Simulation of real-time scheduling policies in multi-product, make-to-order semiconductor fabrication facilities

by

Meow Seen Yong

B.S. Mechanical Engineering, Cornell University, NY(USA), 1999

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2001

© Massachusetts Institute of Technology 2001. All rights reserved.

Author .....
Department of Mechanical Engineering
January 31, 2001

Certified by .....
Stanley B. Gershwin
Senior Research Scientist, Department of Mechanical Engineering
Thesis Supervisor

Accepted by .....
Ain A. Sonin
Chairman, Department Committee of Graduate Students

Simulation of real-time scheduling policies in multi-product, make-to-order semiconductor fabrication facilities

by

Meow Seen Yong

Submitted to the Department of Mechanical Engineering
on January 31, 2001, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

The performance of scheduling policies in multi-product, make-to-order semiconductor fabrication facilities is studied. The three scheduling policies analyzed are the Earliest-Due-Date (EDD), Critical Ratio (CR) and Control Point (CPP) policies. In particular, the CPP and its parameters are explored in detail through simulation experiments. The results obtained provide insights into the roles of the static buffer priority, the hedging time and buffer size parameters.

In the analysis of buffer priority, the resulting service levels are strongly characterized by its discrete nature. For the cases studied, when the buffer of a part type is assigned the lowest priority, the service level of this part type is worse than that of the other. Attempts to use hedging times to offset this effect are not successful. It is found that hedging time, as used in the readiness logic, does not lead to significant improvement in the minimal service level. Instead, the readiness logic results in significant inventory reduction when applied at the raw material buffer. Compared with the two other policies, the CPP does not perform as well when minimal service level is used as a performance measure due to the effects of static priority scheme. However, the CPP is the policy of choice in situations where one type of customer order has priority over the others.

Thesis Supervisor: Stanley B. Gershwin
Title: Senior Research Scientist, Department of Mechanical Engineering

Acknowledgments

This thesis is dedicated to my parents, my wonderful sister and Ching Ching; their love and support have never failed to inspire and motive me.

Special thanks to Dr. Stanley B. Gershwin for being a great teacher whose invaluable advice and guidance will always stay with me; to Dr. Peter L. Jackson of Cornell University whom the Lord has placed in my life to kindle my interest in manufacturing. My gratitude also goes to Omar Gzouli, Young Jae Jang, Loren Werner, Nicola Maggio, Jason Finch and Francis de Vericourt for offering help in my research in one way or another.

I would also like to take this opportunity to thank friends from Ashdown House for having such positive impact on my life; and to the brothers and sisters from BCBSG who have helped constructed many of my wonderful memories of Boston.

Most importantly, to my God and my Lord Jesus Christ who place all these wonderful people in my life: I can only spend my lifetime grasping and understanding the depth of your grace!

Contents

List of Figures

List of Tables

Chapter 1

Introduction

There are many incentives to implement a good production scheduling policy. In addition to cost-savings, a good policy enables a manufacturer to meet early delivery dates consistently. The latter is especially essential to the survival of companies in markets that resemble the make-to-order environments.

This thesis studies the Control Point Policy (CPP) and compares the performances of other real-time scheduling policies in multi-product, make-to-order semiconductor fabrication facilities (fabs). This chapter provides the background and the objectives of the study. They are explained in Section 1.1. A literature review is provided in Section 1.2, followed by a synopsis of the thesis in Section 1.3.

1.1 Background and Objectives

The CPP, a real-time production scheduling policy, is introduced by Gershwin [6]. Through extensive simulation experiments, Gzouli [7] shows that the CPP has good performance in single part-type systems. In addition, some important insights are offered on the understanding of some CPP parameters. Despite these insights, the analysis is not exhaustive. This is because the CPP is still in its early stage of development. Its behavior is not fully understood. Neither is a good set of guidelines for choosing policy parameter values established. The analysis has to be carried out through simulation experiments since numerical/analytical methods are still under

development.

With multiple part-type systems as the focus of this study, the number of parameters increases exponentially with the addition of each part type and machine. Tremendous computing efforts are required to carry out an exhaustive analysis of the parameters. Nonetheless, we attempt to understand the role of each parameter in the CPP scheduling logic through many simulation experiments on two different manufacturing systems. Experiments are also carried out using the Critical Ratio and Earliest-Due-Date policies for comparison purposes.

It is important to bear in mind that this is only the first attempt to study the performance the CPP in multi part-type systems. More extensive simulation experiments need to be carried out to reach a deeper understanding of the policy.

1.2 Literature Review

Insights on the importance of production scheduling and inventory control can be found in Silver and Pyke [14]. Hopp and Spearman [8] offers a good explanation of the underlying behavior of manufacturing systems. Extensive discussions of factory design and production scheduling issues can be found in Gershwin [5]. This book also presents the decomposition method, a quick and reliable analytical technique to predict the performance of production lines.

Gershwin [6] explains the formulation of the CPP, the study of which forms the basis of the thesis. It also summarizes the observations of Deshpande [4] on the implementation of the policy in Boeing. Readers should also refer to Gzouli [7] for simulation experiments of the CPP in single part-type systems.

A classic study of simulation experiments on semiconductor fabrication process is found in Wein [15]. He determines that a good release control of materials into the system results in far greater reduction in mean throughput time than from parts sequencing at other workstations. But Lu et al [13] and Kumar [10] subsequently show that good lot sequencing rules can also lead to significant reduction in queuing time and the standard deviation of cycle time. Experiments on a multi-product

semiconductor fabrication facility with different process flows are carried out by Kim et al [9]. They introduce new dispatching rules that give smaller mean tardiness and are better in terms of average and/or standard deviation of cycle times.

1.3 Thesis Outline

Chapter 2 explains the type of manufacturing environment analyzed in the thesis as well as the applicable scheduling policies. The findings from a study of commercial scheduling software are also included in the same chapter.

The Critical Ratio and Control Point policies are explained in detail in Chapter 3. A description of the toy fab, which is a small-scale representation of an actual fabrication facility, is also provided. The toy fab forms the basis of our investigation into some of the CPP parameters from chapters four to six.

Chapter 4 analyses the effects of buffer priority parameter on the service levels. Chapter 5 investigates the possibility of using hedging time to offset the effects of static buffer priority to adapt to another performance measure. A more detailed study of hedging time is carried out in Chapter 6. Chapter 7 investigates the performance of the CPP in a model of an LSI Logic (LSIL) fab. The buffer size parameter is also studied using this model.

The results of other scheduling policies applied to both the toy fab and the LSIL fab models are presented in Chapter 8. Finally, concluding remarks and a list of suggested future research directions are provided in Chapter 9.

Chapter 2

Background on Production Scheduling

Many production scheduling policies come into existence today to serve different needs. Policies that are applicable in one type of manufacturing environment may not be suited to another. This chapter describes the type of manufacturing environment that is analyzed in the thesis as well as the applicable scheduling policies. Since many scheduling decisions are made through software applications, the findings from a study of commercial software are also provided.

Section 2.1 summarizes the characteristics of many different manufacturing environments and their differences. Section 2.2 explains the different classes of scheduling policies and their applicability in different manufacturing systems. The chapter ends with a discussion of the characteristics of scheduling software found in the market today.

2.1 Manufacturing Environments

The manufacturing environments have been undergoing many different changes due to both internal and external factors. Internal factors consist of the introduction of new physical processes. Externally, customers' behaviors are also changing.

2.1.1 Changes in Manufacturing Processes

Traditional manufacturing,such as the automobile production, has machines lined up in series in the job floor. Raw materials enter through one end of the line and the finished products through another. The materials visit each machine only once and proceed downstream after each operation.

In recent decades, semiconductor fabrication facilities (fabs) have brought attention to a host of new manufacturing issues. Due to process requirements, wafer lots re-visit each machine several times. This is called the re-entrant flow of materials. Further, no longer is each operation associated with a single machine. A workstation that consists of multiple identical machines is commonly found in a fab. As a result, many new scheduling policies have been developed.

2.1.2 Changes in Consumer Behavior

Products in the market are increasingly customized to the needs of consumers. Such a change in consumer behavior has a major impact on manufacturing decisions. A manufacturer has to decide whether to produce in a make-to-stock (MTS), make-to-order (MTO) or in a hybrid environment. The reader is referred to Buzacott and Shanthikumar [3] for further reference.

In a pure MTS system, all parts are interchangeable. This implies that the final products have little or no customizations. The release of materials for production is closely linked to what is happening in the shop, and the associated output or input stores. The production goal, then, is to replenish the stock in the finished goods buffer as quickly as possible. If backlogging is not allowed, a customer who arrives only to find the finished goods buffer empty will be turned away. Typical products of such an environment are commodity goods such as light bulbs.

In a pure MTO environment, the raw material release process is independent of whatever is happening in the job floor, and its associated input and output stores. It depends only on the arrival of a customer order externally. Further, each order is unique. Associated with an order is customer lead time, the amount of waiting

time the customer is willing to put up with before delivery. This type of environment arises when each product is made to comply with some very specific customer requirements. An example of such an environment is found in that of a tailor shop. A good scheduling policy in this environment enables the manufacturer to meet due dates frequently.

Along a supply chain, upstream members are more likely to belong to MTS environments while the downstream members, being closer to end-customers, are likely to be in the MTO environments. In general, however, most companies share the characteristics of both environments.

2.2 Scheduling Policy Classifications

There exists many different ways to classify scheduling policies. One method that has been discussed is classification according to the manufacturing environment. Another possibility is to differentiate between scheduling policies that respond to the state of the system and those that do not.

The system state is made up of all the events happening in the plant in real time. This includes the level of inventory, machine failures and repairs. Scheduling policies that make decisions in response to these events are called real-time policies. Several examples of real-time policies include the Critical Ratio (CR) and the Control Point (CPP) policies. On the other hand, traditional material requirement planning is an example of policies that do not take into account the state of the system. It schedules production based only on demand forecasts.

Only real time scheduling policies are applicable to the MTO environment. This is because manufacturers have to determine feasible due dates based on the state of the system. Once the due dates for different orders are determined, scheduling decisions have to take into account the remaining customer lead time (i.e. the time until the due date).

In this thesis, only real-time scheduling policies are used in the simulation experiments. To arrive at a fair comparison of these policies, we standardize a scheduling

decision that does not belong specifically to a particular policy. It is the blocking-before-service (BBS) logic in which a part is not loaded into a machine if the downstream buffer is full. We standardize the analyses of various scheduling policies under the BBS logic.

2.3 Production Scheduling Software

A study carried out by Werner and Yong [16] focused on the evaluation of commercial scheduling software. It was found that most of them were intended to be tools that facilitate production planning. Many packages are not based on any specific scheduling algorithm. For some of those that are, the exact scheduling procedure is kept confidential by the vendor. Only a general description of the underlying philosophy is made known to the customers. Often, this type of packages is treated like a black box that customers are persuaded to purchase mostly due to good marketing.

Chapter 3

Production Scheduling Policies and Experiments

In Chapter 2, we discuss the characteristics of make-to-stock (MTS) and make-to-order (MTO) environments. The scheduling policies that can be used are different for each environment.

This chapter presents two scheduling policies that can be implemented in an MTO environment. They are the Critical Ratio (CR) and Control-point (CPP) policies. Together with the Earliest-Due-Date policy, these scheduling policies are used in the simulation experiments performed for this thesis.

Section 3.1 introduces the CPP and its scheduling parameters, followed by a discussion of the CR policy in Section 3.2. Performance measures that are used to differentiate scheduling policies are discussed in Section 3.3. Section 3.4 describes the toy fab model that is used as a basis of simulation experiments for a large part of this thesis. The chapter ends with a discussion of the statistical method used to analyze the results of the experiments.

3.1 The Control Point Policy

The CPP is a real-time scheduling policy introduced by Gershwin [6]. The scheduling decisions are made in response to the state of the system, such as the inventory level,

machine failure and repair. This section presents a summary of the CPP scheduling logic. It starts with the explanations of the terms and definitions that are necessary for the understanding of the CPP scheduling logic.

3.1.1 Terminologies and Definitions

A control point is a resource where the CPP scheduling logic is applied. Usually, a machine is selected as a control point. All other resources can be called non-control points.

The policy has three parameters: buffer priority, buffer size and hedging time. Each buffer at the control point is assigned a priority. The priority scheme aids the part-selection process in the scheduling logic. The term buffer refers to the amount of inventory between two consecutive control points. Since not all workstations are control points, and the next control point is likely to be several workstations away, buffer size places a limit on the maximum amount of inventory among all workstations that are between two consecutive control points.

The hedging time parameter is used in the readiness condition of the CPP scheduling logic. Its main purpose is to prevent the flow of materials that are too early compared to their due dates. Consequently, resources are freed up in case parts that are more urgent arrive in the immediate future. Hedging time is closely related to the remaining cycle time, which refers to the amount of time remaining until the end of the manufacturing process.

A part is considered available if it is physically present at a control point, and if the buffer immediately downstream is not full. It is locally ready at a control point if the sum of the current time and the hedging time is later than the due date for the part.

The last definition refers to the availability of a machine. A machine in a workstation is considered available if it is operational and idle. Machines undergoing repair and maintenance are not considered available.

3.1.2 Scheduling algorithm

There are three alternative versions of the CPP scheduling logic: the time-based, token-based and surplus-based versions. The time-based version is presented here since it is more appropriate for an MTO environment. The discussion of the other versions can be found in Gershwin [6].

As mentioned in the previous section, the scheduling logic is only applied at the control point. At all other resources, a simple, sensible policy such as EDD can be used.

In real time, the scheduling logic is listed below.

When a control point becomes available:

(a) Select the part with the earliest due date from the highest priority buffer.

(b) Determine if a part is *available*. This is known as the availability logic.

(c) If a part is available, determine if it is *ready*. An available part that is ready satisfies the readiness logic.

(d) When a part from the highest-ranking buffer satisfies both the availability and readiness criteria, it can be loaded into the available machine.

(e) Otherwise, go to the next highest-priority buffer and repeat the first three steps.

(f) If no available parts are locally ready, let the control point be idle. It can also be used for some other important activities such as preventive maintenance.

3.1.3 Policy Implementation

Two activities are required to implement the policy: setting it up in advance, and executing it. The method of execution has been discussed in the previous section. In the preparatory phase, the followings are carried out:

- Select the number of control points and their locations.
- Rank order all the buffers in front of each control point.

- Determine the appropriate hedging times.
- Determine the buffer sizes between consecutive control points.

Even in the preparatory phase, these decisions are made chronologically. The last two activities are carried out only after the first two have been determined. This is because the hedging time and buffer size parameters can only be determined after the control point locations have been chosen. These parameter values will also be influenced by the buffer priority scheme selected.

The CPP provides a clear way of selecting which part to work on at any time. It tells the operators what to do even when the schedule is disrupted due to machine failures and other variability. Moreover, it requires no elaborate calculations in real time.

Another advantage of the CPP is the ease of implementation. Good performance may only require scheduling at the control points. Since not all the workstations need to be control points, the implementation cost can be significantly lower compared to other policies.

3.2 The Critical Ratio Scheduling Policy

The Critical Ratio scheduling policy (CR) is popular among multi-product, MTO environments. Unlike the earliest-due-date policy (EDD), the CR policy relies on both the due date and the processing time information to rank order parts. However, there is no standard definition on critical ratio. Reader can refer to Hopp [8] and Silver et al. [14] for some of the common definitions.

3.2.1 Scheduling Algorithm

We define critical ratio as the ratio of the remaining customer lead time (i.e. due date minus the current time) to the remaining processing time. This ratio is computed for all the parts waiting in front of the workstation. A part with the lowest CR value is selected for the next operation. When parts are late, their CR values will fall below

zero. Late parts have priority over non-late parts. It is also important to note that when parts have the same lateness, this definition gives priority to the part with the shortest remaining processing time.

The definition we use in the thesis includes a critical ratio factor. This factor, which when multiplied by the remaining processing time in the denominator, is supposed to give priority to certain part types. A high priority part type has a factor value of more than one.

3.2.2 Policy Implementation

The CR policy is dynamic because the ratios change as time progresses. For the same reason, however, it has the disadvantage that new computations have to be performed each time a job completes processing. The computation effort can become enormous in a multiple part-type, re-entrant system such as a semiconductor fab.

In addition, the CR policy is purely a lot-sequencing rule. It does not specify how materials should be released into the system. Nonetheless, the scheduling logic is easy to follow. It offers a clear way of selecting which part to work on at each workstation.

3.3 Performance Measures

In an MTO environment, the ability of a manufacturer to meet the delivery date is crucial. Customers are likely to be turned away if the due dates cannot be met. This calls for a due date based performance measure that differentiates scheduling policies in an MTO environment.

We propose the use of service level as a performance measure. It is defined as the percentage of the finished parts that meet their due dates. A policy that leads to a better service level performance is preferred.

Such a performance measure becomes harder to be defined in a multi-product system. In a production system where all customer orders are equally important, a manufacturer is interested in maximizing the service levels for all orders. In this case, the minimum service level of all the product types can be a suitable measure

to differentiate scheduling policies. We call this performance measure the minimal service level measure. In other cases, a manufacturer may be more interested in maximizing the service level of one type of customer order if it has priority over the rest. As a result, a policy that does well in one measure may not be as good when another measure is used. Reader should note that the extent of earliness or lateness is not reflected in this measure however.

Often, there exists a trade-off between the amount of inventory in the system as well as the service levels of all part types. A policy that leads to high service level may do so at the expense of high inventory level. This situation calls for the need to define another performance measure that is based on the amount of work-in-progress inventory (WIP) in the system. A policy that leads to high service levels with low WIP in the system is preferred.

For all the simulation experiments carried out, both the service levels as well as the amount of WIP are used as performance measures. We define the WIP as all materials in the system, except those at the raw material buffers. The raw material buffer is not included in the computation because the inventory level is influenced more by business decisions between supply chain members rather than the scheduling rules.

3.4 Simulation Model Description

The CPP is still in its early stage of development. Its behavior is not fully understood. Neither is a good set of guidelines for choosing policy parameter values established. Simulating a real fab with hundreds of workstations as well as process routes for many different part types do not allow for many choices of the values of each parameter. For example, in a simple two-machine line with 10 buffers, five choices of buffer size for each buffer results in 5^{10} simulation runs. Thus, optimal results for each of the parameter cannot be practically found through the analysis of real fab data.

To facilitate the analysis of the CPP, a small representation of a wafer fab is developed, which we refer to as a toy fab. Before it is introduced, the simulation software

and the computing resources used are discussed in Section 3.4.1. Assumptions made about the toy fab and the scheduling policies are explained in Section 3.4.2. A detailed description of the toy fab and its demand model can be found in Sections 3.4.3 and 3.4.4 respectively. Finally, the chapter concludes in Section 3.5 with a description of the statistical method used to analyze research results.

3.4.1 Simulation Software and Computing Resources

Before we could decide on the simulation software to use, there were a few important requirements that the chosen software had to satisfy. Most importantly, the program had to be consisted of three basic components: one representing the factory, the other representing the scheduling policy and finally, a simulation manager.

The factory module is representative of the physical states of the system, such as machine failures and repairs, and whether the due date of each customer order has been reached. These states are reported to the simulation manager. The scheduling policy determines the appropriate response to the events happening in the factory. The simulation manager relays these decisions back to the factory to be implemented.

Another important feature of the simulation software is its ability to run many simulations with different policy parameters in a short amount of time. This is a valuable feature since the purpose of the activity is to test many parameter values.

We were not able to find a commercial package that satisfied our requirements. As mentioned in Section 2.3, many of the commercial packages act like black boxes, where users are not allowed to find out the algorithm used in production planning and scheduling. Others that allow the flexibility for the configuration of new policies often come with elaborate graphical-user-interfaces (GUI) that slow down a simulation run. As a result, Gzouli (200) decided to build a multi-platform JAVA-based simulation program for his research. However, the program could only analyze a single part-type, one-failure-mode system with one process flow. We have extended his program to enable the analysis of multiple part types, process flows and machine failure modes. These modifications are necessary for the generation of the results in this thesis.

The computing resources used were *Cell3* (Pentium III 700 MHz with 256 MB

RAM), *Hierarchy* (Linux machine equipped with dual processors of Pentium III 700 MHz) and seven Windows NT workstations with 800 MHz, 256 MB RAM each in the MIT Course 1.00 NT Lab.

3.4.2 Model Assumptions and Limitations

Listed here are some common assumptions made across all scheduling policies and for the toy fab model.

- No setups. This means that no setup time is incurred when different part types or different processing steps of the same part type are carried out in each workstation.
- Material is transported in units of one without delay. This means that there are no batches in the system, and transportation time is assumed zero.
- Machines operate asynchronously. Parts can be loaded when the proper authorization is received according to the algorithm of the scheduling policies. There is no clock that synchronizes events.
- Machines can only fail while operating. This is known as *operation dependent failures* (Buzacott and Hanifin [2]).
- The information of the system state is transmitted instantaneously to the scheduling module. In return, the decisions are transmitted back to the system without delay.
- The raw material supply is also assumed infinite. That is, the factory does not have to wait for raw materials.
- All the scheduling policies operate under blocking-before-service (BBS) algorithm. This means that a part will only be loaded when there is space in the downstream buffer to which it can proceed after processing.
- Machine processing times are modeled using lognormal distributions.

- Machine failure and repair times are modeled as exponential variables.

There are several reasons for making these assumptions: most importantly, they simplify the analysis by reducing the number of variables considered. Secondly, the BBS assumption standardizes all the control policies so that fair comparison can be made. The assumption of infinite raw material supply is to allow the analysis to concentrate on the factory, without being influenced by some of the disruptive effects of upstream supply chain members. These assumptions, except the last two, are also applicable to the simulation analysis of LSI Logic fab in Chapter 7.

3.4.3 Toy Fab Formulation

Originally developed by Gzouli [7] for a single part type, single process flow, the toy fab has now been modified to incorporate two different process flows for two part types. The toy fab retains many of the essential characteristics associated with a semiconductor fab such as re-entrant process flow. The toy fab also consists of multi-server stations made up of several identical machines working in parallel. Both single-server and multi-server stations are collectively called *workstations* in this thesis.

A pictorial representation of the model and the process flows are shown in Figure 3-1. The first part type visits workstations 1 and 2 three times each while the second part type visits each workstation twice in the sequences shown.

The buffers upstream of each workstation are labeled B_{sq}, where the subscript s represents the process step to be expected at the workstation while q indicates the part type. The raw materials as well as the final goods buffers are labeled as B_{0q} and B_{fq} respectively. Each buffer is distinct in terms of the part type and its process step. As such, even though there are only two part types and workstations in the system, there are more than two buffers because of the re-entrant process flows. The hedging times associated with each buffer are labeled H_{sq}, where the definitions of the subscripts are the same as those for buffers. These notations are used extensively in later chapters.

Workstation 1 (WS 1) consists of one machine. Workstation 2 (WS 2), on the

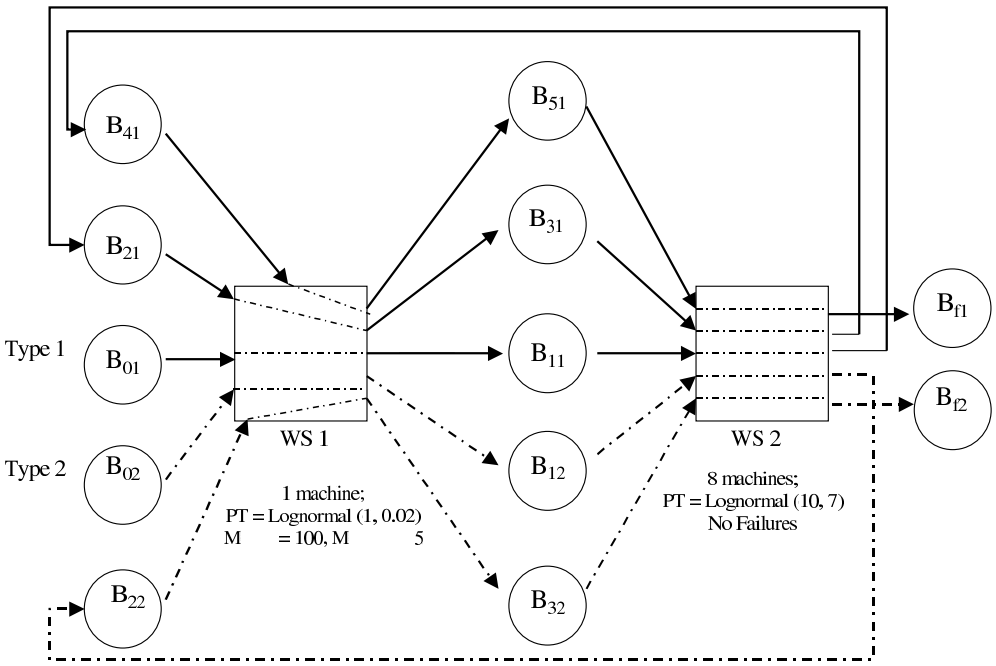

Figure 3-1: Toy Fab

other hand, is made up of eight identical machines capable of processing a total of eight parts simultaneously. WS 2 represents all the operations that a part receives before returning to WS 1. It can include, for example, photolithography and etching operations after the wafers are cleaned at WS 1 before returning for further cleaning operations.

The mean processing time for the operation at workstation 1 is chosen to be one order of magnitude smaller than that of the machines in workstation 2. Further, the processing time at workstation 2 has a coefficient of variation equal to 0.7, modeling the aggregate variability of all the operations a part goes through downstream of workstation 1. The machines in workstation 2 do not explicitly fail. The failure and repair times are included in the processing time variance.

To simplify the analysis, the processing time at each workstation is the same for

both part types in each visit. The processing time follows a lognormal distribution, with the parameters for each workstation shown in Figure 3-1. The total inventory in the system used as part of our performance measures includes all parts found within the dotted line. This includes everything in the fab except the raw material buffers.

The failure, repair, and processing times as well as the number of visits to each workstation are chosen such that the line is roughly balanced. That is, the capacity of the first workstation is very close to that of the second workstation. Thus, there is no bottleneck that will affect the selection of control points. This is an important criterion for the simulation experiments.

3.4.4 Demand Model

In a make-to-order (MTO) environment where customer arrivals as well as due-dates for orders are uncertain, many interesting phenomena arise. Orders that arrive later may have earlier due-dates than orders that have arrived earlier. The priority of the existing orders in the system may often have to be re-sequenced and careful scheduling is required to meet the due dates. Further, priorities of orders often change over time. Business issues may cause an order to rise in importance. In fact, a commonly used term in the semiconductor industry to describe such orders is "hot". Even if all orders start up being equally important, a specific order may become hot when an important customer urgently demands order fulfillment.

The demand model used in Gzouli [7] is used in this analysis to model the order arrival process as well as due date assignment for each order. The *due date* of an order is defined as the sum of the arrival time of an order and the customer lead time as defined in Section 2.1.2. Gzouli considered several models of demand rate and customer lead time combinations. The performance of cases of low demand rate and long customer lead time did not distinguish among scheduling policies.

Following the recommendations of Gzouli, short customer arrival interval and lead time are used in the toy fab model. Both of these random variables are assumed normal distributions with a coefficient of variation of 0.5 each. On average, the customer lead time is set to be 10 times the processing time of a part type. This amount of

customer lead time falls within the range of cycle times commonly experienced in the semiconductor industry.

Some important parameters of the model are provided in Table I. The utilizations for each workstation are 94.6% and 93.3% for WS 1 and WS 2 respectively. Readers should refer to Wein (1988) on the computation of utilization.

