufacturing facility has four small and five large ovens. Ovens can be loaded only in specific combinations or kits. A small oven can be loaded with any combination of one to two OC-48s and none to two REGEN s. As a rule, REGENs cannot be tested alone. Each large oven can be wired into any of the following three configurations, each allowing different combinations of logical devices: either one to eight OC-12s, one to two OC-48s and none to four REGENs, or one to four OC-48s. It takes over eight hours to change the wiring configuration of a large oven. Finally, variations in the composition of a transmission device are allowed in the burn-in operation; these are summarized by a set of minimum and maximum numbers of subassemblies that are allowed in a particular device. A PCB tested as part of a given product may later be delivered to a customer as part of another device.

Figure I is a representation of the complete line. However, not all PCBs go through all workstations. Some exit at locations as shown in the figure. At each station except burn-in, PCBs are processed either in the form of a panel or in the form of a subassembly. However, in the oven, only *configurations* of devices can be tested. As a result, each time an oven is to be run, a decision has to be made regarding the configuration of the "logical" device (e.g., four units OC-48 or two units OC-48) to be loaded in the oven. This "logical" product is made from the product mix available as WIP in the burn-in section. Needless to mention, the "logical" product being tested in the oven can be different from the product demanded by a customer or from the mix of separate subassemblies that may be in demand this month. To provide flexibility, captive subassemblies of different types could be used to complete a "logical" mix. Such captives are used only for testing purposes and are never shipped to customers. After going through burn-in a certain number of times, captives are scrapped. As PCBs are expensive, one aim of the research was to analyze the sensitivity of lead times to the number of available captive PCBs. After burn in, subassemblies are taken out of the oven and either mounted into bays according to customer demand for complete devices or sent to a warehouse as independent subassemblies/PCBs to satisfy demand for individual parts from different customers. Table 2 gives the processing and setup times at different stations in the manufacturing line. Note that significant setups exist mainly at the screen printer and at the in-circuit test. The burn-in ovens have long process times and need to be carefully examined in developing any planning methodology.

3. The decision framework: wiring, order release, and burn-in plans

The decision framework shown in figure 2 is closely modeled by the planning methodology shown in figure 3 which we now describe. For this study, the detailed operations of the shop are simulated by a GPSS/H program that computes the WIP levels, lead times, and other statistics of interest. The simulation requires as input a monthly *wiring plan* (which determines which transmission devices may be loaded in each oven during that month), a daily *order release plan* (each day at 8:00 A.M. it specifies the quantity and sequence of the different panels released in the simulated shop), and a daily *burn-in plan* (each day at 8:00 A.M. it specifies which subassemblies to load in each burn-in oven during that day).

In addition, a few decision rules are required at some intermediate workstations to specify batching and priority decisions.



For this study, all capacities were taken as given, assuming that the line was balanced. This was confirmed by direct observation as well as by our results. Had this not been the case, our methodology would have identified locations with insufficient capacity.

The wiring plan ensures that the ovens are adequately set up to handle next month's demand. The daily release plan tries to ensure that the demand requirements are being met while trying to release a logical mix of PCBs that could eventually form a load in the ovens. If all the PCBs released over a one-day period arrive at the burn-in section on the same day (some days later), the release plan for those boards automatically would become the burn-in plan. However, because of unequal processing times and of random events such as test failures and machine breakdowns, the subassemblies ready for burn-in test at the end of a day may not exactly match the release plan of some previous day. It then is important to redesign a burn-in plan, daily, based on the actual available mix of PCBs in the burn-in section and those that may soon be arriving there. Most often, the burn-in plan will be different from the release plan.

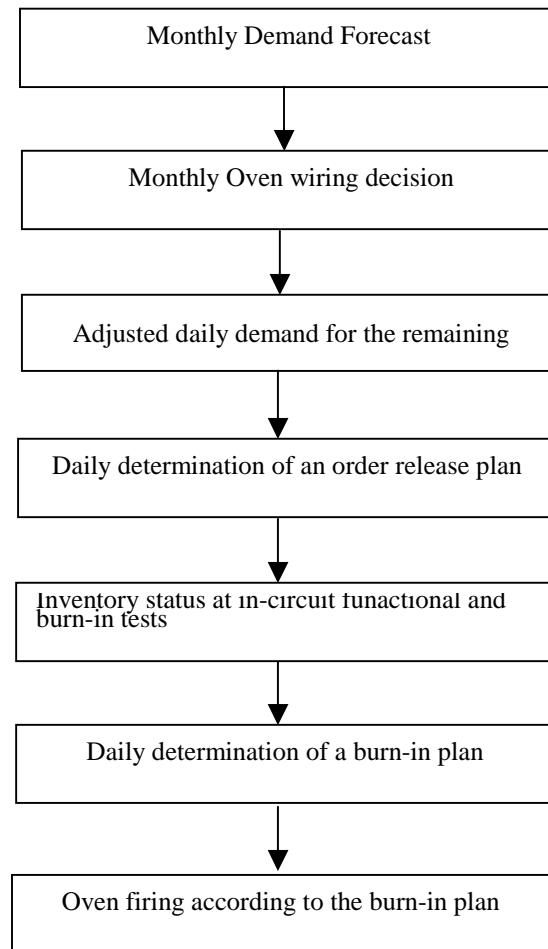


Figure 2. Decision sequence.

We now describe each of the three plans (the wiring, the release, and the burn-in plans) in some detail, followed by several remarks concerning the other decisions at the intermediate stations. All three plans are modeled as integer programs, solved using GAMS/LAMPS on a VAX 3100. The GPSS/H program dynamically calls GAMS at appropriate times.

