Protective Capacity and Time Buffer Design In Theory Of Constraints Controlled Discrete Flow Production Systems

by

Louis Louw

B.Eng (Industrial), University of Stellenbosch

Dissertation presented for the degree Doctor of Philosophy at the University of Stellenbosch.

Promoter: Dr. D.C. Page

April 2003

Stellenbosch University http://scholar.sun.ac.za

Declaration

I, the undersigned, hereby declare that the work contained in this dissertation is my own original work and has not previously in its entirety, or in part, been submitted at any university for a degree.

Signature: - Date: --

Stellenbosch University http://scholar.sun.ac.za

Synopsis

To maximise the throughput of a production system the capacity constrained resource needs to be protected from variation and uncertainty. In the Theory of Constraints philosophy such protection is provided by means of time buffers and protective capacity. Time buffers are protective time that is allowed in the production schedule to buffer against disruptions, whereas protective capacity is defined as a given amount of extra capacity at non-constraints above the system constraint's capacity.

In this research an analytical procedure was developed to more accurately determine the required time buffer lengths. This procedure uses an open queuing network modelling approach where workstations are modelled as GI/G/m queues. A simulation experiment was performed to evaluate the time buffer estimation procedure on the operations of an actual fifteen station flow shop. The results from the study suggest that the analytical procedure is sufficiently accurate to provide an initial quick estimate of the needed time buffer lengths at the design stage of the line.

This dissertation also investigated the effect of protective capacity levels at a secondary constraint resource as well as at the other non-constraint resources on the mean flow time, the bottleneck probability of the primary constraint resource, as well as the output of flow production systems using simulation models and ANOVA. Two different types of flow production systems were investigated: (1) a flow shop with a fixed number of stations and unlimited queue or buffer space between stations, and (2) an assembly line where a total work content is distributed among stations in a certain fashion and the number of stations are not fixed.

The experimental studies show that flow shop performance in the form of flow time and line output is not that much influenced by low protective capacity levels at the secondary constraint resource. Low protective capacity levels at a single station however can significantly reduce the bottleneck probability for the primary constraint resource when it is located before and relatively close or near to the primary constraint in the process flow, or after but relatively far from the primary constraint. An after-far secondary constraint location also causes slightly longer job flow times, and should therefore be avoided if possible. The research further shows that quite high protective capacity levels at the non-constraint resources are needed to ensure a more stable and therefore manageable primary constraint. However low average levels of protective capacity at non-constraint resources are sufficient to ensure that the maximum designed output level as determined by the utilisation of the primary constraint resource is obtained.

Stellenbosch University http://scholar.sun.ac.za

The results for the assembly line experiment showed that an unbalanced line configuration where less work is assigned to the non-constraint stations than to the primary constraint station (but non-constraint stations have an equal work content) can lead to significant reductions in the mean flow time while maintaining the same line output, without resulting in too many additional stations. Low protective capacity levels in the range of 2% to 5% are sufficient to cause substantial improvements in flow time without resulting in too many additional stations in the line.

Stellenbosch University http://scholar.sun.ac.za

Opsomming

Om die finale uitset van 'n produksiestelsel te maksimeer is dit noodsaaklik dat die bottelnek beskerm word teen fluktuasies en onderbrekings. In die "Theory of Constraints" filosofie word van twee soorte beskermingsmeganismes gebruik gemaak: tydbuffers en beskermende kapasiteit. Tydbuffers is beskermende tyd wat in die produksieskedule gevoeg word om sodoende die bottelnek teen onderbrekings te beskerm, terwyl beskermende kapasiteit ekstra produksiekapasiteit relatief tot die bottelnek se kapasiteit is wat by nie-bottelnekke gevoeg word.

In hierdie navorsing is 'n analitiese prosedure ontwikkel om meer akkurate berekenings van tydbuffergroottes te verkry in produksiestelsels wat volgens 'n "Theory of Constraints" filosofie bestuur word. Die prosedure maak gebruik van oop toustaan netwerk modellering waar werkstasies gemodelleer word as GI/G/m toue. Die analitiese prosedure is ge-evalueer met 'n simulasie eksperiment op 'n werklike vyftien stasie vloeiwinkel. Die resultate dui aan dat die analitiese prosedure akkuraat genoeg is om vinnig aanvanklike beramings vir die benodigde tydbuffergroottes tydens die ontwerpsfase van die produksiestelsel te verskaf.

Verder is ook ondersoek ingestel na die effek van beskermende kapasiteitsvlakke by die sekondêre bottelnek asook die ander nie-bottelnekke op die gemiddelde deurvloeityd, die totale uitset, asook die bottelnek waarskynlikheid vir die primêre bottelnek in vloei produksiestelsels deur gebruik te maak van simulasie modelle en ANOVA. Twee verskillende tipes vloei produksiestelsels is ondersoek: (1) 'n vloeiwinkel met 'n vaste aantal stasies en 'n onbeperkte buffer spasie tussen stasies, en (2) 'n monteerlyn waar 'n totale werksinhoud op 'n bepaalde wyse onder stasies verdeel moet word en die aantal stasies nie vas is nie.

Die eksperimentele studies dui aan dat die deurvloeitye en totale uitset van 'n vloeiwinkel nie noemenswaardig beïnvloed word deur lae beskermende kapasiteitsvlakke by die sekondêre bottelnek nie. Hierdie maatstawwe word meer beïnvloed word deur die gemiddelde beskermende kapasiteitsvlakke by al die nie-bottelnekke. Lae beskermende kapasiteit by 'n enkele werkstasie kan egter die bottelnek waarskynlikheid vir die primêre bottelnek aansienlik verlaag indien dit voor en relatief na aan die primêre bottelnek in die prosesvloei geleë is, of na, maar relatief ver, vanaf die primêre bottelnek. 'n Sekondêre bottelnek ligging na maar relatief ver vanaf die primêre bottelnek in die prosesvloei veroorsaak ook langer deurvloeitye, en moet dus vermy word. Verder dui die navorsing aan dat redelike hoë gemiddelde beskermende kapasiteit by nie-bottelnekke benodig word om 'n meer stabiele primêre bottelnek te verseker. In vloeiwinkels met lae vlakke van variasie en onderbrekings is egter lae gemiddelde vlakke van beskermende kapasiteit by nie-bottelnekke

Stellenbosch University http://scholar.sun.ac.za

voldoende om te verseker dat die maksimum ontwerpte uitset soos bepaal deur die benutting van die primêre bottelnek behaal word.

Die resultate vir die monteerlyn eksperiment dui aan dat 'n ongebalanseerde lynkonfigurasie waar minder werk aan die nie-bottelnek stasies as aan die primêre bottelnek stasie toegeken word (maar nie-bottelnek stasies het 'n gelyke werksinhoud), aansienlike verlagings in deurvloeityd teweeg kan bring terwyl dieselfde lyn uitset behou word. Dit is moontlik sonder te veel addisionele stasies in die lyn. Die eksperimentele resultate dui aan dat lae beskermende kapasiteitsvlakke van tussen 2% tot 5% voldoende is om beduidende verlagings in deurvloeityd teweeg te bring sonder te veel addisionele stasies.

Stellenbosch University http://scholar.sun.ac.za

Acknowledgements

I would like to express my sincere gratitude to the following people for their guidance and assistance in completing this dissertation:

- Dr. D.C. Page, my promoter, for his time, guidance and valuable feedback, and for encouraging me to do my PhD.
- Dr. J.F. Cox and Dr. J.R. Holt, the external examiners, and Prof. N. du Preez, the internal examiner, for their assistance and valuable comments.
- Mr. A.C. Nel and K. Sinclair, for their time and assistance in gathering the necessary data for the flow shop simulation model used to evaluate the time buffer estimation procedure.
- Scheduling Technologies Group (STG), UK, for allowing me to visit their head office in London and providing me with accommodation.

Finally, I would like to thank my parents, to whom this dissertation is dedicated, for their love, support and encouragement.

Stellenbosch University http://scholar.sun.ac.za

Table of Contents

Stellenbosch University http://scholar.sun.ac.za

Stellenbosch University http://scholar.sun.ac.za

Stellenbosch University http://scholar.sun.ac.za

Stellenbosch University http://scholar.sun.ac.za

List of Figures

Stellenbosch University http://scholar.sun.ac.za

Stellenbosch University http://scholar.sun.ac.za

List of Tables

Stellenbosch University http://scholar.sun.ac.za

Stellenbosch University http://scholar.sun.ac.za

Glossary

ANOVA: Analysis of Variance.

Batch Processing: A manufacturing technique in which parts are accumulated and processed together in a lot.

Bottleneck: A facility, function, department, or resource whose capacity, if properly scheduled, is less than, or exactly equal to, the demand placed upon it in a comparable period of time.

Buffer Management: A theory of constraints process in which all expediting and remedial action in a shop is driven by what is scheduled to be in the buffers (constraint, assembly, and shipping buffers). Also used to monitor and fine-tune time buffer lengths.

Capacity constraint resource: Any resource whose available capacity limits the organisation's ability to meet the product volume, product mix, or demand fluctuation required by the market place.

CCR: See Capacity constraint resource.

Constraint: Any element or factor that prevents a system from achieving a higher level of performance relative to its goal. Constraints can be physical/logistical, managerial/procedural or behavioural/psychological.

DBR: See Drum-buffer-rope.

Drum-buffer-rope: The production logistical system of the Theory of Constraints used to plan and control work flow.

Inventory: All the money invested in things purchased for sale, valued at purchase price with no value added for applied labour or allocated overhead.

Just-in-Time (JIT): A philosophy of manufacturing based on planned elimination of all waste and continuous improvement of productivity.

Kanban: Scheduling system developed and used by Toyota. A method of Just-In-Time production that uses standard containers or lot sizes with a single card (Kanban) attached to each. It is a pull

Stellenbosch University http://scholar.sun.ac.za

system in which work centres signal with a card that they wish to withdraw parts from feeding operations or suppliers.

MTBF: Mean time between failures.

MTTR: Mean time to repair.

Non-bottleneck: A facility, function, department, or resource whose capacity, if properly scheduled, is greater than the demand placed upon it in a comparable period of time.

Operating Expense: All the money spent by the system to convert inventory into throughput.

PC: Protective capacity.

Primary Constraint/CCR: The resource with the lowest capacity and therefore the most heavily utilised resource.

Process batch: The number of units made between sequential setups at a work centre.

Protective capacity: A given amount of extra capacity at non-constraints above the system constraint's capacity, used to protect against statistical fluctuation (breakdowns, late receipts of materials, quality problems, etc.). Protective capacity provides non-constraints with the ability to catch up to "protect" throughput and due date performance.

Secondary Constraint/CCR: The resource with the second lowest capacity and therefore the second most heavily utilised resource.

Theory of Constraints: A management philosophy developed by Dr. Eliyahu M. Goldratt that can be viewed as three separate but interrelated areas - logistics, performance measurement, and logical thinking. Focuses on the identification and management of constraints.

Time buffer: Time buffers represent the additional planned lead time allowed, beyond the required set-up and run times, for materials to reach a specific point in the product flow. Strategically placed, time buffers are designed to protect the system throughput from the internal disruptions that are inherent in any process.

TOC: See Theory of Constraints.

Stellenbosch University http://scholar.sun.ac.za

Throughput: The rate at which the system generates money through sales of its products or services.

Transfer batch : The quantity of an item moved between sequential work centres during production.

WIP: Work-in-process.

Chapter 1

Introduction

1.1 Purpose of the Research

This dissertation examines two of the main protective mechanisms found in the production logistical branch of the Theory of Constraints philosophy: time buffers and protective capacity. Time buffers are used to protect a production system's output from the disruptions that occur in the system. These buffers are not excess inventory that are released into the system. Instead it is planned time that is allowed in the schedule to buffer against disruptions. Protective capacity is defined as a given amount of extra capacity at non-constraints above the system constraint's capacity. Protective capacity serves to move jobs faster to the constraint operation. WIP build-up therefore naturally occurs mainly in front of the constraint operation, which helps to protect it from disruptions that could negatively impact the output for the system. Protective capacity also reduces the flow time for jobs through the system, which means that less protection in the form of time buffers are needed. There is therefore an interaction between the protective capacity levels in the system and the amount of time buffer protection needed.

The purpose of this dissertation is twofold: first to develop an analytical procedure for the design of time buffers in Theory of Constraints controlled production systems, and second to investigate the effect of protective capacity levels at a secondary constraint resource as well as at the other non-constraint resources on the mean flow time, the bottleneck probability of the primary constraint resource, as well as the output of flow production systems. This will present production managers with a tool as well as design guidelines when designing new or re-designing existing flow production systems.

The analytical procedure was developed by investigating analytical techniques in the literature for modelling production or manufacturing systems. This literature search on analytical queuing network modelling techniques is not presented in this dissertation, since it does not quite fit in with the main focus of the dissertation. The objective was not to exhaustively investigate different modelling techniques, but rather how to apply a proven technique on the analysis of time buffers. The literature on queuing modelling techniques that was studied is however listed in the bibliography. From the literature search the decision was made to use an open queuing network modelling technique based on parametric decomposition of a production network. This technique was chosen based on its relatively simple and flexible modelling approach which lends itself to be easily modified and incorporated in a computerised software tool, and experimental results from the literature that showed it to be fairly accurate compared to simulation models. The developed analytical procedure for estimating time buffer lengths was incorporated in a software tool developed in Visual Basic. The analytical procedure was also evaluated on a simulation model of an actual serial flow shop under a realistic drum-buffer-rope shop control system, which is the manufacturing logistical system of the Theory of Constraints philosophy. For this purpose a computerised drum-buffer-rope production scheduling system was developed in order to generate the material release schedules and primary constraint schedules for the evaluation experiment. The purpose of the evaluation was therefore to determine whether the analytical procedure would be usable in practice.

The effect of protective capacity levels at the secondary constraint as well as non-constraint resources was investigated using simulation models and analysis of variance (ANOVA). Two different types of flow production systems were investigated: (1) a flow shop with a fixed number of stations and unlimited queue or buffer space between stations, and (2) an assembly line where a total work content is distributed among stations in a certain fashion and the number of stations are not fixed. Previous research studies have not investigated the impact of protective capacity levels at the secondary constraint on system performance. One of the objectives therefore was to investigate whether low protective capacity levels at a single station that is not the primary constraint can have a detrimental impact on flow shop performance measures such as the mean flow time, the total output, as well as the bottleneck probability for the primary constraint resource, and how it interacts with other factors such as the location of the secondary constraint relative to the primary constraint, the variability in the system caused by station downtimes, and the amount of protective capacity at the other non-constraint stations. The bottleneck probability indicates the stability of the location of the primary constraint, which is important for planning and scheduling purposes in Theory of Constraints controlled production systems. None of the previous studies on protective capacity have focused specifically on the bottleneck probability of the primary constraint. Studies by Lawrence and Buss (1994), Craighead et al (2001) and Patterson et al (2002) have used a bottleneck shiftiness measure that is based on a bottleneck probability calculation. In all of these studies however a bottleneck(s) was calculated as the station(s) with the most jobs in queue. This is not an accurate indication of a bottleneck, since a number of jobs lying in front of a station with a small throughput rate take much longer to process than the same number of jobs lying in front of a faster station. With the bottleneck definition used in the above mentioned studies both stations would be taken as the bottleneck, whereas the slower station is in fact the true bottleneck or constraint. In this study a station flow time measure was used instead to indicate a bottleneck. With this measure the number of jobs together with the processing rate of the station is used to identify the true bottleneck. Using the bottleneck probability measure together with the flow time and total output measure give a good indication of the impact of protective capacity on some of the most important performance measures for Theory of Constraints controlled production systems. Another objective of the study then was also to investigate what levels of protective capacity would typically be required to ensure a stable primary bottleneck location in flow shops. None of the previous studies on protective capacity have also investigated the effect of different protective capacity levels on assembly type of flow lines where the number of stations are not fixed. Another objective therefore was to study the impact of different protective capacity levels on the flow time measure for assembly lines, and to investigate what levels of protective capacity would be sufficient in such lines to ensure significant flow time reductions without creating too many stations.

The diagram in Figure 1-1 presents an overview framework for this research study.

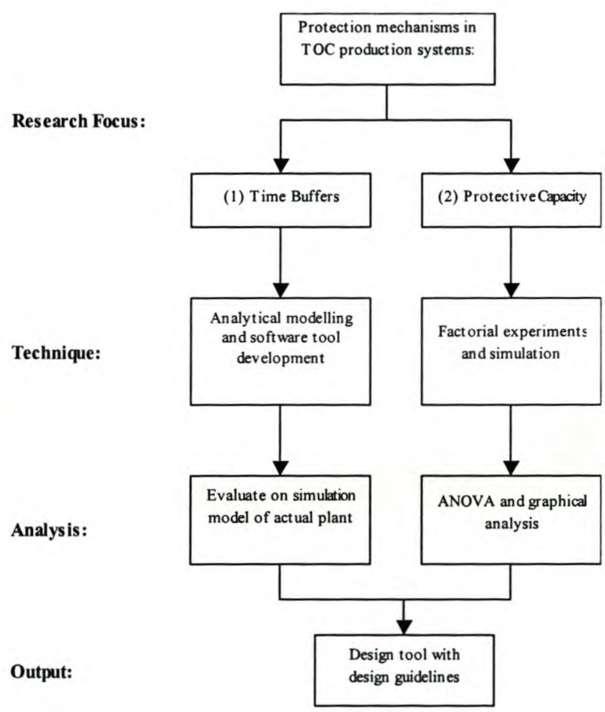

Figure 1-1. Overview framework for the study

1.2 Relevance of Research

In the last 20 years, the importance of the capacity constrained resource or bottleneck in the performance of a manufacturing company has been more widely recognised. It is this critical resource that determines the throughput for the whole plant, and therefore the ability of the organisation to achieve its goal of making more money now as well as in the future. The importance of the bottleneck in the performance of a system was especially highlighted by Eli Goldratt. Goldratt (1990, p. 53; 1997, p. 87-90) gives a useful analogy where the manufacturing process is compared to a chain. The strength of any chain is determined by its weakest link. Strengthening any of the other links will not improve the strength of the chain as a whole. Only by improving the strength of the weakest link will the chain be made stronger. The same concept applies in a manufacturing process. The process is essentially a chain of interdependent events. The weakest link is the resource with the least capacity to perform the load placed on it within a certain time frame. The pace of the manufacturing process is in fact determined by its slowest resource (the weakest link). This weakest link is the constraint of the process.

Any manufacturing process can further be characterised by two basic phenomena, namely statistical fluctuations and dependent events (Goldratt 1986, pp. 95-101). Statistical fluctuations refer to the inherent variability within every process as well as the occurrence of unforeseen events, e.g. variability in task times, station break downs, scrap, etc. Dependent events refer to the fact that certain operations or activities cannot take place until certain other operations or activities have been completed. It is the combination of these two basic phenomena that makes it difficult to achieve a smooth product flow that is in concert with market demand. Due to the combined effect of dependency and variability, disruptions in the manufacturing process will not average out. Negative variances accumulate more rapidly than positive variances, disrupting the planned product flow for the entire plant.

In order to protect the throughput of the plant, the capacity constrained resource or bottleneck therefore needs to be protected from this negative influence of statistical fluctuations and dependent events. Three possible ways of protecting the capacity constrained resource are:

1. extra inventory
2. safety lead-time
3. protective capacity

All of these methods protect the capacity constrained resource from starvation caused by disruptions at upstream workstations. Protective inventory just means placing extra inventory in the system. It is defined in the APICS Dictionary (2002, p. 95) as "The amount of inventory required relative to the protective capacity in the system to achieve a specific throughput rate at the constraint". Physical inventory however is tied to specific product types and specific stages in the production process, which makes it to difficult to determine the amount and location of protective inventory necessary. Safety lead time are time buffers or extra time allowed in the planned lead time of a product. This causes work to be released earlier, which in turn leads to longer lead times and more work-in-process inventory (due to Little's law). In this case, the inventory is however allowed to naturally accumulate in front of the critical resources. In recent years protective capacity has gained more recognition as a protection mechanism. Protective capacity is defined in the APICS dictionary as "a given amount of extra capacity at non-constraints above the system constraint's capacity" (2002, p.95). By providing capacity slack at non-constraint resources, workstations recover any lost time caused by statistical fluctuations and disruptions quickly, protecting the material flows to the constraint resource. Protective capacity also leads to shorter lead times and lower inventory, which in turn provides a competitive edge. Since the efficient design of protective capacity and safety lead time or time buffers can lead to lower WIP levels, they appear to be more effective protection mechanisms than extra inventory.

Goldratt and Fox (1986, p.36), as well as Srikanth and Umble (1997a, p.12), suggest that a firm can achieve a competitive advantage in one of three ways: (1) by having better products, (2) by providing better customer service, or (3) by being the low cost producer. The impact that lower work-in-process inventory (WIP) has on these three competitive elements stems mainly from the shorter manufacturing lead time that results from a lower WIP. In most manufacturing plants the actual processing time contribute very little to the total manufacturing lead time and the majority of lead time for materials is actually spend waiting in queues (Goldratt 1990, p.128) (Srikanth & Umble 1997a, p.12). The manufacturing lead time is therefore roughly proportional to the amount of WIP, and lower WIP-levels will cause shorter manufacturing lead times.

The effects of lower WIP levels on the competitive elements are the following:

- *Product Quality* – In a high inventory environment quality problems are obscured by high inventories and fire fighting, and there is not enough time available to fix the problem. In a low inventory environment problems can be identified and rectified more quickly and easily (Goldratt and Fox 1986, p.44) (Srikanth & Umble 1997a, p.16). Take for example a situation where a defect caused by the first operation is only detected at the last operation. In a high inventory environment where lead times are high, the first operation had probably already finished processing the order making it difficult to detect the cause of the defect. In a low inventory environment with short lead times, the first operation is however probably still busy processing the order making it easier to determine the cause of the problem. New products and product design changes can also be introduced more quickly to the market in a low inventory environment with short manufacturing lead times.
- *Customer Service* – In a low inventory environment with short manufacturing lead times, quoted delivery times can be shorter than those of competitors giving a competitive advantage. With less WIP there is also better control of the work and less opportunity for working on the wrong parts. This combined with the fact that product forecasts can be more accurate in a short lead time environment leads to better due-date performance (Goldratt and Fox 1986, p.60) (Srikanth & Umble 1997a, p.13)

- *Product Cost* – In plants with long manufacturing lead times due-date performance is often poor. This leads to constant expediting and working overtime causing operational costs to rise. Lower inventory plants also have less inventory carrying costs and losses due to obsolescence (Goldratt and Fox 1986, p.52) (Srikanth & Umble 1997a, p.17).

The fundamental relationship between WIP, lead time and the throughput potential of a production line is described by Little's Law (Hopp & Spearman 1996):

Throughput = WIP/Lead Time

This indicates the inverse relationship between WIP and lead time. Lenz (1989 cited Atwater 1991) describes this relationship in his so-called triangle of integration (refer to Figure 1-2).

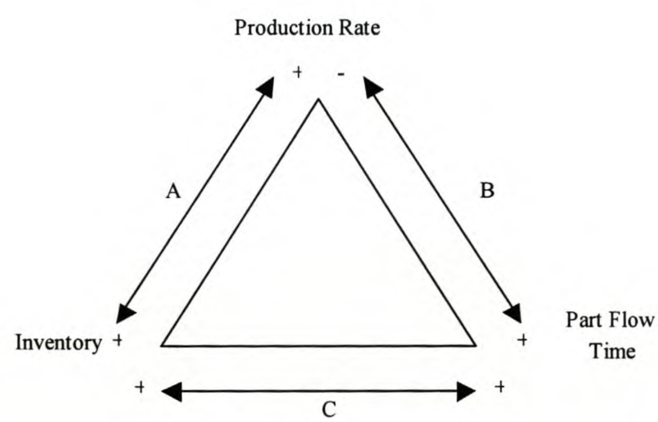

Figure 1-2. The triangle of integration

According to this triangle a three-way link exists between production output, inventory levels and part flow times:

1. A certain amount of inventory is required to achieve a specified output rate. With no inventory the output of the line is at a minimum, but as the inventory levels increase the output rate also increases until the desired output rate is reached. This relationship is represented by arrow A in Figure 1-2.
2. As part flow times decrease the output rate increase. This relationship is represented by arrow B in Figure 1-2.
3. As inventory levels increase the part flow times also increase. This is indicated by arrow C in Figure 1-2.

There is however a conflict in this three way relationship, because increasing inventory levels lead to an increase in both production rate and flow time, whereas an increase in flow time in turn leads to a decrease in output rate. Lenz explains this conflict with his so-called Work-in-Process Against Capacity (WIPAC) curve (refer to Figure 1-3). In case 1 the WIPAC curve has a constant slope and as the amount of inventory increases, the output rate of the line also increases. The relationship between flow time and WIP dictates that the flow time also increases during this part of the curve. In case 2 the WIPAC curve has a decreasing slope and although an increase in inventory still causes the output rate to increase, the increase in output rate per additional unit of WIP is less than in case 1. At this point the increase in flow times start to counter effect the increase in output rate. In case 3 the curve has constant slope and the increase in output rate from increasing inventory levels is completely offset by the increase in flow times. Adding inventory therefore no longer causes an increase in output rate. In case 4 adding inventory causes other detrimental effects such as scrap and rework that start to decrease the output rate.

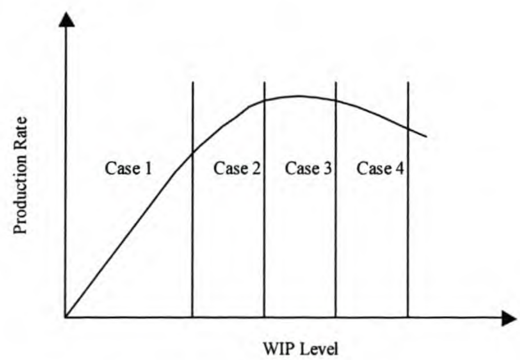

Figure 1-3. WIPAC curve

It is therefore clear that the efficient design and management of flow time (or lead time) and inventory have an important influence on a firm's competitive edge, which therefore calls for efficient techniques to design and manage the logistical flow of materials. Srikanth and Umble (1997a, p.6) also conclude "the degree to which the material and product flow of a plant can be synchronised determines the degree to which waste can be eliminated and, therefore, the degree to which the plant can be competitive in the marketplace." One business philosophy that seems to provide efficient techniques for managing material flow is the Theory of Constraints philosophy. Although time buffers and protective capacity are such an important part of this philosophy, very little research has been done on these two protection mechanisms. The previous discussion on flow times and WIP has however clearly indicated their important influence on a manufacturing plant's competitiveness. This study therefore makes an important contribution to broaden the research knowledge on time buffers and protective capacity, and to provide practitioners with more scientific based guidelines and tools for their design and implementation.

1.3 Organisation of the Dissertation

This dissertation consists of the following chapters:

- Chapter 1 – This is the introductory chapter, specifying the purpose of the research, the relevance of the research, as well as the organisation of the dissertation.
- Chapter 2 – This chapter presents the relevant literature review to explain the Theory of Constraints philosophy, and to review previous research on time buffers, balanced and unbalanced line design, and protective capacity.
- Chapter 3 – This chapter explains the developed analytical procedure for estimating time buffer lengths, and it also discusses the simulation experiment performed in order to evaluate the developed procedure.
- Chapter 4 – The experimental studies performed to investigate protective capacity in two types of discrete flow production systems (flow shops and assembly lines) are presented and discussed in this chapter. The chapter presents the research questions posed, it explains the details of how the experiments were performed, and also discusses the experimental findings.
- Chapter 5 – This chapter provides a summary and conclusions for the research study as a whole, as well as suggestions for further research.

Chapter 2

Review Of Relevant Literature

2.1 Introduction

Constraint management, as well as protective capacity and time buffer protection mechanisms, are central to the management philosophy known as the Theory of Constraints (TOC). This philosophy was developed by Eli Goldratt and includes a manufacturing logistical system known as Drum-Buffer-Rope (DBR) for the effective management of constraints and material flow. The following section describes a literature review on TOC and DBR, followed by sections examining the literature on the two protection mechanisms of TOC, i.e. time buffers and protective capacity.

2.2 The Theory Of Constraints

The Theory of Constraints philosophy was developed by the Israeli physicist Eliyahu Goldratt. Goldratt has a doctorate in physics and became involved with production scheduling through a friend who had scheduling problems at a plant that built chicken coops. Goldratt was very interested in the problem and developed a scheduling software product called OPT which was launched in 1978 (Noreen et al 1995, p.3). OPT stood for optimised production timetables and was offered in the United States by a company called Creative Output. The name was later changed to optimised production technology. In 1986 in a book called "The Race" Goldratt presented a logistical system for the material flow in a manufacturing plant called the drum-buffer-rope (DBR) system. Gradually the focus of this concept moved from the production floor to encompass all aspects of business, and by 1987 the overall concept became known as the Theory of Constraints (TOC) (Rahman 1998). This is viewed as an overall theory for running an organisation and consists of the so-called "five focusing steps" of continuous improvement (which will be discussed later in this chapter). Under this overall theory it is recognised that the main constraint in most organisations may not be physical, but rather related to management policies (Spencer 1991). In order to address such policy constraints and effectively implement a process of on-going improvement, Goldratt developed a generic approach called the "thinking process" (TP) that presents a set of tools to help identify and solve problems. These TP-tools can be seen as "the current paradigm of TOC" (Rahman 1998).

UNIVERSITEIT STELLENBOSCH
BIBLIOTEEK

Spencer and Cox (1995) present a useful schematic of the Theory of Constraints which shows the three main branches of TOC (i.e. the logistics branch, the performance system branch and the problem solving /TP branch) along with their different components (refer to Figure 2-1).

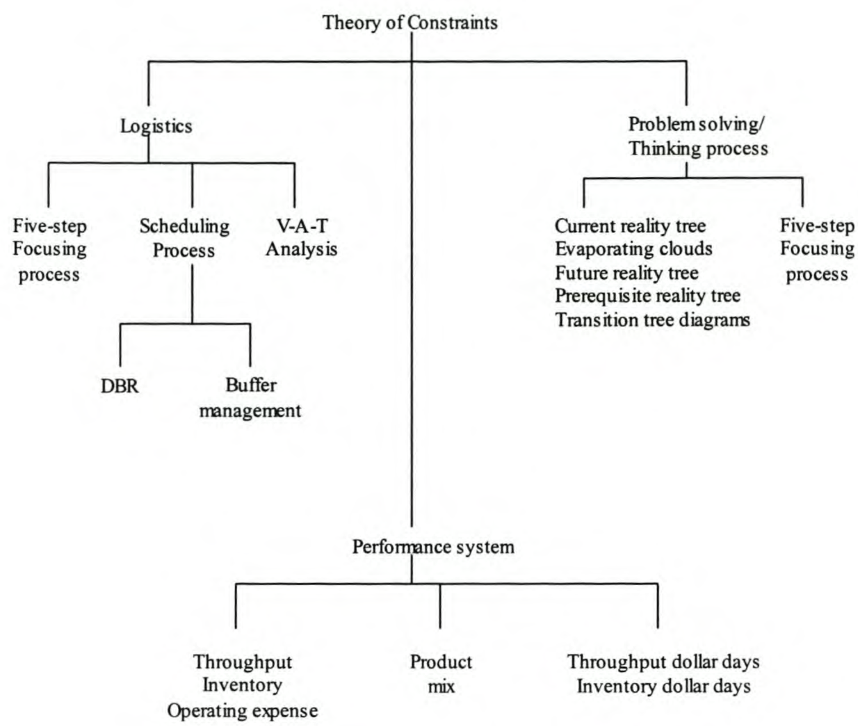

Figure 2-1. Schematic of the Theory of Constraints (Spencer & Cox, 1995)

Since 1995 the Theory of Constraints philosophy however has found a wider application than only in production or logistics. Based on this wider application, the TOC philosophy can be thought of today as basically comprising of two components: one is the Thinking Processes themselves, and the second is the applications that have been derived from applying TOC in various areas in an organisation, whether it be in production, project management, distribution, or marketing (Houle, 1998). In production, the application is the Drum-Buffer-Rope logistical system. In project management the application of the TOC's five step improvement philosophy is called Critical Chain. In distribution the application is basically a concept of replenishment as opposed to pushing products into the market, whereas in sales or marketing there is the so-called Six Phases of Buy-In that can help along with the Thinking Processes to elevate marketing constraints. The focus in this dissertation and literature study is however the production application of Drum-Buffer-Rope. A new schematic for the current TOC paradigm is displayed in Figure 2-2.

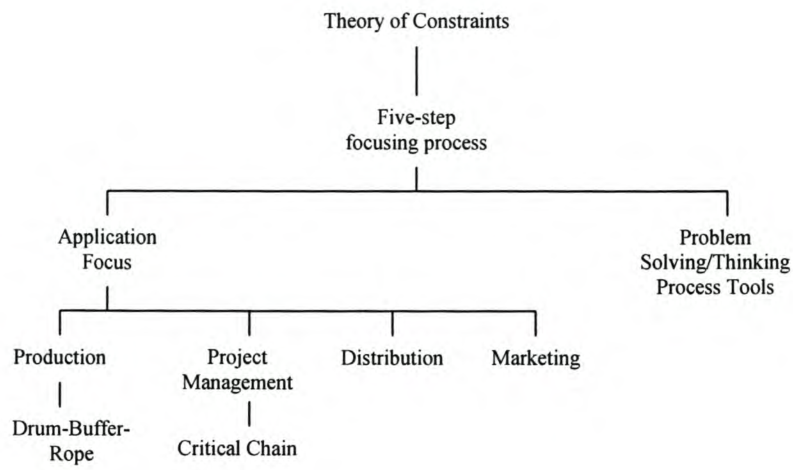

Figure 2-2. TOC focus areas

2.2.1 The performance measurement system of the theory of constraints

Under the TOC philosophy it is believed that the goal of any manufacturing company is "to make more money now as well as in the future" (Goldratt & Fox 1986, p. 18). In order to measure a company's performance in achieving this goal, Goldratt and Fox (1986, pp. 20-31) prescribe two sets of measurements: global (financial measurements) and operational measurements. The operational measurements defined are (Goldratt & Fox 1986, p. 28):

- *Throughput* (T): this the rate at which the system generates money through sales (output which is not sold is not throughput but inventory). Throughput is represented as sales minus totally variable cost.
- *Inventory* (I): all the money invested in things the system intends to sell (or could be sold). Inventory includes physical inventories such as raw material, work in process and unsold finished goods, as well as investments such as tools, buildings, capital equipment and furnishings. Inventory can therefore be seen as all the money tied up within the system (Goldratt 1986, p.72).
- *Operating expense* (OE): this is all the money the system spends in turning inventory into throughput. It includes expenditures such as direct and indirect labour, supplies, outside contractors and interest payments.

The global measurements measuring a company's performance are the following (Goldratt & Fox 1986, p. 20):

- *Net profit* (NP): this is an absolute measurement in monetary terms expressed as total throughput minus operating expense, indicating how much money was made.

Stellenbosch University http://scholar.sun.ac.za

- *Return on investment* (ROI): this is a relative measurement which equals net profit divided by inventory, indicating how much money was made relative to the money invested in the system.
- *Cash flow* (CF): this can be seen as a "survival measurement" indicating whether the company has enough cash.

The three operational measurements defined can be used to assess the impact of actions taken in the company on the global (or bottom line) measurements (Goldratt & Fox 1986, p. 30):

- Increasing throughput without adversely affecting inventory and operating expense will directly increase net profit, return on investment and cash flow.
- Decreasing operating expense without adversely affecting inventory and throughput will directly increase net profit, return on investment and cash flow.
- Decreasing inventory directly increases return on investment and cash flow.
- The indirect impact of inventory on the three global measurements is estimated through the use of carrying charges. Lowering inventory reduces some of the operating expenses such as interest charges, storage space, scrap, obsolescence, material handling and rework, which in turn increases the three global measurements.

Traditionally the emphasis of management has been on reducing operating expense first, followed by increasing throughput, and finally reducing inventory. Goldratt however suggests that the biggest gains can be realised by first increasing throughput, then by reducing inventory, and lastly by reducing operating expense. The reason for this is because the reward from decreasing cost is finite (the theoretical lower limit is zero, whereas a realistic limit is considerably higher), but theoretically increased profit from improved sales is unrestricted (Rahman 1998).

Two of the performance measurements typically used to evaluate production workers are utilisation and efficiency. These measures however encourage workers to maximise the output at each resource (Gardiner et al 1993). According to the Theory of Constraints philosophy non-constraint resources should produce only in quantities sufficient to supply the constraint. These local performance measures therefore only lead to an increase in WIP and lead times.

In order to get the shop floor results congruent with the goal of maximising the throughput of the plant as a whole, Goldratt developed two local performance measurements as replacements for utilisation and efficiency (Gardiner et al 1993):

1. *Inventory dollar-days* – This measures the extent to which a department or worker contributes to the early finish of an order. Goods finished early incur opportunity cost and decreased market responsiveness and should therefore be penalised (Gardiner et al 1993). This measurement is

computed as the number of days of early release of material into the system times the money value of the inventory.

2. *Throughput dollar-days* – This measures the lateness of an order. It is computed as the money value of the order times the number of days the order is late.

The effect of these two measurements is therefore to ensure the timely release and completion of work at all work centres. As Goldratt comments on the importance of local performance measurements: "Tell me how you measure me, and I will tell you how I will behave" (Goldratt 1990, p.28).

2.2.2 The problem solving/thinking process branch of TOC

In order to also address managerial or policy constraints, Goldratt developed a generic approach for diagnosing and solving problems called the "Thinking Process". This process focuses on the following three questions faced by managers in general (Rahman 1998):

1. Decide what to change.
2. Decide what to change to.
3. Decide how to cause the change.

The Thinking Process uses basically a cause-and-effect reasoning to deduce the core problems underlying observed symptoms or problems. Various types of diagrams are used to find answers to the three generic questions mentioned above. Table 2-1 presents the three generic questions along with the purpose of each question or step and the tools used to find answers to the questions asked. For a more thorough description of the Thinking Process and its different tools (along with examples of the tools), refer to Noreen et al (1995, pp.153-187), and Cox and Spencer (1998).

Table 2-1. Thinking Process (TP) tools and their roles (Rahman 1998)

Generic questions	Purpose	TP Tools
1. What to change?	Identify core problems	Current reality tree
2. What to change to?	Develop simple, practical solutions	Evaporating cloud Future reality tree
3. How to cause the change?	Implement solutions	Prerequisite tree Transition tree

2.2.3 The Drum-Buffer-Rope logistical system

The drum-buffer-rope system is a logistical system that is grounded in the concepts of constraints and constraint management. The basic essentials of the drum-buffer-rope system can be described as follows (Srikanth & Umble 1997, p.184):

- *Drum* – Considers the constraints in the system and sets the pace for the entire system.
- *Buffer* – Protects the system from the disruptions inherent in any process.
- *Rope* – Mechanism for synchronising all resources in the system to the drum.

Refer to Figure 2-3. The planned rate of production should be set according the pace of the constraint (the drum). To protect this constraint resource against disruptions, a buffer of materials should be provided in front of the constraint resource to ensure that the constraint always has material to work on. The amount of slack in the rope is therefore analogous to the size of the buffer. Tying the rope to the first station refers to controlling the release of materials according to the schedule of the constraint resource.

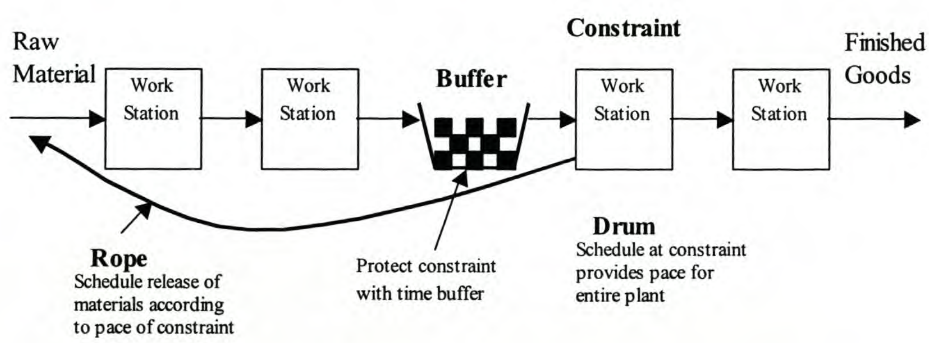

Figure 2-3. The Drum-Buffer-Rope system (Srikanth & Umble, 1997, p. 190)

The working principle of the drum-buffer-rope system is based on the following five focusing steps (Goldratt 1990, p.5):

1. *Identify the system's constraint(s).*
2. *Decide how to exploit the system's constraints* – physical constraints should be made as effective as possible, whereas managerial constraints should be eliminated and replaced with a policy that will increase throughput.
3. *Subordinate everything else to the constraint* – all other resources should be synchronised with the constraint resource.
4. *Elevate the system's constraints* – improvement efforts should be made to remove the constraint.
5. *If in any of the previous steps a constraint is broken, go back to step 1. Do not let inertia become the next constraint.* – If a constraint has been removed, the system will eventually encounter a new

constraint and therefore the five steps have to be repeated. The warning "do not let inertia become the next constraint" is a reminder that no policy or solution is appropriate or correct for all time or in every situation. As the business environment changes, business policy has to be refined to take account of those changes.

These five steps can be seen as the philosophy of the Theory of Constraints and provides the focus for a process of continuous improvement. In the following sections a few of the concepts applied in the drum-buffer-rope system will be explained.

2.2.4 Constraints and "wandering" bottlenecks

A constraint is defined as "Any element that limits the organisation from achieving higher levels of performance, where performance is measured in terms of the organisation's goal." (Srikanth & Umble 1997, p. 117).

Constraints can be seen as belonging to one of three categories (Srikanth & Umble 1997, p. 121):

1. Physical constraints – this includes resource capacity, material availability and quality, space availability, etc.
2. Market constraint – this constraint exists when the demand for the company's products and services is less than or equal to the capacity of the organisation.
3. Policy constraint – This is a constraint that is not physical in nature. It includes the entire system of measures and methods and even the mindset that governs the strategic and tactical decisions of the organisation.

Srikanth & Umble (1997, p. 122) define a capacity constrained resource (CCR) as "any resource whose available capacity limits the organisation's ability to meet the product volume, product mix, or demand fluctuation required by the marketplace."

A distinction should be made between a bottleneck and a capacity constrained resource. A bottleneck resource can be defined as "any resource whose capacity is equal to or less than the market demand placed upon it" (Srikanth & Umble 1997, p. 92). At a bottleneck resource the total load placed upon it is therefore more than the total available capacity. Any bottleneck resource can be a capacity constrained resource, but a capacity constrained resource is not necessarily a bottleneck. A capacity constrained resource or CCR's capacity is not necessarily less than or equal to market demand, but is the most heavy utilised resource which presents the most danger to disrupting throughput.

A frequently reported concern from production managers is the phenomenon of "wandering" or temporary bottlenecks, which is seen as an obstacle to the implementation of drum-buffer-rope (Hurley & Kadipasaoglu 1998). A common explanation for wandering bottlenecks, especially in job-shop environments, is that changes in product mix cause different work stations to become the bottleneck. According to Srikanth and Umble (1997, p.130) wandering bottlenecks are more the result of wrong managerial policies and actions and less the result of limited capacity. "Wandering bottlenecks are typically created by managers actively chasing piles of inventory from one resource to another. In effect, what appears to be a physical constraint is actually the result of dysfunctional policies." (Srikanth and Umble 1997, p.130).

Hurley and Kadipasaoglu (1998) discuss through a case study of an actual production facility how certain performance measurements (such as labour and machine efficiencies) can cause the appearance of wandering bottlenecks. These measurements cause management to balance capacity with current workload by removing any spare capacity from workstations. Depending on the current product mix, processing variability may be high at different workstations at different times. In order to absorb this variability, some protective capacity is necessary at the different resources. By stripping all idle capacity from resources, bottlenecks appear to wander across the plant, while in fact it is caused by an inability to cope with some processing variability. A focus on high worker and machine utilisation causes an order-release rate that is greater than the capacity of the true bottleneck, as well as a focus on large batch sizes in order to minimise set-ups. All this create large amounts of WIP and long lead times which contribute to the appearance of wandering bottlenecks. According to Hurley and Kadipasaoglu (1998) it is only in a minority of production facilities that product mix changes cause the bottleneck location to move. This tends to be facilities that produce products that are seasonal in nature.

Lawrence and Buss (1994) examined the shifting bottleneck phenomenon from an analytic perspective. They define a bottleneck as that work station having the most number of jobs waiting in queue, the greatest proportion of time, and created a bottleneck shiftiness measure (β) which is a function of the variance in the probabilities that each work station will be a bottleneck. In this measure Lawrence and Buss first calculate the bottleneck probability for each resource (i.e. the probability that that resource is the bottleneck a given percent of the time). By looking at the bottleneck probabilities at all resources they then calculate an overall coefficient of variation across all bottleneck probabilities. This coefficient of variation in the bottleneck probabilities is then used in the following formula to calculate the bottleneck shiftiness measure:

$$\beta = 1 - \frac{cv}{\sqrt{n}}$$

where

cv = the coefficient of variation of the bottleneck probabilities

n = the number of work stations

$\sqrt{n}$ = the coefficient of variation of a vector of one 1 and $(n-1)$ zeros

The bottleneck shiftiness measure compares the coefficient of variation of bottleneck probabilities across all stations with the coefficient of variation of bottleneck probabilities if one station is the dominant bottleneck with probability one. The measure varies between zero and one. If the bottleneck changes continuously, this measure is 1, and if the shop has only one constant bottleneck, the measure is zero. They also investigated different policies for managing shifting bottlenecks. Two of these policies (a chasing policy by adding temporary capacity to the current short-term bottleneck, and a policy of permanently increasing the capacity at the long-term bottleneck) were found to improve shop performance in terms of mean flow time, but have the undesirable effect of significantly increasing bottleneck shiftiness. Conversely, a third policy of increasing the capacity of non-bottleneck work stations (adding protective capacity) were found to simultaneously reduce bottleneck shiftiness and improve system performance in terms of mean job flow time.

2.2.5 Process and transfer batches

In a manufacturing environment two types of batches can be defined:

- *Process Batch* – "the quantity of a product processed at a resource before that resource changes over to make a different product" (Srikanth & Umble 1997, p. 163).
- *Transfer Batch* – "the quantity of units that are moved at the same time from one resource to the next" (Srikanth & Umble 1997, p. 163).

According to the drum-buffer-rope methodology the transfer batch need not, and many times should not, be equal to the process batch (Srikanth & Umble 1997, p. 163). The advantage of using transfer batches that are smaller than the process batch is that smaller transfer batches allows for operations to be overlapped, causing shorter manufacturing lead times and therefore less work-in-process inventory (Chase & Aquilano 1992, p. 927).

At bottleneck resources (especially those which require significant set-ups) larger process batches should be used. Larger process batches require fewer set-ups and can therefore generate more processing time and more output. At non-bottleneck resources however smaller process batches should be used in order to keep the flow of materials smooth and balanced (Chase & Aquilano 1992, p. 927).

2.2.6 Comparative studies between DBR, assembly lines, MRP and JIT

Several studies have compared the performances of DBR, MRP and JIT by using computer simulations. Fogarty et al (1991, pp. 649-656) constructed a model of a simple shop that has a two station assembly line and produces one product. Monte Carlo simulations were conducted using a spreadsheet. They conclude that DBR gives superior performance with less effort. In their study both DBR and JIT outperformed the MRP approach. DBR produced greater output than JIT and MRP with only marginally greater lead time and inventory. For larger plants with more operations and therefore more variability they however claim that DBR will require less WIP than JIT because of the different buffering method used in DBR.

Cook (1994) conducted a simulation study on a hypothetical flow shop consisting of five work stations. His study indicates that both DBR and JIT outperform the traditional MRP way, whereas DBR outperforms JIT on a number of critical performance measures, including total output and standard deviation of flow time.

Miltenburg (1997) conducted a study using a two station production line and a Markov chain model. The performance measurements used were output, inventory, cycle time and shortages. Both DBR and JIT outperformed MRP on all these measures. DBR outperformed JIT on output and shortages.

All of these three studies indicate that both DBR and JIT outperform the traditional MRP approach. While JIT achieves lower inventory, DBR produces more output. For JIT to achieve the same output as DBR, the flow time and standard deviation of flow time must be drastically reduced. JIT would have to virtually eliminate all variability across the whole system to make it comparable to DBR, which is an unlikely possibility (Cook, 1994).

Chakravorty and Atwater (1996) conducted a study comparing the relative performance of three approaches to designing and operating serial production lines. They used simulation to study the output rates of a balanced line, a JIT line and a TOC (or DBR) line. Each line is subjected to various combinations of station downtime and variability in station processing time in order to identify which line design is most affected by these forms of variability. Different inventory levels were also used to test the impact of inventory on the relative output performance of the three line designs. Their study indicates the following:

- DBR lines perform best (highest output) when station variance is relatively high, whereas JIT lines perform best when station variance is relatively low.
- DBR lines are least affected by changes in station variance, whereas JIT lines are the most heavily affected by changes in station variance.
- DBR lines perform best when station downtime is relatively high, whereas JIT lines perform best when station downtime is relatively low.
- DBR lines will significantly out-produce both JIT and balanced lines at relatively low levels of total system inventory.
- DBR lines will achieve their maximum output level with much lower levels of inventory in the system than JIT lines.
- With sufficient inventory, the JIT line will significantly out-produce both DBR and balanced lines.

This study therefore also suggests that DBR systems perform better in environments subjected to variability. A higher output level will also be achieved with fewer inventories.

2.3 Time Buffers

Time buffers are used in the drum-buffer-rope methodology to protect the system's output from the disruptions that occur in the system. These buffers represent the additional planned lead time allowed, beyond the required setup and run times, for materials to reach a specified point in the product flow (Srikanth & Umble 1997, p.298). Strategically placed, time buffers are designed to protect the system throughput from the internal disruptions that are inherent in any process. Three types of buffers are used in the DBR methodology (Rahman 1998):

1. *Constraint buffer* - This buffer is used to protect the schedule of the constraint and is inserted just before the constraint resource. The operation immediately preceding the constraint resource should therefore complete the parts a time buffer before they are scheduled on the constraint resource. The length of the time buffer together with the processing time for the work stations preceding the constraint are used to determine the time at which material is to be released to the shop floor (Noreen et al 1995, p.33).
2. *Shipping buffer* – This buffer is used to protect the delivery dates of the orders (the market constraint) and is therefore located at the end of the process. Products should therefore be finished a time period equal to the length of the shipping buffer before the due date of the order (Rahman 1998).
3. *Assembly buffer* – This buffer is located in front of an assembly operation that is fed by both CCR and non-CCR parts. The purpose of this buffer is to ensure that parts coming from a constraint resource do not have to wait for parts coming from non-constraints (Goldratt, 1990, p. 131).

Figure 2-4 shows how the time buffers and processing times are used in placing the load on the constraint and determining the material release date. The rope is only a mechanism whereby the release of materials is controlled to ensure their timely arrival at the respective operations protected by time buffers. It is a deterministic lead time offset from the specific protected operation (which can be a CCR, an assembly using both CCR and non-CCR parts, as well as the final shipping operation) to material release (Gardiner et al 1992).

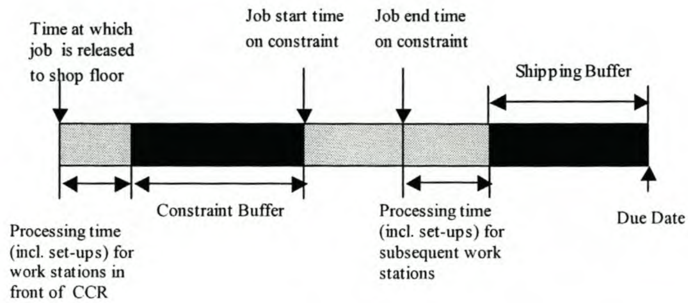

Figure 2-4. Placement of time buffers and constraint load on time axis

Figure 2-5 presents a flow diagram of a simple manufacturing process indicating the flow of raw materials from the first operation to the finished order. It also indicates the placement of the different time buffers, as well as the ropes that are used to calculate material release. The length of a specific rope is the length of the respective time buffer to which it is connected plus the processing times of the operations included between the two end points of the rope.

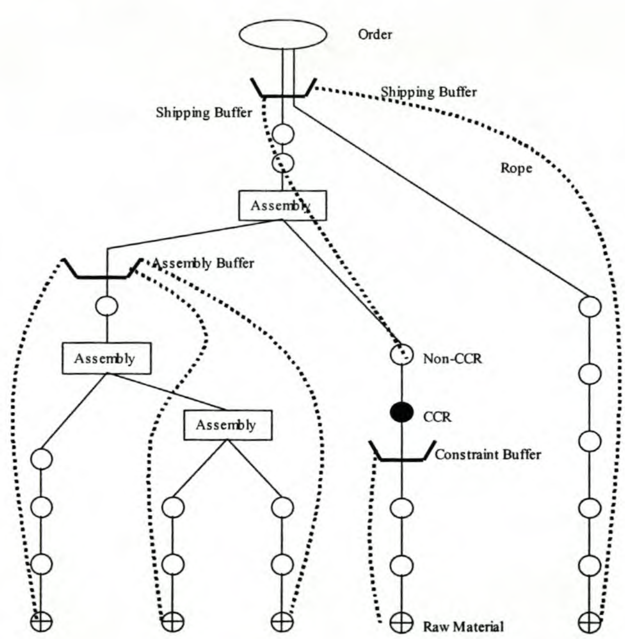

Figure 2-5. Flow diagram indicating placement of time buffers and ropes

2.3.1 Determining the size of the time buffers

The objective of the time buffers is to protect the schedule against the typical disruptions experienced by the manufacturing plant. The length of the time buffers therefore depends on the magnitude of these disruptions and fluctuations (Schragenheim & Ronen 1990). According to the TOC literature the determination of the time buffer lengths is a trial and error approach which consists of first determining the initial size of the time buffers through simple empirical rules (Tu & Li 1998), (Srikanth & Umble 1997, p. 235). The buffer lengths are then monitored and adjusted through a process known as buffer management (Schragenheim & Ronen 1991), (Goldratt 1990, p.238).

The determination of the time buffer lengths depends on whether a manual drum-buffer-rope system is used, or a computerised system. For a manual system, all time buffers are fixed in length (Gardiner et al 1998). The length of the time buffer should therefore include an estimation of the lead time to the specific buffer origin, as well as estimations of disruptions. According to Goldratt (1990, p. 242) the direct contribution of process times to the total lead time is usually very small. The majority of time is spent in queues. [This is more true of job shops than repetitive flow lines (Spencer 1991)]. To determine estimations for individual queue times is however very difficult. Queue times are mainly a function of the load that is placed on the resource. Orders however do not arrive in an even stream, consequently the load placed on each type of resource can fluctuate considerably.

For computerised systems Goldratt (1990, pp. 235-240) proposes a concept known as dynamic buffering. Dynamic buffering allows a buffer to extend in length when non-constraint resources feeding the buffer are overloaded in the short term. The buffer length should therefore only include an estimation of the impact of disruptions and fluctuations, and not the total lead time. The system will add to the buffer the influence of processing times and queue times caused by the specific load situation. The concept of dynamic buffering is discussed in more detail in appendix C.4.

The following empirical rules can be used to determine the initial time buffer lengths:

- Goldratt (1990, p.238) presents a simple empirical rule for determining the initial buffer lengths: estimate the current average lead time of the tasks to the specific buffer origin and divide it by five.
- Srikanth and Umble (1997, p.235) suggest the following empirical rule as a convenient starting point in determining the time buffer lengths: the total time buffer for any product should be approximately one-half the firm's current manufacturing lead time. This total buffer length should be redistributed between the different time buffers (i.e. shipping, constraint and assembly buffers).

- Schragenheim and Ronen (1990) suggest a constraint buffer size of three times the minimum cumulative processing time to the constraint. This calculation therefore includes processing time and not just protection against disruptions.

After the initial time buffers lengths have been determined, the time buffers should be monitored and adjusted through a process known as buffer management, which will be discussed in section 2.3.2.

Other techniques that can be used to determine the lengths of the time buffers include computer simulation studies (Tu & Li 1998). Tu and Li (1998) also developed a constraint time buffer determination model. In this model a tree structure is developed to represent the relationship between the constraint machine and its feeder machines. By incorporating the Mean-Time-Between-Failure (MTBF) of each feeder machine, a mathematical relationship can be formulated and used to determine the constraint buffer length. A simulation experiment was conducted on a six station job shop producing four different products. Different constraint time buffer lengths were simulated and the average flow times and throughput rates monitored. The results indicated that the throughput rate for the system did not increase much for constraint time buffers larger than the computed time buffer length, while the average flow time increased sharply. For constraint time buffer lengths shorter than the computed value, the average flow times decreased but at the same time the throughput rate also significantly decreased. These results illustrated the validity of the constraint time buffer calculation procedure. The time buffer determination model from Tu and Li however only accounts for variability caused by machine downtimes, and not other sources of variability such as processing time variability and transfer batches.

Radovilsky (1998) presents a queuing analysis approach to estimate the size of the time buffers. The constraint resource is modelled as a M/M/1/K system (poisson arrivals, exponential service time, single server with finite queue of maximum K units). The optimal number of units waiting in the queue in front of the constraint (i.e. the optimal size of the time buffer) is then determined that will maintain the highest operational profits while protecting the constraint from becoming idle. Apart from the restrictions of his approach through the assumptions of exponential service times, poisson arrivals and a single server, he also does not model the effect that disruptions at resources upstream from the constraint have on arrivals at the constraint.

2.3.2 Time buffer management

The performance and progress of orders in the production plan can be monitored by monitoring the contents of the different buffers. This process is known as buffer management (Gardiner et al 1993). The monitoring can be done either by visual inspection or by using a computerised system. Basically buffer management is the monitoring of the inventory in front of the protected resources and comparing the actual results to the planned performance (Schragenheim & Ronen 1991). Material that is significantly late in arriving at the specific buffer warns against potential disruptions in the production plan.

If disruption and fluctuations exist in the plant, then the actual buffer content must be smaller than the planned content, otherwise there is no need for a buffer (Goldratt & Fox 1986, p.123). If the buffer is always full, it is a sign that there are no disruptions significant enough to affect the planned material flow and the buffer can be eliminated. The desired buffer content profile is illustrated in Figure 2-6.

The buffer content must be sufficient to protect the throughput of the plant. Material planned to be in more or less the first third of the time buffer should be present almost all of the time (Goldratt and Fox 1986, p.122), because they represent the material the CCR is to work on first. Because the CCR must be exploited, it should not stand idle waiting for material. On the other hand, most of the material planned for the last third of the buffer should be missing. The content profiles for the middle part of the buffer should fall somewhere between the extremes of the first and last parts.

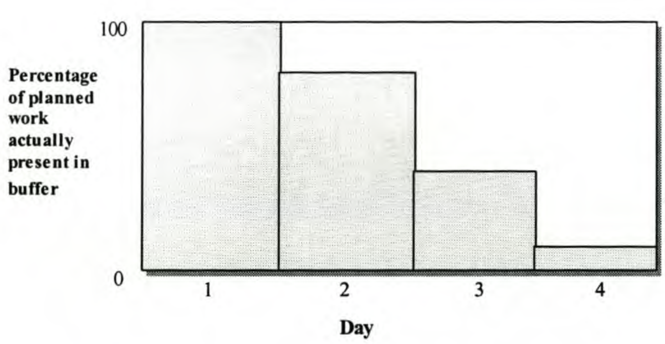

Figure 2-6. Desired Buffer Content Profile

By monitoring the buffer content profiles, the appropriate lengths for the time buffers can be managed. Refer to Figure 2-7. It displays three different possible buffer content profiles. In case 1 the actual buffer content stretches beyond the planned contents. This is an indication that materials are being released prematurely and worked on earlier than required at the workstations preceding the CCR (Goldratt & Fox 1986, p. 124).

In case 2 most of the required materials are already present in the planned buffer. This is an indication that the planned buffer is too large and the plant is paying for this over-protection through unnecessary high lead times, inventory and operating expense (Goldratt & Fox 1986, p. 124), (Umble & Srikanth 1997, p. 232). The buffer length should be reduced.

In case 3 the actual buffer content is too low. Too many of the materials needed within the first third of the buffer are not present and the CCR runs the risk of standing idle waiting for material. The planned buffer size should therefore be increased until the first third of the buffer is completely full (Goldratt & Fox 1986, p. 124).

Holes that appear in the buffer content can be used to warn against potential disruptions to the production plan. Whenever a hole in a buffer is spotted, three lines of action might be adopted: disregard the missing material, track the missing material, or expedite. To help determine the correct line of action a time buffer can be divided into three regions (refer to Figure 2-8) (Schragenheim & Ronen 1991):

- Region 3 – This is the safety zone. Holes that appear in this region of the buffer do not present any real concern and may be disregarded. There is still enough time left until these parts are needed by the constraint.
- Region 2 – This is the tracking zone. Some holes in this region are normal. No expediting is necessary yet, but missing parts should be tracked in order to identify the causes of the delays. This will help to identify and eliminate chronic sources of disruption and therefore provides a focused process of continuous improvement.
- Region 1 – This is the expediting zone. Holes that appear in this region are very near the scheduled processing time and therefore threaten the throughput of the plant. Missing parts should therefore be expedited.

The relative size of each region is dependent on the specific application and can change from buffer to buffer, as well as from time to time (Srikanth & Umble 1997, p. 237). Region 1 has to be large enough so that in most cases the missing parts can be rushed to the constraint in time. Determining the "right" lengths of the different regions "reflects the balance between lead time and protection of the key areas of the whole organisation" (Schragenheim & Ronen 1991).

The value of buffer management therefore lies in the fact that it "enables management to focus on the right corrective actions to keep the system performance intact, monitor the trade-off between protection and lead time, and assess the impact of major changes and/or improvements which are to be implemented" (Schragenheim & Ronen 1991).

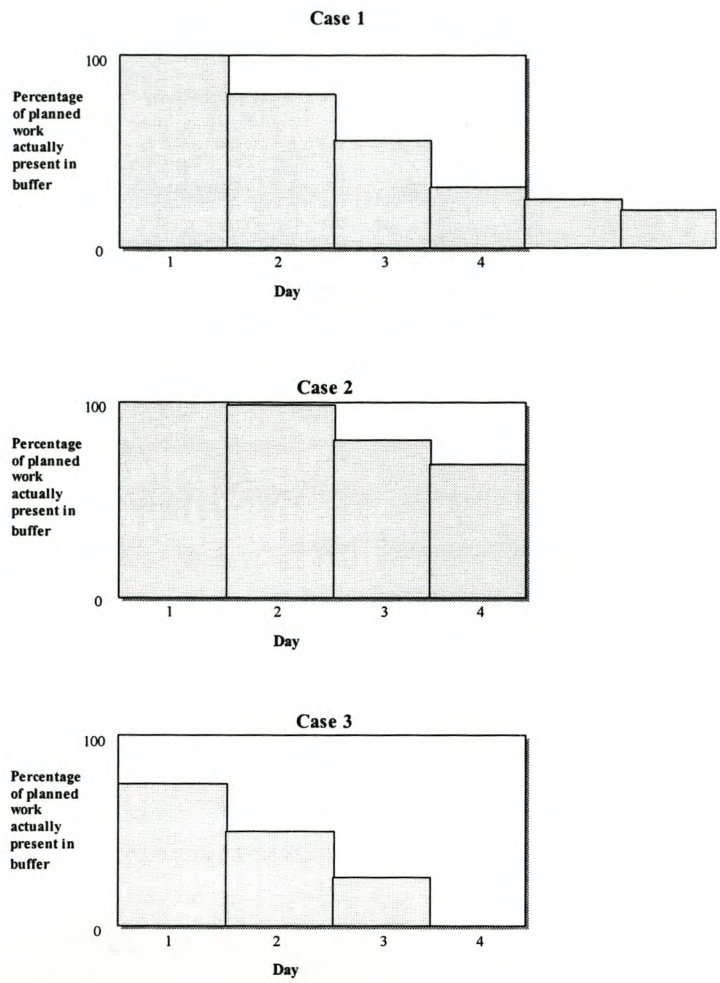

Figure 2-7. Different Buffer Content Profiles

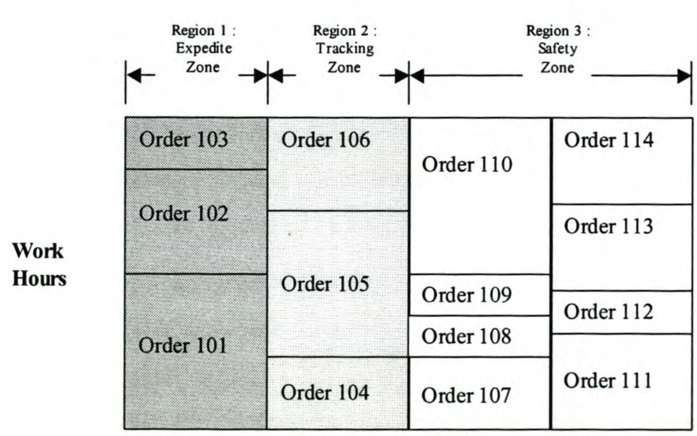

Figure 2-8. The different buffer regions

2.4 Unbalanced Lines and Protective Capacity

As mentioned in section 1.2, protective capacity is defined in the APICS dictionary as "a given amount of extra capacity at non-constraints above the system constraint's capacity" (2002, p.95). The traditional view on line design is that workstations in a line should have balanced capacity. The Theory of Constraints philosophy however advocates unbalanced lines through the addition of protective capacity at non-constraint resources. This section examines the literature on balanced and unbalanced line designs, as well as studies investigating the use of protective capacity.