Demand model parameters	Part Type 1	Part Type 2
Average Takt time, t	6.4	7.2
Standard Deviation of order arrival interval. σ_t	3.2	3.6
Average Customer lead time, ℓ	330	220
Standard Deviation of lead time. σ_ℓ	165	110

Table I: Toy Fab Demand Model

One reason for choosing a higher demand rate for part type 1 is to represent an increasing demand for more complex products in the semiconductor industry. Part type 1, which requires more processing steps, represents wafers with more metal layers. Part type 2 represents the current product produced in a fab.

The length of the transient period is also determined through several pilot runs using the demand rates and customer lead times for each product defined above. For scheduling policies studied, a run of 300,000 time units with statistics collected over the last 150,000 time units are sufficient to capture the steady state performance. It is important to run the simulation long enough so that all possible events, such as machine failures and repairs, occur a large number of times to depict the steady state behavior of the system.

3.5 Statistical Analysis

Due to the random nature of the simulation output, independent replications are needed to obtain credible estimates of the results. The procedure employed is to make enough replications until the half-length of the 95% confidence interval divided by the output of interest is less than 3%. In this case, the output is said to have a

relative error of 3%. Law and Kelton [11] provides a good discussion of statistical analysis using relative error.

In each simulation run, the service levels of both part types as well as the respective work-in-progress inventory (WIP) are the outputs of interest. Often the service level of one part type is much higher than that of the other. In such cases, the relative error of the higher service level easily satisfies the statistical requirement merely after ten replications while that of the other does not. One possible reason is due to the constraints of finite capacity. When most of the capacity is allocated to one part type, very little is left for the other and the resultant service level variability of the less important part is high.

The guideline used in this thesis is to achieve the 3% relative error requirement for the higher service level. This is because the corresponding type of customer order has priority over the other. On average, this requires about fifteen replications of each run.

Chapter 4

The Analysis of Buffer Priority Parameter

The CPP and its scheduling algorithm have been discussed in Chapter 3. A fair comparison of the CPP with any other scheduling policy requires an understanding of each of the CPP parameters. As the manufacturing system becomes more complex, many sets of CPP parameters are needed for each part type, workstation and re-entrant flow. Their interactions become harder to analyze.

This chapter presents some of the observations of the effects of buffer priority schemes on service level. The main reason to begin the analysis of the CPP parameters with buffer priority assignment is that it is partly a business decision determined in advance, depending on the importance of each type of customer order. Hedging time and buffer size parameters can only play their roles after the control points and the buffer priority scheme in the system have been defined. However, good guidelines for determining the location of the control points have yet been determined and thus, hedging time and buffer size parameters cannot be easily analyzed.

Prior to the discussion of experimental procedure in the analysis of buffer priority schemes, Section 4.1 discusses some intuitive understanding of the nature of this parameter. Section 4.2 explains the experimental procedure, beginning with an explanation of the performance measures used. Both sections 4.3 and 4.4 discuss the experimental results and the chapter concludes with a summary of the findings in

Section 4.5.

4.1 Intuitive Discussion

In the CPP scheduling algorithm, part-selection begins from the highest priority buffer. Then, the availability and readiness criteria determine if a part from that buffer should be loaded for the next operation. If no top priority part can be loaded, the policy requires the second highest priority buffer to be checked.

In a single-part-type transfer line where there is no re-entrant flow, there is no any priority scheme since there is only one buffer at each workstation. A multiple-part-type transfer line, however, requires the assignment of buffer priority depending on the relative importance of each part type.

In a re-entrant process flow system, the same part type revisits the same workstation multiple times. Since a buffer is limited to one part type and one processing step, there are multiple buffers at a workstation even for a single part-type system. As such, the buffer priority needs to be carefully assigned even for one-part-type system in order to improve the performance measures.

Many possible priority schemes arise in multiple part-type, re-entrant process flow systems. Among all the possible combinations, careful selection has to be made depending on the performance measure of interest.

The selection of a suitable priority scheme becomes harder when the performance of the fab is determined by minimal service level. This is because the buffer priority parameter cannot assign the same priority to different buffers. It is not clear how the assignment of a low priority for one part type at one process step can be offset by the designation of a higher priority at other process steps so that its service level is not affected badly. The other CPP parameters such as hedging times and buffer sizes can possibly offset this effect, but it is not known at this point what these parameter values should be.

4.2 Experimental Procedure

The effects of buffer priority schemes are studied through computer simulations. The performance measures used are limited to the service levels of both part types. A full analysis of the effects of buffer priority scheme requires the selection of good hedging time values for each buffer and the buffer sizes in between consecutive control points.

As mentioned in Chapter 1, even though a full analysis is not possible, the effects of buffer priority schemes can be observed by making a few simplifications on the other CPP parameters. First, the buffer sizes are assumed infinite. The second simplification is the assumption of infinite hedging time for each buffer. This means that when a part is released into the system, it always satisfies the readiness criterion. It is important to note that this hedging time assumption is limited only to the study of buffer priority schemes in this chapter. Chapter 5 explores the possibility of good hedging time values offsetting the effects of buffer priority schemes.

In the toy fab model, there are five buffers at each workstation. As a result, the number of possible priority schemes is $(5!)(5!) = 14400$. A simulation run on the fastest available computer resource described earlier would require about 15 seconds. This means that a total of 60 hours is needed to finish only one simulation of all possible schemes.

This calls for further simplification of our proposed analysis. Consequently, four sets of buffer priority schemes are chosen and they are summarized in the Table I. A detailed listing of all the priority schemes in these four cases are provided in Appendix A.

The first two cases consider all possible buffer priority schemes at workstation 2 (WS 2) by fixing the scheme at workstation 1 (WS 1) constant. The last two cases, on the other hand, consider all possible schemes at workstation 1 (WS 1) while fixing that of workstation 2. Across from left to right in each row, the buffers are listed in decreasing order of importance. The priority values are labeled from 0 to 4, with a value of 0 being the highest priority.

There are several reasons for the choice of the priority scheme in each case. Gzouli

Buffer Priority Scheme	Workstation (WS)	Buffer Priority, in decreasing order of importance				
		0	1	2	3	4
Case 1	WS 1	B_{41}	B_{22}	B_{21}	B_{02}	B_{01}
	WS 2	All 120 possible permutations				
Case 2	WS 1	B_{22}	B_{41}	B_{21}	B_{01}	B_{02}
	WS 2	All 120 possible permutations				
Case 3	WS 1	All 120 possible permutations				
	WS 2	B_{51}	B_{32}	B_{31}	B_{11}	B_{12}
Case 4	WS 1	All 120 possible permutations				
	WS 2	B_{12}	B_{11}	B_{32}	B_{31}	B_{51}

Table I: Priority Schemes

[7] observed in a single part-type system that priority given to returning parts in the first workstation result in better service level and lower inventory. As for the second workstation, priority given to parts in their first visits to WS 2 results in better performance.

Taking into account his observations, priority is given to the returning parts in both cases 1 and 2 at WS 1. However, they differ in the following way: if the WIP of one part type has the highest priority, then the raw material buffer of that part type receives the lowest priority. The main difference between cases 3 and 4 is that priority is given to a different part type in each case. In case 3, the returning parts of part type 1 (p1) are given priority over those of part type 2 (p2). In case 4, priority is given to p2 early in their visits to workstation 2.

In each of the cases, twenty replications are made to achieve the statistical requirement discussed in Section 3.5. Since hedging time values are kept very large for each individual buffer, the scheduling algorithm behaves like the Earliest-Due-Date policy with a buffer priority scheme.

4.3 Results for Constant Priority Scheme at Workstation 1

In case 1 where priority is given to the returning WIP for p1, with p1 raw material receiving the lowest priority, the results are plotted in Figure B-6.

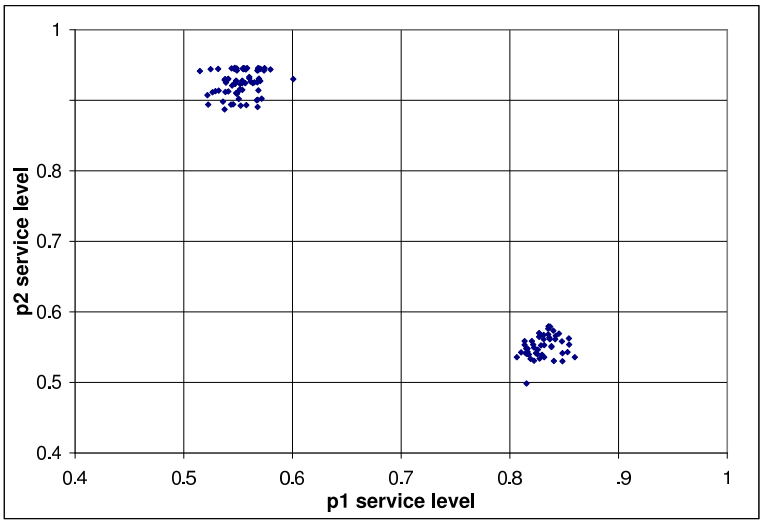

Figure 4-1: Service levels for Case 1, with all possible priority schemes at WS 2

The figure shows two distinct, very tight, clusters. The priority schemes in each cluster share a common characteristic: when the lowest priority at WS 2 is assigned to the buffer of a part type, this part type has worse service level as compared to that of the other. The upper cluster in which p2 service levels are higher corresponds to buffer priority schemes where one of the p1 buffers at WS 2 is assigned the lowest priority. On the other hand, when a p2 buffer is assigned the lowest priority at WS 2, p1 service levels fare much better, as shown in the lower cluster.

This characteristic is also observed for case 2 where a different buffer priority scheme is applied to WS 1. Results from the analysis of case 2 are shown in Figure 4-2.

Again, each of the clusters is characterized by the designation of the lowest buffer priority to the buffers of one part type. In this case, however, the service level of p1 is better even when the lowest priority for the buffers at WS 2 is assigned to a p1 WIP. Nonetheless, p1 performs much better when the lowest priority is assigned to a p2 WIP.

Another observation can be derived from the comparison of cases 1 and 2. In case 1, where p1 is assigned the lowest priority at WS 1, the best p1 service level is very much less than that of p2. When p2 is assigned the lowest priority at WS 1 in case 2,

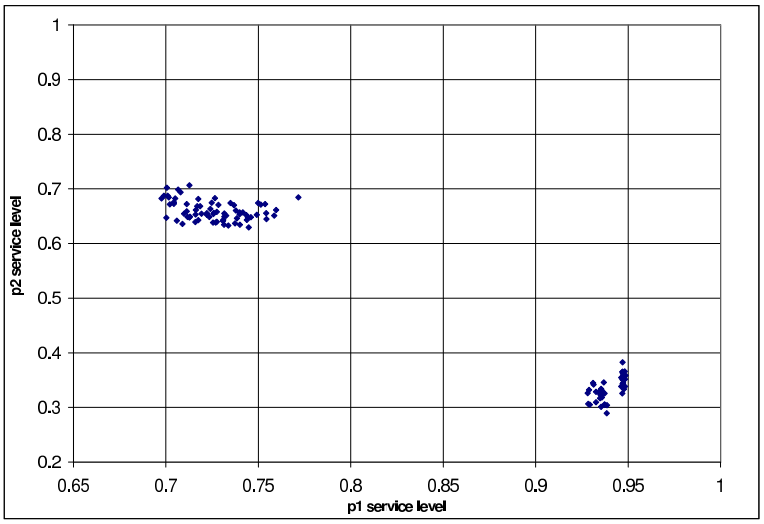

Figure 4-2: Service levels for Case 2, with all possible priority schemes at WS 2

p1 has a much better performance than p2 in all priority combinations. It seems that the assignment of the lowest priority at WS 1 also affects the service levels of either part type under the case of infinite hedging time and buffer size. This observation leads to the analysis of cases 3 and 4 and the results are discussed in Section 4.4.

When the service level of each part type is plotted individually against the priority schemes, more interesting phenomena are observed. Figure 4-3 is a graph of the service levels of p2 only, corresponding to all the buffer priority schemes at WS 2.

There are four distinct regions plotted in Figure 4-3. As explained, when a p2 WIP is assigned the lowest priority (i.e. B_{22} or B_{22} has priority value equal to 4, p2 service level is much lower than when a p1 WIP has the lowest priority. This corresponds to all p2 service levels below 0.62. Otherwise, the performance of p2 is better than that of p1. A closer look reveals that the assignment of priority to different p1 buffers has an influence on p2 service levels. When buffer B11 is assigned a priority of value 4, which is the lowest priority, p2 service levels are below 0.92 as depicted by the circles in the graph. As the lowest priority is assigned to p1 WIP returning for more advanced processing step at WS 2, p2 performance improves. In fact, p2 has the best service level when B_{51} is assigned the lowest priority. It seems that the performance is mostly influenced by the assignment of the lowest buffer priority. The designation

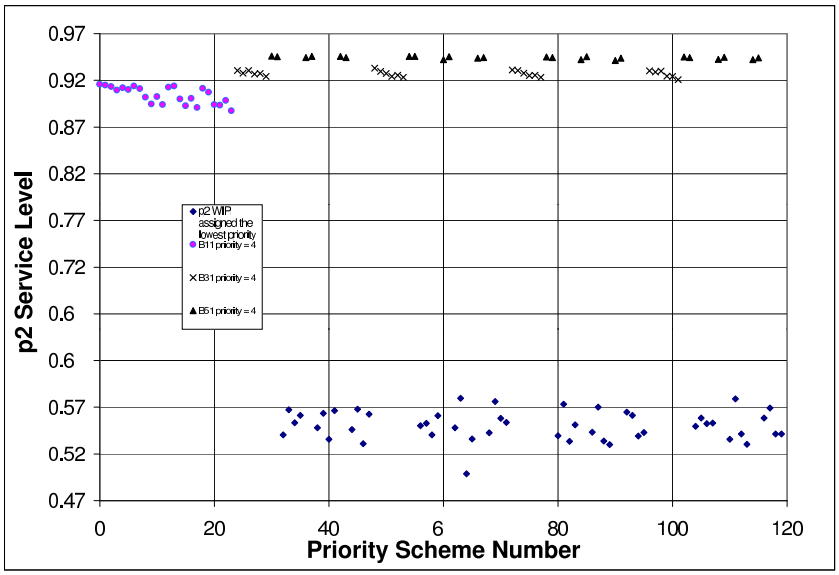

Figure 4-3: p2 service levels for Case 1

of other priority values does not change the service levels substantially.

This observation suggests strongly that the assignment of priority to p1 WIP has an influence on the performance of p2. Further, it also agrees with earlier observations by Gzouli [7] for the single part-type system where priority given to parts visiting WS 2 for the first few times leads to better service level. The service levels of p1 for case 1 are shown in Figure 4-4. There is no significant improvement in p1 service level when priority is given to B11 as compared to B51. In fact, the service levels of p1 are equally low whether the lowest priority is assigned to B_{11}, B_{31} or B_{51}.

The results of case 2 reinforce the observation made in case 1. When priority is given to p2 parts in their early visits at WS 2, p1 service level improves. As shown in Figure 4-5, the service levels of p1 are the highest when B_{32} is assigned the lowest priority.

There is a small difference, however. The service levels of p2 are distinctly higher when a higher priority is given to its own WIP visiting WS 2 for the first time. The service levels of p2 for case 2 are shown in Figure 4-6.

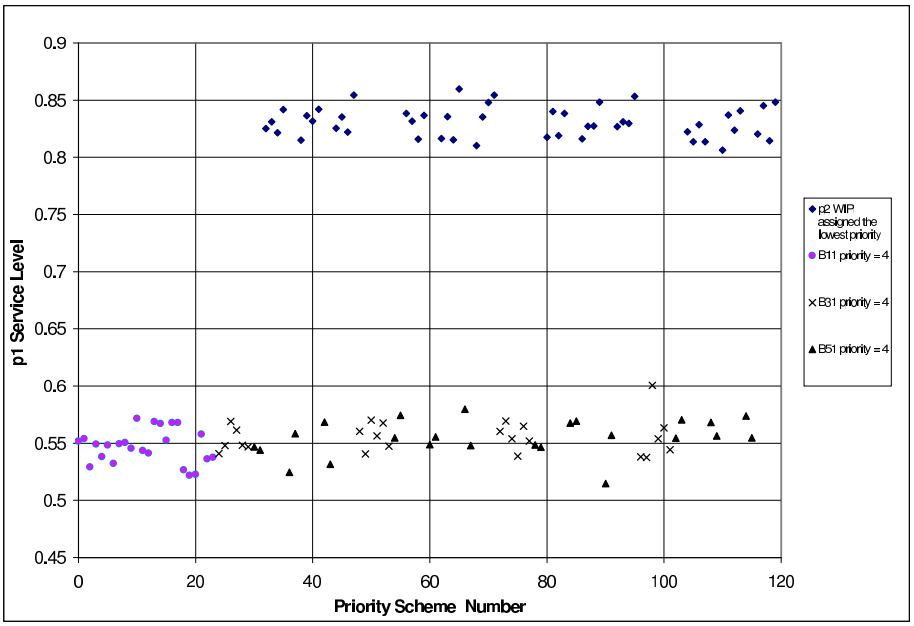

Figure 4-4: p1 service levels for Case 1

4.4 Results for a Constant Priority Scheme at Workstation 2

Keeping a constant priority scheme at WS 2, this section explores the effects of buffer priority schemes at WS 1, the entrance into the system. This leads to the analysis of cases 3 and 4.

In case 3, the service levels of both part types corresponding to all possible priority schemes at WS 1 are shown in Figure 4-7.

The same clustering of points is observed. As in cases 1 and 2, when the lowest priority at WS 1 is assigned to the buffer of one part type, this part type has worse service level compared to that of the other. This time, however, the clusters are not as tight. In the upper cluster, where p2 service levels are above 0.35, p1 service levels range from 0.64 to 0.842. In addition, each of the two clusters seems to be made up of a number of smaller clusters. Similarly, in case 4, where the results are plotted in Figure 4-8, the clusters are rather spread out, each consisting of smaller clusters. Both cases 3 and 4 suggest that the service levels of both part types are not simply

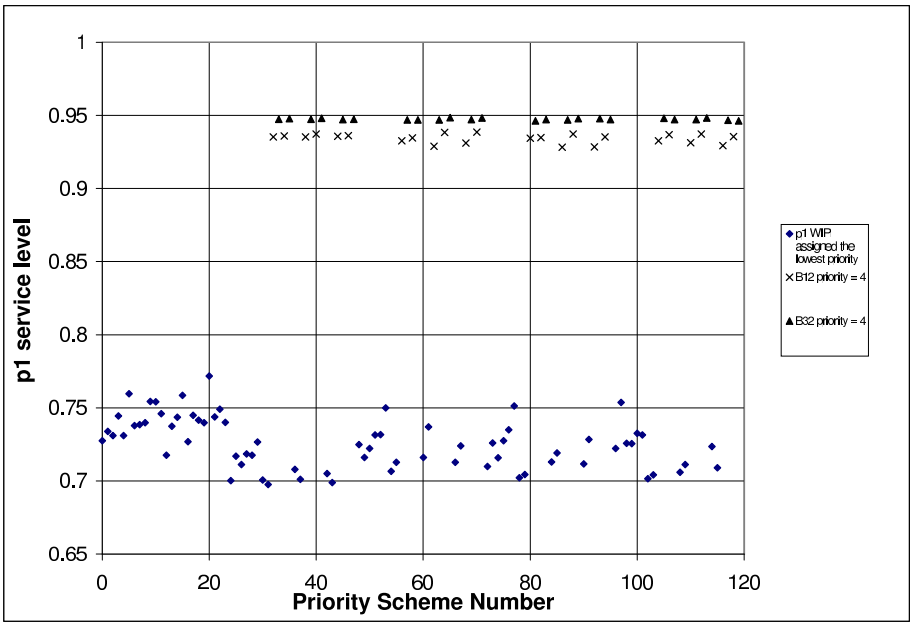

Figure 4-5: p1 service levels for Case 2

dependent on the assignment of the lowest buffer priority. The designation of other buffer priority values also has strong influence on the performance.

The appearance of the smaller clusters requires a more detailed breakdown in the analysis of the priority schemes. Figure 4-9 shows the service levels of p1 corresponding to each priority scheme for case 3. The service levels are plotted according to the assignment of the lowest priority scheme to each of the five buffers in front of WS 1.

Contrary to the observations made earlier in cases 1 and 2, priority given to p1 parts that return for more advanced process steps at WS 1 improves the service level of p1. In Figure 4-9, when B_{41} is assigned the lowest priority, p1 service levels are the lowest. On the other hand, when p1 raw material buffer B_{01} is assigned the lowest priority, p1 service level improves by more than 14%.

Another difference from cases 1 and 2 is that p1 service level does not seem to depend on whether B_{02} or B_{22} is assigned the lowest priority. As long as the lowest priority is assigned to either one of the p2 buffers, p1 service levels will do equally well.

Similarly for p2, priority given to B_{22} over B_{02} leads to a better service level.

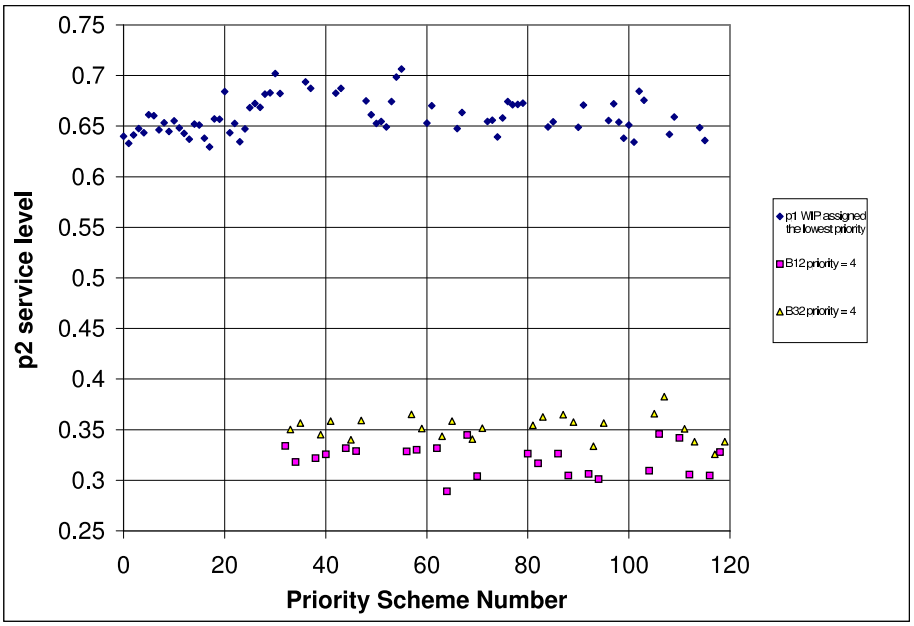

Figure 4-6: p2 service levels for Case 2

Again, p2 service levels are equally good as long as the lowest priority is assigned to one of the p1 buffers. The service levels of p2 are plotted in Figure **??**.

In case 4, where a different priority scheme is kept constant at WS 2, the same observations are drawn as in case 3. Priority given to parts returning for more advanced process steps lead to better service levels for both part types. The service levels of both part types corresponding to different priority schemes at WS 1 are shown individually in Figures 4-11 and **??**.

4.5 Summary

In the analysis of buffer priorities, the resultant performance measure is strongly characterized by its discreteness. This is evident in the distinctive clustering of service levels in all cases. In cases 1 and 2 where we fix the priority scheme at WS 1 while varying that of WS 2, results show that the service levels of both part types are strongly influenced by the designation of the lowest buffer priority but not so much by other priority values. In addition, priority given to WIP in its early visits to WS

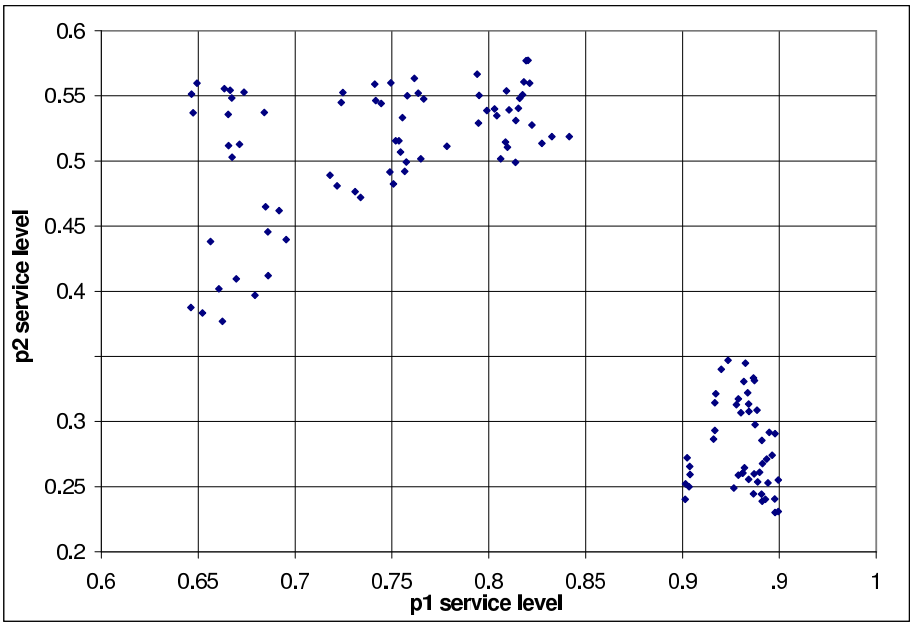

Figure 4-7: Service levels for Case 3, with all possible priority schemes at WS 1

2 improves the performance. In cases 3 and 4 where priority scheme at WS 2 is kept constant while varying that of WS 1, priority given to parts of later visits improve the service levels of both part types. These observations agree with those of Gzouli [7]. The performance is strongly influenced by the assignment of the lowest buffer priority. Moreover, the designation of other priority values for buffers at WS 1, the entrance into the system, also plays a significant role in influencing the service levels.