3.1 The Wiring Plan

Each burn-in oven is wired to test a subset of three different devices: OC-12 only, OC-48 only, and OC-48 + REGEN. Once a wiring is selected, the oven can test only the

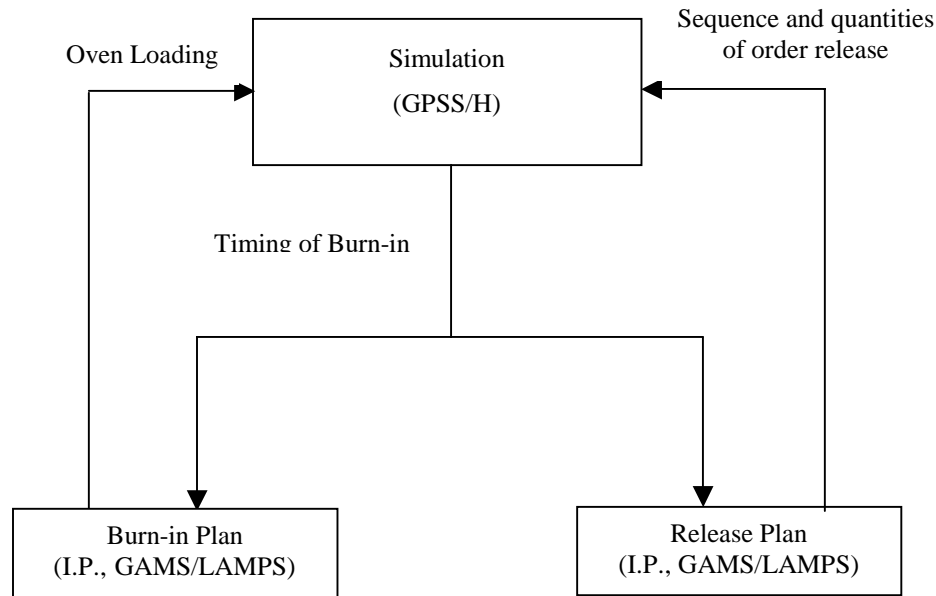


Figure3. Planning methodology

corresponding devices (however, the exact number of devices per oven is determined daily by the burn-in plan). Since rewiring ovens involves substantial loss of capacity and worker power, it is not effective or efficient to use rewiring as a strategy to overcome the short-term fluctuations in the mix of subassemblies available before burn-in. On the other hand, since the monthly demand mix for the subassemblies may vary considerably, monthly rewiring sometimes becomes desirable or even necessary to ensure that ovens are capable of testing all PCB types required in the coming month. The wiring decision is made at the beginning of each month and determines how many ovens will be wired using each of the three basic wirings. If more than one wiring option is feasible to satisfy the monthly demand, the plan chooses the one that requires the least number of ovens to be rewired (as compared to the current wiring). It also assumes that the wiring process does not consume productive capacity, although the time of engineers and the like is sought to be minimized. The model for determining the wiring plan follows.

Indices:

i = index of PCB types

m = index for wiring options

Parameters:

M = number of wiring options ($M = 3$)

J = number of ovens available ($J = 9$)

$C(i, m)$ = maximum number of PCB i required for the wiring option m



$D(i)$ = monthly demand for PCB i

H = number of burn-in cycles available in a month
 = (22 working days per month) x (24 h per day)/(30 h burn-in cycle)

$N^0(m)$ = number of ovens currently wired with option m

Decision Variable:

$N(m)$ = number of ovens to be wired with option m

The model for the wiring plan can be written as this:

$$\begin{aligned} & \text{Minimize} && \sum_m |N(m) - N^0(m)| \\ & \text{subject to} && \\ & H \cdot \sum_m N(m) \cdot C(i, m) \geq D(i) && \forall i \\ & \sum_m N(m) = J && \\ & N(m) \geq 0 \text{ and integer} && \forall m \end{aligned}$$

The objective function represents the total number of ovens to be rewired. Constraint (1) ensures that the capacity for each type of PCB is sufficient to satisfy its monthly demand. Constraint (2) states that, altogether, J ovens are available and each oven has to be wired with some option.

3.2. The release plan

A good release plan should ensure that demands for various subassemblies are met and that WIP is minimized throughout the line. Because a large WIP is observed at the burn-in section, it is important to ensure that the daily release corresponds to a mix of subassemblies that would be "loadable" into the burn-in oven, once they arrive there. The basic premise underlying this is that all the PCBs released on day t_1 indeed will reach the burn-in section on some later day, t_2 . Although this is not likely to occur exactly, it constitutes a good basis for a release policy. An ideal release plan therefore should satisfy the just-described mix criterion and meet the daily demand.

The order release integer program aims at reconciling these two objectives when necessary; that is, the release plan will take into consideration the captives available for burn-in and ensure that the least number of excess PCBs will be released. In case more than the daily demand has to be released to meet the minimum limit, even after considering the captive PCBs, the daily demand for the remaining days in the month will be adjusted to reflect the amount overproduced. Overproduction also may take place because the number of panels released for some PCB types may contain more than the ideal number of PCBs. In this case, the daily demand for this PCB for the remaining days in the month will be reduced accordingly. Since all rejects in the testing section are repaired and there is practically no scrap, no additional parts over this amount are released. The model for determining the release plan follows.