2.4.1 Balanced vs unbalanced lines

One of the first studies to suggest that an unbalanced line may outperform a balanced line was conducted by Patterson (Atwater 1991). Patterson used a Markov process to analyse the flow of jobs through a three station assembly line. He compared the performance of a line with infinite queue space against a line with finite queues using exponential job inter-arrival and service time distributions. No station downtimes were considered. His study demonstrates that allowing enough inventory to accumulate between stations results in a production rate for the line equal to the rate of the slowest station, whereas insufficient WIP protection causes the line's production rate to be slower than the slowest station. Patterson also found that by intermingling fast and slow stations and using a fixed amount of inventory an unbalanced line outperformed a balanced line in terms of production rate.

Davis (1966) investigated pacing effects on manned assembly lines. In his study he also examined different arrangements of an unbalanced line. A three station line was simulated with six different arrangements of processing times: (1) fast-medium-slow; (2) fast-slow-medium; (3) medium-fast-slow; (4) medium-slow-fast; (5) slow-fast-medium; and (6) slow-medium-fast. A fast station is defined as 3.5 minutes per unit, a medium station as 4.5 minutes per unit, and a slow station as 5.5 minutes per unit. All processing times were modelled using the lognormal distribution with a variance of 1.45 minutes. The main performance measures used by Davis were average idle time % for the line, average inter-discharge time from the last station, and the total number of items completed (throughput). With respect to minimising the average system idle time and average inter-discharge time performance measures, the fast-medium-slow configuration was found to be superior. With respect to the maximum number of items completed or total throughput, the slow-medium-fast configuration was found to outperform the other configurations. As highlighted by the Theory of Constraints philosophy, the total system throughput measure is a more important performance measure than the other two measures used in this study.

Stellenbosch University http://scholar.sun.ac.za

A well-known study on the performance of unbalanced lines versus balanced lines is the work of Hillier and Boling (1966 cited Atwater 1991). They discovered the so-called "bowl phenomenon" regarding the optimal allocation of work in some unpaced production line systems with variable processing speeds. They found that unbalancing a line in a certain way would outperform a balanced line in terms of total line throughput. Hillier and Boling used queuing models to study two, three, and four station lines with exponential processing times, limited queue space (buffer storage capacity varied from zero to four units) and no station downtimes. They compared a balanced line with a mean service time of one minute with unbalanced lines where the middle stations had faster processing times than the end stations. The balanced line was unbalanced by deducting service time from the middle station and adding equal amounts to each of the end stations. This allocation produced a bowl shape and their results indicated that the line output could be increased with such an allocation (compared to a balanced line). Their results opened up many research questions, e.g. will the bowl phenomenon hold for service times other than the exponential distribution, or will it hold for longer lines?

Payne, Slack, and Wild (1972) used computer simulation to investigate the operating characteristics of a balanced line versus an unbalanced line. They simulated a 20 station line with normally distributed service times, unlimited buffer storage capacity between stations, and no service down times. They compared a balanced line (where each station had a service time of 10 time units) with four other cases: (1) increasing the coefficient of variation as the station numbers increased while keeping the service times the same for all stations; (2) decreasing the coefficient of variation as the station numbers increased while holding the service times constant; (3) increasing the service times as the station number increased while holding the coefficient of variation constant (service times ranged from 9.525 time units to 10.475 time units) ; and (4) decreasing the service times as the station number increased while holding the coefficient of variation constant. In all these cases the imbalance was achieved by subtracting service time or CV from one end station and adding it to a station at the other end of the line (service times ranged from 10.475 time units to 9.525 time units). In all four of the unbalanced cases a lower line output was achieved than with the balanced line. In their discussion of the results Payne, Slack and Wild emphasize that their results do not apply to steady state conditions since the time period the simulation was run was not long enough. They however argue that few practical flow lines operate at steady state conditions.

Rao (1976) investigated the bowl phenomenon for a three station line with no buffer storage where different stations have different variability. He considered all combinations of deterministic and exponential service times at the three stations. For each combination he developed analytical formulas for the mean cycle time, which he used to determine the production rate for a balanced line where the mean service time equals one at each station. The optimum production rate from unbalancing the line was then found. Rao investigated two ways of arranging the imbalance, i.e (1) the bowl arrangement suggested by Hillier and Boling (1966 cited Atwater 1991) where the mean service time for the middle station is decreased while increasing the service times for the two end stations, and (2) increasing the service times for the less variable stations while decreasing the service times for the more variable stations. Rao concluded that unbalancing a serial line can lead to substantial improvements in throughput rate when the variability of the stages differ from one another. When the coefficient of variation of the stations is less than 0.5 the bowl arrangement achieves the highest production rate, whereas the variability imbalance works best for lines with station coefficients of variation greater than 0.5. Criticisms to Roa's study are that he only considered very short lines, no station down times were considered, and thirdly he only used exponential and deterministic service times which are unrealistic for most real world situations.

De La Wyche and Wild (1977) conducted a simulation study to investigate the effect of imbalance in service time variability and imbalance in buffer storage capacity between stations on mean line idle time. Three, four and twelve station lines with normally distributed processing times were simulated. Different arrangements of variability and buffers storage capacity were considered. Their results indicate that for longer lines, putting the low variance stations in the centre of the line reduced line efficiency. On the other hand, separating relatively variable stations by steadier stations improves line efficiency. A main criticism against their study is the use of line idle time as the performance measure, since high efficiency measures can lead to actions that result in unacceptable performance with respect to line output and lead time.

In a follow up study Hillier and Boling (1979) examined whether the bowl phenomenon will hold for longer lines, increases in the amount of storage space between stations, and less variance of service times. They incorporated Erlang service times and again used queuing models and numerical analysis to investigate the line balancing problem for lines up to 6 stations long. In their study they showed that the bowl phenomenon held for longer lines and different shapes of the Erlang distribution. The improvement using the bowl distribution however is small and the line must be unbalanced in the appropriate way before improvements can be achieved.

El-Rayah (1979) became the first researcher to use simulation to confirm the existence of the bowl phenomenon. He studied 3, 4 and 12 station lines using normal, lognormal and exponential service times. Different buffer storage capacity was also considered. Four different arrangements of imbalance were used, i.e. (1) a balanced line with identical service times, (2) a random pattern where slow and fast service stations are intermingled, (3) an increasing pattern where service times increase as you go down the line, and (4) the bowl arrangement as proposed by Hillier and Boling (1966). Imbalance in all the arrangements was achieved by subtracting service time from some stations and adding it to others. El-Rayah concluded that when lines have limited inter-stage buffer storage capacity an unbalanced line with the bowl arrangement will achieve a higher output than a balanced line.

Another simulation experiment to verify the existence of the bowl phenomenon in longer lines and lines with less variability and more buffer storage space was conducted by Smunt and Perkins (1985). They tested the bowl phenomenon by varying two conditions, i.e. the variance of service times and the capacity of inter-stage buffer storage. The performance measure used was mean cycle time. Lines with three, four and eight stations were modelled using truncated normal distributions for the service times. Buffer storage capacity was varied at four levels, i.e. 1, 3, 5 and 999. The coefficient of variation of service times was modelled at 1.0, 0.5 and 0.2. The mean service time for the balanced line was 100 minutes, and the unbalanced bowl shape lines were obtained by subtracting service time from the centre stations and distributing it symmetrically to the end stations. Four different bowl arrangements were examined. Smunt and Perkins concluded that the bowl distribution of service times "does not significantly increase (and often decrease) the output rate when variance is low or moderate. Second, any benefit of using the bowl distribution is quickly dissipated when buffer storage capacity is increased by only small amounts for low to moderate levels of variance." They suggest using the bowl distribution only when task times exhibit large variances. Hillier and Boling (1979) stated that benefits from using the bowl distribution depends on the correct degree and allocation of the imbalance. The results from Smunt and Perkins's study could therefore be ascribed to inappropriate bowl allocations.

In order to investigate the results from Smunt and Perkins, So (1989) conducted a similar simulation experiment using different bowl allocations than those used by Smunt and Perkins (1985). He used the same experimental design with the same variables and parameters as Smunt as Perkins, except the buffer storage capacity was only varied at levels of 1,3, and 5. Unlimited buffer storage capacity therefore was not considered. He also used different bowl allocations based on the following guidelines as illustrated by the results of Hillier and Boling (1979):

1) "The optimal bowl allocation is symmetric."
2) "The optimal bowl allocation is relatively flat in the middle and is very steep towards the end of the line."
3) "The degree of unbalancing decreases with the inter-station buffer storage capacity."
4) "The degree of unbalancing increases with the length of the line."
5) "The degree of unbalancing increases with the coefficient of variation of the processing times."

So compared his simulation results to the bowl allocation in Smunt and Perkins's study that achieved the best results, as well as to a balanced line. He concluded that although the improvement is generally small (1% to 2% improvement in throughput), the mean cycle time for a production line with finite

Stellenbosch University http://scholar.sun.ac.za

buffer storage can be improved by unbalancing it appropriately according to the bowl shape, even when the variance of the processing times is small. The way and degree to which the imbalance is allocated is however important, otherwise a decrease in line performance would result.

Hillier and So (1996) examined the robustness of the bowl phenomenon to inaccuracy in estimating the optimum amount of unbalance. They again used the analytical queuing model of a production line as in their previous studies and investigated 3 to 7 station lines. Their results show that errors of even up to 50% in the amount of unbalance compared to the optimum bowl arrangement can still provide most of the potential improvement in throughput over a balanced line.

From the literature review on unbalanced versus balanced line designs, it can be seen that most of the research studies created an unbalanced line by redistributing work content from the balanced line among the fixed stations in the line in some fashion. Of these unbalanced arrangements of work the so-called "bowl distribution" received the most attention. The bowl distribution of unbalance appears to improve line performance in terms of total throughput. The improvement to be achieved is however small (1%- 1.6%). All of the unbalanced arrangements used in these studies are more applicable to assembly type flow lines where work can be redistributed more easily along the workstations. By keeping the total number of stations fixed, the unbalanced line arrangements also create stations with longer processing times than those in the balanced line. This explains why the potential increase in throughput is quite small and is only achieved at low degrees of imbalance.

2.4.2 Protective capacity

More applicable to flow shops where work content is not that easily redistributed among the production stages or stations, are the studies examining the use of protective capacity in flow lines. In these studies work content are not redistributed, but rather capacity is added to stations to create an imbalance in processing rates.

One of the first studies to examine the use of protective capacity is the one by Atwater (1991). He examined the impact of protective capacity on a flow shop using a simulation model of a nine station flow line with unlimited buffer capacity between stations. The specific research questions he investigated are the following:

1) What is the effect of different amounts of protective capacity on the output of a flow line?
2) What is the effect of the arrangement of the protective capacity around the slowest station on the output of a flow line?
3) Is the significance of the arrangement of protective capacity affected by the amount of protection provided by the average protective capacity and work-in-process (WIP) inventory?

4) What is the relationship between WIP inventory and average protective capacity?
5) What is the relationship between average protective capacity and the coefficient of variation of the service times at the stations?
6) What is the relationship between average protective capacity and the variability of repair time for stations in the line?
7) What is the relationship between average protective capacity and the ratio of repair time to processing time at the constraint?

In the simulation model the location of the capacity constrained resource or CCR was fixed at the last station in the line, and the total WIP in the line was held constant by routing finished entities at the end of the line back to the first station. All processing times were modelled using the lognormal distribution. The following independent variables were used:

- Service time at the CCR varied at two levels (1 minute and 6 minutes).
- Service times coefficient of variation varied at two levels, i.e. 0.01 and 0.5. All stations were assigned the same coefficient of variation.
- Average protective capacity at non-CCR stations varied at three levels (5%, 20% and 50%). Protective capacity was modelled as an average of the protective capacity shared across all non-constraint operations. The percentage protective capacity at a station is calculated as the complement of the ratio of the constraint output rate to the non-constraint output rate.
- Arrangement of protective capacity was varied at three configurations, i.e. a flat arrangement where all non-CCR stations have the same percentage of protective capacity; a descending arrangement where the amount of protective capacity at a station decreases as the station number increases; and a random arrangement where protective capacity is randomly assigned to stations.
- Different combinations of mean time between failures (MTBF) and mean time to repair (MTTR) resulting in four levels of percentage station downtime, i.e. 1.18%, 2.91%, 4.58% and 10.7%. All stations were assigned the same downtime percentages and the exponential distribution was used for the mean time between failures.
- WIP inventory levels varied at three levels, i.e. 1*(MTTR/Constraint Service Time), 2*(MTTR/Constraint Service Time), and 3*(MTTR/Constraint Service Time).

The dependent variable or performance measure studied was the total line output. Atwater made the following conclusions from his research:

1) Adding protective capacity to a line with a fixed level of WIP will increase the output rate of the line.

2) A flat arrangement of protective capacity seems to outperform the descending and random arrangements. One possible explanation for the higher performance of the flat arrangement could be that it provides the largest minimum level of protective capacity.
3) There is an inverse relationship between WIP inventory and average protective capacity for a given level of output. This means that by designing more protective capacity into a line less WIP inventory will be required.
4) As processing time coefficient of variation increases the amount of protective capacity required to maintain a given output level increases.
5) With enough protective capacity both lines with high and low processing time coefficient of variation can maintain the same output level with the same level of WIP.
6) When downtime percentages or variability in the repair time distribution increases more average protective capacity is required to maintain the line output.

In another simulation study Atwater and Chakravorty (1994) compared the cycle time (time between units coming off the line) of a line with protective capacity to that of a line without protective capacity. A 40 minute job was divided across stations in two different ways to obtain a balanced line and a line with protective capacity. The balanced line consisted of 4 stations with an average service time of 10 minutes each. The line with protective capacity consisted of 5 stations with a 6 minute processing time at the first station, 7 minute processing time at the second station, 8 minute processing time at the third station, 9 minute processing time at the fourth station, and a 10 minute processing time at the last station. This resulted in a line with a descending arrangement of protective capacity and the constraint resource located at the last station in the line. All processing times were modelled using the lognormal distribution. Unlimited queue sizes were provided between stations to prevent any blocking, and WIP inventory was held constant by routing entities at the end of the line back to the first station. A factorial research design was performed with independent variables processing time coefficient of variation, station downtime, total inventory in the system, and the type of line (with or without protective capacity). Processing time coefficient of variation was varied at two levels (5% and 50%), station downtime also at two levels (10% and 30%), and WIP inventory at eight levels ranging from low to high. The results of their experiment indicated the following:

1) The line with protective capacity achieved a faster cycle time at both low and high variability situations, while operating with lower inventory levels.
2) The line with protective capacity required less inventory in the system to achieve relatively stable cycle times.

Their simulation study therefore shows that lines with protective capacity can achieve both shorter lead times and more reliable performance, thus enabling better on-time delivery. This can be achieved with less WIP inventory than lines without protective capacity.

Kadipasaoglu et al (2000) investigated the effect of the level of protective capacity relative to the location of the constraint resource through a simulation experiment. A four station production line with unlimited queue capacity between stations was modelled. A processing time of 40 minutes per unit was maintained at the constraint resource, while protective capacity was added to the non-constraint stations by reducing their processing times. All non-constraint stations were assigned the same amount of protective capacity (thus a flat arrangement was used). The lognormal distribution was used to model the processing times. The amount of protective capacity was varied at four levels (0%, 12.5%, 25%, and 37.5%) while the constraint location was varied respectively at the first, second, third, and fourth station. Processing time coefficient of variation was also varied at three levels (0.1, 0.2 and 0.3), while the constraint downtime and non-constraint downtimes were individually varied at 10%, 15% and 20%. Performance measures studied were flow time, work-in-process, and waiting time in queues. Their study showed that the lowest flow time, WIP and waiting time were achieved with the constraint located at the first station in the line. Moving the constraint further down the line increases the probability of starving the constraint, which degrades line performance. There is however an interaction between constraint location and protective capacity. The higher the average protective capacity, the bigger the improvement in line performance by moving the constraint to the front of the line. A high constraint downtime however lessens the effect of constraint location. When non-constraint downtime is low the effect of constraint location on flow time is also insignificant, whereas at high non-constraint downtimes it is quite significant. Adding protective capacity also has a positive effect of dampening the effect of constraint downtime. This effect is particularly stronger at lower levels of constraint downtime. Protective capacity also reduces the effect of non-constraint downtime, particularly at high levels of non-constraint downtime.

Atwater and Chakravorty (2000) investigated the effect of different Order Release and Review (ORR) methodologies, primary CCR utilisation, as well as the protective capacity at the secondary CCR (second most heavily utilised resource) on the mean flow time and mean percentage orders that were finished late. They used a simulation model of a thirteen station job shop producing ten different product types. The two ORR methodologies investigated were an immediate release method (IMM) using a standard quoted lead time and a drum-buffer-rope (DBR) methodology. The utilisation of the constraint resource was varied at three levels (94%, 96%, and 98.5%), while the protective capacity at the secondary CCR was varied at five different ranges (<%, 1%-3%, 3%-5%, 5%-7%, and 7%-9%). The different levels of constraint utilisation and protective capacity at the secondary CCR were achieved by changing the arrival rate of selected products. The results of the study showed that at all levels of protective capacity the DBR system had a lower mean flow time than the IMM system. The advantage in mean flow time provided by the DBR system however diminishes at lower levels of constraint utilisation. In the study it was observed that the drop in mean flow time from increasing the protective capacity at the secondary CCR is largest for the smaller amounts of protective capacity. The impact from increasing protective capacity is also the biggest at the higher levels of constraint utilisation. In this study however the protective capacity level at the secondary constraint was changed by changing the arrival rate of certain products. This had the effect that changing the protective capacity level at the secondary constraint also changed the protective capacity at the other non-constraint resources. The secondary constraint was therefore not investigated separately.

Craighead et al (2001) investigated the impact that positioning of protective capacity has on manufacturing cell performance. They conducted a simulation experiment on a five station flow line and examined the effect that different levels of protective capacity, different configurations of protective capacity, as well as a bias in the placement of protective have on the performance of a line. Station 3 was the bottleneck station. The levels of total protective capacity used were 10%, 20%, 30%, 40% and 50%. Five different configurations were used, i.e. a level arrangement (where all non-bottleneck stations had the same amount of protective capacity), a valley arrangement (where stations 1 and 5 had more protective capacity than stations 2 and 4), a peak arrangement (where stations 2 and 4 had more protective capacity than stations 1 and 5), a sawtooth arrangement (where stations 1 and 4 had more protective capacity than stations 2 and 5), and a reverse sawtooth arrangement (where stations 2 and 5 had more protective capacity than stations 1 and 4). Apart from these different configurations of protective capacity, the placement of the protective capacity was also biased to either before or after the bottleneck. Three biased patterns were included, i.e. 65%/35% (where the total amount of added capacity was higher before the bottleneck), 35%/65% (where the total amount of added capacity was higher after the bottleneck), and 50%/50% (where it was evenly distributed). The average processing time at the bottleneck was set at 100 seconds and the lognormal distribution with a 40% coefficient of variation was used for all processing times. WIP was held constant by routing finished parts from station 5 back to station 1. The output from the simulation runs was analysed with ANOVA. The performance measures investigated were mean flow time and bottleneck shiftiness. The bottleneck shiftiness measure used was the one developed by Lawrence and Buss (1994) as discussed in section 2.2.4. This measure indicates the inclination for the bottleneck to move or change and therefore provides a measure of the difficulty in managing the cell. The results confirmed that increasing protective capacity significantly decrease mean flow time. With respect to the bias experimental factor the shortest flow time was obtained with the equal bias (where equal amounts of protective capacity were placed before and after the bottleneck). Although significant, the difference in flow time between the three bias patterns was however very small. The lowest bottleneck shiftiness was obtained with the before bias, followed by the even bias. With respect to the protective capacity configuration the level, valley and reverse sawtooth patterns had the shortest flow times (and were not significantly different from each other), whereas the peak and sawtooth patterns had the longest flow times (and were not significantly different from each other). The peak configuration however obtained the lowest bottleneck shiftiness for all levels of protective capacity and bias. Craighead et al (2001) conclude from their study that the mean flow time appears to be largely unaffected by either the configuration pattern or bias of the protective capacity. A different cell configuration where for instance variation is not distributed equally at all workstations may however produce different results. The configuration and bias of protective capacity conversely do affect bottleneck shiftiness, with the peak configuration and before bias achieving the lowest shiftiness.

Patterson et al (2002) investigated the impact of the quantity and location of processing time variance in a manufacturing cell. A full factorial experiment with a simulation model was conducted using a five station manufacturing cell. The third station was the bottleneck and WIP was held constant by routing finished parts from station 5 back to station 1. The average processing time at the bottleneck was set at 100 seconds and the lognormal distribution was used for all processing times. The independent variables used in the study were the pattern of variance with respect to the bottleneck station, the amount of variance, and the amount of protective capacity. Three patterns of variance were used: an even pattern where variance was evenly distributed around the bottleneck, an away pattern where a higher proportion of the variance was away from the bottleneck, and a close pattern where a higher proportion of the variance was close to the bottleneck. For the amount of variance the coefficient of variation was set at five levels: 10%, 20%, 30%, 40%, and 50%. The level of total protective capacity at the non-bottleneck stations was varied at two levels: 10% and 50% of the bottleneck's capacity. The dependent variables studied were mean flow time and bottleneck shiftiness. The bottleneck shiftiness measure developed by Lawrence and Buss (1994) was used to measure the amount of shiftiness. Patterson et al (2002) conclude from their study that the amount as well as pattern of variation affect cell performance in terms of both mean flow time and bottleneck shiftiness at both levels of protective capacity. As variation increases the mean flow time also increases. The increase in mean flow time is however less when the variation is close to the bottleneck, rather than when evenly distributed or away from the bottleneck. No significant difference in the mean flow time was found for the lower levels of variation, suggesting that the placement of the variance in relationship to the bottleneck is not critical when the coefficient of variation is below 30%. Further it was found that the lower level of protective capacity is more sensitive to increased levels of variation than the higher level of protective capacity. With respect to the bottleneck shiftiness measure the importance of protective capacity is even more highlighted. With low levels of protective capacity even a 10% level of variation creates a high enough bottleneck shiftiness measure. At the 50% protective capacity level the coefficient of variation needs to be at least 30% before a high bottleneck shiftiness is reached. The pattern of variance placement has no significant impact on bottleneck shiftiness when the level of protective capacity is low and variation is high (greater than 30%). However at the lower levels of variation having the higher non-bottleneck variance close to the bottleneck reduces bottleneck shiftiness. When the level of protective capacity is high the close pattern of variance distribution again results in the least bottleneck shiftiness when the variation is greater than 10%.

2.5 Conclusions From Literature Review

The literature review in this chapter focused on the Theory of Constraints (TOC) philosophy and its two protection mechanisms, i.e. time buffers and protective capacity. The different concepts of the TOC philosophy and its DBR logistical system were explained. The importance of time buffers and protective capacity in protecting the constraint (and therefore system) throughput was explained, as well as their impact on the manufacturing lead time. It was shown that the levels of protective capacity and the lengths of the time buffers determine the total manufacturing lead time or flow time, and due to the relationship between lead times and WIP also the amount of WIP inventory in the system. The important influence that WIP has on a company's competitive elements was also shown. The efficient design of the time buffers and protective capacity is therefore an important element in the implementation of a Theory of Constraints based production system.

From the literature review it was found that very few studies have been performed on the efficient design of the time buffer lengths. Most of the literature present empirical rules for setting the initial time buffer lengths. Only two studies (Tu and Li 1998, and Radovilsky 1998) attempted to develop analytical approaches for estimating the necessary time buffer lengths. A limitation of Tu and Li's model is that they assume deterministic processing times and therefore only account for variability caused by machine breakdowns. Their procedure was also only evaluated on a simulation model with stations experiencing very high downtimes (between 20%-50%). It is unclear how the model would perform in lower downtime situations. Radovilsky's model makes the simplifying assumptions of poisson arrivals and exponential service times. It has however been found in the research literature that the lognormal distribution with less variance than the exponential distribution is more representative of many real-world processing times (Conrad 1954, Dudley 1963, Muralidhar et al 1992). Radovilsky further assumes single server stations, and he also does not model the effect that disruptions at resources upstream from the constraint have on arrivals at the constraint.

Research on the design of protective capacity in production flow lines is also limited. Most of the research on unbalanced line design focused on the so-called "bowl phenomenon" where work content is redistributed from the centre stations to the end stations while keeping the number of stations fixed. The unbalance is therefore not achieved by adding additional protective capacity to the line, and therefore the potential improvement in throughput is quite small. None of these unbalanced line studies also considered station downtime and service time variability in the same model. It has been shown in the literature (Hopp and Spearman 1996, pp.260-263)(Kadipasaoglu et al, 2000) that variability from downtimes have a much more detrimental effect on line performance than normal processing time variability. Furthermore, the main performance measure studied in these studies was the production rate. It has been shown that flow time also has an important influence on a company's competitive position.

The more recent studies that specifically focused on protective capacity design showed that adding protective capacity to a line can significantly improve both the production rate and flow time. The findings from these studies can be summarised as follows:

1) Adding protective capacity to a line with a fixed level of WIP will increase the output rate of the line.
2) Lines with protective capacity can achieve both shorter lead times and more reliable performance than lines without protective capacity, thus enabling better on-time delivery. This can be achieved with less WIP inventory.
3) In lines with protective capacity the lowest flow times and WIP are achieved by placing the constraint resource at the first station in the line. Moving the constraint further down the line increases the probability of starving the constraint, which degrades line performance.
4) Adding protective capacity reduces the adverse effect of constraint as well as non-constraint downtime on flow time and WIP.
5) The extent of benefits derived from protective capacity or constraint location depends on the level of non-constraint variability and non-constraint downtime percentage (determined as a combination of the mean time between failures and the mean time to repair a resource breakdown). Adding protective capacity is particularly helpful at the higher levels of non-constraint variability and downtime.
6) Adding protective capacity has diminishing returns on line performance. The biggest improvement in line performance from increasing protective capacity appears to occur at the lower levels of protective capacity, with the benefits from adding protective capacity diminishing at higher levels of protective capacity.
7) With regard to the arrangement or positioning of total protective capacity an equal distribution of protective capacity among non-constraint stations appears to have a slight advantage in reducing flow times. However a peak arrangement of placing the largest portion of the protective capacity immediately before and after the bottleneck appears to produce a lower bottleneck shiftiness. The research literature further suggests that placement of protective capacity at upstream workstations that feed the constraint resource is more effective as a means to reduce bottleneck shiftiness than placing the protective capacity at downstream workstations.

8) There is an inverse relationship between WIP inventory (and therefore the required time buffer size) and average protective capacity for a given level of output. This means that by designing more protective capacity into a line less WIP inventory will be required.

While all these studies proved the beneficial effect of protective capacity in improving the flow time, production rate and bottleneck shiftiness in both flow lines and job shop systems, several questions regarding the design of protective capacity still remain unanswered. One of the questions remaining is how much protective capacity is necessary. The studies by Atwater and Chakravorty (2000) as well as Kadipasaoglu et al (2000) both show that adding protective capacity has diminishing returns. It is however still unclear what minimum levels of protective capacity are necessary to protect throughput while keeping flow times and WIP low, or what levels of protective capacity can provide the biggest improvements in line performance. The effect on line performance of especially the amount of protective capacity at the secondary constraint (i.e. the second most utilised resource), as well as the location of the secondary constraint relative to the constraint resource is also uncertain. Furthermore, it is also not clear what levels of protective capacity are necessary to reduce bottleneck shiftiness and ensure a stable bottleneck location. Another shortcoming of current literature is that all the research studies on protective capacity examined flow shop or production cell situations with a fixed number of stations. None investigated assembly type of flow lines where the number of stations varies according to the distribution of the total work content among the stations.

This research study attempts to build on the existing knowledge of the protection mechanisms (time buffers and protective capacity) used in TOC based production systems by investigating certain issues regarding the design of time buffers and protective capacity. Efficient techniques for estimating the sizes of the different time buffers were identified as an important gap in current TOC research. It was therefore decided to develop and evaluate an analytical procedure that could be used in the design phase of a production line to provide quick but accurate estimates of the necessary time buffer lengths. This is described in the next chapter. The effect of protective capacity in a flow line at the secondary constraint resource was identified from the research literature as another area of TOC protection mechanisms requiring further study, as well as what minimum levels of protective capacity are needed to ensure a stable primary constraint location or protecting throughput while keeping flow times short. Chapter 4 describes an investigative study using simulation to answer some of these remaining questions regarding protective capacity design.

Chapter 3

Analytical Model For Estimating Time Buffer Lengths

3.1 Introduction

In spite of the importance of time buffers in protecting the plant throughput as well as determining the final production lead time and WIP (which can be used as a competitive advantage), it could been seen from the literature review in the previous chapter that very little research has been done to develop quick approaches for determining more accurate estimates of the time buffers. This is especially needed in situations where completely new production systems have to be designed. In such a situation no historic production lead time figures are available, and buffer management also cannot be applied since the system is not yet operational. Since the time buffer lengths also have an important influence on the final production lead time and WIP, more accurate estimates of their sizes are necessary in order to evaluate different production resource or capacity configurations. Although simulation models could be used for estimating the time buffer sizes, the development of a simulation model is expensive in time and monetary terms. Analytical models on the other hand are relatively inexpensive with respect to time and money, and are useful for quick initial estimates in order to compare alternatives and to gain insight into the network studied.

The analytical procedure was developed by investigating analytical techniques in the literature for modelling production or manufacturing systems. Although the literature on this search is not presented in this dissertation, the decision was made to use an open queuing network modelling technique based on parametric decomposition of a production network to model manufacturing networks and calculate the average flow times and variability of the flow times for parts through the network. These values could then be used for estimating the time buffer sizes. Open queuing networks are networks where customer arrivals from and departures to the outside are allowed, as opposed to closed networks where customers are not allowed to enter the system from the outside or leave the system (Papadopoulos et al 1993, p.49). One of the well-known and popular open queuing network analysis approaches is the QNA method developed by Whitt (1983a), which was used in this study to estimate the sizes of the necessary time buffers. This technique was chosen based on its relatively simple and flexible modelling approach which lends itself to be easily modified and incorporated in a computerised software tool, and experimental results from the literature that showed it to be fairly accurate compared to simulation models (Whitt 1983b; Bitran & Tirupati 1988; Segal & Whitt 1989; and

Suresh & Whitt 1989). The QNA method has been modified to enable the modelling of production networks with machine failures, batch service and varying transfer batch sizes. The modelling approach has also been incorporated in a computerised tool that uses product specific information such as Bill-of-Material (BOM) and routing data, and production network information such as resource data to estimate the sizes and location of the necessary time buffers for each product. A drum-buffer-rope production scheduling tool was also developed based on the system described by Goldratt in his book "The Haystack Syndrome" (1990). This was done to evaluate the time buffer estimation procedure when used within a DBR scheduling tool. The following section describes the developed time buffer estimation procedure. This is followed by a discussion of a simulation experiment conducted in order to evaluate the time buffer estimation procedure.

3.2 Time Buffer Estimation Procedure

3.2.1 Overview

The procedure for estimating the sizes of the time buffers is based on an open queuing network analysis of a production system. Since the time buffers are estimates of the lead time or flow time to the specific buffer origin, the queuing network analysis is used to obtain estimates of the average flow time and flow time variance to the specific buffer origin. A software program was written in Visual Basic 6.0 that collects the necessary input data, performs the queuing network analysis, and then identifies and calculates the necessary time buffer lengths.

The production operations are modelled as GI/G/m servers in a multi-product open queuing network. A parametric-decomposition approximation method is used (the QNA method developed by Whitt, 1983a) to analyse the steady-state performance of an open queuing network with non-poisson arrival processes and non-exponential service time distributions. In this approximation method the network dependence or interaction between stations are analysed by characterizing the arrival processes to each queue by two arrival parameters (mean arrival rate λ and arrival squared coefficient of variation c^2_a). The network is then decomposed into subsystems of individual stations and the individual queues at each station are analysed separately to determine the queue waiting times. After the congestion in each queue has been described, the total network flow time and flow time variance for each product are determined by assuming that the individual stations are mutually independent and summing for each product the queue times and service times at each station along its route.

The time buffer estimation procedure starts with the necessary data input. First the demand and demand variability for each product is specified by an average external arrival rate and arrival squared coefficient of variation for each product. The routing for each product through the production network is characterised by a production flow network which is determined by a computerised procedure that combines bill-of-material data and routing data given for each product as input by the user. Workstation data specifying the number of each resource as well as its Mean-Time-Between-Failures (MTBF) and Mean-Time-To-Repair (MTTR) are also provided by the user. (The specific user input requirements are specified in section 3.2.2). The user also needs to specify the capacity constrained resource (CCR).

Data for the different product classes are aggregated into a single product class. A computerised procedure then performs the open queuing network analysis on the aggregated input for this single product class. At the end of the procedure the resulting performance measures are again specified for each individual product class. Based on the identification of the capacity constrained resource by the user, the computerised procedure determines the different types of time buffers that are necessary in the production network and where (at which operations) they should be located. Using the results from the queuing network analysis the required sizes for the different time buffers are also determined by computing the average flow times and standard deviation of the flow times in the network to the specific buffer origin, and incorporating a chosen service level or confidence level. A flow diagram of the time buffer estimation procedure is displayed in Figure 3-1.

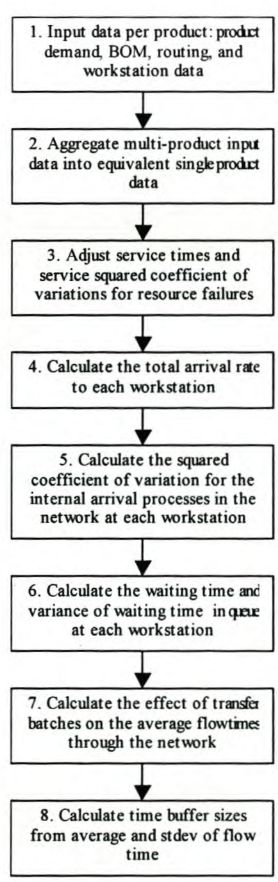

Figure 3-1. Flow diagram of time buffer estimation procedure

Each of these steps in the time buffer estimation procedure is discussed separately in the following sub-sections.

3.2.2 Data input

A graphical user interface was developed in Visual Basic 6.0 that collects the necessary input data from the user. Screenshots from the user interface are displayed in Appendix A. The following input data are needed by the time buffer estimation procedure:

a) Product demand data

- f = the different end products
- λ_{of} = the average demand rate or outside arrival rate for each product f.
- c^2_{oaf} = the outside arrival squared coefficient of variation for each product f.
- c^2_{oak} = the outside arrival squared coefficient of variation for each part k of product f (computed by the computerised procedure). $c^2_{oak} = c^2_{oaf}$ for all parts of product f.
- λ_{ok} = the average demand rate or outside arrival rate for each part k (computed by the computerised procedure).

b) Bill-of-material data

- k = the different parts a product consists of.
- The quantities needed of each part per product.
- r = the total number of parts a product consists of (computed by the computerised procedure).

c) Routing data (deterministic and not Markovian routing is used)

- τ_{kl} = the mean service time of part k at the lth operation on its route.
- c^2_{skl} = the service time squared coefficient of variation of part k at the lth operation on its route.
- b_{kl} = the process batch size of part k at the lth operation on its route (this is for workstations with batch service).
- t_{kl} = the transfer batch size of part k after the lth operation on its route.
- n_{kl} = lth operation visited by part k on its route.

- n_k = total number of operations visited by part k on its route (computed by the computerised procedure).

d) Workstation data

- j = the name of each available workstation or resource.
- N = the total number of workstations.
- m_j = the number of servers at workstation j.
- $MTBF_j$ = the mean time between failures (uptime) at workstation j.
- $MTTR_j$ = the mean time to repair (downtime) at workstation j.

3.2.3 Aggregating the input data

After the input data has been provided, the next step is to aggregate the input data for the different parts of the different products into input data for an equivalent single part. The multi-part model with deterministic routing is therefore converted into a single-part model with Markovian routing. Starting with the different end products f, the computerised procedure creates a product flow diagram for each product that combines the product, BOM, routing and workstation information for each product into a single structure. During this procedure the demand rates for the different parts of an end product are computed by multiplying the end product demand rate with the quantity per for each part at each level of the BOM.

The aggregation procedure is described below. The formulas used are obtained from Whitt (1983a). These formulas were however modified in order to account for batch process times, as well as for the different parts in the BOM structure of a product.

Let $1H(x) = 1$ if $x \in H$ and 0 otherwise.

First the aggregate external arrival rates are obtained by:

$$\lambda_{oj} = \sum_{f} \sum_{k=1}^{r} \lambda_{ok} 1\{k : n_{k1} = j\} \tag{3.1}$$

The flow rate from station i to j is then obtained by:

$$\lambda_{ij} = \sum_{f} \sum_{k=1}^{r} \sum_{l=1}^{n_k - 1} \lambda_{ok} 1\{(k,l) : n_{kl} = i, n_{k,l+1} = j\} \tag{3.2}$$

These values are used to develop the routing matrix with the proportion of parts that go from i to j :

$$q_{ij} = \lambda_{ij} \Big/ \left(\lambda_{io} + \sum_{z=1}^{N} \lambda_{iz} \right) \tag{3.3}$$

The individual service times per part are then aggregated into a single mean service time at each workstation j :

$$\tau_j = \frac{\sum_f \sum_{k=1}^{r} \sum_{l=1}^{n_k} \left(\frac{\lambda_{ok}}{b_{kl}} \right) \tau_{kl} 1\{(k,l) : n_{kl} = j\}}{\sum_f \sum_{k=1}^{r} \sum_{l=1}^{n_k} \left(\frac{\lambda_{ok}}{b_{kl}} \right) 1\{(k,l) : n_{kl} = j\}} \tag{3.4}$$

The individual per part squared coefficient of variation (scv) of service times are also aggregated into a single scv of the service time at workstation j. This formula was obtained by Whitt by using the property that the second moment of a mixture of distributions is the mixture of the second moments. The second moment around the origin [$E(X^2)$] is equal to $var(X) + E(X)^2 = \tau^2 c^2 + \tau^2$:

$$\tau_j^2 \left(c_{sj}^2 + 1 \right) = \frac{\sum_f \sum_{k=1}^{r} \sum_{l=1}^{n_k} \left(\frac{\lambda_{ok}}{b_{kl}} \right) \tau_{kl}^2 \left(c_{skl}^2 + 1 \right) \{(k,l) : n_{kl} = j\}}{\sum_f \sum_{k=1}^{r} \sum_{l=1}^{n_k} \left(\frac{\lambda_{ok}}{b_{kl}} \right) 1\{(k,l) : n_{kl} = j\}} \tag{3.5}$$

Finally the aggregated scv (c_{oj}^2) of the external arrivals to each workstation is calculated by:

$$c_{oj}^2 = (1 - w_j) + w_j \left[\sum_f \sum_{k=1}^{r} c_{oak}^2 \left(\lambda_{ok} 1\{k : n_{k1} = j\} \Big/ \sum_{l=1}^{r} \lambda_{ol} 1\{k : n_{k1} = j\} \right) \right] \tag{3.6}$$

where

$$w_j = \left[1 + 4(1 - \rho_j)^2 (v_j - 1) \right]^{-1}$$

$$v_j = \left[\sum_f \sum_{k=1}^{r} \left(\lambda_{ok} 1\{k : n_{k1} = j\} \Big/ \sum_f \sum_{l=1}^{r} \lambda_{ol} 1\{l : n_{l1} = j\} \right)^2 \right]^{-1}$$

A hybrid approximation for the superposition of arrival processes is used here by Whitt because the external arrival process to station j is the superposition of the external arrival processes to station j from the different product groups. The variables w_j and v_j are weights used in convex combinations arising in hybrid approximations for superpositions.

3.2.4 Adjust the service times and service squared coefficient of variations for resource failures

In order to model the effect of resource failures or other service interruptions, the service times and service squared coefficient of variations at each workstation are adjusted with the server availability parameters at each workstation. The approach for modelling resources failures as presented by Segal and Whitt (1989) is used in the time buffer estimation procedure. In this approach it is assumed that the up and down times for all workstations and all servers at a workstation are mutually independent. Further it is assumed that the up times for all servers at a workstation have a common distribution, partially characterised by its mean. Similarly, it is also assumed that the down times for all servers at a workstation have a common distribution, partially characterised by its mean and squared coefficient of variation. In the time buffer estimation procedure exponential up and down times are assumed, therefore the squared coefficient of variation for the down times are one. Although the server availability parameters must be the same for all servers at a workstation, it can vary from workstation to workstation.

The model further assumes that down times are triggered by service times. Upon starting service at a workstation, each part causes a down time with probability *h*. The part therefore experiences an expanded service time equal to the original service time plus an independent down time, and has an ordinary service time with probability 1 – *h*. The expanded service times τ_{ej} and service squared coefficients of variation c^2_{sej}, as well as the new traffic intensity or utilisation ρ_{ej} at workstation *j* are calculated as follows:

$$\rho_{ej} = \rho_j + MTTR_j / \left(MTTR_j + MTBF_j\right) \tag{3.7}$$

$$\tau_{ej} = \tau_j + h_j MTTR_j \tag{3.8}$$

$$\tau_{ej}^2\left(c_{sej}^2 + 1\right) = \left(c_{sj}^2 + 1\right)\tau_j^2 + h_j\left[2MTTR_j^2 + 2MTTR_j\tau_j\right] \tag{3.9}$$

where

$$h_j = \min\left\{m_j / \left(\lambda_j\left(MTTR_j + MTBF_j\right)\right), 1\right\}$$

3.2.5 Calculate the total arrival rate to each workstation

After the input data have been aggregated for an equivalent single part (product) and adjusted for resource failures, the total arrival rate λ_j to each workstation is calculated. This is done by solving the following linear equations (the traffic rate equations):

$$\lambda_j = \lambda_{oj} + \sum_{i=1}^{N} \lambda_i q_{ij} \qquad (3.10)$$

Using the total arrival rate λ_j to each workstation, the traffic intensities or utilisations ρ_j at each workstation can then be calculated:

$$\rho_j = \lambda_j \tau_j / m_j \qquad (3.11)$$

The proportion of arrivals to workstation *j* that come from workstation *i* can also be calculated:

$$\mathrm{p}_{ij} = \lambda_{ij} / \lambda_j \qquad (3.12)$$

3.2.6 Calculate the squared coefficient of variation for the internal arrival processes in the network at each workstation

In order to capture the effects of variability and disruptions at upstream workstations on downstream workstations, an arrival squared coefficient of variation c_{aj}^2 is calculated. This measure captures the variability of the flow from one workstation to another. The formula used to calculate the c_{aj}^2's is a combination of the formulas proposed in Whitt (1983a) and Segal and Whitt (1989). These formulas are approximations developed by Whitt based on two methods for approximating a point process by a renewal process: the asymptotic method and the stationary-interval method (Whitt, 1982). If c_{di}^2 is the variability parameter of the overall departure process from workstation *i*, q_{ij} is the proportion of departures from workstation *i* to workstation *j*, and deterministic routing is used, then the variability parameter c_{ij}^2 for the flow going from workstation *i* to *j* is:

$$c_{ij}^2 = q_{ij} c_{di}^2 + (1 - q_{ij}) q_{ij} c_{ai}^2 + (1 - q_{ij})^2 c_{ei}^2 \qquad \text{(from Segal and Whitt, 1989)} \qquad (3.13)$$

where

$$c_{ei}^2 = \frac{\sum_f \sum_{k=1}^{r} g_{ki} c_{oak}^2}{\sum_f \sum_{k=1}^{r} g_{ki}}$$

$$c_{di}^2 = 1 + (1 - \rho_i^2)(c_{ai}^2 - 1) + \frac{\rho_i^2}{\sqrt{m_i}} (\max\{c_{si}^2, 0.2\} - 1) \qquad (3.14)$$

g_{ki} is the number of visits of part *k* to station *i*

The system of equations yielding the scv for the arrival processes c_{aj}^2 at each workstation is obtained as follows:

$$c_{aj}^2 = 1 - w_j + w_j \sum_{i=0}^{N} \mathrm{p}_{ij} c_{ij}^2 \qquad \text{(from Whitt, 1983a)} \qquad (3.15)$$

Substituting the above expressions for c_{ij}^2 and c_{di}^2, the following set of linear equations is obtained after some manipulation:

$$c_{aj}^2 = \alpha + \sum_{i=1}^{N} \beta c_{ai}^2 \qquad (3.16)$$

where

$$\alpha = 1 - w_j + w_j \mathrm{p}_{oj} c_{oj}^2 + w_j \sum_{i=1}^{N} \mathrm{p}_{ij} \left(q_{ij} \rho_i^2 \left[1 + m_i^{-0.5} \left(\max\left\{ c_{si}^2, 0.2 \right\} - 1 \right) \right] + (1 - q_{ij})^2 c_{ei}^2 \right)$$

$$\beta = w_j \mathrm{p}_{ij} q_{ij} \left[\left(1 - \rho_i^2\right) + (1 - q_{ij}) \right]$$

$$w_j = \left[1 + 4\left(1 - \rho_j\right)^2 \left(v_j - 1\right) \right]^{-1}$$

$$v_j = \left[\sum_{i=0}^{N} \mathrm{p}_{ij}^2 \right]^{-1}$$

These linear equations are then solved using matrix algebra in order to obtain c_{aj}^2. Again the variables v_{ij} and w_j are weights or probabilities that are used in convex combinations arising in hybrid approximations for departure and superpositions respectively.

3.2.7 Calculate the waiting time and variance of waiting time in queue at each workstation

After the arrival variability at each workstation has been determined, the next step is to calculate the average waiting time in queue a part experience at each workstation, as well as the variance of the waiting time. A workstation is modelled as a *GI/G/m* queue, which has *m* identical servers in parallel, unlimited waiting room, first-come-first-served queue discipline, and inter-arrival and service times that come from independent identically distributed random variables with general distributions. The approximation used in the time buffer estimation procedure for calculating the waiting times and the variances of the waiting times in *GI/G/m* queues, is the one developed by Whitt (1993). According to this approximation the average expected waiting time in the queue at workstation j (EW_j) is determined by:

$$EW_j = \phi\left(\rho_j, c_{aj}^2, c_{sj}^2, m_j\right)\left(\frac{c_{aj}^2 + c_{sj}^2}{2}\right)EW_j\left(M / M / m\right) \tag{3.17}$$

where

$$\phi\left(\rho_j, c_{aj}^2, c_{sj}^2, m_j\right)$$

$$= \begin{cases} \left(\frac{4\left(c_{aj}^2 - c_{sj}^2\right)}{4c_{aj}^2 - 3c_{sj}^2}\right)\phi_1\left(m_j, \rho_j\right) + \left(\frac{c_{sj}^2}{4c_{aj}^2 - 3c_{sj}^2}\right)\psi\left(\frac{c_{aj}^2 + c_{sj}^2}{2}, m_j, \rho_j\right), & c_{aj}^2 \geq c_{sj}^2 \\ \left(\frac{c_{sj}^2 - c_{aj}^2}{2c_{aj}^2 + 2c_{sj}^2}\right)\phi_3\left(m_j, \rho_j\right) + \left(\frac{c_{sj}^2 + 3c_{aj}^2}{2c_{aj}^2 + 2c_{sj}^2}\right)\psi\left(\frac{c_{aj}^2 + c_{sj}^2}{2}, m_j, \rho_j\right), & c_{aj}^2 \leq c_{sj}^2 \end{cases}$$

$$\psi\left(c^2, m.\rho\right) = \begin{cases} 1, & c^2 \geq 1 \\ \phi_4\left(m, \rho\right)^{2\left(1-c^2\right)} & 0 \leq c^2 \leq 1 \end{cases}$$

$$\phi_4\left(m, \rho\right) = \min\left\{1, \frac{\phi_1\left(m.\rho\right) + \phi_3\left(m, \rho\right)}{2}\right\}$$

$$\phi_3\left(m, \rho\right) = \phi_2\left(m, \rho\right)\exp\left(-2\left(1-\rho\right)/3\rho\right)$$

$$\phi_2\left(m, \rho\right) = 1 - 4\gamma\left(m, \rho\right)$$

$$\phi_1\left(m, \rho\right) = 1 + \gamma\left(m, \rho\right)$$

$$\gamma\left(m, \rho\right) = \min\left\{0.24, \frac{\left(1-\rho\right)\left(m-1\right)\left(\left(4+5m\right)^{0.5} - 2\right)}{16m\rho}\right\}$$

$$EW_j\left(M / M / m\right) = \frac{\tau_j P\left(N \geq m_j\right)}{m_j\left(1-\rho\right)}$$

$$P\left(N \geq m_j\right) = \left[\frac{\left(m_j \rho_j\right)^{m_j}}{m_j!\left(1-\rho\right)}\right]\xi$$

$$\xi = \left[\frac{\left(m_j \rho_j\right)^{m_j}}{m_j!\left(1-\rho_j\right)} + \sum_{h=0}^{m_j - 1}\frac{\left(m_j \rho_j\right)^h}{h!}\right]^{-1}$$

The variance of the expected waiting time in the queue at workstation j [Var(W_j)] is determined by:

$$\mathrm{Var}\left(W_j\right) = \left(EW_j\right)^2 c_{W_j}^2 \tag{3.18}$$

where

$$c_{W_j}^2 = \text{squared coefficient of variation of the waiting time} = \frac{c_{D_j}^2 + 1 - P\left(W_j > 0\right)}{P\left(W_j > 0\right)}$$

$$c_{D_j}^2 = 2\rho_j - 1 + \frac{4(1-\rho_j)d_{sj}^3}{3(c_{sj}^2+1)^2} \quad \text{and } D_j = (W_j | W_j > 0)$$

$$d_s^3 = \begin{cases} 3c_{sj}^2(1+c_{sj}^2) & c_{sj}^2 \geq 1 \\ (2c_{sj}^2+1)(c_{sj}^2+1) & c_{sj}^2 < 1 \end{cases}$$

$$P(W_j > 0) \approx \min\{\pi, 1\}$$

where

$$\pi = \begin{cases} \pi_1, & \text{if } m_j \leq 6 \text{ or } \gamma_j \leq 0.5 \text{ or } c_{aj}^2 \geq 1 \\ \pi_2, & \text{if } m_j \geq 7 \text{ and } \gamma_j \geq 1 \text{ and } c_{aj}^2 < 1 \\ \pi_3, & \text{if } m_j \geq 7 \text{ and } c_{aj}^2 < 1 \text{ and } 0.5 < \gamma_j < 1 \end{cases}$$

$$\gamma_j = (m_j - m_j\rho_j - 0.5) / \sqrt{m_j \rho_j z_j}$$

$$z_j = (c_{aj}^2 + 1)/2$$

$$\pi_1 = \rho_j^2 \pi_4 + (1-\rho_j^2)\pi_5$$

$$\pi_2 = c_{aj}^2 \pi_1 + (1-c_{aj}^2)\pi_6$$

$$\pi_3 = 2(1-c_{aj}^2)(\gamma_j - 0.5)\pi_2 + (1 - [2(1-c_{aj}^2)(\gamma_j - 0.5)])\pi_1$$

$$\pi_4 = \min\left\{1, \frac{1-\Phi((1+c_{sj}^2)(1-\rho_j)m^{0.5}/(c_{aj}^2+c_{sj}^2))}{1-\Phi((1-\rho_j)m^{0.5})} P(W(M/M/m)>0)\right\}$$

$$\pi_5 = \min\left\{1, \frac{1-\Phi(2(1-\rho_j)m_j^{0.5}/(1+c_{aj}^2))}{1-\Phi((1-\rho_j)m_j^{0.5})} P(W(M/M/m)>0)\right\}$$

$$\pi_6 = 1 - \Phi((m_j - m_j\rho_j - 0.5)/\sqrt{m_j \rho_j z_j})$$

Φ is the standard normal CDF

3.2.8 Calculate the effect of transfer batches on the average flow times through the network

A transfer batch can be defined as "the quantity of units that are moved at the same time from one resource to the next" (Srikanth & Umble, 1997, p. 163). The effect of transfer batches on the average flow times is modelled based on the approach presented by Hopp & Spearman (1996, pp.290-291). According to this approach there are two main influences on the flow time caused by a transfer batch between two stations:

1. The first effect is the waiting time for a transfer batch to form after the first operation.
2. The second effect is caused by the arrival of batches at the second station.

3.2.8.1 Batch waiting time after first operation

When parts are moved in a transfer batch between two stations, the parts have to wait after the first station for a transfer batch of size t to form. The first part arriving at the transfer batch has to wait for t-1 other parts, while the last part arriving does not wait at all. The average number of parts waiting in the batch is therefore $(t-1)/2$. Because of the conservation of flow in queuing networks, parts arrive to the transfer batch at the same rate as they arrive to the station itself. The average time spent forming the transfer batch is therefore the average number of parts waiting in a batch times the average time between arrivals of parts:

Average waiting time for part k to form transfer batch t_{kl} after the lth operation at station i

$$= \lambda_i^{-1}(t_{kl}-1)/2$$

This analysis however assumes that parts arrive at the transfer batch as individual units. When there is a process batch operation at the station feeding the transfer batch, parts arrive at the transfer batch in batches of size b_{kl}. The effective transfer batch size is therefore t_{kl}/b_{kl}. When this effective transfer batch size is ≤ 1, then the number of parts waiting for a transfer batch is zero and there is no batching time. Otherwise this effective transfer batch size is used. Since the parts arrive in batches of size b_{kl}, the average number of parts waiting for a transfer batch to form is now

$$\frac{b_{kl}\left(\dfrac{t_{kl}}{b_{kl}}-1\right)}{2}$$

The average transfer batching time is now calculated as:

Average waiting time for part k to form effective transfer batch size t_{kl}/b_{kl} after the lth operation at station i

$$= \lambda_i^{-1} b_{kl}\left(\frac{t_{kl}}{b_{kl}}-1\right)\Big/2 = \lambda_i^{-1}(t_{kl}-b_{kl})/2 \tag{3.19}$$

3.2.8.2 Batch arrivals at second operation

To calculate the effect of transfer batch arrivals at the second operation, it can be viewed as a queue of whole batches, a queue of single parts (a partial batch), and a server (refer to Figure 3-2).

The queue of whole batches can be treated as a queue of single parts with an effective process time of $t\tau_j$. Hopp and Spearman (1996, pp.290-291) show that the waiting time in the queue of whole batches is the same as the waiting time in queue for single parts. The only addition to the total flow time at the second operation due to transfer batch arrivals is therefore caused by the mean time spent in the partial batch. The first part arriving in a batch of size t at an idle server would not have to wait, whereas the last part in the batch would have to wait for t-1 other parts to finish processing. The average number of parts therefore that has to wait in the partial batch is $(t-1)/2$, and the average time waiting in the partial batch is $(t-1)\tau_j/2$.

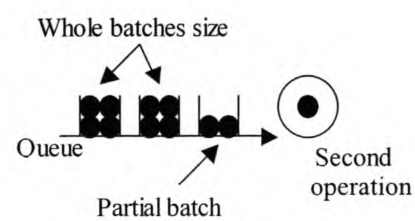

Figure 3-2. Transfer batch arrivals at second operation

The above analysis however assumes that the second operation processes parts individually. When the second operation is in fact a batch service operation with process batch size b_{kl+1}, then an effective partial batch size needs to be calculated. This effective partial batch size can be calculated as t_{kl}/b_{kl+1}. If this effective partial batch is ≤ 1, then there is no waiting time in the partial batch and the partial batch size is taken as 1. If the effective partial batch size however is > 1, then the waiting time in the partial batch is calculated as :

$$\text{Waiting time in partial batch} = \frac{(t_{kl}/b_{kl+1})-1}{2}\tau_j \qquad (3.20)$$

3.2.9 Calculate the sizes of the different time buffers for each product

After all the queuing and batching times have been calculated in the network, the next step is to calculate the sizes of the different time buffers. For each end product, the computerised procedure goes through the product flow network and identifies the types and locations of the necessary time buffers, as well as the longest operational paths (in terms of total flow time) feeding the specific time buffer. Refer to Figure 3-3. The dotted lines indicate all the operations along the longest flow time path feeding the different time buffers. All the calculated queuing times and batching times, as well as the service times at all the operations along the longest flow time path feeding the specific time buffer, are added. This gives the average total flow time for a part along the longest flow time path to the specific time buffer origin. To calculate the variances for these flow times it is assumed that the sojourn times (service time plus waiting time) at successive stations are independent (an approximation), and therefore the variances of the service times as specified in the model input plus the calculated variances of the queuing times at each station along the longest flow time path feeding the specific time buffer, are added.

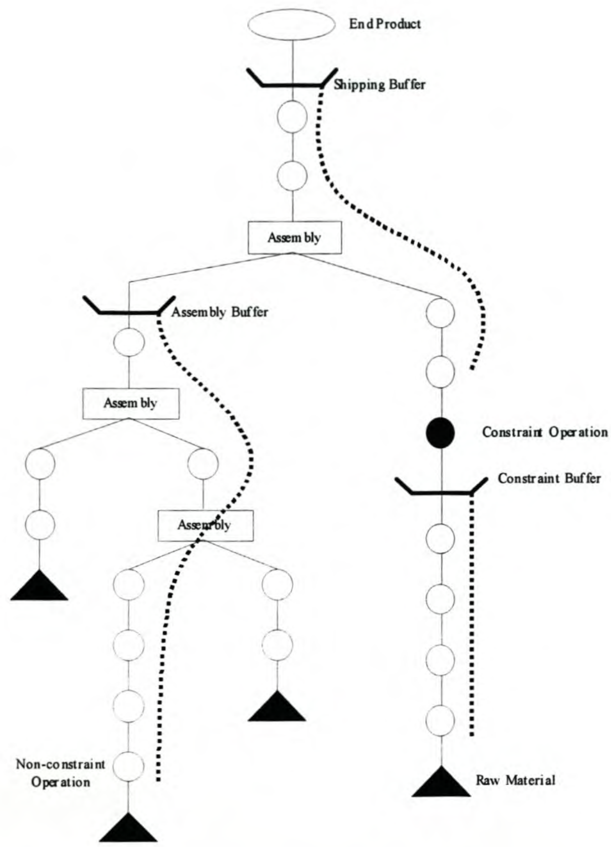

Figure 3-3. Production flow network indicating the different time buffer computation paths

Using only the average flow times for calculating the time buffer lengths are not sufficient. Since these are only average values, the variability of these flow times also need to be taken into account in order to ensure a high protection against disruptions. Using the central limit theorem, the flow time for a product can be approximated by the normal distribution with mean the average flow time and variance the square of the standard deviation of the flow time. By deciding on a certain service level s and reading the corresponding z-value from a standard normal table, the total time buffer length can be computed as:

$$\textit{Time buffer length} = \textit{AverageFlowTimeToBufferOrigin} + z_s(\textit{StdDevOfFlowTimeToBufferOrigin}) - \textit{SumOfProcessingTimesToBufferOrigin}$$

The sum of the processing times is subtracted from the total flow time calculation in order to obtain the variability component of the flow time. This is necessary for time buffers used in a computerised scheduling procedure that performs a capacity analysis and therefore adds the deterministic processing times to the time buffer. The time buffer therefore only needs to include the variability component of the flow time. For manual DBR implementations the time buffer length can include the sum of the processing times.

3.3 Evaluating the Time Buffer Estimation Procedure

Several studies have shown that the parametric decomposition queuing network analysis approach developed by Whitt compare relatively well to simulation models (Whitt, 1983b; Bitran & Tirupati, 1988; Segal & Whitt, 1989; and Suresh & Whitt, 1989). A simulation study by Segal and Whitt (1989) shows that it performs quite well in estimating the average flow time in a stochastic manufacturing network. Segal and Whitt evaluated the approach on two simulation models of manufacturing lines. Both lines were single product lines with the one line consisting of 67 workstations and the other line of 30 workstations. Some of the stations experienced breakdowns. For the one line, the percentage difference between the average flow time calculated by the queuing network analysis approach and the value obtained from the simulation experienced was only 5.6%, while for the other line the percentage difference was only 2.4%. All these simulation models were however rather simplistic in that stochastic arrival rates were used for releasing jobs (and not material release schedules); jobs immediately proceeded to the next station after being processed (no transfer batches); and all stations consisted of single servers.

In order to evaluate the performance of the time buffer estimation procedure in a more complex real life industrial setting, a simulation experiment on an actual multi-product 16 station flow shop was performed. The simulation model was made more realistic by releasing products to the flow shop model according to a material release schedule developed by a DBR scheduling program, and not using stochastic arrival rates as done in most simulation studies. The simulation model is described in the following section.

3.3.1 Simulation model description

The model used in the simulation experiment is based on data from an actual production facility. This plant is a medium-sized manufacturer of electrical equipment with two main product groups. The product flows or routings for the two product groups are presented in Table 3-1 and Table 3-2. Because all the products basically follow the same sequence of operations, the plant can be classified as a flow shop.

Table 3-1. Product 1 routing information

Operation Nr	Workstation	Mean Service Time (min)	Std Dev (min)	Process Batch	Transfer Batch
1	1	1.5	0.9	1	1
2	2	0.65	0.52	1	8
3	3	0.7	0.28	1	16
4	4	0.467	0.2335	1	48
5	5	2	0.4	8	1
6	6	0.65	1.04	1	1
7	7	0.6	0.5367	1	25
8	8	0.8	0.48	1	1
9	9	1.565	0.7825	6	50
10	10	240	0	50	50
11	11	0.6	0.536	1	1
12	12	2.123	1.2738	1	1
13	13	0.5	1	1	1
14	14	0.25	0.125	1	1
15	15	0.3	0.24	1	1
16	16	0.28	0.448	1	1

Table 3-2. Product 2 routing information

Operation Nr	Workstation	Mean Service Time (min)	Std Dev (min)	Process Batch	Transfer Batch
1	1	1.5	0.9	1	1
2	2	0.65	0.52	1	8
3	3	0.7	0.28	1	16
4	4	0.467	0.2335	1	48
5	5	2	0.4	8	8
6	6	0.65	1.04	1	1
7	7	0.6	0.72	1	100
8	10	240	0	100	100
9	8	0.63	0.378	1	1
10	9	2.12	1.06	6	6
11	12	2.123	1.2738	1	1
12	13	0.5	1	1	1
13	14	0.25	0.125	1	1
14	15	0.3	0.24	1	1
15	16	0.28	0.448	1	1

The Arena 3.0 simulation software was used to model the flow shop. The shop was modelled with resource failures, different transfer batch sizes between workstations, as well as batch processing at certain workstations. Some workstations also consist of more than one server that was modelled individually. Material handling times between workstations were assumed negligible. Unlimited queue sizes were assigned between stations. This prevented any blocking between stations to occur because of unavailable waiting room in front of a station. A flow diagram of the developed simulation model of the flow shop is displayed in Figure 3-4.

The main entities flowing through the model are the two end product groups. Entities are released to the model according to a material release schedule developed with a DBR scheduling program. The material release schedule was exported to an Excel spreadsheet file and re-organised in a format suitable for the simulation model to be read as input. The material release schedule specifies the daily orders and order quantities that need to be released. At the start of each simulation day the simulation model reads the orders to be released for the day, and releases an amount of entities equal to the release quantity to the model. Entities are first sent to a buffer and then released to station 1 at the fixed maximum throughput rate of the capacity constrained resource. Each released entity is assigned an attribute specifying the end product ID, the part type, the order ID, and the time of entry into the model.