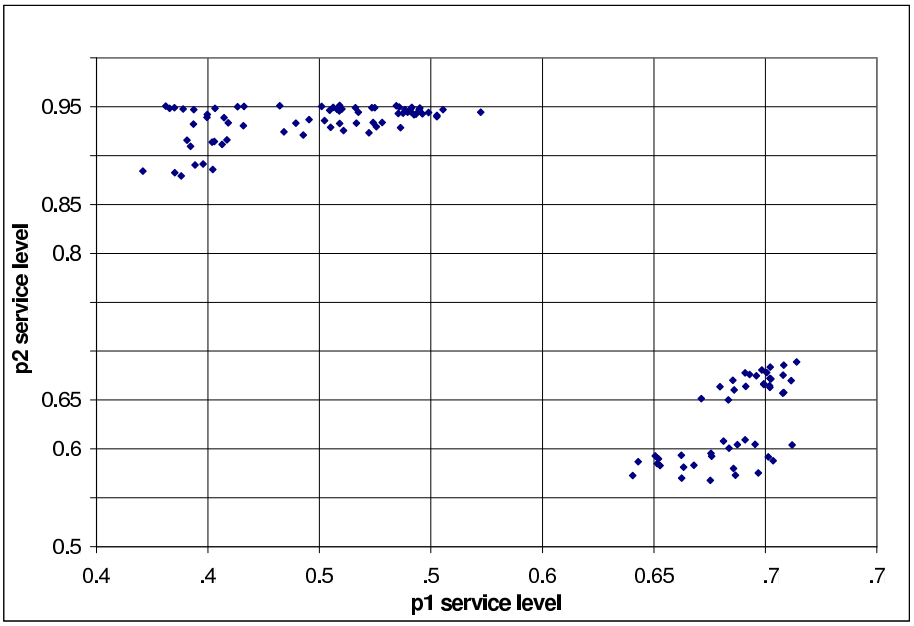

Figure 4-8: Service levels for Case 4, with all possible priority schemes at WS 1

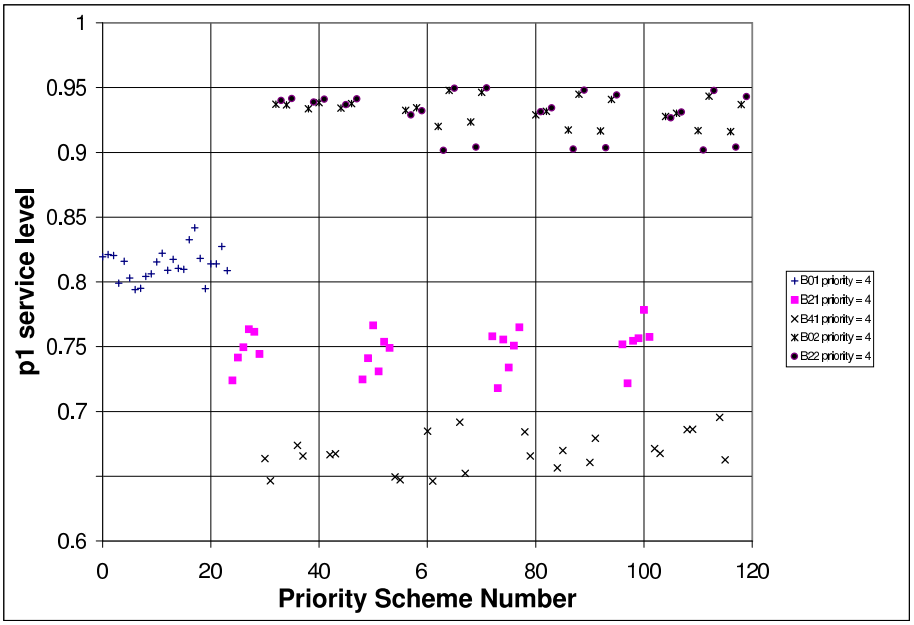

Figure 4-9: p1 service levels for Case 3

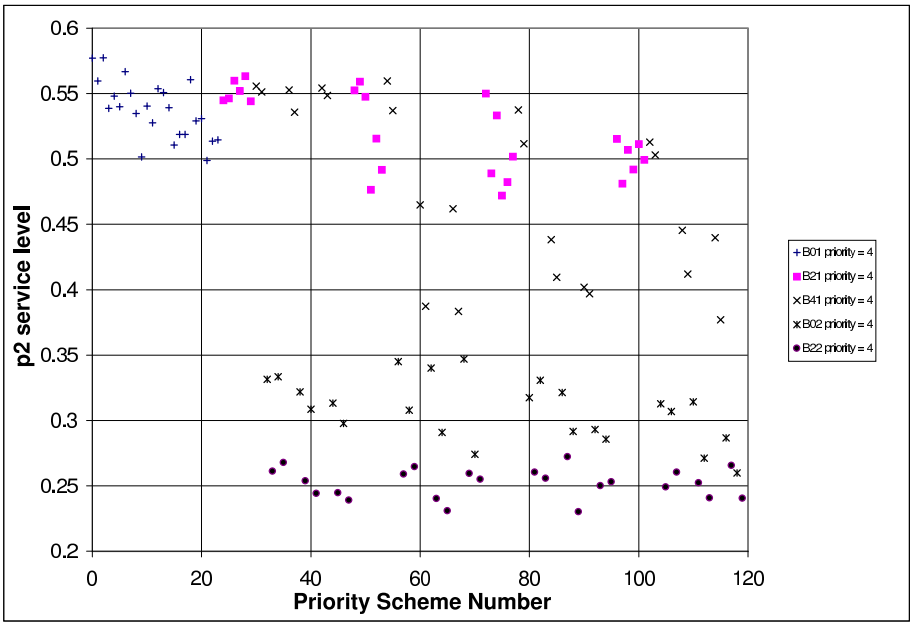

Figure 4-10: p2 service levels for Case 3

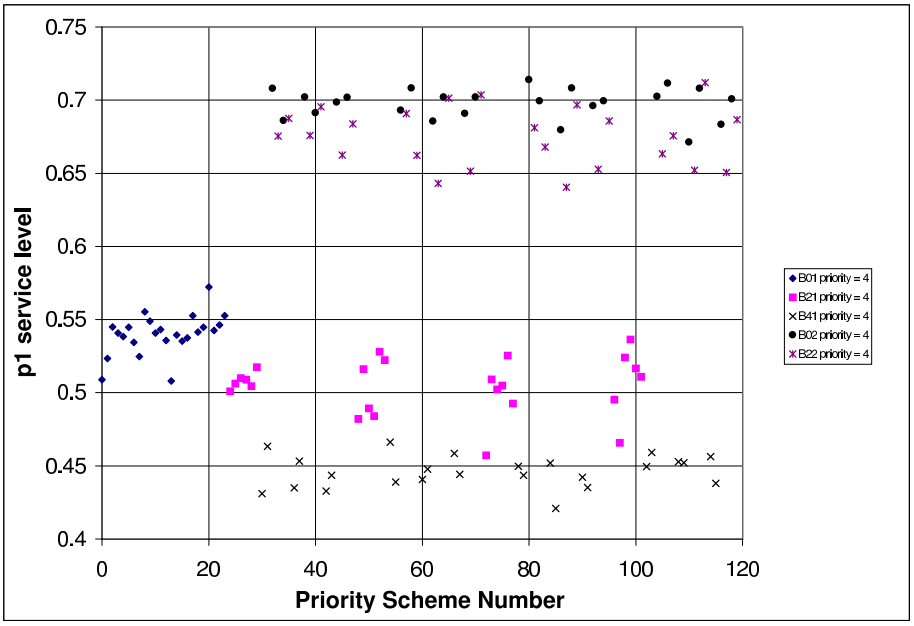

Figure 4-11: p1 service levels for Case 4

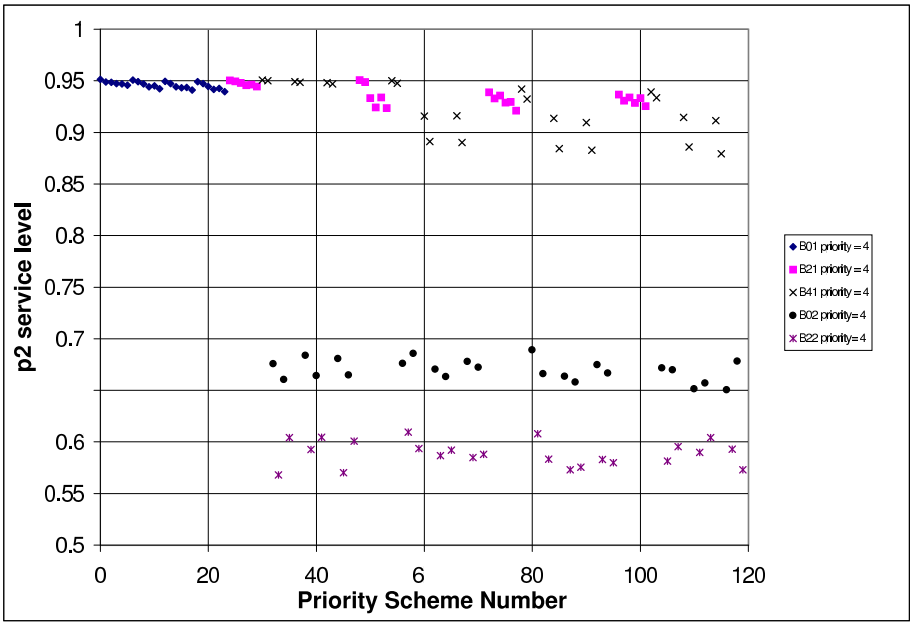

Figure 4-12: p2 service levels for Case 4

Chapter 5

Buffer Priority and Hedging Time Parameters

As seen from the results in Chapter 4, buffer priority schemes alone causes clustering of service levels. In each of the clusters, the service level of one part type usually exceeds that of the other. This phenomenon is desirable if one class of customer order is much more important than the others. However, the service level of the less important part type becomes so low that the corresponding customers may have to find another manufacturer. In some of the cases analyzed, the service level of the more important part type reaches as high as 94.78% while that of the less important part type is only 23%.

With constraints imposed by a finite capacity, this phenomenon comes as no surprise. Often much more capacity is needed to increase the service level when it is already high than when it is at a low value. Thus, it is likely that a slight decrease in the service level of the more important part type can result in a significant increase in that of the other. Consequently, there are incentives for companies to allocate slightly more capacity to the less important part type.

This chapter investigates the possibility of increasing the service level of the less important part type through the use of the hedging time parameter. Indirectly, minimal service level is used as the performance measure. Section 5.1 restates the concept of the hedging time parameter and how it is used in the scheduling logic of the

CPP. The experimental procedure is discussed in detail in the subsections of Section 5.2. Section 5.2.1 explains the selection of the control point in the toy fab. The priority schemes to be analyzed is discussed in Section 5.2.2, followed by an outline of the experimental procedure in Section 5.2.3. Experimental results are discussed in Section 5.3. The chapter concludes with a summary of the findings in Section 5.4.

5.1 The Readiness Logic

In the CPP scheduling logic, availability and readiness criteria are applied at the control points. A part from the highest priority buffer will be operated on next when it has satisfied both criteria. Thus, a full analysis of the scheduling logic cannot do away with any of these criteria.

However, when both conditions are used as specified by the scheduling logic, it will be hard to understand the role that the readiness criterion plays. Further, the complexity of the toy fab does not allow many choices about the appropriate buffer sizes and hedging times that will lead to good performance. These reasons call for the need to limit the CPP scheduling logic to include only the readiness condition.

As explained in Chapter 3, the hedging time parameter is used only in the readiness condition. A part from a buffer is considered ready when the sum of its hedging time and the current time is later than the due date for that part.

If the hedging time for a buffer is infinitely high, this implies that a part from that buffer always satisfies the readiness condition. On the other hand, a low hedging time means that a part will have to wait before it is to be processed. Whether hedging times should be high or low are strongly influenced by the choice of the buffer priority scheme and ultimately, the performance measure of interest.

5.2 Simulation Experiments

5.2.1 The Control Point Location

Since the readiness criterion is only applied at the control point, it is essential to first determine how many control points are needed and where they should be located in the toy fab model. With only two workstations in the model, it is not hard to make such decisions.

In general, it may be expensive if the CPP is executed in every workstation. This is due to amount of data that needs to be collected and analyzed, and to the effort of implementing the policy. Earlier simulation experiments for a single part-type re-entrant system indicate that the performance with one control point can be as good as that with two control points (Gzouli [7]). Therefore, only one control point in the system is chosen.

The control point should be located at the workstation where the CPP scheduling logic will result in the greatest performance. Wein [15] determines that a good release control of materials into the system results in far greater reduction in the mean throughput time than from parts sequencing at other workstations. Subsequently, Lu et al [13] and Kim et al [9] verify and base their simulation experiments on the findings of Wein. Moreover, Gzouli also determines that having a control point at the entrance into the toy fab results in a much higher service level and lower inventory than having the control point at workstation 2. Following these earlier results, workstation 1, the entrance into the toy fab, is chosen to be the control point.

5.2.2 Formulation of Experimental Cases

In each of the clusters of service levels shown in the results of Chapter 4, the service level of one part type far exceeds that of the other simply due to the designation of the lowest priority buffer. The purpose of the experiment is to determine how much improvement can be made in the lesser service level through the use of the hedging time parameter. The effects on WIP are also explored.

Several guidelines are used to determine the choices of buffer priority schemes to which the readiness logic is applied. First, they should consist of cases where buffer priority scheme is kept constant at the control point (WS1) while that at the non-control point (WS 2) is varied, and vice-versa. Consider Case 1 in Chapter 4 where the priority scheme at WS 1 is kept constant. A priority scheme corresponding to each of the service level clusters can be used to compare the resulting minimal service levels when the readiness logic is applied.

Instead of simply selecting a scheme corresponding to each cluster, a second guideline for the selection is imposed. The two schemes should differ only in the assignment of the two lowest priority buffers. Taking Case 1 in Chapter 4 as an example, if the second lowest priority is assigned to B_{11} and the lowest to B_{12} at WS 2, the second case used for comparison should have the two buffers reversed in their priorities. Under this guideline, it can be more conclusively determined whether the hedging time parameter alone can be used to offset the effect of the lowest priority assignment.

The buffer priority schemes used in the experiments are listed in Table I.

Buffer Priority Scheme	Workstation (WS)	Buffer Priority, in decreasing order of importance				
		0	1	2	3	4
Case 1	WS 1	B_{41}	B_{22}	B_{21}	B_{02}	B_{01}
	WS 2	B_{51}	B_{32}	B_{31}	B_{11}	B_{12}
Case 2	WS 1	B_{41}	B_{22}	B_{21}	B_{02}	B_{01}
	WS 2	B_{51}	B_{32}	B_{31}	B_{12}	B_{11}
Case 3	WS 1	B_{41}	B_{22}	B_{21}	B_{01}	B_{02}
	WS 2	B_{51}	B_{32}	B_{31}	B_{11}	B_{12}
Case 4	WS 1	B_{22}	B_{41}	B_{21}	B_{01}	B_{02}
	WS 2	B_{12}	B_{11}	B_{31}	B_{32}	B_{51}
Case 5	WS 1	B_{22}	B_{41}	B_{21}	B_{01}	B_{02}
	WS 2	B_{12}	B_{11}	B_{31}	B_{51}	B_{32}
Case 6	WS 1	B_{22}	B_{41}	B_{21}	B_{02}	B_{01}
	WS 2	B_{12}	B_{11}	B_{31}	B_{32}	B_{51}

Table I: Buffer Priority Schemes

These six buffer priority schemes can be separated into two major groups. Cases

1, 2 and 3 form the first group and the subsequent cases form the second. Members within each group differ only in the assignment of the two lowest buffer priorities, either at the control point or the non-control point.

5.2.3 Experimental Procedure and Guidelines

The buffer priority schemes are first analyzed using very large hedging times. This means that the system is operating on a Earliest-Due-Date policy with buffer priority. Subsequently, the hedging times are modified to improve the minimum service level.

One method of searching for good hedging times that improve the minimal service level is to simply make guesses of random positive numbers. However, the number of simulation runs will easily become too large to handle. In the toy fab with 5 buffers in front of the control point, ten guesses of hedging times corresponding to each buffer will lead to $10^5 = 100,000$ simulation runs in just one replication.

As a result, a set of guidelines used in the search for good hedging times are proposed and listed below. Explanations for each of the guideline are given in the ensuing paragraphs.

(a) The hedging times are varied one buffer at a time, starting with the raw material buffer and working downstream according to the process flow.

(b) When upstream hedging times are varied, downstream values are kept at infinity.

(c) More guesses are allowed for hedging times of upstream than downstream buffers.

(d) The measure that is used to determine good hedging times is minimal service level. The issue of whether the improvement is worth the percentage decrease in the service level of the other part type is ignored.

The reasoning behind the first two guidelines comes from the intuition that the flow rate of the WIP in the downstream process depends on the rate of release upstream. In fact, the readiness logic controlling upstream buffers can possibly nullify

the effects of the control downstream. This situation will arise when materials are forced to wait for a long time upstream such that they are likely to be late when they leave the system.

The third proposed guideline is based on the simulation results of Wein [15]. He concludes that controlling the release of materials into the system results in the greatest reduction in the mean throughput time. This observation alludes to the importance of upstream control. As a result, more guesses of good hedging times are allowed for buffers upstream in the process flow.

Finally, the fourth guideline is needed in the presence of many possible performance measures in a multi-product system. Different performance measures will lead to different definitions of good hedging times, and thus, different sets of hedging times will be chosen. It is likely that a set of good hedging times for one measure will be different from that of another measure.

About ten to twelve choices are made for the hedging time of each buffer, with more good guesses retained for the raw material buffers. During the trials to search for good guesses, three replications of each guess are made. On the average, the best five choices are kept for H_{01} and H_{02}, which are the raw material buffers, and four for the other buffers at the control point. Fifteen replications are made for the three final sets of guesses that result in the best minimal service level.

5.3 Experimental Results

The simulation results are presented in the following tables. Table II presents the three sets of hedging times for each priority scheme that lead to the best minimal service level. The resulting service levels and the WIP are contained in Table IV. For comparison purposes, the results of the performance measures with infinite hedging times are shown in Table III.

As seen from the results in Table IV, in the attempt to raise the minimal service level by applying the readiness logic to the part type with a higher service level, this part type suffers a significant decrease in service level. The minimal service level,

Buffer Priority Scheme		Hedging times for buffers at WS 1				
		H_{01}	H_{21}	H_{41}	H_{02}	H_{22}
Case 1	A	290	90	100	Infinite	Infinite
	B	280	125	100	Infinite	Infinite
	C	270	160	80	Infinite	Infinite
Case 2	A	Infinite	Infinite	Infinite	240	80
	B	Infinite	Infinite	Infinite	190	50
	C	Infinite	Infinite	Infinite	140	60
Case 3	A	300	190	90	Infinite	Infinite
	B	260	180	85	Infinite	Infinite
	C	240	170	85	Infinite	Infinite
Case 4	A	360	330	Infinite	280	Infinite
	B	360	330	Infinite	240	Infinite
	C	340	330	Infinite	280	Infinite
Case 5	A	280	90	85	Infinite	Infinite
	B	260	130	95	Infinite	Infinite
	C	260	130	45	Infinite	Infinite
Case 6	A	Infinite	Infinite	Infinite	220	60
	B	Infinite	Infinite	Infinite	120	120
	C	Infinite	Infinite	Infinite	120	90

Table II: The top 3 sets of hedging times for each case

however, has little or no improvement at all. Several reasons may account for this. First, these hedging times are not likely to be optimal due to the insufficient number of choices for good hedging times for each buffer. Secondly, the proposed guidelines for determining good hedging times may need improvements. These two reasons are significant but hard to overcome without further development of analytical work. It will be increasingly difficult to estimate good hedging times as a system becomes more complex. More simulation experiments need to be carried out in order to develop a better intuition of the hedging time parameter.

A more likely reason is due to the specification of the readiness logic. In these experiments with minimal service level as the performance measure, the hedging time parameter is used as a control to limit parts of the better performance part type from going through the system. Thus, the corresponding hedging times are decreased so that materials of the other part type can be processed. However, once a part in the system is late in terms of its due date, it will be ready at all points downstream in the

Buffer Priority Scheme	p1 service level	p1 WIP	p2 service level	p2 WIP	min service level	total WIP
Case 1	0.831	39.71	0.550	37.39	0.550	77.10
Case 2	0.562	44.56	0.903	27.56	0.562	72.11
Case 3	0.936	48.49	0.308	35.26	0.308	83.74
Case 4	0.703	58.47	0.645	19.37	0.645	77.83
Case 5	0.947	49.96	0.362	38.31	0.362	88.27
Case 6	0.583	51.32	0.946	29.26	0.583	80.57
Note: Results are obtained assuming infinite hedging times						

Table III: Performance under infinite hedging times

system. As a result, hedging time is not able to limit parts from going downstream.

Nonetheless, the readiness logic does lead to a significant reduction in the inventory in the system. On average, more than 10% reduction in total inventory is observed across all the cases. Most of the reduction comes from the WIP of the part type to which the readiness control logic is applied. It seems that the hedging time parameter is indeed effective in limiting the flow of materials through the system.

Depending on the priority scheme analyzed, the hedging times of the part type with lower service level are kept very large. Consequently, materials of these part types are always made ready when the readiness logic is applied. This is necessary since the static priority scheme has resulted in more capacity being allocated to the other part type. Any further control to limit the flow of materials of the lower service level part type will only lead to a worse service level. In Table II, these high hedging times are labeled "Infinite".

5.4 Summary

This chapter explores the effectiveness of the hedging time parameter, under a predetermined static buffer priority scheme, to improve the minimal service level of a two

Buffer Priority Scheme		p1 service level	p1 WIP	p2 service level	p2 WIP	min service level (SV)	total WIP	% Improvement	
								in min SV	in total WIP
Case 1	A	0.592	24.98	0.542	36.26	0.542	61.24	-1.511	20.570
	B	0.631	24.43	0.542	37.26	0.542	61.69	-1.375	19.987
	C	0.656	24.28	0.547	36.83	0.547	61.12	-0.552	20.730
Case 2	A	0.592	44.30	0.866	22.48	0.592	66.78	5.341	7.397
	B	0.579	44.30	0.809	19.16	0.579	63.46	3.073	12.005
	C	0.607	44.45	0.797	14.97	0.607	59.41	7.984	17.611
Case 3	A	0.892	34.79	0.292	34.50	0.292	69.29	-5.058	17.254
	B	0.885	31.14	0.292	34.03	0.292	65.17	-5.266	22.184
	C	0.877	29.29	0.317	33.70	0.317	62.99	3.081	24.785
Case 4	A	0.644	51.80	0.616	17.16	0.616	68.95	-4.460	11.407
	B	0.639	51.98	0.586	15.68	0.584	67.67	-9.422	13.062
	C	0.635	50.13	0.613	17.23	0.602	67.35	-6.630	13.467
Case 5	A	0.817	34.86	0.341	37.28	0.341	72.15	-5.707	18.268
	B	0.875	32.48	0.331	38.37	0.331	70.85	-8.540	19.736
	C	0.697	33.26	0.337	37.08	0.337	70.34	-6.885	20.312
Case 6	A	0.578	51.54	0.881	23.00	0.578	74.54	-0.862	7.484
	B	0.553	50.84	0.919	13.94	0.553	64.79	-5.193	19.593
	C	0.574	49.89	0.909	13.87	0.574	63.75	-1.472	20.875

Table IV: The results of readiness logic control

part-type system. In these simulation experiments, the best hedging times do little to improve the service level. However, they do result in significant reduction in the total work-in-progress. The failure of the readiness logic to control the flow of late parts is the main reason that limits the extent to which the hedging time parameter can offset the effects of a static buffer priority scheme.

Chapter 6

The Analysis of The Hedging Time Parameter

The simulation results in Chapter 5 show that the hedging time does little to improve the minimal service level. It is important then to understand the role the readiness control logic plays at the control point.

This chapter presents simulation results in an attempt to understand hedging time at a greater depth. It contains experiments performed to understand the effects of the readiness logic controlling one buffer at a time, as well as an interactive control of all the buffers at the control point.

Section 6.1 provides an intuitive discussion of hedging time, and what good values mean. Simulation experiments to understand the effects of the readiness logic applied to one buffer at a time are explained in Section 6.2, followed by a discussion of the results in Section 6.3. Results from Chapter 5 are analyzed in greater depth to form the discussions of the effects of an interactive readiness logic control in Section 6.4. The chapter concludes with a summary.

6.1 Intuitive Discussion

The readiness logic functions like a floodgate control. It prevents an overflow of materials downstream. As a result, it has the ability to reduce inventory and possibly

free up resources in case more urgent or important materials arrive later at the control point. However, if the control is over conservative, not only are resources wasted, but service levels will be low also.

With the potential benefits but also high risks if the readiness logic is not applied properly, it becomes very important to determine good hedging times. It is believed that good hedging times should be conservative estimates of the remaining cycle time for parts in a particular buffer (Gershwin, 1999). One main reason is that if the hedging time is less than the remaining cycle time, a part is likely to be late. Variability in the system is likely to cause delay, and thus, a conservative estimate is necessary.

6.2 Experimental Procedure on Singular Readiness Logic Control

Simulation experiments are formulated to test the intuition as well as to shed light on the effects of the readiness logic. To achieve this, several steps are needed. First, it will be necessary to investigate the effects of controlling each individual buffer using the readiness logic.

The experiments are carried out assuming infinite buffer sizes. The priority scheme chosen is that of Case 3 in Chapter 5. One control point is chosen and is located at WS 1, the entrance into the system. For reasons discussed in Section 5.2.3, 26 values of H_{01} are studied, compared to 19 and 15 values for H_{21} and H_{41} respectively. The hedging times for the p2 buffers are kept at infinity since p2 already has very low service level.

On average, ten replications are made for each hedging time value to achieve the statistical requirement. The WIP and service level of each part type, as well as the resulting expected remaining cycle time are plotted and shown in the next section. All possible performance measures are shown in order to understand the effects of the readiness logic.

6.3 Results

6.3.1 Service Levels

The service levels of p1 and p2 corresponding to various values of H_{01}, H_{21} and H_{41} are plotted in Figures 6-1, 6-2 and 6-3. These results are separated into two regions: one with the hedging times being greater than the expected remaining cycle time (E(RCT)) of that buffer while the other corresponds to values less than E(RCT).

In each of the plots, a wide range of hedging times is covered. Large hedging times such that a part from the buffer is always ready are included for comparison purposes. Taking the Figure of H_{01} as an example, with the average customer lead time of 330 time units for p1 parts, H_{01} values of 330 and above represent cases where parts are always considered ready.