Indices:

i =: index for PCB types

k = index for burn-in ovens

j = index for configurations

Parameters:

$A(i, j)$ = maximum required number of PCB i for configuration j

$B(i, j)$ = minimum required number of PCB i for configuration j

$CAP(i)$ = the total number of captive board i available for burn-in

$d(i)$ = daily demand for PCB i

W_r = weight assigned to PCBs ($W_r = 10$ is used in all our experiments)

W_c = weight assigned to captive boards ($W_c = 1$ is used in all our experiments)

Decision variables:

$X(i, k)$ = number of PCB i processed by oven k

$Y(i, k)$ = number of captive board i required by oven k

$Z(k, j) = \{1, \text{ if oven } k \text{ uses configuration } j, 0 \text{ otherwise}\}$

The model for the release plan can be written as:

$$\text{Minimize } \sum_i \sum_k [W_r \cdot X(i, k) + W_c \cdot Y(i, k)]$$

subject to

$$\sum_j B(i, j) \cdot Z(k, j) \leq X(i, k) + Y(i, k) \leq \sum_j A(i, j) \cdot Z(k, j) \quad \forall i, k \quad (4)$$

$$\sum_j Z(k, j) \leq 1 \quad \forall k \quad (5)$$

$$\sum_k Y(i, k) \leq CAP(i) \quad \forall i \quad (6)$$

$$\sum_k X(i, k) \geq d(i) \quad \forall i \quad (7)$$

$$Z(k, j) \in \{0, 1\} \quad \forall k, j \quad (8)$$

$$X(i, k), Y(i, k) \geq 0 \quad \forall i, k \quad (9)$$

This model minimizes a composite objective function with two terms: the number of real PCBs released and the number of captive boards. The first term has a higher weight (i.e., 10) than the second (i.e., 1), reflecting the higher priority of releasing as closely as possible to the adjusted daily demand. There are four groups of constraints. The first group (constraint (4)) ensures that the quantity of boards (including the captive PCBs) in any oven is between its required minimum and maximum limits as determined by the wiring of that oven and its selected configuration. A *configuration* is defined as the number of OC-12, OC-48, and REGEN to load for a firing of the oven. The second group of constraints (5) asserts that each oven may take at most one configuration for that day. If the result of the model indicates 0 as the configuration number for an oven, this means that the oven is not used for this release. The third group of constraints (6) ensures that the total number of captive PCBs used for each type does not exceed the quantity available. The fourth group of constraints (7) ensures that the total number of PCBs released of each type is no less than its adjusted daily demand. The adjusted daily demand is computed for



each PCB type as the monthly demand minus the quantity released in previous days minus the number of rejects, divided by the remaining number of workdays in the current month.

3.3. The burn-in plan

The goal of the burn-in plan is to reduce the waiting time for the subassemblies that already have arrived at the burn-in section. The observed wait in the burn-in buffer is large, so that substantial lead-time savings may be secured by reducing the WIP there. As explained earlier, because of unequal processing times and the occurrence of random events such as test failures and machine breakdowns, the subassemblies ready for burn-in test at the end of a day may not exactly match the release plan of some previous day. In this case, one could either hold off firing the ovens and wait until all the remaining subassemblies have been collected or construct a burn-in plan based on how many of each type currently are or soon will be available for burn-in. We do not expect the first approach to perform well, since the delay of perhaps a few subassemblies would hold up the entire batch and provoke long waiting time before the burn-in test.

Instead, we choose the second approach and construct a burn-in plan that will in general differ from the release plan. The challenge of such a burn-in plan is to balance three goals that often contradict one another: (1) to burn in the subassemblies that are ready and available as soon as possible, (2) to respect the minimum and maximum limits of each burn-in oven, and (3) to wait for subassemblies that have yet to arrive to minimize their waiting time.

The minimum requirements of burn-in ovens prevent the subassemblies from being sent for burn-in as soon as they arrive. They have to wait at least for the minimum requirements to be collected. However, prompt firing of any oven with just the minimum requirements may underutilize the oven capacity and incur long waiting time for the boards that are to come later on. Our burn-in plan attempts to strike a balance between the short-term view and the long-term view by taking a planning horizon of 24h when determining the loading of the ovens. Therefore, the decision is based not only on what subassemblies currently are available but also what subassemblies are likely to become available in the next 16h. The automated insertion, manual insertion and assembly, and the testing sections work for two shifts while the burn-in section operates over three shifts. As can be seen from Table 2, the total processing times of all operations before the burn-in is between 2 and 4h. Each oven takes 3h for loading, 24h for testing, and 3h for unloading. At the beginning of each day (i.e., 8:00 A.M.), the burn-in department determines the following: which subassemblies are expected to be ready for the burn-in test in the next 16h (i.e., by midnight) and which ovens and captive boards will become available in the next 24h. Note that the identities of and times at which the ovens will become free are perfectly predictable, whereas the information on the availability of subassemblies must be secured via an imperfect estimation procedure, described later. Based on this information, the burn-in plan is created.

The objective of the estimation procedure is to find out at 8:00 A.M. every day how many subassemblies of each type are likely to finish the functional and in-circuit tests and become available for the burn-in test within the next 16h. We refer to those subassemblies that could be available as admissible. Thus, all the subassemblies currently sitting in the burn-in stock are admissible. We then look at the subassemblies waiting for the functional and in-circuit tests to see how many of them also can become admissible.

This is tantamount to predicting how many subassemblies of each type will clear the two tandem queuing systems making up the in-circuit and the functional tests. The precise estimation procedure used is given in the appendix to this article.



The model for the burn-in plan follows.