Upon arrival at a station with more than one server, entities are sent to the available server with the shortest queue. A batch module is provided after stations with a transfer batch of greater than one. Entities wait at this batch module for a transfer batch to form before being sent to the next station. Upon arrival at the following station the batch is split up again, except if it is a batch processing server that process batches of entities instead of individual entities. To prevent entities from being stuck in a batch module waiting for other entities to make up the transfer batch when resource breakdowns occur or when it is the last entity in an order, the transfer batch modules are checked every 30 minutes and if entities have been waiting in the batch module for longer than 120 minutes, they are moved to the next station.

Station 5 is modelled as a conveyor with a certain speed and length on which batches of parts are moved, whereas stations 9 and 10 are batch processing servers that process batches of entities instead of individual entities.

Processing times for all workstations were assumed to be lognormal. It has been shown in other studies (Atwater and Chakravorty 1994; Muralidhar et al 1992) that the lognormal distribution is a good representation of real-world processing times. The mean service times, standard deviations of service times, process batch sizes and transfer batch sizes for the two product flows are presented in Table 3-1 and Table 3-2.

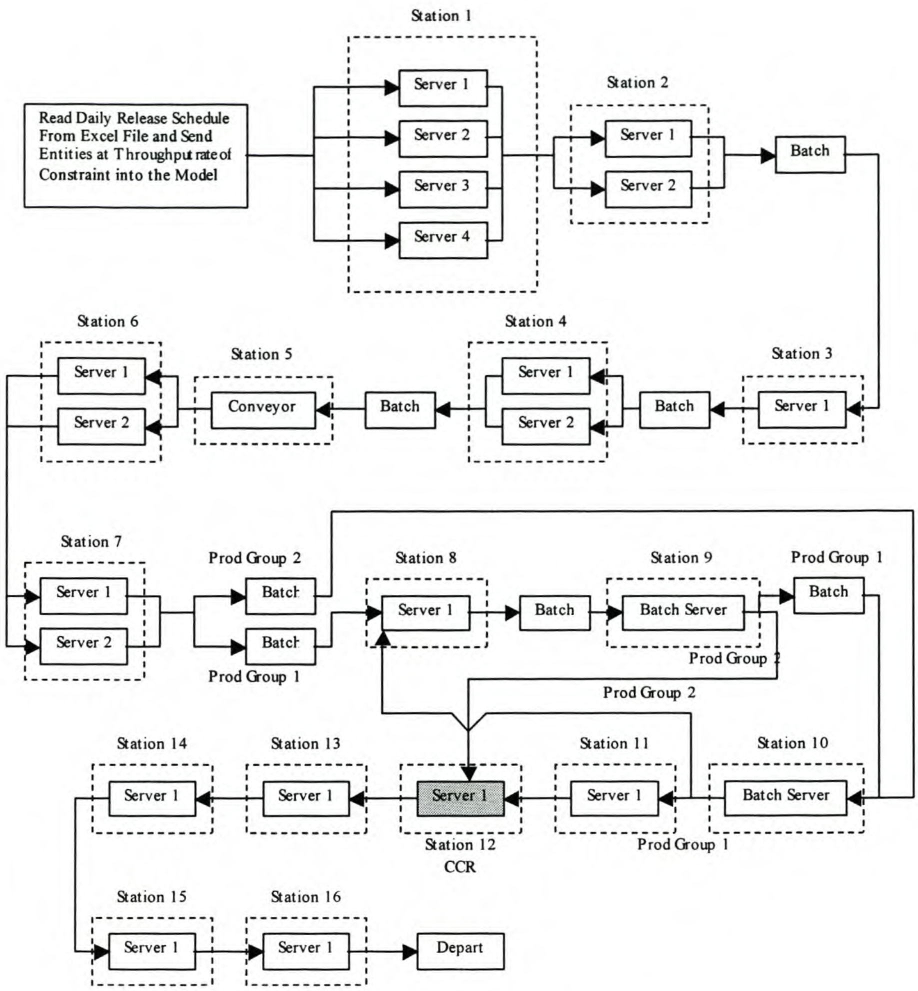

Figure 3-4. Simulation model of flow shop

Three different variability situations were modelled, i.e. where the resources have high downtimes, medium downtimes, and low downtimes (refer to Table 3-3). This was done to evaluate the accuracy of the model under different levels of variability. Customer orders were created for a three and half month period. Order quantities were chosen to ensure an average utilisation of 95% at the capacity constrained resource. Workstation 12 had the highest utilisation (and therefore the lowest throughput rate) and was used as the capacity constrained resource or CCR. Average per minute demand rates per product were calculated from the order data, as well as the squared coefficient of variation of the demand rates. These were used as the arrival rates and variability of the arrival rates in the time buffer queuing analysis procedure (refer to Table 3-4).

Table 3-3. Workstation data

		High Downtime			Medium Downtime			Low Downtime		
Workstation	Nr of Servers	MTBF (min)	MTTR (min)	% Downtime	MTBF (min)	MTTR (min)	% Downtime	MTBF (min)	MTTR (min)	% Downtime
1	4	-	-	0	-	-	0	-	-	0
2	2	10000	960	8.76	15000	480	3.10	25000	480	1.88
3	3	-	-	0	-	-	0	-	-	0
4	2	20000	960	4.58	20000	960	4.58	30000	480	1.57
5	1	-	-	0	-	-	0	-	-	0
6	2	-	-	0	-	-	0	-	-	0
7	2	13000	1440	9.97	15000	480	3.10	20000	360	1.77
8	2	-	-	0	-	-	0	-	-	0
9	1	18000	960	5.06	30000	960	3.10	30000	300	0.99
10	10	-	-	0	-	-	0	-	-	0
11	2	40000	300	0.74	40000	300	0.74	40000	300	0.74
12	5	40000	200	0.50	40000	200	0.50	40000	200	0.50
13	2	-	-	0	-	-	0	-	-	0
14	1	25000	1200	4.58	20000	480	2.34	22000	360	1.61
15	1	-	-	0	-	-	0	-	-	0
16	1	-	-	0	-	-	0	-	-	0

Table 3-4. Product arrival rates

Product	Arrival Rate	Arrival SCV
1	0.5392	1.136139
2	1.699142	0.168004

In order to accurately represent an actual production facility, the model was simulated with material release schedules as well as production schedules for the capacity constrained resource that were generated with a computerised drum-buffer-rope production scheduler. This drum-buffer-rope production scheduler was also developed in Visual Basic 6.0 and is based on the computerised drum-buffer-rope scheduling procedure described by Goldratt in his book "The Haystack Syndrome" (1990). The developed DBR scheduling program is described in more detail in Appendix C. This procedure schedules the orders on the capacity constrained resource based on an earliest due date rule. The scheduling procedure first places the orders on the constraint resource in order of ascending due dates, while allocating any end product or work in process inventory. No capacity is considered at this stage. Loads are placed at a certain time offset before the order's due date. The length of this time offset is equal to the corresponding shipping buffer plus the sum of the processing times for one unit from the constraint resource to the shipping buffer. The next step involves a backward pass where the loads are levelled backwards in time according to the available capacity. This is followed by a forward pass to remove any conflicts, such as loads been pushed into the past (past the schedule start date). The developed primary CCR schedule used in the simulation experiment is presented in Appendix B.4.

The material release schedules are then determined by a subordination procedure. The procedure can be described as a simple bin-loading algorithm treating the loads as a fluid that flows back from one day to the next. The procedure starts at the scheduled start time on the constraint of the latest customer order and moves consistently backward in time. A constraint time buffer is subtracted from the order's start time on the constraint resource and the load is placed on the feeding non-constraint resources. Taking into account the capacity of the non-constraint resources, any overloads are shifted back in time. The material release schedules developed for the different downtime situations are presented in Appendix B.

All the time buffers used in the DBR scheduling procedure were calculated with the time buffer queuing analysis procedure described in section 3.2. The data for each downtime situation of the flow shop were used as input to the time buffer queuing analysis software tool and the average flow times, standard deviation of the flow times and the final time buffer lengths for each product type calculated. Initial runs with the simulation model showed that the time buffer analysis procedure seems to overestimate the average flow time and standard deviation of the flow time for the high downtime situation, whereas it underestimates these values for the medium and low downtime situations. Different z-values were therefore used for the high downtime situation and the medium and low downtime situations in calculating the time buffer lengths. To ensure a high protection it was decided to use a 99% z-value for the high downtime situation and a 99.98% z-value for the medium and low downtime situations. In the simulation study the time buffers were used in a computerised DBR scheduling procedure that during subordination adds the processing times along a route to the time buffer lengths provided as input to the system. For use in a computerised DBR scheduling procedure the sum of the processing times along a route to the constraint buffer origin is therefore subtracted from the total flow time to the buffer origin in order to obtain the time buffer length. The final time buffer lengths used therefore only consisted of a variability component. The calculated constraint buffers for the different downtime situations are displayed in Table 3-5, along with corresponding inventory levels for the time buffers when converted through the use of Little's formula. This table also shows the ratio of the calculated time buffer length to the sum of the raw processing times to the specific buffer origin. This gives an indication of the inflation in flow time protection needed because of variability present. These values clearly show the higher levels of protection needed for higher variability situations.

Table 3-5. Constraint buffer lengths

	Product 1			Product 2		
	High Downtime	Medium Downtime	Low Downtime	High Downtime	Medium Downtime	Low Downtime
Constraint Buffer (minutes)	4600 min	1700 min	525 min	4620 min	1720 min	546 min
Constraint Buffer (days)	9.6 days	3.5 days	1.1 day	9.6 days	3.6 days	1.1 day
Constraint Buffer (inventory units)	10290 units	3804 units	1175 units	10337 units	3848 units	1222 units
Constr. Buffer/ Raw Process Time Ratio	18.44	6.81	2.1	18.55	6.91	2.19

3.3.2 Model time span

Since a particular set of orders with a specific material release schedule was used, it was decided to analyse the flow shop as a terminating system. An empty plant is used as the starting condition, and the completion of all the scheduled orders for a three and a half month period is the terminating condition. Data is therefore collected by using independent replications of this terminating system. The three and half month period was arbitrarily chosen to ensure a long enough time span for resource breakdowns to occur and to process a sufficient number of orders.

3.3.3 Performance measures studied

Four performance measures or dependent variables were studied, i.e. the average flow time to the constraint resource (per product), the standard deviation of the flow time to the constraint resource (per product), the number of late arrivals at the constraint resource, as well as the average lateness of late orders at the constraint resource. Only the constraint buffer length was therefore evaluated. The flow time for a part was measured as the time between the arrival of a part at the queue of station 1 and the processing start time for the part at the constraint station (station 12). The percentage of late orders as well as the average lateness of late orders at the constraint resource were determined by writing the process finish time at the constraint resource for the last part in each order to an Excel spreadsheet file and comparing it to the scheduled order completion time at the constraint resource as determined by the DBR scheduling procedure. The average flow times to the constraint resource as well as the standard deviation of the flow times were compared to the corresponding measures as calculated by the time buffer queuing analysis procedure.

3.3.4 Preliminary statistical analysis

After the simulation models were developed tests were done to establish the length of the warm-up period. This was done because the simulation model starts out empty and idle. The truncated-replication approach suggested by Kelton et al (1998, pp.219-224) for terminating simulations was used. The high down time simulation model was chosen because of its high variability and 10 runs of 48000 minutes (nearly 3.5 months) were made. The flow time measures were plotted and the time at which the flow time appears to stabilise marked. Using this procedure the warm-up period was determined to be 10000 minutes for each downtime situation. Simulations were therefore run for 10000 minutes after which statistics were cleared and data was collected over the next 38000 minutes.

Because a terminating simulation was used, the number of replications needed to obtain a certain accuracy in the dependent variables had to be obtained. Ten initial replications were run and the number of replications needed was computed using the confidence interval half-width formula from Kelton et al (1998, p.185):

$$n^{*} = n \times \left(\frac{h}{h^{*}} \right)^{2}$$

where

n^{*} = The required number of replications.

n = The number of initial replications (10 in this case).

h = Calculated confidence interval half width obtained from the initial run.

h^{*} = Desired confidence interval half width based on the variable that was studied and the accuracy wanted for that specific variable.

This formula is applied for each of the performance measures under study. The final number of replications to perform is then determined as the largest required number of replications from the set of output measures. Using this approach it was determined that to be 95% confident of a 2% precision level on dependent variables 500 replications were required for the high and medium downtime situations respectively, and 300 replications for low downtime situation.

3.3.5 Simulation results

The results from the simulation study are displayed in Table 3-6. The average flow time and standard deviation calculations are not always that accurate compared to the simulation values. It should however be remembered that the simulation model used in the study was a very close representation of a real-life plant's operations, therefore this was to be expected. It seems as if the time buffer estimation procedure overestimates the mean flow times and standard deviations for high downtime situations. In terms of the average flow times to the constraint resource, the best results were obtained for the medium downtime situation. The difference between the simulated values and the values calculated by the time buffer estimation procedure is only 5.35% for product 1 and 5.63% for product 2. Less accurate results were obtained for the standard deviation of the flow times to the constraint. The maximum percentage difference between the simulated and calculated values are 55% for the product 1 medium downtime situation, whereas the minimum percentage difference is 27% for the low downtime situation. In terms of the percentage of late orders and average lateness of late orders completed at the constraint resource, the results are quite promising.

Table 3-6. Simulation results and comparison with time buffer analysis procedure

Performance Measure	Downtime Situation	Simulation (min)	Time Buffer Analysis Procedure (min)	% Difference
Product 1 average flow time to the constraint resource	High	1740	2172	+ 24.83%
	Medium	916	867	- 5.35%
	Low	582	461	- 20.79%
Product 2 average flow time to the constraint resource	High	1750	2193	+ 25.31%
	Medium	941	888	- 5.63%
	Low	566	482	- 14.84%
Std Dev of product 1 average flow time to the constraint resource	High	843	1150	+ 36.42%
	Medium	681	310	- 54.48%
	Low	124	90	- 27.42%
Std Dev of product 2 average flow time to the constraint resource	High	875	1150	+ 31.43%
	Medium	623	310	- 50.24%
	Low	123	90	- 26.83%
% of late orders at the constraint resource	High	1.68%		
	Medium	8.50%		
	Low	10.00%		
Average lateness of late orders at the constraint resource	High	1110		
	Medium	711		
	Low	274		

For the high downtime situation the percentage late orders completed at the constraint resource is 1.68%, with an average lateness of 1110 minutes (which is about 64% of the average flow time to the constraint). This good performance was obtained because the time buffer estimation procedure overestimated the average flow time and standard deviation of the flow time to the constraint. Some

Stellenbosch University http://scholar.sun.ac.za

orders were however still late, which shows that even conservative levels of protection cannot fully guard against high variability. The percentage late orders and average lateness for the medium downtime situation are 8.5% and 711 minutes (± 78% of the average flow time to the constraint), whereas for the low downtime situation the values are 10% and 274 minutes (± 48% of the average flow time to the constraint) respectively. The percentage late orders are more than the desired service level chosen for the z-value when the time buffer lengths were calculated. This can be ascribed to the fact that the average flow time and standard deviation of the time buffer procedure's flow time calculations were less accurate. The relatively low percentage late orders and average lateness however indicate that the calculated time buffers provides quite good throughput protection at the constraint, even though the average flow time and standard deviation calculations are not always that accurate. Inaccuracies in the average flow time and flow time standard deviation estimations seem to be offset by the safety level (z-value) incorporated in the final time buffer lengths. For adequate protection the maximum z value should therefore be used. The simulation results also illustrate that it is not only sufficient to have accurate average flow time estimations, but that it is also important to accurately estimate the variance in the flow times. This is demonstrated by the medium downtime situation that had quite accurate average flow time estimations, but less accurate standard deviation estimations. This resulted in a higher percentage of late orders at the constraint.

3.4 Conclusions on the Time Buffer Estimation Procedure

This chapter presented an open queuing network analysis model developed to estimate the size of the time buffers in Theory of Constraints controlled production systems. Although the model could also be applied to other shop types such as job shops, it was only evaluated on a flow shop. The results from the simulation study show that such an analytical model is sufficiently accurate to be used as a quick estimator of the needed time buffer sizes. This was demonstrated over different levels of variability ranging from high to low. Although the average flow time and flow time standard deviation estimations of the queuing analysis procedure are not always that accurate in more realistic network models, the relatively low percentage late orders and average lateness measures from the simulation results indicate that the final calculated time buffer sizes are still useful estimators of the amount of protection necessary. The time buffer lengths can always later be fine tuned with simulation experiments or through the application of buffer management techniques after the line is in operation. It should also be remembered that no buffer management techniques were implemented in the simulation. In practice these techniques should greatly improve the delivery performance.

The time buffer estimation procedure presented in this study can therefore be quite useful in the design phase of production flow lines in order to compare different line configurations with respect to the amount of protection in terms of time (which results in a higher or lower WIP according to Little's formula) required by the specific line design. It can also be used to estimate the size of the time buffers to be used when putting the line into operation. It is however still essential that buffer management techniques be applied while operating the line to monitor the time buffers and fine-tune their lengths. The time buffer estimation procedure presented in this study however enables a manager to build the initial time buffer on a more scientific estimate than the approaches currently described in the Theory of Constraints literature.

Stellenbosch University http://scholar.sun.ac.za

Chapter 4

Investigating Protective Capacity Levels In Flow Production Systems

4.1 Introduction

From the literature review in Chapter 2 it could be seen that adding protective capacity to a line can significantly improve the production rate, flow time and bottleneck shiftiness. While all these studies proved the beneficial effect of protective capacity in improving the flow time, production rate and bottleneck shiftiness in both flow lines and job shop systems, several questions regarding the design of protective capacity still remain unanswered. One of the important questions remaining is what minimum levels of protective capacity are necessary to ensure a stable primary capacity constraint location. Another question is what effect the amount of protective capacity at the secondary constraint (i.e. the second most utilised resource) has on line performance, as well as the effect of the location of the secondary constraint relative to the primary constraint resource.

This chapter describes an investigative study using simulation to answer some of these questions regarding protective capacity design. Simulation models of both a flow shop and an assembly line were developed in order to represent different types of production lines. A flow shop refers to a discrete production system where workstations are arranged along the line of flow of the parts (product layout), and all jobs follow mostly the same processing sequence (Chase & Aquilano 1992, p.455). An assembly line is a special case of flow shop where resources are dedicated to continual production of a narrow product line, and operations entail the combination of several components to form another component, or a product. In this research the term flow shop therefore refers to a production system with a product layout and non-assembly operations. This research further only considers discrete manufacturing systems with asynchronous part transfer (parts therefore move independently of other parts).

All research studies up to date investigating protective capacity have focused solely on flow shops where the number of stations are fixed and there are infinite buffer capacity between stations. Increasing protective capacity therefore simply means increasing the processing capacity at stations without increasing the number of stations. It was therefore decided to also investigate assembly line configurations where there is a total work content that needs to be distributed among variable workstations. The number of workstations therefore is not fixed, but vary according to the specific distribution of the work content and the protective capacity level.

In the next section the specific research questions posed are presented. Section 4.3 describes the simulation experiment conducted on the flow shop, and section 4.4 the simulation experiment for the assembly line.

4.2 Research Questions and Hypotheses

The main performance measures for discrete flow production systems investigated in this study are the total output for the line, the mean flow time for jobs through the line, and the percentage of time the primary CCR is the only constraint or bottleneck in the line (the exclusive bottleneck probability). The specific research questions posed with respect to the design of protective capacity in flow production systems are the following:

1) Does the location of the secondary CCR relative to the primary CCR affect the average flow time for jobs through a flow shop; the total output for a flow shop; and the exclusive bottleneck probability of the primary CCR in a flow shop?
2) What is the effect of the amount of protective capacity at the second most heavily utilised resource (the secondary constraint resource or secondary CCR) on the average flow time for jobs through a flow shop; on the total output for a flow shop; and on the exclusive bottleneck probability of the primary constraint resource or CCR (the resource with the lowest throughput rate) in a flow shop? Are there any interactions between the amount of protective capacity at the secondary CCR and the location of the secondary CCR relative to the primary CCR, the amount of protective capacity at the other non-constraint stations, and the amount of variability in the flow shop characterised by machine breakdowns?
3) Using an equal distribution of protective capacity at all non-constraint resources, what is the effect of the amount of protective capacity at the non-constraint stations on the average flow time for jobs through a flow shop; on the total output for a flow shop; and on the exclusive bottleneck probability of the primary CCR in a flow shop? What minimum levels of average protective capacity are needed to ensure a long-term stable primary CCR that does not move or shift too often between stations in a flow shop?
4) In an assembly line using an equal distribution of protective capacity at all non-constraint resources, what is the effect of the amount of protective capacity at the non-constraint stations on the mean flow time for jobs through the line?

5) In an assembly line, does the same relationship between the amount of protective capacity at the non-constraint stations and the mean flow time for jobs through the line hold for different line lengths and different levels of processing time variation at stations?
6) In an assembly line, does the same relationship between the amount of protective capacity at non-constraint stations and the mean flow time hold for lines with limited buffer capacity and lines with unlimited buffer capacity between stations?
7) In an assembly line, what levels of protective capacity can provide significant improvements in flow time over a balanced capacity line without resulting in too many stations?

The first three research questions were investigated with a full factorial experiment performed on a simulation model of a 15 station flow shop. A 6x4x3x2 factorial ANOVA was performed with the amount of protective capacity at the secondary CCR (SecCCR PC), the amount of protective capacity at the non-constraint resources (NonCCR PC), the location of the secondary constraint relative to the primary constraint (SecCCRLoc), and the amount of downtime experienced by non-constraint resources (Downtime) as the independent variables, and the mean flow time, total output and primary CCR bottleneck probability as the dependent variables.

Research questions four to seven were investigated with simulation models of different configurations of an assembly type of flow line. Single factor ANOVAS were performed using the amount of protective capacity at the non-constraint stations as the independent variable and the mean flow time as the dependent variable. Simulation models were developed and the results analysed for each of six different assembly line configurations: (1) a short line (with 6 initial stations) with infinite buffer capacity between stations and a low processing time variability; (2) a short line (with 6 initial stations) with infinite buffer capacity between stations and a high processing time variability; (3) a short line (with 6 initial stations) with a finite or limited buffer capacity of one between stations and a low processing time variability; (4) a long line (with 20 initial stations) with infinite buffer capacity between stations and a low processing time variability; (5) a long line (with 20 initial stations) with infinite buffer capacity between stations and a high processing time variability; and (6) a long line (with 20 initial stations) with a finite or limited buffer capacity of one between stations and a low processing time variability.

The observations for the flow shop experiment can be described by the following statistical model:

$$y_{ijklm} = \mu + \tau_i + \beta_j + \gamma_k + \omega_l + (\tau\beta)_{ij} + (\tau\gamma)_{ik} + (\tau\omega)_{il} + (\beta\gamma)_{jk} + (\beta\omega)_{jl} + (\gamma\omega)_{kl} + (\tau\beta\gamma)_{ijk} + (\tau\beta\omega)_{ijl}$$

$$+ (\tau\gamma\omega)_{ikl} + (\beta\gamma\omega)_{jkl} + (\tau\beta\gamma\omega)_{ijkl} + \varepsilon_{ijklm} \qquad \begin{cases} i = 1,2,\ldots,6 \\ j = 1,2,\ldots,4 \\ k = 1,2,3 \\ l = 1,2 \\ m = 1,2,\ldots,150 \end{cases} \qquad (4.1)$$

where μ is the overall mean effect, τ_i is the effect of the ith level of the SecCCR PC factor, β_j is the effect of the jth level of the SecCCRLoc factor, γ_k is the effect of the kth level of the NonCCR PC factor, ω_l is the effect of the lth level of the Downtime factor, ε_{ijklm} is the random error , m is the number of observations per cell, and the other terms are all the two factor, three factor and four factor interactions. The errors are assumed to be normally and independently distributed random variables with mean zero and constant variance σ^2. The hypotheses for main effects and interactions to be tested are:

H_o: the factor or interaction has no effect

H_a: the factor or interaction does have an effect

The decision rule is:

if P-value ≥ 0.05, conclude H_o

if P-value < 0.05, conclude H_a

The observations for each of the assembly line experiments can be described by the following statistical model

$$y_{ij} = \mu + \tau_i + \varepsilon_{ij} \qquad \begin{cases} i = 1,2,\ldots,6 \\ j = 1,2,\ldots,n \end{cases} \qquad (4.2)$$

where μ is the overall mean effect, τ_i is the effect of the ith level of the NonCCR PC factor, ε_{ij} is the random error , and n is the number of observations under the ith treatment. The hypotheses for treatment effects to be tested are:

H_o: $\tau_1 = \tau_2 = \ldots = \tau_6 = 0$

H_a: $\tau_1 \neq 0$ for at least one i

The decision rule is:

if P-value ≥ 0.05, conclude H_o

if P-value < 0.05, conclude H_a

Section 4.3 discusses the flow shop experiment whereas section 4.4 discusses the experiment for the assembly line. For all the statistical tests the level of significance used is 0.05.

4.3 Flow Shop Experiment

4.3.1 Simulation model

The Arena 3.0 simulation software was used to develop a simulation model of a single product 15 station flow shop. Refer to Figure 4-1 for a flow diagram of the simulation model. An open system was modelled where jobs arrive according to a certain arrival distribution, and finished jobs leave the system. WIP was therefore not held constant in the line by routing finished units back to the start of the line. This was done because average flow time was one of the dependent variables studied, and should be allowed to vary by keeping non-constant WIP (because of the relationship between WIP and flow time according to Little's Law). In practice where jobs are released to a flow shop according to a schedule developed by a drum-buffer-rope scheduler, WIP is also not usually held constant but allowed to vary. Jobs are still released according to the throughput rate of the constraint, but WIP is allowed to vary based on the variability conditions on the shop floor. Preliminary experimental simulation runs suggested such a non-constant WIP flow shop to obtain better delivery performance than one in which the WIP is held constant. Unlimited queue sizes were assigned to each station. This prevented blocking between stations, allowing inventory to flow freely in the system. When a job completed processing at a station, it immediately proceeds to the following station. No transfer batches were therefore modelled. It was further assumed that transfer times between stations are negligible. All stations consisted of a single server. Station 8 was designated as the primary constraint or CCR (or the resource with the lowest possible throughput rate). The primary constraint was therefore located at the centre of the line. This position was chosen to enable different locations of the secondary constraint relative to the primary constraint to be evaluated both before and after the primary constraint. The position of this primary constraint resource was kept constant during the simulation.

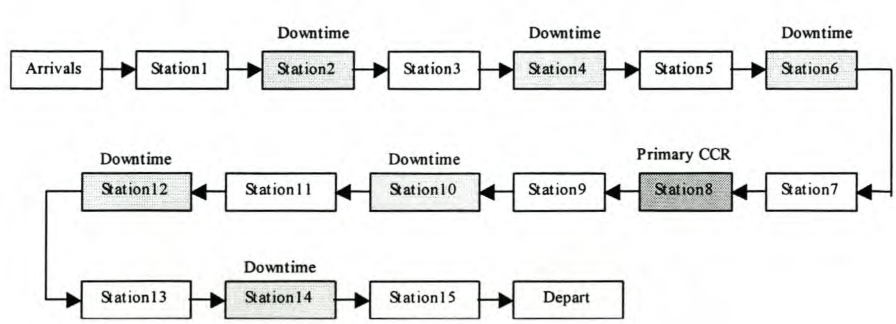

Figure 4-1. Simulation model of 15 station flow shop

Processing times for all workstations were assumed to be lognormal. It has been shown in other studies (Atwater and Chakravorty 1994; Muralidhar et al 1992) that the lognormal distribution is a good representation of real-world processing times. It was decided to use a mean processing time of 4 minutes at the primary constraint station. This was arbitrarily chosen as a processing time that is neither too fast nor too slow. Furthermore all processing times were assigned a coefficient of variation of 0.4. This represents a low processing time variability. According to Hopp and Spearman (1996) variability caused by resource breakdowns or resource unavailability is more damaging to line performance than normal processing time variability. As pointed out by Dudley (1963) processing times tend to follow a lognormal distribution with relatively low variance.

Resource failures were modelled only at stations 2, 4, 6, 10, 12, and 14. The primary constraint station (station 8) therefore experiences no failures. Resource failures are time-based and when a failure occurs, the resource immediately stops processing the job. When the resource is back in operation again, it finishes the job it was busy with when the failure occurred.

An arrival rate was chosen to ensure a 95% utilisation at the primary CCR. This utilisation level for the primary CCR was chosen based on a study by Atwater and Chakravorty (2000) who suggested that constraint utilisation levels of around 94% appear to be best for DBR systems. Because of variability, too high utilisation of the constraint endangers line performance such as flow time and due date delivery performance. A certain amount of protective capacity is therefore also needed at the primary constraint. Based on this 95 % utilisation measure and the service time at station 8 (the primary CCR), the arrival rate was calculated as 0.2375 jobs per minute. Arrivals were modelled with the lognormal distribution using a coefficient of variation of 0.4. This represented a low arrival variability.

4.3.2 Model time span

The flow shop was modelled as a terminating system. An empty shop is used as the starting condition, and the simulation is terminated after a run length of 48000 minutes, which relates to 100 production days if a production day consists of 8 production hours. Data is therefore collected by using independent replications of this terminating system. The 48000 minute simulation period was arbitrarily chosen to ensure a long enough time span for resource breakdowns to occur and to process a sufficient number of jobs.

4.3.3 The independent variables

Four independent variables were used in this study: (1) the amount of protective capacity at the secondary constraint resource or CCR, (2) the location of the secondary constraint relative to the primary constraint, (3) the amount of protective capacity at the non-constraint resources, and (4) the amount of downtime experienced by non-constraint resources.

4.3.3.1 The amount of protective capacity at the secondary CCR (SecCCR PC)

The secondary CCR is defined as the resource with the second lowest output rate as calculated by the following formula: *Output rate in jobs per hour = 60 / Mean service time in minutes.*

Protective capacity is defined as the complement of the percentage difference in the output rate of the primary constraint relative to the secondary constraint or the non-constraint. For example, if the processing time at the primary constraint is 4 minutes and at the non-constraint 3 minutes, their respective output rates are 15 jobs per hour and 20 jobs per hour. The amount of protective capacity at the non-constraint is then calculated as (1- 15/20) = 25%.

The amount of protective capacity was modelled at six levels: 2%, 5.25%, 10.5%, 21%, 31.5%, and 52.5%. These levels were chosen to present a wide enough range of protective capacity ranging from low to high. More protective capacity levels were chosen to fall in the low to medium range, since preliminary experiments suggested that the largest relative improvements in line performance occur within the low to medium protective capacity range.

4.3.3.2 The location of the secondary CCR relative to the primary CCR (SecCCRLoc)

The location of the secondary CCR relative to the primary CCR was varied at four different locations: at workstation 7 which is before but near the primary CCR (BN), at workstation 3 which is before but far from the primary CCR (BF), at workstation 9 which is after but near the primary CCR (AN), and at workstation 13 which is after but far from the primary CCR (AF). These locations were chosen to examine whether it is more desirable to have a secondary constraint located before or after the primary constraint, and how the distance between the secondary and primary CCRs affect line performance. The locations were also chosen to ensure that at all locations the secondary CCR is affected by upstream station failures.

4.3.3.3 The amount of protective capacity at the non-constraint resources (NonCCR PC)

This variable was varied at three different levels: 10.5%, 21%, and 52.5%. These levels were arbitrarily chosen to reflect low to high protective capacity levels. Protective capacity is again defined and calculated as in section 4.3.3.1 (the protective capacity at the secondary CCR). For the non-constraint resources, protective capacity was equally assigned to all stations (a level arrangement) that are neither a primary nor secondary CCR. For example, if the protective capacity level is 10.5%, then all non-constraint resources have a protective capacity of 10.5%.

4.3.3.4 The amount of downtime experienced by non-constraint resources (Downtime)

This variable was chosen to examine how protective capacity requirements are influenced by different levels of variability in the line. Two levels of downtime were used: 1.57% and 10.04%. These were selected arbitrarily to represent low and high downtime situations. Station downtimes were calculated by using the Mean-Time-Between-Failures (MTBF) and the Mean-Time-To-Repair (MTTR) measures:

Low level % downtime = MTTR/(MTTR + MTBF) = 120min/(120min + 7500min) *100 = 1.57%
High level % downtime = MTTR/(MTTR + MTBF) = 480min/(480min + 4300min) *100 = 10.04%

The MTTR and MTBF values were also arbitrarily chosen. Although frequent short breakdowns and long infrequent breakdowns can have the same downtime percentage, it is actually the infrequent long breakdowns that hold the greater danger of starving the constraint and therefore disrupting its utilisation. Care was therefore taken to ensure that when choosing the MTTR values these values were sufficiently long with respect to the average flow time to the specific buffer origin, and also that the higher downtime situations experience a longer MTTR than the lower downtime situations. The MTBF was modelled using the exponential distribution with mean 4300 minutes for the high downtime situation and 7500 minutes for the low downtime situation. It has been shown that actual equipment failures commonly follow the exponential distribution (Atwater & Chakravorty 1994). The MTTR has been modelled using the lognormal distribution with mean values 120 minutes for the low downtime situation and 480 minutes for the high downtime situation, and a coefficient of variation of 0.4 for both downtime situations. The lognormal distribution has been used since repairs can be viewed as a process (Atwater & Chakravorty 1994), and it has been shown in other studies that the lognormal distribution is a good representation of real-world processing times (Muralidhar et al 1992). The coefficient of variation is the same as the measure used for all processing times in this experiment and represents a low processing time variability.

4.3.4 The dependent variables

Three dependent variables or performance measures were studied: (1) the total output for the line, (2) the mean flow time for jobs through the line, and (3) the exclusive bottleneck probability of the primary CCR.

4.3.4.1 The total output for the line

Total output is measured as the total number of jobs completed by the line within the simulation time period. This measure is an important indication of the throughput capability of the line.

4.3.4.2 The mean flow time for jobs through the line

Mean flow time is measured as the average time between job arrivals at the first station and job departures at the last station. The mean flow time indicates the ability of the line to process orders quickly and compete against time-based dimensions. This can lead to increases in future throughput, as well as protecting current throughput. From Little's law it is also an indication of the amount of WIP in the line. The advantages of low WIP levels were discussed in chapter 1.2.

4.3.4.3 The exclusive bottleneck probability for the primary CCR

The exclusive bottleneck probability measures the probability for the primary CCR to be the only long run bottleneck. This in turn indicates the inclination for the primary CCR or bottleneck to move or shift in the line, and therefore also the difficulty in managing the line. Since production scheduling and planning, as well as strategic planning within the Theory of Constraints philosophy focus on constraints and constraint management, it is important that the location of the constraint remain relatively stable over the long run. Shifting bottlenecks create control problems for shop floor personnel, since reactive measures such as expediting, extra labour and overtime to ameliorate the bottleneck requires management attention and cause disruption at other workstations (Lawrence and Buss 1994). Using the bottleneck probabilities for stations, Lawrence and Buss (1994) developed a bottleneck shiftiness measure that indicates the propensity for bottlenecks to shift between stations (refer to section 2.2.4). In a simulation experiment they demonstrated that all workstations (regardless of their utilisation rates) are the bottleneck some of the time, even when there is clearly a dominant long-run bottleneck (i.e. has the highest utilisation rate).

Lawrence and Buss defined the bottleneck probability for each station as the long-run proportion of time a given workstation has more jobs in queue than any other. In their simulation experiment they therefore measured the proportion of time each station has the largest queue, which they used to calculate the probability of each station being the long run bottleneck. The problem with their definition of the bottleneck probability however is that they do not take into account the processing rate of the station. A given amount of WIP lying in front of a slow station takes much longer to process than the same amount of WIP in front of a fast station, and therefore the slower station should be viewed as the constraint or bottleneck. With Lawrence and Buss's definition both stations would be recorded as being the bottleneck.

In this study the exclusive bottleneck probability for each station is defined as the long-run proportion of time a given station has a longer station flow time than any other station, where station flow time is the time to move jobs through the station. In the simulation model the exclusive bottleneck probability is calculated by counting the frequency each station has the longest station flow time. Each five minutes of simulation time the number of jobs in queue in front of each station is checked. The flow time for each station is then calculated by multiplying the number of jobs in queue at a station with the average service time per job for the station. The station with the longest flow time is then recorded as the bottleneck. At the end of the simulation run the exclusive bottleneck probability for each station is then determined by calculating the proportion of time it was the exclusive bottleneck. An ANOVA is however only performed on the bottleneck probability for the primary CCR.

Preliminary simulation experiments conducted showed that compared to this study's definition of the bottleneck probability, Lawrence and Buss's definition of the bottleneck probability significantly underestimates the bottleneck probability for the long-run bottleneck station with the highest utilisation. A summary of the all the experimental factors and levels of the factors used, as well as other research parameters is provided in Table 4-1.

Table 4-1. Experimental factors and parameters

Parameter	Values
Amount of protective capacity at the secondary CCR	2%, 5.25%, 10.5%, 21%, 31.5%, and 52.5%.
Location of the secondary CCR relative to the primary CCR	Before-Near, Before-Far, After-Near, and After-Far relative to primary CCR
Amount of protective capacity at the non-CCRs	10.5%, 21%, and 52.5%
Amount of downtime experienced by non-CCRs	1.57% and 10.04%
Processing time CV	0.4
Processing time distribution	Lognormal
CCR utilisation	95%
Line Length	15 Stations

4.3.5 Preliminary statistical analysis

The preliminary statistical analysis involved determining the warm up period and the number of replications needed to draw valid statistical conclusions from the model. Since the model starts out in an empty and idle state, the length of the warm up period had to be determined in order to eliminate the effects of the transient phase when gathering the statistics for analysis. The truncated-replication approach suggested by Kelton et al (1998, pp.219-224) for terminating simulations was used. To determine the length of the warm up period, a single run of 10 replications was made using the following levels of the different factors: 2% protective capacity at the secondary CCR, before-near (BN) secondary CCR location, 10.5% protective capacity at the non-CCR stations, and 10.04% downtime. This represented a worst-case scenario. The different performance measure values were saved to data files using the Statistics module of the Arena software package. Arena's Output Analyser was then used to draw plots of the different performance measure values recorded against time. The values for the 10 replications were superimposed on one graph for each performance measure. A visual inspection of the different plots showed the values to stabilise and therefore the transient phase to end after more or less 6000 minutes. The length of the warm up period was therefore chosen as 6000 minutes. Simulations were therefore run for 6000 minutes after which statistics were cleared and data was collected over the next 42000 minutes.

The next step was to determine the actual number of replications needed to achieve the necessary confidence intervals for the performance measures under study. Ten initial replications were run using the specified warm up period of 6000 minutes, and the number of replications needed was computed using the confidence interval half-width formula from Kelton et al (1998, p.185):

$$n^* = n \times \left(\frac{h}{h^*} \right)^2$$

Where

n^* = The required number of replications.

n = The number of initial replications (10 in this case).

h = Calculated confidence interval half width obtained from the initial run.

h^* = Desired confidence interval half width based on the variable that was studied and the accuracy wanted for that specific variable.

This formula is applied for each of the performance measures under study. The final number of replications to perform is then determined as the largest required number of replications from the set of output measures. Using a 95% confidence level, it was determined that to obtain a 5 % precision level in the performance measures 150 replications had to be performed. For the final simulation runs the actual precision level achieved was better than 5% (in the order of 2%).

4.3.6 Simulation results and analysis

The output from the simulation runs were analysed utilising ANOVA. A 6x4x3x2 ANOVA model was constructed for each of the three dependent variables or performance measures individually, using the amount of protective capacity at the secondary CCR, the amount of protective capacity at the non-constraint resources, the location of the secondary constraint relative to the primary constraint, and the amount of downtime experienced by non-constraint resources as the independent variables. The simulation results for each of the 144 cases by the 4 factors and 3 dependent variables are listed in Appendix F. Each combination of factors was replicated 150 times, which resulted in a total of 150x6x4x3x2 = 21600 simulation runs. All main effects, two-way interactions, three-way interactions, and four-way interactions were computed. Graphical analysis and Tukey's HSD comparison tests were used to further analyse interactions between independent variables. All statistical analyses are based on a 95% confidence level. The large number of replications per combination of factors (150) resulted in a very large number of degrees of freedom for the error term in the ANOVA experiment. This also had the effect that the power for the experiment was very high, which results in small factor effects to be statistically significant. The number of replications in the simulation experiment was determined to obtain a 5% precision level in the confidence interval half widths for the means at a worst-case scenario (therefore at a factor combination resulting in the largest variability). Because the number of replications calculation was based on a worst-case scenario, in the simulation experiment the confidence interval half-width precision level achieved for the means at the other factor levels with less variation was much higher (between 1% and 2%). Because of the high precision level, the ANOVA analyses showed small factor effects to be statistically significant. All significant factor effects from the ANOVA results were however further analysed with graphical analysis and comparison of means. This was used to distinguish between results that are statistically significant and practically useful or meaningful. The Tukey comparisons of means showed differences between means of more or less 2% and lower to be statistically non-significant, and differences of more than around 2% as significant. In the discussions of the experimental results it was however indicated whether such small effects are significant or meaningful from a practical viewpoint, and this was used to make the final decision about the significance of factor effects in the experiments.

In order for the ANOVA procedure to be valid, certain assumptions need to be satisfied. Two of these assumptions are that the errors in the model are normally and independently distributed with mean zero and constant but unknown variance (Montgomery 1997, p.79). In order to test for the normality assumption a normal probability plot of the residuals was constructed for each of the ANOVA models. If the underlying error distribution is normal, this plot will resemble a straight line. The constant variance assumption was checked by plotting the residuals versus the predicted values for the dependent variable in each of the three ANOVA models. This plot should not reveal any obvious pattern. In this experiment the constant variance assumption was intentionally violated, because of the design of the experiment where different resource downtime levels were investigated. This represented different variability levels. The normal probability plots and the scatter plots for the residuals confirmed the violation of both the normality and constant variance assumptions. The normal probability plots did not resemble a straight line, and the scatter plot for the residuals showed an obvious pattern of increasing variances (refer to Appendix D). The non-constant variances for especially the flow time measure can also be explained by the fact that lower protective capacity levels cause more variability. This was confirmed by calculating the standard deviations for the flow times at the different protective capacity levels. Montgomery (1997, pp.143-146) as well as Conover (1999, pp. 419-420) suggest that in situations where the assumptions for an ANOVA are violated, a rank transformation be performed by replacing the observations by their ranks. The observations are first ranked in ascending order and each observation is then replaced by its rank, with the smallest observation having rank 1. In the case of ties, the average rank is assigned to each of the tied observations. An analysis of variance is then performed on both the original data and the ranks. When both procedures give nearly identical results, then the assumptions underlying the usual analysis of variance are reasonably met and the parametric analysis is valid. In this study ANOVAs based on the ranks for all three dependent measures were consistent with that performed on the raw data, therefore the ANOVAs based on the original performance values were used to analyse the data.

In the following sub-sections the analysis for each of the three dependent variables are discussed separately. When higher order interactions are present, analysis of results are focused on the highest order factor interactions that are significant, since lower order interactions can give misleading results in the presence of higher order interactions. For example if in an experiment all main effects, two-way, three-way and four-way interactions are significant, analysis is focused on the four-way interactions. The main effects for these cases are however also discussed, since they confirm certain observations made for the higher order interactions.

4.3.6.1 The mean flow time

The output for the ANOVA performed on the mean flow time values obtained from the simulation runs are displayed in Table 4-2. The ANOVA performed on the corresponding ranks are displayed in Table 4-3. The rows for non-significant effects are shaded. From the two tables it can be seen that the two ANOVA's give basically similar results. The only difference is for the SecCCRPC*NonCCRPC*Downtime interaction and the SecCCRLoc*SecCCRPC*NonCCRPC interaction. The ANOVA based on the ranks gives a non-significant three-way interaction for the SecCCRPC*NonCCRPC*Downtime interaction, whereas the ANOVA based on the original values gives a significant interaction (although very small). For the SecCCRLoc*SecCCRPC*NonCCRPC interaction the ANOVA based on the original values gives a non-significant three-way interaction, whereas the ANOVA based on the ranks gives a significant interaction (although very small). Because of the small differences between the two ANOVAs it is therefore safe to continue the analysis based on the ANOVA for the original non-ranked observations.

From Table 4-2 it can be seen that all main effects and two-way interactions are significant (p-value < 0.05). Furthermore all three-way interactions except the SecCCRLoc*SecCCRPC*NonCCRPC interaction were found significant, whereas the four-way interaction was not significant. Graphs for all the main effects are displayed in Figure 4-2, and for the significant three-way interactions in Figure 4-3 to Figure 4-5. Although in the ANOVA based on the original non-ranked data no significant interaction was found between the location of the secondary CCR, the amount of protective capacity at the secondary CCR, and the amount of protective capacity at the non-CCR resources, it is included because the ranked ANOVA suggested a significant interaction. The Tukey HSD comparisons of the means for the main effects (based on a 95% confidence level) are displayed in Table 4-4 and Table 4-5. The numbered columns indicate groups of means that do not differ significantly.

Figure 4-2a shows the main effect of the amount of protective capacity at the secondary CCR on the mean flow time. This graph suggests that a slight improvement in the mean flow time can be obtained by increasing the amount of protective capacity at the secondary CCR. The Tukey HSD comparisons of the means in Table 4-5 however shows that only the mean flow times for the 2% and 5.25% secondary CCR protective capacity levels differ significantly from the means at the other protective capacity levels. No statistically significant difference exist between the mean flow times of the 10.5%, 21%, 31.5% and 52.5% protective capacity levels. The difference between the best and worst flow times is however relatively small [100*(1146-1073)/1146 = 6.4%].

The main effect for the secondary CCR location in Figure 4-2b suggests that the worse flow time is obtained at the After-Far secondary CCR location. From this figure as well as the Tukey HSD comparisons of the means in Table 4-4 it can be seen that the After-Far secondary CCR location has the longest flow time, and there is a statistically significant difference between this mean flow time and the mean flow times for the other locations. The difference between the best and worst mean flow time performances obtained by the secondary CCR location is however relatively small. Only a 4.7% [100*(1130 -1077)/1130] improvement in mean flow time is obtained by moving the secondary CCR from the After-Far location to the After-Near location. The shortest flow time is obtained for the Before-Near and After-Near secondary CCR locations, and there is not a statistically significant difference between these two means.

Table 4-2. ANOVA table: original non-ranked flow time output measures

	SS	Degr. Of Freedom	MS	F	p
Intercept	2.598258E+10	1	2.598258E+10	182193.4	0.000000
{1}SecCCRLoc	9.808688E+06	3	3.269563E+06	22.9	0.000000
{2}SecCCR PC	1.457163E+07	5	2.914326E+06	20.4	0.000000
{3}NonCCR PC	4.491823E+09	2	2.245912E+09	15748.6	0.000000
{4}Downtime	1.875535E+10	1	1.875535E+10	131515.0	0.000000
SecCCRLoc*SecCCR PC	1.040076E+07	15	6.933841E+05	4.9	0.000000
SecCCRLoc*NonCCR PC	4.421149E+06	6	7.368582E+05	5.2	0.000026
SecCCR PC*NonCCR PC	5.563910E+06	10	5.563910E+05	3.9	0.000026
SecCCRLoc*Downtime	7.779117E+06	3	2.593039E+06	18.2	0.000000
SecCCR PC*Downtime	7.437678E+06	5	1.487536E+06	10.4	0.000000
NonCCR PC*Downtime	3.642610E+09	2	1.821305E+09	12771.2	0.000000
SecCCRLoc*SecCCR PC*NonCCR PC	6.237718E+06	30	2.079239E+05	1.5	0.050513
SecCCRLoc*SecCCR PC*Downtime	8.009696E+06	15	5.339797E+05	3.7	0.000001
SecCCRLoc*NonCCR PC*Downtime	3.598282E+06	6	5.997137E+05	4.2	0.000311
SecCCR PC*NonCCR PC*Downtime	5.234979E+06	10	5.234979E+05	3.7	0.000064
1*2*3*4	5.950119E+06	30	1.983373E+05	1.4	0.075769
Error	3.059838E+09	21456	1.426099E+05		

Table 4-3. ANOVA table: ranked flow time output measures

	SS	Degr. Of Freedom	MS	F	p
Intercept	2.519657E+12	1	2.519657E+12	1409164	0.000000
{1}SecCCRLoc	5.880074E+08	3	1.960025E+08	110	0.000000
{2}SecCCR PC	2.020392E+09	5	4.040783E+08	226	0.000000
{3}NonCCR PC	1.667635E+11	2	8.338173E+10	46633	0.000000
{4}Downtime	6.298560E+11	1	6.298560E+11	352258	0.000000
SecCCRLoc*SecCCR PC	6.402414E+08	15	4.268276E+07	24	0.000000
SecCCRLoc*NonCCR PC	4.174260E+08	6	6.957101E+07	39	0.000000
SecCCR PC*NonCCR PC	2.710850E+08	10	2.710850E+07	15	0.000000
SecCCRLoc*Downtime	3.585886E+07	3	1.195295E+07	7	0.000166
SecCCR PC*Downtime	3.680247E+08	5	7.360495E+07	41	0.000000
NonCCR PC*Downtime	5.028097E+07	2	2.514049E+07	14	0.000001
SecCCRLoc*SecCCR PC*NonCCR PC	2.512937E+08	30	8.376456E+06	5	0.000000
SecCCRLoc*SecCCR PC*Downtime	5.470053E+07	15	3.646702E+06	2	0.009994
SecCCRLoc*NonCCR PC*Downtime	4.105315E+07	6	6.842192E+06	4	0.000814
SecCCR PC*NonCCR PC*Downtime	1.756800E+07	10	1.756800E+06	1	0.455996
1*2*3*4	6.817022E+07	30	2.272341E+06	1	0.146783
Error	3.836444E+10	21456	1.788052E+06		

Stellenbosch University http://scholar.sun.ac.za

a. Effect of protective capacity at SecCCR on flow time

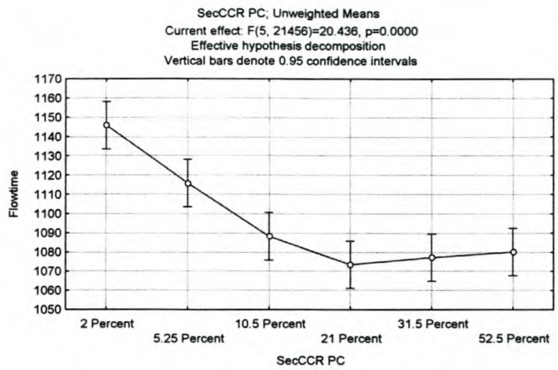

b. Effect of location of SecCCR on flow time

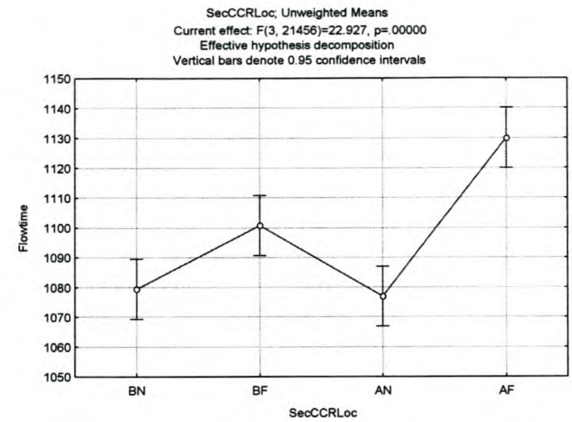

c. Effect of protective capacity at NonCCRs on flow time

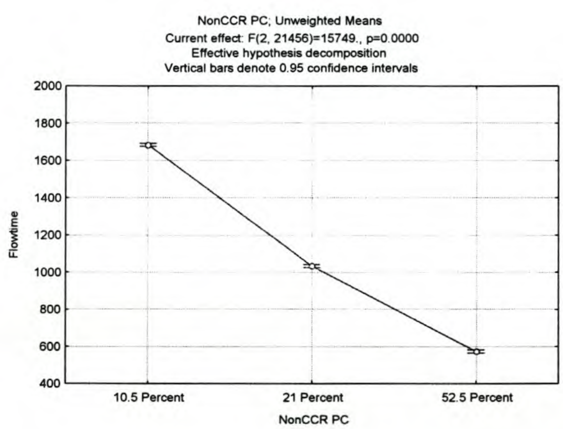

d. Effect of downtime on flow time

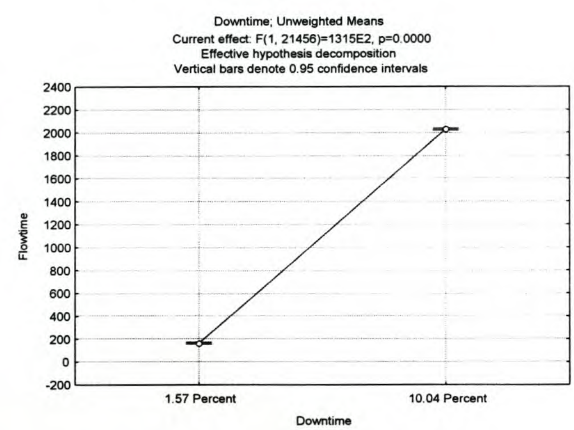

Figure 4-2. Main effects of experimental factors on flow time

Table 4-4. Tukey HSD comparisons of mean flow times for SecCCRLoc main effect

SecCCRLoc	Flowtime	1	2	3
AN	1076.941	**		
BN	1079.375	**		
BF	1100.741		**	
AF	1130.010			**

Table 4-5. Tukey HSD comparisons of mean flow times for SecCCR PC main effect

SecCCR PC	Flowtime	1	2	3
21 Percent	1073.459	**		
31.5 Percent	1077.183	**		
52.5 Percent	1080.081	**		
10.5 Percent	1088.221	**		
5.25 Percent	1115.844		**	
2 Percent	1145.813			**

(means within a group that do not significantly differ are indicated with **)

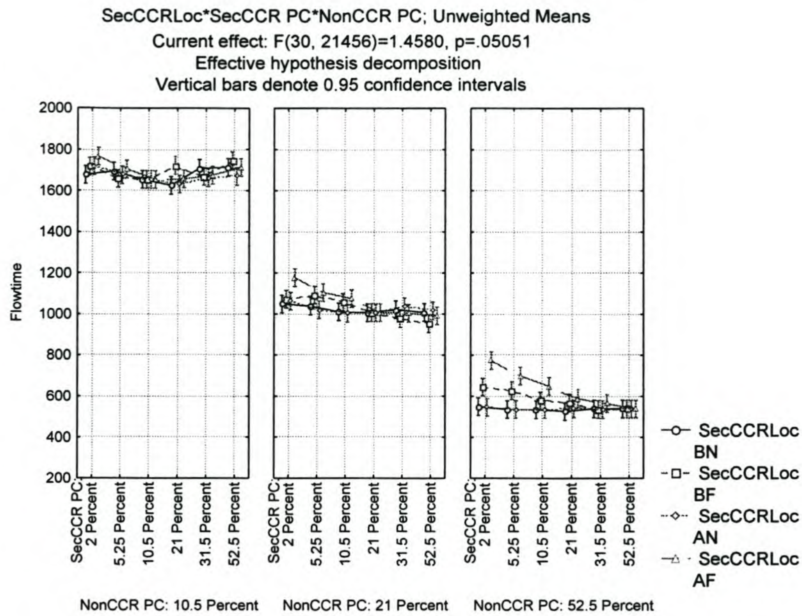

Figure 4-3. Three-way interaction of SecCCR Loc, SecCCR PC, and NonCCR PC on flow time

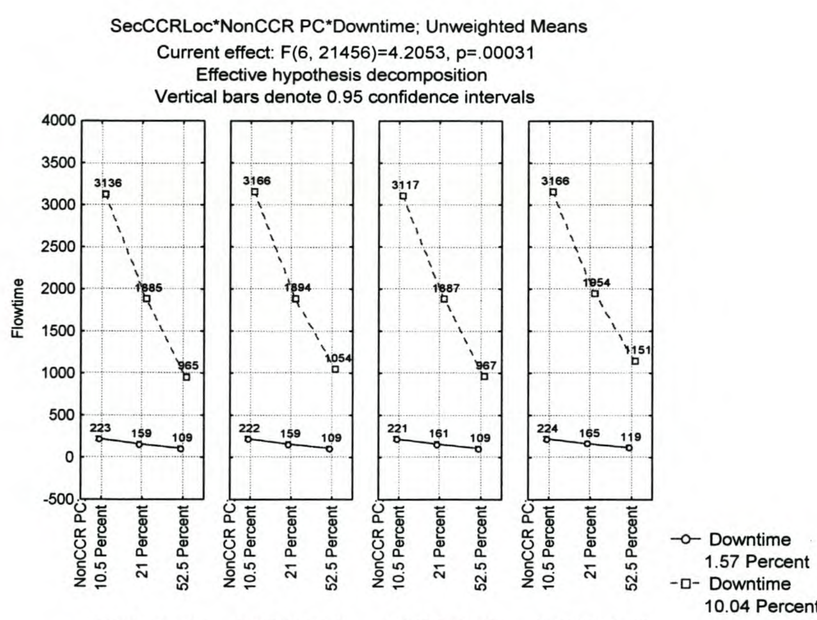

Figure 4-4. Three-way interaction of SecCCR Loc, NonCCR PC, and Downtime on flow time

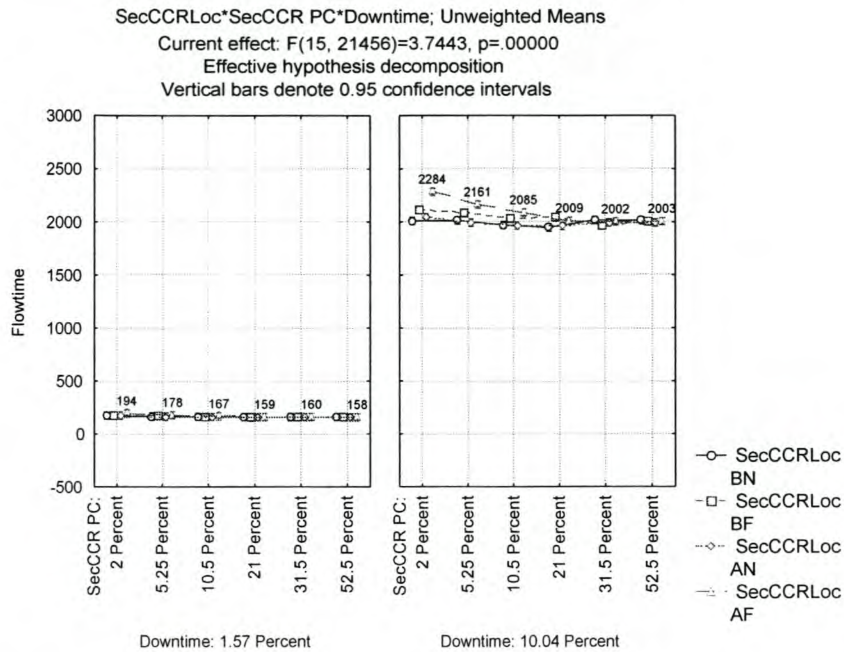

Figure 4-5. Three-way interaction of SecCCR Loc, SecCCR PC, and Downtime on flow time

The three-way interaction between the secondary CCR location, the amount of protective capacity at the secondary CCR, and the amount of protective capacity at the non-CCRs in Figure 4-3 however indicates that the worse flow time performance for the After-Far secondary CCR location and low secondary CCR protective capacity levels only apply when there is high levels of protective capacity at the non-CCRs, since no statistically significant differences exist between the mean flow times for the different secondary CCR locations at the 10.5% and 21% non-CCR protective capacity levels. From Figure 4-3 it can be seen that when the non-CCRs have 52.5% protective capacity and the secondary CCR only 2% protective capacity, an After-Far secondary CCR location results in the longest flow time. This is however only true for the 52.5% non-CCR protective capacity level. Tukey HSD comparisons of the means indicate that at the other non-CCR protective capacity levels the mean flow time for an After-Far location with 2% protective capacity at the secondary CCR is not significantly different from the means of the other secondary CCR locations. Furthermore the three-way interaction in Figure 4-5 suggests that this worse flow time performance for the After-Far secondary CCR location at low secondary CCR protective capacity levels also only applies for the high downtime situation. No statistically significant differences between the mean flow times exist for the 1.57% downtime level, whereas for the 10.04% downtime level the After-Far secondary CCR location has the highest flow time at the 2% secondary CCR protective capacity level.

From the ANOVA results in Table 4-2 it can be seen that the amount of protective capacity at the non-CCRs, as well as the amount of downtime experienced by resources, have the biggest main effects on the mean flow time performance measure (they have the largest F-values of the main effects). Figure 4-2c indicates that increasing the amount of protective capacity at the non-CCRs significantly reduces the mean flow time. The three-way interaction between the secondary CCR location, the amount of protective capacity at the non-CCRs, and the downtime level as graphed in Figure 4-4 however shows that the biggest improvement in mean flow time by increasing the amount of protective capacity at the non-CCRs is obtained for the high downtime situation. For the high downtime situation the flow time improves on average from more or less 3150 minutes to about 1900 minutes when the protective capacity at the non-constraints is increased from 10.5% to 21% (which is a 40% improvement). When the protective capacity is at the 52.5% level, the flow time is reduced to more or less 1030 minutes (which is a 67% improvement over the 10.5% protective capacity level). For the low downtime situation, the flow time reduces from more or less 220 minutes to 160 minutes (a 27% improvement) when the protective capacity at the non-constraints is increased from 10.5% to 21%, and to 110 minutes when the protective capacity is further increased to 52.5% (a 50% improvement over the 10.5% protective capacity level). The graph showing the main effect of the downtime level on the mean flow time is displayed in Figure 4-2d. It can be seen that a significantly lower mean flow time is obtained for the low downtime situation.

In the preceding analysis of the secondary constraint it should be noted that at situations where the secondary CCR protective capacity equals or surpasses the protective capacity of the non-constraints, it actually no longer can be called a secondary CCR, since the other stations actually have more capacity than the so-called secondary CCR. To see if this would change much of the preceding analysis, graphs showing the interaction for all independent variables or factors together were constructed. Figure 4-6 presents the 3-way interaction between the amount of protective capacity at the secondary CCR, the amount of protective capacity at the non-CCRs, and the location of the secondary CCR when stations experience low downtimes, whereas Figure 4-7 displays this interaction for the high downtime situation. The vertical dotted lines in these graphs indicate for each non-CCR protective capacity level the secondary CCR protective capacity level where the identified secondary constraint's protective capacity equals the non-CCR protective capacity and therefore ceases to be the secondary constraint. For secondary CCR protective capacity levels above these identified levels the secondary constraint is by definition no longer the secondary constraint. In Figure 4-6, for the 10.5% non-CCR protective capacity this protective capacity level where the secondary CCR is by definition no longer the secondary constraint is determined as 10.5% - 1.57% (the downtime level) = 8.93%. For the 21% non-CCR protective capacity level this is determined as 21% - 1.57% = 19.43%, and for the 52.5% non-CCR protective capacity level it is 52.5% - 1.57% = 50.93%.