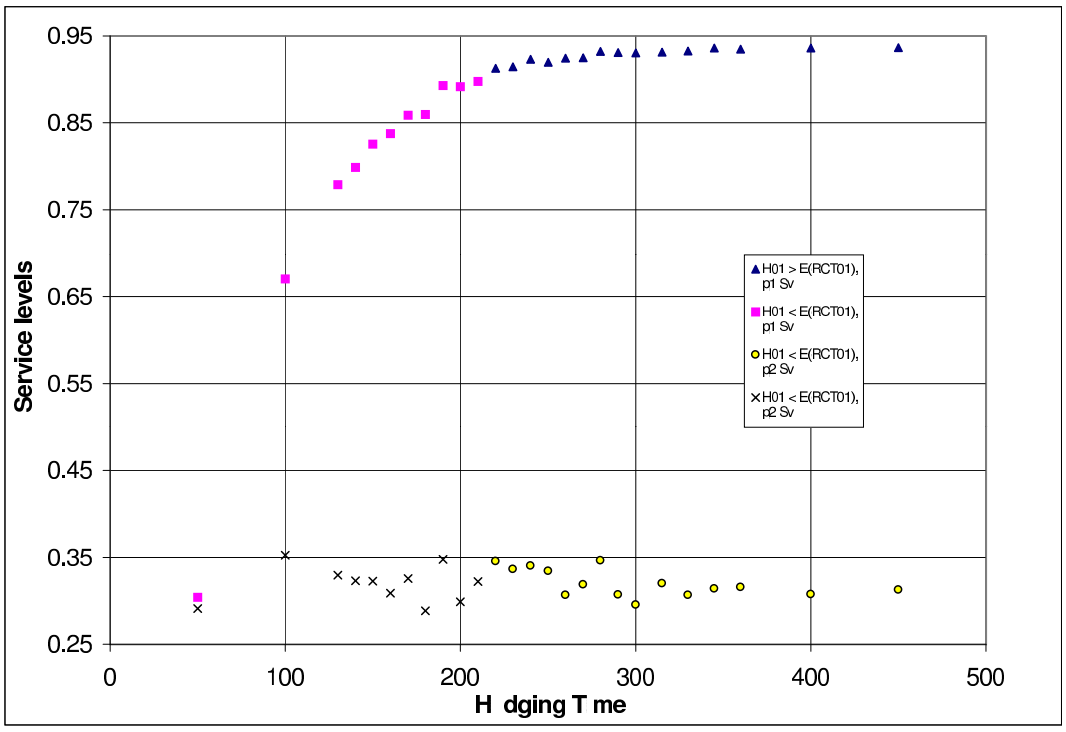

Figure 6-1: p1 and p2 service levels corresponding to H_{01}

All of the plots share the same characteristic: no hedging times for any of the p1 buffers are able to significantly improve the service level of p2. Even when H_{01} is

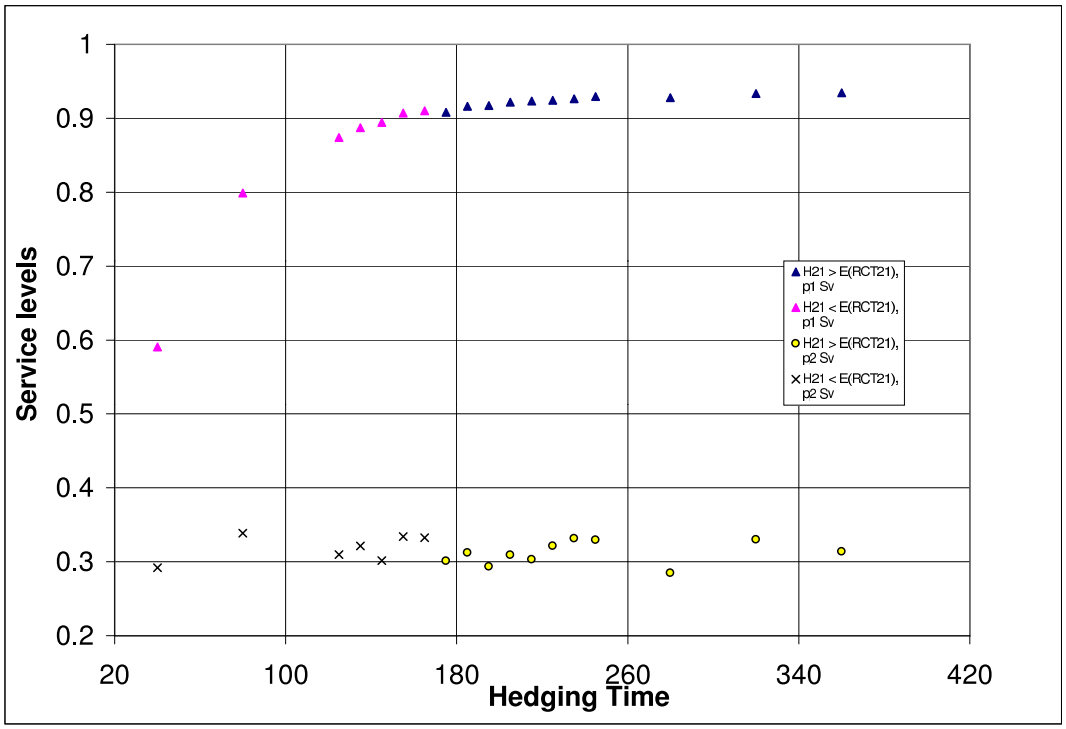

Figure 6-2: p1 and p2 service levels corresponding to H_{21}

small, implying that p1 parts have to wait for a long time before being released into the system, p2 parts still suffer from the effects of the static priority scheme. In fact when the hedging times become too small such that all p1 parts become late, both p1 and p2 service levels decrease, as shown by $H_{01} = 50$ in Figure 6-1. This phenomenon may again be due to the failure of the readiness logic to limit the flow of late parts.

All three figures also shows that when the hedging time value for a p1 buffer falls below the expected remaining cycle time of that buffer, p1 service level decreases steeply. This is especially true for the raw material buffer. A manufacturer may not want to have such low hedging times since the decrease in service level may not be worth the very small service level improvement of the other part type.

6.3.2 Expected Remaining Cycle Time and WIP

Even though the minimal service level does not improve significantly using the hedging time parameter, the results in Chapter 5 do show that there is significant reduction

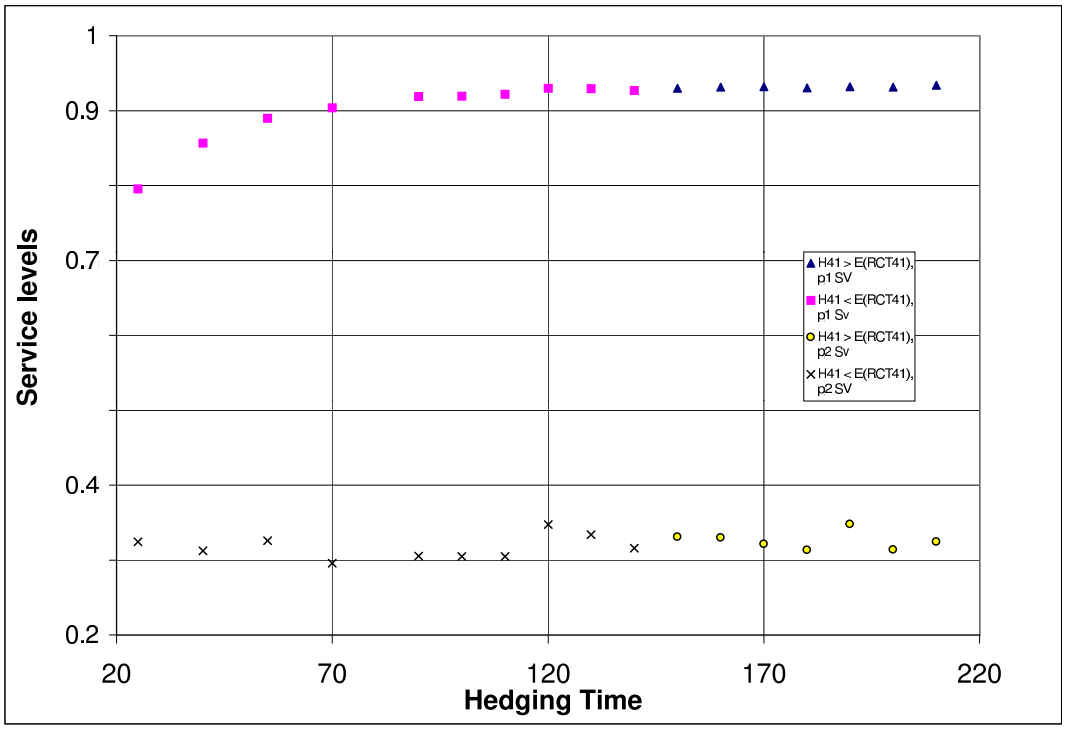

Figure 6-3: p1 and p2 service levels corresponding to H_{41}

in the WIP. Figures 6-4, 6-5 and 6-6 show the resulting WIP in the system when the readiness logic is applied to each p1 buffer at the control point.

These three figures show that only H_{01} causes a reduction in the p1 WIP. The hedging times at the other p1 buffers do not result in any increase or decrease in p1 WIP. The main reason is due to the definition of the WIP used that includes the entire inventory in the system, except those at the raw material buffers. The readiness logic, when applied to downstream buffers, only serves to move the bulk of materials further upstream within the system. This observation may help explain which hedging time parameter causes most of the WIP reduction in the results of Chapter 5.

In addition, there is neither a significant decrease nor increase in the p2 WIP when the readiness logic is applied to the p1 buffers. The reduction in the total WIP of the system comes only from the decrease in the p1 WIP.

A plot of the expected remaining cycle time as a function of the hedging times

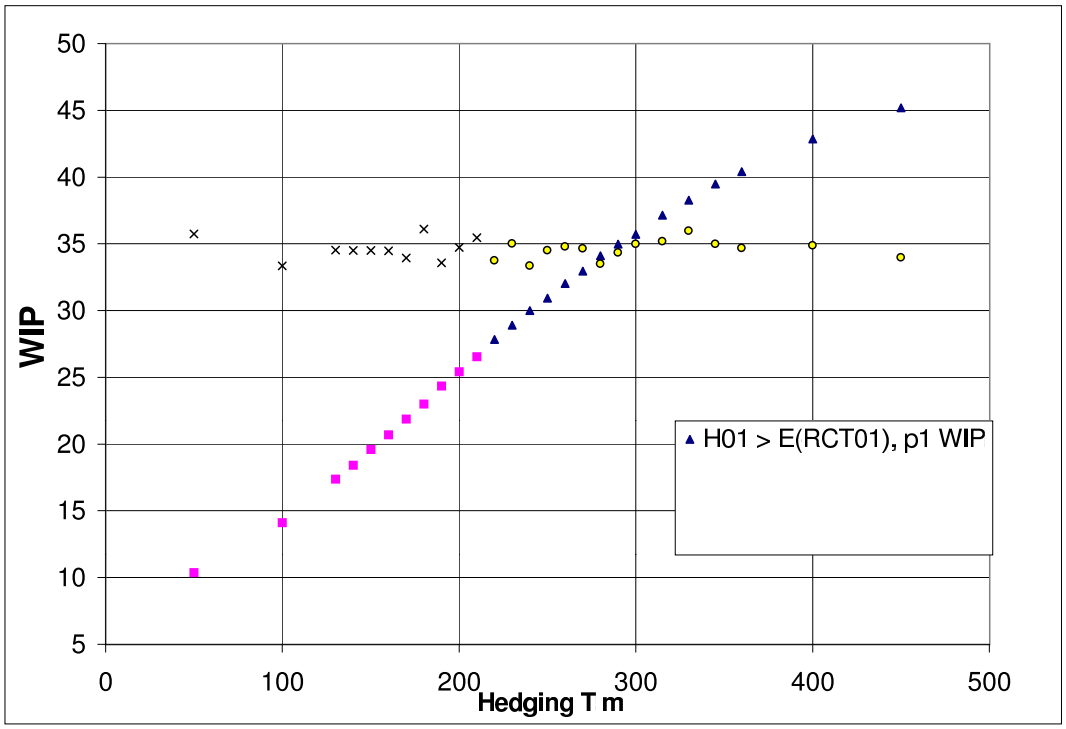
50
45
40
35
30
25
20
15
10
5
WIP
0
100
200
300
400
500
Hedging T m
H01 > E(RCT01), p1 WIP

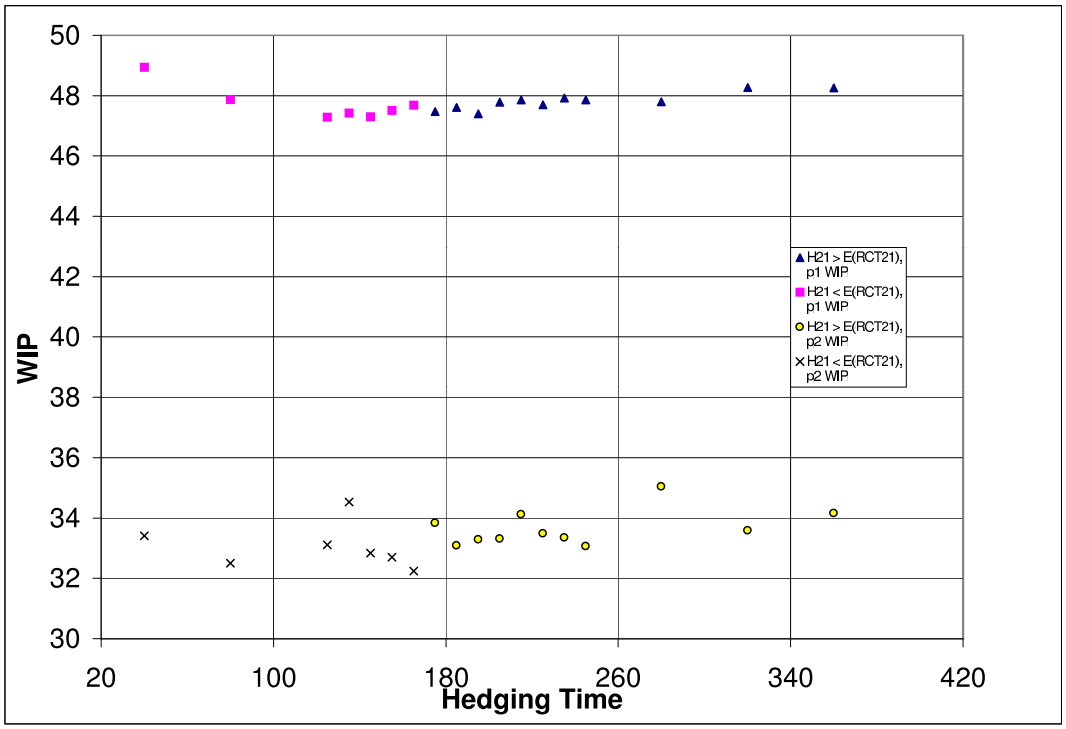

Figure 6-5: Relationship between H_{21} and WIP

time for p1 parts is 84.6 time units. With an average customer lead time of 330 time units, most of the p1 parts will spend large amount of time waiting in the finished goods buffer before they are due. Lowering the hedging times merely distributes part of this waiting time upstream. This is only true for a certain range of hedging times that are greater than the resulting average remaining cycle time.

The remaining cycle times for all the p2 buffers are not affected even when the values of H_{01}, H_{21} and H_{41} are lowered. This is another reason why p2 service levels do not improve.

6.4 The Effects of Interactive Readiness Logic Control

So far, the study of the hedging time parameter has been limited to applying the readiness logic to one buffer at a time. The results suggest that good hedging times

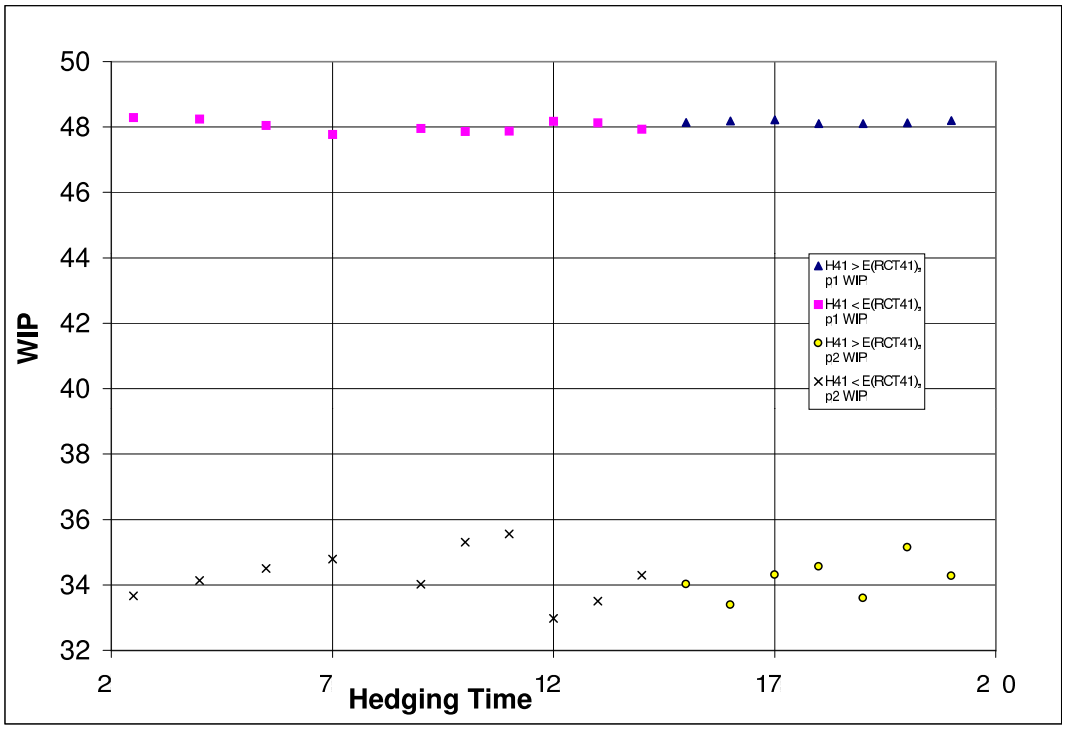

Figure 6-6: Relationship between H_{41} and WIP

should be above the expected remaining processing time since the resulting decrease in the inventory is not worth the sharp drop in service level. A more detailed analysis of the best hedging times used and the resulting expected remaining cycle time from the cases in Chapter 5 will determine what the good hedging times are when the readiness logic is applied to multiple buffers at a time.

Figure **??** is a plot of the three best hedging times and their resulting expected remaining cycle time for all six cases of Chapter 5. Each point is grouped by the respective buffer to which the readiness logic is applied. The resulting expected remaining cycle time that corresponds to infinite hedging times are not plotted. A 45-degree line is plotted to distinguish the hedging times that are below and above the resulting expected remaining cycle times.

It is interesting to note that quite a large number of points are above the 45-degree line. This means that some good hedging times are less than the resulting expected remaining cycle time. Most of these points correspond to the hedging times of H_{01}

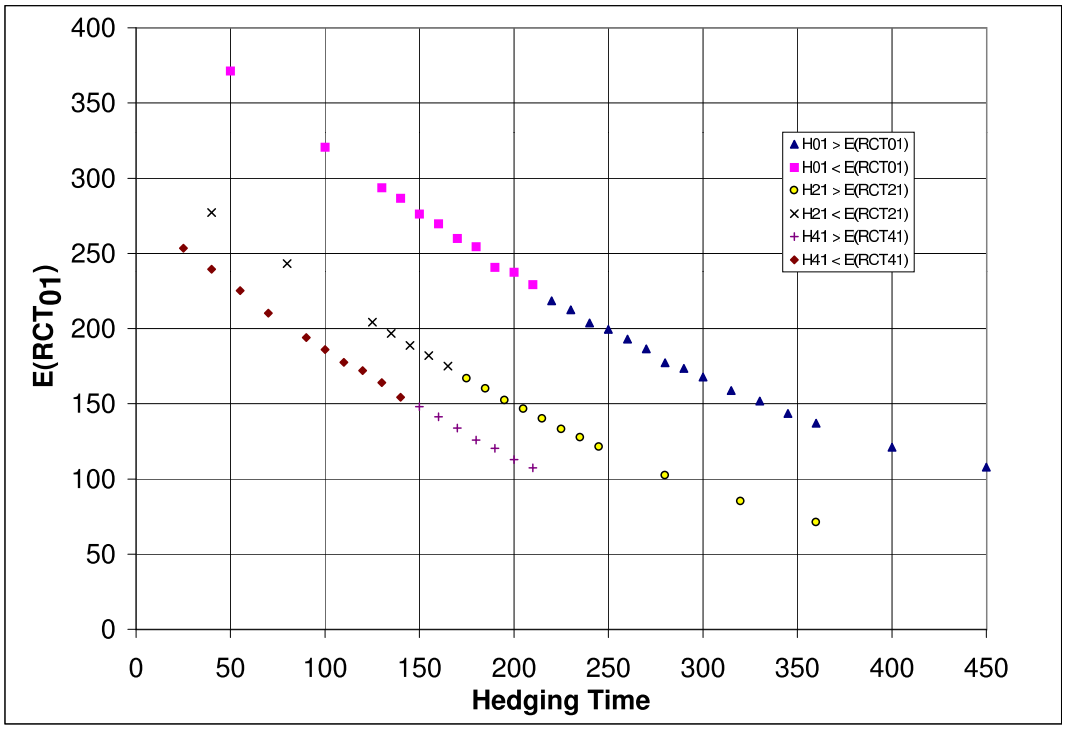

Figure 6-7: Relationship between hedging times and the expected remaining cycle time

and H_{02}, the raw material buffers for p1 and p2 respectively. On the other hand, the hedging times of H_{41} and H_{22} tend to stay close to or below the 45-degree line.

The relationship between hedging time and E(RCT) are presented in a different form in Table I. In each of the three best sets of hedging times for each priority scheme, an "N" is assigned to hedging times that are below E(RCT) and the reverse is assigned the letter "Y". Buffers with infinite hedging times are labeled "Infinite".

In almost all the best sets of hedging times for each priority scheme, there is a mixture of values higher and lower than the resulting expected remaining cycle time in each set. These results do seem to be contrary to the intuition of what good hedging times should be, as discussed in the Section 5.1.

Several reasons may explain why a significant proportion of good hedging times are less than the resulting expected remaining cycle time. It is important to recall that in the process of determining good hedging times, the issue of whether the improvement

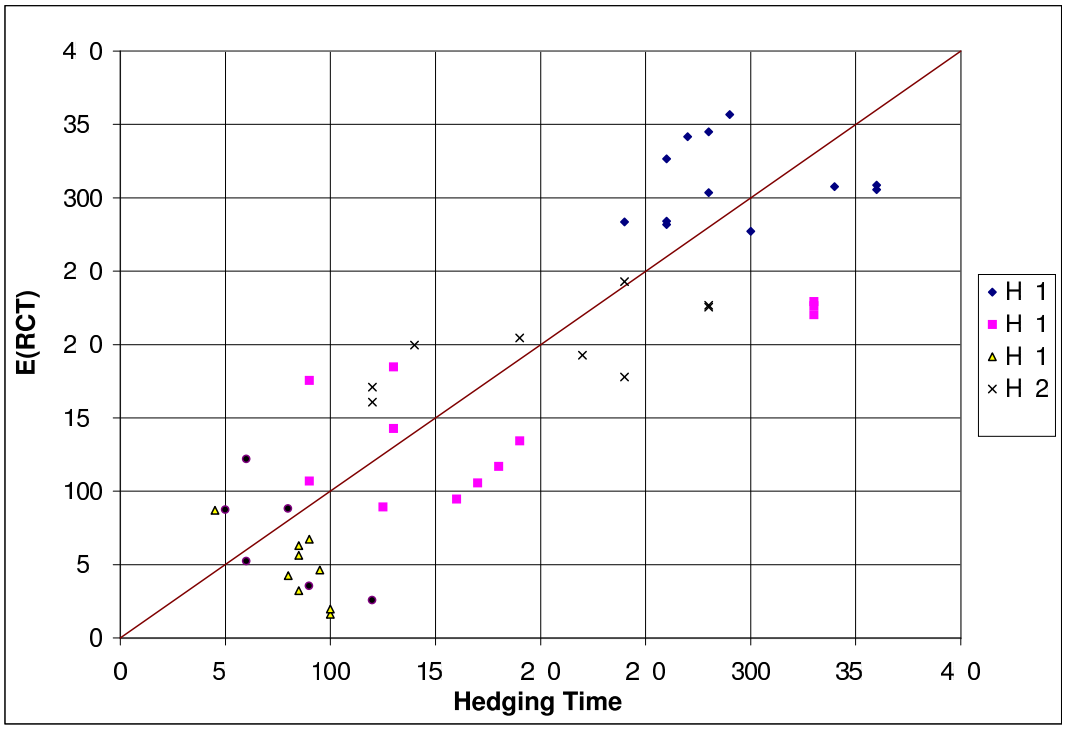

Figure 6-8: Relationship between hedging times and expected remaining cycle times for all 6 cases in Chapter 5

in the minimal service level is worth the percentage decrease in the service level of the other part type is ignored. Even though some hedging times are very low and cause dramatic decrease in the better service level, they are still considered good if they happen to raise the minimal service level. It is unlikely that these values will actually be used in the implementation of the CPP.

Another reason is due to the variability in the results from the simulation runs. As shown in the results of singular and interactive hedging time control, the improvement in the minimal service level through the modification of the hedging times is not significant. The minimal service level may not be the right measure to use in the determination of good hedging times.

Table II shows the expected remaining cycle time for all the buffers at the control point corresponding to the three best sets of hedging times. The expected remaining cycle times for these buffers before applying the readiness logic are also shown for

Buffer Priority Scheme		Comparison between hedging time and E(RCT)				
		B_{01}	B_{21}	B_{41}	B_{02}	B_{22}
Case 1	A	N	N	Y	Infinite	Infinite
	B	N	Y	Y	Infinite	Infinite
	C	N	Y	Y	Infinite	Infinite
Case 2	A	Infinite	Infinite	Infinite	Y	N
	B	Infinite	Infinite	Infinite	N	N
	C	Infinite	Infinite	Infinite	N	Y
Case 3	A	Y	Y	Y	Infinite	Infinite
	B	N	Y	Y	Infinite	Infinite
	C	N	Y	Y	Infinite	Infinite
Case 4	A	Y	Y	Infinite	Y	Infinite
	B	Y	Y	Infinite	N	Infinite
	C	Y	Y	Infinite	Y	Infinite
Case 5	A	N	N	Y	Infinite	Infinite
	B	N	N	Y	Infinite	Infinite
	C	N	N	N	Infinite	Infinite
Case 6	A	Infinite	Infinite	Infinite	Y	N
	B	Infinite	Infinite	Infinite	N	Y
	C	Infinite	Infinite	Infinite	N	Y
Note: 'Y' = Hedging Time above E(RCT), 'N' = Hedging Time Below E(RCT)						

Table I: Relationship between hedging times and E(RCT) for each set of good hedging times

comparison purposes. Indeed, the readiness logic does not lead to a reduction in the remaining cycle time for the buffers of the lower service level part type.

6.5 Summary

Even though the best set of hedging times show otherwise, it is not conclusively determined if good hedging times can fall below the resulting expected remaining cycle time. This is because the hedging time parameter does not significantly improve the minimal service level. In fact, it is observed that when the hedging times used to limit the flow of a part type are lower than the resulting expected remaining cycle time, the service level of this part type drops sharply. Indeed, under the influence of a static priority scheme, it may not be appropriate to use minimal service level as a performance measure in the determination of good hedging times. The readiness logic does however limit the total WIP in the system when applied to the release.

Buffer Priority Scheme		Expected Remaining Cycle times				
		B_{01}	B_{21}	B_{41}	B_{02}	B_{22}
Case 1	A	356.8	106.9	16.1	228.9	31.4
	B	345.0	89.3	19.8	233.9	32.2
	C	341.6	94.7	42.4	230.2	33.7
	Infinite HT	156.4	37.9	15.7	227.7	31.9
Case 2	A	307.8	58.0	16.5	177.9	88.0
	B	314.0	59.5	16.5	204.6	87.4
	C	293.7	56.0	16.4	199.7	52.3
	Infinite HT	313.1	55.2	16.3	67.6	17.7
Case 3	A	277.2	134.3	67.2	388.2	37.4
	B	281.8	116.8	63.1	374.4	37.4
	C	283.5	105.7	56.3	363.6	36.5
	Infinite HT	84.6	37.6	15.7	383.8	33.5
Case 4	A	305.5	226.5	202.9	226.7	28.1
	B	308.7	229.3	205.7	242.9	28.6
	C	307.7	220.2	201.6	225.5	28.1
	Infinite HT	251.9	226.0	207.6	206.3	27.3
Case 5	A	303.5	175.5	32.1	388.2	215.7
	B	284.1	142.9	46.2	419.6	224.6
	C	326.4	184.7	86.9	404.3	214.8
	Infinite HT	74.0	48.4	30.1	381.4	220.1
Case 6	A	313.8	214.7	195.2	192.8	121.8
	B	328.6	216.6	198.4	160.7	25.5
	C	311.4	207.4	188.7	170.9	35.2
	Infinite HT	315.9	212.2	194.1	47.0	24.2

Table II: The resulting expected remaining cycle times

Chapter 7

Simulation Experiments for LSI Logic Fab

The analysis of the toy fab model has provided insights into the effects of static priority scheme as well as its interaction with the hedging time parameter. These insights offer important guidelines in the study of the buffer size parameter. Moreover, they also enable us to study the performance of the CPP in a realistic model of a fabrication facility. As such, simulation experiments no longer assume infinite buffer size.

The facility under study is the LSI Logic fabrication plant (LSIL fab), a large-scale semiconductor manufacturing facility located in Portland, Oregon. LSIL is a supplier of custom performance semiconductors. It is currently using AutoSched AP, a simulation program provided by AutoSimulations Incorporated, for decisions related to production scheduling.

In addition to the discussion of simulation results, this chapter also provides an explanation of simulation modeling. Section 7.1 offers a background to semiconductor manufacturing. Some of the important modeling issues pertaining to the LSIL fab are raised in Section 7.2. The parameters and other statistics of the wafer fab model are explained in Section 7.3. Section 7.4 discusses the results of the toy fab under the CPP scheduling policy. It is broken down according to the policy parameters used. The chapter concludes with a review of the results in Section 7.5.