Indices:

i = index for PCB types

k = index for burn-in ovens

j = index for configurations

Parameters:

K' = the set of ovens to be available during the next 24h

$A(i, j)$ = maximum required number of PCB i for configuration j

$B(i, j)$ = minimum required number of PCB i for configuration j

$E(i)$ = estimated number of PCB i to be available in the next 16h

$CAP(i)$ = the number of captive board i to be available in the next 24h

W_r = weight assigned to PCBs ($W_r = 30$ is used in all our experiments)

W_c = weight assigned to captive boards ($W_c = 1$ is used in all our experiments)

Decision variables:

$X(i, k)$ = quantity of product i processed in oven k

$Y(i, k)$ = number of captive board i required by oven k

$Z(k) = \{1, \text{ if oven } k \text{ uses configuration } j, 0 \text{ otherwise}\}$

The model for the burn-in plan can be written as:

$$\begin{aligned} & \text{Maximize} && \sum_i \sum_{k \in K'} [W_r \cdot X(i, k) - W_c \cdot Y(i, k)] \\ & \text{subject to} && \\ & \sum_j B(i, j) \cdot Z(k, j) \leq X(i, k) + Y(i, k) \leq \sum_j A(i, j) \cdot Z(k, j) && \forall i, k \in K' \quad (10) \\ & \sum_j Z(k, j) \leq 1 && \forall k \in K' \quad (11) \\ & \sum_k Y(i, k) \leq CAP(i) && \forall i \in K' \quad (12) \\ & \sum_k X(i, k) \leq E(i) && \forall i \quad (13) \\ & Z(k, j) \in \{0, 1\} && \forall k \in K', j \quad (14) \\ & X(i, k), Y(i, k) \geq 0 && \forall i, k \in K' \quad (15) \end{aligned}$$

The objective of this program is to maximize the number of PCBs to the burn-in tested in the next 24h while using the minimum number of captive boards. The interpretation of constraints (10), (11), and (12) are the same as the first three constraints in the model for the release plan, except that the ovens and captive boards to be considered are limited to the ones that will be available during the next 24h. Constraint (13) ensures that we consider only those PCBs expected to be



ready for the burn-in test. The values of $E(i)$ are determined by the estimation procedure presented in the appendix.

3.4. Decisions at intermediate stations

The preceding three plans deal with what appear to be the main operational decisions along the line. However, a few secondary decisions remain to be specified for the simulation program to run unambiguously, as described in the following.

Batching at the wave oven and at in-circuit testing stations. In addition to the setup at the screen printer, setups are also incurred at the wave oven and in-circuit test stations, hence, the need for a robust batching policy at these two locations. Based on a preliminary investigation, the following batching policies were adopted. At the wave oven, each PCB requires one of three different temperature profiles. A setup is needed to switch from one temperature profile to another. To minimize the WIP at the wave machine, the profile selected was the one with the largest number of PCBs waiting. This profile remains active until no more PCBs requiring this profile are waiting. At this point, a setup is incurred for switching to the profile that has the maximum WIP waiting. At the in-circuit test, a sequence-independent setup is incurred for each subassembly type. The batch size used is on the same order of magnitude as the release batch size. The justification for this batch size is that it represents the number of subassemblies required in a feasible burn-in plan. Since only a small fraction of PCBs go through the mold seal operation, we do not batch products here and allow setups to be incurred as the PCB sequence changes.

The queue discipline at the manual insertion stations. The queue discipline at the manual insertion stations is first come, first served. There are 10 parallel lines of workstations for before wave insertion and 12 at the after wave insertion. As already mentioned, workers on the parallel lines are cross-trained so that the arriving job may go to the first available line, after waiting in the common queue.

Rules for assembly. After the after wave insertion, panels are depaneled into individual PCBs (if there are multiple PCBs on the panel) and subassemblies are generated. Several procedures are used to generate subassemblies. In some cases, PCBs go to in-circuit test or functional test as individual PCBs and assembly is done later; in other cases, subassemblies are generated first and then sent for testing. In still other cases, testing is done both for the PCB and the subassembly. To avoid the complexity of including all individual rules in the simulation, the following approximation was used: PCBs requiring depaneling were depaneled after the after wave insertion and the assembly was done as soon as the entire assembly set was available. In-circuit and functional testing was performed at the level of the subassemblies. Wherever necessary, the test times were inflated to reflect test times at both PCB and subassembly levels. This rule is conservative in terms of lead times since it forces the generation of subassemblies earlier than required.

4. Results and managerial implications

We now present experimental results based on the study of the NT Fiber world production facility and derive insights for the management of lines with batch processing units. We begin with the observation made on the shop floor to the effect that a large portion of the wait occurs in front of those stations that require logical sets of boards to be present; that is, assembly and burn-in. We observed that these two stations alone account for 65% of the total wait experienced by the boards.

To compare our approach, as presented in Section 3, with others, we designed three policies, which are given in Table 3. In Policy 1, orders are released based only on the daily demand; that



is, at 8:00 A.M. every day each assembly's adjusted daily demand is released on to the shop floor in a single batch. The burn-in plan then is applied daily as described in Section 3.3. Policy 3 is a dynamic-reactive policy that uses the release plan of Section 3.2 for releasing boards and the reactive burn-in plan of Section 3.3 for determining, daily, the oven loading. Finally, Policy 2 uses the release plan of Section 3.2 both as the release plan and as the burn-in plan. In this policy, the ovens wait for the predetermined mix to arrive before they are fired.

A comparison of Policies 1 and 3 sheds light on the impact of release batch size on lead time, since the batches released by Policy 3 are smaller than those by Policy 1. A comparison of Policies 2 and 3 reveals the impact of a reactive burn-in policy as opposed

Table 3. Shop synchronization policies.