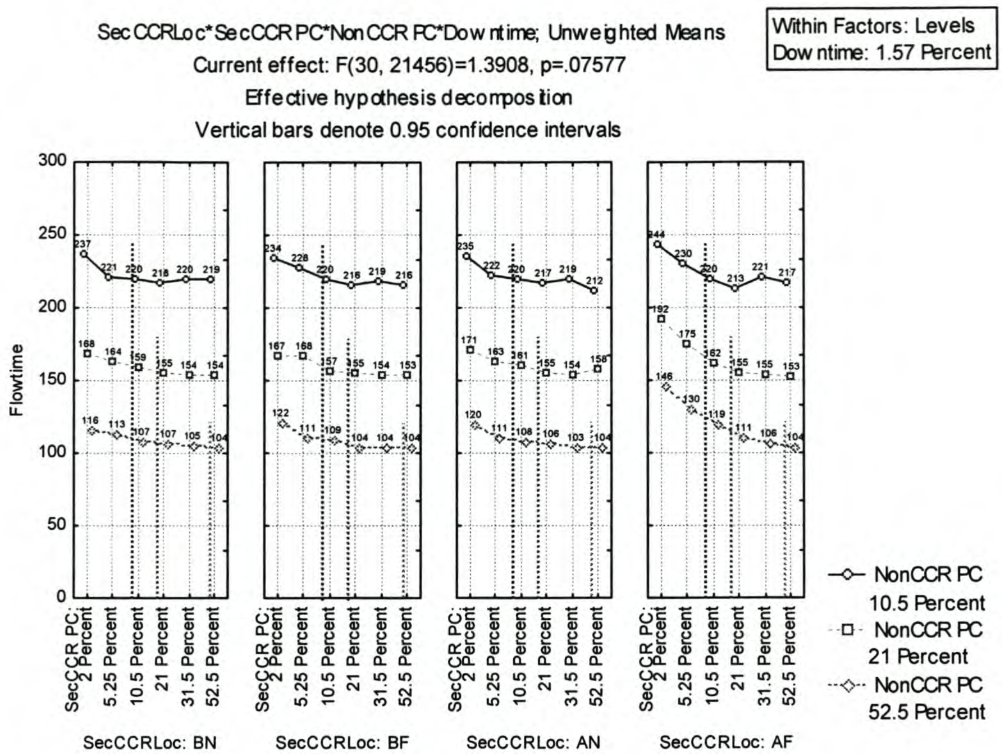

Figure 4-6. Three-way interaction of SecCCR Loc, SecCCR PC and NonCCR PC on the mean flow time: low downtime situation

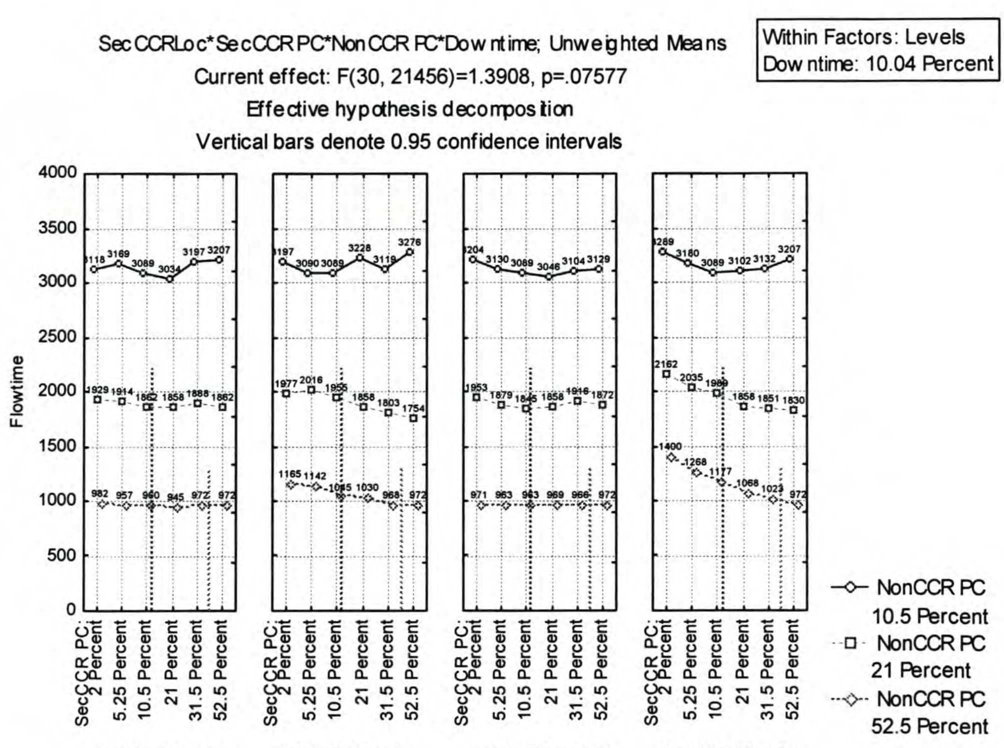

Figure 4-7. Three-way interaction of SecCCR Loc, SecCCR PC and NonCCR PC on the mean flow time: high downtime situation

For the high downtime situation in Figure 4-7 it is actually not necessary to analyse the 10.5% non-CCR protective capacity level, because at this level the so-called secondary CCR is actually never a secondary constraint. In this case those non-CCR stations that experience breakdowns are in fact the secondary constraints. For the 21% non-CCR protective capacity level the secondary CCR protective capacity level where it ceases to be the secondary constraint is determined as 21% - 10.04% = 10.96%, and for the 52.5% non-CCR protective capacity level it is 52.5%-10.04% = 42.46%. Both Figure 4-7 and Figure 4-6 indicate that the earlier analyses made regarding the protective capacity level at the secondary CCR and its location relative the primary CCR are still valid. These figures confirm that an After-Far secondary CCR location produces slightly longer flow times at low secondary CCR protective capacity levels. The effect of the secondary CCR location and protective capacity level on the mean flow time is however small.

4.3.6.2 The total output for the line

The ANOVA results for the total line output measure are displayed in Table 4-6. The ANOVA performed on the corresponding ranks of the line output measures are displayed in Table 4-7. From the two tables it can be seen that the two ANOVA's give basically similar results. The ANOVA based on the original data however shows that all main effects are significant, whereas the ANOVA based on the ranks suggests that the main effect of the amount of protective capacity at the secondary CCR is not significant. Furthermore the ANOVA based on the ranks shows a significant three-way interaction between the amount of protective capacity at the secondary CCR, the amount of protective capacity at the non-CCRs, and the downtime level, whereas the ANOVA based on the original data suggests that this three-way interaction is not significant. Because of the small difference between the two ANOVAs, it is therefore safe to continue the analysis based on the ANOVA for the original non-ranked observations. Since the ANOVA results based on the original data indicates that no statistically significant four-way and three-way interactions exist, only the main effects and significant two-way interactions are further analysed using graphs and Tukey HSD comparisons of the means. The graphs for the main effects are displayed in Figure 4-8, whereas the graphs for the significant two-way interactions are displayed in Figure 4-9.

The graph for the main effect of the location of the secondary CCR on total line output is displayed in Figure 4-8a. Table 4-8 presents the Tukey HSD comparisons showing the means that differ significantly. From the graph in Figure 4-8a it can be seen that the highest line output is obtained with the After-Near secondary CCR location, whereas the lowest line output is obtained with the After-Far location. The Tukey comparisons of the means in Table 4-8 however show that only the total line output for the After-Near and After-Far location differ significantly. The difference in the total line output for the two locations is however very small: less than 1%. Although the statistical analysis

Stellenbosch University http://scholar.sun.ac.za

shows there is a statistically significant difference between the outputs for these two locations, from a practical viewpoint this difference is very small so that it can be concluded that the location of the secondary CCR does not seem to have a significant influence on the total line output.

Looking at the graph of the interaction between the location of the secondary CCR and the downtime level in Figure 4-9a, it can be seen that the slight influence of the secondary CCR location on line output as suggested by the main effect depends on the variability in the line in the form of station downtime. At the low downtime level the location of the secondary CCR has no influence on line output, whereas at the high downtime level the After-Near location seems to obtain the highest line output. The Tukey HSD comparisons of the means in Table 4-8 however show that there is no significant difference between the means of the After-Near and Before-Near locations at the high downtime level, but the mean for the After-Near location does differ significantly from the means of the Before-Far and After-Far locations. Again the differences in the means are very small, so that although statistically significant they can be ignored.

Table 4-6. ANOVA table: original non-ranked line output measures

	SS	Degr. Of Freedom	MS	F	p
Intercept	2.091493E+12	1	2.091493E+12	68599757	0.000000
{1}SecCCRLoc	3.691293E+05	3	1.230431E+05	4	0.007035
{2}SecCCR PC	5.814336E+05	5	1.162867E+05	4	0.001870
{3}NonCCR PC	2.898197E+08	2	1.449098E+08	4753	0.000000
{4}Downtime	3.897268E+08	1	3.897268E+08	12783	0.000000
SecCCRLoc*SecCCR PC	4.248853E+05	15	2.832569E+04	1	0.530412
SecCCRLoc*NonCCR PC	1.950365E+05	6	3.250608E+04	1	0.380245
SecCCR PC*NonCCR PC	3.483647E+05	10	3.483647E+04	1	0.325360
SecCCRLoc*Downtime	3.629036E+05	3	1.209679E+05	4	0.007734
SecCCR PC*Downtime	7.439839E+05	5	1.487968E+05	5	0.000183
NonCCR PC*Downtime	2.902435E+08	2	1.451217E+08	4760	0.000000
SecCCRLoc*SecCCR PC*NonCCR PC	5.805628E+05	30	1.935209E+04	1	0.939038
SecCCRLoc*SecCCR PC*Downtime	4.019752E+05	15	2.679834E+04	1	0.588050
SecCCRLoc*NonCCR PC*Downtime	1.393828E+05	6	2.323047E+04	1	0.599806
SecCCR PC*NonCCR PC*Downtime	4.238173E+05	10	4.238173E+04	1	0.177646
1*2*3*4	5.586716E+05	30	1.862239E+04	1	0.952973
Error	6.541578E+08	21456	3.048834E+04		

Stellenbosch University http://scholar.sun.ac.za

Table 4-7. ANOVA table: ranked line output measures

	SS	Degr. Of Freedom	MS	F	p
Intercept	2.519657E+12	1	2.519657E+12	115374.1	0.000000
{1}SecCCRLoc	2.838310E+08	3	9.461034E+07	4.3	0.004652
{2}SecCCR PC	1.113456E+08	5	2.226912E+07	1.0	0.404011
{3}NonCCR PC	8.207349E+10	2	4.103675E+10	1879.1	0.000000
{4}Downtime	2.036559E+11	1	2.036559E+11	9325.3	0.000000
SecCCRLoc*SecCCR PC	4.496573E+08	15	2.997715E+07	1.4	0.150586
SecCCRLoc*NonCCR PC	2.110437E+08	6	3.517395E+07	1.6	0.139611
SecCCR PC*NonCCR PC	1.160405E+08	10	1.160405E+07	0.5	0.869256
SecCCRLoc*Downtime	2.490218E+08	3	8.300725E+07	3.8	0.009749
SecCCR PC*Downtime	3.650052E+08	5	7.300105E+07	3.3	0.005088
NonCCR PC*Downtime	8.164969E+10	2	4.082485E+10	1869.4	0.000000
SecCCRLoc*SecCCR PC*NonCCR PC	5.185465E+08	30	1.728488E+07	0.8	0.783375
SecCCRLoc*SecCCR PC*Downtime	4.420330E+08	15	2.946887E+07	1.3	0.163024
SecCCRLoc*NonCCR PC*Downtime	1.323791E+08	6	2.206319E+07	1.0	0.416359
SecCCR PC*NonCCR PC*Downtime	5.258941E+08	10	5.258941E+07	2.4	0.007413
1*2*3*4	4.384929E+08	30	1.461643E+07	0.7	0.914386
Error	4.685778E+11	21456	2.183901E+07		

Table 4-8. Tukey HSD comparisons of mean line output for SecCCRLoc main effect

SecCCRLoc	Output	1	2
AF	9834.962	**	
BF	9837.783	**	**
BN	9841.958	**	**
AN	9845.857		**

Table 4-9. Tukey HSD comparisons of mean line output for SecCCR PC main effect

SecCCR PC	Output	1	2
2 Percent	9832.255	**	
5.25 Percent	9836.652	**	**
52.5 Percent	9837.077	**	**
31.5 Percent	9842.933	**	**
21 Percent	9844.914		**
10.5 Percent	9847.011		**

(means within a group that do not significantly differ are indicated with **)

Table 4-10. Tukey HSD comparisons of mean line output for NonCCR PC main effect

NonCCR PC	Output	1	2	3
10.5 Percent	9682.094	**		
21 Percent	9881.847		**	
52.5 Percent	9956.479			**

Table 4-11. Tukey HSD comparisons of mean line output for Downtime main effect

Downtime	Output	1	2
10.04 Percent	9705.816	**	
1.57 Percent	9974.464		**

(means within a group that do not significantly differ are indicated with **)

a. Effect of location of SecCCR on line output

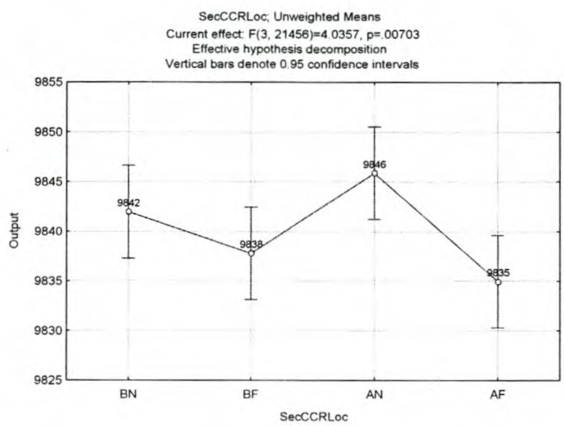

b. Effect of protective capacity at SecCCR on line output

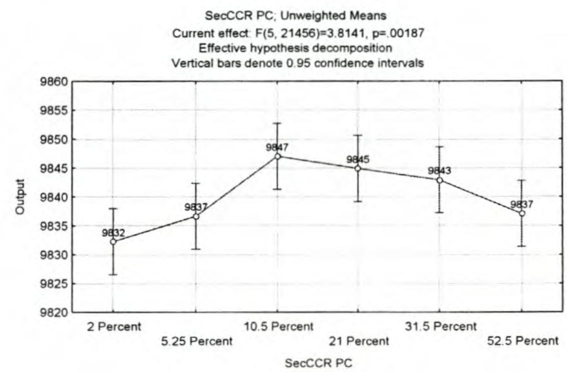

c. Effect of protective capacity at NonCCRs on line output

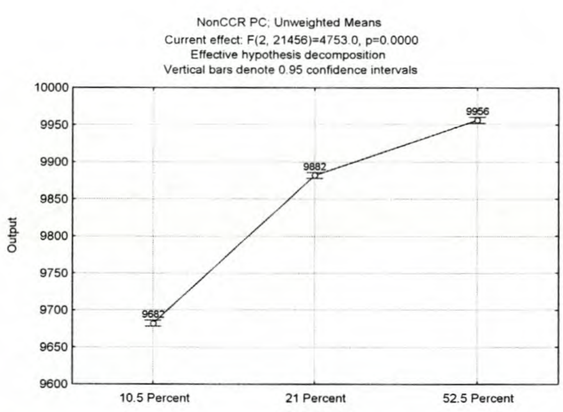

d. Effect of downtime on line output

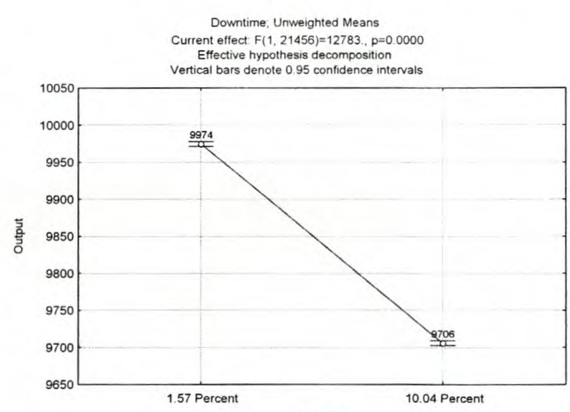

Figure 4-8. Main effects of experimental factors on line output

The graph showing the main effect of the amount of protective capacity at the secondary CCR on line output is displayed in Figure 4-8b. From this graph it seems as if the line output increases up to the 10.5% protective capacity level, and thereafter drops again with further increase in protective capacity. The Tukey HSD comparisons of the means in Table 4-9 however shows that there is no significant difference in the means of the 2%, 5.25%, 52.5% and 31.5% protective capacity levels. Furthermore there is also no significant difference in the means of the 5.25%, 52.5%, 31.5%, 21%, and 10.5%

protective capacity levels. The amount of protective capacity at the secondary CCR therefore does not seem to have much of an influence on the total line output. This was also indicated by the ANOVA based on the ranks in Table 4-7, which shows no significant main effect for the amount of protective capacity at the secondary CCR. The null hypothesis of no effect is therefore not rejected.

a. Interaction of SecCCR location and downtime with respect to line output

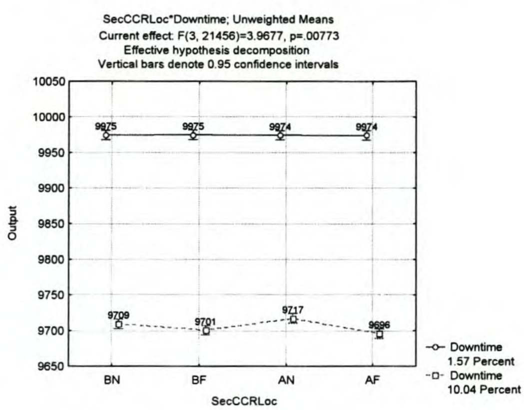

b. Interaction of protective capacity at SecCCR and downtime with respect to line output

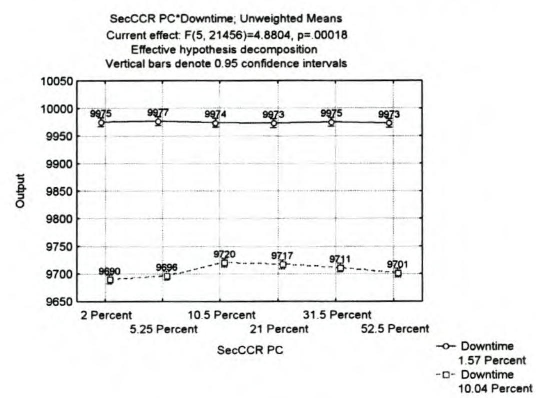

c. Interaction of protective capacity at NonCCRs and downtime with respect to line output

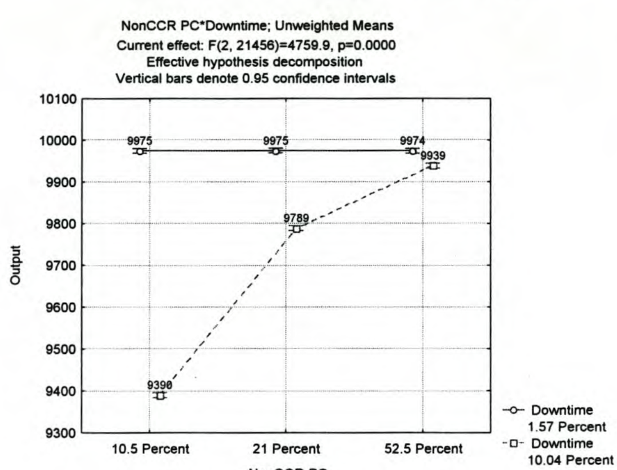

Figure 4-9. Significant two-way interactions between factors with respect to line output

Table 4-12. Tukey HSD comparisons of mean line output for SecCCR PC and Downtime interaction

SecCCR PC	Downtime	Output	1	2	3	4
2 Percent	10.04 Percent	9689.570	**			
5.25 Percent	10.04 Percent	9696.403	**	**		
52.5 Percent	10.04 Percent	9701.140	**	**	**	
31.5 Percent	10.04 Percent	9710.828		**	**	
21 Percent	10.04 Percent	9716.903			**	
10.5 Percent	10.04 Percent	9720.053			**	
21 Percent	1.57 Percent	9972.924				**
52.5 Percent	1.57 Percent	9973.013				**
10.5 Percent	1.57 Percent	9973.968				**
2 Percent	1.57 Percent	9974.941				**
31.5 Percent	1.57 Percent	9975.037				**
5.25 Percent	1.57 Percent	9976.901				**

(means within a group that do not significantly differ are indicated with **)

Table 4-13. Tukey HSD comparisons of mean line output for NonCCR PC and Downtime interaction

NonCCR PC	Downtime	Output	1	2	3	4
10.5 Percent	10.04 Percent	9389.670	**			
21 Percent	10.04 Percent	9789.036		**		
52.5 Percent	10.04 Percent	9938.742			**	
52.5 Percent	1.57 Percent	9974.216				**
10.5 Percent	1.57 Percent	9974.518				**
21 Percent	1.57 Percent	9974.658				**

(means within a group that do not significantly differ are indicated with **)

The interaction graph between the amount of protective capacity at the secondary CCR and the downtime level in Figure 4-9b shows that this small influence of the amount of protective capacity at the secondary CCR on line output as observed in the main effect graph in Figure 4-8b, is only visible at the high downtime level. The Tukey HSD comparisons of the means for the protective capacity and downtime interaction are displayed in Table 4-12. This comparison of the means confirms that there is no significant difference in the means for the low downtime (1.57%) situation. For the high downtime situation there is also not any clear group of means that significantly differs form the other groups. The null hypothesis of no interaction effect between the protective capacity at the secondary CCR and the downtime level is therefore not rejected.

The largest main effects on the line output are caused by the amount of protective capacity at the non-CCRs and the downtime level. This is evident from the large F-values for the main effects of these two factors in the ANOVA results of Table 4-6. The graph for the main effect of the amount of protective capacity at the non-CCRs is displayed in Figure 4-8c, and the Tukey comparison of the means is displayed in Table 4-10. From the graph in Figure 4-8c it is evident that an increase in the amount of protective capacity at the non-CCRs leads to an increase in line output. Increasing the protective capacity however has diminishing returns. The increase in line output by increasing the protective capacity from 10.5% to 21% is 2.07% [100*(9882-9682)/9682], whereas the increase in line output by increasing the protective capacity from 21% to 52.5% is only 0.75% [100*(9956-9882)/9882]. A larger increase in line output is therefore obtained by increasing the protective capacity at the non-CCRs from 10.5% to 21%, than from 21% to 52.5%.

By looking at the two-way interaction between the amount of protective capacity at the non-CCRs and the level of downtime (refer to Figure 4-9c), it is evident that this effect of non-CCR protective capacity on line output is only visible at the high downtime situation. This is confirmed by the Tukey HSD comparisons of the means in Table 4-13, which shows that there is no significant difference in the means for the low downtime (1.57%) situation, whereas all the means differ significantly for the high downtime situation. Again the largest increase in line output for the high downtime situation is obtained by increasing the amount of protective capacity at the non-CCRs from 10.5% to 21%. A 4% increase in line output is obtained when the protective capacity at the non-CCRs is increased from 10.5% to 21%, whereas a 1.5% increase in line output is obtained when the protective capacity at the non-CCRs is increased from 21% to 52.5%. The maximum theoretical output for the line with a 95% constraint utilisation is 9975 jobs [0.95*(simulation time – warmup time)/constraint process time = 0.95*42 000 min/4min]. For the low downtime situation this output level was reached at all non-CCR protective capacity levels, as well as all secondary CCR protective capacity levels. This is an indication that at low variability situations sufficient WIP protection in front of the primary CCR occurs at low protective capacity levels to prevent significant starvation of the primary CCR. For the high downtime situation however this output level was only reached at the 52.5% non-CCR protective capacity level. This suggests that protective capacity as a means for protecting line output is only necessary in lines with resources that experience long breakdowns. This could be explained by the fact that loss in output is caused when the primary constraint is starved for work. At low protective capacity levels WIP is more spread out among the different stations. When a resource breakdown occurs upstream from the primary CCR, not enough WIP protection therefore exists between the failed resource and the primary CCR and significant starvation occurs at the primary CCR. At high protective capacity levels jobs are moved faster to the primary CCR so that most of the WIP build-up takes place in front of the primary CCR. This reduces the number and duration of starvations at the primary CCR due to upstream resource failures.

The graph for the main effect of the level of downtime on line output is displayed in Figure 4-8d. From this graph it is evident that a high amount of variability in the form of station downtime significantly decreases line output, which confirms previous research results.

As in the case of the flow time dependent variable, again the secondary CCR protective capacity levels where it ceases to be the secondary constraint should be identified. For the low downtime situation in Figure 4-10 the three-way interaction between the amount of protective capacity at the secondary CCR, the amount of protective capacity at the non-CCRs, and the location of the secondary CCR in indicates that these factors have no influence on the line output under low downtime conditions. For the high downtime situation in Figure 4-11 it is again not necessary to analyse the 10.5% non-CCR protective capacity level, because at this level the so-called secondary CCR is actually never a secondary constraint. For the 21% non-CCR protective capacity level the secondary CCR protective capacity level where it ceases to be the secondary constraint is determined as 21% - 10.04% = 10.96%, and for the 52.5% non-CCR protective capacity level it is 52.5%-10.04% = 42.46% (indicated by the vertical dotted lines). Both Figure 4-10 and Figure 4-11 indicate that the earlier analyses made regarding the protective capacity level at the secondary CCR and its location relative the primary CCR are still valid. These figures confirm that the protective capacity level at the secondary CCR and its location relative the primary CCR do not have a significant influence on line output.

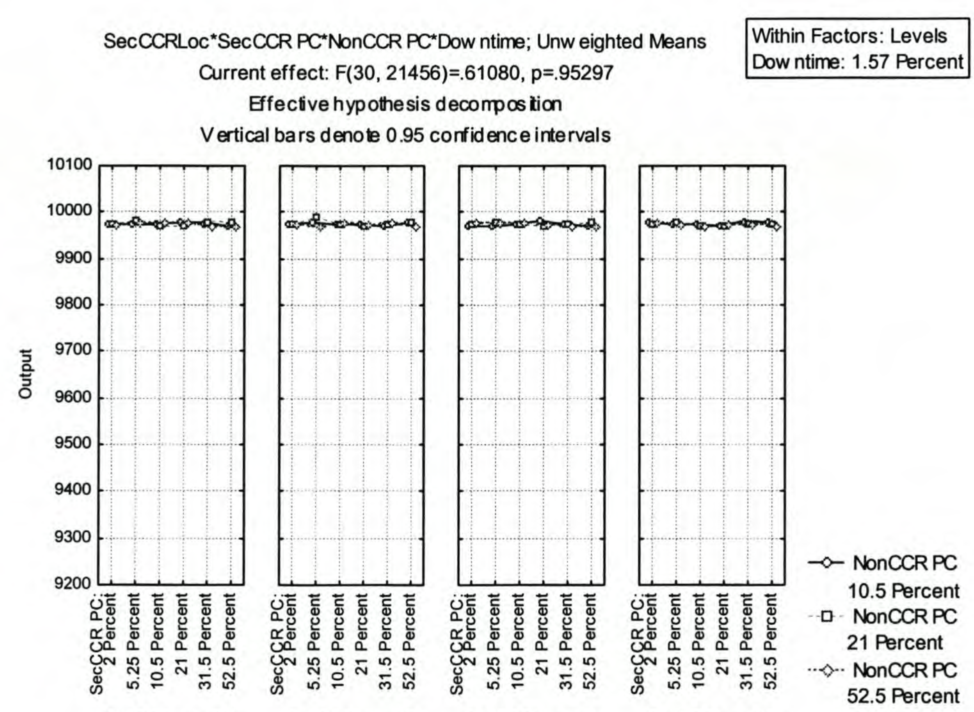

Figure 4-10. Three-way interaction of SecCCR Loc, SecCCR PC and NonCCR PC on the line output: low downtime situation

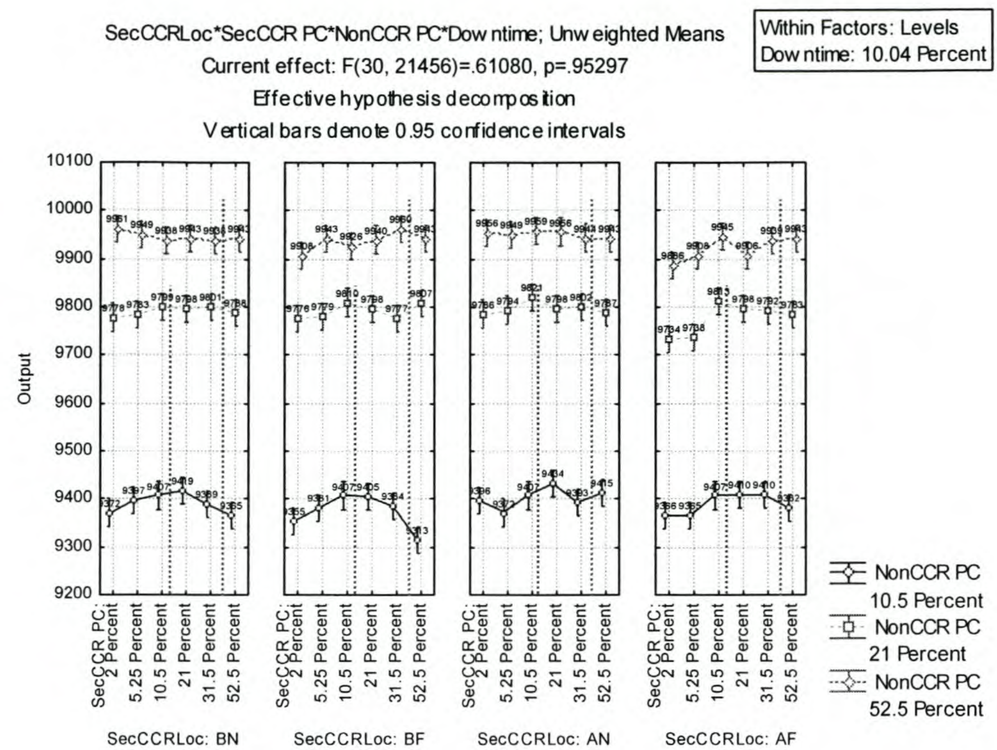

Figure 4-11. Three-way interaction of SecCCR Loc, SecCCR PC and NonCCR PC on the line output: high downtime situation

4.3.6.3 The bottleneck probability for the primary CCR

The ANOVA results for the bottleneck probability of the primary CCR measure are displayed in Table 4-14. The ANOVA performed on the corresponding ranks of the bottleneck probability values are displayed in Table 4-15. From the two tables it can be seen that both ANOVAs give exactly the same results. The ANOVA tables indicate that all main effects, two-way interactions, three-way interactions, as well as four-way interactions are significant. The graphs displaying the four-way interactions are presented in Figure 4-13 and Figure 4-14. Figure 4-13 displays the interaction between the amount of protective capacity at the secondary CCR, the amount of protective capacity at the non-CCRs, and the location of the secondary CCR when stations experience low downtimes, whereas Figure 4-14 displays this interaction for the high downtime situation.

Table 4-14. ANOVA table: original non-ranked primary CCR bottleneck probability measures

	SS	Degr. Of Freedom	MS	F	p
Intercept	58450315	1	58450315	941975.5	0.000000
{1}SecCCRLoc	446715	3	148905	2399.7	0.000000
{2}SecCCR PC	526588	5	105318	1697.3	0.000000
{3}NonCCR PC	6534201	2	3267101	52652.0	0.000000
{4}Downtime	2333432	1	2333432	37605.2	0.000000
SecCCRLoc*SecCCR PC	525222	15	35015	564.3	0.000000
SecCCRLoc*NonCCR PC	198768	6	33128	533.9	0.000000
SecCCR PC*NonCCR PC	129209	10	12921	208.2	0.000000
SecCCRLoc*Downtime	92718	3	30906	498.1	0.000000
SecCCR PC*Downtime	3090	5	618	10.0	0.000000
NonCCR PC*Downtime	882816	2	441408	7113.7	0.000000
SecCCRLoc*SecCCR PC*NonCCR PC	103055	30	3435	55.4	0.000000
SecCCRLoc*SecCCR PC*Downtime	75241	15	5016	80.8	0.000000
SecCCRLoc*NonCCR PC*Downtime	75204	6	12534	202.0	0.000000
SecCCR PC*NonCCR PC*Downtime	10953	10	1095	17.7	0.000000
1*2*3*4	53401	30	1780	28.7	0.000000
Error	1331362	21456	62		

Table 4-15. ANOVA table: ranked primary CCR bottleneck probability measures

	SS	Degr. Of Freedom	MS	F	p
Intercept	2.519657E+12	1	2.519657E+12	599455.5	0.000000
{1}SecCCRLoc	3.017588E+10	3	1.005863E+10	2393.1	0.000000
{2}SecCCR PC	4.036258E+10	5	8.072516E+09	1920.5	0.000000
{3}NonCCR PC	4.699812E+11	2	2.349906E+11	55907.0	0.000000
{4}Downtime	1.010052E+11	1	1.010052E+11	24030.3	0.000000
SecCCRLoc*SecCCR PC	2.750305E+10	15	1.833536E+09	436.2	0.000000
SecCCRLoc*NonCCR PC	1.945377E+10	6	3.242295E+09	771.4	0.000000
SecCCR PC*NonCCR PC	1.457118E+10	10	1.457118E+09	346.7	0.000000
SecCCRLoc*Downtime	3.515445E+09	3	1.171815E+09	278.8	0.000000
SecCCR PC*Downtime	2.034589E+09	5	4.069178E+08	96.8	0.000000
NonCCR PC*Downtime	2.288149E+10	2	1.144074E+10	2721.9	0.000000
SecCCRLoc*SecCCR PC*NonCCR PC	8.511723E+09	30	2.837241E+08	67.5	0.000000
SecCCRLoc*SecCCR PC*Downtime	1.707258E+09	15	1.138172E+08	27.1	0.000000
SecCCRLoc*NonCCR PC*Downtime	4.925752E+09	6	8.209586E+08	195.3	0.000000
SecCCR PC*NonCCR PC*Downtime	4.195186E+08	10	4.195186E+07	10.0	0.000000
1*2*3*4	2.560335E+09	30	8.534450E+07	20.3	0.000000
Error	9.018479E+10	21456	4.203243E+06		

a. Effect of location of SecCCR on primary CCR bottleneck probability

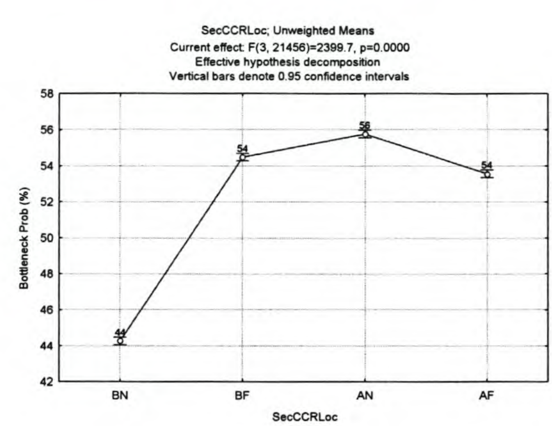

b. Effect of protective capacity at SecCCR on primary CCR bottleneck probability

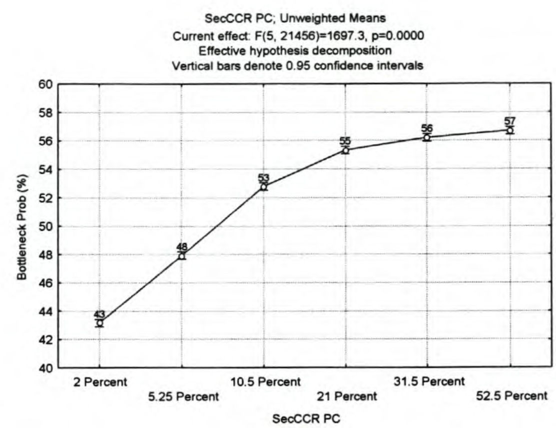

c. Effect of protective capacity at NonCCRs on primary CCR bottleneck probability

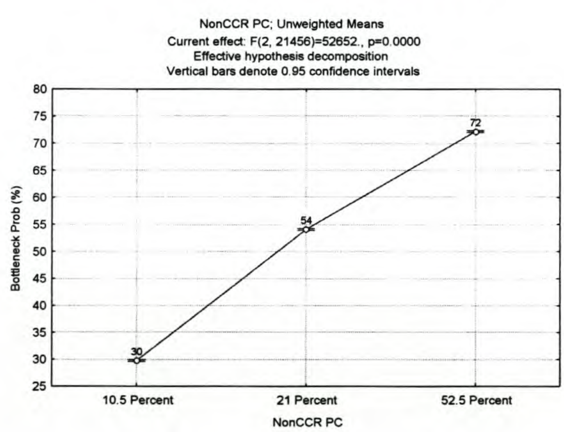

d. Effect of downtime on primary CCR bottleneck probability

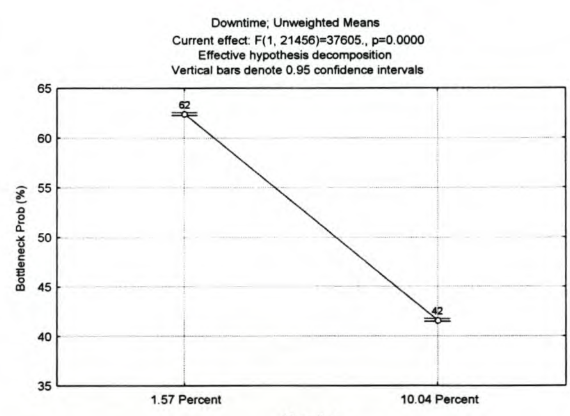

Figure 4-12. Main effects of experimental factors on primary CCR bottleneck probability

Table 4-16. Tukey HSD comparisons of primary CCR bottleneck probability for SecCCRLoc main effect

SecCCRLoc	Bottl. Prob	1	2	3	4
BN	44.26027	**			
AF	53.56709		**		
BF	54.48083			**	
AN	55.76998				**

Table 4-17. Tukey HSD comparisons of primary CCR bottleneck probability for SecCCR PC main effect

SecCCR PC	Bottl. Prob	1	2	3	4	5
2 Percent	43.16903	**				
5.25 Percent	47.91646		**			
10.5 Percent	52.77672			**		
21 Percent	55.33763				**	
31.5 Percent	56.22195					**
52.5 Percent	56.69547					**

(means within a group that do not significantly differ are indicated with **)

Table 4-18. Tukey HSD comparisons of primary CCR bottleneck probability for NonCCR PC main effect

NonCCR PC	Output	1	2	3
10.5 Percent	29.75500	**		
21 Percent	54.09746		**	
52.5 Percent	72.20618			**

Table 4-19. Tukey HSD comparisons of primary CCR bottleneck probability for Downtime main effect

Downtime	Output	1	2
10.04 Percent	41.62583	**	
1.57 Percent	62.41326		**

(means within a group that do not significantly differ are indicated with **)

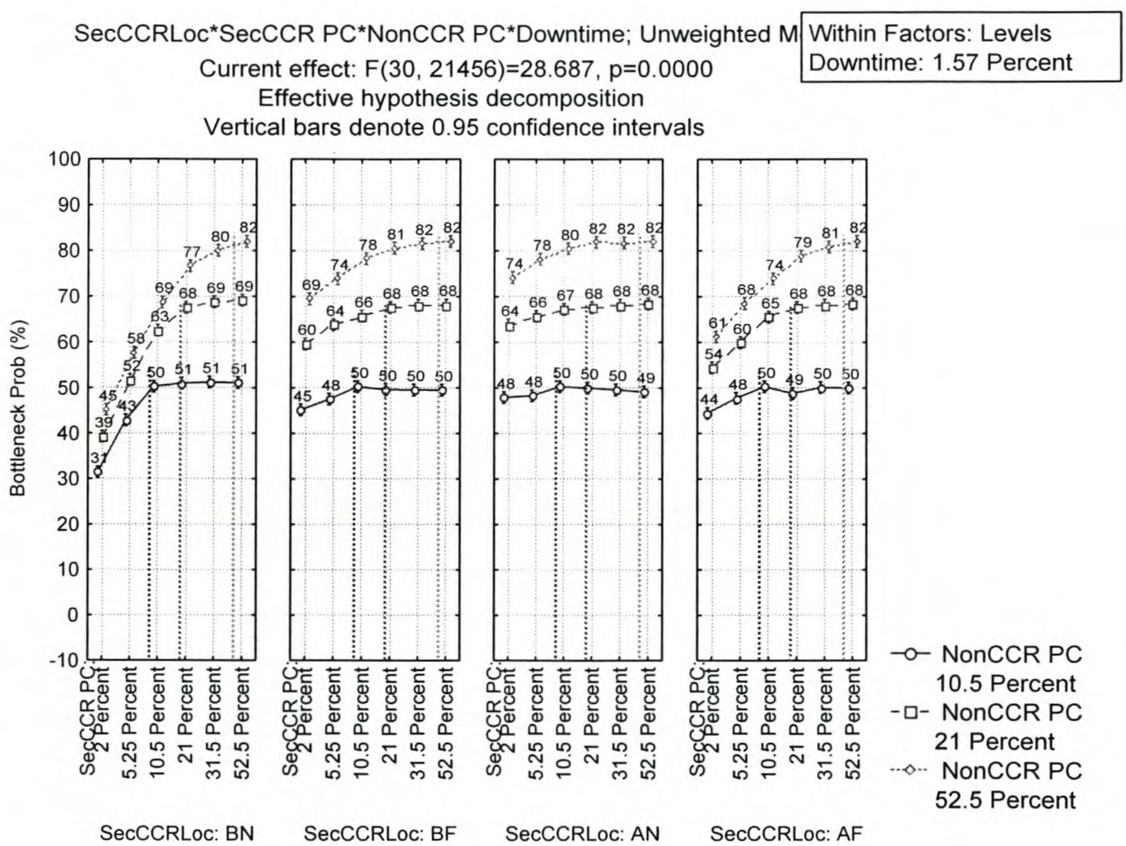

Figure 4-13. Three-way interaction of SecCCR Loc, SecCCR PC and NonCCR PC on the primary CCR bottleneck probability: low downtime situation

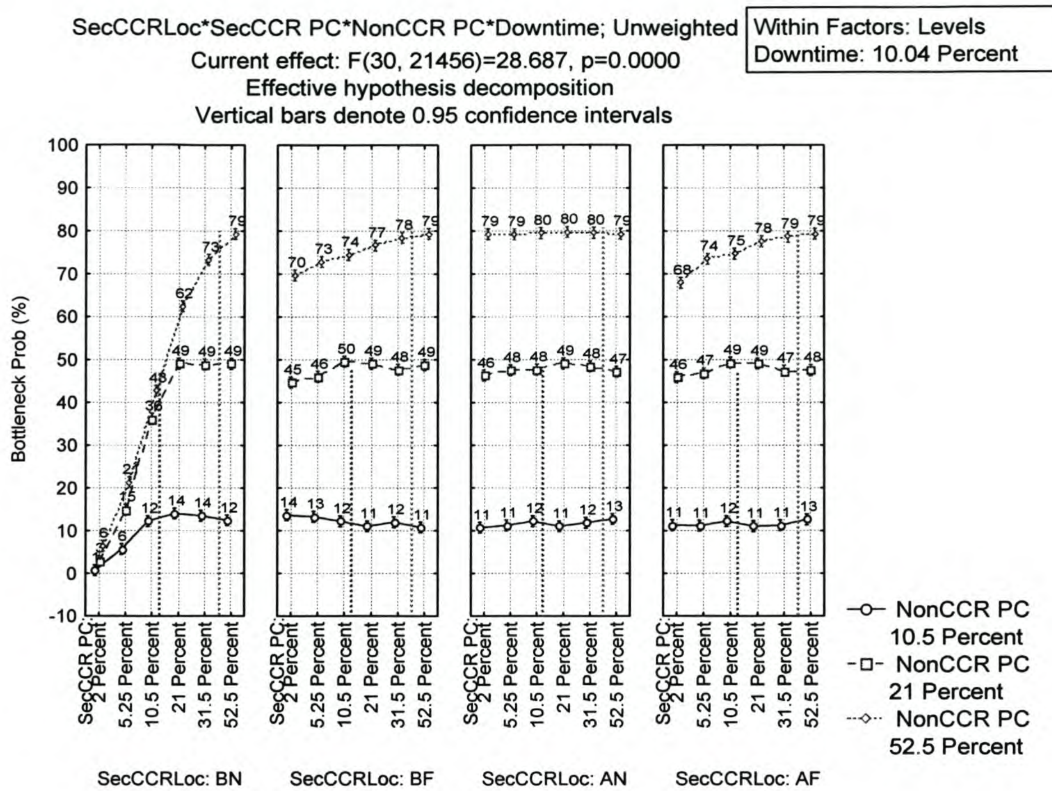

Figure 4-14. Three-way interaction of SecCCR Loc, SecCCR PC and NonCCR PC on the primary CCR bottleneck probability: high downtime situation

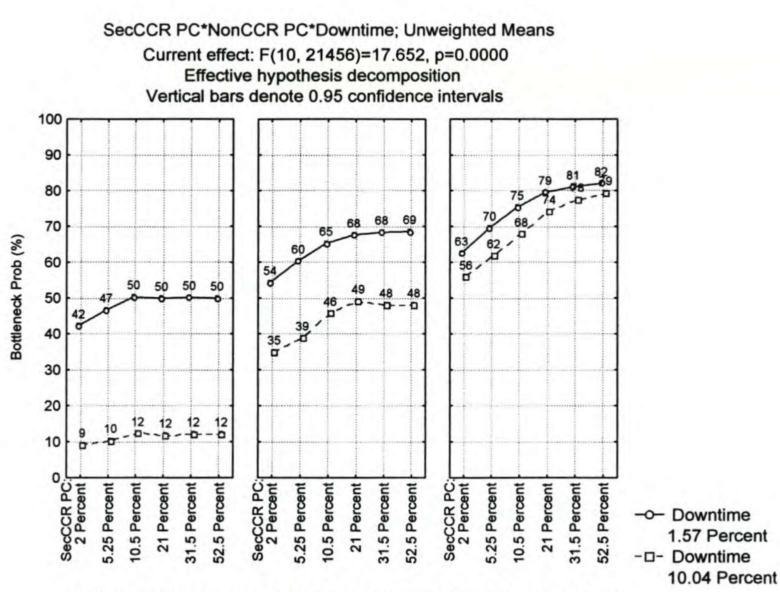

Figure 4-15. Three-way interaction of SecCCR PC, NonCCR PC, and Downtime on the primary CCR bottleneck probability

Stellenbosch University http://scholar.sun.ac.za

The vertical dotted lines in these graphs indicate for each non-CCR protective capacity level the secondary CCR protective capacity level where the identified secondary constraint's protective capacity equals the non-CCR protective capacity and therefore ceases to be the secondary constraint. These four-way interaction figures indicate that under both the high and low downtime situations, the Before-Near secondary CCR location has the largest influence on the bottleneck probability of the primary CCR when the amount of protective capacity at the secondary CCR is relatively low. This is because at this location the secondary CCR paces the flow of materials into the primary constraint resource. Refer to Figure 4-14. For this high downtime situation, at all three non-CCR protective capacity levels the bottleneck probability for the primary CCR is less than 7% when the amount of protective capacity at the secondary CCR is 2% and the secondary CCR is located before but near the primary CCR. This is because statistical fluctuation accumulates and 2% protection is inadequate, which therefore starves the constraint. Increasing the amount of protective capacity at the secondary CCR for the Before-Near location causes an increase in the bottleneck probability. The bottleneck probability however appears to reach a certain maximum level where further protective capacity increases at the secondary CCR for a Before-Near location do not significantly influence the bottleneck probability any more. As the capacity of the secondary CCR approaches the capacity of the non-constraints then the impact of the secondary CCR disappears because it is no longer the secondary constraint. Increasing the protective capacity at the secondary CCR for a Before-Near location therefore has diminishing returns. For the 10.5% non-CCR protective capacity level this maximum bottleneck probability appears to be around 12% and is reached at a 10.5% secondary CCR protective capacity level. For the 21% non-CCR protective capacity level this maximum bottleneck probability appears to be around 49% and is reached near the 21% secondary CCR protective capacity level. When the non-CCR protective capacity level is 52.5%, this maximum bottleneck probability appears to be around 79% and is only reached at a 52.5% secondary CCR protective capacity level.

For all four secondary CCR locations the maximum bottleneck probability reached for the high downtime situation seems to be around 80% for high protective capacity levels at the non-CCRs (52.5%), 50% for medium protective capacity levels at the non-CCRs (21%), and 12% for low protective capacity levels at the non-CCRs (10.5%). For the 10.5% non-CCR protective capacity level the secondary constraint is actually the second last workstation (station 14). This station has a slightly higher bottleneck probability than the other stations that experience breakdowns. Station 14 is the secondary constraint because of the breakdowns it experience and because of its position near the end of the line and the principle of statistical fluctuations accumulating across resources. In order to ensure that under high variability situations (characterised by high resource breakdowns) the primary CCR is the only bottleneck or constraint for more than 50% of the time, the results suggest that the average protective capacity level at non-CCRs need to be at least greater than 20% with a level arrangement. However, to ensure a relatively stable primary CCR with an 80% bottleneck probability, the average protective capacity at non-CCRs need to be around 50%. The After-Near secondary CCR location appears to have the least influence on the bottleneck probability for different protective capacity levels at the secondary CCR under high variability conditions. Tukey HSD comparisons showed no difference between the bottleneck probability means for different secondary CCR protective capacity levels at the 10.5% non-CCR protective capacity level, between the means at the 21% non-CCR protective capacity level, or between means at the 52.5% non-CCR protective capacity level for this secondary CCR location and high station downtimes. This suggests that when the secondary CCR is located After-Near the primary CCR under high variability conditions, the bottleneck probability for the primary CCR is unaffected by low, medium or high protective capacity levels at the secondary CCR. The bottleneck probability for the primary CCR is also relatively unaffected by different secondary CCR protective capacity levels when the secondary CCR is located Before-Far and After-Far relative to the primary CCR, and the amount of protective capacity at the non-CCRs is either 10.5% or 21%. When the level of protective capacity at the non-CCRs is however high (52.5%), lower protective capacity levels at the secondary CCR for a Before-Far or After-Far location causes a decrease in the bottleneck probability. For the 2% secondary CCR protective capacity level the bottleneck probability is around 69% for both the Before-Far and After-Far locations when the non-CCR protective capacity level is 52.5%. This represents a nearly 14% decrease in the bottleneck probability from the maximum bottleneck probability level of 80%.

When looking at the low downtime situation in Figure 4-13, the Before-Near secondary CCR location again has the largest influence on the bottleneck probability when the amount of protective capacity at the secondary CCR is relatively low. The minimum bottleneck probability in this case is however 31% for a 2% protective capacity level at the secondary CCR and a 10.5% protective capacity level at the non-CCRs. Increasing the amount of protective capacity at the secondary CCR when located Before-Near the primary CCR again increases the bottleneck probability until a certain maximum level is reached. Increasing the protective capacity at the secondary CCR therefore has diminishing returns. For the 10.5% non-CCR protective capacity level this maximum bottleneck probability appears to be around 51% and is reached at a 10.5% secondary CCR protective capacity level. For the 21% non-CCR protective capacity level this maximum bottleneck probability appears to be around 69% and is reached near the 21% secondary CCR protective capacity level. When the non-CCR protective capacity level is 52.5%, this maximum bottleneck probability appears to be around 80%-82% and is reached near a 31.5% secondary CCR protective capacity level.

The Before-Far and After-Near secondary CCR locations appear to have more or less the same influence on the bottleneck probability for different secondary CCR protective capacity levels at the low downtime situation. These two locations also seem to have less of an influence on the bottleneck probability than the Before-Near and After-Far locations.

When the protective capacity level at the non-CCRs is low (10.5%), the amount of protective capacity at the secondary CCR does not seem to have much of an influence on the bottleneck probability for the Before-Far, After-Near and After-Far locations under the low variability situation. This is because at this low non-CCR protective level the so-called secondary CCR is actually not the secondary constraint. However for the medium (21%) and high (52.5%) non-CCR protective capacity levels, low protective capacity levels at the secondary CCR cause a significant decrease in the bottleneck probability. The fastest decrease however appears to occur when the amount of protective capacity at the secondary CCR is less than 10.5%. The largest decrease in the bottleneck probability also occurs for the After-Far secondary CCR location.

For all four secondary CCR locations at the low downtime situation the maximum bottleneck probability reached seems to be around 80% for high protective capacity levels at the non-CCRs (52.5%), 69%-70% for medium protective capacity levels at the non-CCRs (21%), and 50% for low protective capacity levels at the non-CCRs (10.5%). This suggests that in order to ensure that the primary CCR is the only bottleneck or constraint for more than 50% of the time, the average protective capacity level at non-CCRs need to be at least greater than 10% for low downtime situations. However, to ensure a relatively stable primary CCR with a 70% bottleneck probability, the average protective capacity at non-CCRs need to be at least 20%.

For both the low and high downtime situations in Figure 4-13 and Figure 4-14 the maximum bottleneck probability appears to be around 80% when the amount of protective capacity at the non-CCRs is very high (52.5%). This suggest that no matter what the variability in the line, quite a substantial amount of protective capacity is needed at non-CCRs in order to ensure a very stable primary CCR. The question however is what percentage of bottleneck probability is adequate for planning and operational purposes in the line. For example under the low downtime situation a bottleneck probability of around 70% can be reached with medium levels of protective capacity at the non-CCRs (more or less 20%), which could be sufficient.

Graphs for the main effects of the different factors are displayed in Figure 4-12. These main effects confirm some of the observations made while investigating the four-way interaction. The main effect of the location of the secondary CCR is displayed in Figure 4-12a. From this graph it can be seen that the Before-Near secondary CCR location produces the lowest bottleneck probability, whereas the After-Near location produces the highest bottleneck probability. The Tukey HSD comparisons of the means in Table 4-16 confirm that the means for all four locations differ significantly. The means for the Before-Far, After-Near and After-Far locations however do not differ that much. It is therefore safe to assume that their influence on the bottleneck probability of the primary CCR is more or less the same. This was also indicated by the four-way interaction of the factors.

The main effect for the amount of protective capacity at the secondary CCR on the bottleneck probability is displayed in Figure 4-12b. This graph also indicates that reducing the protective capacity at the secondary CCR below 10% causes the fastest drop in the bottleneck probability. As indicated by the investigation of the four-way interaction this effect however depends largely on the location of the secondary CCR relative to the primary CCR, and to a lesser extent on the amount of protective capacity at the non-CCRs.

The main effect for the amount of protective capacity at the non-CCRs in Figure 4-12c highlights the fact that increasing the protective capacity at all non-CCR stations causes a significant increase in the bottleneck probability. It also shows that in order to ensure that the primary CCR remains the only constraint or bottleneck for more than half the time, the protective capacity at the non-CCRs need to be at least around 20% or higher. The investigation of the four-way interaction however reveals that such a protective capacity level is only needed at the high downtime situation. For the low downtime situation, the protective capacity at non-CCRs need to be only larger than 10 % in order to ensure a 50% or more bottleneck probability.

The main effect for the downtime level in Figure 4-12d shows that the bottleneck probability is substantially lower for the high downtime situation than for the low downtime situation. The four-way interaction however revealed that with enough protective capacity at non-CCRs the bottleneck probabilities for the two downtime situations could be brought closer together. This is also illustrated by the three-way interaction between the amount of protective capacity at the secondary CCR, the amount of protective capacity at the non-CCRs, and the downtime level in Figure 4-15. This figure shows that the graphs for the two downtime levels lie much closer together under higher non-CCR protective capacity levels.

4.3.7 Answers to research questions for the flow shop experiment

To summarise the discussion of the simulation results in section 4.3.6 and answer the research questions posed in section 4.2, the following conclusions can be made:

1) *Research question 1: Does the location of the secondary CCR relative to the primary CCR affect the mean flow time, the total output and the primary CCR bottleneck probability?*
    - The location of the secondary constraint relative to the primary constraint does not have that much of an effect on the mean flow time for jobs through the flow shop. If possible, it could

however be tried to avoid having a secondary constraint located after but relatively far from the primary constraint station (or near the end of the flow line), since this could lead to slightly inflated flow times when there is relatively high variability in the line. This could be explained by the fact that variability does not average out, but actually propagates through the line. A constraint or secondary constraint located near the end of the line is therefore subjected to a higher variability in job arrivals than when located near the beginning of the line (especially when there is high variability present at stations). As evident from the queuing analysis presented in the previous chapter a higher arrival variability causes longer queue wait times, which in turn inflates the total flow time.

- The location of the secondary constraint relative to the primary constraint does not have any significant effect on the total output for the line. It was expected that the location of the secondary constraint would have an effect. In this study the secondary constraint however did not experience any breakdowns. The location of a secondary constraint experiencing breakdowns could possibly have a more significant impact on line output.
- The location of the secondary constraint relative to the primary constraint does have a significant effect on the bottleneck probability of the primary constraint. To ensure a stable primary constraint that does not move too often in the line it should be avoided having the secondary constraint located before and near the primary constraint operation. The other secondary constraint locations do not have that much of an influence on the bottleneck probability of the primary constraint. It however appears to be best to have the secondary constraint located after but near the primary constraint operation in order to ensure a stable primary constraint location. At a before-near location jobs seem to queue at the secondary constraint instead of at the primary constraint causing the lower bottleneck probability for the primary constraint. It was expected that this should lead to a drop in total line output, since there is less WIP protection in front of the constraint. The simulation analysis for the line output measure however showed no significant drop in line output for a before-near secondary constraint location. This could be explained by the fact that the secondary constraint did not experience any resource downtimes. The only variability was the normal processing time variability, which was relatively low. Because of this low variability at the secondary constraint the primary constraint was not significantly starved due to the fact that jobs queued at the secondary constraint instead of at the primary constraint. The WIP protection was therefore still relatively close to the primary constraint.

2) *Research question 2: What is the effect of the amount of protective capacity at the secondary CCR?*

- The amount of protective capacity at the secondary constraint do not have too much of an influence on the mean flow time. The largest influence is observed when the secondary

Stellenbosch University http://scholar.sun.ac.za

constraint is located after but relatively far from the primary constraint. Very low protective capacity levels for the secondary constraint at this location causes longer flow times than for the other locations. Increasing the amount of protective capacity at the secondary constraint for this location improves the flow time measure. This however comes with diminishing returns. It seems that a 10 % secondary constraint protective capacity level is sufficient when the secondary constraint is located after but relatively far from the primary constraint. Higher levels of protective capacity at the secondary constraint do not cause much more improvement in the mean flow time. The higher flow time at the after-far location can again be ascribed to the fact that variability propagates through the line, therefore stations near the end of the line experience higher job arrival variability. A secondary constraint with low protective capacity located near the end of the line therefore experiences longer queue wait times because of the higher arrival variability, which in turn leads to longer flow times. Increasing the protective capacity at the secondary constraint diminishes the negative effect of the higher arrival variability. For the other secondary constraint locations the amount of protective capacity at the secondary constraint is not that critical with respect to mean flow time. Flow time therefore does not seem to be much affected by low protective capacity levels at a single non-primary constraint station.

- The amount of protective capacity at the secondary constraint has no significant effect on the total output of the line. Line output therefore does not seem to be affected by low protective capacity levels at a single station that is not the primary constraint, irrespective of the location of this station relative to the primary constraint.
- The amount of protective capacity at the secondary constraint does have a significant effect on the bottleneck probability of the primary constraint. Increasing the amount of protective capacity at the secondary constraint from low levels increases the bottleneck probability of the primary constraint up to a certain maximum level. This maximum level seems to be determined by the average protective capacity available at all other stations besides the primary and secondary constraints. The effect of the amount of protective capacity at the secondary constraint on the bottleneck probability depends on the location of the secondary constraint relative to the primary constraint. Low protective capacity levels at the secondary constraint significantly reduce the bottleneck probability of the primary constraint when the secondary constraint is located before but near the primary constraint operation. For all the other locations the detrimental effect of low protective capacity levels at the secondary constraint is more evident for lines with low variability in the form of station downtimes, and when there is medium to high levels of average protective capacity available at non-constraint stations.

3) *Research question 3: What is the effect of the amount of protective capacity at the non-constraint stations?*

- The amount of protective capacity at the non-constraint resources has a very significant effect on the mean flow time measure. Evenly increasing the protective capacity at all non-constraint resources drastically reduces the mean flow time, especially when there is high variability in the line due to high station downtimes. Increasing the amount of protective capacity at the non-constraint resources however has diminishing returns on the flow time. These results agree with previous research.
- Evenly increasing the amount of protective capacity at the non-constraint resources significantly improves the total line output when there is high variability in the line due to high resource downtimes. For low downtime situations the protective capacity level at the non-constraints have no significant effect on the total line output. Under high downtime situations the primary constraint station is significantly "starved" for work, which causes a drop in line output. Increasing the protective capacity at the non-constraints however help to move jobs faster to the primary constraint operation. Jobs therefore mainly accumulate in front the constraint operation and are not spread out. When breakdowns occur at stations upstream from the primary constraint, less WIP is therefore held up at the failed station, so that the primary constraint is less starved for work. Increasing the amount of protective capacity at the non-constraint resources however has diminishing returns on the line output. For the high downtime situation very high protective capacity levels at the non-constraints are needed to obtain the maximum designed line output as specified by the primary constraint. At low downtime situations however low levels of protective capacity are sufficient to obtain this maximum designed line output. The specific level of protective capacity needed to obtain the maximum line output therefore depends on the amount of variability in the line. For very low variability situations the results suggest that average protective capacity levels of around 10% seem to be sufficient. More extensive experimentation with a wider range of low to high protective capacity levels at non-constraints and a wider range of variability levels is however necessary to see whether levels of protective capacity of even lower than 10% would be sufficient to obtain the maximum line output.
- The bottleneck probability for the primary constraint is significantly influenced by the amount of protective capacity at the non-constraint stations. Increasing the amount of protective capacity at the non-constraint stations increases the probability for the primary constraint station of being the only long-term constraint. For low station downtime situations a protective capacity level of at least 10% is needed at non-constraint stations to ensure that the primary constraint station will be the only constraint for more than half the time. In order to have a more stable primary constraint that does not move too often in the line, it appears as if protective capacity levels of more than 20% are needed at non-constraint stations (for low

downtime situations). For a bottleneck probability of around 80% as much as 50% protective capacity is needed at non-constraint stations. When there is high variability in the line due to high station downtimes, a protective capacity level of around 20% is needed to ensure a bottleneck probability of more than 50%. For a more stable primary constraint with a bottleneck probability of around 80%, again a protective capacity level of as much as 50% is needed. These levels of protective capacity are for an even assignment of protective capacity at non-constraint stations. As illustrated by the amount of protective capacity at the secondary constraint factor, different levels of protective capacity at a few individual stations (i.e. different from the average protective capacity level at non-constraints) can have a significant impact on the bottleneck probability, especially when these stations are located before and near the primary constraint, or after but relatively far from the primary constraint.

4.4 Assembly Line Experiment

In this experiment an assembly type of flow line was investigated where a total work content needs to be distributed among the stations in the line. The number of stations in the line depends on the specific distribution of the work content. Starting with a balanced line where all stations have the same work content, the addition of protective capacity in this configuration requires additional stations over and above the original number of stations in the balanced line. In this study it was investigated how the flow time is influenced by increases in protective capacity when the number of stations is also increased correspondingly, and what levels of protective capacity are typically needed to obtain a good balance between flow time improvements and an increase in the number of stations.

Six different line configurations were investigated: (1) a short line (with 6 initial stations) with infinite buffer capacity between stations and a low processing time variability; (2) a short line (with 6 initial stations) with infinite buffer capacity between stations and a high processing time variability; (3) a short line (with 6 initial stations) with a finite or limited buffer capacity of one between stations and a low processing time variability; (4) a long line (with 20 initial stations) with infinite buffer capacity between stations and a low processing time variability; (5) a long line (with 20 initial stations) with infinite buffer capacity between stations and a high processing time variability; and (6) a long line (with 20 initial stations) with a finite or limited buffer capacity of one between stations and a low processing time variability. Arena 3.0 was used to develop simulation models of the different line configurations, and each of the different line configurations was analysed separately with single factor ANOVAs. The mean flow time for jobs through the line was the dependent variable. Graphical analysis was used to further analyse the results. All line configurations were modelled as open systems where jobs arrive according to a certain arrival distribution, and finished jobs leave the system. Jobs immediately proceed to the next station after finished processing, therefore no transfer batches were modelled. All stations consisted of a single server, and transfer times between stations were assumed negligible.

The independent variable for all the different line configuration experiments was the amount of protective capacity at the non-constraint stations. Protective capacity is again defined as the complement of the percentage difference in the output rate of the constraint station relative to the non-constraint station. Protective capacity was varied at six levels: 0%, 2%, 5.25%, 10.5%, 16%, and 31.5% for the infinite buffer lines, and 0%, 2%, 5%, 10%, 15%, and 30% for the finite buffer lines. The 0% level represented a balanced line. A flat arrangement of protective capacity was used for all experiments, therefore equal amounts of protective capacity was assigned to all non-constraint stations. Processing times for all workstations were assumed to be lognormal. It has been shown in other studies (Atwater and Chakravorty, 1994; Muralidhar et al, 1992) that the lognormal distribution is a good representation of real-world processing times, especially manual assembly type of operations where the variability is relatively low. A mean processing time of 2 minutes at each station in the balanced line case was chosen arbitrarily. When protective capacity was added to the stations in the line, the processing time at the first station in the line was kept at 2 minutes. This station was therefore designated as the primary constraint station, and its position was kept constant. The beginning station was chosen as the constraint position, since a study by Kadipasaoglu et al (2000) showed that the lowest WIP and shortest flow time and waiting time were achieved with the constraint located at the first station in the line. For the addition of protective capacity it was further assumed that job content assignment was flexible and could be assigned in any amount to a station. No resource downtimes were modelled, but only normal processing time variability through the specification of the coefficient of variation for the processing times.