7.1 Background

Semiconductor manufacturing is rather complex not simply because of the re-entrant flow characteristics but also the many process constraints that are experienced in the wafer fabrication process. Existing literature provides a good background on the wafer fabrication process, as well as the production scheduling issues in a fab. Among other resources, readers are encouraged to refer to Atherton [12] and for explanations of the semiconductor manufacturing processes, and Wein [15] and Lu et al [13] regarding simulation experiments to compare production scheduling policies in a fab.

7.2 Fab Modeling Issues

Many challenges arise in the modeling of a full-scale fab. This section highlights some of the important issues and the approximations made to form a feasible simulation model. Discussion of these issues begin with the general demand and product information through detailed process modeling.

7.2.1 Demand and Product Information

In the LSIL fab, each customer order is unique. The orders differ in terms of the specifications such as reticle make-up, number of metal layers, wire width and the quality level of the wafers. As a result, there are multiple process flows in the system. Currently, there are more than 74 different wafer types produced in the fab.

In addition to meeting customer orders, LSIL is also exploring new technologies through the fabrication of more complex wafers. These wafers are collectively termed "proto". Even though it does not satisfy current demand, a proto lot has the highest priority compared to all other lots in the system.

Proto lots are not included in the simulation model. This is due mainly to their high priority and the subsequent influence on the fairness of policy comparison. Only wafer types L29_3Rocket_G10 (3Rocket) and L29_4Rocket_G11 (4Rocket) are included in the model. They make up the two highest production volumes in the fab. Since

they differ in terms of both the number of metal layers and wire width, the lots of these wafer types go through different process flows in the fab.

In LSIL, certain lots will become "hot" during the fabrication process due to customers expediting their orders. These lots will then have higher priority over other non-proto lots. This phenomenon is also ignored in the simulation model.

7.2.2 Material Release and Scheduling Policies

New lots are released into the system according to the takt time, the average time interval between the arrivals of two consecutive customer orders. The takt time is computed based on the demand for the month. This type of release is considered periodic and is used in the simulation model.

Once the lots are released into the system, the lot-sequencing rule used is based on the critical ratio. This chapter explores the performance of the CPP. The following chapter presents the performance of the fab under Critical Ratio and Earliest-Due-Date scheduling policies.

7.2.3 Fabrication Process

In the fab, setup is only required for the furnace operations. Due to the periodic release policy, almost no two identical lots arrive at the furnace operations at the same time. As a result, setup is always required. Further, the setup time is not sequence dependent. To simplify the analysis, setup time is included as part of the processing time in the fab model.

Another phenomenon that is ignored is batching. This is because none of the scheduling policies studied has a well-defined policy on batch operations. Modeling the batching operations will only distort the fairness of policy comparison.

In spite of this, the multi-capacity characteristic of a batch machine is captured in the model. It is done through the addition of virtual machines to match the capacity of the batch machine.

There are certain operations that are performed only on a fraction of each customer

order. In light of the purpose of the analysis, these process characteristics will not have much bearing on the conclusions and thus, they are not included in the study.

7.2.4 Multiple Failure Modes and Preventive Maintenance

Many workstations in the etching operation have multiple failure modes. These arise due to the different failure and repair probability distributions for different machine parts.

In the simulation model, the multiple failure and repair modes belonging to the same machine are modeled together. The resultant probability distributions for both failure and repair are assumed exponential. The resultant means of each distribution are computed using the availability theory as described in Barlow and Proschan [1].

From the data given, the parameters for all the failure probability distributions are in terms of calendar time. The simulation model, on the other hand, assumes an operation-dependent failure. Given information on machine utilization rate, the conversion could be easily done. However, such information is not readily available and cannot be accurately determined based on the data given. The parameters of the failure distributions are nonetheless used in the model. Readers have to be mindful of the disagreement between the data given and our model assumption.

The last issue to highlight is the modeling of preventive maintenance (PM). PM takes places in multiple modes for many different machines. There are daily, weekly, monthly, quarterly as well as semi-annual PM, and the operations carried out for different time intervals are not the same. During the scheduled time, not all the machines in the workstation undergo PM simultaneously.

In the simulation model, virtual workstations are introduced to represent the different PM modes. These workstations have the same number of identical machines as those of the real workstations. There are two important differences: these virtual workstations have zero processing times and zero buffer sizes. Their failure and repair probability distributions reflect the time intervals between consecutive PM and their durations respectively.

With the addition of the virtual workstations, the resulting manufacturing line

becomes too large. Consequently, only limited set of PM modes can be captured in the model. Two guidelines are developed for the selection of the appropriate PM modes. First, the PM interval cannot be too long relative to the simulation period. This guideline helps eliminate the modeling of PM modes with interval longer than a month. The second guideline comes from Gershwin [5]. He shows that an infrequent failure mode with a long repair period is more destructive to the production rate than a frequent failure mode with short repair time. Since we are interested in capturing the destructive effects of PM on production rate, PM mode that has long interval and repair time forms the criterion of the second guideline. As a result, only the weekly PM mode is modeled.

7.3 Simulation Model

The demand model is similar to that of the toy fab. A short customer arrival interval and lead time is used. Both of these random variables are assumed normal distributions with a coefficient of variation of 0.5 each. On average, the customer lead time is set to be 10 times the processing time of a lot. This factor is henceforth called the customer lead time factor. Demand information and other important statistics of the model are listed in Table I.

LSIL Fab Model Statistics		4Rocket (Part Type 1)	3Rocket (Part Type 2)
Average Takt time, t (hrs)		2.74	3.95
Standard Deviation of order arrival interval. σ_t		1.37	1.975
Average Customer lead time, ℓ (hrs)		1747.5	1563.3
Standard Deviation of lead time. σ_ℓ		873.75	781.65
Processing Steps	(without virtual workstations)	254	218
	(with virtual workstations)	385	331
Total Processing Time / lot		174.75 hrs	156.33 hrs
Longest-Processing-Time Workstations		Furnace_SiN2 (7.7 Hrs)	Furnace_SiN2 (7.7 Hrs)

Table I: LSIL Fab Model Statistics

A higher demand rate was chosen for 4Rocket lots. This is because LSIL is facing an increasing demand for this type of wafers, which are more complex than 3Rocket. Due to the importance of 4Rocket wafers, they are given priority in the simulation

runs under the CPP.

The length of the transient period was also determined through several pilot runs. A run of 7 years, assuming round the clock operations, with statistics collected over the last 2.5 years are sufficient to capture the steady state performance. Readers can refer to Law and Kelton [11] on how to determine the steady state behavior of the system.

The entrance into the system is the virtual workstation W_DNS_PRECLN that is introduced to capture the PM mode of workstation DNS_PRECLN. The definition of the WIP includes the entire inventory in the fab except the raw material buffers at this virtual workstation. Subsequent analysis with the CPP focuses on the scheduling logic applied at the entrance into the system.

A description of the equipment is given in Tables II and III. A process flow chart for 4Rocket lots is also provided in Figure 7-1. The names of the workstations in italic are the virtual workstations used to capture the PM mode of the preceding workstation. The processing time information is not listed here. It is given for each operation of each wafer type, and it is deterministic. Readers are referred to Appendix B for a detailed description of each operation. Also provided in the appendix are several diagrams describing 4Rocket process flow in detail.

7.4 Discussion of the CPP Results

This section discusses the results of simulation runs using the CPP. Both the hedging times and buffer sizes are first assumed infinite. Subsequent analyses do away these assumptions. The buffer priority scheme is held constant throughout the analysis. It gives a higher priority to lots returning for more advanced visits at all workstations. A higher priority is also given to 4Rocket lots among returning lots of both wafer types that have visited the workstation for the same number of times.

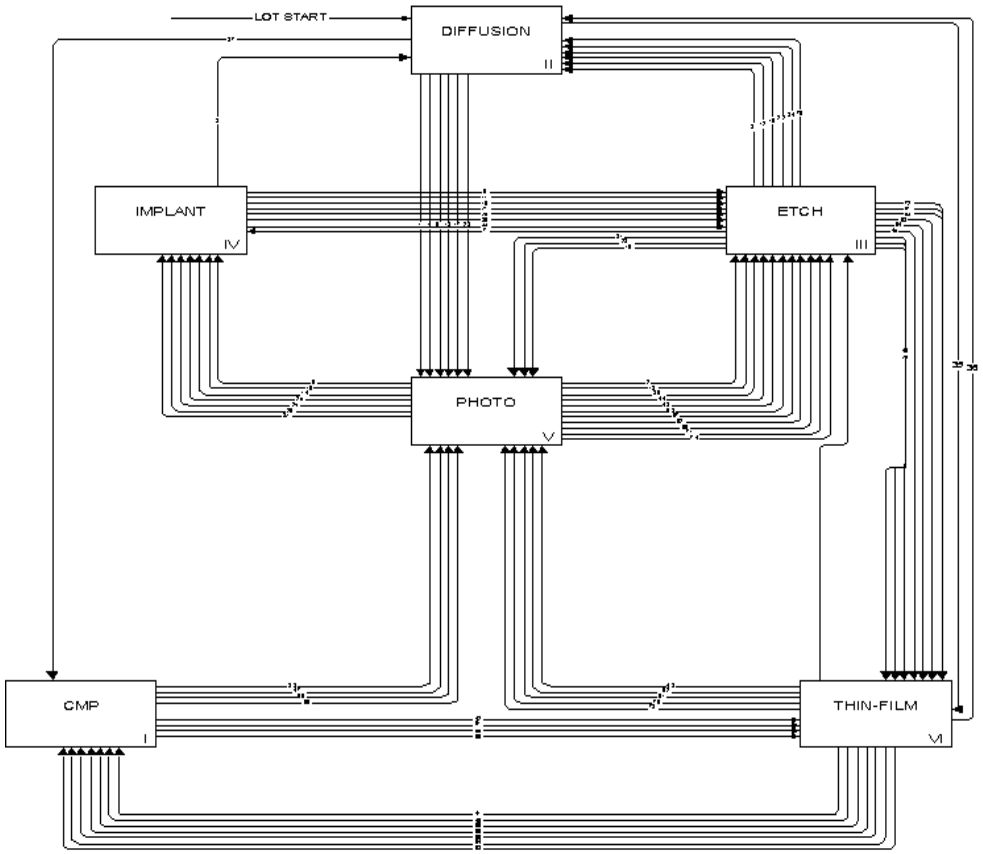

Figure 7-1: Process Flow for 4Rocket lots

7.4.1 Part-type Priority and Lead Time

Often experienced in an MTO environment is that a high priority order may be the result of a shorter customer lead time requirement. Before accepting this order, a manufacturer will have to consider several issues. Among these are the effects of this order on the performance of other existing orders in the fab and the determination of a feasible delivery date for this customer.

This phenomenon is first studied with the fab running under the CPP. The performance of the system using other scheduling policies is discussed in the following

chapter.

Table **??** below shows the performance of the CPP with infinite hedging times and buffer sizes for all the buffers in the fab. The results in each row correspond to a different average customer lead time for 4Rocket lots, expressed in terms of the product of the lead time factor and the processing time. The average customer lead time for 3Rocket lots is kept constant throughout the analysis.

As the average customer lead time decreases, both the service level and the WIP of 4Rocket decrease. Indeed, the percentage drop in the WIP is higher than that of the service level. Since the order arrival rate remains constant, what accounts for a decrease in the WIP is the result of a lower inventory level in the finished goods buffer. This is because when the lead time factor is high (i.e. customer lead time is far into the future), 4Rocket lots are finished too soon compared to the due dates. As a result, most of the 4Rocket lots accumulate in the finished goods buffer. In the subsequent analyses, the lead time factor for a 4Rocket lot is reduced from 10 to 7.

The results also demonstrate the effects of a static priority scheme. To our surprise, the service level of 3Rocket is zero. None of the 3Rocket lots produced can meet the due date. It seems to us that the effects of the lowest priority assignment in Chapter 5 have amplified in the LSIL fab model. We conjecture that as the number of buffers at the control point increases, the probability of selecting a lot from the lowest priority buffer decreases. By the time a lot is selected from the lowest priority buffer, which holds the 3Rocket raw materials, it might not have been possible for the lot to finish processing before its due date. The service level of 4Rocket lots, on the other hand, only suffers a 9.28% decrease when the average customer lead time is reduced to 524.25 hours (3 times the total processing time).

7.4.2 Buffer Priority Scheme and Hedging Time

The previous analysis is carried out under an infinite hedging time assumption. This section experiments with the hedging time parameter to test the possibility of raising the minimal service level by applying the readiness logic to 4Rocket buffers.

Earlier experiments with the toy fab model have lent important insights into the

role of the readiness logic. It is shown in Chapter 5 that the hedging time parameter does not lead to a significant improvement in the minimal service level. Instead, when the readiness logic is applied to the entrance into the system, it reduces the average WIP. In the analysis of the LSIL fab model, the effects of the hedging time parameter are explored by applying the readiness logic to the 4Rocket raw material buffer. Buffer size is again assumed infinite in the analysis of the hedging time parameter.

The results are tabulated and shown in Table VI. Results in italics correspond to hedging time values below the resulting expected remaining cycle time. Similar to the observations of the toy fab model in Section 5.3, even when 4Rocket lots are forced to wait until they are late before being released into the fab, there is still no improvement in the service levels of 3Rocket lots. The hedging time parameter does manage to reduce the amount of 4Rocket WIP without sacrificing its service level.

7.4.3 The Availability Logic

The analysis of the availability logic has so far been excluded due to the number of simulation runs needed. Having gained some insights into the other CPP parameters, it is appropriate to include the study of the availability logic in this section. The CPP parameter involved in the availability logic is the buffer size.

In the availability logic, a lot is loaded only if it is present in the upstream buffer and there is a buffer space available immediately downstream to receive this lot after operation. When the downstream buffer size is small, the buffers are more likely to be full most of the time. Once the buffer space is filled up, the flow of materials will be blocked, regardless of whether they are late or not. Thus, the buffer size parameter can possibly offset the effects of a static priority scheme.

In order to determine where the availability logic is to be applied, we look at the results that have been presented. The static priority scheme has resulted in zero service level for wafer type 3Rocket. Readiness logic control has failed to improve its service level. Consequently, 4Rocket buffers are the candidates to which the availability logic control is applied.

The control point in the system is located at workstation W_DNS_PRECLN, the

entrance into the fab. Both types of wafers are released through this workstation. In fact, they make the same number of visits to this workstation also. The availability logic applied here limit the buffer size after a 4Rocket lot has had its first operation at W_DNS_PRECLN.

The results of buffer size control are tabulated below. As the buffer size for 4Rocket is decreased, 3Rocket service level improves suddenly from zero to 96%, while that of 4Rocket decreases to zero. This calls for the need to have an interactive control of both the hedging time and buffer size parameters for both wafer types.

7.4.4 Application of Both Availability and Readiness Logic

With such an interactive control at the release, there are 4 CPP parameters that can be controlled. They are H_{01} and B_{01} for 4Rocket lots, H_{02} and B_{02} for 3Rocket lots. A detailed analysis requires controlling all these parameters at the same time. However, insights from earlier results help simplify the analysis to three major forms. These parameters and their results are shown in Tables VIII, IX and X.

It is important to note that the trials stop after only three replications for each of the parameter values. This is because the results have shown that the objective of improving the minimal service level is not possible through the application of the CPP logic at the release.

The simulation runs are carried out for B_{01} values equal to 8 and 9. This is because previous results have shown that the service levels for the two wafer types change drastically for these two buffer size values. The test of the parameter values take into account the followings: when the B_{01} value is such that 4Rocket lots have zero service level, then the scheduling logic is applied to 3Rocket raw material buffer to limit the its material flow, and vice versa.

As seen in Tables VIII, IX and X, some of the resulting minimal service levels are non-zero. However, their relative errors are on the order of a hundred percents. We cannot conclude that the corresponding parameters improve the minimal service levels. In summary, applying the CPP scheduling logic at the release does not increase the minimal service level to a non-zero value.

7.5 Summary

Summary This chapter studies the effect of the availability logic when applied to the workstation that is the entrance into the fab. Simulation results show that the buffer size parameter has the ability to offset the effects of a static priority scheme. Moreover, the availability logic actually causes the service levels of one wafer type to reduce to zero while the other to increase to a very high, non-zero value. These observations arise because both the priority scheme and buffer size parameters remain static throughout a simulation run.

Workstation No.	Workstation Name	Operation	Total # Visits	MTTF	VTTF	F Distri	MTTR	VTTR	R Distri	Utilization	# of MACHINES
1	Asher2	Etch	41	20	Exp	1	1.5	0	0	58.8%	11
2	*D_Matrix_1188*	Etch	41	22.3	0.4	5	0.5	0.0	5	N.A	11
3	BPSG	Thin-films	3	50.64	Exp	1	6	0	0	21.4%	4
4	*W_QTI*	Thin-films	3	23.5	0.5	5	2	0	0	N.A	4
5	CVD_TiN	Thin-films	7	28	Exp	1	3	0	0	21.3%	6
6	*BW_Nov_CVD*	Thin-films	7	14	0.2	4	2	0	0	N.A	6
7	CVP_Etch	Etch	13	66.7	Exp	1	8	0	0	19.6%	18
8	*W_CVP_4520*	Etch	13	166	0.7	4	5	0	0	N.A	18
9	DNS_Preclean	Diffusion	27	150	Exp	1	5	0	0	102.7%	8
10	*W_DNS_Precln*	Diffusion	27	72	16	5	1.7	0.1	4	N.A	8
11	Furnace_Allov1	Diffusion	3	250	Exp	1	3	0	0	28.9%	6
12	*W_STI_Furn_Allov1*	Diffusion	3	492	2	5	10	0	0	N.A	6
13	Furnace_Anneal1	Diffusion	12	225	Exp	1	5	0	0	43.6%	30
14	*W_STI_Furn_Anneal1*	Diffusion	12	162	2	5	10	0	0	N.A	30
15	Furnace_Dens1	Diffusion	3	80	Exp	1	1	0	0	41.0%	6
16	*W_STI_Furn_Dens1*	Diffusion	3	162	2	5	10	0	0	N.A	6
17	Furnace_Oxidation2	Diffusion	12	160	Exp	1	5	0	0	35.5%	42
18	*W_SVG_Furn_Oxidation2*	Diffusion	12	162	2	5	10	0	0	N.A	42
19	Furnace_Poly2	Diffusion	3	14.5	Exp	1	2	0	0	42.9%	12
20	*W_SVG_Furn_Poly2*	Diffusion	3	5834	2	5	10	0	0	N.A	12
21	Furnace_SiN2	Diffusion	3	55	Exp	1	5	0	0	34.8%	18
22	*W_SVG_Furn_SiN2*	Diffusion	3	162	2	5	10	0	0	N.A	18
23	High_Energy	Implant	12	75	Exp	1	3	0	0	82.8%	2
24	*W_Eaton_HE*	Implant	12	163.7	2.7	4	3	0.2	4	N.A	2
25	IMD_Dep	Thin-films	7	16.6	Exp	1	3	0	0	33.5%	10
26	BW_Nov_IMD	Thin-films	7	14	0.2	4	1.2	0	0	N.A	10
27	Inspect_CMP4	CMP	20	Inf	Exp	1	0	0	0	N.A	5
28	Inspect_CMP5	CMP	14	Inf	Exp	1	0	0	0	N.A	7
29	Inspect_Diff2	Diffusion	21	Inf	Exp	1	0	0	0	N.A	2
30	Inspect_Etch1	Etch	13	Inf	Exp	1	0	0	0	N.A	2
31	Inspect_Etch2	Etch	21	Inf	Exp	1	0	0	0	N.A	4
32	Inspect_Etch4	Etch	5	Inf	Exp	1	0	0	0	N.A	1
33	Inspect_Etch5	Etch	1	Inf	Exp	1	0	0	0	N.A	1
34	Inspect_Etch6	Etch	4	Inf	Exp	1	0	0	1	N.A	1
35	Inspect_Imp	Implant	19	Inf	Exp	1	0	0	0	N.A	2
36	Inspect_Photo1	Photo	19	Inf	Exp	1	0	0	0	N.A	4
37	Inspect_Photo4	Photo	22	Inf	Exp	1	0	0	0	N.A	3
38	Inspect_Photo5	Photo	42	Inf	Exp	1	0	0	0	N.A	6
39	Inspect_TF1	Thin-films	13	Inf	Exp	1	0	0	0	N.A	2
40	Inspect_TF2	Thin-films	8	Inf	Exp	1	0	0	0	N.A	1
41	Inspect_TF3	Thin-films	7	Inf	Exp	1	0	0	0	N.A	2
42	Island_Spacer_Etch	Etch	6	47.368	Exp	1	4	0	0	19.3%	6
43	*W_IS_4520*	Etch	6	166	0.7	4	1	0	0	N.A	6

Table II: LSIL Fab Equipment List

No.	Workstation Name	Operation	Total # Visits	MTTF	VTTF	F Distri	MTTR	VTTR	R Distri	Utilization	# of MACHINES
44	LTO	Thin-films	6	54	Exp	1	3	0	0	15.2%	4
45	*BW_Nov_LTO*	Thin-films	6	14	0.2	4	0.75	0	0	N.A	4
46	Medium_Current	Implant	20	44	Exp	1	2	0	0	33.1%	4
47	*W_Varian*	Implant	20	163.7	2.7	4	5	0.2	4	N.A	4
48	Metal_Dep	Thin-films	10	35	Exp	1	2.8	0	0	29.7%	8
49	*Kit_Chg_AMAT_AlCu*	Thin-films	10	252	1	4	7	0.2	4	N.A	8
50	Metal_Etch	Etch	10	60	Exp	1	6.8	0	0	22.0%	15
51	*W_9600*	Etch	10	166	0.7	4	2	0	0	N.A	15
52	Nitride_Strip	Diffusion	3	100	Exp	1	8	0	0	101.5%	2
53	*W_DNS_Nitride*	Diffusion	3	72	16	4	1.7	0.1	4	N.A	2
54	Oxide_CMP	CMP	10	70	Exp	1	5	0	0	56.6%	5
55	*W_IPEC*	CMP	10	166	4.7	4	2	1	5	N.A	5
56	Passivation	Thin-films	3	65	Exp	1	2	0	0	9.5%	4
57	*BW_Nov_Pass*	Thin-films	3	15.3	1.1	4	1	0	0	N.A	4
58	Poly_Etch	Etch	3	13.3	Exp	1	2	0	0	26.8%	6
59	*W_Poly_9400*	Etch	3	166	0.7	4	1	0	0	N.A	6
60	PreMetal_Sink	Thin-films	7	38	Exp	1	5	0	0	37.7%	4
61	*D_DNS_Pre_Metal*	Thin-films	7	72	16	4	1.7	0.1	4	N.A	4
62	Resist_Strip	Etch	28	67	Exp	1	10	0	0	38.5%	16
63	*D_STI_Resist_Strip*	Etch	28	23	0.2	4	1.5	0.0	4	N.A	16
64	RTA	Thin-films	9	46	Exp	1	4	0	0	30.5%	8
65	*W_DNS_RTA*	Thin-films	9	168	4	4	1	0	0	N.A	8
66	Solvent_Clean	Etch	23	300	Exp	1	3	0	0	17.7%	16
67	*D_STI_Solvent*	Etch	23	22	0.7	4	1.5	0.0	4	N.A	16
68	Sort_CMP	CMP	10	Inf	Exp	1	0	0	0	N.A	4
69	Stepper_2	Photo	39	75	Exp	1	3	0	0	66.2%	8
70	*W_Nikon_I14*	Photo	39	168	2.7	4	1.1	0.13	5	N.A	8
71	Stepper_3	Photo	10	34	Exp	1	3	0	0	31.7%	6
72	*W_Nikon_Deep_UV*	Photo	10	168	2.7	4	1.1	0.13	5	N.A	6
73	Ti_Dep	Thin-films	13	70	Exp	1	5	0	0	27.9%	8
74	*Kit_Chg_AMAT_TiN*	Thin-films	13	252	1	4	8	0.2	4	N.A	8
75	Tungsten_CMP	CMP	10	53	Exp	1	4	0	0	22.0%	8
76	*W_IPEC_W*	CMP	10	166	4.7	4	2	1	5	N.A	8
77	UV_Cure	Implant	8	400	Exp	1	4	0	0	13.1%	4
78	*S_UV_Cure*	Implant	8	7	0.2	4	0.25	0	4	N.A	4
79	W_CVD	Thin-films	10	470	Exp	1	2	0	0	10.3%	8
80	*BW_Nov_W*	Thin-films	10	14	0.2	4	1	0	0	N.A	8

Table III: LSIL Fab Equipment List

Explanations of Equipment Description Table:	
MTTF = Mean time to fail	*Probability Distributions*
VTTF = Variance of time to fail	*0 = Deterministic*
MTTR = Mean time to repair	*1 = Exponential*
VTTR = Variance of time to repair	*2 = Gamma*
F Distri = Failure probability distribution	*3 = Lognormal*
R Distri = Repair probability distribution	*4 = Normal*
Exp = Square of the mean for exponential variables.	*5 = Uniform*
N.A = % Utilization for Inspection operations and virtual workstations	
Inf = Infinite mean time to fail	
Notes	
1. Total number of visits include visits by 3ProtoG11, 4RocketG11, 3RocketG11 Only the last two part types are modeled in the simulation.	
2. The number of machines in each workstation takes into account the multi-capacity characteristics of the batch machines.	