Policy 1
Release plan = daily demand
Burn-in plan = results of the burn-in plan model
Policy 2
Release plan = results of the release plan model
Burn-in plan = release plan
Policy 3
Release plan = results of the release plan model
Burn-in plan = results of the burn-in plan model

To rigid planning. Furthermore, for each policy, two different cases are investigated allowing for a high and a low number of captives, respectively. For the "high captive" case, we use 500 captive PCBs for each subassembly type; while for the "low captive" case, we use 30 captive PCBs for all subassembly types except for one subassembly type, where this number is set to 75. The numbers in the low captive case represent approximately 50% of the maximum number of captives actually utilized on any given day in the case with high captives. Shortage of captive PCBs therefore is sure to occur in the low captive case.

Table 4 contains the main results of our simulation. Each column corresponds to the combination of a policy and a level of captive boards. Each row corresponds to one month. Two initial months were used to "warm up" the simulation and results were obtained for the next 11 months. In each cell, the monthly average lead-time (in days) is shown along with the estimated standard deviation (in days). We did not attempt to compute rigorous confidence intervals for the mean lead times of Table 4 due to the excessive computing time this would have required. However, the relatively small size of the overall standard deviations and the large number of boards processed in a month (typically several thousands) indicate that the error on the monthly lead-time average will be very small as compared to the mean. (Neglecting autocorrelation, one would evaluate a typical value for the error at about 0.01 day for Policy 1 with high captives, a very small number indeed.)



Table 4. Mean and standard deviation of production lead times.

Month	Mean (standard deviation) of production lead times (in days)					
	Policy 1		Policy 2		Policy 3	
	High captive	Low captive	High captive	Low captive	High captive	Low captive
1	3.32	3.86	5.82	8.65	2.84	3.54
	0.61	1.31	2.22	2.84	0.52	1.26
2	3.24	3.97	6.64	10.83	2.8	3.51
	0.58	1.48	2.28	2.8	0.48	1.59
3	3.16	3.71	10.04	16.26	2.91	3.6
	0.6	1.34	3.26	2.91	0.53	1.82
4	3.27	3.57	10.96	23.29	2.83	3.48
	0.59	0.83	4.55	2.83	0.5	1.36
5	3.17	3.41	10.26	28.67	2.76	3.63
	0.58	0.72	4.37	2.76	0.5	1.2
6	3.19	3.34	8.87	33.67	2.88	3.15
	0.6	0.66	5.33	2.88	0.53	0.94
7	3.02	3.13	11.66	37.28	2.88	3.27
	0.68	0.74	7.1	2.88	0.61	1.09
8	3.26	3.37	5.86	41.94	2.87	3
	0.65	0.71	4.12	2.87	0.56	0.64
9	3.31	3.93	7.1	31.35	3.03	3.44
	0.64	1.28	3.14	3.03	0.64	1.05
10	3.24	4.77	8.5	30.98	3.02	4.33
	0.68	3.13	4.55	3.02	0.64	2.71
11	3.17	4.85	7.41	36.29	2.89	4.38
	0.6	3.64	5.79	2.89	0.58	3.3

Table 5. Detailed simulation results using Policy 3 with high captives.

Month	Number of boards processed	Mean lead-time	Standard deviation of lead-time	Average WIP	Utilization of ovens
1	4657	2.84	0.52	606.81	33.03%
2	4436	2.8	0.48	555.64	30.60%
3	3603	2.91	0.53	478.65	25.48%
4	4385	2.83	0.5	557.9	30.22%
5	3940	2.76	0.5	488.28	27.77%
6	4039	2.88	0.53	532.74	28.17%
7	1684	2.88	0.61	202.8	11.29%
8	3817	2.87	0.56	541.86	27.40%
9	5130	3.03	0.64	696.8	35.75%
10	4392	3.02	0.64	590.59	30.37%
11	4381	2.89	0.58	588.8	31.75%

Table 5 gives additional information on the average monthly WIP and the burn-in oven utilization for Policy 3 with high captives. It of course is well known that average WIP is proportional to average lead-time, due to Little's law.

4.1. Effect of burn-in policy

As can be seen from Table 4, the dynamic-reactive approach of Policy 3 outperforms the other two policies for each of the two captive PCB scenarios. Policy 2 consistently performs poorly compared to the others, chiefly due to increased waiting for the released orders to be collected before the burn-in. This situation is aggravated because, due to variability in the line (e.g., breakdowns, poor yield), ovens have to wait longer before a mix can be collected for a burn-in. This is evident when we compare Policies 2 and 3 (with high captives). The waiting times up to burn-in are identical for the two policies, but the waiting at burn-in under Policy 2 increases considerably (ranging 80-160 h for different months) as compared to that with Policy 3 (which is fairly stable at approximately 18-22h for all the months). For Policy 2, the wait at the burn-in represents 80-90% of the total waiting time. Our burn-in plan addresses this shortcoming. It reacts to such variations by adjusting the burn-in mix based on the subassemblies in the buffer and the expected arrivals. As a result, Policy 3 (with both high and low captives) yields vastly improved lead time and consequently lower WIP than other policies.

We also observe that the unavailability of captive PCBs has a much greater impact on the lead times in Policy 2 than in the others. This points to another advantage of the reactive approach in determining the burn-in mix.

4.2. Effect of release policy

Comparing Policies 1 and 3 can see the role of the release plan. The known trade-off between the increased numbers of setups versus the reduced waiting time for the collection of batches comes into play here. Policy 3 outperforms Policy 1 by about 0.5 days in terms of lead times, indicating the relative importance of reducing waiting times through smaller batches.

