ALTERNATIVE LOGISTIC CONCEPTS AND THEIR EFFECT WITHIN THE COMPANY AND THE SUPPLY CHAIN

by

Richard Farr (University of Nottingham)
Sofia Holm (Volvo Aero Corporation)

Abstract:

This document presents findings from the simulation of a focused factory at Volvo Aero, Trollhättan. A series of experiments were carried out, aimed at testing a range of alternative logistic concepts. Each concept describes an approach to the control of workflow within the business, for the purpose of determining its effect within the company, and in the wider context of the supply chain.

Dissemination:

Public

Deliverable/Output n°:	D2.5.3_1	Issue n°:	1.0

Keywords:

Discrete event simulation, logistics, logistic control, supply chain

EXECUTIVE SUMMARY

This work follows a state-of-the-art-review [Kim et al, 2004], describing current techniques and methods to evaluate, simulate and optimise a range of logistic systems, and including a software tool review. Within this task, two use case facilities have been modelled; the focused factory at VAC where turbine exhaust cases are produced, and part of the turbine blades and vanes facility at MTU. A model of each was constructed, and validated.

This report describes the development and validation of the VAC model, and presents experimental results to-date, covering a range of scenarios agreed among the partners. It demonstrates the usefulness of discrete event simulation as a tool for the analysis of logistic performance within a facility, and as a part of a wider supply chain.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

GLOSSARY OF TERMS AND ACRONYMS

Arena	Discrete event simulation software tool produced by Rockwell Inc., used during this study.
CNC	Computer Numerical Control. A means of automating production equipment such as machine tools.
CR	Common resource. An area outside the focused factory, where components from other business units are also processed.
ERP	Enterprise Resource Planning: A broad range of activities supported by a software tool that assists in the management of important business processes, including product planning, purchasing, inventory strategy, interactions with suppliers, management of customer service activity, order tracking, etc.
Extended Enterprise	An enterprise where companies are interdependent and integrated collaboratively in the design, development, manufacturing and delivery of a product to an end user.
Fill rate	The percentage of customer orders that are immediately satisfied.
Inventory	Material held within the business.
Lead time	Lead time is used within this report to mean the time elapsed from when a product first enters a manufacturing system, until it leaves in a completed state. In some situations, lead time might include delays for order processing, etc. We have used lead time to mean *makespan* with the focused factory.
Makespan	Closely related to lead time, although referring solely to time within manufacturing.
MRP	Material Requirements Planning: Process (and supporting software) for determining material, labour and machine requirements in a manufacturing environment. Now superseded by MRPII.
MRPII	Manufacturing Resource Planning: The consolidation of material requirements planning (MRP), capacity requirements planning (CRP), and master production scheduling (MPS).
MTU	MTU Aero Engines.
OPT	xxx
Safety stock	Inventory held for the purpose of absorbing fluctuations (in supply from upstream, or demand downstream).
Simulation	The imitation of the operation of a real world process or system over time. Simulation involves the generation of an artificial history of the system and the observation of that artificial history to draw inferences concerning the operational characteristics of the real system that is represented. [Banks, 1998].
Supply Chain:	A network of connected and interdependent organisations mutually and co-operatively working together to control, manage and improve the flow of material and information from suppliers to end users.
TEC	Turbine Exhaust Casing.

UNOTT	The University of Nottingham.
VAC	Volvo Aero Corporation.
Virtual Enterprise	An enterprise created to add value by selecting business resources from different companies and integrating them into a single business entity.
WIP	Work-in-progress: the amount of material found within the manufacturing system. May be measured in numeric terms, or as a monetary value.

1. INTRODUCTION

A manufacturing system can be created, through investment in machines and the employment of staff, with the aim of meeting a given production target. That system will receive raw materials, and process them to produce desired products. In reality, however, the demands made of the manufacturing system are likely to change both in the short-term (in response to fluctuations within in the supply chain) and in the long-term (as market conditions change, or new products are introduced).

While the level of investment in manufacturing infrastructure is a major influence to its performance, this cannot be varied in a responsive manner, to address immediate issues such as a component shortage or a quality problem. While investment in the form of over-capacity within the manufacturing system – or in high levels of safety stocks – offers a way to ensure that the system can deliver when needed, it may be possible to achieve similar levels of service from a less expensive system if the logistic controls within the system are matched to the demands that are likely to be made.