All the lines were modelled as terminating systems. An empty line is used as the starting condition, and the simulation is terminated after a run length of 48000 minutes, which relates to 100 production days if a production day consists of 8 production hours. Data is therefore collected by using independent replications of this terminating system. The 48000 minute simulation period was arbitrarily chosen to ensure a long enough time span for start-up conditions to disappear and flow time statistics to be collected.

All the ANOVA results for the different line configurations had to be checked for the normality and constant variance assumptions. Normal probability plots of the residuals as well as scatter plots of the residuals versus the predicted values were therefore drawn for all line configurations. These are displayed in Appendix E. In all the cases the normal probability plots do not resemble a straight line, and the residual scatter plots show that the residuals increase with larger predicted values. Both the normality and constant variance assumptions are therefore violated. The violation of the constant variance assumptions can be ascribed to the fact that lower levels of protective capacity cause more flow time variability. This was tested by calculating the standard deviation of the flow times at the different protective capacity levels. These calculations showed a higher standard deviation for flow times at lower protective capacity levels. ANOVAs based on the rank transformations of the flow time measures (refer to Appendix E) however give the same results as the ANOVAs based on the original non-ranked data. It is therefore safe to continue the analysis based on the ANOVAs for the original non-ranked data.

Each of the different line configurations is discussed separately in the following sections.

4.4.1 Six station infinite buffer line

In this line unlimited queue sizes were assigned to each station. This prevented blocking between stations, allowing inventory to flow freely in the system. A total work content of 12 minutes was distributed among the stations. For the balanced line with zero protective capacity at stations and using a maximum average processing time of 2 minutes at stations, this resulted in a six-station line. The flow diagram for the simulation model of the balanced line is displayed in Figure 4-16.

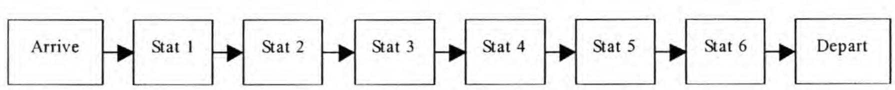

Figure 4-16. Flow diagram for simulation model of six station line

An arrival rate was chosen to ensure a 95% utilisation at the primary constraint station (station 1). Based on this 95% utilisation measure and the two minute service time at station 1 (the primary CCR), the arrival rate was calculated as 0.475 jobs per minute. Arrivals were modelled with the lognormal distribution using a coefficient of variation of 0.3. This represented a low arrival variability.

The effect of protective capacity on the mean flow time was studied at two levels for the coefficient of variation of the processing times: 0.3 and 1.0. The 0.3 level represented a low variability line, whereas the 1.0 level represented a high variability line. ANOVAs were developed for each of these two variability situations.

The preliminary statistical analysis involved determining the warm up period and the number of replications needed to draw valid statistical conclusions from the model. Again the truncated-replication approach suggested by Kelton et al (1998, pp.219-224) for terminating simulations was used. Since the infinite buffer experiments for the 6 station and 20 station lines were performed together, the 20 station line with a 1.0 coefficient of variation of the processing time and zero protective capacity was used to determine the warm-up period and the number of replications. This represented a worst-case scenario. To determine the length of the warm up period, a single run of 10 replications was made and a plot of the flow time performance measure values recorded against time was made. A visual inspection of the plot showed the values to stabilise and therefore the transient phase to end after more or less 10000 minutes. The length of the warm up period was therefore chosen as 10000 minutes. Simulations were therefore run for 10000 minutes after which statistics were cleared and data was collected over the next 38000 minutes.

The number of replications needed was computed using the confidence interval half-width formula from Kelton et al (1998, p.185) as explained in section 3.3.4. Using a 95% confidence level, it was determined that to obtain a 2 % precision level in the performance measures 50 replications had to be performed. For each level of protective capacity the number of stations required as well as the processing times at each non-constraint station also had to be determined. These values are shown in Table 4-20.

Table 4-20. Line configuration at different protective capacity levels for 6 station infinite buffer line

		Processing Times (min)								
Protective Capacity Level	Nr of Stations	Stat 1	Stat 2	Stat 3	Stat4	Stat5	Stat6	Stat 7	Stat 8	Stat 9
0%	6	2	2	2	2	2	2			
2%	7	2	1.96	1.96	1.96	1.96	1.96	0.2		
5.25%	7	2	1.895	1.895	1.895	1.895	1.895	0.525		
10.5%	7	2	1.79	1.79	1.79	1.79	1.79	1.05		
16%	7	2	1.68	1.68	1.68	1.68	1.68	1.6		
31.5%	9	2	1.37	1.37	1.37	1.37	1.37	1.37	1.37	0.41

The single factor ANOVA performed on the flow time performance measure for the six-station line with low variability (0.3 processing time coefficient of variation) is displayed in Table 4-21. The p-value is smaller than 0.05 whereas the F- value is quite large. It can therefore be concluded that there are significant differences in the means for the flow time values obtained at the different protective capacity levels.

Table 4-21 . ANOVA table: flow time output measures for 6-station infinite buffer line with 0.3 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	144436.0	1	144436.0	353194.0	0.00
PC Level	7847.9	5	1569.6	3838.1	0.00
Error	120.2	294	0.4		

Table 4-22. Tukey HSD comparisons of mean flow times for Protective Capacity level main effect (6-station infinite buffer line with 0.3 cv)

PC Level	Flow Time	1	2	3	4	5	6
31.5 Percent	16.08018	**					
16 Percent	17.80270		**				
10.5 Percent	18.93902			**			
5.25 Percent	22.00537				**		
2 Percent	25.89197					**	
0 Percent	30.93302						**

(means within a group that do not significantly differ are indicated with **)

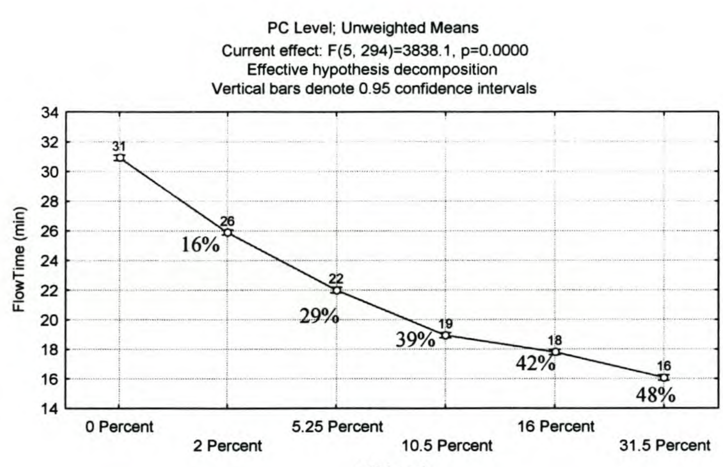

Figure 4-17. Effect of protective capacity at non-CCRs on mean flow time for 6-station infinite buffer line with 0.3 cv (percentages beneath graph indicate percentage improvement over flow time for zero protective capacity line)

To further investigate which means differ significantly, graphical analysis and Tukey HSD comparisons at a 95% significance level were used. The graph of the mean flow time measure at the different levels of protective capacity for a 0.3 processing time coefficient of variation is displayed in Figure 4-17, and the Tukey comparison of the means are displayed in Table 4-22. The Tukey comparison shows that all means differ significantly. The graph in Figure 4-17 indicates that the flow time measure improves with increasing levels of protective capacity. The graph also displays the percentage improvement in flow time at a specific protective capacity level from the zero protective capacity level. This illustrates that increases in protective capacity has diminishing returns. What is interesting from these results is that even a low protective capacity level of only 2% can cause a flow time improvement as large as 16% from that of the balanced line. When looking at the number of stations at the different protective capacity levels in Table 4-20, it can be seen that a protective capacity level of up to around 16% causes only one additional station for this specific line. In turn a 16% protective capacity level causes a 42% [100*(31-18)/31] reduction in flow time from the balanced six station line, which is quite significant. The ANOVA for the six station line with high variability (1.0 processing time coefficient of variation) is displayed in Table 4-23. The F and p-values also indicate a significant difference in the means.

Table 4-23. ANOVA table: flow time output measures for 6-station infinite buffer line with 1.0 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	3586108	1	3586108	21123.07	0.00
PC Level	775655	5	155131	913.76	0.00
Error	49913	294	170		

The graph of the mean flow time measure at the different levels of protective capacity for a 1.0 processing time coefficient of variation is displayed in Figure 4-18, and the Tukey comparison of the means are displayed in Table 4-24. The Tukey comparison confirms that all means differ significantly. Increasing the protective capacity from 0% to 2% causes a 23% decrease in flow time, whereas a protective capacity increase from 0% to 16% results in a 64% reduction in flow time. Low levels of protective capacity in this case are therefore also sufficient to cause significant improvement in the mean flow time. The improvement in flow time for protective capacity increases is however greater in this high variability line than for the low variability line.

Table 4-24. Tukey HSD comparisons of mean flow times for Protective Capacity level main effect (6-station infinite buffer line with 1.0 cv)

PC Level	Flow Time	1	2	3	4	5	6
31.5 Percent	48.6073	**					
16 Percent	69.8221		**				
10.5 Percent	79.4622			**			
5.25 Percent	110.5672				**		
2 Percent	150.6776					**	
0 Percent	196.8614						**

(means within a group that do not significantly differ are indicated with **)

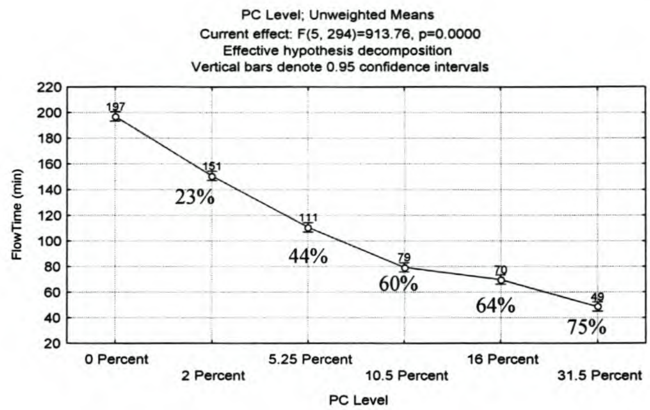

Figure 4-18. Effect of protective capacity at non-CCRs on mean flow time for 6-station infinite buffer line with 1.0 cv (percentages beneath graph indicate percentage improvement over flow time for zero protective capacity line)

4.4.2 Six station finite buffer line

For the finite buffer experiment the same simulation model was used as for the infinite buffer six-station line (refer to Figure 4-16), except that the buffer or queue space in front of stations was limited to one job. If a station finished processing a job and the downstream queue or buffer already had a job in queue, the upstream station is blocked and cannot start processing the next job until the downstream buffer is emptied.

Again the total work content that was distributed among the stations was 12 minutes. The balanced or zero protective capacity line consisted of 6 stations with an average processing time of 2 minute each, whereas for the unbalanced line configurations the constraint was kept at the first station with a processing time of 2 minutes. The required number of stations as well as the processing times at each non-constraint station for each level of protective capacity are displayed in Table 4-25.

Table 4-25. Line configuration at different protective capacity levels for 6 station finite buffer line

		Processing Times (min)								
Protective Capacity Level	Nr of Stations	Stat 1	Stat 2	Stat 3	Stat4	Stat5	Stat6	Stat 7	Stat 8	Stat 9
0%	6	2	2	2	2	2	2			
2%	7	2	1.96	1.96	1.96	1.96	1.96	0.2		
5%	7	2	1.9	1.9	1.9	1.9	1.9	0.5		
10%	7	2	1.8	1.8	1.8	1.8	1.8	1		
15%	7	2	1.7	1.7	1.7	1.7	1.7	1.5		
30%	9	2	1.4	1.4	1.4	1.4	1.4	1.4	1.4	0.2

For the finite buffer situation only the low processing time variability situation with a coefficient of variation of 0.3 was investigated, since the finite buffers itself also contributed to the variability. The required job arrival rate could not be determined directly from a pre-specified required utilisation level at the constraint station, since the finite buffers reduces the maximum obtainable output level for the line. In order therefore to ensure that not too high a utilisation level at a specific station is reached and that steady state conditions can be achieved, preliminary simulation runs were used to determine a suitable arrival rate. Different arrival rates were simulated using the 6 station balanced line with zero protective capacity until plots of the flow time values showed that steady state conditions could be reached and a maximum utilisation of about 96% was obtained at a single station. The arrival rate in this case that was also used for the experiment was 0.4425 jobs per minute. Arrivals were modelled with the lognormal distribution using a coefficient of variation of 0.3. This represented a low arrival variability.

The warm-up period was determined using the truncated-replication approach suggested by Kelton et al (1998, pp.219-224) for terminating simulations. A single run of 10 replications was made with the six-station zero protective capacity line configuration and the flow time performance measure values were recorded. These values were plotted against time, and a visual inspection of the plot showed the values to stabilise and therefore the transient phase to end after more or less 1000 minutes. The length of the warm up period was therefore chosen as 1000 minutes, which meant that simulations were run for 1000 minutes after which statistics were cleared and data was collected over the next 47000 minutes. Using the confidence interval half-width formula from Kelton et al (1998, p.185) as explained in section 3.3.4, it was determined that to obtain a 2 % precision level in the performance measures at a 95% confidence level, 70 replications had to be performed.

The ANOVA results are presented in Table 4-26. These results indicate that there are significant differences in the mean flow times caused by the different protective capacity levels.

Table 4-26. ANOVA table: flow time output measures for 6-station finite buffer line with 0.3 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	138212.2	1	138212.2	125644.6	0.00
PC Level	7660.4	5	1532.1	1392.8	0.00
Error	455.4	414	1.1		

Table 4-27. Tukey HSD comparisons of mean flow times for Protective Capacity level main effect (6-station finite buffer line with 0.3 cv)

PC Level	Flow Time	1	2	3	4	5	6
30 Percent	14.00691	**					
15 Percent	15.03366		**				
10 Percent	15.71845			**			
5 Percent	17.42307				**		
2 Percent	19.98040					**	
0 Percent	26.68031						**

(means within a group that do not significantly differ are indicated with **)

Figure 4-19 displays the graph of the mean flow time values at the different protective capacity levels. The Tukey HSD comparisons in Table 4-27 confirm that all the mean flow time values differ significantly. For the finite buffer case the increases in protective capacity have much the same effect on the mean flow time as with the infinite buffer case. The flow time decrease with increasing protective capacity levels, but with diminishing returns. Increasing the protective capacity from 0% to 2% decreases the flow time with 26%, while a protective capacity increase from 0% to 15% results in a 44% decrease in flow time with only one additional station needed above the six station balanced capacity line.

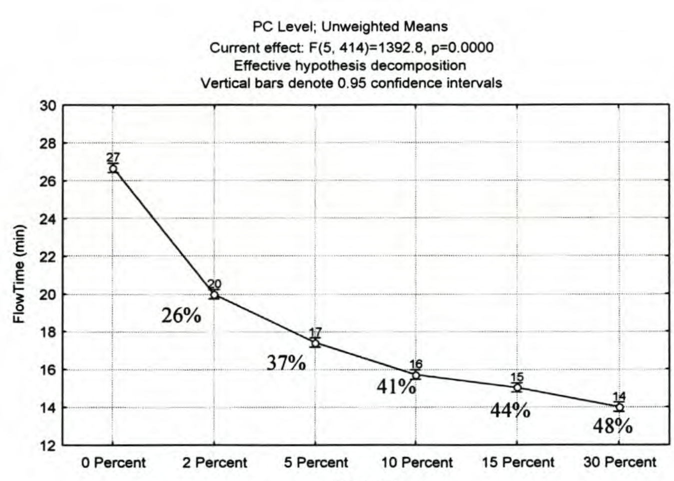

Figure 4-19. Effect of protective capacity at non-CCRs on mean flow time for 6-station finite buffer line with 0.3 cv (percentages beneath graph indicate percentage improvement over flow time for zero protective capacity line)

The finite buffer lines are characterised by a phenomenon known as blocking. When two workstations function in series and the buffer or queue space between the stations is limited, it could happen that the downstream workstation is busy with service and the queue in front of the second workstation is full. When the upstream workstation finishes its service and such a situation is taking place, the part at the upstream workstation can't leave the workstation until the downstream workstation finishes its service. The upstream workstation is therefore temporarily blocked from further processing. Blocking has the effect of reducing the maximum attainable output of a line (refer to Du Preez 1984, pp. B.22 – B.30). Although not officially investigated in this study, the simulation results showed that adding protective capacity not only reduces the flow time, but also reduces the blocking percentages at stations in the line. This means that protective capacity actually also increases the maximum attainable output for finite buffer serial lines. Protective capacity also reduces the buffer sizes required between stations to reach the maximum output for the line.

4.4.3 Twenty station infinite buffer line

The twenty station infinite buffer line was modelled in exactly the same way as the six station infinite buffer line, except that the total work content in this case was 40 minutes. For the balanced line with a 2 minute processing time at each station this resulted in a 20 station line. The flow diagram for the simulation model of the 20 station balanced capacity line is displayed in Figure 4-20.

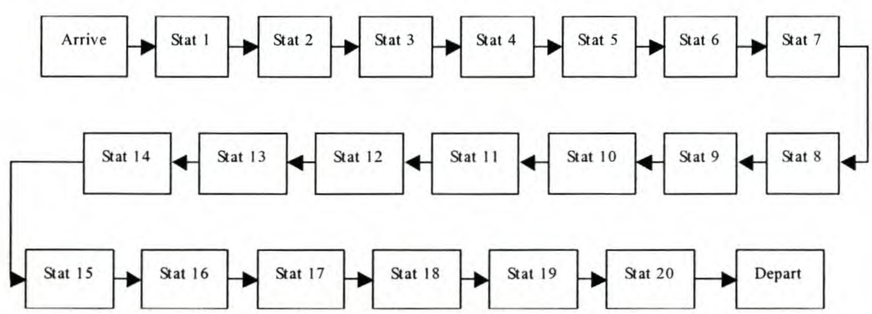

Figure 4-20. Flow diagram for 20 station balanced line simulation model

The same arrival rate, arrival and processing time distributions, and protective capacity levels were used as for the 6 station infinite buffer line. The effect of protective capacity levels on flow time was also investigated at two processing times coefficient of variability levels: 0.3 (low variability) and 1.0 (high variability). The same preliminary statistical analysis was performed as for the six station infinite buffer line (the 20 station line was in fact used for the preliminary statistical analysis in the 6 station line). The warm-up period in this case was therefore also 10000 minutes, and the number of simulation replications was 50. The required number of stations as well as the processing times at each non-constraint station for each level of protective capacity are displayed in Table 4-28.

Table 4-28. Line configuration at different protective capacity levels for 20 station infinite buffer line

		Processing Times (min)										
Protective Capacity Level	Nr of Stations	Stat 1	Stat 2 to 20	Stat 21	Stat 22	Stat 23	Stat 24	Stat 25	Stat 26	Stat 27	Stat 28	Stat 29
0%	20	2	2									
2%	21	2	1.96	0.76								
5.25%	22	2	1.895	1.895	0.1							
10.5%	23	2	1.79	1.79	1.79	0.41						
16%	24	2	1.68	1.68	1.68	1.68	1.04					
31.5%	29	2	1.37	1.37	1.37	1.37	1.37	1.37	1.37	1.37	1.37	1.01

The ANOVA results for the low variability situation (0.3 processing time coefficient of variation) are displayed in Table 4-29. This table indicates that there are significant differences in the mean flow time values at the different protective capacity levels. The Tukey HSD comparisons in Table 4-30 confirm that all the mean flow time values differ significantly.

Table 4-29. ANOVA table: flow time output measures for 20-station infinite buffer line with 0.3 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	1481323	1	1481323	557711.6	0.00
PC Level	111886	5	22377	8424.9	0.00
Error	781	294	3		

Table 4-30. Tukey HSD comparisons of mean flow times for Protective Capacity level main effect (20-station infinite buffer line with 0.3 cv)

PC Level	Flow Time	1	2	3	4	5	6
31.5 Percent	47.9044	**					
16 Percent	54.0938		**				
10.5 Percent	59.7369			**			
5.25 Percent	70.6953				**		
2 Percent	85.0636					**	
0 Percent	104.1204						**

(means within a group that do not significantly differ are indicated with **)

The graph of the mean flow time measure at the different levels of protective capacity for a 0.3 processing time coefficient of variation is displayed in Figure 4-21. Increasing the protective capacity from 0% to 2% causes an 18% decrease in the mean flow time. From Table 4-28 it can be seen that a 2% protective capacity level only requires one additional station above the original 20 stations for the balanced capacity line. A 5.25% protective capacity level requires 2 additional stations above the 0%

protective capacity line, whereas this level causes a 32% drop in flow time from the zero protective capacity level. The 10.5 % protective capacity level requires 3 additional stations with a 42% drop in flow time, whereas the 16% level requires 4 additional stations above the zero protective capacity line with a 48% drop in flow time. Increasing protective capacity levels again display diminishing returns.

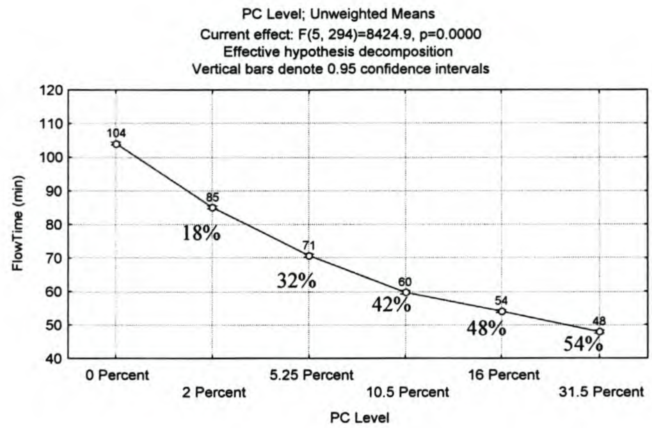

Figure 4-21. Effect of protective capacity at non-CCRs on mean flow time for 20-station infinite buffer line with 0.3 cv (percentages beneath graph indicate percentage improvement over flow time for zero protective capacity line)

The ANOVA results for the 1.0 processing time coefficient of variation are presented in Table 4-31. There is a significant difference in the mean flow time values, and the Tukey HSD comparisons in Table 4-32 indicate that all the means differ significantly.

The graph for the main effect of the protective capacity level on the mean flow time is displayed in Figure 4-22. A 2% protective capacity level causes a 25% reduction in the flow time from the zero capacity line, while only requiring one additional station. When comparing the flow time reduction percentages for the 1.0 coefficient of variation case to the 0.3 coefficient of variation case, it appears as if increasing protective capacity levels causes larger reductions in flow time for the high variability situation than for the low variability situation.

Table 4-31. ANOVA table: flow time output measures for 20-station infinite buffer line with 1.0 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	43154519	1	43154519	68195.28	0.00
PC Level	12346272	5	2469254	3902.06	0.00
Error	186046	294	633		

Table 4-32. Tukey HSD comparisons of mean flow times for Protective Capacity level main effect (20-station infinite buffer line with 1.0 cv)

PC Level	Flow Time	1	2	3	4	5	6
31.5 Percent	132.0270	**					
16 Percent	209.7418		**				
10.5 Percent	272.4615			**			
5.25 Percent	395.7812				**		
2 Percent	541.1470					**	
0 Percent	724.4825						**

(means within a group that do not significantly differ are indicated with **)

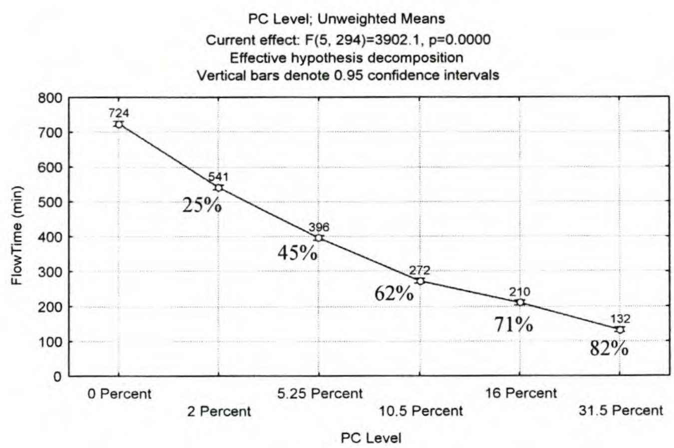

Figure 4-22. Effect of protective capacity at non-CCRs on mean flow time for 20-station infinite buffer line with 1.0 cv (percentages beneath graph indicate percentage improvement over flow time for zero protective capacity line)

4.4.4 Twenty station finite buffer line

The simulation model for the finite buffer line was similar to the infinite buffer line model displayed in Figure 4-20, except the buffer capacity between stations was limited to one. Again a total work content of 40 minutes was distributed among the stations. The balanced or zero protective capacity line consisted of 20 stations with an average processing time of 2 minutes each, whereas for the unbalanced line configurations the constraint was kept at the first station with a processing time of 2 minutes. The required number of stations as well as the processing times at each non-constraint station for each level of protective capacity are displayed in Table 4-33.

Table 4-33. Line configuration at different protective capacity levels for 20 station finite buffer line

		Processing Times (min)										
Protective Capacity Level	Nr of Stations	Stat 1	Stat 2 to 20	Stat 21	Stat 22	Stat 23	Stat 24	Stat 25	Stat 26	Stat 27	Stat 28	Stat 29
0%	20	2	2									
2%	21	2	1.96	0.76								
5%	22	2	1.9	1.9								
10%	23	2	1.8	1.8	1.8	0.2						
15%	24	2	1.7	1.7	1.7	1.7	0.6					
30%	29	2	1.4	1.4	1.4	1.4	1.4	1.4	1.4	1.4	1.4	0.2

Only the low processing time variability situation with a coefficient of variation of 0.3 was investigated. As with the six station finite buffer line, preliminary simulation runs were used to determine a suitable arrival rate. Different arrival rates were simulated using the 20 station balanced line with zero protective capacity until plots of the flow time values showed that steady state conditions could be reached and a maximum utilisation of about 96.5% was obtained at a single station. The arrival rate in this case was 0.435 jobs per minute. Arrivals were modelled with the lognormal distribution using a coefficient of variation of 0.3. This represented a low arrival variability.

Using the same procedures as with the previous simulation experiments, the warm-up period was determined to be 2000 minutes, and the number of replications needed was determined at 70 (for a 2 % precision level in the performance measures at a 95% confidence level). The ANOVA results are presented in Table 4-34. These results indicate that there are significant differences in the mean flow times caused by the different protective capacity levels. The Tukey HSD comparisons in Table 4-35 confirm that all the mean flow time values differ significantly.

Table 4-34. ANOVA table: flow time output measures for 20-station finite buffer line with 0.3 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	1349949	1	1349949	81485.81	0.00
PC Level	51160	5	10232	617.62	0.00
Error	6859	414	17		

Figure 4-23 displays the graph of the mean flow time values at the different protective capacity levels. From the graph it can be seen that the flow time decreases with increasing protective capacity levels, but with diminishing returns. Increasing the protective capacity from 0% to 5% decreases the flow time with 29% with only one additional station needed above the twenty station balanced capacity line.

Table 4-35. Tukey HSD comparisons of mean flow times for Protective Capacity level main effect (20-station finite buffer line with 0.3 cv)

PC Level	Flow Time	1	2	3	4	5	6
30 Percent	44.94230	**					
15 Percent	48.76605		**				
10 Percent	51.37492			**			
5 Percent	55.65767				**		
2 Percent	60.71913					**	
0 Percent	78.70155						**

(means within a group that do not significantly differ are indicated with **)

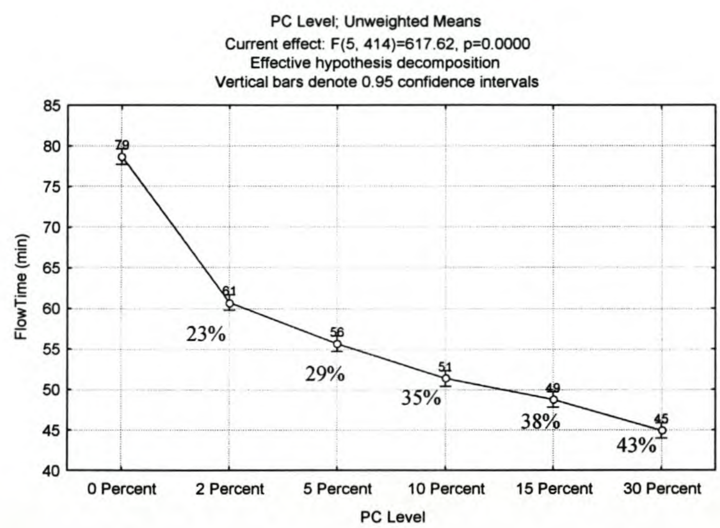

Figure 4-23. Effect of protective capacity at non-CCRs on mean flow time for 20-station finite buffer line with 0.3 cv (percentages beneath graph indicate percentage improvement over flow time for zero protective capacity line)

4.4.5 Answers to research questions for the assembly line experiment

To answer the research questions posed in section 4.2, the following conclusions can be summarised for the assembly line experiment:

a) *Research question 4: In an assembly line, what is the effect of the amount of protective capacity at the non-constraint stations on mean flow time?*

Increasing protective capacity causes significant reductions in flow time, however with diminishing returns. In all of the six line configurations investigated in the previous section the graphs plotting the flow time performance measure against the protective capacity level showed significant reductions in flow time with increasing protective capacity levels. In the infinite buffer lines this can be explained by the fact that protective capacity moves the jobs faster through the line after being served at the constraint (the constraint was located at the beginning of the line). Less work-in-process build-up therefore occurs at the non-constraint stations. Because of the relationship between lead times or flow times and work-in-process the lower work-in-process levels lead to faster flow times. The additional improvements however start to decrease with higher protective capacity levels. This happens because the WIP levels in the line start to reach the minimum WIP levels that can be reached because of the pacing effect of the constraint resource located at the beginning of the line that releases jobs to the rest of the line. For the finite buffer capacity lines the reduction in flow time can be explained by the fact that the higher protective capacity levels reduces the effect of blocking in the line. Because of the lower blocking, resources don't have to wait that long for downstream resources to complete their task before a job can be released to a downstream station. The shorter wait times therefore contribute to the faster flow times.

b) *Research question 5: In an assembly line, does the same relationship between the amount of protective capacity at the non-constraint stations and the mean flow time for jobs through the line hold for different line lengths and different levels of processing time variation at stations?*

As was expected, the same relationship between protective capacity levels and flow time were observed for both short lines and long lines. For example when comparing the flow time performance measure graphs for both the 6 and 20 stations lines under low processing time variability (Figure 4-17 and Figure 4-21) it can be seen that more or less the same order of improvements in flow time are obtained with higher protective capacity levels. The 6-station line for example has a 16% improvement in flow time at the 2% protective capacity level and a 29% improvement at the 5% protective capacity level. The 20-station line has an 18% improvement in flow time at the 2% protective capacity level and a 32% improvement at the 5% level. When comparing the 6 and 20-station lines under high variability (Figure 4-18 and Figure 4-22) the improvements in flow time for the 6-station line are 23% and 44% respectively at the 2% and 5% protective capacity levels, whereas for the 20-station line it is 25% and 45% respectively. Higher protective capacity levels also cause significant improvements in flow time for both low and high variability lines. For higher variability lines however the improvement in flow times is larger than for low variability lines. For example when comparing the low and high variability situations of the 6-station line, the flow time improvements for the low variability line are 16% and 29% respectively for the 2% and 5% protective capacity levels (Figure 4-17), whereas for the high variability line the improvements are 23% and 44% respectively (Figure 4-18). Since the higher variability lines have more initial WIP build-up, there is more room for WIP reduction and therefore flow time reduction. Higher protective capacity levels also serve to dissipate some of the variability accumulation that occurs under high processing time variability conditions.

c) *Research question 6: In an assembly line, does the same relationship between the amount of protective capacity at non-constraint stations and the mean flow time hold for lines with limited buffer capacity and lines with unlimited buffer capacity between stations?*

In lines with limited buffer capacity between stations, the same order of improvements in flow time for higher protective capacity levels are observed as with infinite buffer capacity lines. For example compare the 20-station finite buffer line in Figure 4-23 with the 20-station infinite buffer line under low processing time variability in Figure 4-22. For the finite buffer line the flow time improvements are 23% and 29% respectively for the 2% and 5% protective capacity levels, whereas the improvements are 18% and 32% respectively for the infinite buffer line. As stated before the flow time improvements in the finite buffer line are mainly due to less blocking and therefore less waiting time, whereas the improvements in the infinite buffer line are due to less WIP build-up. When comparing the 6-station finite buffer line in Figure 4-19 with the 6-station low processing time variability line in Figure 4-17, it can be seen that the finite buffer line achieves higher flow time improvements at the low protective capacity levels. The flow time improvements at the higher protective capacity levels are however very similar for the two line configurations.

d) *Research question 7: In an assembly line, what levels of protective capacity can provide significant improvements in flow time over a balanced capacity line without resulting in too many stations?*

The results suggest that low protective capacity levels at non-constraint stations of between 2%-5% are sufficient to cause significant reductions in flow time without resulting in too many additional stations above the balanced or zero capacity configuration. This is true for both finite and infinite buffer lines. At the 2% protective capacity level lines with up to 50 stations in the balanced capacity case will only result in one additional station, whereas for a 5% level lines with up to 20 stations in the balanced capacity case will only result in one additional station. Flow time reductions in the order of 30% over a balanced or zero protective capacity line configuration is however possible for a 5% protective capacity level in lines with low processing time variability (refer for example Figure 4-17 and Figure 4-21), whereas flow time improvements in the order of 45% are attainable in high variability lines (refer Figure 4-18 and Figure 4-22). These results are for a level distribution of protective capacity (an equal amount of protective capacity at all non-constraint stations).

Stellenbosch University http://scholar.sun.ac.za

Chapter 5

Summary, Conclusions and Future Research

5.1 Summary and Conclusions

This research study investigated the design of time buffers and protective capacity in discrete flow production systems. Time buffers and protective capacity are the two main protection mechanisms in Theory of Constraints controlled production systems to protect system output or throughput while keeping flow times and WIP low. The first part of the research study presented an analytical procedure to be used for estimating the lengths of the time buffers needed in production systems controlled with the drum-buffer-rope control system of the Theory of Constraints. This analytical procedure is based on an open queuing network analysis of a production flow network.

Such an analytical procedure is valuable to quickly determine the lengths of the needed time buffers. Compared to simulation, the analytical model is inexpensive with respect to time and money and is useful for quick initial estimates in order to compare alternatives and to gain insight into the network studied. The analytical model enables more realistic modelling of production networks by enabling the modelling of the influence of transfer batches on time buffer lengths, as well as variability caused by normal processing time variations and downtimes of resources. Evaluations of the analytical procedure on a simulation model of an actual flow shop under a drum-buffer-rope scheduling and job release methodology, showed that the time buffer length estimations calculated with the procedure are sufficiently accurate to be used in practice. Although the average flow time and flow time standard deviation estimations of the queuing analysis procedure are not always that accurate in more realistic network models, the relatively low percentage late orders and average lateness measures from the simulation results indicate that the final calculated time buffer sizes are still useful estimators of the amount of protection necessary. It should be remembered that no buffer management techniques were implemented in the simulation. In practice these techniques should greatly improve the delivery performance. The simulation model used in the study was also a very close representation of a real-life plant's operations simulated with material release schedules and constraint schedules, therefore larger differences between the analytical model and the simulation model's average flow time and flow time standard deviation calculations were expected. Inaccuracies in the average flow time and flow time standard deviation estimations seem to be offset by the safety level (z-value) incorporated in the final

Stellenbosch University http://scholar.sun.ac.za

time buffer lengths. For adequate protection the maximum z value should therefore be used. The time buffer lengths can always later be fine tuned with simulation experiments or through the application of buffer management techniques after the line is in operation.

The time buffer estimation procedure presented in this study can therefore be quite useful in the design phase of production flow lines in order to compare different line configurations with respect to the amount of protection in terms of time (which results in a higher or lower WIP according to Little's formula) required by the specific line design. This is because of the interaction that exists between time buffers and protective capacity. With more protective capacity, less WIP protection is needed which means that the required time buffer lengths that need to be implemented can be shorter. The resulting flow times for parts are therefore also shorter, which can be used as a competitive advantage. By modelling different protective capacity scenarios with the analytical model, the time buffer estimation procedure can be used to calculate the corresponding time buffer requirements for a specific protective capacity scenario. In this way different protective capacity line designs can be compared and evaluated. It can also be used to estimate the size of the time buffers to be used when putting the line into operation. It is however still essential that buffer management techniques be applied while operating the line to monitor the time buffers and fine-tune their lengths, since any analytical procedure and even simulation models are still only abstractions of reality and can never be 100% accurate. The time buffer estimation procedure presented in this study however enables a manager to build the initial time buffer on a more scientific estimate than the approaches currently described in the Theory of Constraints literature.

The second part of the research study investigated the effect of protective capacity levels on the performance of both flow shops and assembly type of flow line configurations. There is a relationship between time buffer design and protective capacity design, since (as illustrated by the experimental research in this study) protective capacity influences the flow times for jobs through the system, whereas the mean flow times in turn determine the lengths of the time buffers needed for protection of order delivery. The experimental results on the use of protective capacity in flow shops and infinite buffer assembly lines can therefore be used by managers together with the analytical procedure for calculating time buffer lengths when designing new production configurations. Using the developed time buffer analysis tool, the design rules suggested by the experimental results can be evaluated on different production configurations to determine their effect on the flow times and therefore the buffer protection required by the specific system configuration.

The first experiment for investigating the use of protective capacity focused on a flow shop where the number of stations is fixed and there is unlimited queue space in front of stations. One of the observations from these experimental results is that flow shop performance in the form of flow time and line output is not that much influenced by low protective capacity levels at the secondary constraint resource. Adding higher levels of protective capacity at a single secondary constraint resource also will not do much to improve flow time and line output. These performance measures seem to be more influenced by the average level of protective capacity available at all non-constraint resources. These results could be important for production managers in that it suggest that when only a few secondary constraints are present, not too much effort need to be wasted on planning and scheduling these secondary constraint resources in order to obtain more capacity at these resources. This finding supports the Theory of Constraints philosophy that most of the effort should be concentrated on the primary constraint. It should however be remembered that this study only considered a simple serial flow line with one product type. In more complex production systems such as a job shop with multiple part types and multiple routings the capacity of the secondary constraint could have a larger influence on performance measures such as the job flow time and the total line output. In such situations short term product mix changes could necessitate more careful planning and scheduling of the secondary constraint. Future research studies could perhaps look at the effect of the protective capacity at the secondary constraint under more complex production systems such as multi-part job shops. This study however showed that low protective capacity levels at a single station can significantly reduce the bottleneck probability for the primary constraint resource when it is located before and relatively close or near to the primary constraint in the process flow, or after but relatively far from the primary constraint. An after-far secondary constraint location also causes higher job flow times, and should therefore be avoided. The reason for the worse primary constraint bottleneck probability performance of the before-near secondary constraint location is that the secondary constraint in this case actually paces the flow of jobs to the primary constraint. This prevents significant WIP build-up in front of the primary constraint. More WIP build-up takes place in front of the secondary constraint. When the secondary constraint is located very close to the primary constraint and the resources with high breakdown levels are located before the secondary constraint in the process flow, such a secondary constraint will however not cause much starvation at the primary constraint and therefore throughput or output to be lost. This is because starvation and lost throughput is more caused by resource breakdowns or downtimes than by normal low processing time variability. In this situation the resources with high breakdowns are however located before the secondary constraint and not between the secondary constraint and the primary constraint. The WIP at the secondary constraint therefore also serves as a protection buffer for the primary constraint in the case of breakdowns at stations earlier in the process flow. If the secondary constraint however experiences significant breakdowns, a before-near secondary constraint location will definitely cause more starvation of the primary constraint and therefore throughput to be lost. The reason for the worse primary constraint bottleneck probability performance of the after-far secondary constraint location is because of the accumulation of variability in a line. A station located near the end of the process flow therefore experiences much more part arrival variability as stations earlier in the process flow. With low protective capacity levels at such a secondary constraint location the higher arrival variability causes more WIP build-up at the secondary constraint, which reduces the bottleneck probability for the primary constraint.

The experimental results for the 15-station flow shop also showed that a secondary constraint located after but near the primary constraint in the process flow do not have much of an effect on the flow time, line output as well as the bottleneck probability for the primary constraint. In this study an infinite buffer capacity was however allowed between stations. Because of this no blocking could occur. In lines with limited buffer capacity between stations, the blocking effect would definitely cause an after-near secondary constraint location to have more of a detrimental effect on the line output and bottleneck probability performance measures.

The results of this study further suggest that for low variability flow shops it seems as if low average levels of protective capacity at non-constraint resources are sufficient to ensure that the maximum designed output level as determined by the utilisation of the primary constraint resource is obtained. This has positive implications for production managers, because it indicates that the protective capacity requirements need not be too expensive when considering the cost of capacity. This study only considered a minimum average protective capacity level of 10.5% at the non-constraint resources. Further experimentation with lower levels is needed to determine at what level of protective capacity a significant drop in output occurs for the low variability line configuration. An average protective capacity level of around 10% at non-constraint resources would however be a good representation of low protective capacity levels found in actual flow shops.

For low variability flow shops the decision of what level of protective capacity to use when designing the shop therefore depends more on the strategic decision of what average production flow time is acceptable, and the required stability in the position of the primary constraint resource. The cost of the additional protective capacity should be weighed against the benefits of shorter flow times and more manageable shops. Although faster flow times will not improve the current output or throughput for the shop, it will have other benefits such as lower WIP (and the associated benefits of lower WIP such as less obsolescence, scrap, re-work, carrying costs), shorter planning horizons, and better due-date performance, which relates to lower cost, higher quality and rapid delivery. The shorter manufacturing flow or lead time can also be used as a competitive advantage for increasing future product demand and therefore throughput, or protecting current demand. Care should however be taken to ensure that a higher demand will not overload the primary constraint resource. If the primary constraint becomes overloaded the cost implications of increasing its capacity should be considered.

One of the original contributions of the study was the investigation of the effect that protective capacity has on the bottleneck probability of the primary constraint. The study demonstrated that short term bottlenecks that move around are an inevitable result of the variability present in production systems, and by adding protective capacity the bottleneck probability of the primary CCR can be improved, thereby reducing the bottleneck shiftiness problem. When designing lines the impact of protective capacity on the probability of the primary constraint resource of remaining the only long-term constraint should therefore also be considered. The experimental results suggest that for low variability situations an average protective capacity level at non-constraint resources of around 20% (where protective capacity is defined relative to the primary constraint's capacity) would be needed to ensure a relatively stable primary constraint resource location with a probability of around 70%. Very high protective capacity levels of around 50% would only result in a primary constraint bottleneck probability of around 80%. A primary constraint bottleneck probability between 60% to 70% therefore appears to be a sufficient target to obtain a balance between protective capacity levels and shop manageability for low variability situations. The relatively low primary constraint bottleneck probability also indicates that it is inevitable that other resources besides the primary constraint will also become short term constraints some of the time. The so-called "wandering bottleneck" phenomenon observed in some plants could therefore be ascribed to low protective capacity levels. Care should therefore be taken to simply declare a new primary constraint resource based on WIP observations on the shop floor. For scheduling and throughput planning purposes the primary constraint should instead be identified with more careful output capability calculations. In this study a more accurate measure was developed for identifying short-term bottlenecks. This measure is based on the flow time for parts at a workstation and not just the amount of WIP in front of a station. This measure could be helpful in practice for more accurately identifying the primary CCR.

For shops with high variability due to station downtimes the experimental results indicate that quite high levels of protective capacity are needed in order to attain the maximum global output for the shop. In this experimental study protective capacity levels of around 50% were needed in order to reach the maximum output for a shop with high variability resources (10% downtimes characterised by long breakdown times). The chosen level of protective capacity therefore has a significant effect on the final output to be achieved by the shop under high variability conditions. Higher protective capacity levels also cause quite substantial improvements in the mean flow time, as well as improving the bottleneck probability for the primary constraint. For the high variability shop in this experiment an average protective capacity level of more or less 20 % at non-constraints was needed to ensure a bottleneck probability of around 50%. Quite high protective capacity levels are therefore also needed to ensure a more stable and predictable primary constraint. A limitation of this study was that it did not consider the effect of different buffer lengths on the line performance measures (flow time, output and bottleneck probability). For the high resource breakdown situation quite high protective capacity levels were needed to obtain the designed output for the line or to ensure a higher primary constraint bottleneck probability. The flow time measures resulting from these experiments are an indication of the minimum time buffer protection that would be required for that specific line configuration. There is however an inverse relationship between time buffers and protective capacity. By using longer time buffers for protection, lower protective capacity levels would be required to obtain the designed output for the line. The study also showed that by increasing protective capacity levels the bottleneck probability measure can be improved. An increase in the time buffer size however would also be able to increase the primary constraint bottleneck probability.

The second experiment investigating the use of protective capacity focused on an assembly type of flow line where a total work content is distributed among the stations in the line. The number of stations therefore varies according to the level of protective capacity. The results for the assembly line experiment showed that an unbalanced line configuration where less work is assigned to the non-constraint stations than to the primary constraint station (but non-constraint stations have an equal work content) can lead to significant reductions in the mean flow time while maintaining the same line output, without resulting in too many additional stations. Higher protective capacity levels also reduce the flow time variability, which is important for better delivery performance. These results hold for both short and longer lines, as well as lines with limited queue space between stations and lines with unlimited queue space between stations. The results also showed that increasing protective capacity levels have diminishing returns on mean flow time. Low protective capacity levels in the range of 2% to 5% are however sufficient to cause substantial improvements in flow time (30% to 40% reductions) without resulting in too many additional stations in the line. These flow time improvements can be realised in both low and high variability lines, although larger improvements are possible for high variability lines.

Another observation made when conducting the assembly line experiments was that higher protective capacity levels not only improve the flow time, but also reduce the blocking factor in limited buffer capacity serial lines. In limited buffer capacity lines blocking has the effect of reducing the maximum output level to be obtained from a line. The final output level obtained is therefore far less than the throughput rate of the constraint station. Although not officially investigated in this study, initial experimentation with the 6 and 20 station limited buffer capacity lines showed that the higher protective capacity levels reduced the blocking percentages for the stations in the line. This means that the maximum output level to be obtained from the line can be increased, and with enough protective capacity the output rate can be brought very close to the maximum throughput rate of the primary CCR. Protective capacity in these types of lines therefore not only reduces the flow times, but also increases the output or throughput.

These results are important for line managers designing assembly types of flow lines where job content is distributed among stations. Traditional literature and practice treat assembly line design in a balanced way where the objective is to minimise station idle time and the number of stations by assigning an equal job content to all stations. The experimental results however indicate that reduced flow times and flow time variability is possible by assigning equal amounts of protective capacity at non-constraint stations, without resulting in too many additional stations. Very low protective capacity levels (2% to 5%) are sufficient to obtain substantial reductions in flow time, which means that the additional cost caused by the extra capacity (because of more stations) could be quite low. These line performance improvements in turn lead to less WIP and better due date delivery performance. Again when making the decision of whether to add protective capacity to the line and what level, the benefits of faster flow times should be weighed against the increased cost (if any) of a larger amount of stations. Costs that could result from extra stations are when additional equipment or machinery are needed or additional workers need to be hired, or where conveyors are used and additional conveyor length is needed. These costs should be compared to the advantages of faster flow times, such as less WIP (and the accompanying advantages of lower WIP as discussed in section 1.2) and the competitive advantage of shorter production lead times that can lead to improved future throughput or the protection of current throughput.

The contributions made by this study to the production component of the Theory of Constraints are therefore an analytical procedure for more accurate estimates of the time buffer lengths, as well as an expansion of the knowledge surrounding the role and efficient design of protective capacity in discrete flow production systems. This could help production managers in practice by providing general guidelines for the more efficient design of production systems by using protective capacity. The analytical model provides a tool that can be used for design purposes in practice, and also in further studies investigating the efficient design of time buffers. The experimental studies on protective capacity design showed that generally low levels of protective capacity are sufficient to provide the necessary throughput protection while improving performance measures such as flow time. This could serve to better promote the use of protective capacity in practice, because many production managers still believe that the cost of such protective capacity is greater than its benefits.

5.2 Limitations and Further Research

Based on the assumptions, limitations and findings of this research, the following recommendations for further research can be made:

- With respect to the time buffer estimation procedure, further research could focus on the use of other queuing network analysis approaches or GI/G/m queuing formulas in the time buffer calculation procedure, especially procedures that more accurately estimate the amount of variance in the flow times. The accuracy of the time buffer estimation procedure under different line utilisation levels caused by different demand levels can also be investigated. The time buffer estimation procedure can also be evaluated on other shop configurations such as job shops, or production networks with a larger product mix. Because the MTTR value actually determines the danger of the constraint being starved during a resource breakdown, further studies could also investigate the accuracy of the time buffer estimation procedure under different combinations of MTTR and MTBF failure values that actually produce the same downtime percentage.

- One of the limitations of the study was that it did not consider the effect of blocking when investigating the effect of the location and protective capacity level of the secondary constraint. Since blocking could also cause lost throughput on the constraint, an after-near secondary constraint location with a low secondary constraint protective capacity level would have a much larger influence on the flow times and line output due to the effect of blocking. Future studies could therefore include limited buffer capacity between stations as another factor when investigating the effect of the secondary constraint on line performance.
- The experiment investigating the protective capacity levels in the 15-station flow shop showed that for the low variability situation a 10% protective capacity level at non-constraints were sufficient to attain the global maximum output for the shop, whereas a 50% protective capacity level was needed for the high variability situation to ensure an output rate near the global maximum. Further experimentation with a wider range of protective capacity levels at non-constraint stations, as well as a wider range of station downtime percentages could be performed in order to better establish the relationship between protective capacity and variability and its effect on flow time, total output and the primary constraint bottleneck probability. Regression models could then be developed to better investigate this relationship. Such future studies could specifically investigate at what minimum protective capacity levels will the maximum throughput or output for the line start to drop for different resource breakdown levels and configurations. This could help to better determine the degree of unbalance required when designing production systems.
- Another limitation of this research is that the experimental studies on protective capacity were limited to serial flow production systems. It would be of value to investigate other shop configurations such as job shops that are more subject to product mix changes. It would especially be interesting to see how protective capacity should be applied in such shop configurations in order to reduce bottleneck shiftiness. In such more complex environments the protective capacity level of the secondary constraint could have a larger influence on line performance measures such as flow time and line output.
- The flow shop experimental study in this research for investigating protective capacity levels was limited to a 95% primary constraint utilisation. Further experimental studies could include different utilisation levels at the primary constraint as another independent variable. This would

help to indicate if protective capacity requirements are dependent on the utilisation level of the primary constraint. Such a study could also help to determine a "best" utilisation level for the primary constraint, and help to investigate what protective capacity levels are required at the constraint resource in order to protect market demand and delivery.

- In this study the secondary constraint did not experience any breakdowns. In order to better understand the effect of the secondary constraint in line design, it would be beneficial to investigate the impact of the secondary constraint when it is actually identified as the secondary constraint because of the breakdowns it experiences. The relative location of a secondary constraint that experience high breakdowns could have a significant effect on line performance measures such as flow time and line output. The protective capacity level of such a secondary constraint would have more of an impact than a secondary constraint that does not experience any breakdowns.
- One of the contributions of the study was the investigation of the primary constraint bottleneck probability measure. The study showed that this measure could be increased by increasing the protective capacity. This reduces short-term bottlenecks that move around, which makes it easier to identify and manage the true primary capacity constrained resource that determines the maximum throughput rate. The bottleneck probability measure of the primary constraint could maybe also be improved by increasing the time buffer size. A further study could investigate the potential increase of the bottleneck probability of the primary constraint through the use of time buffers. Such a study could also investigate the interaction effect between the time buffer (or WIP level) and the protective capacity level on the bottleneck probability measure. This could help to better understand the relationship between protective capacity and buffer levels.

Stellenbosch University http://scholar.sun.ac.za

References

1. *APICS Dictionary*, 2002, 10th ed., American Production and Inventory Control Society, Inc., Falls Church.
2. Atwater, J.B., 1991. *The impact of protective capacity on the output of a typical unblocked flow shop*. Ph.D. Dissertation. University of Georgia.
3. Atwater, J.B., and Chakravorty, S.S., 1994. Does Protective Capacity Assist Managers In Competing Along Time-Based Dimensions? *Production and Inventory Management Journal*, Third Quarter, pp.53-59.
4. Atwater, J.B., and Chakravorty, S.S., 2000. The Impact of Capacity Utilisation On The Performance Of Drum-Buffer-Rope (DBR) Systems. Working paper, Coles College of Business, Kennesaw State University.
   Available from: - http://coles.kennesaw.edu/coles/working2000.asp.
5. Bitran, G.B., and Tirupati, D., 1988. Multiproduct Queueing Networks With Deterministic Routing: Decomposition Approach And The Notion Of Interference. *Management Science*, Vol.34, Nr.1, pp. 75-100.
6. Chakravorty, S.S., and Atwater, J.B., 1996. A Comparative Study of Line Design Approaches For Serial Production Systems. *International Journal of Operations & Production Management*, Vol.16, Nr.6, pp.91-109.
7. Chase, R.B., and Aquilano, N.J., 1992. *Production & Operations Management*, Sixth Edition. Burr Ridge: Irwin.
8. Conover, W. J., 1999. *Practical Nonparametric Statistics*. USA: John Wiley & Sons.
9. Conrad, R., 1954. The rate of paced man-machine systems. *Journal of the Institute of Production Engineers*, Vol.33, p. 562.
10. Cook, D.P., 1994. A Simulation Comparison of Traditional, JIT, and TOC Manufacturing Systems in a Flow Shop With Bottlenecks. *Production and Inventory Management Journal*, First Quarter, pp. 73-78.
11. Cox, J.F., and Spencer, M.S., 1998. *The Constraints Management Handbook*. Boca Raton, Florida: CRC Press LLC.
12. Craighead, C.W., Patterson, J.W., and Fredendall, L.D., 2001. Protective capacity positioning: Impact on manufacturing cell performance. *European Journal of Operational Research*, Vol.134, pp. 425-438.
13. Davis, L.E., 1966. Pacing Effects On Manned Assembly Lines. *International Journal of Production Research*, Vol.4, Nr.3, pp.171-184.

14. De La Wyche, W., and Wild, R., 1977. The Design of Imbalanced Series Queue Flow Lines. *Operational Research Quarterly*, Vol.28. Nr.3, pp.695-702.

15. Dudley, N.A., 1963. Work-Time Distributions. *International Journal of Production Research*, Vol.2, No.2, pp.137-144.

16. Du Preez, N.D., 1984. *Ontleding Van Die Lewering Van LandbouProdukte By Opbergings- en Verwerkingsaanlegte Asook Die Operasionele Gedrag Van So 'n Aanleg*, Masters Thesis, University of Stellenbosch.

17. El-Rayah, T.E., 1979 The efficiency of balanced and unbalanced production lines. *International Journal of Production Research*, Vol.17, No.1, pp.61-75.

18. Fogarty, D.W., Blackstone, J.H., and Hoffmann, T.R., 1991. *Production & Inventory Management*, 2nd Edition, Cincinnati: South-Western Publishing Co.

19. Gardiner, S.C., and Blackstone, J.H., 1998 Dynamic Buffering. *International Journal of Production Research*, Vol.36, No.2, pp.333-342.

20. Gardiner, S.C., Blackstone, J.H., and Gardiner, L.R., 1993. Drum-Buffer-Rope and Buffer Management: Impact on Production Management Study and Practices. *International Journal of Operations & Production Management*, Vol.13, No.6, pp.68-78.

21. Goldratt, E.M., 1986. *The Goal: a Process of Ongoing Improvement*. Revised Edition. Croton-on-Hudson: North River Press.

22. Goldratt, E.M., 1990. *The Haystack Syndrome: Sifting Information Out of The Data Ocean.* Croton-on-Hudson : North River Press.

23. Goldratt, E.M., and Fox, R.E., 1986. *The Race.* Great Barrington: North River Press.

24. Hillier, F.S., and Boling, R.W., 1979. On the Optimal Allocation of Work in Symmetrically Unbalanced Production Line Systems With Variable Operation Times. *Management Science*, Vol. 25, No.8, pp. 721-728.

25. Hillier, F.S., and So, K.C., 1996. On the robustness of the bowl phenomenon. *European Journal of Operational Research*, Vol.89, pp. 496-515.

26. Hopp, W.J., and Spearman, M.L., 1996. *Factory Physics: Foundations For Manufacturing Management.* Chicago: Irwin.

27. Houle, D., 1998. Introduction to the Theory of Constraints [online]. Available from: http://www.goldratt.com/juk4t.htm

28. Hurley, S.F., and Kadipasaoglu, S., 1998. Wandering Bottlenecks: Speculating on the True Causes. *Production and Inventory Management Journal*, Fourth Quarter, pp. 1-4.

29. Kadipasaoglu, S.N., Xiang, W., Hurley, S.F., and Khumawala, B.M., 2000. A study on the effect of the extent and location of protective capacity in flow systems. *International Journal of Production Economics*, Vol.63, pp.217-228.

30. Kelton, W.D., Sadowski, R.P., and Sadowski, D.A., 1998. *Simulation with Arena.* Boston: McGraw-Hill.

31. Lawrence, S.R., and Buss, A.H., 1994. Shifting Production Bottlenecks: Causes, Cures, and Conundrums. *Production and Operations Management*, Vol. 3, No.1, pp. 21-37.
32. Miltenburg, J., 1997. Comparing JIT, MRP and TOC, and embedding TOC into MRP. *International Journal of Production Research*, Vol.35, No.4, pp.1147-1169.
33. Montgomery, D.G., 1997. *Design And Analysis Of Experiments*. Fourth Edition. USA: John Wiley & Sons.
34. Muralidhar, K., Swenseth, S.R., and Wilson, R.L., 1992. Describing Processing Times In JIT Environments. *International Journal of Production Research*, Vol.30, No.1, pp.1-11.
35. Muralidhar, K., Swenseth, S.R., and Wilson, R.L., 1992. Describing Processing Times in JIT Environments. *International Journal of Production Research*, Vol.30, No.1, pp.1-11.
36. Noreen E.N, Smith, D., and Mackey, J.T., 1995. *The Theory of Constraints and Its Implications for Management Accounting*. Great Barrington: North River Press.
37. Papadopolous, H.T., Heavy, C., and Browne, J., 1993. *Queueing Theory In Manufacturing Systems Analysis And Design*. Cambridge: University Press.
38. Patterson, J.W., Fredendall, L.D., and Craighead, C.W., 2002. The impact of non-bottleneck variation in a manufacturing cell. *Production Planning and Control*, Vol.13, No.1, pp.76-85.
39. Payne, S., Slack, N., and Wild, R., 1972. A note on the operating characteristics of 'balanced' and 'unbalanced' production flow lines. *International Journal of Production Research*, Vol.10, No.1, pp.93-98.
40. Radovilsky, Z.D., 1998. A quantitative approach to estimate the size of the time buffers in the theory of constraints. *International Journal of Production Economics*, Vol.55, pp.113-119.
41. Rahman, S., 1998. Theory of Constraints: A review of the philosophy and its applications. *International Journal of Operations & Production Management*, Vol.18, No.4, pp.336-355.
42. Rao, N.P., 1976. A generalization of the 'bowl phenomenon' in series production systems. *International Journal of Production Research*, Vol.14, No.4, pp.437-443.
43. Schragenheim, E., and Ronen, B., 1990. Drum-Buffer-Rope Shop Floor Control. *Production and Inventory Management Journal*, Third Quarter, pp. 18-23.
44. Schragenheim, E., and Ronen, B., 1991. Buffer Management: A Diagnostic Tool For Production Control. *Production and Inventory Management Journal*, Second Quarter, pp. 74-79.
45. Segal, M., and Whitt, W., 1989. A Queueing Network Analyser For Manufacturing. *Proceedings of the 12th International Teletraffic Conference*, pp.1146-1152.
46. Simons, J.V., and Simpson, W.P., 1995. An Explanation and Characterization Of The Goal System (Formerly DISASTER) Approach To Multiple Constraint Scheduling For Drum-Buffer-Rope (DBR) Production Management Systems. *Technical Memorandum*, AFIT-LA-TM-95-1, Graduate School of Logistics and Acquisition Management, Air Force Institute of Technology (AU), Wright-Patterson.

47. Simons, J.V., Simpson, W.P., Carlson, B.J., James, S.W., Lettiere, C.A., and Mediate, B.A., 1996. Formulation and solution of the drum-buffer-rope constraint scheduling problem (DBRCSP). *International Journal of Production Research*, Vol.34, No.9, pp.2405-2420.
48. Smith, J.J., 1994. TOC and MRP II, From Theory to Results [online]. Available from: http://www.rogo.com/cac/JJSmith.html.
49. Smunt, T. L., and Perkins, W.C., 1985. Stochastic Unpaced Line Design: Review and Further Experimental Results. *Journal of Operations Management*, Vol.5. No.3, pp.351-373.
50. So, K.C., 1989. On the efficiency of unbalancing production lines. *International Journal of Production Research*, Vol.27, No.4, pp.717-729.
51. Spencer, M.S., 1991. Using 'The Goal' in an MRP System. *Production and Inventory Management Journal*, Fourth Quarter, pp. 22-27.
52. Spencer, M.S., and Cox, J.F., 1995. Master Production Scheduling Development in a Theory of Constraints Environment. *Production and Inventory Management Journal*, First Quarter, pp. 8-14.
53. Spencer, M.S., and Cox, J.F., 1995. Optimum production technology (OPT) and the theory of constraints (TOC): analysis and genealogy. *International Journal of Production Research*, Vol.33, No.6, pp.1495-1504.
54. Srikanth, M.L., and Umble, M.M., 1997a. *Synchronous Management: Profit Based Manufacturing for the 21st Century Volume 1.* Guilford: The Spectrum Publishing Company.
55. Srikanth, M.L., and Umble, M.M., 1997b. *Synchronous Management: Profit Based Manufacturing for the 21st Century Volume 2.* Guilford: The Spectrum Publishing Company.
56. Stein, R.E., 1996. *Re-Engineering the Manufacturing System: Applying the Theory of Constraints.* New York: Marcel Dekker Inc.
57. Suresh, S., and Whitt, W., 1989. Arranging Queues In Series: a simulation experiment [online]. Available from: http://www.research.att.com.
58. Tu, Y.M., and Li, R.K., 1998. Constraint time buffer determination model. *International Journal of Production Research*, Vol.36, No.4, pp.1091-1103.
59. Umble, M.M., and Srikanth, M.L., 1990. *Synchronous Manufacturing: Principles for World Class Excellence*, Cincinnati: South-Western Publishing Co.
60. Whitt, W., 1982. Approximating a Point Process by a Renewal Process, I: Two Basic Methods. *Operations Research*, Vol. 30, No. 1, pp. 125-147.
61. Whitt, W., 1983a. The Queueing Network Analyzer. *The Bell System Technical Journal*, Vol.62, No.9, pp.2779-2815.
62. Whitt, W., 1983b. Performance of the Queueing Network Analyzer. *The Bell System Technical Journal*, Vol.62, No.9, pp.2817-2843.