It is worth noting that locations where coordination of product mix was required (e.g., burn-in and assembly) exhibited higher contributions to the lead-time. In one example (for Policy 3 with high captives), we observed that 50 and 17% of the waiting time was for burn-in and assembly, respectively. We also observed that too many subassemblies of the same type queue up at a particular location (e.g., in-circuit and functional tests) while stations for other types of products are idle. This contributes to congestion at some stations. Once again, the superiority of the release plan as determined by the model in Section 3.2 is obvious. It takes into consideration potential oven loadings in determining order releases, thereby alleviating this situation. The practice of ignoring batch processors while releasing orders onto the shop floor leads to excessive WIP due to the poor coordination of the mix. Similarly, if we do not control the impact of variability, coordination across the product mix (which is essential for operating batch processors) will

deteriorate. (Needless to mention, management's priority should be toward eliminating the sources of variability.) Our dynamic-reactive policy (Policy 3) takes these issues into account. In releasing orders, it considers the potential burn-in loading mix, but as PCBs/subassemblies are affected by variability in the line, it reacts by adjusting the burn-in mix that is used in oven runs. The robustness of this policy is exhibited by stable lead times (mean and standard deviation) fairly independent of the demand, and relatively consistent levels of capacity utilizations. It can



be observed from Table 5 that average lead times and standard deviations change by less than 5%, even when the demand is decreased by 60% or increased by 30%.

4.3. Effect of forced utilization

Another managerial practice that often leads to higher WIP relates to resource utilization. We observed at several plants (including the current one) a propensity to operate scarce resources at high levels of utilization. Karmarkar (1987a) has shown the perverse effect of utilization of a single machine on congestion. This impact is accentuated in complex, multistation lines, especially those with batch processors. Table 6 shows the impact on lead-time when capacity utilization requirements are imposed on batch processors; that is, when a burn- in oven is run only when subassemblies representing a certain fraction of the capacity of the ovens have been collected. (Note that this "space" utilization is different from "time"

Table 6: Mean and Standard deviation of production lead times with oven loading restrictions using policy3 with high captives

Mean and Standard Deviation of production lead-times			
Month	75% restriction	50% restriction	No restriction
1	6.7(6.11)	2.80(0.54)	2.84(0.52)
2	7.76(8.09)	2.84(0.50)	2.80(0.48)
3	9.59(11.39)	2.88(0.61)	2.91(0.53)
4	11.68(14.94)	2.85(0.53)	2.83(0.50)
5	13.05(16.84)	2.98(0.61)	2.76(0.50)
6	14.51(19.87)	2.84(0.58)	2.88(0.53)
7	19.38(28.51)	3.69(6.97)	2.88(0.61)
8	21.48(30.68)	2.93(0.76)	2.87(0.56)
9	21.62(31.37)	3.19(0.92)	3.03(0.64)
10	21.58(31.31)	3.85(1.90)	3.02(0.64)
11	21.93(31.59)	3.16(1.36)	2.89(5.79)

utilization. The effect of maximizing the latter also could lead to delays as more runs of the oven may be needed to burn-in the entire demand.) The high capacity utilization restriction can be seen to lead to large lead times. In our case, both the minimum and maximum requirement on subassemblies and the capacity utilization restrictions interact to yield poor lead times. This effect is relatively unimportant in up to 50% forced minimum utilization of burn-in ovens (in fact we observe that average lead-times may be marginally reduced for such a policy). However the performance deteriorates considerably when the minimum restriction is increased to 75% and both the lead times and standard deviations are observed to increase substantially (often in excess of 500%). Moreover, often unbalanced demands for different subassembly types can prevent the formation of logical mixes for burn-in for a given utilization level, thereby causing excessive delays at the burn-in.

5. Conclusion,

In this research, we develop a fairly robust methodology for managing manufacturing lines with batch processors. Order releases are determined based on product mix requirements of the batch processors while the loading of the batch processors takes into consideration the variability in the manufacturing line. This mechanism dynamically reacts to the state of the shop and the loading



patterns of the batch processors. As a guideline, first identify the locations where WIP accumulates, determine why WIP tends to accumulate at these locations, and then design policies accordingly rather than using general recipes for managing such lines.

The procedure we develop is very useful in several respects-apart from being used as a control mechanism for synchronizing production on the shop floor to reduce lead-time and WIP, it can be used for resource planning over a planning horizon. Some of the key conclusions could be summarized as follows:

- Considerable reduction in lead times was achieved using our procedure. For our test data, the observed average lead times in the plant were close to 15 days per subassembly and the mandate of the project was to reduce them to near 3 days. As can be seen from the results, this objective clearly has been achieved.
- The logical constraints governing the batch processors drive the planning process and should be explicitly considered in order release and loading policies.
- Given the variability in the manufacturing line, it is always better to develop a separate reactive burn-in loading plan than to use a fixed plan.
- The managerial objective of maximizing capacity utilization (i.e., space used in the batch processors) could lead to long lead times, as it tends to postpone the firing of the batch processors.
- The procedure allows one to determine periods where additional capacity would be needed to maintain low WIP and lead times.

Although our procedure was developed for a specific production line, these conclusions and the general approach remains valid for manufacturing lines with complex batch processes.

The solution from the burn-in plan also gives the number of captive PCBs needed for each type per month. This is useful information for planning, especially since the sunk cost associated with each captive PCB is quite high. Similarly, a manager can make capacity- related decisions (i.e., number of shifts, worker levels, additional machine requirements) using lead times in different demand scenarios determined by the model. The varying lead-time figures (as well as WIP levels) also could point toward periods (and locations) where extra resources could be added to reduce lead-time to desired levels.

Several additional interesting issues arise in the context of this research that should be addressed in the future. A direct extension of this work would be to develop release policies dependent on the state of the shop. This would make order releases more reactive.

Appendix I. Estimation procedure

The estimation procedure involves four steps.