Table IV: Explanations on the Equipment List

Performance of the CPP, with infinite buffer size and hedging times						
Lead-time factors [4Rocket, 3Rocket]	**4Rocket Sv**	**4Rocket WIP**	**3Rocket Sv**	**3Rocket WIP**	**Min. Service Level (Sv)**	**Total WIP**
[10, 10]	0.959	640.3	0.000	579.1	0.000	1219.3
[9.5, 10]	0.958	607.7	0.000	579.6	0.000	1187.3
[8.5, 10]	0.956	544.6	0.000	575.9	0.000	1120.5
[7.5, 10]	0.951	479.1	0.000	583.1	0.000	1062.2
[6.5, 10]	0.945	416.7	0.000	583.6	0.000	1000.3
[6, 10]	0.940	300.1	0.000	582.5	0.000	882.5
[5, 10]	0.931	278.6	0.000	575.6	0.000	854.1
[4, 10]	0.911	244.6	0.000	577.1	0.000	821.7
[3, 10]	0.870	195.1	0.000	579.6	0.000	774.7

Table V: Explanations on the Equipment List

Performance of the CPP, with infinite buffer size.						
H_{01}, for 4Rocket raw material buffer	4Rocket Sv	4Rocket WIP	3Rocket Sv	3Rocket WIP	Min. Service Level (Sv)	Total WIP
2000	0.951	466.8	0.00	580.8	0.00	1047.6
1850	0.951	453.0	0.00	588.0	0.00	1041.0
1700	0.949	441.5	0.00	592.9	0.00	1034.3
1550	0.952	424.7	0.00	577.9	0.00	1002.6
1400	0.951	401.0	0.00	582.0	0.00	983.0
1250	0.953	373.5	0.00	589.1	0.00	962.6
1100	0.949	343.9	0.00	586.7	0.00	930.5
950	0.951	306.8	0.00	566.6	0.00	873.5
800	*0.952*	*264.8*	*0.00*	*589.0*	*0.00*	*853.9*
650	*0.949*	*221.0*	*0.00*	*589.4*	*0.00*	*810.4*
500	*0.952*	*173.4*	*0.00*	*584.5*	*0.00*	*757.9*
350	*0.950*	*124.1*	*0.00*	*579.7*	*0.00*	*703.7*
210	*0.063*	*84.2*	*0.00*	*577.0*	*0.00*	*661.3*
200	*0.068*	*81.1*	*0.00*	*560.0*	*0.00*	*641.2*
160	*0.000*	*84.7*	*0.00*	*584.3*	*0.00*	*668.9*

Table VI: The CPP performance with only readiness logic control

Performance of the CPP, with infinite hedging time values						
B_{01}	4Rocket Sv	4Rocket WIP	3Rocket Sv	3Rocket WIP	Min. Service Level (Sv)	Total WIP
10	0.950563	452.7	0.00	580.7	0.00	1033.4
9	0.908901	212.5	0.00	485.9	0.00	698.4
8	0	70.5	0.96	227.4	0.00	297.8
7	0	68.3	0.96	227.8	0.00	296.1

Table VII: The CPP performance with only availability logic control

Performance of the CPP							
B_{01}	H_{01}	4Rocket Sv	4Rocket WIP	3Rocket Sv	3Rocket WIP	Min. Service Level (Sv)	Total WIP
9	2300	0.86	204.3	0.000	488.6	0.000	692.8
9	1800	0.71	176.3	0.145	500.9	0.145	677.2
9	1300	0.56	136.1	0.005	478.9	0.005	615.0
9	800	0.36	96.5	0.000	517.4	0.000	614.0
9	350	0.00	78.9	0.000	409.1	0.000	488.0
9	230	0.00	77.9	0.000	512.0	0.000	589.9

Table VIII: Readiness Logic Control when $B_{01} = 9$

Performance of the CPP							
B_{01}	H_{02}	4Rocket Sv	4Rocket WIP	3Rocket Sv	3Rocket WIP	Min. Service Level (Sv)	Total WIP
8	1800	0.00	70.0	0.96	345.0	0.00	415.0
8	1400	0.00	69.9	0.96	298.6	0.00	368.5
8	1000	0.00	70.3	0.96	228.5	0.00	298.8
8	600	0.00	70.3	0.96	143.1	0.00	213.4
8	200	0.00	69.9	0.037	67.4	0.00	137.2

Table IX: Readiness Logic Control when $B_{01} = 8$

Performance of the CPP							
B_{01}	B_{02}	4Rocket Sv	4Rocket WIP	3Rocket Sv	3Rocket WIP	Min. Service Level (Sv)	Total WIP
8	11	0.00	70.2	0.960	304.0	0.00	374.3
8	10	0.00	70.5	0.960	302.7	0.00	373.2
8	9	0.00	70.5	0.959	299.0	0.00	369.5
8	8	0.00	70.8	0.956	196.1	0.00	267.0
8	7	0.00	71.5	0.000	58.7	0.00	130.3

Table X: Availability Logic Control when $B_{01} = 8$

Chapter 8

Scheduling Policy Comparison and New Performance Measures

This chapter presents the results of other scheduling policies applied to both the toy fab and the LSIL fab models. They are the Earliest-Due-Date (EDD) and Critical Ratio (CR) policies. Their results form a basis of comparison to assess how well the CPP is performing in each model. Both of these policies are run with the same model parameters for a fair comparison.

In addition, new performance characterization is also presented. It explores the effects on the system performance when order arrival rates are increased.

Section 8.1 discusses the results of each model under the EDD scheduling policy. The results of the CR scheduling policy are presented in Section 8.2. Comparison with the CPP is discussed in Section 8.3. The investigation of new performance measures is explained in Section 8.4. The chapter concludes with a summary.

8.1 Earliest-Due-Date Policy

In an MTO environment, the due date of a part is often used in the scheduling logic. The simplest form, then, is a scheduling policy that is based solely on the due date information of a part. As the name suggests, the EDD policy gives priority to parts that have the earliest due dates for processing operations.

Results of the toy fab and LSIL fab models are given in Sections 8.1.1 and ?? respectively. It is important to note that the results obtained do not satisfy the statistical requirement as explained in Chapter 3. Instead, the 95% confidence intervals (95% C.I.) are given for the results of each performance measure. This is because the relative errors obtained are more than 3% even after thirty replications. Variability in the results is considerably greater than that obtained from the CPP. Nonetheless, at least 20 replications are used to compute the 95% confidence intervals for all the EDD results.

The variability in the results implies that the EDD performance is hard to predict. It becomes hard to co-ordinate downstream operations on the completed wafers, such as assembly, etc.

8.1.1 Toy Fab Model

Table I below shows the results of the EDD policy used in the toy fab model. If minimal service level is used as a performance measure, the EDD policy does better than the CPP for the CPP parameters analyzed in Chapter 5. However, the EDD policy is not suitable in environments where a certain type of customer orders has priority over others. This is because its scheduling logic is not able to give priority to a certain set of customer orders.

Toy Fab Model Performance Under EDD					
p1 Sv ± 95% C.I	p2 WIP ± 95% C.I	p2 Sv ± 95% C.I	p2 WIP ± 95% C.I	Min Sv ± 95% C.I	Total WIP ± 95% C.I
0.734 ± 0.023	41.4 ± 0.4	0.759 ± 0.022	26.1 ± 0.2	0.734 ± 0.023	67.6 ± 0.56

Table I: Toy Fab Performance Under EDD

8.1.2 LSIL Fab Model

As explained in Section 7.4.1, a high priority order may be the result of a shorter customer lead time requirement. The CPP study has shown that even when the lead

time of the high priority order is short, the CPP is still able to maintain a very high service level. A similar analysis is carried out in the LSIL fab model using the EDD policy. The results are shown in Table II.

LSIL Fab Model Performance Under EDD						
Lead-time factors [4Rocket, 3Rocket]	4Rocket Sv ± 95% C.I	4Rocket WIP ± 95% C.I	3Rocket Sv ± 95% C.I	3Rocket WIP ± 95% C.I	Min Sv ± 95% C.I	Total WIP ± 95% C.I
[10, 10]	0.845 ± 0.08	631.49 ± 3.17	0.859 ± 0.076	394.5 ± 1.87	0.845 ± 0.08	1026 ± 4.3
[9.5, 10]	0.818 ± 0.09	606.1 ± 2.12	0.836 ± 0.084	396.8 ± 1.55	0.817 ± 0.09	1002.9 ± 1.8
[8.5, 10]	0.799 ± 0.05	548.19 ± 2.03	0.833 ± 0.05	397.35 ± 2.34	0.799 ± 0.052	945.54 ± 3.82
[7.5, 10]	0.555 ± 0.114	505.71 ± 11.1	0.587 ± 0.111	411.66 ± 6.49	0.555 ± 0.114	917.4 ± 17.4
[6.5, 10]	0.41 ± 0.116	455.28 ± 12.87	0.442 ± 0.117	420.5 ± 8.22	0.41 ± 0.116	875.8 ± 21

Table II: LSIL Fab Performance Under EDD

As the customer lead time factor is reduced, the service levels of both part types decrease significantly. 4Rocket service level drops below 0.5 when its lead time factor is 6.5. For the same lead time factor, the CPP is still able to produce a service level of 0.945 for 4Rocket lots. Thus, the CPP is a better policy in these situations.

8.2 Critical Ratio Policy

Similar to the results of the EDD policy, there is large variability in the results of the performance measures under the CR policy. Nonetheless, twenty replications are used to compute the 95% confidence intervals for the CR policy results in both the toy fab and LSIL fab models. The variability implies that production completion date is hard to predict.

8.2.1 Toy Fab Model

In the CR scheduling logic, we include a critical ratio factor to give priority to a certain set of customer orders. It is thought that a high factor value will lead to a better service level. Readers are referred to Section 3.2 for explanations of the CR scheduling logic.

In the toy fab model, priority is given to part type 1. This is achieved in the CR policy by increasing the value of the critical ratio factor corresponding to p1. The value of p2 factor remains at one for all the cases. The results of the analysis are shown in Table III.

Toy Fab Model Performance Under Critical Ratio (CR)						
CR Factor	p1 Sv ± 95% C.I	p1 WIP ± 95% C.I	p2 Sv ± 95% C.I	p2 WIP ± 95% C.I	Min Sv ± 95% C.I	Total WIP ± 95% C.I
1.5	0.82 ± 0.025	46.3 ± 0.5	0.795 ± 0.028	24.5 ± 0.5	0.795 ± 0.028	70.8 ± 0.8
5.5	0.824 ± 0.021	48.2 ± 0.6	0.663 ± 0.028	21 ± 0.7	0.663 ± 0.028	69.2 ± 1
10.5	0.791 ± 0.025	48.7 ± 0.4	0.458 ± 0.027	20.4 ± 0.7	0.458 ± 0.027	69.1 ± 0.8
100	0.806 ± 0.028	48.7 ± 0.4	0.256 ± 0.024	21.1 ± 0.6	0.256 ± 0.024	69.8 ± 0.9

Table III: Toy Fab Performance Under CR

The results show that a higher value of the critical ratio factor does not improve p1 service level significantly. In fact, the service level suffers a slight decrease when the factor value is too high. In addition, p2 service level is decreasing with increasing values of p1 CR factor. As such, there is no incentive to raise the CR factor value. It also implies that the CR scheduling logic is not able to give priority to certain type of customer order through the CR factor.

8.2.2 LSIL Fab Model

The results of LSIL fab model analysis with various lead time factor for 4Rocket lots are shown in Table IV.

Comparing with the EDD results, the CR policy leads to higher service levels when the lead time factors of both wafer types are comparable. However, when the lead time factor for 4Rocket lots decreases to less than 8.5, EDD policy leads to higher service levels for both wafer types.

It is also observed that for both EDD and CR scheduling policies, the service levels of both wafer types are very close to one another. Even when the average customer lead time for 4Rocket lots is shorter than that of 3Rocket lots (i.e. 4Rocket lead time factor of 8.9 and less), the service levels are still comparable.

LSIL Performance Under Critical Ratio						
Lead-time factors [4Rocket, 3Rocket]	4Rocket Sv ± 95% C.I	4Rocket WIP ± 95% C.I	3Rocket Sv ± 95% C.I	3Rocket WIP ± 95% C.I	Min Sv ± 95% C.I	Total WIP ± 95% C.I
[10, 10]	0.906 ± 0.034	635.9 ± 2.6	0.905 ± 0.035	393 ± 1.6	0.904 ± 0.035	1029 ± 3.7
[9.5, 10]	0.841 ± 0.065	603.8 ± 1.3	0.842 ± 0.066	394.9 ± 1.9	0.84 ± 0.066	998.7 ± 2
[8.5, 10]	0.772 ± 0.069	552.2 ± 3.8	0.776 ± 0.069	400 ± 1.8	0.772 ± 0.069	952.2 ± 5.2
[7.5, 10]	0.565 ± 0.082	505.5 ± 8.2	0.568 ± 0.083	410.1 ± 4.2	0.564 ± 0.082	915.6 ± 12
[6.5, 10]	0.345 ± 0.085	466.8 ± 11.2	0.347 ± 0.087	427.6 ± 6.7	0.344 ± 0.086	894.4 ± 17.8

Table IV: LSIL Fab Performance Under CR

The results of the CR policies from both the toy fab and LSIL fab analyses lead to an interesting observation. The service levels of both part types are always so close to one another, even when we attempt to give priority to one part type over the other. This may be due to the way the CR scheduling logic treats late materials. A late part has priority over non-late parts. Therefore, even when priority is given to one part type such that the WIP of this part type are always processed, it will come to a point where the WIP of the low priority part type becomes late. When this happens, the scheduling logic will allow the late parts to go through, but not the high priority, non-late WIP. As such, the CR policy is likely to do well under the minimal service level performance measure.

It is also observed that the CR policy does better than the EDD policy. This may be because the computation of critical ratio also makes use of the remaining processing time information. We feel that critical ratio is a better reflection of the urgency of the parts and this may account for the better performance over the EDD policy.

8.3 Comparison of Scheduling Policies

The performance of the CPP is easier to predict compared to that of the EDD and CR. The CPP allows a more accurate prediction of production completion time, which improves downstream coordination of further operations on completed wafers.

Among the buffer priority schemes studied, the CPP does not perform as well as the EDD and CR policies when minimal service level is used as the performance measure. This is due to the effects of static priority scheme. However, the CPP is the policy of choice in situations where one type of customer order has priority over the others.

8.4 Alternative Method to Assess Policy Performance

Manufacturers always face the problems of accepting new customer orders. They have to determine if there is sufficient capacity to accommodate new orders and meet their due dates. There is great incentive to accept new orders since they represent more revenue. At the same time, however, they place more strain on the finite capacity. In fact, a higher demand rate may possibly cause the service levels of existing orders to decrease, and this will turn away existing customers.

It is desirable to have information on the relationship between demand rates and the corresponding service levels. A manufacturer is then able to make better decisions on whether to accept new customer orders.

8.4.1 Relationship between Demand Rates and Service Levels

We try to determine such a relationship using the toy fab model. It is run under the CPP scheduling policy with buffer priority given to returning p1 parts at both workstations. All buffer sizes and hedging times are assumed infinite. The analysis is carried out by fixing the demand rates for p1 while varying the rates for p2. The resulting service levels for both part types are shown in Figure 8-1.

With the buffer priority scheme used, the p1 service level is much higher than that of p2. In fact, for a large range of p2 demand rates, p1 service levels remain above 0.90 even when the p2 demand rates are increased. The results agree with the

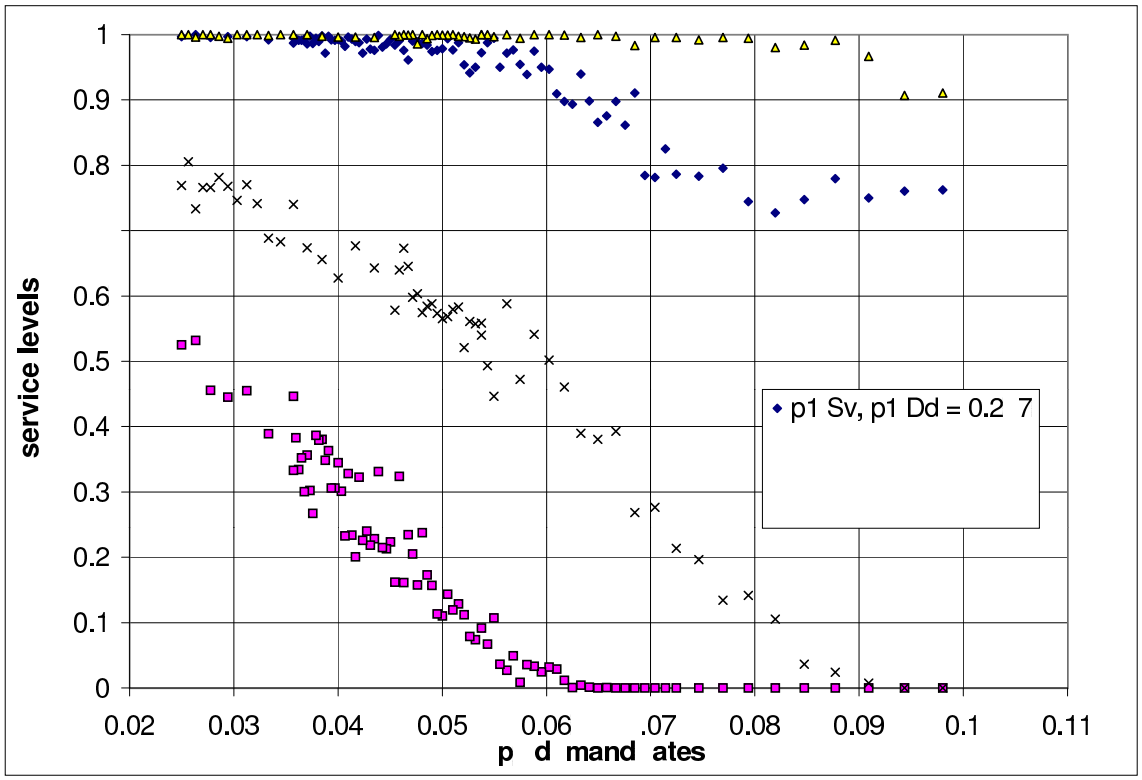
service levels
1
0.9
0.8
0.6
0.5
0.4
0.3
0.2
0.1
0
0.02
0.03
0.04
0.05
0.06
.07
0.08
0.09
0.1
0.11
p d mand ates
p1 Sv, p1 Dd = 0.2 7

a cluster is not observed when the demand rates are high.

8.4.2 New Performance Characterization

The insights gained from the analysis of the relationship between demand rates and service levels lead us to suggest an alternative performance measure. We can plot iso-service-level curves corresponding to a scheduling policy. Along each of these curves are the demand rates for all part types corresponding to the same service level. Such a curve could enable a manufacturer to align his/her decisions to accept customer orders based on a pre-determined goal. If the company aims to achieve a minimal service level of 90%, the manufacturer will then look at the 90% iso-service curve to determine the acceptable demand rate for each product.

Such an approach is explored using the toy fab model. We use the CPP with buffer priority given to the retuning p1 parts at both workstations. Buffer sizes and hedging times are again assumed infinite. The iso-service curve corresponds to the minimal service level of the toy fab. However, it is not possible to fix a value for the minimal service level and determine the corresponding demand rates due to variability in simulation runs. Instead, a 0.5% offset is allowed above and below the minimal service value of interest. Figure 8-2 shows the demand rates corresponding to two minimal service values, namely, 90% and 95%. The results obtained satisfy the relative error statistical requirement.

The results agree with our intuition. A higher minimal service level is attainable only with lower demand rates under the same scheduling policy and customer lead time. Fixing the demand rate for one part type, there are multiple demand rates for the other part type that also lead to roughly the same service level. This observation is also consistent with the results from the analysis of demand rates and service levels.

This performance measure can also be a basis of scheduling policy comparison. A good scheduling policy is able to accommodate high demand rates while maintaining high service levels for all the products. In a plot of iso-service curves, the better scheduling policy will have a higher iso-service curve.

Such a performance measure does have disadvantages. First, the warm-up as

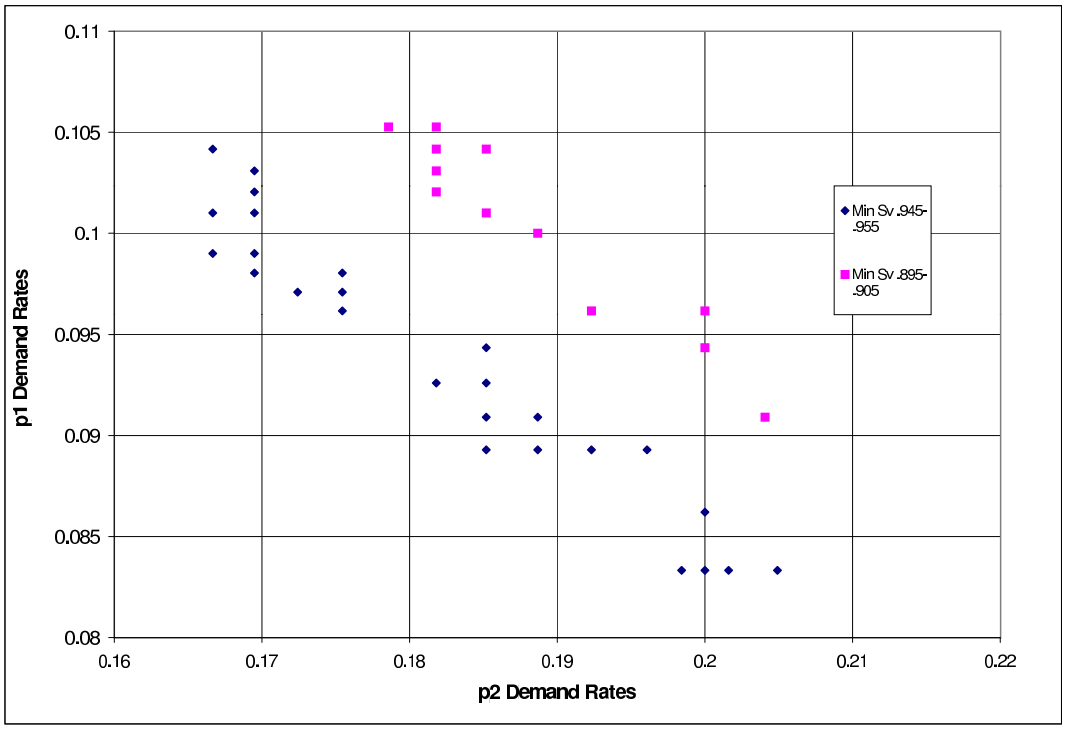

Figure 8-2: Iso-minimal service levels using the CPP

well as simulation time periods are functions of the demand rates. In the process of searching demand rates that lead to the same service level, it is necessary to ensure the simulation run-time is long enough. This calls for a very conservative estimate of the minimum warm-up time. As a result, each simulation run becomes very time consuming.

There is another reason that calls for long warm-up period. The search for feasible demand rates is carried out in a trial-and-error process. As such, many demand rate values tested are very close to or higher than the capacity of the system. We know from the analysis of a simple M/M/1 queue that the delay and the average number of parts in the system increase dramatically when the material arrival rate approaches the capacity. Readers can refer to Gershwin [5], Section 2.3.6 for a discussion of this phenomenon.

In the toy fab model, which is a more complex queuing system, an insufficient warm-up period can cause large fluctuations in the service levels when the demand

rates are very high. In one replication, the service level can be as high as 90%; in another, it can drop to 30%. However, as the warm-up period is increased, the service levels drop to zero consistently in all the replications. Due to these reasons, the results generated from the toy fab model are based on a warm-up period of a million time units in order to capture steady state performance.

Another disadvantage is that a policy with scheduling parameters, such as hedging times or the number of production authorization cards between workstations, will not be optimal generally. This is because this performance measure requires varying the order arrival rates. However, we know that the policy parameters are functions of the order arrival rates. Good parameter values can only be determined upon knowing the arrival rates. Nonetheless, scheduling policies such as EDD can utilize this measure for comparison purposes since they make use of only the due date and processing time information.

8.5 Summary

This chapter presents the results of the toy fab and LSIL fab systems under the EDD and CR scheduling policies. When the minimal service level is used as the performance measure, both policies fare better than the CPP for the CPP parameters analyzed in Chapter 5 and 7. However, both the EDD and CR scheduling policies are not able to perform well in environment where one set of customer orders is preferred over the rest. In this case, the CPP performance is much better than those of the EDD and CR policies. A performance measure that uses iso-service level as a basis of policy comparison is explained. It has some practical advantages but it allows a fair comparison of only some scheduling policies.

Chapter 9

Conclusion

This thesis studies the performance of scheduling policies in multi-product, make-to-order semiconductor fabrication facilities. The three scheduling policies analyzed are the Earliest-Due-Date, Critical Ratio and Control Point policies. In particular, the CPP and its parameters are explored in detail through simulation experiments. The results obtained provide insights into the roles of the static buffer priority, the hedging time and buffer size parameters.

The research results are summarized by first discussing the effects of the CPP parameters in Section 9.1. The conclusions of policy comparison are stated in Section 9.2. Further research ideas are discussed in Section 9.3.

9.1 The CPP Parameters

In the analysis of buffer priority, the resulting service levels are strongly characterized by its discrete nature. This is evident by the distinctive clustering of service levels in the priority schemes studied. In general, when the buffer of a part type is assigned the lowest priority, the service level of this part type is worse than that of the other. The absolute difference between the service levels can be as high as sixty percentage points. In situations where there are many buffers at the control point, the undesirable effects of the lowest priority assignment are amplified, as seen in the results of LSI Logic fab.

The results also reveal that priority given to parts returning for more advanced process steps at workstation 1 (WS 1), the entrance into the system, leads to better service level performance. At workstation 2, on the other hand, priority should be given to parts in their early visits for better performance.

We also attempt to use hedging time to offset the effects of static priority scheme. With the control point located at WS 1, the readiness logic does not lead to significant improvement in the minimal service level. The failure of the readiness logic to control the flow of late parts is the main reason that limits the extent to which hedging time can offset the effects of a static buffer priority scheme. On the other hand, it significantly reduces the amount of WIP in the system when the control is applied to the raw material buffers.

In the study of hedging time, we cannot conclusively determine if good hedging times can be below the resulting expected remaining cycle time. This is because the performance measure used to determine good hedging times may not be appropriate as the results have indicated. However, it is found that when the hedging time for the buffer of one part type falls below the expected remaining cycle time, the corresponding service level drops sharply.

Buffer size is able to offset the effects of a static buffer priority scheme. In fact, as the buffer size is reduced, the availability logic can cause the service level of one part type to drop to zero while the other increases to a very high value.

More work has to be done to understand the CPP and its parameters. Even though these results are not conclusive since they are only based on two models, they do provide insights into the roles of each policy parameter. We recommend that all three of the policy parameters, namely buffer priority, hedging time and buffer size, should be chosen in such a way as to suit a specified goal. In a multi-product system, if the buffer priority scheme is assigned to suit one performance measure, we should not use the hedging time and buffer size parameters to offset the effects of the buffer priority scheme in order to meet a different objective.

9.2 Comparison with other Scheduling Policies

Due to the many policy parameters, a full analysis of the CPP was not carried out. As such, the comparison of the CPP with other scheduling policies is not conclusive. Despite an only partial analysis of the CPP, there are some lessons to be learned from the cases of the CPP analyzed. We have found that due to the effects of static priority scheme, the CPP does not perform as well as the EDD and CR policies when minimal service level is used as the performance measure. However, the CPP is the policy of choice in situations where one type of customer orders has priority over the others.

In terms of minimal service level, the CR policy is a better choice over the EDD policy. The main reason is that the value of critical ratio is a better representation of the urgency of a part. In addition to due date, critical ratio also takes into account the remaining process time.

9.3 Recommendations for Further Research

This research can be extended in several directions. Some of the possible areas include:

- Developing guidelines to select buffer priority schemes that perform well when minimal service level is used as the performance measure.
- Exploring the performance of the CPP using metrics such as the mean and variance of cycle time.
- Introducing procedures to deal with setup, yield and batch operation.
- Comparing the performance of the CPP with other scheduling policies such as Least Slack, etc.