Stellenbosch University http://scholar.sun.ac.za

63. Whitt, W., 1993. Approximations for the GI/G/m Queue. *Production and Operations Management*, Vol.2, 114-161.

Bibliography

1. Adams, J., Balas, E., and Zawack, D., 1988. The Shifting Bottleneck Procedure For Job Shop Scheduling. *Management Science*, Vol. 34, No.3, pp. 391-401.
2. Altiok, T., and Ranjan, R., 1989. Analysis Of Production Lines With General Service Times And Finite Buffers: A Two-Node Decomposition Approach. *Engineering Costs and Production Economics*, Vol.17, pp. 155-165.
3. *APICS Dictionary*, 2002, 10th ed., American Production and Inventory Control Society, Inc., Falls Church.
4. Ashton, J.E., Johnson, M.D., and Cook, F.X., 1990. Shop Floor Control in a System Job Shop: Definitely Not MRP. *Production and Inventory Management Journal*, Second Quarter, pp. 27-31.
5. Atwater, J.B., 1991. *The impact of protective capacity on the output of a typical unblocked flow shop*. Ph.D. Dissertation. University of Georgia.
6. Atwater, J.B., and Chakravorty, S.S., 1994. Does Protective Capacity Assist Managers In Competing Along Time-Based Dimensions? *Production and Inventory Management Journal*, Third Quarter, pp.53-59.
7. Atwater, J.B., and Chakravorty, S.S., 2000. The Impact of Capacity Utilisation On The Performance Of Drum-Buffer-Rope (DBR) Systems [online]. Working paper, Coles College of Business, Kennesaw State University.
   Available from: - http://coles.kennesaw.edu/coles/working2000.asp.
8. Bitran, G.B., and Tirupati, D., 1988. Multiproduct Queueing Networks With Deterministic Routing: Decomposition Approach And The Notion Of Interference. *Management Science*, Vol.34, Nr.1, pp. 75-100.
9. Chakravorty, S.S., 1996. Robert Bowden Inc.: A Case Study of Cellular Manufacturing and Drum-Buffer-Rope Implementation. *Production and Inventory Management Journal*, Third Quarter, pp. 15-19.
10. Chakravorty, S.S., and Atwater, J.B., 1996. A Comparative Study of Line Design Approaches For Serial Production Systems. *International Journal of Operations & Production Management*, Vol.16, Nr.6, pp.91-109.
11. Chase, R.B., and Aquilano, N.J., 1992. *Production & Operations Management*, Sixth Edition. Burr Ridge: Irwin.
12. Conover, W. J., 1999. *Practical Nonparametric Statistics*. USA: John Wiley & Sons.
13. Conrad, R., 1954. The rate of paced man-machine systems. *Journal of the Institute of Production Engineers*, Vol.33, p. 562.

14. Conway, R., Maxwell, W., McClain, J.O., and Thomas, L.J., 1988. The Role Of Work-In-Process Inventory In Serial Production Lines. *Operations Research*, Vol.36, No.2, pp.229-241.

15. Cook, D.P., 1994. A Simulation Comparison of Traditional, JIT, and TOC Manufacturing Systems in a Flow Shop With Bottlenecks. *Production and Inventory Management Journal*, First Quarter, pp. 73-78.

16. Cox, J.F., and Spencer, M.S., 1998. *The Constraints Management Handbook*. Boca Raton, Florida: CRC Press LLC.

17. Craighead, C.W., Patterson, J.W., and Fredendall, L.D., 2001. Protective capacity positioning: Impact on manufacturing cell performance. *European Journal of Operational Research*, Vol.134, pp. 425-438.

18. Dallery, Y., and Frein, Y., On Decomposition Methods for Tandem Queueing Networks With Blocking. *Operations Research*, Vol.41, No.2, pp. 386-399.

19. Davel, H.M., 1992. *Die Inkorporering Van OPT Beginsels in 'n Voorwaartse Skeduleringsprosedure Wat Van Diskrete Gebeurtenis Simulasie Gebruik Maak*, M.Eng Thesis, University of Stellenbosch.

20. Davis, L.E., 1966. Pacing Effects On Manned Assembly Lines. *International Journal of Production Research*, Vol.4, Nr.3, pp.171-184.

21. De La Wyche, W., and Wild, R., 1977. The Design of Imbalanced Series Queue Flow Lines. *Operational Research Quarterly*, Vol.28. Nr.3, pp.695-702.

22. Demeulemeester, E.L., and Herroelen, W.S., 1996. Modelling setup times, process batches and transfer batches using activity network logic. *European Journal of Operational Research*, Vol. 89, pp. 355-365.

23. Dudley, N.A., 1963. Work-Time Distributions. *International Journal of Production Research*, Vol.2, No.2, pp.137-144.

24. Du Preez, N.D., 1984. *Ontleding Van Die Lewering Van LandbouProdukte By Opbergings- en Verwerkingsaanlegte Asook Die Operasionele Gedrag Van So 'n Aanleg*, Masters Thesis, University of Stellenbosch.

25. El-Rayah, T.E., 1979 The efficiency of balanced and unbalanced production lines. *International Journal of Production Research*, Vol.17, No.1, pp.61-75.

26. Fogarty, D.W., Blackstone, J.H., and Hoffmann, T.R., 1991. *Production & Inventory Management*, 2nd Edition, Cincinnati: South-Western Publishing Co.

27. Franks, S., and Keith, E., 1995. Material Requirements Planning – Rest in Peace. *SAPICS 17th International Conference*, Earn Your Place in the Global Economy, Port Elizabeth, South Africa.

28. Gardiner, S.C., and Blackstone, J.H., 1998 Dynamic Buffering. *International Journal of Production Research*, Vol.36, No.2, pp.333-342.

29. Gardiner, S.C., Blackstone, J.H., and Gardiner, L.R., 1993. Drum-Buffer-Rope and Buffer Management: Impact on Production Management Study and Practices. *International Journal of Operations & Production Management*, Vol.13, No.6, pp.68-78.
30. Gargeya, V.B., and Deane, R.H., 1996. Scheduling research in multiple resource constrained job shops: a review and critique. *International Journal of Production Research*, Vol.34, No.8, pp.2077-2097.
31. Gelenbe, E., and Pujole, G., 1999. *Introduction To Queueing Networks*. Second Edition. Chichester: John Wiley & Sons.
32. Gershwin, S.B., 1987. An Efficient Decomposition Method For The Approximate Evaluation Of Tandem Queues With Finite Storage Space and Blocking. *Operations Research*, Vol.35, No.2, pp.291-305.
33. Goldratt, E.M., 1986. *The Goal: a Process of Ongoing Improvement*. Revised Edition. Croton-on-Hudson: North River Press.
34. Goldratt, E.M., 1988. Computerised shop floor scheduling. *International Journal of Production Research*, Vol.26, No.3, pp.443-455.
35. Goldratt, E.M., 1990. *Essays On The Theory Of Constraints*. Great Barrington: North River Press.
36. Goldratt, E.M., 1990. *The Haystack Syndrome: Sifting Information Out of The Data Ocean*. Croton-on-Hudson : North River Press.
37. Goldratt, E.M., 1994. *It's Not Luck*. Great Barrington: North River Press.
38. Goldratt, E.M., 1997. *Critical Chain*. Great Barrington: North River Press.
39. Goldratt, E.M., and Fox, R.E., 1986. *The Race*. Great Barrington: North River Press.
40. Goldratt, E.M., Schragenheim, E., and Ptak, C.A. 2000. *Necessary But Not Sufficient*. Great Barrington: North River Press.
41. Guide, V.D.R., 1995. A Simulation Model of Drum-Buffer-Rope for Production Planning and Control at a Naval Aviation Depot. *Simulation*, Vol. 65, No.3, pp.157-168.
42. Guide, V.D.R., 1996. Scheduling using drum-buffer-rope in a remanufacturing environment. *International Journal of Production Research*, Vol.34, No.4, pp.1081-1091.
43. Guide, V.D.R., and Ghiselli, G.A., 1995. Implementation of Drum-Buffer-Rope at a Military Rework Depot Engine Works. *Production and Inventory Management Journal*, Third Quarter, pp. 79-83.
44. Gupta, S.M., and Kavusturucu, A., 1998. Modelling of Finite Buffer Cellular Manufacturing Systems With Unreliable Machines. *International Journal of Industrial Engineering*, Vol.5, No.4, pp. 265-277.
45. Harrison. J.M., and Pich, T., 1996. Two-Moment Analysis of Open Queueing Networks With General Workstation Capabilities. *Operations Research*, Vol.44, No.6, pp.936-950.

46. Hillier, F.S., and Boling, R.W., 1979. On the Optimal Allocation of Work in Symmetrically Unbalanced Production Line Systems With Variable Operation Times. *Management Science*, Vol. 25, No.8, pp. 721-728.

47. Hillier, F.S., and So, K.C., 1996. On the robustness of the bowl phenomenon. *European Journal of Operational Research*, Vol.89, pp. 496-515.

48. Hopp, W.J., and Spearman, M.L., 1996. *Factory Physics: Foundations For Manufacturing Management.* Chicago: Irwin.

49. Houle, D., 1998. Introduction to the Theory of Constraints [online]. Available from: http://www.goldratt.com/juk4t.htm

50. Hurley, S.F., and Kadipasaoglu, S., 1998. Wandering Bottlenecks: Speculating on the True Causes. *Production and Inventory Management Journal*, Fourth Quarter, pp. 1-4.

51. Ivens, P., and Lambrecht, M., 1996. Extending the shifting bottleneck procedure to real-life applications. *European Journal of Operational Research*, Vol.90, pp. 252-268.

52. Kadipasaoglu, S.N., Xiang, W., Hurley, S.F., and Khumawala, B.M., 2000. A study on the effect of the extent and location of protective capacity in flow systems. *International Journal of Production Economics*, Vol.63, pp.217-228.

53. Kamenetsky, R.D., 1985. Successful MRP II implementation can be complemented by smart scheduling, sequencing systems. Industrial Engineering, Vol.17, No.10, pp. 44-52.

54. Kanet, J.J., 1988. MRP 96: Time to Rethink Manufacturing Logistics. *Production and Inventory Management Journal*, Second Quarter, pp. 57-61.

55. Kelton, W.D., Sadowski, R.P., and Sadowski, D.A., 1998. *Simulation with Arena.* Boston: McGraw-Hill.

56. Kerbache, L., and Smith, J.M., 1987. The generalized expansion method for open finite queueing networks. *European Journal of Operational Research*, Vol.32, pp. 448-461.

57. Kouvatsos, D.D., 1994. Entropy maximisation and queueing network models. *Annals of Operations Research*, Vol.48, pp. 63-126.

58. Lambrecht, M.R, and Decaluwe, L., 1988. JIT and Constraint Theory: The Issue of Bottleneck Management. *Production and Inventory Management Journal*, Third Quarter, pp. 61-65.

59. Law, M.L., and Kelton, W.D., 1982. *Simulation Modelling and Analysis*, New York: McGraw-Hill.

60. Lawrence, S.R., and Buss, A.H., 1994. Shifting Production Bottlenecks: Causes, Cures, and Conundrums. *Production and Operations Management*, Vol. 3, No.1, pp. 21-37.

61. Lockamy, A., and Cox, J.F., 1991. Using V-A-T analysis for determining the priority and location of JIT manufacturing techniques. *International Journal of Production Research*, Vol.29, No.8, pp.1661-1672.

62. Miltenburg, J., 1997. Comparing JIT, MRP and TOC, and embedding TOC into MRP. *International Journal of Production Research*, Vol.35, No.4, pp.1147-1169.

63. Montgomery, D.G., 1997. *Design And Analysis Of Experiments*. Fourth Edition. USA: John Wiley & Sons.

64. Moss, H.K., and Yu, W.B., 1999. Toward The Estimation Of Bottleneck Shiftiness In a Manufacturing Operation. *Production and Inventory Management Journal*, Second Quarter, pp. 53-58.

65. Muralidhar, K., Swenseth, S.R., and Wilson, R.L., 1992. Describing Processing Times in JIT Environments. *International Journal of Production Research*, Vol.30, No.1, pp.1-11.

66. Noreen E.N, Smith, D., and Mackey, J.T., 1995. *The Theory of Constraints and Its Implications for Management Accounting*. Great Barrington: North River Press.

67. Papadopolous, H.T., Heavy, C., and Browne, J., 1993. *Queueing Theory In Manufacturing Systems Analysis And Design*. Cambridge: University Press.

68. Patterson, J.W., Fredendall, L.D., and Craighead, C.W., 2002. The impact of non-bottleneck variation in a manufacturing cell. *Production Planning and Control*, Vol.13, No.1, pp.76-85.

69. Payne, S., Slack, N., and Wild, R., 1972. A note on the operating characteristics of 'balanced' and 'unbalanced' production flow lines. *International Journal of Production Research*, Vol.10, No.1, pp.93-98.

70. Ptak, C.A., 1991. MRP, MRP II, OPT, JIT, and CIM – Succession, Evolution, Or Necessary Combination. *Production and Inventory Management Journal*, Second Quarter, pp. 7-11.

71. Raban, S., and Nagel, R.N., 1991. Constraint-based control of flexible flow lines. *International Journal of Production Research*, Vol.29, No.10, pp.1941-1951.

72. Radovilsky, Z.D., 1998. A quantitative approach to estimate the size of the time buffers in the theory of constraints. *International Journal of Production Economics*, Vol.55, pp.113-119.

73. Rahman, S., 1998. Theory of Constraints: A review of the philosophy and its applications. *International Journal of Operations & Production Management*, Vol.18, No.4, pp.336-355.

74. Rao, N.P., 1976. A generalization of the 'bowl phenomenon' in series production systems. *International Journal of Production Research*, Vol.14, No.4, pp.437-443.

75. Ronen, B., and Starr, M.K., 1990. Synchronized Manufacturing as in OPT: From Practice to Theory. *Computers & Industrial Engineering*, Vol.18, No.4, pp.585-600.

76. Sadre, R,, Boudewijn, R.H., and Ost, A., 1999. An Efficient and Accurate Decomposition Method for Open Finite- and Infinite-Buffer Queueing Networks [online]. Available from: http://www-lvs.informatik.rwth-aachen.de.

77. Scheduling Technologies Group (STG), 1998. OPT21 User Manual [online]. Available from: http://www.stg.co.uk.

78. Scheduling Technologies Group (STG), 1998. ST-POINT User Manual [online]. Available from: http://www.stg.co.uk.

79. Schneider, D.I., 1998. *An Introduction to Programming Using Visual Basic 5.0*. Upper Saddle River: Prentice Hall.

80. Schragenheim, E., and Dettmer, H., 2001. *Manufacturing at Warp Speed: Optimizing Supply Chain Financial Performance*. Boca Raton: St. Lucie Press.

81. Schragenheim, E., and Ronen, B., 1990. Drum-Buffer-Rope Shop Floor Control. *Production and Inventory Management Journal*, Third Quarter, pp. 18-23.

82. Schragenheim, E., and Ronen, B., 1991. Buffer Management: A Diagnostic Tool For Production Control. *Production and Inventory Management Journal*, Second Quarter, pp. 74-79.

83. Schragenheim, E., Cox, J., and Ronen, B., 1994. Process flow industry – scheduling and control using theory of constraints. *International Journal of Production Research*, Vol.32, No.8, pp.1867-1877.

84. Segal, M., and Whitt, W., 1989. A Queueing Network Analyser For Manufacturing. *Proceedings of the 12th International Teletraffic Conference*, pp.1146-1152.

85. Simons, J.V., and Simpson, W.P., 1995. An Explanation and Characterization Of The Goal System (Formerly DISASTER) Approach To Multiple Constraint Scheduling For Drum-Buffer-Rope (DBR) Production Management Systems. *Technical Memorandum*, AFIT-LA-TM-95-1, Graduate School of Logistics and Acquisition Management, Air Force Institute of Technology (AU), Wright-Patterson.

86. Simons, J.V., Simpson, W.P., Carlson, B.J., James, S.W., Lettiere, C.A., and Mediate, B.A., 1996. Formulation and solution of the drum-buffer-rope constraint scheduling problem (DBRCSP). *International Journal of Production Research*, Vol.34, No.9, pp.2405-2420.

87. Smith, J.J., 1994. TOC and MRP II, From Theory to Results [online]. Available from: http://www.rogo.com/cac/JJSmith.html.

88. Smunt, T. L., and Perkins, W.C., 1985. Stochastic Unpaced Line Design: Review and Further Experimental Results. *Journal of Operations Management*, Vol.5. No.3, pp.351-373.

89. So, K.C., 1989. On the efficiency of unbalancing production lines. *International Journal of Production Research*, Vol.27, No.4, pp.717-729.

90. Spencer, M.S., 1991. Using 'The Goal' in an MRP System. *Production and Inventory Management Journal*, Fourth Quarter, pp. 22-27.

91. Spencer, M.S., 1994. Economic theory, cost accounting and theory of constraints: an examination of relationships and problems. *International Journal of Production Research*, Vol.32, No.2, pp.299-308.

92. Spencer, M.S., and Cox, J.F., 1995. Master Production Scheduling Development in a Theory of Constraints Environment. *Production and Inventory Management Journal*, First Quarter, pp. 8-14.

93. Spencer, M.S., and Cox, J.F., 1995. Optimum production technology (OPT) and the theory of constraints (TOC): analysis and genealogy. *International Journal of Production Research*, Vol.33, No.6, pp.1495-1504.

Stellenbosch University http://scholar.sun.ac.za

94. Srikanth, M.L., and Umble, M.M., 1997a. *Synchronous Management: Profit Based Manufacturing for the 21st Century Volume 1.* Guilford: The Spectrum Publishing Company.

95. Srikanth, M.L., and Umble, M.M., 1997b. *Synchronous Management: Profit Based Manufacturing for the 21st Century Volume 2.* Guilford: The Spectrum Publishing Company.

96. Steele, D.C, and Papke-Shields, K.E., 1993. Capacity Slack: Strategic Alternative To Lead Time. *Production and Inventory Management Journal*, Fourth Quarter, pp. 1-5.

97. Stein, R.E., 1996. *Re-Engineering the Manufacturing System: Applying the Theory of Constraints.* New York: Marcel Dekker Inc.

98. Suresh, S., and Whitt, W., 1989. Arranging Queues In Series [online]. Available from: http://www.research.att.com.

99. Tu, Y.M., and Li, R.K., 1998. Constraint time buffer determination model. *International Journal of Production Research*, Vol.36, No.4, pp.1091-1103.

100. Turner, V.K., 1992. *The Development of a Computerised Heuristic Production Scheduling Algorithm Based on the Synchronous Management (OPT) Principles*, M.Eng Thesis, University of Stellenbosch.

101. Umble, M.M., and Srikanth, M.L., 1990. *Synchronous Manufacturing: Principles for World Class Excellence*, Cincinnati: South-Western Publishing Co.

102. Whitt, W., 1982. Approximating a Point Process by a Renewal Process, I: Two Basic Methods. *Operations Research*, Vol. 30, No. 1, pp. 125-147.

103. Whitt, W., 1983a. The Queueing Network Analyzer. *The Bell System Technical Journal*, Vol.62, No.9, pp.2779-2815.

104. Whitt, W., 1983b. Performance of the Queueing Network Analyzer. *The Bell System Technical Journal*, Vol.62, No.9, pp.2817-2843.

105. Whitt, W., 1993. Approximations for the GI/G/m Queue. *Production and Operations Management*, Vol.2, 114-161.

106. Wu, S-Y., Morris, J.S., and Gordon, T.M., 1994. A Simulation Analysis of the Effectiveness of Drum-Buffer-Rope Scheduling in Furniture Manufacturing. *Computers & Industrial Engineering*, Vol.26, No.4, pp. 757-764.

Stellenbosch University http://scholar.sun.ac.za

Appendix A

Computerised Tool For Analytical Time Buffer Estimation Procedure

This appendix contains screenshots of the computerised time buffer estimation tool developed in Visual Basic 6.0. It is organised according to the different functionalities of the tool.

A.1 Data entry

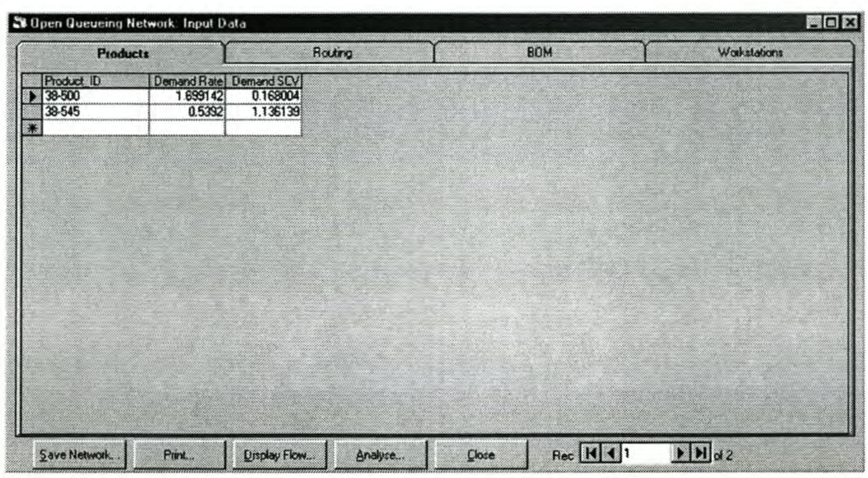

Product_ID	Demand Rate	Demand SCV
38-500	1.699142	0.168004
38-545	0.5392	1.136139

Figure A- 1. Product demand data entry screen

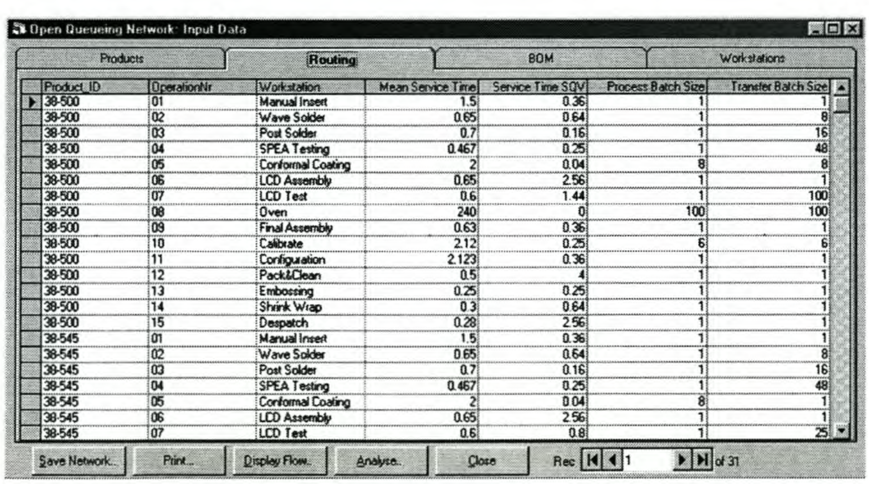

Product_ID	OperationNr	Workstation	Mean Service Time	Service Time SQV	Process Batch Size	Transfer Batch Size
38-500	01	Manual Insert	1.5	0.36	1	1
38-500	02	Wave Solder	0.65	0.64	1	8
38-500	03	Post Solder	0.7	0.16	1	16
38-500	04	SPEA Testing	0.467	0.25	1	48
38-500	05	Conformal Coating	2	0.04	8	8
38-500	06	LCD Assembly	0.65	2.56	1	1
38-500	07	LCD Test	0.6	1.44	1	100
38-500	08	Oven	240	0	100	100
38-500	09	Final Assembly	0.63	0.36	1	1
38-500	10	Calibrate	2.12	0.25	6	6
38-500	11	Configuration	2.123	0.36	1	1
38-500	12	Pack&Clean	0.5	4	1	1
38-500	13	Embossing	0.25	0.25	1	1
38-500	14	Shrink Wrap	0.3	0.64	1	1
38-500	15	Despatch	0.28	2.56	1	1
38-545	01	Manual Insert	1.5	0.36	1	1
38-545	02	Wave Solder	0.65	0.64	1	8
38-545	03	Post Solder	0.7	0.16	1	16
38-545	04	SPEA Testing	0.467	0.25	1	48
38-545	05	Conformal Coating	2	0.04	8	1
38-545	06	LCD Assembly	0.65	2.56	1	1
38-545	07	LCD Test	0.6	0.8	1	25

Figure A- 2. Routing data entry screen

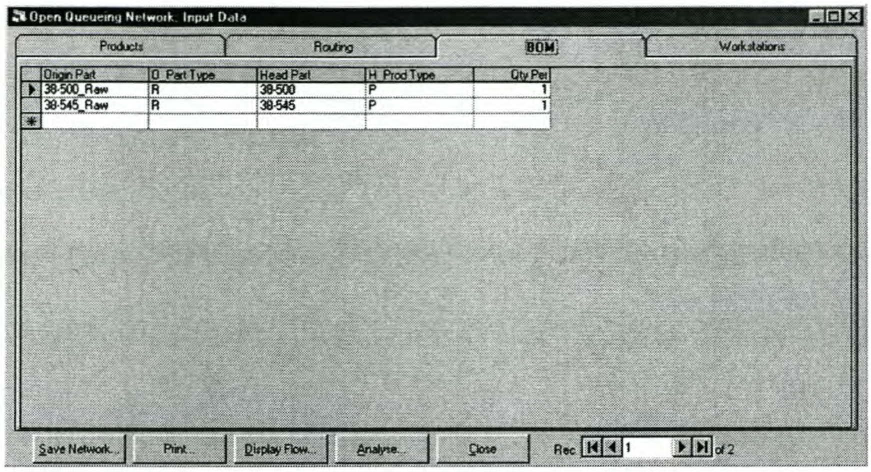

Figure A- 3. Bill-of-Material data entry screen

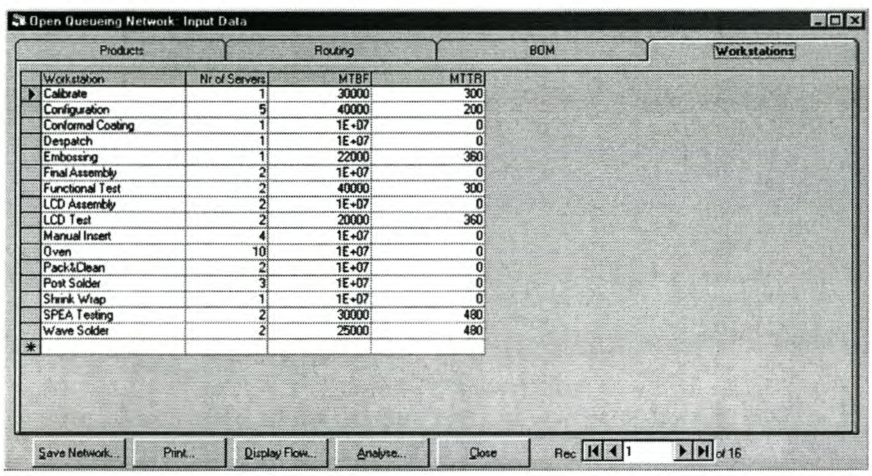

Figure A- 4. Workstation data entry screen

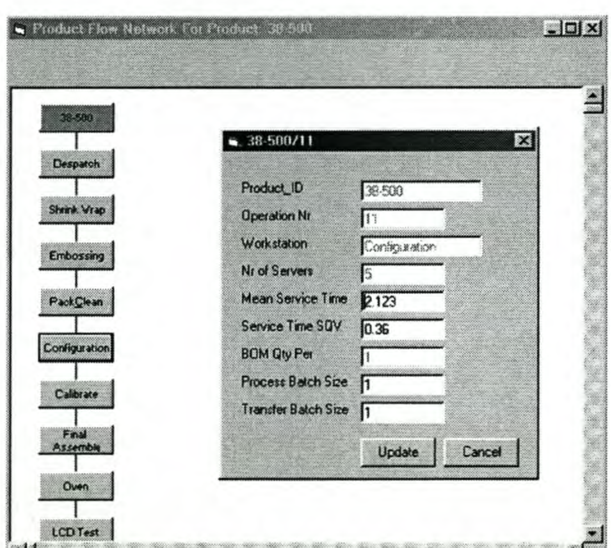

Figure A- 5. Production flow network

Figure A- 5 presents a visual flow diagram of the production flow network for a specific product that can be generated from the BOM, routing and workstation input data. Input data can also be modified from within this product flow diagram by clicking on a specific station in the flow network. A dialog box is then displayed as shown in Figure A- 5 that allows certain input data for the station to be modified.

A.2 Calculating the time buffer lengths

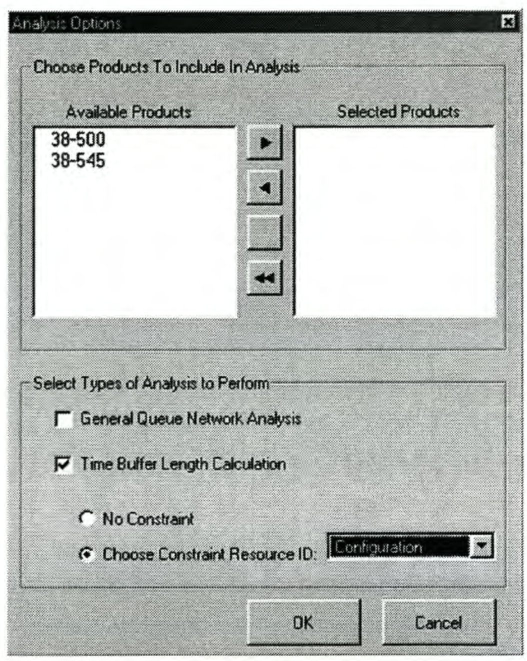

Figure A- 6. Time buffer analysis options

Figure A- 6 presents the options that can be chosen for the time buffer analysis. The specific products to be included in the analysis should be chosen. An option is also available to distinguish whether a general queue network analysis should be performed, or whether a time buffer analysis should be performed. Choosing the queuing analysis options generates the two reports presented in Figure A- 7 and Figure A- 8. These reports display general results for an open queuing network analysis, such as the mean flow times, queue times, throughput rates, WIP, arrival squared coefficient of variation, mean service times, service time squared coefficient of variation. Choosing the time buffer calculation option in Figure A- 6 requires either choosing no primary constraint and therefore only the shipping buffer length is calculated, or choosing the primary constraint workstation from a list of all the available workstations. A time buffer calculation report is generated as displayed in Figure A- 9. This report identifies the locations of all the required time buffers, the type of time buffer, the length of the time buffer, as well as the mean flow time, flow time variance, and the total raw processing time that was used to calculate the final time buffer lengths.

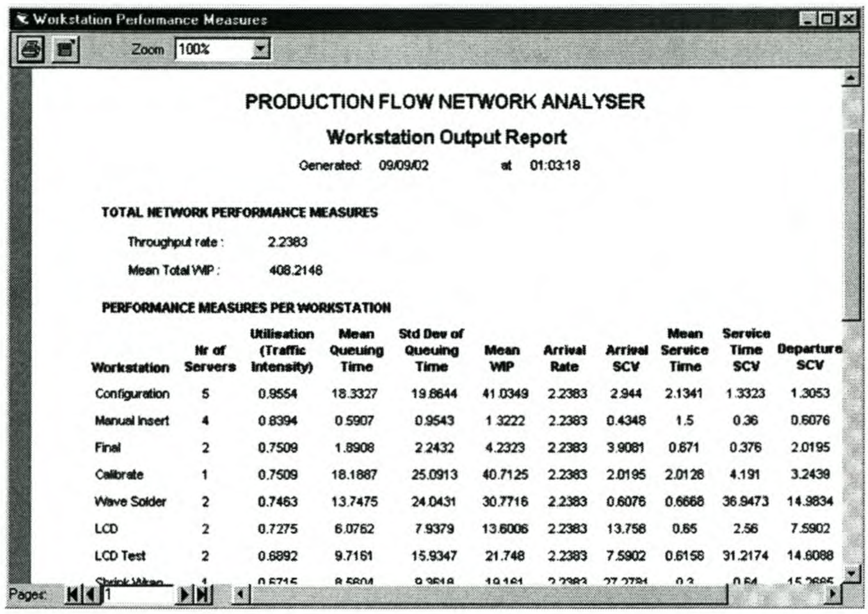

PRODUCTION FLOW NETWORK ANALYSER

Workstation Output Report

Generated: 09/09/02 at 01:03:18

TOTAL NETWORK PERFORMANCE MEASURES

Throughput rate : 2.2383

Mean Total WIP : 408.2146

PERFORMANCE MEASURES PER WORKSTATION

Workstation	Nr of Servers	Utilisation (Traffic Intensity)	Mean Queuing Time	Std Dev of Queuing Time	Mean WIP	Arrival Rate	Arrival SCV	Mean Service Time	Service Time SCV	Departure SCV
Configuration	5	0.9554	18.3327	19.8644	41.0349	2.2383	2.944	2.1341	1.3323	1.3053
Manual Insert	4	0.8394	0.5907	0.9543	1.3222	2.2383	0.4348	1.5	0.36	0.6076
Final	2	0.7509	1.8908	2.2432	4.2323	2.2383	3.9081	0.671	0.376	2.0195
Calibrate	1	0.7509	18.1887	25.0913	40.7125	2.2383	2.0195	2.0128	4.191	3.2439
Wave Solder	2	0.7463	13.7475	24.0431	30.7716	2.2383	0.6076	0.6668	36.9473	14.9834
LCD	2	0.7275	6.0762	7.9379	13.6006	2.2383	13.758	0.65	2.56	7.5902
LCD Test	2	0.6892	9.7161	15.9347	21.748	2.2383	7.5902	0.6158	31.2174	14.6088

Figure A- 7. Workstation performance report

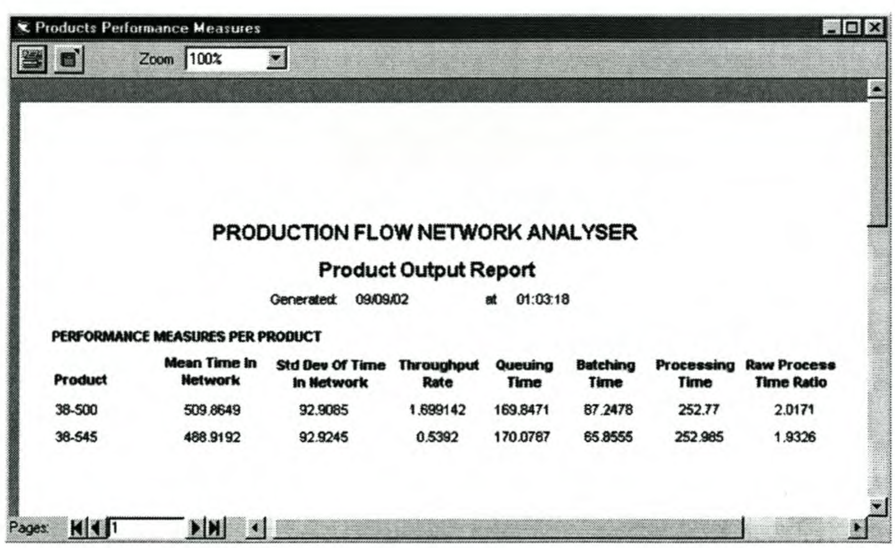

PRODUCTION FLOW NETWORK ANALYSER

Product Output Report

Generated: 09/09/02 at 01:03:18

PERFORMANCE MEASURES PER PRODUCT

Product	Mean Time In Network	Std Dev Of Time In Network	Throughput Rate	Queuing Time	Batching Time	Processing Time	Raw Process Time Ratio
38-500	509.8649	92.9085	1.699142	169.8471	87.2478	252.77	2.0171
38-545	488.9192	92.9245	0.5392	170.0787	65.8555	252.985	1.9326

Figure A- 8. Product performance report

Stellenbosch University http://scholar.sun.ac.za

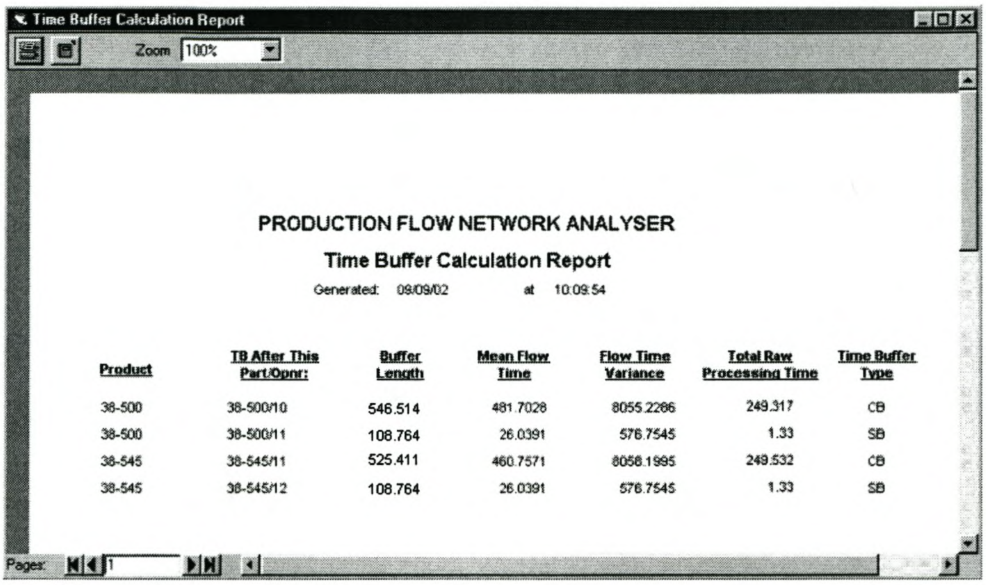
Time Buffer Calculation Report

Zoom 100%

PRODUCTION FLOW NETWORK ANALYSER

Time Buffer Calculation Report

Generated: 09/09/02 at 10:09:54

Product	TB After This Part/Opnr:	Buffer Length	Mean Flow Time	Flow Time Variance	Total Raw Processing Time	Time Buffer Type
38-500	38-500/10	546.514	481.7028	8055.2286	249.317	CB
38-500	38-500/11	108.764	26.0391	576.7545	1.33	SB
38-545	38-545/11	525.411	460.7571	8058.1995	249.532	CB
38-545	38-545/12	108.764	26.0391	576.7545	1.33	SB

Pages: 1

Figure A- 9. Time buffer calculation report

Stellenbosch University http://scholar.sun.ac.za

Appendix B

Time Buffer Evaluation Simulation Study: Data Used In Fifteen Station Flow Shop Model

B.1 Material release schedule developed with DBR program: High Downtimes

OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date
ORDER1/38-500	18	6/23/99	ORDER19/38-545	350	7/10/99	ORDER43/38-545	550	7/23/99
ORDER1/38-500	1092	6/24/99	ORDER21/38-545	504	7/10/99	ORDER45/38-500	352	7/23/99
ORDER2/38-500	18	6/24/99	ORDER22/38-500	124	7/10/99	ORDER42/38-500	500	7/24/99
ORDER2/38-500	1092	6/25/99	ORDER24/38-500	283	7/10/99	ORDER44/38-545	550	7/24/99
ORDER3/38-500	18	6/25/99	ORDER22/38-500	436	7/11/99	ORDER46/38-500	18	7/24/99
ORDER3/38-500	1092	6/26/99	ORDER23/38-545	550	7/11/99	ORDER45/38-500	208	7/24/99
ORDER4/38-500	18	6/26/99	ORDER25/38-500	18	7/11/99	ORDER46/38-500	1092	7/25/99
ORDER4/38-500	1092	6/27/99	ORDER24/38-500	272	7/11/99	ORDER47/38-500	560	7/26/99
ORDER5/38-545	550	6/28/99	ORDER25/38-500	1092	7/12/99	ORDER48/38-545	200	7/26/99
ORDER7/38-545	550	6/29/99	ORDER26/38-500	18	7/12/99	ORDER50/38-545	46	7/26/99
ORDER8/38-500	18	6/29/99	ORDER26/38-500	1092	7/13/99	ORDER49/38-500	268	7/26/99
ORDER6/38-500	500	6/29/99	ORDER27/38-500	500	7/14/99	ORDER48/38-545	350	7/27/99
ORDER8/38-500	1092	6/30/99	ORDER28/38-545	200	7/14/99	ORDER50/38-545	504	7/27/99
ORDER9/38-500	18	6/30/99	ORDER31/38-500	92	7/14/99	ORDER51/38-500	18	7/27/99
ORDER9/38-500	1092	7/1/99	ORDER28/38-545	350	7/15/99	ORDER49/38-500	292	7/27/99
ORDER10/38-500	18	7/1/99	ORDER29/38-500	500	7/15/99	ORDER51/38-500	1092	7/28/99
ORDER10/38-500	1092	7/2/99	ORDER30/38-545	200	7/15/99	ORDER52/38-500	18	7/28/99
ORDER11/38-500	18	7/2/99	ORDER32/38-545	46	7/15/99	ORDER52/38-500	1092	7/29/99
ORDER11/38-500	1092	7/3/99	ORDER31/38-500	176	7/15/99	ORDER53/38-545	200	7/30/99
ORDER12/38-500	560	7/4/99	ORDER30/38-545	350	7/16/99	ORDER55/38-545	194	7/30/99
ORDER14/38-500	18	7/4/99	ORDER32/38-545	504	7/16/99	ORDER53/38-545	350	7/31/99
ORDER13/38-545	550	7/4/99	ORDER33/38-500	18	7/16/99	ORDER54/38-500	550	7/31/99
ORDER14/38-500	1092	7/5/99	ORDER31/38-500	292	7/16/99	ORDER56/38-500	18	7/31/99
ORDER15/38-500	18	7/5/99	ORDER33/38-500	1092	7/17/99	ORDER55/38-545	356	7/31/99
ORDER15/38-500	1092	7/6/99	ORDER34/38-500	18	7/17/99	ORDER56/38-500	1092	8/1/99
ORDER16/38-500	18	7/6/99	ORDER34/38-500	1092	7/18/99	ORDER57/38-500	18	8/1/99
ORDER16/38-500	1092	7/7/99	ORDER35/38-500	18	7/18/99	ORDER57/38-500	1092	8/2/99
ORDER17/38-500	500	7/8/99	ORDER35/38-500	1092	7/19/99	ORDER58/38-500	18	8/2/99
ORDER18/38-545	200	7/8/99	ORDER36/38-500	500	7/20/99	ORDER58/38-500	1092	8/3/99
ORDER18/38-545	350	7/9/99	ORDER37/38-545	550	7/20/99	ORDER59/38-545	200	8/4/99
ORDER20/38-500	560	7/9/99	ORDER39/38-545	550	7/21/99	ORDER61/38-545	37	8/4/99
ORDER19/38-545	200	7/9/99	ORDER38/38-500	500	7/21/99	ORDER59/38-545	350	8/5/99
ORDER21/38-545	46	7/9/99	ORDER41/38-545	550	7/22/99	ORDER61/38-545	513	8/5/99
ORDER24/38-500	5	7/9/99	ORDER40/38-500	500	7/22/99	ORDER60/38-500	102	8/5/99

Stellenbosch University http://scholar.sun.ac.za

OrderID/ProdID	Rel Qty	Rel Date
ORDER63/38-545	238	8/5/99
ORDER60/38-500	448	8/6/99
ORDER62/38-500	500	8/6/99
ORDER64/38-500	18	8/6/99
ORDER63/38-545	312	8/6/99
ORDER64/38-500	1092	8/7/99
ORDER65/38-500	500	8/8/99
ORDER67/38-500	18	8/8/99
ORDER66/38-545	550	8/8/99
ORDER67/38-500	1092	8/9/99
ORDER68/38-500	18	8/9/99
ORDER68/38-500	1092	8/10/99
ORDER69/38-500	18	8/10/99
ORDER69/38-500	1092	8/11/99
ORDER70/38-500	560	8/12/99
ORDER71/38-545	200	8/12/99
ORDER72/38-545	182	8/12/99
ORDER71/38-545	350	8/13/99
ORDER73/38-500	560	8/13/99
ORDER72/38-545	368	8/13/99
ORDER75/38-545	550	8/14/99
ORDER76/38-545	342	8/14/99
ORDER74/38-500	500	8/15/99
ORDER77/38-500	550	8/15/99
ORDER78/38-500	18	8/15/99
ORDER76/38-545	208	8/15/99
ORDER78/38-500	1092	8/16/99
ORDER80/38-500	18	8/17/99
ORDER79/38-545	550	8/17/99
ORDER80/38-500	1092	8/18/99
ORDER81/38-500	18	8/18/99
ORDER81/38-500	1092	8/19/99
ORDER82/38-500	18	8/19/99
ORDER82/38-500	1092	8/20/99
ORDER83/38-500	18	8/20/99
ORDER83/38-500	1092	8/21/99
ORDER84/38-500	18	8/21/99
ORDER84/38-500	1092	8/22/99
ORDER86/38-500	18	8/23/99
ORDER85/38-545	550	8/23/99
ORDER86/38-500	1092	8/24/99
ORDER87/38-500	18	8/24/99
ORDER87/38-500	1092	8/25/99
ORDER88/38-500	18	8/25/99
ORDER88/38-500	1092	8/26/99
ORDER89/38-500	560	8/27/99
ORDER91/38-500	18	8/27/99
ORDER90/38-545	550	8/27/99
ORDER91/38-500	1092	8/28/99
ORDER93/38-545	550	8/29/99
ORDER94/38-500	18	8/29/99
ORDER92/38-500	500	8/29/99
ORDER94/38-500	1092	8/30/99
ORDER95/38-500	550	8/31/99
ORDER97/38-500	18	8/31/99
ORDER96/38-545	550	8/31/99
ORDER97/38-500	1092	9/1/99
ORDER98/38-500	18	9/1/99
ORDER98/38-500	1092	9/2/99
ORDER99/38-500	560	9/3/99
ORDER101/38-500	18	9/3/99
ORDER100/38-545	550	9/3/99
ORDER101/38-500	1092	9/4/99
ORDER102/38-500	18	9/4/99
ORDER102/38-500	1092	9/5/99
ORDER103/38-500	18	9/5/99
ORDER103/38-500	1092	9/6/99
ORDER104/38-500	560	9/7/99
ORDER105/38-545	200	9/7/99
ORDER107/38-545	150	9/7/99
ORDER105/38-545	350	9/8/99
ORDER106/38-500	500	9/8/99
ORDER108/38-500	18	9/8/99
ORDER107/38-545	400	9/8/99
ORDER108/38-500	1092	9/9/99
ORDER109/38-500	18	9/9/99
ORDER109/38-500	1092	9/10/99
ORDER110/38-500	18	9/10/99
ORDER110/38-500	1092	9/11/99
ORDER111/38-500	560	9/12/99
ORDER112/38-545	200	9/12/99
ORDER114/38-545	202	9/12/99
ORDER112/38-545	350	9/13/99
ORDER113/38-500	560	9/13/99
ORDER115/38-500	18	9/13/99
ORDER114/38-545	348	9/13/99
ORDER115/38-500	1092	9/14/99
ORDER116/38-500	560	9/15/99
ORDER117/38-545	200	9/15/99
ORDER118/38-545	45	9/15/99
ORDER117/38-545	350	9/16/99
ORDER118/38-545	505	9/16/99
ORDER119/38-500	124	9/16/99
ORDER121/38-545	286	9/16/99
ORDER119/38-500	436	9/17/99
ORDER120/38-500	560	9/17/99
ORDER122/38-500	18	9/17/99
ORDER121/38-545	264	9/17/99
ORDER122/38-500	1092	9/18/99
ORDER124/38-545	550	9/19/99
ORDER123/38-500	124	9/19/99
ORDER127/38-500	236	9/19/99
ORDER123/38-500	436	9/20/99
ORDER125/38-500	560	9/20/99
ORDER126/38-545	200	9/20/99
ORDER128/38-545	45	9/20/99
ORDER127/38-500	35	9/20/99
ORDER126/38-545	350	9/21/99
ORDER128/38-545	505	9/21/99
ORDER129/38-500	18	9/21/99
ORDER127/38-500	289	9/21/99
ORDER129/38-500	1092	9/22/99
ORDER131/38-545	550	9/23/99
ORDER132/38-500	18	9/23/99
ORDER130/38-500	550	9/23/99
ORDER132/38-500	1092	9/24/99
ORDER133/38-500	18	9/24/99
ORDER133/38-500	1092	9/25/99
ORDER134/38-500	560	9/26/99
ORDER135/38-545	200	9/26/99
ORDER136/38-545	45	9/26/99
ORDER139/38-500	7	9/26/99
ORDER135/38-545	350	9/27/99
ORDER136/38-545	505	9/27/99
ORDER137/38-500	124	9/27/99
ORDER139/38-500	281	9/27/99
ORDER137/38-500	436	9/28/99
ORDER138/38-545	550	9/28/99
ORDER140/38-500	18	9/28/99
ORDER139/38-500	272	9/28/99
ORDER140/38-500	1092	9/29/99
ORDER141/38-500	18	9/29/99
ORDER141/38-500	1092	9/30/99
ORDER142/38-545	550	10/1/99
ORDER144/38-500	18	10/1/99
ORDER143/38-500	560	10/1/99
ORDER144/38-500	1092	10/2/99
ORDER145/38-500	18	10/2/99
ORDER145/38-500	1092	10/3/99

B.2 Material release schedule developed with DBR program: Med Downtimes

OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date
ORDER1/38-500	966	7/1/99	ORDER22/38-500	560	7/18/99	ORDER44/38-545	115	8/1/99
ORDER1/38-500	144	7/2/99	ORDER23/38-545	435	7/18/99	ORDER45/38-500	560	8/1/99
ORDER2/38-500	966	7/2/99	ORDER25/38-500	168	7/18/99	ORDER46/38-500	604	8/1/99
ORDER2/38-500	144	7/3/99	ORDER23/38-545	115	7/19/99	ORDER46/38-500	144	8/2/99
ORDER3/38-500	966	7/3/99	ORDER24/38-500	560	7/19/99	ORDER47/38-500	416	8/2/99
ORDER3/38-500	144	7/4/99	ORDER25/38-500	604	7/19/99	ORDER48/38-545	18	8/2/99
ORDER4/38-500	966	7/4/99	ORDER25/38-500	144	7/20/99	ORDER50/38-545	45	8/2/99
ORDER4/38-500	144	7/5/99	ORDER26/38-500	966	7/20/99	ORDER51/38-500	309	8/2/99
ORDER8/38-500	185	7/5/99	ORDER26/38-500	144	7/21/99	ORDER47/38-500	144	8/3/99
ORDER5/38-545	550	7/6/99	ORDER27/38-500	288	7/21/99	ORDER48/38-545	532	8/3/99
ORDER7/38-545	382	7/6/99	ORDER28/38-545	54	7/21/99	ORDER50/38-545	390	8/3/99
ORDER8/38-500	169	7/6/99	ORDER33/38-500	132	7/21/99	ORDER51/38-500	53	8/3/99
ORDER7/38-545	168	7/7/99	ORDER27/38-500	212	7/22/99	ORDER50/38-545	115	8/4/99
ORDER6/38-500	500	7/7/99	ORDER28/38-545	496	7/22/99	ORDER49/38-500	560	8/4/99
ORDER8/38-500	612	7/7/99	ORDER29/38-500	288	7/22/99	ORDER51/38-500	604	8/4/99
ORDER8/38-500	144	7/8/99	ORDER30/38-545	54	7/22/99	ORDER51/38-500	144	8/5/99
ORDER9/38-500	966	7/8/99	ORDER32/38-545	45	7/22/99	ORDER52/38-500	966	8/5/99
ORDER9/38-500	144	7/9/99	ORDER33/38-500	177	7/22/99	ORDER52/38-500	144	8/6/99
ORDER10/38-500	966	7/9/99	ORDER29/38-500	212	7/23/99	ORDER56/38-500	64	8/6/99
ORDER10/38-500	144	7/10/99	ORDER30/38-545	496	7/23/99	ORDER53/38-545	550	8/7/99
ORDER11/38-500	966	7/10/99	ORDER32/38-545	390	7/23/99	ORDER54/38-500	394	8/7/99
ORDER11/38-500	144	7/11/99	ORDER33/38-500	53	7/23/99	ORDER55/38-545	24	8/7/99
ORDER12/38-500	416	7/11/99	ORDER32/38-545	115	7/24/99	ORDER56/38-500	306	8/7/99
ORDER13/38-545	18	7/11/99	ORDER31/38-500	560	7/24/99	ORDER54/38-500	156	8/8/99
ORDER14/38-500	362	7/11/99	ORDER33/38-500	604	7/24/99	ORDER55/38-545	526	8/8/99
ORDER12/38-500	144	7/12/99	ORDER33/38-500	144	7/25/99	ORDER56/38-500	596	8/8/99
ORDER13/38-545	532	7/12/99	ORDER34/38-500	966	7/25/99	ORDER56/38-500	144	8/9/99
ORDER14/38-500	604	7/12/99	ORDER34/38-500	144	7/26/99	ORDER57/38-500	966	8/9/99
ORDER14/38-500	144	7/13/99	ORDER35/38-500	966	7/26/99	ORDER57/38-500	144	8/10/99
ORDER15/38-500	966	7/13/99	ORDER35/38-500	144	7/27/99	ORDER58/38-500	966	8/10/99
ORDER15/38-500	144	7/14/99	ORDER36/38-500	288	7/27/99	ORDER58/38-500	144	8/11/99
ORDER16/38-500	966	7/14/99	ORDER37/38-545	54	7/27/99	ORDER59/38-545	550	8/12/99
ORDER16/38-500	144	7/15/99	ORDER36/38-500	212	7/28/99	ORDER61/38-545	426	8/12/99
ORDER17/38-500	288	7/15/99	ORDER37/38-545	496	7/28/99	ORDER64/38-500	135	8/12/99
ORDER18/38-545	54	7/15/99	ORDER39/38-545	382	7/28/99	ORDER61/38-545	124	8/13/99
ORDER17/38-500	212	7/16/99	ORDER39/38-545	168	7/29/99	ORDER60/38-500	550	8/13/99
ORDER18/38-545	496	7/16/99	ORDER38/38-500	500	7/29/99	ORDER62/38-500	288	8/13/99
ORDER20/38-500	416	7/16/99	ORDER41/38-545	382	7/29/99	ORDER63/38-545	54	8/13/99
ORDER19/38-545	18	7/16/99	ORDER41/38-545	168	7/30/99	ORDER64/38-500	259	8/13/99
ORDER21/38-545	45	7/16/99	ORDER40/38-500	500	7/30/99	ORDER62/38-500	212	8/14/99
ORDER25/38-500	43	7/16/99	ORDER43/38-545	382	7/30/99	ORDER63/38-545	496	8/14/99
ORDER20/38-500	144	7/17/99	ORDER46/38-500	186	7/30/99	ORDER64/38-500	572	8/14/99
ORDER19/38-545	532	7/17/99	ORDER43/38-545	168	7/31/99	ORDER64/38-500	144	8/15/99
ORDER21/38-545	390	7/17/99	ORDER42/38-500	500	7/31/99	ORDER65/38-500	288	8/15/99
ORDER25/38-500	151	7/17/99	ORDER44/38-545	435	7/31/99	ORDER66/38-545	54	8/15/99
ORDER21/38-545	115	7/18/99	ORDER46/38-500	176	7/31/99	ORDER67/38-500	394	8/15/99

OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date
ORDER65/38-500	212	8/16/99	ORDER91/38-500	362	9/3/99	ORDER115/38-500	190	9/19/99
ORDER66/38-545	496	8/16/99	ORDER89/38-500	144	9/4/99	ORDER111/38-500	144	9/20/99
ORDER67/38-500	572	8/16/99	ORDER90/38-545	532	9/4/99	ORDER112/38-545	532	9/20/99
ORDER67/38-500	144	8/17/99	ORDER91/38-500	604	9/4/99	ORDER113/38-500	416	9/20/99
ORDER68/38-500	966	8/17/99	ORDER91/38-500	144	9/5/99	ORDER114/38-545	18	9/20/99
ORDER68/38-500	144	8/18/99	ORDER93/38-545	382	9/5/99	ORDER115/38-500	172	9/20/99
ORDER69/38-500	966	8/18/99	ORDER94/38-500	354	9/5/99	ORDER113/38-500	144	9/21/99
ORDER69/38-500	144	8/19/99	ORDER93/38-545	168	9/6/99	ORDER114/38-545	532	9/21/99
ORDER70/38-500	416	8/19/99	ORDER92/38-500	500	9/6/99	ORDER115/38-500	604	9/21/99
ORDER71/38-545	18	8/19/99	ORDER94/38-500	612	9/6/99	ORDER115/38-500	144	9/22/99
ORDER70/38-500	144	8/20/99	ORDER94/38-500	144	9/7/99	ORDER116/38-500	416	9/22/99
ORDER71/38-545	532	8/20/99	ORDER95/38-500	394	9/7/99	ORDER117/38-545	18	9/22/99
ORDER73/38-500	416	8/20/99	ORDER96/38-545	24	9/7/99	ORDER118/38-545	45	9/22/99
ORDER72/38-545	18	8/20/99	ORDER97/38-500	370	9/7/99	ORDER116/38-500	144	9/23/99
ORDER73/38-500	144	8/21/99	ORDER95/38-500	156	9/8/99	ORDER117/38-545	532	9/23/99
ORDER72/38-545	532	8/21/99	ORDER96/38-545	526	9/8/99	ORDER118/38-545	390	9/23/99
ORDER75/38-545	382	8/21/99	ORDER97/38-500	596	9/8/99	ORDER122/38-500	190	9/23/99
ORDER78/38-500	178	8/21/99	ORDER97/38-500	144	9/9/99	ORDER118/38-545	115	9/24/99
ORDER75/38-545	168	8/22/99	ORDER98/38-500	966	9/9/99	ORDER119/38-500	560	9/24/99
ORDER74/38-500	500	8/22/99	ORDER98/38-500	144	9/10/99	ORDER120/38-500	416	9/24/99
ORDER77/38-500	394	8/22/99	ORDER99/38-500	416	9/10/99	ORDER121/38-545	18	9/24/99
ORDER76/38-545	24	8/22/99	ORDER100/38-545	18	9/10/99	ORDER122/38-500	172	9/24/99
ORDER78/38-500	192	8/22/99	ORDER101/38-500	362	9/10/99	ORDER120/38-500	144	9/25/99
ORDER77/38-500	156	8/23/99	ORDER99/38-500	144	9/11/99	ORDER121/38-545	532	9/25/99
ORDER76/38-545	526	8/23/99	ORDER100/38-545	532	9/11/99	ORDER122/38-500	604	9/25/99
ORDER78/38-500	596	8/23/99	ORDER101/38-500	604	9/11/99	ORDER122/38-500	144	9/26/99
ORDER78/38-500	144	8/24/99	ORDER101/38-500	144	9/12/99	ORDER124/38-545	435	9/26/99
ORDER80/38-500	238	8/24/99	ORDER102/38-500	966	9/12/99	ORDER129/38-500	186	9/26/99
ORDER79/38-545	550	8/25/99	ORDER102/38-500	144	9/13/99	ORDER124/38-545	115	9/27/99
ORDER80/38-500	728	8/25/99	ORDER103/38-500	966	9/13/99	ORDER123/38-500	560	9/27/99
ORDER80/38-500	144	8/26/99	ORDER103/38-500	144	9/14/99	ORDER125/38-500	416	9/27/99
ORDER81/38-500	966	8/26/99	ORDER104/38-500	416	9/14/99	ORDER126/38-545	18	9/27/99
ORDER81/38-500	144	8/27/99	ORDER105/38-545	18	9/14/99	ORDER128/38-545	45	9/27/99
ORDER82/38-500	966	8/27/99	ORDER108/38-500	135	9/14/99	ORDER129/38-500	123	9/27/99
ORDER82/38-500	144	8/28/99	ORDER104/38-500	144	9/15/99	ORDER125/38-500	144	9/28/99
ORDER83/38-500	966	8/28/99	ORDER105/38-545	532	9/15/99	ORDER126/38-545	532	9/28/99
ORDER83/38-500	144	8/29/99	ORDER106/38-500	288	9/15/99	ORDER128/38-545	390	9/28/99
ORDER84/38-500	966	8/29/99	ORDER107/38-545	54	9/15/99	ORDER129/38-500	53	9/28/99
ORDER84/38-500	144	8/30/99	ORDER108/38-500	259	9/15/99	ORDER128/38-545	115	9/29/99
ORDER86/38-500	238	8/30/99	ORDER106/38-500	212	9/16/99	ORDER127/38-500	560	9/29/99
ORDER85/38-545	550	8/31/99	ORDER107/38-545	496	9/16/99	ORDER129/38-500	604	9/29/99
ORDER86/38-500	728	8/31/99	ORDER108/38-500	572	9/16/99	ORDER129/38-500	144	9/30/99
ORDER86/38-500	144	9/1/99	ORDER108/38-500	144	9/17/99	ORDER131/38-545	426	9/30/99
ORDER87/38-500	966	9/1/99	ORDER109/38-500	966	9/17/99	ORDER132/38-500	362	9/30/99
ORDER87/38-500	144	9/2/99	ORDER109/38-500	144	9/18/99	ORDER131/38-545	124	10/1/99
ORDER88/38-500	966	9/2/99	ORDER110/38-500	966	9/18/99	ORDER130/38-500	550	10/1/99
ORDER88/38-500	144	9/3/99	ORDER110/38-500	144	9/19/99	ORDER132/38-500	604	10/1/99
ORDER89/38-500	416	9/3/99	ORDER111/38-500	416	9/19/99	ORDER132/38-500	144	10/2/99
ORDER90/38-545	18	9/3/99	ORDER112/38-545	18	9/19/99	ORDER133/38-500	966	10/2/99

Stellenbosch University http://scholar.sun.ac.za

OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date
ORDER133/38-500	144	10/3/99						
ORDER134/38-500	416	10/3/99						
ORDER135/38-545	18	10/3/99						
ORDER136/38-545	45	10/3/99						
ORDER140/38-500	43	10/3/99						
ORDER134/38-500	144	10/4/99						
ORDER135/38-545	532	10/4/99						
ORDER136/38-545	390	10/4/99						
ORDER140/38-500	151	10/4/99						
ORDER136/38-545	115	10/5/99						
ORDER137/38-500	560	10/5/99						
ORDER138/38-545	435	10/5/99						
ORDER140/38-500	168	10/5/99						
ORDER138/38-545	115	10/6/99						
ORDER139/38-500	560	10/6/99						
ORDER140/38-500	604	10/6/99						
ORDER140/38-500	144	10/7/99						
ORDER141/38-500	966	10/7/99						
ORDER141/38-500	144	10/8/99						
ORDER142/38-545	435	10/8/99						
ORDER144/38-500	362	10/8/99						
ORDER142/38-545	115	10/9/99						
ORDER143/38-500	560	10/9/99						
ORDER144/38-500	604	10/9/99						
ORDER144/38-500	144	10/10/99						
ORDER145/38-500	966	10/10/99						
ORDER145/38-500	144	10/11/99						