Step 1. Calculation of expected processing times

The expected processing time, E_i^r , for PCB i ($r = F$ for functional and $r = C$ for in-circuit (I/C) tests) is computed as follows:

$$E_i^r = T_i^r + P_i^r (TS_i^r + RP_i^r + RT_i^r)$$

where T_i^r , P_i^r , TS_i^r , RP_i^r , and RT_i^r represent, respectively, the test time, probability of failure, trouble shooting time, repair time, and retest time.



Step 2. Estimation of the number of PCBs at the functional test that will become ready for the burn-in test within the next 16h

Denote by $j(f)$ the j th test station at the functional test location f , $1 < f < 8$. Step 2 is accomplished by performing a sequence of operations that follows:

1. Obtain from the simulation model the values of B_i and $i(f,n)$, representing, respectively, the number of PCB i waiting at the burn-in ovens and the PCB type of the n th PCB waiting at the functional test location f , $1 < f < 8$, $n=1,2,\dots$ (PCBs at this time are waiting in a single queue before their designated functional test location)
2. Initialize to 0 the queue waiting time, denoted by $Q_{j(f)}^f$, of each functional test station $j(f)$ (in terms of the cumulative processing time of PCBs to be processed)
3. According to their waiting sequence, PCBs are allocated one by one to the test station with the shortest queue. Once a PCB is allocated to a test station (note that there can be multiple test stations at each location) the queue waiting time at that station is updated by to it the PCBs expected process time for the functional test. Clearly the updated queue waiting time represents the expected sojourn time (waiting plus processing experience by the PCB before it is ready for the burn-in test.
4. The allocation process described in operation 3 for a functional test location stops when either the shortest queue waiting time at any test station is equal to or greater than 16h or all the PCBs waiting at that test location have been allocated to a test station.
5. At the end of the allocation process, let $L(i)$ be the number of type i PCBs whose expected sojourn times are less than or equal to 16h and S be the set of PCB types for which the expected sojourn times at all designated test stations are equal to or greater than 16h. If the set S contains all the PCB types (which means that no more than $L(i)$ for PCB i is expected to complete its functional tests within the next 16h), go to step 4, otherwise, go to step 3. In other words, if for some PCB of type I the expected sojourn time of all the units and the functional test is less than 16h, only then is there a need to examine the I/C test to check if any more parts of that type could arrive at the functional test within the specified 16h.

Step 3. Estimation of the number of PCBs at the I/C test that will become ready for the burn-in test within the next 16h.

Denote by $j(c)$ the j th test station at the I/C test location c , $c=1$ and 2 . Step 3 is accomplished by performing the sequence of operations that follows.

1. Obtain from the stimulation model the value of $i(c,m)$ and $b(c,m)$, representing, respectively, the type and number of PCBs in the m th batch waiting before I/C test location $c=1$ and 2 ; $m=1$ and 2 :
2. Initialize to 0 the q th length, denoted by $Q_{j(c)}^c$, of each I/C station $j(c)$ in terms of the cumulative times of PCBs to be processed.

3. According to the waiting sequence of their batch, PCBs in each batch at the I/C test are allocated one by one to the test station with the shortest queue. The expected sojourn time, $T_{(c,m),k}^c$ of the k th PCB in the m th batch for the I/C test is given by



$$T_{(c,m),k}^C = \min_{j(c)} \{ Q_{j(c)}^C \} + k \cdot E_{i(c,m)}^C + ST$$

where $j(c)$ is an I/C test station designated to that PCB type $i(c, m)$, $Q_{j(c)}^C$ is the queue length of the test stations $j(c)$ just before the first PCB in this batch is allocated, and ST is the setup time when an I/C test station switches from one batch to another.

4. Once the sojourn time of a PCB for the I/C test is obtained, we immediately compute its expected total sojourn time, $T_{(c,m),k}$, before it becomes ready for the burn-in test. $T_{(c,m),k}$ is given by

$$T_{(c,m),k} = \max \left\{ Q_{(c,m),k}^C, \min_{j(f)} \{ Q_{j(f)}^F \} \right\} + E_{i(c,m)}^F$$

where $Q_{j(f)}^F$ is the queue length of its designated functional test station $j(f)$. Notice that the quantity

$$\max \left\{ Q_{(c,m),k}^C, \min_{j(f)} \{ Q_{j(f)}^F \} \right\}$$

is the point in time when a functional test station first becomes free after the PCB's expected completion time for the I/C test. If $T_{(c,m),k}$ is greater than 16 h, the PCB type $i(c, m)$ is added to set S; otherwise, $L[i(c, m)] = L[i(c, m)] + 1$.

5. The queue length, $Q_{j(f)}^F$, of the functional test station $j(f)$,

$$j(f) = \arg \min_{j(f)} \{ Q_{j(f)}^F \}$$

is immediately updated by

$$Q_{j(f)}^F = Q_{j(f)}^F + \max \{ 0, Q_{(c,m),k}^C - Q_{j(f)}^F \} + E_{i(c,m)}^F$$

6. The queue length of the I/C test station, however, is not updated until the expected total sojourn times for all the PCBs in the batch are computed. Once the full batch is allocated to a test station $j(c)$, the queue length of that test station is updated as follows:

$$Q_{j(c)}^C = Q_{j(c)}^C + b(c, m) \cdot E_{i(c,m)}^C + ST$$

7. Step 3 continues until either PCBs in all the batches waiting at the I/C test have been allocated to a test station or the set S contains all the PCB types.