Appendix A

Buffer Priority Schemes in Chapter 4

Case 1 Buffer Prirority Schemes

Note: *Due to the requirement of the simulation program, the higher the priority value, the higher the priority of that buffer.
*This definition is in exact opposite to that introduced in Chapter 4

Priority	Part Type 1						Part Type 2			
Scheme No.	B01	B11	B21	B31	B41	B51	B02	B12	B22	B32
0	0	0	2	1	4	2	1	3	3	4
1	0	0	2	1	4	2	1	4	3	3
2	0	0	2	1	4	3	1	2	3	4
3	0	0	2	1	4	3	1	4	3	2
4	0	0	2	1	4	4	1	2	3	3
5	0	0	2	1	4	4	1	3	3	2
6	0	0	2	2	4	1	1	3	3	4
7	0	0	2	2	4	1	1	4	3	3
8	0	0	2	2	4	3	1	1	3	4
9	0	0	2	2	4	3	1	4	3	1
10	0	0	2	2	4	4	1	1	3	3
11	0	0	2	2	4	4	1	3	3	1
12	0	0	2	3	4	1	1	2	3	4
13	0	0	2	3	4	1	1	4	3	2
14	0	0	2	3	4	2	1	1	3	4
15	0	0	2	3	4	2	1	4	3	1
16	0	0	2	3	4	4	1	1	3	2
17	0	0	2	3	4	4	1	2	3	1
18	0	0	2	4	4	1	1	2	3	3
19	0	0	2	4	4	1	1	3	3	2
20	0	0	2	4	4	2	1	1	3	3
21	0	0	2	4	4	2	1	3	3	1
22	0	0	2	4	4	3	1	1	3	2
23	0	0	2	4	4	3	1	2	3	1
24	0	1	2	0	4	2	1	3	3	4
25	0	1	2	0	4	2	1	4	3	3
26	0	1	2	0	4	3	1	2	3	4
27	0	1	2	0	4	3	1	4	3	2
28	0	1	2	0	4	4	1	2	3	3
29	0	1	2	0	4	4	1	3	3	2
30	0	1	2	2	4	0	1	3	3	4
31	0	1	2	2	4	0	1	4	3	3
32	0	1	2	2	4	3	1	0	3	4
33	0	1	2	2	4	3	1	4	3	0
34	0	1	2	2	4	4	1	0	3	3
35	0	1	2	2	4	4	1	3	3	0
36	0	1	2	3	4	0	1	2	3	4
37	0	1	2	3	4	0	1	4	3	2
38	0	1	2	3	4	2	1	0	3	4
39	0	1	2	3	4	2	1	4	3	0
40	0	1	2	3	4	4	1	0	3	2
41	0	1	2	3	4	4	1	2	3	0
42	0	1	2	4	4	0	1	2	3	3
43	0	1	2	4	4	0	1	3	3	2

Table I: Detailed breakdown of all Case 1 priority schemes

Case 1 Buffer Prirority Schemes

Note: *Due to the requirement of the simulation program, the higher the priority value, the higher the priority of that buffer.
*This definition is in exact opposite to that introduced in Chapter 4

Priority Scheme	Part Type 1						Part Type 2			
	B01	B11	B21	B31	B41	B51	B02	B12	B22	B32
44	0	1	2	4	4	2	1	0	3	3
45	0	1	2	4	4	2	1	3	3	0
46	0	1	2	4	4	3	1	0	3	2
47	0	1	2	4	4	3	1	2	3	0
48	0	2	2	0	4	1	1	3	3	4
49	0	2	2	0	4	1	1	4	3	3
50	0	2	2	0	4	3	1	1	3	4
51	0	2	2	0	4	3	1	4	3	1
52	0	2	2	0	4	4	1	1	3	3
53	0	2	2	0	4	4	1	3	3	1
54	0	2	2	1	4	0	1	3	3	4
55	0	2	2	1	4	0	1	4	3	3
56	0	2	2	1	4	3	1	0	3	4
57	0	2	2	1	4	3	1	4	3	0
58	0	2	2	1	4	4	1	0	3	3
59	0	2	2	1	4	4	1	3	3	0
60	0	2	2	3	4	0	1	1	3	4
61	0	2	2	3	4	0	1	4	3	1
62	0	2	2	3	4	1	1	0	3	4
63	0	2	2	3	4	1	1	4	3	0
64	0	2	2	3	4	4	1	0	3	1
65	0	2	2	3	4	4	1	1	3	0
66	0	2	2	4	4	0	1	1	3	3
67	0	2	2	4	4	0	1	3	3	1
68	0	2	2	4	4	1	1	0	3	3
69	0	2	2	4	4	1	1	3	3	0
70	0	2	2	4	4	3	1	0	3	1
71	0	2	2	4	4	3	1	1	3	0
72	0	3	2	0	4	1	1	2	3	4
73	0	3	2	0	4	1	1	4	3	2
74	0	3	2	0	4	2	1	1	3	4
75	0	3	2	0	4	2	1	4	3	1
76	0	3	2	0	4	4	1	1	3	2
77	0	3	2	0	4	4	1	2	3	1
78	0	3	2	1	4	0	1	2	3	4
79	0	3	2	1	4	0	1	4	3	2
80	0	3	2	1	4	2	1	0	3	4
81	0	3	2	1	4	2	1	4	3	0
82	0	3	2	1	4	4	1	0	3	2
83	0	3	2	1	4	4	1	2	3	0
84	0	3	2	2	4	0	1	1	3	4
85	0	3	2	2	4	0	1	4	3	1
86	0	3	2	2	4	1	1	0	3	4
87	0	3	2	2	4	1	1	4	3	0
88	0	3	2	2	4	4	1	0	3	1
89	0	3	2	2	4	4	1	1	3	0
90	0	3	2	4	4	0	1	1	3	2
91	0	3	2	4	4	0	1	2	3	1

Table II: Detailed breakdown of all Case 1 priority schemes

Case 1 Buffer Prirority Schemes

Priority Scheme No.	Part Type 1						Part Type 2			
	B01	B11	B21	B31	B41	B51	B02	B12	B22	B32
92	0	3	2	4	4	1	1	0	3	2
93	0	3	2	4	4	1	1	2	3	0
94	0	3	2	4	4	2	1	0	3	1
95	0	3	2	4	4	2	1	1	3	0
96	0	4	2	0	4	1	1	2	3	3
97	0	4	2	0	4	1	1	3	3	2
98	0	4	2	0	4	2	1	1	3	3
99	0	4	2	0	4	2	1	3	3	1
100	0	4	2	0	4	3	1	1	3	2
101	0	4	2	0	4	3	1	2	3	1
102	0	4	2	1	4	0	1	2	3	3
103	0	4	2	1	4	0	1	3	3	2
104	0	4	2	1	4	2	1	0	3	3
105	0	4	2	1	4	2	1	3	3	0
106	0	4	2	1	4	3	1	0	3	2
107	0	4	2	1	4	3	1	2	3	0
108	0	4	2	2	4	0	1	1	3	3
109	0	4	2	2	4	0	1	3	3	1
110	0	4	2	2	4	1	1	0	3	3
111	0	4	2	2	4	1	1	3	3	0
112	0	4	2	2	4	3	1	0	3	1
113	0	4	2	2	4	3	1	1	3	0
114	0	4	2	3	4	0	1	1	3	2
115	0	4	2	3	4	0	1	2	3	1
116	0	4	2	3	4	1	1	0	3	2
117	0	4	2	3	4	1	1	2	3	0
118	0	4	2	3	4	2	1	0	3	1
119	0	4	2	3	4	2	1	1	3	0

Table III: Detailed breakdown of all Case 1 priority schemes

Appendix B

LSIL Fab

B.1 4Rocket Process Flow

B.2 3Rocket Process Flow

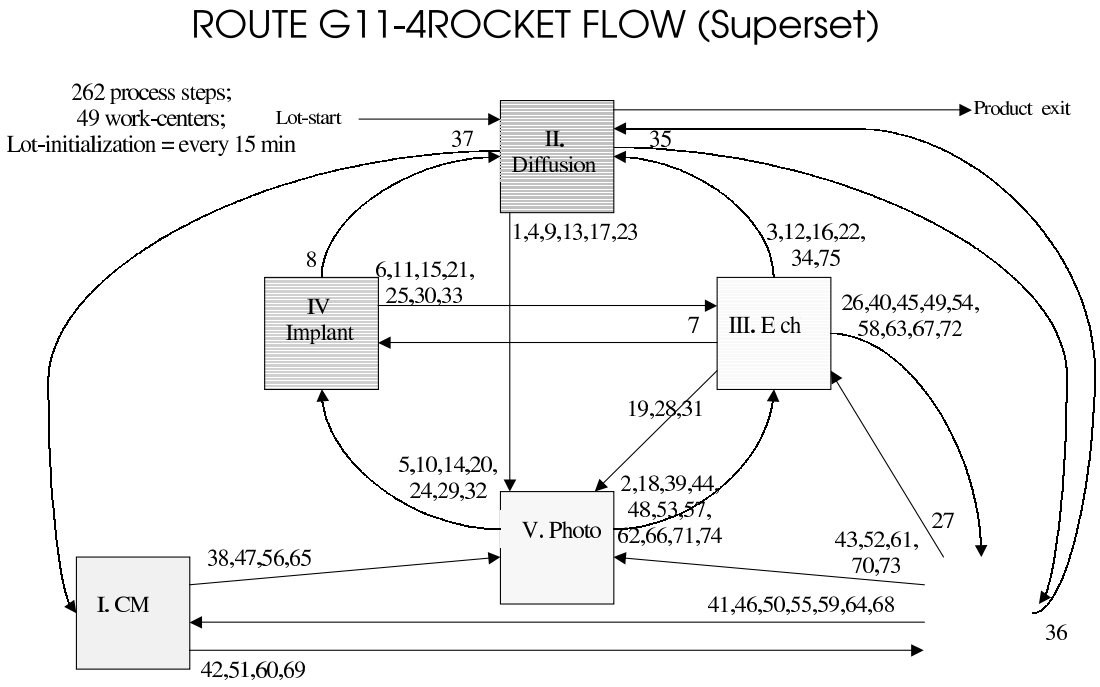

Figure B-1: General Overview of 4Rocket Process Flow

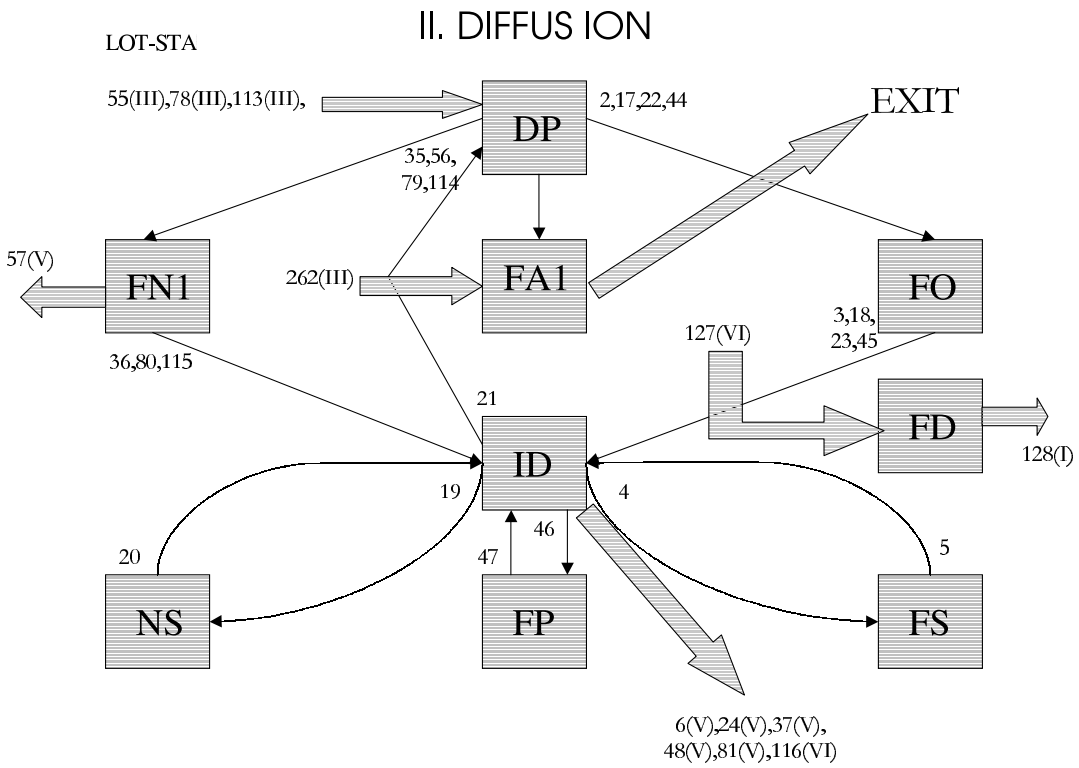
II. DIFFUS ION
LOT-STA
55(III),78(III),113(III),
DP
2,17,22,44
EXIT
35,56,
79,114
57(V)
FN1
262(III)
FA1
FO
3,18,
23,45
127(VI)
36,80,115
21
ID
FD
128(I)
19
4
46
20
47
5
NS
FP
FS
6(V),24(V),37(V),
48(V),81(V),116(VI)

Figure B-2:

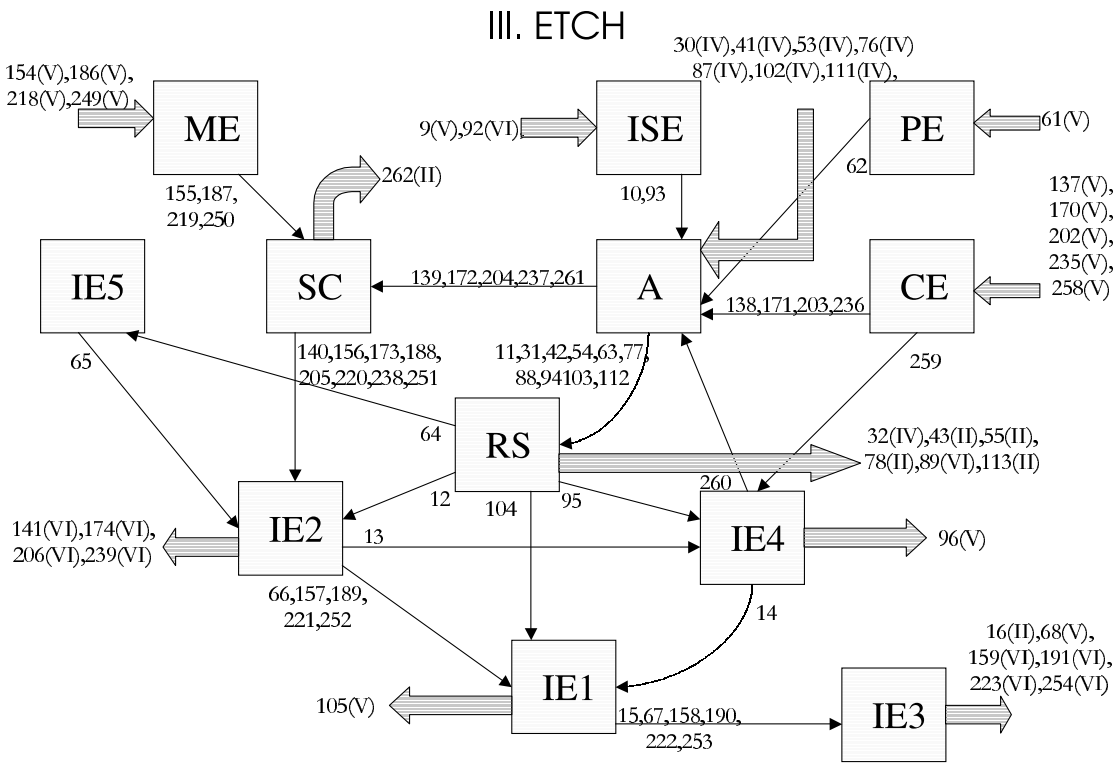

Figure B-3: Service levels for Case 1, with all possible priority schemes at WS 2

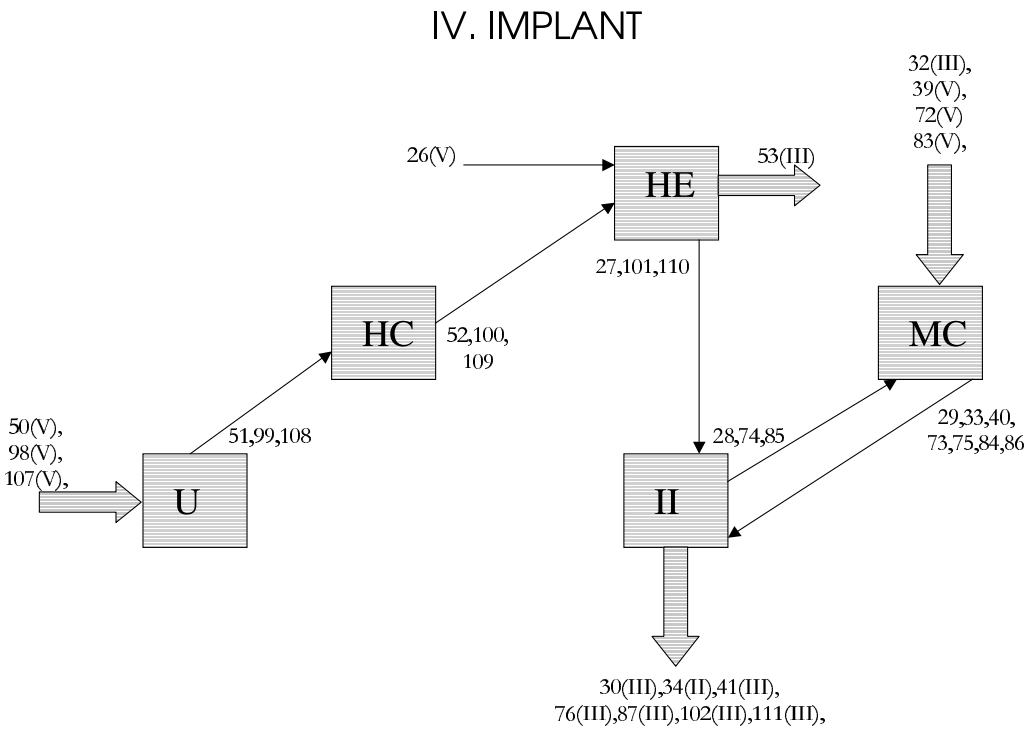

Figure B-4: Service levels for Case 1, with all possible priority schemes at WS 2

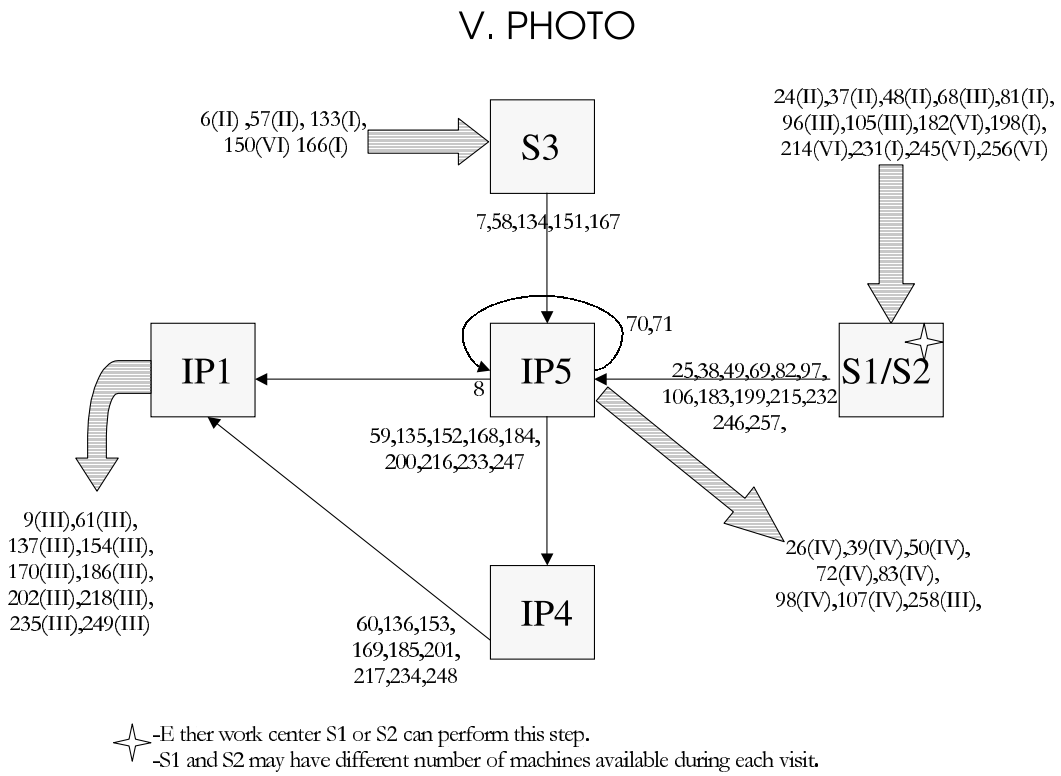

Figure B-5: Service levels for Case 1, with all possible priority schemes at WS 2

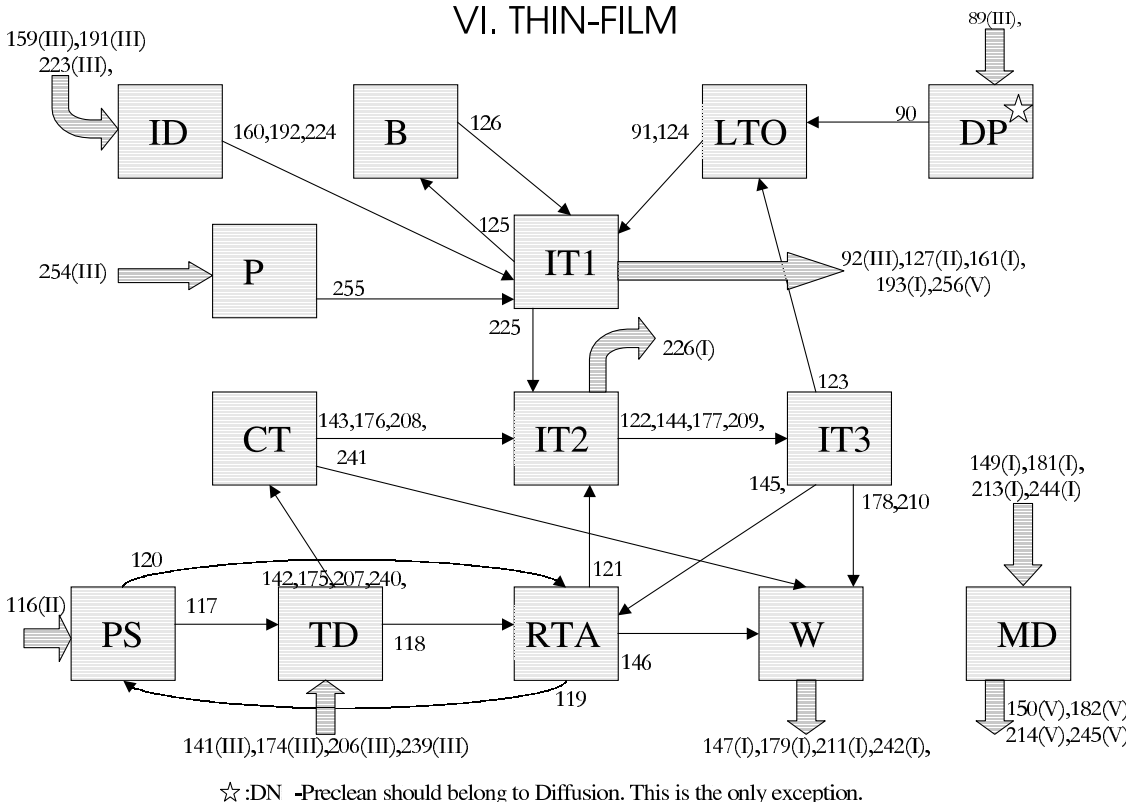

Figure B-6: Service levels for Case 1, with all possible priority schemes at WS 2

Process Flow for 4Rocket Lots

4Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
1	W_DNS_PreIn	10	0
2	DNS_Preclean	9	1.605
3	W_SVG_Furn_Oxidation2	18	0
4	Furnace_Oxidation2	17	5.000
5	Inspect_Diff2	29	0.117
6	W_SVG_Furn_SiN2	22	0
7	Furnace_SiN2	21	7.700
8	Inspect_Diff2	29	0.117
9	W_Nikon_Deep_UV	72	0
10	Stepper_3	71	0.776
11	Inspect_Photo5	38	0.167
12	Inspect_Photo1	36	0.333
13	W_IS_4520	43	0
14	Island_Spacer_Etch	42	0.678
15	D_Matrix_1188	2	0
16	Asher2	1	0.558
17	D_STI_Resist_Strip	63	0
18	Resist_Strip	62	0.472
19	Inspect_Etch2	31	0.333
20	Inspect_Etch4	32	0.333
21	Inspect_Etch1	30	0.167
22	W_DNS_PreIn	10	0
23	DNS_Preclean	9	1.593
24	W_SVG_Furn_Oxidation2	18	0
25	Furnace_Oxidation2	17	7.333
26	Inspect_Diff2	29	0.117
27	W_DNS_Nitride	53	0
28	Nitride_Strip	52	3.050
29	Inspect_Diff2	29	0.117
30	W_DNS_PreIn	10	0
31	DNS_Preclean	9	1.600
32	W_SVG_Furn_Oxidation2	18	0
33	Furnace_Oxidation2	17	4.700
34	Inspect_Diff2	29	0.117
35	W_Nikon_I14	70	0
36	Stepper_2	69	0.488
37	Inspect_Photo5	38	0.167
38	W_Eaton_HE	24	0
39	High_Energy	23	1.559
40	Inspect_Imp	35	0.250
41	W_Varian	47	0
42	Medium_Current	46	0.212
43	Inspect_Imp	35	0.250
44	D_Matrix_1188	2	0
45	Asher2	1	0.558
46	D_STI_Resist_Strip	63	0
47	Resist_Strip	62	0.805
48	W_Varian	47	0
49	Medium_Current	46	0.212
50	Inspect_Imp	35	0.250
51	W_DNS_PreIn	10	0
52	DNS_Preclean	9	1.232
53	W_STI_Furn_Anneal1	14	0
54	Furnace_Anneal1	13	7.108
55	Inspect_Diff2	29	0.117
56	W_Nikon_I14	70	0
57	Stepper_2	69	0.428
58	Inspect_Photo5	38	0.167
59	W_Varian	47	0
60	Medium_Current	46	0.212
61	Inspect_Imp	35	0.250
62	D_Matrix_1188	2	0
63	Asher2	1	0.558
64	D_STI_Resist_Strip	63	0
65	Resist_Strip	62	0.805
66	W_DNS_PreIn	10	0
67	DNS_Preclean	9	1.767
68	W_SVG_Furn_Oxidation2	18	0
69	Furnace_Oxidation2	17	5.017

4Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
70	Inspect_Diff2	29	0.117
71	W_SVG_Furn_Poly2	20	0
72	Furnace_Poly2	19	5.967
73	Inspect_Diff2	29	0.117
74	W_Nikon_I14	70	0
75	Stepper_2	69	0.428
76	Inspect_Photo5	38	0.167
77	S_UV_Cure	78	0
78	UV_Cure	77	0.302
79	W_Eaton_HE	24	0
80	High_Energy	23	0.715
81	D_Matrix_1188	2	0
82	Asher2	1	1.054
83	D_STI_Resist_Strip	63	0
84	Resist_Strip	62	0.867
85	W_DNS_PreIn	10	0
86	DNS_Preclean	9	1.383
87	W_STI_Furn_Anneal1	14	0
88	Furnace_Anneal1	13	3.525
89	W_Nikon_Deep_UV	72	0
90	Stepper_3	71	0.776
91	Inspect_Photo5	38	0.167
92	Inspect_Photo4	37	0.183
93	Inspect_Photo1	36	0.333
94	W_Poly_9400	59	0
95	Poly_Etch	58	1.337
96	D_Matrix_1188	2	0
97	Asher2	1	0.558
98	D_STI_Resist_Strip	63	0
99	Resist_Strip	62	0.867
100	Inspect_Etch5	33	0.167
101	Inspect_Etch2	31	0.333
102	Inspect_Etch1	30	0.167
103	W_Nikon_I14	70	0
104	Stepper_2	69	0.428
105	Inspect_Photo5	38	0.167
106	Inspect_Photo5	38	0.167
107	Inspect_Photo5	38	0.167
108	W_Varian	47	0
109	Medium_Current	46	0.212
110	Inspect_Imp	35	0.250
111	W_Varian	47	0
112	Medium_Current	46	0.604
113	Inspect_Imp	35	0.250
114	D_Matrix_1188	2	0
115	Asher2	1	0.558
116	D_STI_Resist_Strip	63	0
117	Resist_Strip	62	0.867
118	W_DNS_PreIn	10	0
119	DNS_Preclean	9	1.232
120	W_STI_Furn_Anneal1	14	0
121	Furnace_Anneal1	13	4.475
122	Inspect_Diff2	29	0.117
123	W_Nikon_I14	70	0
124	Stepper_2	69	0.428
125	Inspect_Photo5	38	0.167
126	W_Varian	47	0
127	Medium_Current	46	0.604
128	Inspect_Imp	35	0.250
129	W_Varian	47	0
130	Medium_Current	46	0.212
131	Inspect_Imp	35	0.250
132	D_Matrix_1188	2	0
133	Asher2	1	0.558
134	D_STI_Resist_Strip	63	0
135	Resist_Strip	62	0.867
136	W_DNS_PreIn	10	0
137	DNS_Preclean	9	1.232
138	BW_Nov_LTO	45	0

4Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
139	LTO	44	0.319
140	Inspect_TF1	39	0.117
141	W_IS_4520	43	0
142	Island_Spacer_Etch	42	0.462
143	D_Matrix_1188	2	0
144	Asher2	1	0.558
145	D_STI_Resist_Strip	63	0
146	Resist_Strip	62	0.867
147	Inspect_Etch4	32	0.117
148	W_Nikon_I14	70	0
149	Stepper_2	69	0.428
150	Inspect_Photo5	38	0.167
151	S_UV_Cure	78	0
152	UV_Cure	77	0.302
153	W_Eaton_HE	24	0
154	High_Energy	23	0.703
155	Inspect_Imp	35	0.250
156	D_Matrix_1188	2	0
157	Asher2	1	1.054
158	D_STI_Resist_Strip	63	0
159	Resist_Strip	62	0.472
160	Inspect_Etch1	30	0.083
161	W_Nikon_I14	70	0
162	Stepper_2	69	0.428
163	Inspect_Photo5	38	0.167
164	S_UV_Cure	78	0
165	UV_Cure	77	0.302
166	W_Eaton_HE	24	0
167	High_Energy	23	0.708
168	Inspect_Imp	35	0.250
169	D_Matrix_1188	2	0
170	Asher2	1	1.054
171	D_STI_Resist_Strip	63	0
172	Resist_Strip	62	0.472
173	W_DNS_PreIn	10	0
174	DNS_Preclean	9	1.232
175	W_STI_Furn_Anneal1	14	0
176	Furnace_Anneal1	13	4.975
177	Inspect_Diff2	29	0.117
178	D_DNS_Pre_Metal	61	0
179	PreMetal_Sink	60	0.842
180	Kit_Chg_AMAT_TiN	74	0
181	Ti_Dep	73	0.509
182	W_DNS_RTA	65	0
183	RTA	64	0.909
184	D_DNS_Pre_Metal	61	0
185	PreMetal_Sink	60	1.148
186	W_DNS_RTA	65	0
187	RTA	64	0.969
188	Inspect_TF2	40	0.167
189	Inspect_TF3	41	0.333
190	BW_Nov_LTO	45	0
191	LTO	44	0.315
192	Inspect_TF1	39	0.117
193	W_QTI	4	0
194	BPSG	3	0.702
195	Inspect_TF1	39	0.117
196	W_STI_Furn_Dens1	16	0
197	Furnace_Dens1	15	3.858
198	Inspect_CMP4	27	0.200
199	W_IPEC	55	0
200	Oxide_CMP	54	1.073
201	Inspect_CMP4	27	0.083
202	Inspect_CMP5	28	0.083
203	D_STI_Solvent	67	0
204	Sort_CMP	68	0.167
205	W_Nikon_Deep_UV	72	0
206	Stepper_3	71	0.776
207	Inspect_Photo5	38	0.167

Note: Workstations with zero processing times are virtual workstations that represent preventive maintenance

Table I: 4Rocket Process Steps, in sequence

Process Flow for 4Rocket Lots

4Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
208	Inspect_Photo4	37	0.183
209	Inspect_Photo1	36	0.333
210	W_CVP_4520	8	0
211	CVP_Etch	7	0.479
212	D_Matrix_1188	2	0
213	Asher2	1	0.529
214	Solvent_Clean	66	0.506
215	Inspect_Etch2	31	0.333
216	Kit_Chg_AMAT_TiN	74	0
217	Ti_Dep	73	0.562
218	BW_Nov_CVD	6	0
219	CVD_TiN	5	0.480
220	Inspect_TF2	40	0.167
221	Inspect_TF3	41	0.333
222	W_DNS_RTA	65	0
223	RTA	64	1.041
224	BW_Nov_W	80	0
225	W_CVD	79	0.345
226	W_IPEC_W	76	0
227	Tungsten_CMP	75	0.483
228	Inspect_CMP5	28	0.083
229	Kit_Chg_AMAT_AlCu	49	0
230	Metal_Dep	48	0.778
231	W_Nikon_Deep_UV	72	0
232	Stepper_3	71	0.776
233	Inspect_Photo5	38	0.167
234	Inspect_Photo4	37	0.183
235	Inspect_Photo1	36	0.333
236	W_9600	51	0
237	Metal_Etch	50	0.789
238	Solvent_Clean	66	0.272
239	Inspect_Etch2	31	0.333
240	Inspect_Etch1	30	0.167
241	BW_Nov_IMD	26	0
242	IMD_Dep	25	1.137
243	Inspect_TF1	39	0.333
244	Inspect_CMP4	27	0.200
245	W_IPEC	55	0
246	Oxide_CMP	54	1.353
247	Inspect_CMP4	27	0.083
248	Inspect_CMP5	28	0.083
249	D_STI_Solvent	67	0
250	Sort_CMP	68	0.167
251	W_Nikon_Deep_UV	72	0
252	Stepper_3	71	0.776
253	Inspect_Photo5	38	0.167
254	Inspect_Photo4	37	0.183
255	Inspect_Photo1	36	0.333
256	W_CVP_4520	8	0
257	CVP_Etch	7	0.443
258	D_Matrix_1188	2	0
259	Asher2	1	0.558
260	Solvent_Clean	66	0.506
261	Inspect_Etch2	31	0.333
262	Kit_Chg_AMAT_TiN	74	0
263	Ti_Dep	73	0.546
264	BW_Nov_CVD	6	0
265	CVD_TiN	5	0.480
266	Inspect_TF2	40	0.167
267	Inspect_TF3	41	0.333
268	BW_Nov_W	80	0
269	W_CVD	79	0.345
270	W_IPEC_W	76	0
271	Tungsten_CMP	75	0.483
272	Inspect_CMP5	28	0.083
273	Kit_Chg_AMAT_AlCu	49	0
274	Metal_Dep	48	0.778
275	W_Nikon_I14	70	0
276	Stepper_2	69	0.807

4Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
277	Inspect_Photo5	38	0.167
278	Inspect_Photo4	37	0.183
279	Inspect_Photo1	36	0.333
280	W_9600	51	0
281	Metal_Etch	50	0.789
282	Solvent_Clean	66	0.272
283	Inspect_Etch2	31	0.333
284	Inspect_Etch1	30	0.167
285	BW_Nov_IMD	26	0
286	IMD_Dep	25	1.137
287	Inspect_TF1	39	0.117
288	Inspect_CMP4	27	0.200
289	W_IPEC	55	0
290	Oxide_CMP	54	1.353
291	Inspect_CMP4	27	0.083
292	Inspect_CMP5	28	0.083
293	D_STI_Solvent	67	0
294	Sort_CMP	68	0.167
295	W_Nikon_I14	70	0
296	Stepper_2	69	0.767
297	Inspect_Photo5	38	0.167
298	Inspect_Photo4	37	0.183
299	Inspect_Photo1	36	0.333
300	W_CVP_4520	8	0
301	CVP_Etch	7	0.527
302	D_Matrix_1188	2	0
303	Asher2	1	0.558
304	Solvent_Clean	66	0.506
305	Inspect_Etch2	31	0.333
306	Kit_Chg_AMAT_TiN	74	0
307	Ti_Dep	73	0.546
308	BW_Nov_CVD	6	0
309	CVD_TiN	5	0.480
310	Inspect_TF2	40	0.167
311	Inspect_TF3	41	0.333
312	BW_Nov_W	80	0
313	W_CVD	79	0.345
314	W_IPEC_W	76	0
315	Tungsten_CMP	75	0.483
316	Inspect_CMP5	28	0.083
317	Kit_Chg_AMAT_AlCu	49	0
318	Metal_Dep	48	0.778
319	W_Nikon_I14	70	0
320	Stepper_2	69	0.807
321	Inspect_Photo5	38	0.167
322	Inspect_Photo4	37	0.183
323	Inspect_Photo1	36	0.333
324	W_9600	51	0
325	Metal_Etch	50	0.789
326	Solvent_Clean	66	0.272
327	Inspect_Etch2	31	0.333
328	Inspect_Etch1	30	0.167
329	BW_Nov_IMD	26	0
330	IMD_Dep	25	1.137
331	Inspect_TF1	39	0.117
332	Inspect_TF2	40	0.167
333	Inspect_CMP4	27	0.200
334	W_IPEC	55	0
335	Oxide_CMP	54	1.353
336	Inspect_CMP4	27	0.083
337	Inspect_CMP5	28	0.083
338	D_STI_Solvent	67	0
339	Sort_CMP	68	0.167
340	W_Nikon_I14	70	0
341	Stepper_2	69	0.767
342	Inspect_Photo5	38	0.167
343	Inspect_Photo4	37	0.183
344	Inspect_Photo1	36	0.333
345	W_CVP_4520	8	0

4Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
346	CVP_Etch	7	0.527
347	D_Matrix_1188	2	0
348	Asher2	1	0.558
349	Solvent_Clean	66	0.506
350	Inspect_Etch2	31	0.333
351	Kit_Chg_AMAT_TiN	74	0
352	Ti_Dep	73	0.546
353	BW_Nov_CVD	6	0
354	CVD_TiN	5	0.480
355	BW_Nov_W	80	0
356	W_CVD	79	0.345
357	W_IPEC_W	76	0
358	Tungsten_CMP	75	0.483
359	Inspect_CMP5	28	0.083
360	Kit_Chg_AMAT_AlCu	49	0
361	Metal_Dep	48	0.846
362	W_Nikon_I14	70	0
363	Stepper_2	69	0.671
364	Inspect_Photo5	38	0.167
365	Inspect_Photo4	37	0.183
366	Inspect_Photo1	36	0.333
367	W_9600	51	0
368	Metal_Etch	50	0.789
369	Solvent_Clean	66	0.272
370	Inspect_Etch2	31	0.333
371	Inspect_Etch1	30	0.167
372	BW_Nov_Pass	57	0
373	Passivation	56	0.423
374	Inspect_TF1	39	0.117
375	W_Nikon_I14	70	0
376	Stepper_2	69	0.508
377	Inspect_Photo5	38	0.167
378	W_CVP_4520	8	0
379	CVP_Etch	7	0.831
380	Inspect_Etch4	32	0.333
381	D_Matrix_1188	2	0
382	Asher2	1	0.313
383	Solvent_Clean	66	0.506
384	W_STI_Furn_Alloy1	12	0
385	Furnace_Alloy1	11	2.693

Note: Workstations with zero processing times are virtual workstations that represent preventive maintenance

Table II: 4Rocket Process Steps, in sequence (continue)

Process Flow for 3Rocket Lots

3Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
1	W_DNS_PreCln	10	0
2	DNS_Preclean	9	1.605
3	W_SVG_Furn_Oxidation2	18	0
4	Furnace_Oxidation2	17	5.000
5	Inspect_Diff2	29	0.117
6	W_SVG_Furn_SiN2	22	0
7	Furnace_SiN2	21	7.700
8	Inspect_Diff2	29	0.117
9	W_Nikon_I14	70	0
10	Stepper_2	69	0.584
11	Inspect_Photo5	38	0.167
12	Inspect_Photo1	36	0.333
13	W_IS_4520	43	0
14	Island_Spacer_Etch	42	0.606
15	D_Matrix_1188	2	0
16	Asher2	1	0.558
17	D_STI_Resist_Strip	63	0
18	Resist_Strip	62	0.472
19	Inspect_Etch2	31	0.333
20	Inspect_Etch4	32	0.333
21	Inspect_Etch1	30	0.167
22	W_DNS_PreCln	10	0
23	DNS_Preclean	9	1.593
24	W_SVG_Furn_Oxidation2	18	0
25	Furnace_Oxidation2	17	7.333
26	Inspect_Diff2	29	0.117
27	W_DNS_Nitride	53	0
28	Nitride_Strip	52	3.033
29	Inspect_Diff2	29	0.117
30	W_DNS_PreCln	10	0
31	DNS_Preclean	9	1.608
32	W_SVG_Furn_Oxidation2	18	0
33	Furnace_Oxidation2	17	4.700
34	Inspect_Diff2	29	0.117
35	W_Nikon_I14	70	0
36	Stepper_2	69	0.488
37	Inspect_Photo5	38	0.167
38	W_Eaton_HE	24	0
39	High_Energy	23	0.093
40	Inspect_Imp	35	0.250
41	D_Matrix_1188	2	0
42	Asher2	1	0.558
43	D_STI_Resist_Strip	63	0
44	Resist_Strip	62	0.805
45	W_Varian	47	0
46	Medium_Current	46	0.212
47	Inspect_Imp	35	0.250
48	W_DNS_PreCln	10	0
49	DNS_Preclean	9	1.232
50	W_STI_Furn_Anneal1	14	0
51	Furnace_Anneal1	13	7.108
52	Inspect_Diff2	29	0.117
53	W_Varian	47	0
54	Medium_Current	46	0.212
55	Inspect_Imp	35	0.250
56	W_Nikon_I14	70	0
57	Stepper_2	69	0.427
58	Inspect_Photo5	38	0.167
59	W_Varian	47	0
60	Medium_Current	46	0.212
61	Inspect_Imp	35	0.250
62	D_Matrix_1188	2	0
63	Asher2	1	0.558
64	D_STI_Resist_Strip	63	0
65	Resist_Strip	62	0.805
66	W_DNS_PreCln	10	0
67	DNS_Preclean	9	1.767
68	W_SVG_Furn_Oxidation2	18	0
69	Furnace_Oxidation2	17	5.017

3Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
70	Inspect_Diff2	29	0.117
71	W_SVG_Furn_Poly2	20	0
72	Furnace_Poly2	19	5.967
73	Inspect_Diff2	29	0.117
74	W_Eaton_HE	24	0
75	High_Energy	23	0.267
76	W_DNS_PreCln	10	0
77	DNS_Preclean	9	1.383
78	W_STI_Furn_Anneal1	14	0
79	Furnace_Anneal1	13	3.525
80	W_Nikon_I14	70	0
81	Stepper_2	69	0.680
82	Inspect_Photo1	36	0.333
83	Inspect_Photo5	38	0.167
84	Inspect_Photo4	37	0.183
85	Inspect_Photo1	36	0.333
86	W_Poly_9400	59	0
87	Poly_Etch	58	1.337
88	D_Matrix_1188	2	0
89	Asher2	1	0.558
90	D_STI_Resist_Strip	63	0
91	Resist_Strip	62	0.866
92	Inspect_Etch6	34	0.333
93	Inspect_Etch2	31	0.333
94	Inspect_Etch1	30	0.167
95	W_Nikon_I14	70	0
96	Stepper_2	69	0.428
97	Inspect_Photo5	38	0.167
98	Inspect_Photo4	37	0.183
99	W_Varian	47	0
100	Medium_Current	46	0.212
101	Inspect_Imp	35	0.250
102	D_Matrix_1188	2	0
103	Asher2	1	1.458
104	D_STI_Resist_Strip	63	0
105	Resist_Strip	62	0.805
106	W_DNS_PreCln	10	0
107	DNS_Preclean	9	1.232
108	W_STI_Furn_Anneal1	14	0
109	Furnace_Anneal1	13	4.475
110	Inspect_Diff2	29	0.117
111	W_Nikon_I14	70	0
112	Stepper_2	69	0.428
113	Inspect_Photo5	38	0.167
114	W_Varian	47	0
115	Medium_Current	46	0.212
116	Inspect_Imp	35	0.250
117	W_Varian	47	0
118	Medium_Current	46	0.212
119	Inspect_Imp	35	0.250
120	D_Matrix_1188	2	0
121	Asher2	1	1.458
122	D_STI_Resist_Strip	63	0
123	Resist_Strip	62	0.805
124	W_DNS_PreCln	10	0
125	DNS_Preclean	9	1.232
126	BW_Nov_LTO	45	0
127	LTO	44	0.335
128	Inspect_TF1	39	0.117
129	W_IS_4520	43	0
130	Island_Spacer_Etch	42	0.466
131	Inspect_Etch4	32	0.117
132	W_Nikon_I14	70	0
133	Stepper_2	69	0.428
134	Inspect_Photo5	38	0.167
135	Inspect_Photo5	38	0.167
136	Inspect_Photo5	38	0.167
137	Inspect_Photo4	37	0.183
138	Inspect_Photo4	37	0.183

3Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
139	Inspect_Photo4	37	0.183
140	S_UV_Cure	78	0
141	UV_Cure	77	0.302
142	W_Eaton_HE	24	0
143	High_Energy	23	0.278
144	Inspect_Imp	35	0.250
145	D_Matrix_1188	2	0
146	Asher2	1	1.054
147	D_STI_Resist_Strip	63	0
148	Resist_Strip	62	0.472
149	Inspect_Etch6	34	0.083
150	W_Nikon_I14	70	0
151	Stepper_2	69	0.428
152	Inspect_Photo5	38	0.167
153	Inspect_Photo4	37	0.183
154	S_UV_Cure	78	0
155	UV_Cure	77	0.302
156	W_Eaton_HE	24	0
157	High_Energy	23	0.287
158	Inspect_Imp	35	0.250
159	D_Matrix_1188	2	0
160	Asher2	1	1.054
161	D_STI_Resist_Strip	63	0
162	Resist_Strip	62	0.472
163	Inspect_Etch6	34	0.083
164	W_DNS_PreCln	10	0
165	DNS_Preclean	9	1.232
166	W_STI_Furn_Anneal1	14	0
167	Furnace_Anneal1	13	4.975
168	Inspect_Diff2	29	0.117
169	D_DNS_Pre_Metal	61	0
170	PreMetal_Sink	60	0.208
171	Kit_Chg_AMAT_TiN	74	0
172	Ti_Dep	73	0.596
173	W_DNS_RTA	65	0
174	RTA	64	0.885
175	D_DNS_Pre_Metal	61	0
176	PreMetal_Sink	60	0.515
177	W_DNS_RTA	65	0
178	RTA	64	0.969
179	Inspect_TF2	40	0.167
180	Inspect_TF3	41	0.333
181	BW_Nov_LTO	45	0
182	LTO	44	0.315
183	Inspect_TF1	39	0.117
184	W_QTI	4	0
185	BPSG	3	0.702
186	Inspect_TF1	39	0.117
187	W_STI_Furn_Dens1	16	0
188	Furnace_Dens1	15	3.858
189	Inspect_CMP4	27	0.200
190	W_IPEC	55	0
191	Oxide_CMP	54	0.679
192	Inspect_CMP4	27	0.083
193	Inspect_CMP5	28	0.083
194	Sort_CMP	68	0.167
195	W_Nikon_I14	70	0
196	Stepper_2	69	0.687
197	Inspect_Photo5	38	0.167
198	Inspect_Photo4	37	0.183
199	Inspect_Photo1	36	0.333
200	W_CVP_4520	8	0
201	CVP_Etch	7	0.508
202	D_Matrix_1188	2	0
203	Asher2	1	0.263
204	D_STI_Solvent	67	0
205	Solvent_Clean	66	0.901
206	Inspect_Etch2	31	0.333
207	D_DNS_Pre_Metal	61	0

Note: Workstations with zero processing times are virtual workstations that represent preventive maintenance

Table III: 4Rocket Process Steps, in sequence

Process Flow for 3Rocket Lots

3Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
208	PreMetal_Sink	60	0.532
209	Kit_Chg_AMAT_TiN	74	0
210	Ti_Dep	73	0.774
211	Inspect_TF2	40	0.167
212	Inspect_TF3	41	0.333
213	W_DNS_RTA	65	0
214	RTA	64	1.041
215	BW_Nov_W	80	0
216	W_CVD	79	0.378
217	W_IPEC_W	76	0
218	Tungsten_CMP	75	0.644
219	Inspect_CMP5	28	0.083
220	Kit_Chg_AMAT_AlCu	49	0
221	Metal_Dep	48	0.818
222	W_Nikon_I14	70	0
223	Stepper_2	69	0.815
224	Inspect_Photo5	38	0.167
225	Inspect_Photo4	37	0.183
226	Inspect_Photo1	36	0.333
227	W_9600	51	0
228	Metal_Etch	50	0.809
229	D_STI_Solvent	67	0
230	Solvent_Clean	66	0.588
231	Inspect_Etch2	31	0.333
232	Inspect_Etch1	30	0.167
233	BW_Nov_IMD	26	0
234	IMD_Dep	25	1.137
235	Inspect_TF1	39	0.333
236	Inspect_CMP4	27	0.200
237	W_IPEC	55	0
238	Oxide_CMP	54	0.896
239	Inspect_CMP4	27	0.083
240	Inspect_CMP5	28	0.083
241	Sort_CMP	68	0.167
242	W_Nikon_I14	70	0
243	Stepper_2	69	0.767
244	Inspect_Photo5	38	0.167
245	Inspect_Photo4	37	0.183
246	Inspect_Photo1	36	0.333
247	W_CVP_4520	8	0
248	CVP_Etch	7	0.479
249	D_Matrix_1188	2	0
250	Asher2	1	0.279
251	D_STI_Solvent	67	0
252	Solvent_Clean	66	0.901
253	Inspect_Etch2	31	0.333
254	Kit_Chg_AMAT_TiN	74	0
255	Ti_Dep	73	0.715
256	BW_Nov_W	80	0
257	W_CVD	79	0.378
258	W_IPEC_W	76	0
259	Tungsten_CMP	75	0.644
260	Inspect_CMP5	28	0.083
261	Kit_Chg_AMAT_AlCu	49	0
262	Metal_Dep	48	0.818
263	W_Nikon_I14	70	0
264	Stepper_2	69	0.815
265	Inspect_Photo5	38	0.167
266	Inspect_Photo4	37	0.183
267	Inspect_Photo1	36	0.333
268	W_9600	51	0
269	Metal_Etch	50	0.809
270	D_STI_Solvent	67	0
271	Solvent_Clean	66	0.588
272	Inspect_Etch2	31	0.333
273	Inspect_Etch1	30	0.167
274	BW_Nov_IMD	26	0
275	IMD_Dep	25	1.137
276	Inspect_TF1	39	0.117

3Rocket Process Flow			
Step	Workstatioin Name	WS	MPT
277	Inspect_CMP4	27	0.200
278	W_IPEC	55	0
279	Oxide_CMP	54	0.896
280	Inspect_CMP4	27	0.083
281	Inspect_CMP5	28	0.083
282	Sort_CMP	68	0.167
283	W_Nikon_I14	70	0
284	Stepper_2	69	0.767
285	Inspect_Photo5	38	0.167
286	Inspect_Photo4	37	0.183
287	Inspect_Photo1	36	0.333
288	W_CVP_4520	8	0
289	CVP_Etch	7	0.479
290	D_Matrix_1188	2	0
291	Asher2	1	0.279
292	D_STI_Solvent	67	0
293	Solvent_Clean	66	0.901
294	Inspect_Etch2	31	0.333
295	Kit_Chg_AMAT_TiN	74	0
296	Ti_Dep	73	0.715
297	Inspect_TF2	40	0.167
298	Inspect_TF3	41	0.333
299	BW_Nov_W	80	0
300	W_CVD	79	0.378
301	W_IPEC_W	76	0
302	Tungsten_CMP	75	0.644
303	Inspect_CMP5	28	0.083
304	Kit_Chg_AMAT_AlCu	49	0
305	Metal_Dep	48	0.818
306	W_Nikon_I14	70	0
307	Stepper_2	69	0.815
308	Inspect_Photo5	38	0.167
309	Inspect_Photo4	37	0.183
310	Inspect_Photo1	36	0.333
311	W_9600	51	0
312	Metal_Etch	50	0.809
313	D_STI_Solvent	67	0
314	Solvent_Clean	66	0.588
315	Inspect_Etch2	31	0.333
316	Inspect_Etch1	30	0.167
317	BW_Nov_Pass	57	0
318	Passivation	56	0.423
319	Inspect_TF1	39	0.117
320	W_Nikon_I14	70	0
321	Stepper_2	69	0.508
322	Inspect_Photo5	38	0.167
323	W_CVP_4520	8	0
324	CVP_Etch	7	0.831
325	Inspect_Etch6	34	0.333
326	D_Matrix_1188	2	0
327	Asher2	1	0.625
328	D_STI_Solvent	67	0
329	Solvent_Clean	66	0.901
330	W_STI_Furn_Alloy1	12	0
331	Furnace_Alloy1	11	2.693

Note: Workstations with zero processing times are virtual workstations that represent preventive maintenance

Table IV: 4Rocket Process Steps, in sequence (continue)

Bibliography

[1] Richard. E. Barlow and Frank Proschan. *Statistical Theory of Reliability and Life Testing.* Holt, Rinehart and Winston, Inc., 1975.

[2] J.A. Buzacott and L.E. Hanifin. Models of automatic transfer lines with inventory banks–a review and comparison. *AIIE Transactions*, 10(2):197–207, 1978.

[3] John A. Buzacott and J. George Shanthikumar. *Stochastic Models of Manufacturing Systems.* Prentice-Hall, Englewood Cliffs, New Jersey, 1993.

[4] S. Deshpende. A scheduling policy experiment for lean implementation. Master's thesis, Massachusetts Institute of Technology, 1999.

[5] S. B. Gershwin. *Manufacturing Systems Engineering.* Prentice Hall, Englewood Cliffs, New Jersey, 1994.

[6] S. B. Gershwin. Design and operation of manufacturing systems – the control-point policy. In *IIE Transactions on Design and Manufacturing, Special Issue on Decentralized Control of Manufacturing Systems.*, 1999.

[7] Omar. Gzouli. Comparison of scheduling policies by simulation. Master's thesis, Massachusetts Institute of Technology, 2000.

[8] W. J. Hopp and M. L. Spearman. *Factory Physics: Foundations of Manufacturing Management.* Irwin, 1996.

[9] Yeong-Dae Kim, Jung-Ug Kim, and Seung-Kil Lim. Due-date based scheduling and control policies in a multiproduct semiconductor wafer fabrication facility. *IEEE transactions on semiconductor manufacturing*, 11(1):155–164, 1998.

[10] P.R. Kumar. Scheduling semiconductor manufacturing plants. *IEEE transactions on semiconductor manufacturing*, 1994.

[11] Averill M. Law and W. David Kelton. *Simulation Modeling and Analysis.* McGraw-Hill, New York, NY, second edition, 1991.

[12] Linda and Robert Atherton. *Wafer Fabrication: Factory Performance and Analysis.* Kluwer, 1995.

[13] Steve C. H. Lu, Deepa Ramaswamy, and P.R. Kumar. Efficient scheduling policies to reduce mean and variance of cycle-time in semiconductor manufacturing plants. *IEEE transactions on semiconductor manufacturing*, 7(3):374–388, 1994.

[14] David F. Pyke, Edward A. Silver, and Rein Peterson. *Inventory Management and Production Planning and Scheduling.* John Wiley and Sons, New York NY, third edition, 1998.

[15] L. M. Wein. Scheduling semiconductor wafer fabrication. *IEEE Transactions on semiconductor manufacturing*, 1(3):115–130, 1988.

[16] Loren Werner and Meow Seen Yong. A survey of commercial production scheduling software. In *Lean Aerospace Initiative, M.I.T.*, 1999.