B.3 Material release schedule developed with DBR program: Low Downtimes

OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date
ORDER1/38-500	202	7/3/99	ORDER32/38-545	550	7/26/99	ORDER64/38-500	908	8/17/99
ORDER1/38-500	908	7/4/99	ORDER33/38-500	202	7/26/99	ORDER65/38-500	500	8/18/99
ORDER2/38-500	202	7/4/99	ORDER31/38-500	524	7/26/99	ORDER67/38-500	202	8/18/99
ORDER2/38-500	908	7/5/99	ORDER33/38-500	908	7/27/99	ORDER66/38-545	550	8/18/99
ORDER3/38-500	202	7/5/99	ORDER34/38-500	202	7/27/99	ORDER67/38-500	908	8/19/99
ORDER3/38-500	908	7/6/99	ORDER34/38-500	908	7/28/99	ORDER68/38-500	202	8/19/99
ORDER4/38-500	202	7/6/99	ORDER35/38-500	202	7/28/99	ORDER68/38-500	908	8/20/99
ORDER4/38-500	908	7/7/99	ORDER35/38-500	908	7/29/99	ORDER69/38-500	202	8/20/99
ORDER5/38-545	550	7/8/99	ORDER36/38-500	500	7/30/99	ORDER69/38-500	908	8/21/99
ORDER7/38-545	550	7/9/99	ORDER37/38-545	550	7/30/99	ORDER70/38-500	560	8/22/99
ORDER8/38-500	202	7/9/99	ORDER39/38-545	550	7/31/99	ORDER71/38-545	550	8/22/99
ORDER6/38-500	500	7/9/99	ORDER38/38-500	500	7/31/99	ORDER73/38-500	560	8/23/99
ORDER8/38-500	908	7/10/99	ORDER41/38-545	550	8/1/99	ORDER72/38-545	550	8/23/99
ORDER9/38-500	202	7/10/99	ORDER40/38-500	500	8/1/99	ORDER75/38-545	550	8/24/99
ORDER9/38-500	908	7/11/99	ORDER43/38-545	550	8/2/99	ORDER74/38-500	500	8/24/99
ORDER10/38-500	202	7/11/99	ORDER42/38-500	500	8/2/99	ORDER76/38-545	26	8/24/99
ORDER10/38-500	908	7/12/99	ORDER45/38-500	36	8/2/99	ORDER77/38-500	550	8/25/99
ORDER11/38-500	202	7/12/99	ORDER44/38-545	550	8/3/99	ORDER78/38-500	202	8/25/99
ORDER11/38-500	908	7/13/99	ORDER46/38-500	202	8/3/99	ORDER76/38-545	524	8/25/99
ORDER13/38-545	34	7/13/99	ORDER45/38-500	524	8/3/99	ORDER78/38-500	908	8/26/99
ORDER12/38-500	560	7/14/99	ORDER46/38-500	908	8/4/99	ORDER80/38-500	202	8/27/99
ORDER14/38-500	202	7/14/99	ORDER47/38-500	560	8/5/99	ORDER79/38-545	550	8/27/99
ORDER13/38-545	516	7/14/99	ORDER48/38-545	550	8/5/99	ORDER80/38-500	908	8/28/99
ORDER14/38-500	908	7/15/99	ORDER49/38-500	36	8/5/99	ORDER81/38-500	202	8/28/99
ORDER15/38-500	202	7/15/99	ORDER50/38-545	550	8/6/99	ORDER81/38-500	908	8/29/99
ORDER15/38-500	908	7/16/99	ORDER51/38-500	202	8/6/99	ORDER82/38-500	202	8/29/99
ORDER16/38-500	202	7/16/99	ORDER49/38-500	524	8/6/99	ORDER82/38-500	908	8/30/99
ORDER16/38-500	908	7/17/99	ORDER51/38-500	908	8/7/99	ORDER83/38-500	202	8/30/99
ORDER17/38-500	500	7/18/99	ORDER52/38-500	202	8/7/99	ORDER83/38-500	908	8/31/99
ORDER18/38-545	550	7/18/99	ORDER52/38-500	908	8/8/99	ORDER84/38-500	202	8/31/99
ORDER20/38-500	560	7/19/99	ORDER53/38-545	550	8/9/99	ORDER84/38-500	908	9/1/99
ORDER19/38-545	550	7/19/99	ORDER55/38-545	26	8/9/99	ORDER86/38-500	202	9/2/99
ORDER21/38-545	550	7/20/99	ORDER54/38-500	550	8/10/99	ORDER85/38-545	550	9/2/99
ORDER22/38-500	560	7/20/99	ORDER56/38-500	202	8/10/99	ORDER86/38-500	908	9/3/99
ORDER24/38-500	36	7/20/99	ORDER55/38-545	524	8/10/99	ORDER87/38-500	202	9/3/99
ORDER23/38-545	550	7/21/99	ORDER56/38-500	908	8/11/99	ORDER87/38-500	908	9/4/99
ORDER25/38-500	202	7/21/99	ORDER57/38-500	202	8/11/99	ORDER88/38-500	202	9/4/99
ORDER24/38-500	524	7/21/99	ORDER57/38-500	908	8/12/99	ORDER88/38-500	908	9/5/99
ORDER25/38-500	908	7/22/99	ORDER58/38-500	202	8/12/99	ORDER90/38-545	34	9/5/99
ORDER26/38-500	202	7/22/99	ORDER58/38-500	908	8/13/99	ORDER89/38-500	560	9/6/99
ORDER26/38-500	908	7/23/99	ORDER59/38-545	550	8/14/99	ORDER91/38-500	202	9/6/99
ORDER27/38-500	500	7/24/99	ORDER61/38-545	550	8/15/99	ORDER90/38-545	516	9/6/99
ORDER28/38-545	550	7/24/99	ORDER60/38-500	550	8/15/99	ORDER91/38-500	908	9/7/99
ORDER29/38-500	500	7/25/99	ORDER62/38-500	500	8/16/99	ORDER93/38-545	550	9/8/99
ORDER30/38-545	550	7/25/99	ORDER64/38-500	202	8/16/99	ORDER94/38-500	202	9/8/99
ORDER31/38-500	36	7/25/99	ORDER63/38-545	550	8/16/99	ORDER92/38-500	500	9/8/99

Stellenbosch University http://scholar.sun.ac.za

OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date	OrderID/ProdID	Rel Qty	Rel Date
ORDER94/38-500	908	9/9/99	ORDER129/38-500	202	10/1/99			
ORDER96/38-545	26	9/9/99	ORDER127/38-500	524	10/1/99			
ORDER95/38-500	550	9/10/99	ORDER129/38-500	908	10/2/99			
ORDER97/38-500	202	9/10/99	ORDER130/38-500	26	10/2/99			
ORDER96/38-545	524	9/10/99	ORDER131/38-545	550	10/3/99			
ORDER97/38-500	908	9/11/99	ORDER132/38-500	202	10/3/99			
ORDER98/38-500	202	9/11/99	ORDER130/38-500	524	10/3/99			
ORDER98/38-500	908	9/12/99	ORDER132/38-500	908	10/4/99			
ORDER100/38-545	34	9/12/99	ORDER133/38-500	202	10/4/99			
ORDER99/38-500	560	9/13/99	ORDER133/38-500	908	10/5/99			
ORDER101/38-500	202	9/13/99	ORDER134/38-500	560	10/6/99			
ORDER100/38-545	516	9/13/99	ORDER135/38-545	550	10/6/99			
ORDER101/38-500	908	9/14/99	ORDER136/38-545	550	10/7/99			
ORDER102/38-500	202	9/14/99	ORDER137/38-500	560	10/7/99			
ORDER102/38-500	908	9/15/99	ORDER139/38-500	36	10/7/99			
ORDER103/38-500	202	9/15/99	ORDER138/38-545	550	10/8/99			
ORDER103/38-500	908	9/16/99	ORDER140/38-500	202	10/8/99			
ORDER104/38-500	560	9/17/99	ORDER139/38-500	524	10/8/99			
ORDER105/38-545	550	9/17/99	ORDER140/38-500	908	10/9/99			
ORDER106/38-500	500	9/18/99	ORDER141/38-500	202	10/9/99			
ORDER108/38-500	202	9/18/99	ORDER141/38-500	908	10/10/99			
ORDER107/38-545	550	9/18/99	ORDER143/38-500	36	10/10/99			
ORDER108/38-500	908	9/19/99	ORDER142/38-545	550	10/11/99			
ORDER109/38-500	202	9/19/99	ORDER144/38-500	202	10/11/99			
ORDER109/38-500	908	9/20/99	ORDER143/38-500	524	10/11/99			
ORDER110/38-500	202	9/20/99	ORDER144/38-500	908	10/12/99			
ORDER110/38-500	908	9/21/99	ORDER145/38-500	202	10/12/99			
ORDER111/38-500	560	9/22/99	ORDER145/38-500	908	10/13/99			
ORDER112/38-545	550	9/22/99						
ORDER114/38-545	34	9/22/99						
ORDER113/38-500	560	9/23/99						
ORDER115/38-500	202	9/23/99						
ORDER114/38-545	516	9/23/99						
ORDER115/38-500	908	9/24/99						
ORDER116/38-500	560	9/25/99						
ORDER117/38-545	550	9/25/99						
ORDER118/38-545	550	9/26/99						
ORDER119/38-500	560	9/26/99						
ORDER121/38-545	34	9/26/99						
ORDER120/38-500	560	9/27/99						
ORDER122/38-500	202	9/27/99						
ORDER121/38-545	516	9/27/99						
ORDER122/38-500	908	9/28/99						
ORDER124/38-545	550	9/29/99						
ORDER123/38-500	560	9/29/99						
ORDER125/38-500	560	9/30/99						
ORDER126/38-545	550	9/30/99						
ORDER127/38-500	36	9/30/99						
ORDER128/38-545	550	10/1/99						

B.4 CCR schedule developed with DBR program

OrderID/ProdID	Part/Opnr	Qty	Start date	Start Time	Fin date	Fin Time
ORDER1/38-500	38-500/11	1110	7/4/99	11:50	7/5/99	11:42
ORDER2/38-500	38-500/11	1110	7/5/99	11:50	7/6/99	11:42
ORDER3/38-500	38-500/11	1110	7/6/99	11:50	7/7/99	11:42
ORDER4/38-500	38-500/11	1110	7/7/99	11:50	7/8/99	11:42
ORDER5/38-545	38-545/12	550	7/8/99	15:48	7/9/99	11:42
ORDER7/38-545	38-545/12	550	7/9/99	12:16	7/10/99	8:09
ORDER6/38-500	38-500/11	500	7/10/99	8:09	7/10/99	11:42
ORDER8/38-500	38-500/11	1110	7/10/99	11:50	7/11/99	11:42
ORDER9/38-500	38-500/11	1110	7/11/99	11:50	7/12/99	11:42
ORDER10/38-500	38-500/11	1110	7/12/99	11:50	7/13/99	11:42
ORDER11/38-500	38-500/11	1110	7/13/99	11:50	7/14/99	11:42
ORDER12/38-500	38-500/11	560	7/14/99	11:50	7/14/99	15:48
ORDER13/38-545	38-545/12	550	7/14/99	15:48	7/15/99	11:42
ORDER14/38-500	38-500/11	1110	7/15/99	11:50	7/16/99	11:42
ORDER15/38-500	38-500/11	1110	7/16/99	11:50	7/17/99	11:42
ORDER16/38-500	38-500/11	1110	7/17/99	11:50	7/18/99	11:42
ORDER17/38-500	38-500/11	500	7/18/99	12:16	7/18/99	15:48
ORDER18/38-545	38-545/12	550	7/18/99	15:48	7/19/99	11:42
ORDER20/38-500	38-500/11	560	7/19/99	11:50	7/19/99	15:48
ORDER19/38-545	38-545/12	550	7/19/99	15:48	7/20/99	11:42
ORDER21/38-545	38-545/12	550	7/20/99	11:50	7/20/99	15:44
ORDER22/38-500	38-500/11	560	7/20/99	15:44	7/21/99	11:42
ORDER23/38-545	38-545/12	550	7/21/99	11:50	7/21/99	15:44
ORDER24/38-500	38-500/11	560	7/21/99	15:44	7/22/99	11:42
ORDER25/38-500	38-500/11	1110	7/22/99	11:50	7/23/99	11:42
ORDER26/38-500	38-500/11	1110	7/23/99	11:50	7/24/99	11:42
ORDER27/38-500	38-500/11	500	7/24/99	12:16	7/24/99	15:48
ORDER28/38-545	38-545/12	550	7/24/99	15:48	7/25/99	11:42
ORDER29/38-500	38-500/11	500	7/25/99	12:16	7/25/99	15:48
ORDER30/38-545	38-545/12	550	7/25/99	15:48	7/26/99	11:42
ORDER32/38-545	38-545/12	550	7/26/99	11:50	7/26/99	15:44
ORDER31/38-500	38-500/11	560	7/26/99	15:44	7/27/99	11:42
ORDER33/38-500	38-500/11	1110	7/27/99	11:50	7/28/99	11:42
ORDER34/38-500	38-500/11	1110	7/28/99	11:50	7/29/99	11:42
ORDER35/38-500	38-500/11	1110	7/29/99	11:50	7/30/99	11:42
ORDER36/38-500	38-500/11	500	7/30/99	12:16	7/30/99	15:48
ORDER37/38-545	38-545/12	550	7/30/99	15:48	7/31/99	11:42
ORDER39/38-545	38-545/12	550	7/31/99	12:16	8/1/99	8:09
ORDER38/38-500	38-500/11	500	8/1/99	8:09	8/1/99	11:42
ORDER41/38-545	38-545/12	550	8/1/99	12:16	8/2/99	8:09
ORDER40/38-500	38-500/11	500	8/2/99	8:09	8/2/99	11:42
ORDER43/38-545	38-545/12	550	8/2/99	12:16	8/3/99	8:09
ORDER42/38-500	38-500/11	500	8/3/99	8:09	8/3/99	11:42
ORDER44/38-545	38-545/12	550	8/3/99	11:50	8/3/99	15:44
ORDER45/38-500	38-500/11	560	8/3/99	15:44	8/4/99	11:42
ORDER46/38-500	38-500/11	1110	8/4/99	11:50	8/5/99	11:42
ORDER47/38-500	38-500/11	560	8/5/99	11:50	8/5/99	15:48
ORDER48/38-545	38-545/12	550	8/5/99	15:48	8/6/99	11:42
ORDER50/38-545	38-545/12	550	8/6/99	11:50	8/6/99	15:44
ORDER49/38-500	38-500/11	560	8/6/99	15:44	8/7/99	11:42

OrderID/ProdID	Part/Opnr	Qty	Start date	Start Time	Fin date	Fin Time
ORDER51/38-500	38-500/11	1110	8/7/99	11:50	8/8/99	11:42
ORDER52/38-500	38-500/11	1110	8/8/99	11:50	8/9/99	11:42
ORDER53/38-545	38-545/12	550	8/9/99	15:48	8/10/99	11:42
ORDER54/38-500	38-500/11	550	8/10/99	11:55	8/10/99	15:48
ORDER55/38-545	38-545/12	550	8/10/99	15:48	8/11/99	11:42
ORDER56/38-500	38-500/11	1110	8/11/99	11:50	8/12/99	11:42
ORDER57/38-500	38-500/11	1110	8/12/99	11:50	8/13/99	11:42
ORDER58/38-500	38-500/11	1110	8/13/99	11:50	8/14/99	11:42
ORDER59/38-545	38-545/12	550	8/14/99	15:48	8/15/99	11:42
ORDER61/38-545	38-545/12	550	8/15/99	11:55	8/15/99	15:48
ORDER60/38-500	38-500/11	550	8/15/99	15:48	8/16/99	11:42
ORDER62/38-500	38-500/11	500	8/16/99	12:16	8/16/99	15:48
ORDER63/38-545	38-545/12	550	8/16/99	15:48	8/17/99	11:42
ORDER64/38-500	38-500/11	1110	8/17/99	11:50	8/18/99	11:42
ORDER65/38-500	38-500/11	500	8/18/99	12:16	8/18/99	15:48
ORDER66/38-545	38-545/12	550	8/18/99	15:48	8/19/99	11:42
ORDER67/38-500	38-500/11	1110	8/19/99	11:50	8/20/99	11:42
ORDER68/38-500	38-500/11	1110	8/20/99	11:50	8/21/99	11:42
ORDER69/38-500	38-500/11	1110	8/21/99	11:50	8/22/99	11:42
ORDER70/38-500	38-500/11	560	8/22/99	11:50	8/22/99	15:48
ORDER71/38-545	38-545/12	550	8/22/99	15:48	8/23/99	11:42
ORDER73/38-500	38-500/11	560	8/23/99	11:50	8/23/99	15:48
ORDER72/38-545	38-545/12	550	8/23/99	15:48	8/24/99	11:42
ORDER75/38-545	38-545/12	550	8/24/99	12:16	8/25/99	8:09
ORDER74/38-500	38-500/11	500	8/25/99	8:09	8/25/99	11:42
ORDER77/38-500	38-500/11	550	8/25/99	11:55	8/25/99	15:48
ORDER76/38-545	38-545/12	550	8/25/99	15:48	8/26/99	11:42
ORDER78/38-500	38-500/11	1110	8/26/99	11:50	8/27/99	11:42
ORDER79/38-545	38-545/12	550	8/27/99	15:48	8/28/99	11:42
ORDER80/38-500	38-500/11	1110	8/28/99	11:50	8/29/99	11:42
ORDER81/38-500	38-500/11	1110	8/29/99	11:50	8/30/99	11:42
ORDER82/38-500	38-500/11	1110	8/30/99	11:50	8/31/99	11:42
ORDER83/38-500	38-500/11	1110	8/31/99	11:50	9/1/99	11:42
ORDER84/38-500	38-500/11	1110	9/1/99	11:50	9/2/99	11:42
ORDER85/38-545	38-545/12	550	9/2/99	15:48	9/3/99	11:42
ORDER86/38-500	38-500/11	1110	9/3/99	11:50	9/4/99	11:42
ORDER87/38-500	38-500/11	1110	9/4/99	11:50	9/5/99	11:42
ORDER88/38-500	38-500/11	1110	9/5/99	11:50	9/6/99	11:42
ORDER89/38-500	38-500/11	560	9/6/99	11:50	9/6/99	15:48
ORDER90/38-545	38-545/12	550	9/6/99	15:48	9/7/99	11:42
ORDER91/38-500	38-500/11	1110	9/7/99	11:50	9/8/99	11:42
ORDER93/38-545	38-545/12	550	9/8/99	12:16	9/9/99	8:09
ORDER92/38-500	38-500/11	500	9/9/99	8:09	9/9/99	11:42
ORDER94/38-500	38-500/11	1110	9/9/99	11:50	9/10/99	11:42
ORDER95/38-500	38-500/11	550	9/10/99	11:55	9/10/99	15:48
ORDER96/38-545	38-545/12	550	9/10/99	15:48	9/11/99	11:42
ORDER97/38-500	38-500/11	1110	9/11/99	11:50	9/12/99	11:42
ORDER98/38-500	38-500/11	1110	9/12/99	11:50	9/13/99	11:42
ORDER99/38-500	38-500/11	560	9/13/99	11:50	9/13/99	15:48
ORDER100/38-545	38-545/12	550	9/13/99	15:48	9/14/99	11:42
ORDER101/38-500	38-500/11	1110	9/14/99	11:50	9/15/99	11:42
ORDER102/38-500	38-500/11	1110	9/15/99	11:50	9/16/99	11:42
ORDER103/38-500	38-500/11	1110	9/16/99	11:50	9/17/99	11:42

Stellenbosch University http://scholar.sun.ac.za

OrderID/ProdID	Part/Opnr	Qty	Start date	Start Time	Fin date	Fin Time
ORDER104/38-500	38-500/11	560	9/17/99	11:50	9/17/99	15:48
ORDER105/38-545	38-545/12	550	9/17/99	15:48	9/18/99	11:42
ORDER106/38-500	38-500/11	500	9/18/99	12:16	9/18/99	15:48
ORDER107/38-545	38-545/12	550	9/18/99	15:48	9/19/99	11:42
ORDER108/38-500	38-500/11	1110	9/19/99	11:50	9/20/99	11:42
ORDER109/38-500	38-500/11	1110	9/20/99	11:50	9/21/99	11:42
ORDER110/38-500	38-500/11	1110	9/21/99	11:50	9/22/99	11:42
ORDER111/38-500	38-500/11	560	9/22/99	11:50	9/22/99	15:48
ORDER112/38-545	38-545/12	550	9/22/99	15:48	9/23/99	11:42
ORDER113/38-500	38-500/11	560	9/23/99	11:50	9/23/99	15:48
ORDER114/38-545	38-545/12	550	9/23/99	15:48	9/24/99	11:42
ORDER115/38-500	38-500/11	1110	9/24/99	11:50	9/25/99	11:42
ORDER116/38-500	38-500/11	560	9/25/99	11:50	9/25/99	15:48
ORDER117/38-545	38-545/12	550	9/25/99	15:48	9/26/99	11:42
ORDER118/38-545	38-545/12	550	9/26/99	11:50	9/26/99	15:44
ORDER119/38-500	38-500/11	560	9/26/99	15:44	9/27/99	11:42
ORDER120/38-500	38-500/11	560	9/27/99	11:50	9/27/99	15:48
ORDER121/38-545	38-545/12	550	9/27/99	15:48	9/28/99	11:42
ORDER122/38-500	38-500/11	1110	9/28/99	11:50	9/29/99	11:42
ORDER124/38-545	38-545/12	550	9/29/99	11:50	9/29/99	15:44
ORDER123/38-500	38-500/11	560	9/29/99	15:44	9/30/99	11:42
ORDER125/38-500	38-500/11	560	9/30/99	11:50	9/30/99	15:48
ORDER126/38-545	38-545/12	550	9/30/99	15:48	10/1/99	11:42
ORDER128/38-545	38-545/12	550	10/1/99	11:50	10/1/99	15:44
ORDER127/38-500	38-500/11	560	10/1/99	15:44	10/2/99	11:42
ORDER129/38-500	38-500/11	1110	10/2/99	11:50	10/3/99	11:42
ORDER131/38-545	38-545/12	550	10/3/99	11:55	10/3/99	15:48
ORDER130/38-500	38-500/11	550	10/3/99	15:48	10/4/99	11:42
ORDER132/38-500	38-500/11	1110	10/4/99	11:50	10/5/99	11:42
ORDER133/38-500	38-500/11	1110	10/5/99	11:50	10/6/99	11:42
ORDER134/38-500	38-500/11	560	10/6/99	11:50	10/6/99	15:48
ORDER135/38-545	38-545/12	550	10/6/99	15:48	10/7/99	11:42
ORDER136/38-545	38-545/12	550	10/7/99	11:50	10/7/99	15:44
ORDER137/38-500	38-500/11	560	10/7/99	15:44	10/8/99	11:42
ORDER138/38-545	38-545/12	550	10/8/99	11:50	10/8/99	15:44
ORDER139/38-500	38-500/11	560	10/8/99	15:44	10/9/99	11:42
ORDER140/38-500	38-500/11	1110	10/9/99	11:50	10/10/99	11:42
ORDER141/38-500	38-500/11	1110	10/10/99	11:50	10/11/99	11:42
ORDER142/38-545	38-545/12	550	10/11/99	11:50	10/11/99	15:44
ORDER143/38-500	38-500/11	560	10/11/99	15:44	10/12/99	11:42
ORDER144/38-500	38-500/11	1110	10/12/99	11:50	10/13/99	11:42
ORDER145/38-500	38-500/11	1110	10/13/99	11:50	10/14/99	11:42

Appendix C

Developed Drum-Buffer-Rope Scheduling Software Program

This appendix presents the drum-buffer-rope (DBR) scheduling software program developed in Visual Basic 6.0. The program developed is largely based on the system described in the book "The Haystack Syndrome" by Goldratt (1990). It applies the techniques of identifying the constraints, "exploiting" or scheduling the constraints, and subordinating other resources to the constraint(s) schedule(s). The overall structure and logic flow of the developed DBR system is illustrated in Figure C- 1. This flow diagram presents the basic structure of the DBR system from a functional point of view. From this diagram it can be seen that the system is composed of four basic functions:

- Database function.
- Scheduling function.
- Printable reports.
- Graphical user interface.

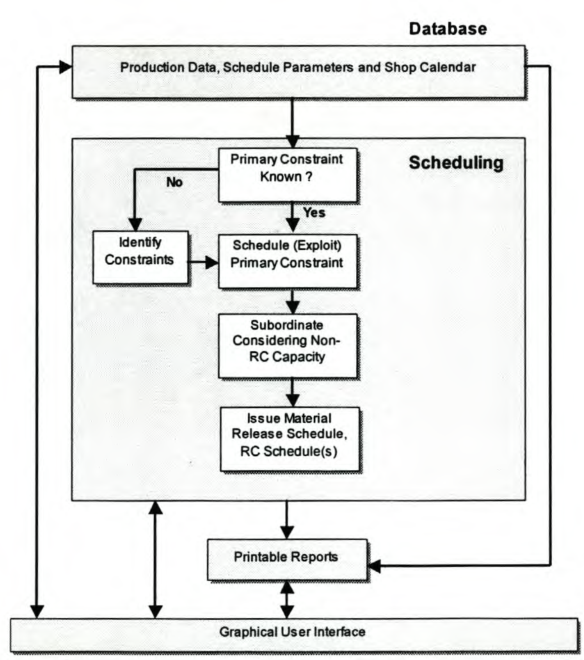

Figure C- 1. DBR scheduling program functional flow diagram

C.1 Database Function

An internal database was created using Visual Basic's database engine. This is the same engine that powers Microsoft Access. Microsoft Access can therefore also be used as the interface to open the scheduling program's database. This broadens the database's functionality, e.g. Access's "Data Import" functions could be used to import manufacturing data from other database formats. The database is used to store the manufacturing data, the scheduling parameters and the shop floor calendar. The following data tables are contained in the database to be used during the scheduling procedure:

- *Customer orders* - This table stores the information regarding the market demand. It basically tells which product is needed, when and how much.
- *Product data* - This table contains the ID's and descriptions of the different end products.
- *Bill of materials (BOM)* - The purpose of this table is to show the individual components of the different products, and how they relate to each other. It is an assembly BOM that shows how a product is put together. It contains the information to identify each item and the quantity used per unit of the item of which it is part. The BOM is structured in a single-level explosion, i.e. each item or component is listed showing only its parent and the number of units needed per unit of its parent.
- *Routings* - This table contains data on the flow of the products through manufacturing. It shows for each component the different manufacturing steps or operations that it has to go through, together with data such as the resource used at each step, the processing times and the set-up times.
- *Work in process (WIP)* - The WIP table identifies and quantifies the inventory that is on the shop floor.
- *Balance on hand (end-product inventory)* - The Balance on Hand table specifies the quantity of each end product currently in finished goods inventory.
- *Resource data* - The Resource Data table contains information on the manufacturing resources used in the production process. The information in this table is used to model the resources on the shop floor. The scheduling program can model batch processing resources, as well as different working time schedules to be assigned to different resources.
- *Raw material data* - The Raw Material data table is used to keep information on the raw material stocks.
- *Calendar* - The different shop floor calendars are stored in the Calendar table. Each calendar is identified by a unique calendar number. The calendars describe the working hours of the resources and can be used to model different working patterns or working shifts for different resources.

Stellenbosch University http://scholar.sun.ac.za

- *Scheduling parameters* – Contains scheduling parameters such as the scheduling period start date and end date.
- *Time buffers* – Contains the different types of time buffers present and their lengths.

Example screenshots of the BOM and Routing tables are presented in Figure C- 2 and Figure C- 3.

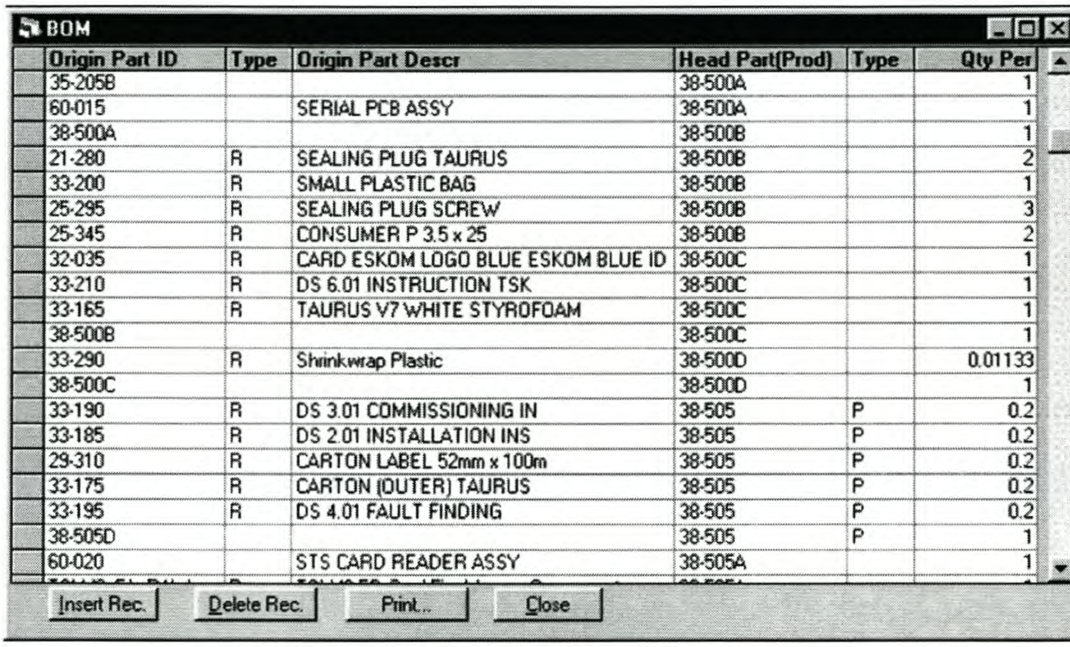

BOM

Origin Part ID	Type	Origin Part Descr	Head Part(Prod)	Type	Qty Per
35-205B			38-500A		1
60-015		SERIAL PCB ASSY	38-500A		1
38-500A			38-500B		1
21-280	R	SEALING PLUG TAURUS	38-500B		2
33-200	R	SMALL PLASTIC BAG	38-500B		1
25-295	R	SEALING PLUG SCREW	38-500B		3
25-345	R	CONSUMER P 3.5 x 25	38-500B		2
32-035	R	CARD ESKOM LOGO BLUE ESKOM BLUE ID	38-500C		1
33-210	R	DS 6.01 INSTRUCTION TSK	38-500C		1
33-165	R	TAURUS V7 WHITE STYROFOAM	38-500C		1
38-500B			38-500C		1
33-290	R	Shrinkwrap Plastic	38-500D		0.01133
38-500C			38-500D		1
33-190	R	DS 3.01 COMMISSIONING IN	38-505	P	0.2
33-185	R	DS 2.01 INSTALLATION INS	38-505	P	0.2
29-310	R	CARTON LABEL 52mm x 100m	38-505	P	0.2
33-175	R	CARTON (OUTER) TAURUS	38-505	P	0.2
33-195	R	DS 4.01 FAULT FINDING	38-505	P	0.2
38-505D			38-505	P	1
60-020		STS CARD READER ASSY	38-505A		1

Insert Rec. | Delete Rec. | Print... | Close

Figure C- 2. BOM table

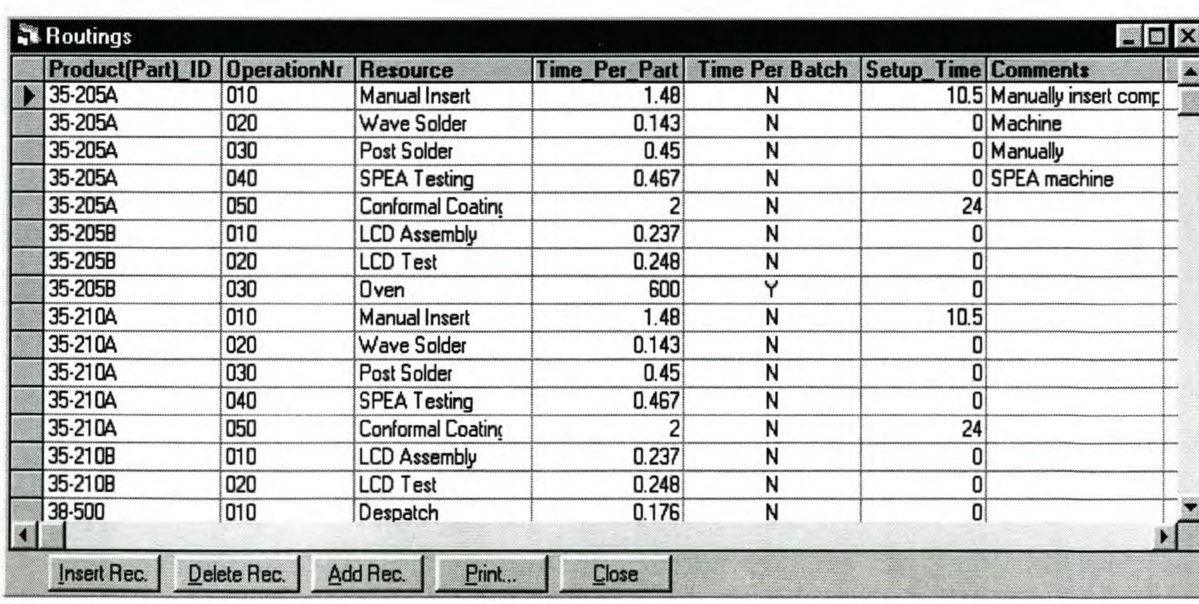

Routings

Product(Part)_ID	OperationNr	Resource	Time_Per_Part	Time Per Batch	Setup_Time	Comments
35-205A	010	Manual Insert	1.48	N	10.5	Manually insert comp
35-205A	020	Wave Solder	0.143	N	0	Machine
35-205A	030	Post Solder	0.45	N	0	Manually
35-205A	040	SPEA Testing	0.467	N	0	SPEA machine
35-205A	050	Conformal Coating	2	N	24	
35-205B	010	LCD Assembly	0.237	N	0	
35-205B	020	LCD Test	0.248	N	0	
35-205B	030	Oven	600	Y	0	
35-210A	010	Manual Insert	1.48	N	10.5	
35-210A	020	Wave Solder	0.143	N	0	
35-210A	030	Post Solder	0.45	N	0	
35-210A	040	SPEA Testing	0.467	N	0	
35-210A	050	Conformal Coating	2	N	24	
35-210B	010	LCD Assembly	0.237	N	0	
35-210B	020	LCD Test	0.248	N	0	
38-500	010	Despatch	0.176	N	0	

Insert Rec. | Delete Rec. | Add Rec. | Print... | Close

Figure C- 3. Routing table

The developed DBR scheduling software provides a function for visually representing the bill-of-material and routing data. This is achieved by means of a product flow diagram, which is a combined network of the BOM table and the Routing table. The product flow diagram visually shows the product flow from the raw material, through the different manufacturing operations, to the final end product. One product flow at a time is displayed, and the different products can be chosen from a drop-down list. Figure C- 4 presents an example of a product flow diagram for a product with an identification number of 9000.

The different blocks in the product flow diagram are linked to the database. Clicking on a certain block will present a form with more details on the specific operation, such as the resource used, processing time, set-up time, the WIP inventory, etc. Some of the data in this query form can also be changed and the corresponding database table will be updated. The advantage of such a product flow diagram is that manufacturing data can very easily be visually inspected for any data entry errors. It also helps to visualise and better understand the whole manufacturing process.

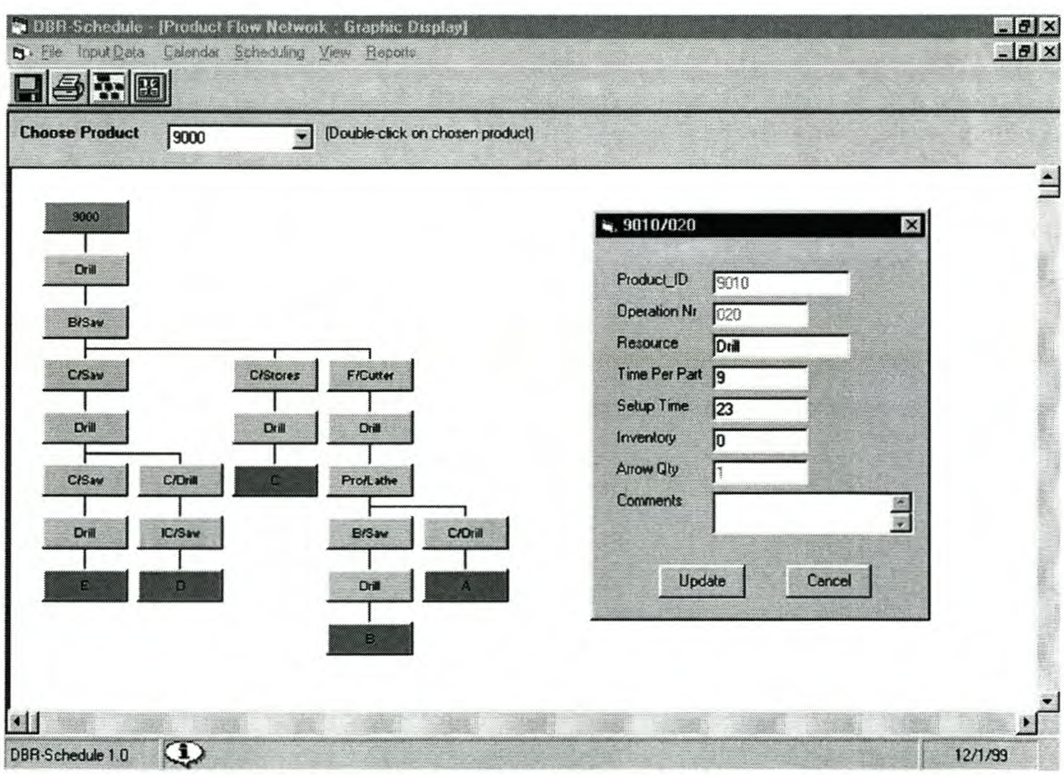

Figure C- 4. Example Product Flow Diagram

C.2 Constraint Identification

The developed DBR scheduling program provides a function for the identification of capacity constraint resources (CCR's). The DBR software program uses a kind of capacity resource profile which takes into account the capacity and availability of the resources during a certain time frame, as well as the placement of the orders on the time line. Simply dividing the total load by the total availability ignores the timing aspect of the order placement. The order is placed on the CCR with a certain lead time offset from the due date. In order to protect the due date, time to the right of the order therefore is not available for processing.

The first step in the constraint identification procedure of the developed DBR software is to start placing the loads generated by the orders on the different resources. Processing begins with the earliest order and moves forward until the latest order in the time horizon specified. Inventory is allocated first-come first-served by moving down the product flow network for each product (from end product down to raw materials). At each operation of the product flow diagram of each product, the required load is calculated and placed on the corresponding resource without considering its capacity. All the loads generated by a specific product flow diagram for a specific order are placed on the time line a shipping buffer length before the order due date. This is illustrated by Figure C- 5, which shows a simple flow diagram for a certain product and the corresponding load placements on the different resources.

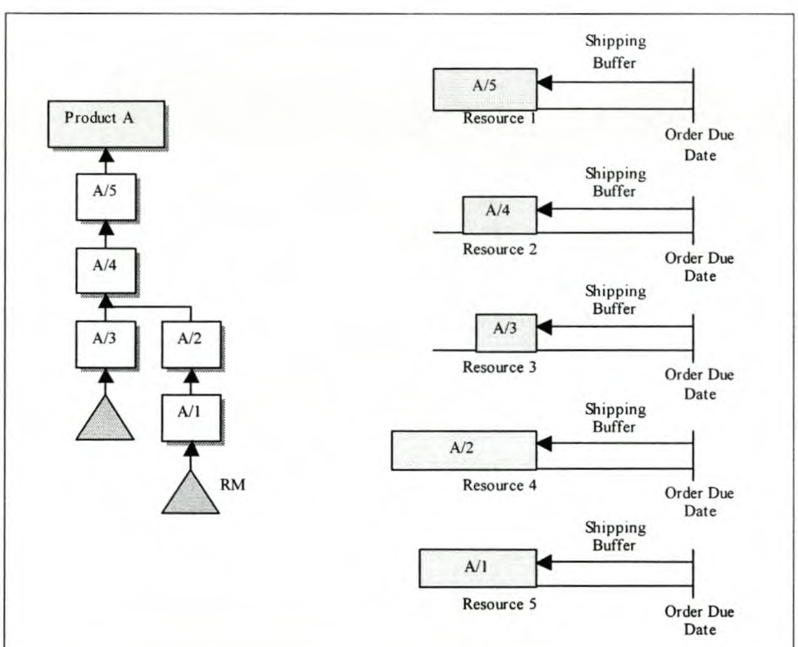

Figure C- 5. Product flow diagram and corresponding initial load placements during constraint identification function

Stellenbosch University http://scholar.sun.ac.za

From Figure C- 5 it can be seen that the load for each part/operation is placed at the same place on the time line. This seems illogical, because the loads for successive operations appear in the exact same place in time. It should be remembered however that the objective at this stage is not to schedule the factory, but to identify the primary CCR. It therefore represents the extreme "optimum" case where the transfer batch size is one. The processing of one part in comparison to the total protective time used is also so small that the difference between where a part would fit on the time line by including the actual processing time or sequencing is very small. After the loads for all the orders have been placed, the load profile on a given resource should look something like that in Figure C- 6.

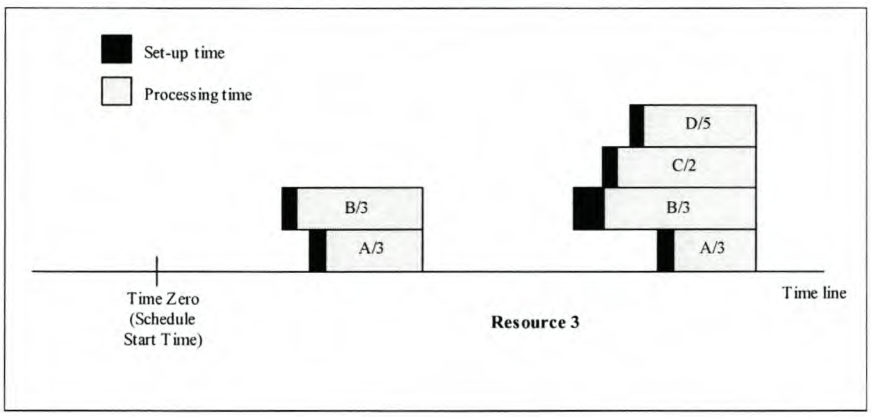

Figure C- 6. Example load profile on a given resource after first step in constraint identification procedure

The next step in the constraint identification procedure is to level the load while considering the resource capacity. This is done through a backward pass that starts at the latest order and moves backward until the earliest order. Since set-up savings can have a significant impact on finding additional capacity, for the purposes of identifying the constraint only one set-up per part-number/operation is considered (unrelated to the number of orders requiring this particular task). The number of set-ups that are going to be needed is also a function of constructing the schedule itself and cannot be predetermined. According to Goldratt (1990, p. 193) set-up time estimations are in any case generally not very accurate, therefore it is best to find a bottleneck without relying too much on set-up data. Assuming Resource 3 in the example of Figure C- 6 has only one available resource unit, the load profile after levelling will look something like that in Figure C- 7.

The primary constraint resource can now be identified by looking at the levelled load profiles of the different resources. That resource whose load has been pushed back the furthest past time zero should be identified as the first candidate for the primary capacity constraint resource (CCR). The loads that have passed time zero represent the amount of protective capacity that has been taken away from the resource. This resource therefore poses the greatest threat to delivering to the market demand on time.

Stellenbosch University http://scholar.sun.ac.za

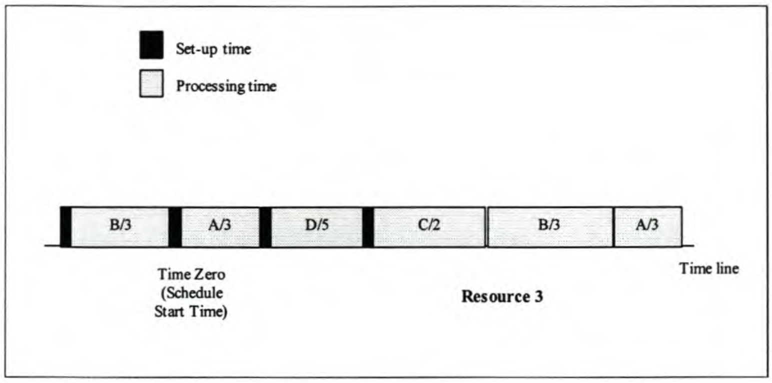

Figure C- 7. Load profile after load levelling during constraint identification process

Figure C- 8 presents an example of an output from the constraint identification procedure of the developed DBR scheduling software. The levelled load profiles of the different resources are shown on a Gantt chart. On the left of the Gantt chart the resources are listed from the resource whose load has been pushed back the furthest on the time line to the resource whose load has been pushed back the least. All the resources whose loads have been pushed past time zero are indicated in red. The resource on top of the list is the candidate for the primary CCR. Once the existence of the primary constraint resource has been validated, the constraint resource ID must be given to the system.

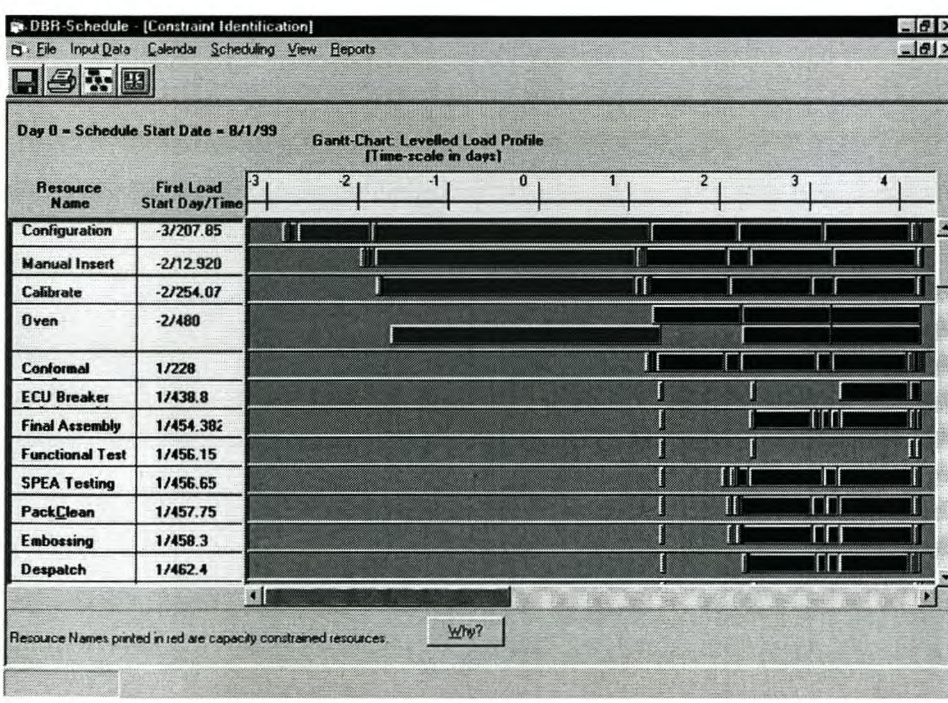

Figure C- 8. Gantt chart output for constraint identification function

C.3 Scheduling The Primary Constraint Resource

The process of scheduling the constraint resource is part of the exploitation step of the five focusing steps of the Theory of Constraints. The purpose is to maximise the amount of throughput generated by the constraint by making effective use of its available time. The scheduling algorithm used in the DBR scheduling software is Goldratt's backward-forward algorithm described in his book "The Haystack Syndrome" (1990). The algorithm was chosen for its simplicity and the fact that it can easily be incorporated in a computerised system. Computation times for rather large data sets would also not be too long, because the algorithm does not perform an exhaustive search procedure trying to find the best combination of jobs within the given time and capacity constraints. Furthermore the algorithm applies a due date based priority rule, which makes it intuitively appealing. Due date based rules are practical and understandable which makes it more believable on the shop floor. The objective of the scheduling algorithm is to minimise the number of late orders as well as the maximum lateness. It will also try to prevent too early completion of the orders, which helps to minimise WIP. The algorithm is therefore geared toward better delivery performance.

Before scheduling can begin, the user must first specify the planning horizon by giving the schedule start and end dates. All orders with due dates within this time frame plus one default shipping buffer length will be included for scheduling. The user must also specify the various time buffer lengths and the specific time buffer option to be used. The first step in the actual scheduling procedure involves the placement of the loads generated by the orders within a given time frame. Processing starts with the earliest order and moves forward in time until the latest order. For each order the loads generated by the order is calculated by moving down the product flow diagram of each product. The net requirement at each operation is calculated by adjusting the order quantity for inventory, scrap percentages and the quantity per specified in the BOM table. Any end product or WIP inventory is allocated on a first-come-first-serve basis. Only the loads for those operations that require processing on the primary constraint resource is calculated and placed. The loads are initially placed without considering the resource's capacity. The end times of the loads after this step therefore represent the ideal or maximum end time for the load. If the load gets shifted to the right of this ideal end time it will endanger the due date of the order.

Loads are placed at a certain time offset before the order's due date. The length of this time offset is equal to the corresponding shipping buffer plus the sum of the processing times for one unit from the constraint resource to the shipping buffer. This ensures that tasks with a longer lead time from the constraint resource to the shipping buffer will be processed first. Situations where a single job requires multiple operations on the same constraint resource with non-constraint operations in-between are handled through the use of time rods. The different backward and forward rod lengths are calculated and "attached" to the corresponding loads.

Refer to Figure C- 9 for an example of a product flow diagram and the corresponding placement of the loads. Operations C/2, B/3, B/1 and E/1 require processing at the primary constraint resource. This is an example of a situation where a job requires multiple operations on the same constraint resource. In order to ensure that the part processed at B/1 will arrive in time at operation B/3, a buffer or time rod of minimum length ½ the constraint buffer must be inserted between B/1 and B/3. This time rod prevents B/1 from being placed too near B/3. A time rod is also inserted between E/1 and B/1.

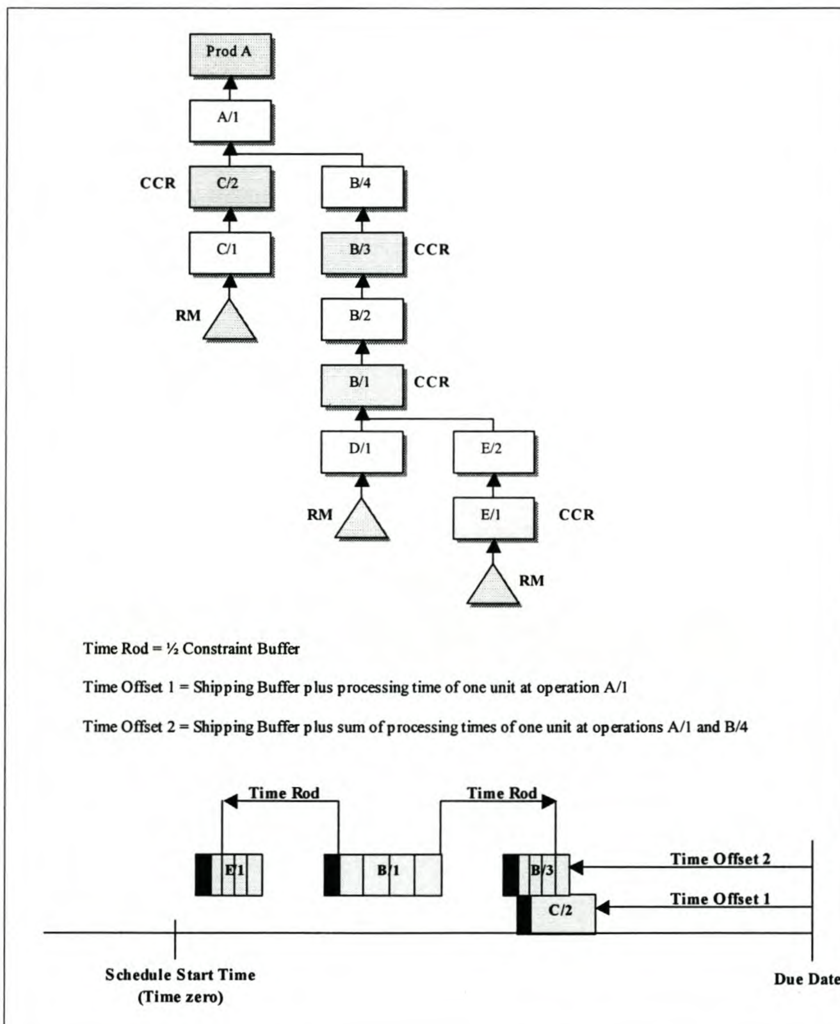

Figure C- 9. Example product flow diagram and placement of loads on time line.

In the previous step the loads were placed on the constraint resource without considering its capacity. As a result of this loads may pile up on top of each other requiring processing at the same point in time. The next step therefore is to level the loads to ensure that the number of jobs required to be worked on at the same point in time will never exceed the number of units available on the resource constraint. This means that whenever loads accumulate, the upper loads need to be shifted. These loads can only be shifted to the left on the time line, because shifting it to the right will endanger the due date of the corresponding order. The levelling process starts at the latest load and moves backward in time until the earliest order. Whenever accumulated loads are encountered, the upper loads are shifted to the left, making them earlier than strictly demanded by the order's due date. This means increasing inventory, but this is better than losing throughput. The sequence of the loads is kept during the levelling process. A load that appears later than another load will therefore keep its relative position. Whenever a load needs to be shifted to the left, the scheduling procedure will look if there are any backward rods attached to the load. If there are and shifting of the load violates the minimum time gap that is needed between the load and its predecessor, the predecessor load and its predecessors will also be shifted to the left.

Figure C- 10 presents the example introduced in Figure C- 9 after levelling of the load. Assume only one constraint resource unit is available. Part/operation B/3 accumulates on top of C/2 and therefore needs to be shifted to the left. The backward rod attached to B/3 forces B/1 also to be shifted to the left, and the backward rod attached to B/1 forces E/1 to be moved to the left.

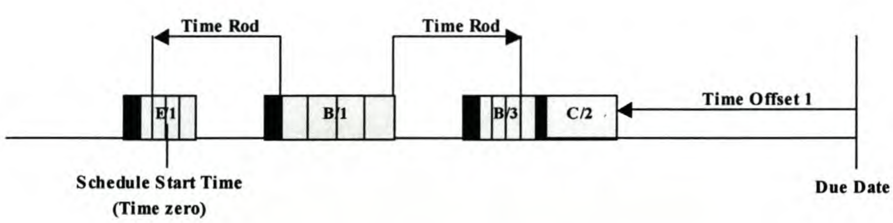

Figure C- 10. Levelling the load

The resource being scheduled is a capacity constraint resource, therefore the result of the backward pass will probably be that some loads will have been pushed to the past (past the schedule start time). Because work cannot be done in the past, the next step is to shift the loads forward in time. Starting at time zero, the scheduling procedure starts placing the earliest order and shifting all corresponding loads forward in time. During the forward pass loads of the same product group or with the same part/operation number appearing next to each other will be combined in order to save set-up time. Since this is a constraint resource, some loads will end up later than their original position (determined during initial placement of the loads). This means that the corresponding order's due date will be endangered.

Orders are considered to have serious timing problems whenever the end time of the latest load within the product flow diagram is later than the load's ideal end time by more than (or equal to) 50% of the corresponding shipping buffer length. The cut-off of 50% of the shipping buffer length is suggested by Goldratt (1990, p. 206). Remember that this latest load was initially placed more or less a shipping buffer before the due date of the order. If the final end time of this load is more than (or equal to) 50% the shipping buffer length passed the ideal end time for the load, it means that a protection of less than 50% of the shipping buffer length is left to finish the order at operations following the constraint resource operation. The chance of missing the due date of the corresponding order is therefore very high and the load can be labelled as late (meaning the corresponding order will most probably be late and action is necessary). If the shift is less than 50% of the shipping buffer length but more than zero, it is considered to be a danger load (meaning the corresponding order is in danger of being late, but no immediate action is necessary).

Refer to Figure C- 11 and Figure C- 12 for an illustration of a "danger" load and a "late" load. Part/operation C/2 is the latest load of the product flow diagram introduced in Figure C- 9. In Figure C- 11 the difference between the load's scheduled end time and its ideal end time is less than 50% the shipping buffer length and is therefore labelled as a "danger" load. In Figure C- 12 the difference is more than 50% the shipping buffer length and the load is labelled as "late".

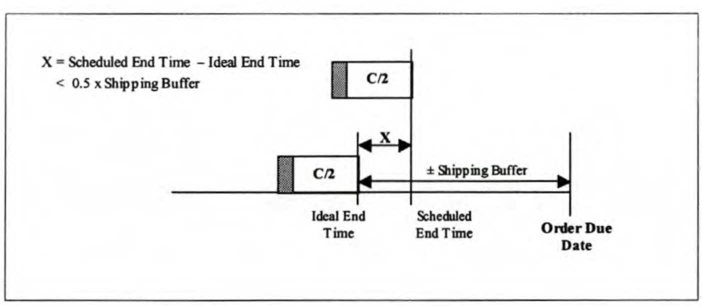

Figure C- 11. "Danger" Load

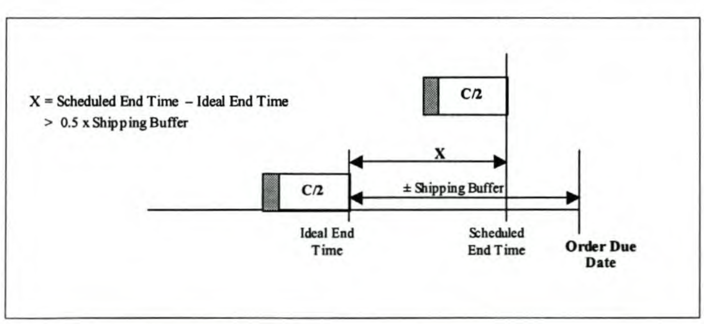

Figure C- 12. "Late" Load

The three steps of placing the loads, levelling the loads through a backward pass and removing conflicts through a forward pass are automatically carried out by the developed DBR scheduling program. The output displayed to the user after the initial scheduling is therefore a Gantt chart of the loads on the primary constraint resource as they appear after the forward pass. Together with this Gantt chart a list is displayed showing all the "danger" and the "late" orders. On the Gantt-chart itself all those loads that belong to a "danger" order and are later than their ideal end times are also shown in pink, and those loads that belong to a "late" order and are later than their ideal end times are shown in red. Figure C- 13 presents an example of a Gantt-chart output of the scheduling procedure of the DBR scheduling program.

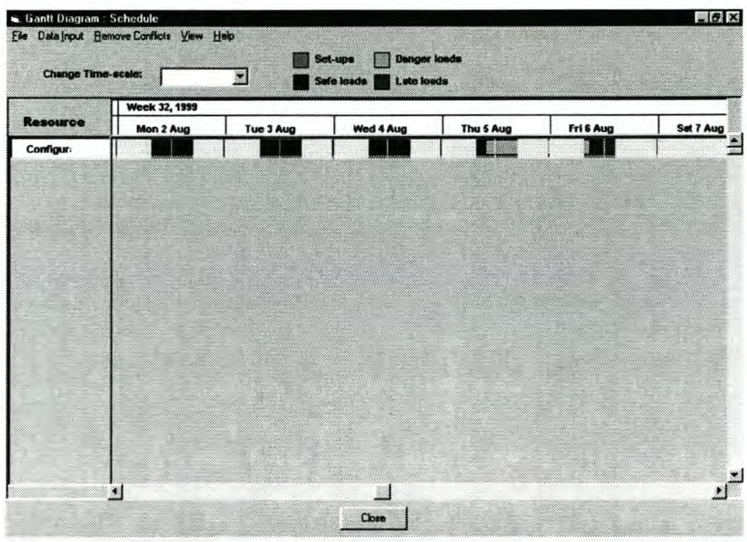

Figure C- 13. Gantt chart showing schedule for primary constraint resource

The user can also view a detailed report on the developed schedule at the primary constraint resource, showing all the start and finish times for the different loads. (Refer to Figure C- 14).

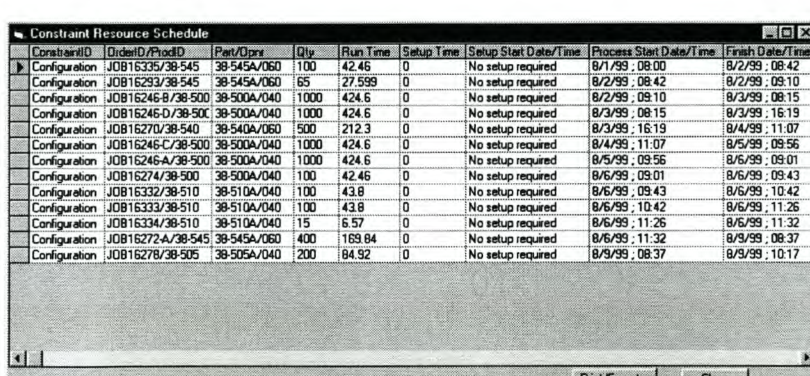

Constraint Resource Schedule

ConstraintID	OrderID/ProdID	Part/Opnr	Qty	Run Time	Setup Time	Setup Start Date/Time	Process Start Date/Time	Finish Date/Time
Configuration	JOB16335/38-545	38-545A/060	100	42.46	0	No setup required	8/1/99 ; 08:00	8/2/99 ; 08:42
Configuration	JOB16293/38-545	38-545A/060	65	27.599	0	No setup required	8/2/99 ; 08:42	8/2/99 ; 09:10
Configuration	JOB16246-B/38-500	38-500A/040	1000	424.6	0	No setup required	8/2/99 ; 09:10	8/3/99 ; 08:15
Configuration	JOB16246-D/38-50(	38-500A/040	1000	424.6	0	No setup required	8/3/99 ; 08:15	8/3/99 ; 16:19
Configuration	JOB16270/38-540	38-540A/060	500	212.3	0	No setup required	8/3/99 ; 16:19	8/4/99 ; 11:07
Configuration	JOB16246-C/38-500	38-500A/040	1000	424.6	0	No setup required	8/4/99 ; 11:07	8/5/99 ; 09:56
Configuration	JOB16246-A/38-500	38-500A/040	1000	424.6	0	No setup required	8/5/99 ; 09:56	8/6/99 ; 09:01
Configuration	JOB16274/38-500	38-500A/040	100	42.46	0	No setup required	8/6/99 ; 09:01	8/6/99 ; 09:43
Configuration	JOB16332/38-510	38-510A/040	100	43.8	0	No setup required	8/6/99 ; 09:43	8/6/99 ; 10:42
Configuration	JOB16333/38-510	38-510A/040	100	43.8	0	No setup required	8/6/99 ; 10:42	8/6/99 ; 11:26
Configuration	JOB16334/38-510	38-510A/040	15	6.57	0	No setup required	8/6/99 ; 11:26	8/6/99 ; 11:32
Configuration	JOB16272-A/38-545	38-545A/060	400	169.84	0	No setup required	8/6/99 ; 11:32	8/9/99 ; 08:37
Configuration	JOB16278/38-505	38-505A/040	200	84.92	0	No setup required	8/9/99 ; 08:37	8/9/99 ; 10:17

Print/Export... Close

Figure C- 14. The detailed primary constraint schedule

If at this stage the developed schedule will cause any of the orders to be late, additional processing time will have to be found on the constraint resource. The DBR scheduling system provides three different options for finding additional processing time on the constraint resource: set-up time savings, off-loading, and overtime. Set-up time savings are performed by combining similar loads, whereas off-loading is performed by moving loads from the primary constraint resource to alternative resources specified by the user. Overtime is inserted by adding additional work hours to late loads.

C.4 Subordination

The objective of the subordination phase is to co-ordinate the activities of all the other resources in the factory to ensure that the schedule developed for the constraint resource is protected, that shipping is completed on time and that raw material is released to the shop floor in synchronisation with the constraint schedule and the order's due dates.

The subordination procedure as applied in the DBR scheduling program performs the following functions:

- It determines whether there is enough protective capacity on the non-constraint resources to protect the constraint resource's schedule and automatically increases the size of the time buffer for those buffer origins that need additional capacity.
- It identifies those resources whose protective capacity is inadequate (secondary or additional CCR's) and which threaten the schedule of the primary constraint resource.
- It prepares the release schedule for raw material into the gating operations.

The heart of the subordination procedure is the buffering system. The concept of time buffers was discussed in section 2.3. The three basic types of buffers used are shipping buffers, constraint buffers and assembly buffers. The DBR scheduling system uses a concept called dynamic buffering. This concept is proposed by Goldratt (1990, pp. 235-240) for use in computerised systems. During subordination, dynamic buffering is only applied to constraint buffers and assembly buffers (if there is a constraint operation), and not the shipping buffer. If there is not a constraint operation and the market is the only constraint, then it is applied to the shipping buffer.

According to dynamic buffering, the initial time buffer lengths (for the constraint and assembly buffers) provided by the user, should only include estimations of the impact of disruptions and fluctuations, and not the total lead time. Estimates for the shipping buffer should however include both processing times and estimations for disruptions. The system will add to the buffer the influence of processing times and queue times caused by the specific load situation. Dynamic buffering is therefore a process that increases the size of the buffer whenever a lack of protective capacity is detected by the system. The time buffer can be separated into two portions, i.e. a fixed portion and a variable portion (Stein 1996, p.115). The fixed portion is given by the user and represents the time protection necessary to handle disturbances created by variability in the system. The variable portion is that portion which increases or decreases depending on the impact of demand on the schedule. If the demand for a given period of time is greater than the capacity of a resource, the demand gets pushed into an earlier time period. Pushed far enough, it extends the size of the buffer.

The dynamic buffering concept can be explained by means of an example. Refer to Figure C- 15.