Step 4. Computation of the total expected number of PCBs that will become ready for the burn-in test within the next 16 h

This number, denoted by $E(i)$ for PCB type i , is given by

$$E(i) = B(i) + L(i)$$

Acknowledgments

This research was supported by Northern Telecom's University Initiative. We gratefully acknowledge the assistance of Chahram Bolouri, Claude Carrier, and Y van Girard of Northern Telecom. We thank the guest editors and the two anonymous referees for careful and extensive comments and suggestions that helped improve the paper.



References

- Ahmadi, J.H., Ahmadi, R.Z., Dasu, S., and Tang, C.S., "Batching and Scheduling Jobs on Batch and Discrete Processors," *Operations Research*, Vol. 40, pp. 750-763 (1992).
- Akella, R., Rajagopalan, S., and Singh, M.R., "Part Dispatch in Random Yield Multistage flexible Test Systems for Printed Circuit Boards," *Operations Research*, Vol. 40, pp. 776-789 (1992).
- Awate, P.G. and Sastry, B.L.N., "Analysis and Decomposition of Transfer and Flow Lines," *Opsearch*, Vol. 24, pp. 175-196 (1987).
- Buzacott, J.A. and Yao, D.D., "On Queueing Network Models of flexible Manufacturing Systems," *Queueing Systems*, Vol. 1, pp. 5-27 (1986).
- Chandra, P. and Gupta, S., "Managing Batch Processors to Reduce Lead Time in a Semiconductor Packaging Line," Working Paper No.92-10-13, Faculty of Management, McGill University, Montreal (1992).
- Chandru, V., Lee, C. Y., and Uzsoy, R., "Minimizing Total Completion Time on Batch Processing Machines," *International Journal of Production Research*, Vol. 31, pp. 2097-2121 (1993).
- Conway, R., Maxwell, W., McClain, J.O., and Thomas, L.J., "The Role of Work-in-Process Inventory in Serial Production Lines," *Operations Research*, Vol. 36, pp. 229-241 (1988).
- Dallery, Y. and Gershwin, S.B., "Manufacturing Flow Lines: A Review of Models and Analytical Results," Technical Report LMP-91-002, Department of Mechanical Engineering, MIT, Cambridge, MA (1991).
- Dobson, G. and Nambimadom, R.S., "The Batch Loading and Scheduling Problem," Working Paper No. QM-92- f03, Simon School of Business Administration, University of Rochester, Rochester, NY (1992).
- Fowler, J.W., Hogg, G.L., and Phillips, D.T., "Control of Multiproduct Bulk Service Diffusion/Oxidation Processes," *IIE Transactions*, Vol. 24, pp. 84-96 (1992).
- GAMS/LAMPS Modeling Software, Release 2.25, GAMS Development Corporation, Washington, DC (1993).
- Glassey, C.R. and Resende, M.G.C., "Closed-Loop Job Release Control for VLSI Circuit Manufacturing," *IEEE Transactions on Semiconductor Manufacturing*, Vol. SM-I, pp. 36-46 (1988).
- Glassey, C.R. and Weng, W. W., "Dynamic Batching Heuristic for Simultaneous Processing," *IEEE Transactions on Semiconductor Manufacturing*, Vol. SM-4, pp. 77-82 (1991).
- GPSS/H Simulation Software, version 3.1, Wolverine Associates, Annendale, VA (1992).
- Gurnani, H., Anupindi, R., and Akella, R., "Control of Batch Processing Systems in Semiconductor Wafer Fabrication Facilities," *IEEE Transactions on Semiconductor Manufacturing*, Vol. SM-5, pp. 319-328 (1992). [Hsu, L.F., Tapiero, C.S., and Lin, C., "Network of Queues Modelling in Flexible Manufacturing Systems: A Survey," *Recherche Operationnelle*, Vol. 27, pp. 201-248 (1993).
- Ikura, Y. and Gimple, M., "Scheduling Algorithms for a Single Batch Processing Machine," *Operations Research Letters*, Vol. 5, pp. 61-65 (1986).
- Karmarkar, U.S., "Lot Sizes, Lead Times and In-Process Inventories," *Management Science*, Vol. 33, pp. 409-418 (1987a).
- Karmarkar, U.S., "Lot Sizing and Sequencing Delays," *Management Science*, Vol. 33, pp. 419-423 (1987b). Kannarkar, U.S., Kekre, S., and Kekre, S., "Lot Sizing in Multi-Item Multi-Machine Job Shops," *IIE Transactions*, Vol. 17, pp. 290-297 (1985a).
- Kannarkar, U.S., Kekre, S., Kekre, S., and Freeman, S., "Lot Sizing and Lead Time Performance in a Manufacturing Cell," *Interfaces*, Vol. 15, pp. 1-9 (1985b).
- Lee, C.Y., Uzsoy, R., and Martin-Vega, L.A., "Efficient Algorithms for Scheduling Semi-Conductor Bum-in



Operations," *Operations Research*, Vol. 40, pp. 764-775 (1992).

Lin, L. and Cochran, J.K., "Optimization of a Complex Flow Line for Printed Circuit Board Fabrication by Computer Simulation," *Journal of Manufacturing Systems*, Vol. 6, pp. 47-57 (1987).

McDowell, E.D. and Randhawa, S.U., "A Simulation-Based Production Planning Support System for Printed Circuit Board Fabrication," *Journal of Manufacturing Systems*, Vol. 8, pp. 225-234 (1989).

Uzsoy, R., "Scheduling Batch Processing Machines with Incompatible Job Families," Research Memorandum No.94-3, School of Industrial Engineering, Purdue University, West Lafayette, IN (1994).

Wein, L.M., "Scheduling Semiconductor Wafer Fabrication," *IEEE Transactions on Semiconductor Manufacturing*, Vol. SM-I, No.3, pp. 115-130 (1988).