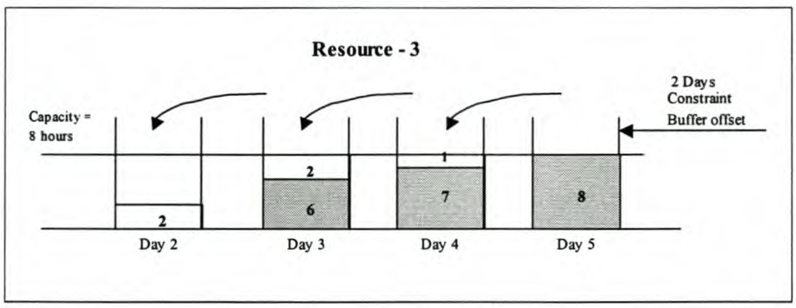

Figure C- 15. Dynamic Buffering

Suppose the fixed constraint buffer length is determined as 2 days. The operation immediately preceding the constraint resource should therefore finish processing two days before the job is to be processed on the constraint resource. Suppose the job is scheduled for processing on the constraint resource on at the start of day 8. The job should therefore be finished at the preceding operation at the end of day 5. Suppose that the resource required for processing at the operation preceding the constraint operation is resource-3. Resource-3's capacity is 8 hours per day. Suppose the load generated by the job at resource-3 is 5 hours. At day 5 resource-3 is already loaded to 100% capacity. The job should therefore be shifted to day 4. On day 4 the resource is already loaded with 7 hours of work. One hour of the job can therefore be performed on day 4, whereas the rest of the load (4 hours) must be shifted to day 3. On day 3 the resource is already loaded with 6 hours of work. Two hours of the job can therefore be performed on day 3, and the rest (2 hours) must be shifted to day 2. All the resource's capacity is available on day 2. The effect of this shifting of the load from day 5 to day 2 means that the buffer for this job was increased from 2 days to 5 days.

The advantage of dynamic buffering is that by increasing the size of the buffer only in those places and times where it is absolutely necessary, the size of the fixed portion of the buffer can be reduced. This has a direct impact on inventory. A smaller buffer will need less inventory to support it, which therefore means a decrease in WIP. Gardiner et al (1998) performed a simulation study comparing the effects of dynamic buffering to a fixed buffer system. Their study showed that the dynamic buffering system used inventory more efficient than the fixed buffer system. Lower WIP levels were achieved and due-date performance was better with the dynamic buffering system.

During the subordination procedure it is very important that the system should be consistent in moving only backward in time. Before the system can assign an operation to a specific date, it must make sure that all operations that should be assigned to a later date have already been dealt with. This is the only way to ensure that if a peak load needs to be shifted to an earlier date, its feeding operations will be placed at a point in time when enough capacity is available and won't create even bigger overloads.

Before the actual subordination process starts, the system first categorises the different part/operations in each product flow diagram corresponding to the different orders into red-lane stations and non-red lane stations:

- Red-lane stations – These are all those part/operations that lie on the routing of a product from a constraint resource forwards. These operations are very important because they process parts that have already consumed valuable constraint time. (Refer to Figure C- 16)
- Non-Red lane stations – These are all those part/operations that lie between material release and a constraint resource, or lie on those routings that do not include any constraint resource. (Refer to Figure C- 16)

The next step is to assign inventory to the non-red lane operations. During the constraint scheduling procedure inventory was only allocated to red-lane operations and CCR operations. It was not necessary then to also allocate inventory to non-red lane operations. It must be remembered that inventory gets allocated according to the orders' due dates - first come, first served. When constructing the CCR schedule, some of the order due dates were likely changed. The correct time for allocating inventory to non-red lane operations is therefore during the subordination stage. Inventory gets allocated by starting at the earliest order and moving forward in time until the latest order. The order due dates are based on the scheduled delivery dates and allocation is carried out first come first served.

Finally the subordination procedure can start. The procedure can be described as a simple bin-loading algorithm treating the loads as a fluid that flows back from one day to the next. The procedure starts at the scheduled delivery date of the latest customer order and moves consistently backward in time. The precision of the procedure depends on the time unit used. It was decided to use daily time buckets. The procedure therefore moves backward in daily time intervals. Since order due dates are usually given without specifying a particular hour, it seems reasonable that as far as movements backward in time are concerned, daily intervals could be used as the unit of maximum sensitivity.

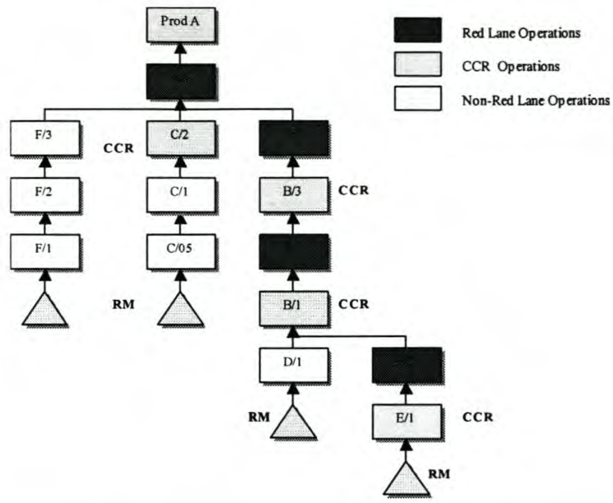

Figure C- 16. Red-Lane and Non-Red Lane Operations

In order to keep track of the movement in time and all the tasks that have to be performed on a particular day, the system uses a "to-do" list. At the top of the list are all the tasks closest to the current date, and at the bottom those tasks that are closest to the present. The current date is defined as the date the system is currently busy with. First of all the initial "to-do" list is created by placing all the orders' scheduled delivery dates, as well as the constraint resource operations at their correct positions on the list. The "to-do" list is organised from the most remote date down to the schedule start date or day zero. The subordination procedure then starts at the highest entry on the list. Each row of the list represents a different day, and tasks to be performed on the same day are therefore all entered in the same row of the "to-do" list. All the entries in a specific row (therefore all the tasks for a specific day) must have been taken care of before the procedure can move to the next earlier day. The type of action that will be performed depends on the type of entry encountered on the "to-do" list:

- An order - If the procedure encounters an order, it dives down the product structure of the product for the corresponding order to determine the immediate feeding operation(s) for the end product. This immediate feeding operation(s) is then placed in the "to-do" list at the scheduled delivery date for the order minus the corresponding shipping buffer. The order is erased from the "to-do" list and the procedure then moves to the next entry at the current row (therefore the current day) of the "to-do" list. If there are no more entries left on the current row, the procedure moves down the list to the next row (therefore it moves one day backward on the time line).
- A constraint operation – If the procedure encounters a constraint operation, it does not again calculate the load represented by the operation, because it had already been taken care of during the constraint scheduling procedure. The subordination procedure only identifies the operation(s) immediately preceding the constraint operation. If the preceding operation is not on a red-lane,

then it is placed on the "to-do" list at a date equal to the start time of the constraint operation minus the corresponding constraint buffer. If the preceding operation however appears on a red-lane, then it gets placed on the "to-do" list at a date equal to the start time of the constraint operation minus only half the corresponding constraint buffer. Because the preceding operation is on a red-lane, it means that further down that particular leg of the product structure is another constraint operation. This is therefore an interactive constraint situation. Goldratt (1990, p. 218) suggests that half a constraint buffer is sufficient protection for the intermediate operations between the two constraint operations. This is to prevent the overall protection in the product flow from becoming over-inflated. Refer for example to Figure C- 16. In this product flow diagram the two immediate feeding operations for constraint operation B/1 are D/1 and E/2. D/1 is not on a red-lane, and will therefore be offset a whole constraint buffer length from the constraint operation's start time. E/2 however is on a red-lane and will there be offset only half a constraint buffer length from the constraint operation.

- A red-lane assembly operation – A red-lane assembly operation is an assembly operation that receives parts from a constraint resource operation and therefore lies on a red-lane. The procedure first calculates the load represented by the assembly operation and places it on the corresponding resource. If the resource's capacity for that day is insufficient, then the surplus load is put on the "to-do" list at the current day minus one and the procedure moves to the next entry on the "to-do" list. If the capacity is sufficient, the load is placed on the resource and the procedure determines the immediate feeding operations of the assembly operation. If one of the feeding operations is a non-red lane operation, it means that this assembly operation is receiving parts from both constraint and non-constraint operations. The non-red lane feeding operations are therefore placed on the "to-do" list at the current date minus an assembly buffer length. For the red-lane feeding operations, the procedure continues diving down the particular leg of the product flow diagram.
- A red-lane operation –Red-lane operations include all the non-assembly operations lying on a red lane. When the procedure encounters such an operation, it calculates the load represented by the operation and checks the current available capacity of the resource that is supposed to perform that operation. If there is enough available capacity, the load is placed on the resource at the current day and the procedure continues diving into the operation's feeding operations. If the available capacity is not enough, the surplus load is placed on the "to-do" list at the current date minus one. The procedure will not continue diving down the operation's feeding operations and will move to the next entry on the "to-do" list at the current date. The problem with red-lane loads is that they are preceded by a constraint resource operation that already has been assigned a place in time during the scheduling procedure. When a surplus red-lane load therefore gets shifted backwards, it can not be shifted past the scheduled end date of the constraint resource operation, because the constraint operation feeds the red-lane operations. Surplus loads will therefore not be shifted past the constraint operation, but will be allowed to build up at the end of the constraint operation.

- A normal operation - This includes all the non-red-lane operations. The procedure calculates the load represented by the operation and checks the current available capacity of the resource that is supposed to perform that operation. If there is enough available capacity, the load is placed on the resource at the current day and the procedure continues diving into the operation's feeding operations. If the available capacity is not enough, the surplus load is placed on the "to-do" list at the current date minus one. The procedure will not continue diving down the operation's feeding operations and will move to the next entry on the "to-do" list at the current date. Surplus loads will not be pushed back past the schedule start date or time zero, but will be allowed to build up at time zero, or the so-called first day.

While the procedure dives down the product structure for a specific operation, all the loads represented by the operation's feeding operations will be placed on their respective resources at the same day as the specific operation on the "to-do" list the procedure is currently dealing with. It is assumed that operations will be overlapped, therefore feeding operations can be run simultaneously. The procedure will continue diving down the product structure (recording every assembly it passes through) until one of three situations is encountered:

- The first situation is reaching a raw material. In this case the procedure will jump back to the nearest higher assembly and dive down any additional legs if they exist (No movement in time has yet occurred). If no assembly is marked, the procedure will return to the next task on the "to-do" list.
- The second situation occurs when in diving the procedure reaches an operation of a constraint resource. The procedure doesn't deal with this operation since it has already been taken care of when the constraint resource was scheduled. The procedure therefore moves to the nearest higher assembly (if it exists) and dives down any additional leg, or returns to the "to-do" list.
- The third situation occurs when the corresponding resource for the specific operation does not have enough capacity. The surplus load needs to be pushed back a day, therefore the operation is placed on the "to-do" list at the current day minus one. The procedure moves to the nearest higher assembly (if it exists) and dives down any additional leg, or returns to the "to-do" list.

The following is a simple example of the subordination process and the use of the "to-do" list. Figure C- 17 presents a simple product flow diagram for a certain product X. An order is created for this product to be due on day 100. The constraint operation for this product flow diagram is operation A/30. This operation has already been scheduled for processing on the constraint resource on day 98. At the start of subordination, the "to-do" list will look like that in Figure C- 18. Only the order and the constraint operation at this stage appear on the list.

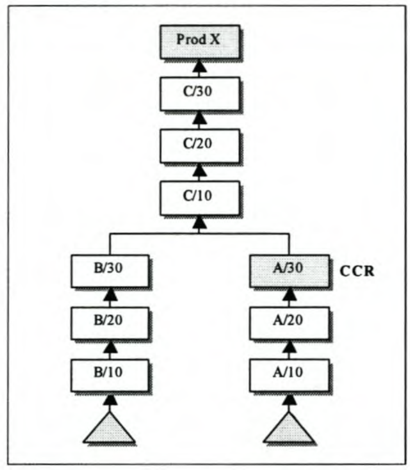

Figure C- 17. Product Flow Diagram

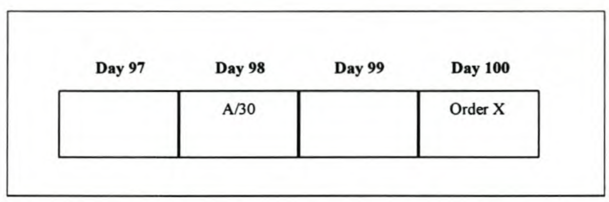

Figure C- 18. Initial "To-do" List

Processing starts at the top of the list with day 100. On this day appears only one entry, which is an order. The immediate feeding operation for the order is part/operation C/30. Assuming the shipping buffer is two days, then this operation is placed on the "to-do" list on day 98 (refer to Figure C- 19).

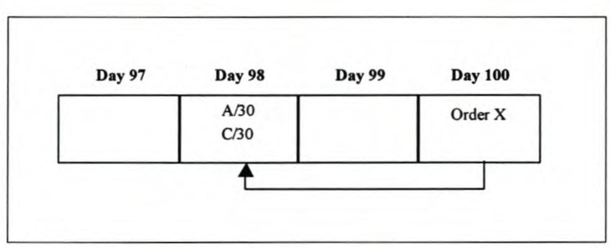

Figure C- 19. Placing Order's Feeding Operation on "To-do" List

No more entries are left on day 100, therefore the procedure moves back a day to day 99. No entries appear on day 99, therefore the procedure moves back to day 98. The first entry on day 98 is a constraint operation. The load for this operation has already been dealt with during the scheduling of the constraint resource. The procedure therefore determines the constraint operation's feeding operation (which is part/operation A/20) and places it on the "to-do" list at the current day minus the constraint buffer length. Assuming a constraint buffer length of one day, operation A/20 is placed on the "to-do" list on day 97 (refer to Figure C- 20). The next entry on the list on day 98 is operation C/30. This is a red-lane operation. The procedure calculates the load represented by the operation and places it on the corresponding resource at the same day. Assuming there is enough capacity on the resource, the procedure continues diving into C/30's feeding operations. The first is operation C/20. The load for this operation is calculated and placed on the corresponding resource, also on day 98. Assuming there is sufficient capacity, the procedure proceeds to operation C/10. This is a red-lane assembly operation. The load for this operation is calculated and placed on the corresponding resource. Assume there is insufficient capacity to process the entire load on day 98. The surplus load would therefore have to be shifted back a day. Operation C/10 however appears on a red-lane, which means it is preceded by a constraint operation, A/30. Shifting C/10's surplus load back a day, would therefore necessitate moving constraint operation A/30 also backwards. A/30 however appears on the constraint schedule and is already fixed in time. The surplus load for operation C/10 would therefore have to be allowed to build up on day 98, causing an overload or peak load on day 98 at the corresponding resource.

The procedure continues determining the feeding operations for operation C/10. The first one is B/30. This is a non-red lane operation, which means the procedure has to place this operation on the "to-do" list at the current day minus an assembly buffer. Assuming an assembly buffer of one day, operation B/30 is placed on the list at day 97. The other feeding operation is A/30. This however is a constraint operation that has already been taken care of. The procedure therefore stops diving down the product flow diagram and returns to the next entry on the "to-do" list. At this stage the "to-do" list will look like that in Figure C- 20.

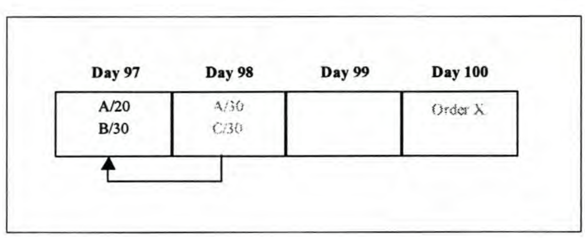

Figure C- 20. "To-do" list after the last entry on day 98 have been processed

No more entries are left on day 98, therefore the procedure moves back a day to day 97. The first entry on the list is A/20. The procedure calculates the load for the operation and places it on the corresponding resource. Assuming sufficient capacity, the procedure moves to its feeding operation (A/10), calculates the load and places it on the corresponding resource. The next operation in the product flow diagram is a raw material. The procedure therefore stops diving down the product structure and moves to the next entry on the "to-do" list, B/30. The procedure continues calculating the loads for this operation and its feeding operation and places the loads on the corresponding resources.

After the system has carried out the subordination procedure, the user is presented with a raw materials release schedule showing which materials to be released on what days, the operation to which it should be released, as well as the quantities. Figure C- 21 shows an example of the raw materials release schedule.

Material Release Schedule

Release Date	Raw Mat ID	To Station	Qty Req	OrderID/ProdID
8/2/99	ED_MIRaw	35-205A/010	1843	JOB16246-D/38-500
8/2/99	35-205	35-205A/010	1843	JOB16246-D/38-500
8/2/99	ED_LCD_Assem_Ra	35-205B/010	1843	JOB16246-D/38-500
8/2/99	60-040_Raw	60-040/010	1843	JOB16246-D/38-500
8/2/99	60-015_Raw	60-015/010	1843	JOB16246-D/38-500
8/2/99	TSK/2_FA_RAW	38-500A/010	1843	JOB16246-D/38-500
8/2/99	33-290	38-545D/010	1.133	JOB16335/38-545
8/2/99	33-355	38-545C/010	100	JOB16335/38-545
8/2/99	32-035	38-545C/010	100	JOB16335/38-545
8/2/99	25-345	38-545B/010	200	JOB16335/38-545
8/2/99	33-200	38-545B/010	100	JOB16335/38-545
8/2/99	21-280	38-545B/010	200	JOB16335/38-545
8/2/99	25-295	38-545B/010	300	JOB16335/38-545
8/2/99	33-165	38-545C/010	100	JOB16335/38-545
8/2/99	33-175	38-545/010	20	JOB16335/38-545
8/2/99	29-310	38-545/010	20	JOB16335/38-545
8/2/99	33-290	38-545D/010	0.73645	JOB16293/38-545
8/2/99	33-355	38-545C/010	65	JOB16293/38-545
8/2/99	32-035	38-545C/010	65	JOB16293/38-545
8/2/99	25-345	38-545B/010	130	JOB16293/38-545

Print/Export... Close

Figure C- 21. The Material Release Schedule

After the subordination procedure, the next step is to identify any peak loads at non-constraint resources. The end of the subordination procedure leaves the user with a capacity diagram showing the loads on the different resources. It may be that some of these resources show peaks of overloads on the first day (because overloads can't be shifted past the schedule start time), or peak loads may appear on the red lanes where overloads can't be shifted past an already scheduled constraint operation. While the largest peak load does not necessarily identify the secondary constraint resource, it is one among several clues that will help to identify the next problem resource.

The system displays a list to the user showing all the overloads, the resource on which it occurs, as well as the ratio between the overload (expressed in terms of hours) and the smallest buffer size preceded by the resource. Refer to Figure C- 22 for an example of the peak loads list. The list also displays the amount of available hours for the resource on the day that the overload occurs, as well as the amount of overload hours. For a normal resource the amount of available hours is calculated as the working hours for that day times the number of identical resource units. For a batch resource it is calculated as the working hours for the day times the number of identical resource units times the number of capacity units.

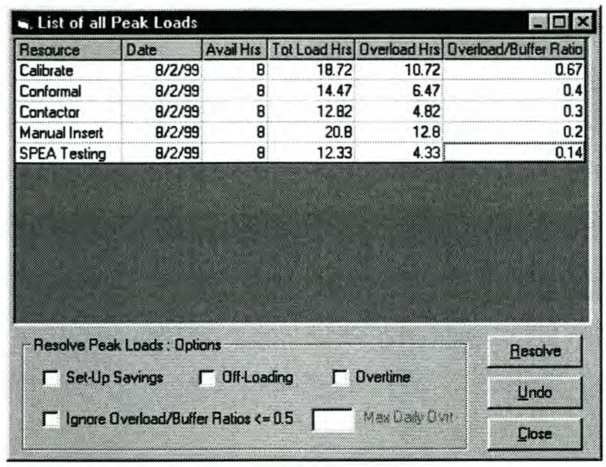

Figure C- 22. The Peak Loads List displaying the overloads

Refer to Figure C- 22. Resources that display a peak load to buffer ratio of less than 50% usually do not present such a greater danger to the protection provided by the system. Half a buffer or more could still be enough time left to absorb the overload. The user therefore has the option to choose whether overload/buffer ratios of less than 50% should be ignored or not during execution of the above three options. Three options are available to resolve peak loads:

- Set-up time savings - Set-up savings are only performed for those loads that appear on the day of the specific overload. It is also only carried out for normal resources, because set-up savings were already taken into account when the loads for batch resources were placed. The set-up savings procedure looks for all similar part/operation numbers or product groups (depending on the set-up option chosen in the Resource Data file) on the day of the specific overload. Only one set-up per similar product group or part/operation number is counted, and the amount of set-up time saved is subtracted from the overload hours.
- Off-loading - If there are still overloads after the set-up saving procedure, the system will try off-loading loads in order to resolve the overloads.
- Overtime - If there are still overloads left, the system will finally look for overtime to be inserted. The maximum amount of overtime that can be inserted on a specific day is given by the user.

As seen from Figure C- 22, the user can choose which combination of the three options to execute. All the chosen options will be executed together. The system will first carry out set-up savings (if the option was chosen), then automatically proceed to off-loading (if the option was chosen), and finally will automatically look for finding additional time through overtime savings (if the option was chosen). It will try and resolve all peak loads on all resources.

Stellenbosch University http://scholar.sun.ac.za

Appendix D

Normal Probability Plots and Scatter Plots of Residuals For 15-Station Flow Shop Protective Capacity Experiment

D.1 Flow Time ANOVA

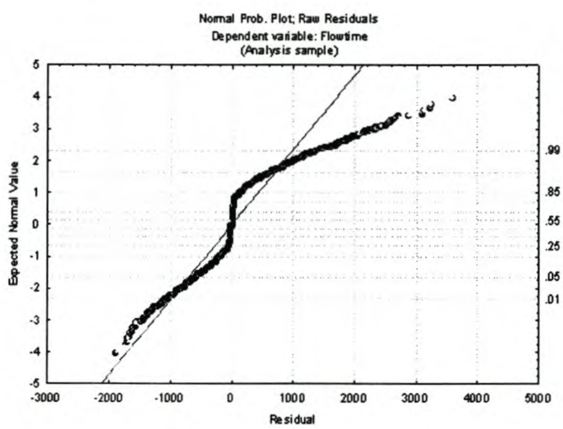

Figure D- 1. Normal Probability Plot

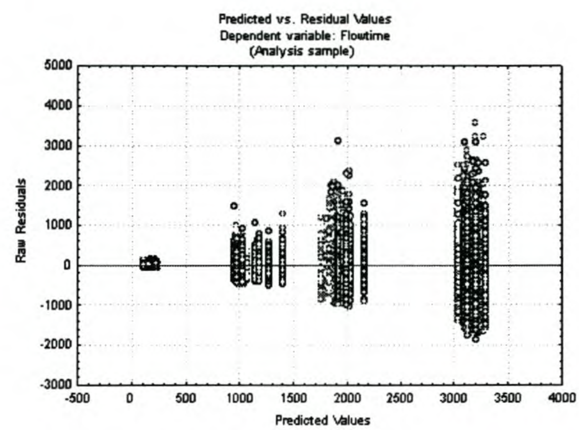

Figure D- 2. Scatter Plot of Residuals vs Predicted or Fitted Values

D.2 Line Output ANOVA

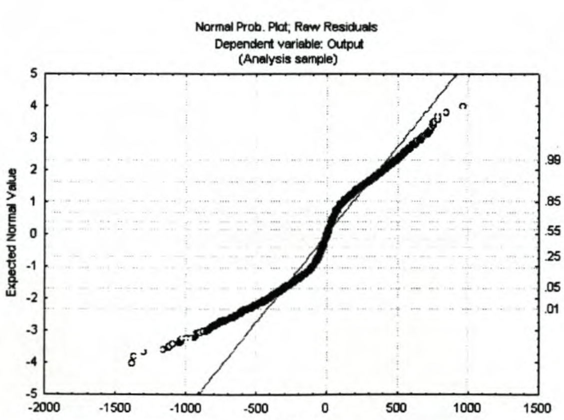

Figure D- 3. Normal Probability Plot

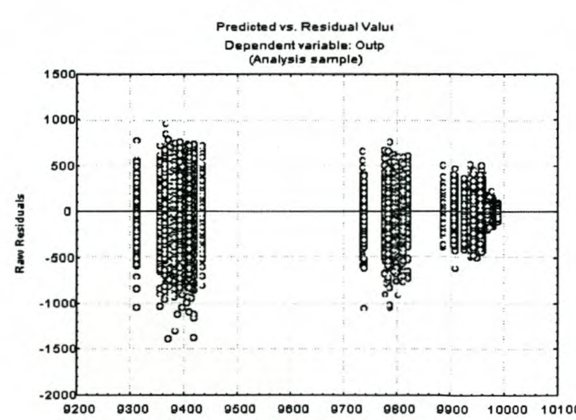

Figure D- 4. Scatter Plot of Residuals vs Predicted or Fitted Values

D.3 Primary CCR Bottleneck Probability ANOVA

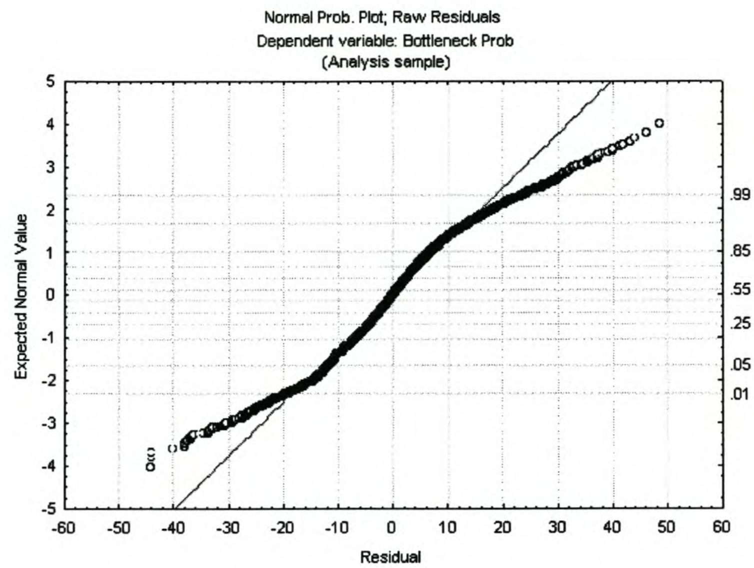

Figure D- 5. Normal Probability Plot

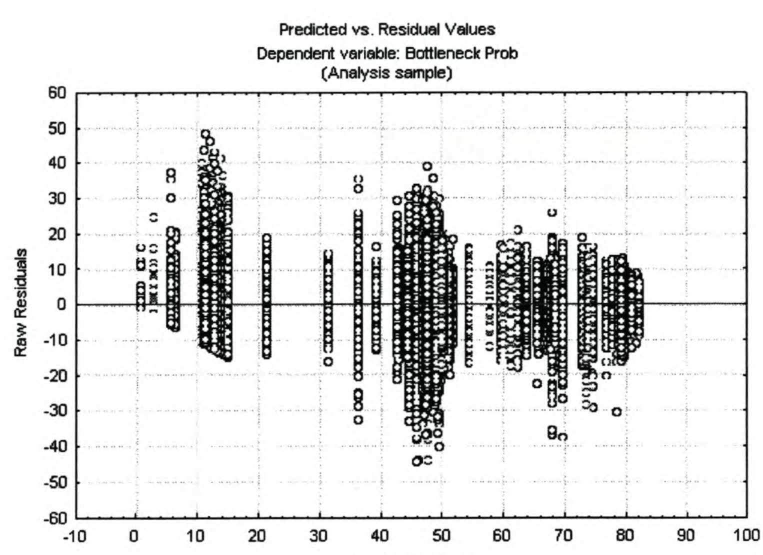

Figure D- 6. Scatter Plot of Residuals vs Predicted or Fitted Values

Stellenbosch University http://scholar.sun.ac.za

Appendix E

Normal Probability Plots, Scatter Plots of Residuals and ranked ANOVA's For the Assembly Line Experiment

E.1 6-Station Infinite Buffer Line with 0.3 CV

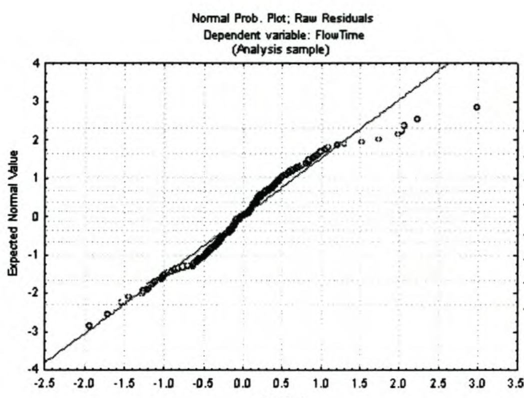

Figure E- 1. Normal Probability Plot: 6-station infinite buffer line with 0.3 cv

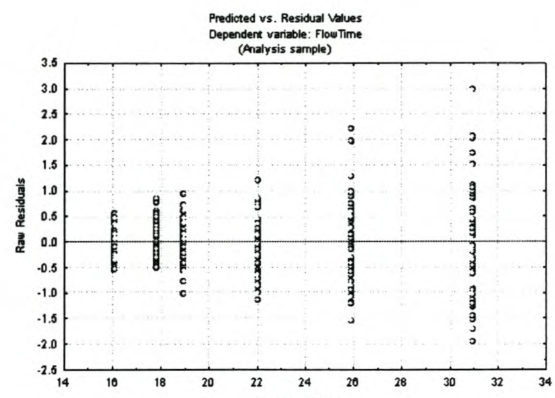

Figure E- 2. Scatter Plot of Residuals vs Predicted or Fitted Values: 6-station infinite buffer line with 0.3 cv

Table E- 1. Ranked ANOVA table: 6-station infinite buffer line with 0.3 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	6795075	1	6795075	30081.64	0.00
PC Level	2183564	5	436713	1933.32	0.00
Error	66411	294	226		

E.2 6-Station Infinite Buffer Line with 1.0 CV

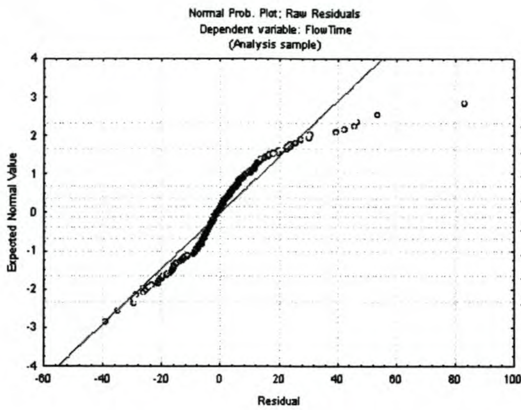

Figure E- 3. Normal Probability Plot: 6-station infinite buffer line with 1.0 cv

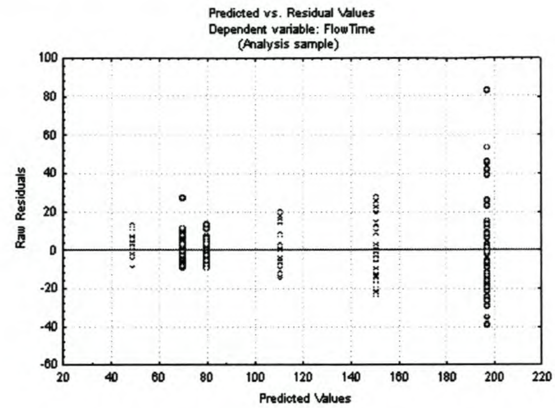

Figure E- 4. Scatter Plot of Residuals vs Predicted or Fitted Values: 6-station infinite buffer line with 1.0 cv

Table E- 2. Ranked ANOVA table: 6-station infinite buffer line with 1.0 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	6795075	1	6795075	20713.00	0.00
PC Level	2153526	5	430705	1312.89	0.00
Error	96449	294	328		

E.3 6-Station Finite Buffer Line

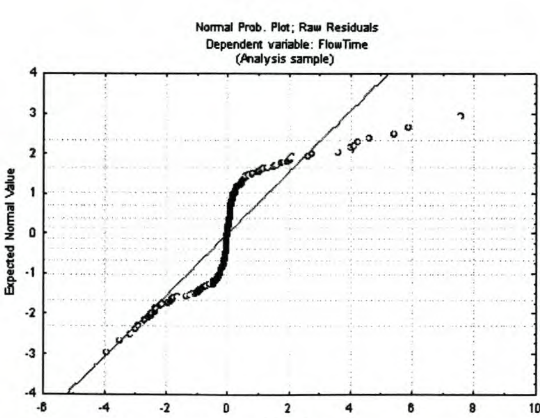

Figure E- 5. Normal Probability Plot: 6-station finite buffer line 1.0 buffer capacity

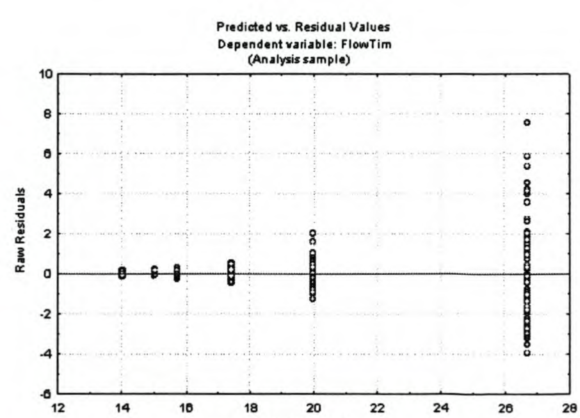

Figure E- 6. Scatter Plot of Residuals vs Predicted or Fitted Values: 6-station finite buffer line with 1.0 buffer capacity

Table E- 3. Ranked ANOVA table: 6-station finite buffer line with 1.0 buffer capacity

	SS	Degr. Of Freedom	MS	F	p
Intercept	18610305	1	18610305	44934.34	0.00
PC Level	6002500	5	1200500	2898.59	0.00
Error	171465	414	414		

E.4 20-Station Infinite Buffer Line with 0.3 CV

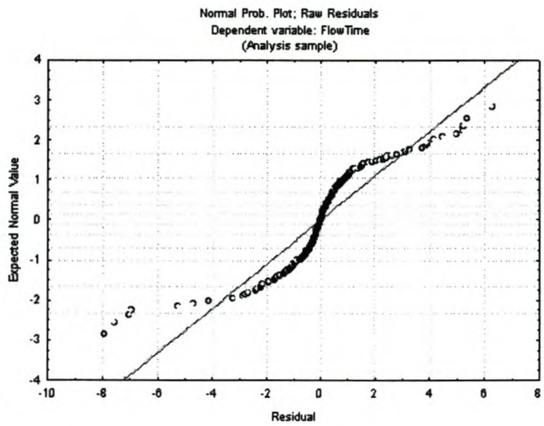

Figure E- 7. Normal Probability Plot: 20-station infinite buffer line with 0.3 cv

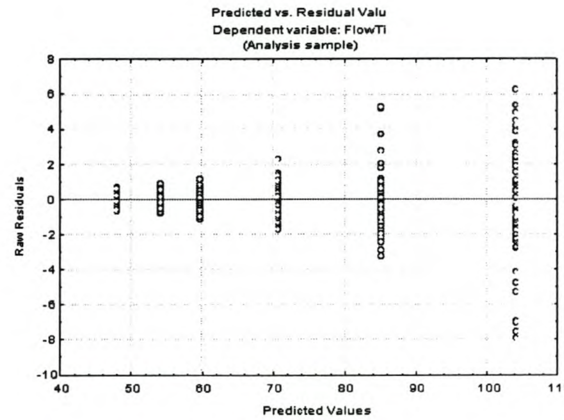

Figure E- 8. Scatter Plot of Residuals vs Predicted or Fitted Values: 20-station infinite buffer line with 0.3 cv

Table E- 4. Ranked ANOVA table: 20-station infinite buffer line with 0.3 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	6795075	1	6795075	31976.82	0.00
PC Level	2187500	5	437500	2058.82	0.00
Error	62475	294	212		

Stellenbosch University http://scholar.sun.ac.za

E.5 20-Station Infinite Buffer Line with 1.0 CV

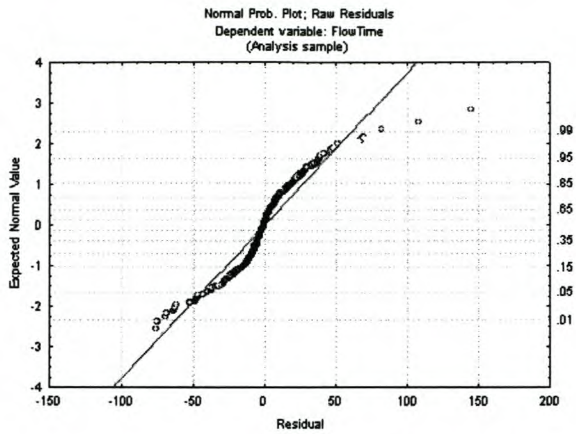

Figure E- 9. Normal Probability Plot: 20-station infinite buffer line with 1.0 cv

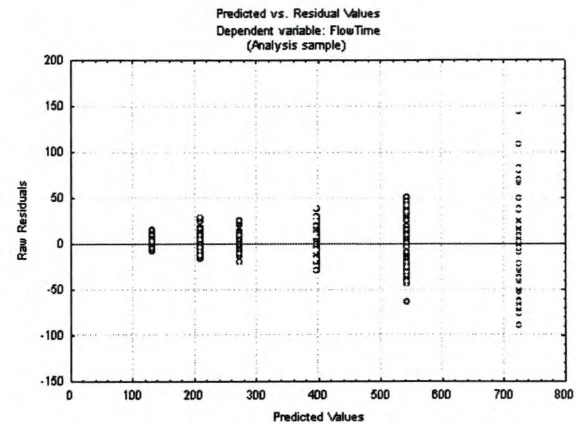

Figure E- 10. Scatter Plot of Residuals vs Predicted or Fitted Values: 20-station infinite buffer line with 1.0 cv

Table E- 5. Ranked ANOVA table: 20-station infinite buffer line with 1.0 cv

	SS	Degr. Of Freedom	MS	F	p
Intercept	6795075	1	6795075	31976.82	0.00
PC Level	2187500	5	437500	2058.82	0.00
Error	62475	294	212		

E.6 20-Station Finite Buffer Line

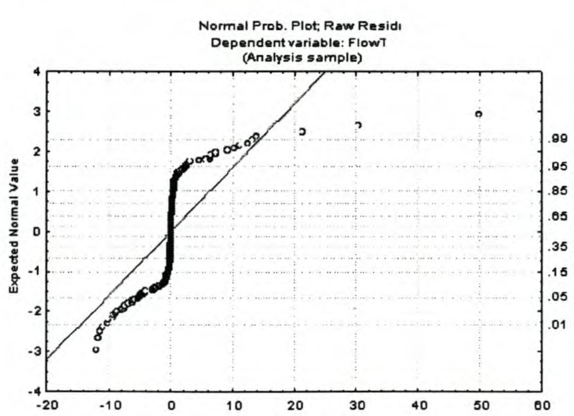

Figure E- 11. Normal Probability Plot: 20-station finite buffer line 1.0 buffer capacity

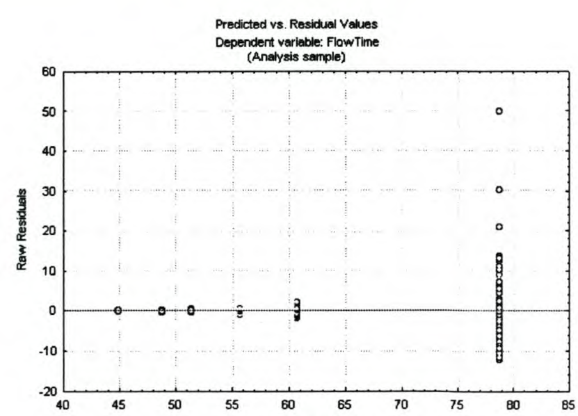

Figure E- 12. Scatter Plot of Residuals vs Predicted or Fitted Values: 20-station finite buffer line with 1.0 buffer capacity

Table E- 6. Ranked ANOVA table: 20-station finite buffer line with 1.0 buffer capacity

	SS	Degr. Of Freedom	MS	F	p
Intercept	18610305	1	18610305	44934.34	0.00
PC Level	6002500	5	1200500	2898.59	0.00
Error	171465	414	414		

Stellenbosch University http://scholar.sun.ac.za

Appendix F

Simulation Results of the 144 Cases By the Four Factors and the Three Dependent Variables For the 15-Station Flow Shop Experiment

F.1 Simulation Results For Total Output Dependent Variable

Table F- 1. Group 1: factor combinations resulting in highest output

SecCCRLoc	SecCCR PC Level	NonCCR PC	Downtime	Throughput
BF	5.25%	21%	1.57%	9989.15
AN	21%	10.5%	1.57%	9981.22
BN	5.25%	21%	1.57%	9980.49
BN	31.5%	21%	1.57%	9979.83
AF	31.5%	10.5%	1.57%	9979.74
AN	10.5%	52.5%	1.57%	9979.69
AF	2%	52.5%	1.57%	9979.61
BF	5.25%	10.5%	1.57%	9979.59
BF	31.5%	52.5%	1.57%	9978.85
BF	52.5%	21%	1.57%	9978.40
BN	21%	10.5%	1.57%	9978.26
AN	2%	52.5%	1.57%	9978.13
AN	5.25%	52.5%	1.57%	9978.13
BF	52.5%	10.5%	1.57%	9977.45
BN	5.25%	52.5%	1.57%	9977.25
AN	5.25%	21%	1.57%	9977.19
AF	5.25%	21%	1.57%	9977.12
BN	52.5%	21%	1.57%	9977.09
BN	10.5%	52.5%	1.57%	9976.82
BN	21%	52.5%	1.57%	9976.78
AF	52.5%	10.5%	1.57%	9976.71
BF	10.5%	52.5%	1.57%	9976.38
AF	2%	10.5%	1.57%	9976.32
AN	52.5%	21%	1.57%	9976.11
BF	2%	52.5%	1.57%	9975.91
AF	31.5%	21%	1.57%	9975.87
BN	31.5%	10.5%	1.57%	9975.83
BN	2%	52.5%	1.57%	9975.80
AF	21%	52.5%	1.57%	9975.65
AN	31.5%	21%	1.57%	9975.35
BF	2%	21%	1.57%	9975.01
AN	2%	21%	1.57%	9975.00
AF	31.5%	52.5%	1.57%	9974.91
AF	5.25%	10.5%	1.57%	9974.85
BN	10.5%	10.5%	1.57%	9974.69
BF	10.5%	10.5%	1.57%	9974.69
AN	10.5%	10.5%	1.57%	9974.69
AF	10.5%	10.5%	1.57%	9974.69
BN	5.25%	10.5%	1.57%	9974.53
AN	10.5%	21%	1.57%	9974.45
BN	2%	21%	1.57%	9974.10
AN	31.5%	10.5%	1.57%	9974.08
AF	52.5%	21%	1.57%	9973.76
BF	31.5%	21%	1.57%	9973.75
AF	2%	21%	1.57%	9973.21
BF	10.5%	21%	1.57%	9972.96
BF	2%	10.5%	1.57%	9972.55
BF	21%	52.5%	1.57%	9972.45
AF	5.25%	52.5%	1.57%	9972.43
AN	21%	52.5%	1.57%	9972.29
BF	21%	10.5%	1.57%	9972.26
BN	2%	10.5%	1.57%	9972.07
AN	31.5%	52.5%	1.57%	9971.76
AN	2%	10.5%	1.57%	9971.58
BN	52.5%	10.5%	1.57%	9971.54
AF	10.5%	52.5%	1.57%	9971.33
AN	5.25%	10.5%	1.57%	9971.29
BN	31.5%	52.5%	1.57%	9970.96
BF	5.25%	52.5%	1.57%	9970.79
AF	21%	10.5%	1.57%	9970.48
AN	52.5%	10.5%	1.57%	9969.82
BF	31.5%	10.5%	1.57%	9969.53
BN	21%	21%	1.57%	9968.93
BF	21%	21%	1.57%	9968.93
AN	21%	21%	1.57%	9968.93
AF	21%	21%	1.57%	9968.93
BN	52.5%	52.5%	1.57%	9968.82
BF	52.5%	52.5%	1.57%	9968.82
AN	52.5%	52.5%	1.57%	9968.82
AF	52.5%	52.5%	1.57%	9968.82
AF	10.5%	21%	1.57%	9968.77
BN	10.5%	21%	1.57%	9968.47
BN	2%	52.5%	10.04%	9961.33
BF	31.5%	52.5%	10.04%	9960.45
AN	10.5%	52.5%	10.04%	9959.34
AN	21%	52.5%	10.04%	9956.29
AN	2%	52.5%	10.04%	9955.66
BN	5.25%	52.5%	10.04%	9949.41
AN	5.25%	52.5%	10.04%	9949.01
AF	10.5%	52.5%	10.04%	9945.16
AN	31.5%	52.5%	10.04%	9944.11
BN	52.5%	52.5%	10.04%	9943.48
BF	52.5%	52.5%	10.04%	9943.48
AN	52.5%	52.5%	10.04%	9943.48
AF	52.5%	52.5%	10.04%	9943.48
BN	21%	52.5%	10.04%	9942.83
BF	5.25%	52.5%	10.04%	9942.52
BF	21%	52.5%	10.04%	9940.13
AF	31.5%	52.5%	10.04%	9938.80
BN	31.5%	52.5%	10.04%	9938.20
BN	10.5%	52.5%	10.04%	9938.15
BF	10.5%	52.5%	10.04%	9925.95
BF	2%	52.5%	10.04%	9908.19
AF	5.25%	52.5%	10.04%	9908.17
AF	21%	52.5%	10.04%	9906.22
AF	2%	52.5%	10.04%	9885.99

Stellenbosch University http://scholar.sun.ac.za

Table F- 2. Group 2: factor combinations resulting in 2nd highest output

SecCCRLoc	SecCCR PC Level	NonCCR PC	Downtime	Throughput
AN	10.5%	21%	10.04%	9820.73
AF	10.5%	21%	10.04%	9812.68
BF	10.5%	21%	10.04%	9810.00
BF	52.5%	21%	10.04%	9807.30
AN	31.5%	21%	10.04%	9801.71
BN	31.5%	21%	10.04%	9801.13
BN	10.5%	21%	10.04%	9799.13
BN	21%	21%	10.04%	9797.51
BF	21%	21%	10.04%	9797.51
AN	21%	21%	10.04%	9797.51
AF	21%	21%	10.04%	9797.51
AN	5.25%	21%	10.04%	9794.09
AF	31.5%	21%	10.04%	9792.35
BN	52.5%	21%	10.04%	9787.67
AN	52.5%	21%	10.04%	9786.78
AN	2%	21%	10.04%	9786.20
AF	52.5%	21%	10.04%	9782.86
BN	5.25%	21%	10.04%	9782.75
BF	5.25%	21%	10.04%	9778.90
BN	2%	21%	10.04%	9778.32
BF	31.5%	21%	10.04%	9776.70
BF	2%	21%	10.04%	9775.73
AF	5.25%	21%	10.04%	9737.71
AF	2%	21%	10.04%	9734.09

Table F- 3. Group 3: factor combinations resulting in lowest output

SecCCRLoc	SecCCR PC Level	NonCCR PC	Downtime	Throughput
AN	21%	10.5%	10.04%	9433.83
BN	21%	10.5%	10.04%	9418.77
AN	52.5%	10.5%	10.04%	9414.63
AF	31.5%	10.5%	10.04%	9410.28
AF	21%	10.5%	10.04%	9409.57
BN	10.5%	10.5%	10.04%	9407.37
BF	10.5%	10.5%	10.04%	9407.37
AN	10.5%	10.5%	10.04%	9407.37
AF	10.5%	10.5%	10.04%	9407.37
BF	21%	10.5%	10.04%	9405.17
BN	5.25%	10.5%	10.04%	9396.75
AN	2%	10.5%	10.04%	9396.33
AN	31.5%	10.5%	10.04%	9393.39
BN	31.5%	10.5%	10.04%	9388.56
BF	31.5%	10.5%	10.04%	9384.25
AF	52.5%	10.5%	10.04%	9382.20
BF	5.25%	10.5%	10.04%	9381.07
AN	5.25%	10.5%	10.04%	9371.86
BN	2%	10.5%	10.04%	9371.55
AF	2%	10.5%	10.04%	9366.39
BN	52.5%	10.5%	10.04%	9365.34
AF	5.25%	10.5%	10.04%	9364.61
BF	2%	10.5%	10.04%	9355.05
BF	52.5%	10.5%	10.04%	9312.98

Stellenbosch University http://scholar.sun.ac.za

F.2 Simulation Results For Flow Time Dependent Variable

Table F- 4. Flow time measures arranged from best to worst

SecCCRLoc	SecCCR PC Level	NonCCR PC	Downtime	Flowtime
AN	31.5%	52.5%	1.57%	103.35905112
BF	31.5%	52.5%	1.57%	103.67484357
BN	52.5%	52.5%	1.57%	103.89173300
BF	52.5%	52.5%	1.57%	103.89173300
AN	52.5%	52.5%	1.57%	103.89173300
AF	52.5%	52.5%	1.57%	103.89173300
BF	21%	52.5%	1.57%	104.39973519
BN	31.5%	52.5%	1.57%	104.72364588
AN	21%	52.5%	1.57%	106.05324745
AF	31.5%	52.5%	1.57%	106.19062631
BN	21%	52.5%	1.57%	106.98111555
BN	10.5%	52.5%	1.57%	107.37450543
AN	10.5%	52.5%	1.57%	108.39496646
BF	10.5%	52.5%	1.57%	109.39632009
AF	21%	52.5%	1.57%	110.57071770
BF	5.25%	52.5%	1.57%	110.96884112
AN	5.25%	52.5%	1.57%	111.21046061
BN	5.25%	52.5%	1.57%	112.82503310
BN	2%	52.5%	1.57%	116.36266731
AF	10.5%	52.5%	1.57%	119.42749076
AN	2%	52.5%	1.57%	119.52444373
BF	2%	52.5%	1.57%	121.59621953
AF	5.25%	52.5%	1.57%	130.11194969
AF	2%	52.5%	1.57%	145.82055889
AF	52.5%	21%	1.57%	152.81376407
BF	52.5%	21%	1.57%	153.37991372
BF	31.5%	21%	1.57%	154.17169457
BN	52.5%	21%	1.57%	154.20295147
AN	31.5%	21%	1.57%	154.40751304
BN	31.5%	21%	1.57%	154.49454566
AF	31.5%	21%	1.57%	154.62389200
BN	21%	21%	1.57%	154.80850167
BF	21%	21%	1.57%	154.80850167
AN	21%	21%	1.57%	154.80850167
AF	21%	21%	1.57%	154.80850167
BF	10.5%	21%	1.57%	157.17083643
AN	52.5%	21%	1.57%	158.32197434
BN	10.5%	21%	1.57%	159.03692717
AN	10.5%	21%	1.57%	160.99485349
AF	10.5%	21%	1.57%	161.65239444
AN	5.25%	21%	1.57%	163.27659427
BN	5.25%	21%	1.57%	163.65586792
BF	2%	21%	1.57%	167.41189542
BF	5.25%	21%	1.57%	167.64405393
BN	2%	21%	1.57%	168.13279109
AN	2%	21%	1.57%	171.28337927
AF	5.25%	21%	1.57%	175.27349357
AF	2%	21%	1.57%	191.67498324
AN	52.5%	10.5%	1.57%	212.00857601
AF	21%	10.5%	1.57%	212.91619709
BF	52.5%	10.5%	1.57%	215.54760582
BF	21%	10.5%	1.57%	215.91866808
AF	52.5%	10.5%	1.57%	216.57944685
AN	21%	10.5%	1.57%	216.93902561
BN	21%	10.5%	1.57%	217.50414829
BF	31.5%	10.5%	1.57%	219.03005743
AN	31.5%	10.5%	1.57%	219.25276395
BN	52.5%	10.5%	1.57%	219.49682313
BN	31.5%	10.5%	1.57%	219.58257760
BN	10.5%	10.5%	1.57%	220.34531392
BF	10.5%	10.5%	1.57%	220.34531392
AN	10.5%	10.5%	1.57%	220.34531392
AF	10.5%	10.5%	1.57%	220.34531392
AF	31.5%	10.5%	1.57%	220.52641302
BN	5.25%	10.5%	1.57%	221.11089885
AN	5.25%	10.5%	1.57%	222.15983303
BF	5.25%	10.5%	1.57%	227.86126615
AF	5.25%	10.5%	1.57%	229.85845106
BF	2%	10.5%	1.57%	234.37383823
AN	2%	10.5%	1.57%	235.20599936
BN	2%	10.5%	1.57%	237.22345784

AF	2%	10.5%	1.57%	243.72870328
BN	21%	52.5%	10.04%	945.39254900
BN	5.25%	52.5%	10.04%	957.44951405
BN	10.5%	52.5%	10.04%	960.42056990
AN	5.25%	52.5%	10.04%	962.74311135
AN	10.5%	52.5%	10.04%	963.29601053
AN	31.5%	52.5%	10.04%	966.39988555
BF	31.5%	52.5%	10.04%	967.89766961
AN	21%	52.5%	10.04%	969.08116716
AN	2%	52.5%	10.04%	970.92360929
BN	52.5%	52.5%	10.04%	971.76655571
BF	52.5%	52.5%	10.04%	971.76655571
AN	52.5%	52.5%	10.04%	971.76655571
AF	52.5%	52.5%	10.04%	971.76655571
BN	31.5%	52.5%	10.04%	972.31908460
BN	2%	52.5%	10.04%	981.85753615
AF	31.5%	52.5%	10.04%	1022.52268023
BF	21%	52.5%	10.04%	1030.18693121
BF	10.5%	52.5%	10.04%	1044.94440388
AF	21%	52.5%	10.04%	1067.53084237
BF	5.25%	52.5%	10.04%	1141.99507495
BF	2%	52.5%	10.04%	1165.15383823
AF	10.5%	52.5%	10.04%	1176.87156795
AF	5.25%	52.5%	10.04%	1268.05075527
AF	2%	52.5%	10.04%	1400.32522537
BF	52.5%	21%	10.04%	1754.14392753
BF	31.5%	21%	10.04%	1803.41010039
AF	52.5%	21%	10.04%	1829.52143581
AN	10.5%	21%	10.04%	1844.82991293
AF	31.5%	21%	10.04%	1850.57164380
BN	21%	21%	10.04%	1857.72219705
BF	21%	21%	10.04%	1857.72219705
AN	21%	21%	10.04%	1857.72219705
AF	21%	21%	10.04%	1857.72219705
BN	52.5%	21%	10.04%	1862.23404307
BN	10.5%	21%	10.04%	1862.29851075
AN	52.5%	21%	10.04%	1872.33058723
AN	5.25%	21%	10.04%	1878.58732580
BN	31.5%	21%	10.04%	1887.67474068
BN	5.25%	21%	10.04%	1913.79731217
AN	31.5%	21%	10.04%	1916.21418595
BN	2%	21%	10.04%	1928.82836900
AN	2%	21%	10.04%	1952.68937680
BF	10.5%	21%	10.04%	1955.33569287
BF	2%	21%	10.04%	1977.48666893
AF	10.5%	21%	10.04%	1988.90800195
BF	5.25%	21%	10.04%	2016.47416637
AF	5.25%	21%	10.04%	2035.23433107
AF	2%	21%	10.04%	2161.70052113
BN	21%	10.5%	10.04%	3033.82061487
AN	21%	10.5%	10.04%	3046.08254320
BN	10.5%	10.5%	10.04%	3088.88957647
BF	10.5%	10.5%	10.04%	3088.88957647
AN	10.5%	10.5%	10.04%	3088.88957647
AF	10.5%	10.5%	10.04%	3088.88957647
BF	5.25%	10.5%	10.04%	3090.43731973
AF	21%	10.5%	10.04%	3101.92487413
AN	31.5%	10.5%	10.04%	3103.56328700
BN	2%	10.5%	10.04%	3117.68451073
BF	31.5%	10.5%	10.04%	3118.83046953
AN	52.5%	10.5%	10.04%	3128.64604060
AN	5.25%	10.5%	10.04%	3130.46176820
AF	31.5%	10.5%	10.04%	3131.68051100
BN	5.25%	10.5%	10.04%	3169.39328840
AF	5.25%	10.5%	10.04%	3179.67727400
BF	2%	10.5%	10.04%	3197.16845940
BN	31.5%	10.5%	10.04%	3197.25990533
AN	2%	10.5%	10.04%	3204.20180287
BN	52.5%	10.5%	10.04%	3206.95297560
AF	52.5%	10.5%	10.04%	3207.49794120
BF	21%	10.5%	10.04%	3227.58765233
BF	52.5%	10.5%	10.04%	3275.64390087
AF	2%	10.5%	10.04%	3289.14369987

Stellenbosch University http://scholar.sun.ac.za

F.3 Simulation Results For Bottleneck Probability Dependent Variable

Table F- 5. Bottleneck probability measures arranged from best to worst

SecCCRLoc	SecCCR PC Level	NonCCR PC	Downtime	BottlProb
BN	52.5%	52.5%	1.57%	82.030794
BF	52.5%	52.5%	1.57%	82.030794
AN	52.5%	52.5%	1.57%	82.030794
AF	52.5%	52.5%	1.57%	82.030794
AN	21%	52.5%	1.57%	81.862857
AN	31.5%	52.5%	1.57%	81.752143
BF	31.5%	52.5%	1.57%	81.59246
AF	31.5%	52.5%	1.57%	80.860317
BF	21%	52.5%	1.57%	80.573651
AN	10.5%	52.5%	1.57%	80.460952
BN	31.5%	52.5%	1.57%	80.050397
AN	21%	52.5%	10.04%	79.772302
AN	31.5%	52.5%	10.04%	79.58
AN	10.5%	52.5%	10.04%	79.508651
BN	52.5%	52.5%	10.04%	79.380397
BF	52.5%	52.5%	10.04%	79.380397
AN	52.5%	52.5%	10.04%	79.380397
AF	52.5%	52.5%	10.04%	79.380397
AN	2%	52.5%	10.04%	79.294683
AN	5.25%	52.5%	10.04%	79.245873
AF	21%	52.5%	1.57%	78.829286
AF	31.5%	52.5%	10.04%	78.647063
BF	31.5%	52.5%	10.04%	78.458333
BF	10.5%	52.5%	1.57%	78.23619
AN	5.25%	52.5%	1.57%	78.200556
AF	21%	52.5%	10.04%	77.627063
BN	21%	52.5%	1.57%	76.669365
BF	21%	52.5%	10.04%	76.648413
AF	10.5%	52.5%	10.04%	74.719206
BF	10.5%	52.5%	10.04%	74.495476
AN	2%	52.5%	1.57%	74.21119
BF	5.25%	52.5%	1.57%	73.917222
AF	10.5%	52.5%	1.57%	73.855952
AF	5.25%	52.5%	10.04%	73.512302
BN	31.5%	52.5%	10.04%	73.333968
BF	5.25%	52.5%	10.04%	72.828016
BF	2%	52.5%	10.04%	69.701111
BF	2%	52.5%	1.57%	69.496984
BN	52.5%	21%	1.57%	69.143095
BN	31.5%	21%	1.57%	68.823492
BN	10.5%	52.5%	1.57%	68.751349
AF	5.25%	52.5%	1.57%	68.481429
AN	52.5%	21%	1.57%	68.381349
AF	52.5%	21%	1.57%	68.364841
BF	52.5%	21%	1.57%	68.195238
BF	31.5%	21%	1.57%	68.185873
AF	31.5%	21%	1.57%	68.183333
AN	31.5%	21%	1.57%	68.167143
AF	2%	52.5%	10.04%	67.946984
BN	21%	21%	1.57%	67.594206
BF	21%	21%	1.57%	67.594206
AN	21%	21%	1.57%	67.594206
AF	21%	21%	1.57%	67.594206
AN	10.5%	21%	1.57%	67.230159
BF	10.5%	21%	1.57%	65.72881
AN	5.25%	21%	1.57%	65.719603
AF	10.5%	21%	1.57%	65.483651
BF	5.25%	21%	1.57%	63.834365
AN	2%	21%	1.57%	63.790159
BN	10.5%	21%	1.57%	62.569048
BN	21%	52.5%	10.04%	62.456667
AF	2%	52.5%	1.57%	61.187937
AF	5.25%	21%	1.57%	59.913175
BF	2%	21%	1.57%	59.630397
BN	5.25%	52.5%	1.57%	57.693254
AF	2%	21%	1.57%	54.439365
BN	5.25%	21%	1.57%	51.775635
BN	31.5%	10.5%	1.57%	51.249127
BN	52.5%	10.5%	1.57%	51.115317
BN	21%	10.5%	1.57%	50.999524
BN	10.5%	10.5%	1.57%	50.279048

Stellenbosch University http://scholar.sun.ac.za

BF	10.5%	10.5%	1.57%	50.279048
AN	10.5%	10.5%	1.57%	50.279048
AF	10.5%	10.5%	1.57%	50.279048
AF	31.5%	10.5%	1.57%	50.12246
AN	21%	10.5%	1.57%	50.014841
AF	52.5%	10.5%	1.57%	49.983968
BF	21%	10.5%	1.57%	49.668651
AN	31.5%	10.5%	1.57%	49.604048
BF	10.5%	21%	10.04%	49.554444
BF	31.5%	10.5%	1.57%	49.547857
BF	52.5%	10.5%	1.57%	49.527857
AF	10.5%	21%	10.04%	49.339444
AN	52.5%	10.5%	1.57%	49.141111
BN	21%	21%	10.04%	49.126984
BF	21%	21%	10.04%	49.126984
AN	21%	21%	10.04%	49.126984
AF	21%	21%	10.04%	49.126984
BN	52.5%	21%	10.04%	49.050476
BN	31.5%	21%	10.04%	48.91246
AF	21%	10.5%	1.57%	48.825238
BF	52.5%	21%	10.04%	48.778333
AN	31.5%	21%	10.04%	48.468333
AN	5.25%	10.5%	1.57%	48.344365
AN	2%	10.5%	1.57%	47.907381
AF	5.25%	10.5%	1.57%	47.776349
AN	10.5%	21%	10.04%	47.678889
BF	31.5%	21%	10.04%	47.655317
AN	5.25%	21%	10.04%	47.556905
BF	5.25%	10.5%	1.57%	47.530556
AF	52.5%	21%	10.04%	47.517302
AF	31.5%	21%	10.04%	47.386111
AN	52.5%	21%	10.04%	47.164048
AF	5.25%	21%	10.04%	46.831032
AN	2%	21%	10.04%	46.313175
BF	5.25%	21%	10.04%	46.050635
AF	2%	21%	10.04%	45.920476
BF	2%	10.5%	1.57%	45.197937
BN	2%	52.5%	1.57%	45.181349
BF	2%	21%	10.04%	44.684048
AF	2%	10.5%	1.57%	44.48119
BN	5.25%	10.5%	1.57%	42.914683
BN	10.5%	52.5%	10.04%	42.750714
BN	2%	21%	1.57%	39.284206
BN	10.5%	21%	10.04%	36.191111
BN	2%	10.5%	1.57%	31.478095
BN	5.25%	52.5%	10.04%	21.399206
BN	5.25%	21%	10.04%	14.974444
BN	21%	10.5%	10.04%	14.063571
BF	2%	10.5%	10.04%	13.678413
BN	31.5%	10.5%	10.04%	13.556508
BF	5.25%	10.5%	10.04%	13.329841
AN	52.5%	10.5%	10.04%	12.759524
AF	52.5%	10.5%	10.04%	12.643651
BN	52.5%	10.5%	10.04%	12.468254
BN	10.5%	10.5%	10.04%	12.242778
BF	10.5%	10.5%	10.04%	12.242778
AN	10.5%	10.5%	10.04%	12.242778
AF	10.5%	10.5%	10.04%	12.242778
BF	31.5%	10.5%	10.04%	12.033571
AN	31.5%	10.5%	10.04%	11.88246
AF	2%	10.5%	10.04%	11.338413
AF	31.5%	10.5%	10.04%	11.274127
AN	5.25%	10.5%	10.04%	11.266111
AF	5.25%	10.5%	10.04%	11.195952
BF	21%	10.5%	10.04%	11.083651
AF	21%	10.5%	10.04%	11.073571
AN	21%	10.5%	10.04%	11.049762
BF	52.5%	10.5%	10.04%	10.812143
AN	2%	10.5%	10.04%	10.736508
BN	2%	52.5%	10.04%	6.4610317
BN	5.25%	10.5%	10.04%	5.7036508
BN	2%	21%	10.04%	2.9231746
BN	2%	10.5%	10.04%	0.7724603