The purpose of the simulation study described in this document was to evaluate a series of different logistic control principles, under a variety of situations, to show how each would perform. The situations modelled are based upon real events that have taken place at VAC; by varying the logistic controls applied within the manufacturing system it was possible to investigate how the presence of a buffer, or an alternative scheduling rule, etc., might have affected with performance of the system.

The goal is to identify where improvements in the performance of the system as a whole can be made, without requiring substantial investments in new machinery and the like. Performance, in this context, must be considered against multiple criteria, such as the fill rate achieved, the robustness of the manufacturing system and the cost of work-in-progress (WIP) held within the system. The reader will see how some improvements can be made at little or no cost.

1.1. AIMS OF THIS WORK

The aims and objectives of this work may be summarised as follows:

- Increase understanding of how alternative production scheduling and control techniques can affect the performance of a business unit
- Allow the evaluation of alternative scheduling rules, inventory holding policies, staffing levels and factory layouts, without disrupting the real system
- Use multiple criteria such as lead time, WIP and weekly fill rate to measure system effectiveness
- Investigate the relationship between company-level policies and their effect upon a larger supply chain

1.2. STRUCTURE OF THE REMAINDER OF THE DOCUMENT

The remainder of this document is organised into six chapters. Chapter 2 presents an overview of the manufacturing system at VAC, and the scope of the model produced. Chapter 3 describes the modelling approach adopted, and Chapter 4 presents an overview of the logistic control principles that were considered to be of interest.

Chapter 5 discusses performance measurement for the simulated system; an issue that was addressed in collaboration with Task 2.5.1 (supply chain logistics modelling). Chapter 6 describes a series of experiments conducted with the model and presents results. These are discussed in Chapter 7, while Chapter 8 presents final conclusions and suggestions for further work.

2. THE USE CASE MANUFACTURING FACILITY

The production of Turbine Exhaust Casings (TECs) at VAC is achieved in a focused factory, located in Trollhättan, Sweden. The focused factory is an approach to manufacturing in which excellence is pursued by concentrating upon a limited, more manageable set of products, technologies and volumes, as originally advocated by Skinner [1974]. This is done to avoid the contradictions and compromises introduced when trying to bid for every business opportunity. Instead, attention is concentrated upon a subset of activities where the company can deliver competitive advantage.

The TEC (sometimes known as a Tail Bearing Housing) is a relatively large engine component, with a set of intricate features that must be produced by casting or fabrication, plus a considerable amount of machining (Figure 1). There are two product types in the focused factory; the V2500 and the PW2000. The International Aero Engines' V2500 is a medium-sized engine, used in aircraft such as the Airbus A320 family, while the Pratt & Whitney 2000 is larger, being used on aircraft such as the Boeing 757 and the C17 Globemaster transport aircraft. Production volumes for the two types of TEC differ as a result.
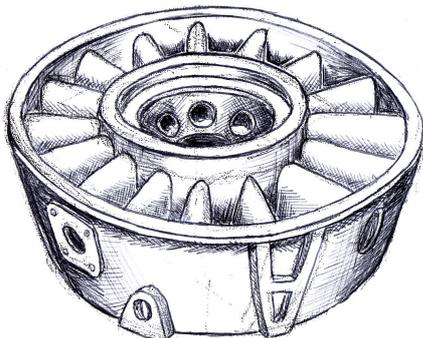
Figure 1: Turbine Exhaust Casing (TEC)

In addition to the two main products, some relatively minor manufacturing operations are carried out on the Trent 500 intercasing. Other tasks may occasionally require the use of a resource within the focused factory, for prototyping purposes or perhaps as a result of a breakdown elsewhere on the site. Similarly, the focused factory can occasionally 'borrow' capacity from another manufacturing system. These events are rare, however, and it was decided not to include them within the model.

A more commonplace form of resource sharking occurs with processes such as inspection tasks, many of which are conducted at shared facilities. Here, components from the focused factory must queue together with components from outside the simulated system. Thus, the focused factory is not entirely isolated from the rest of the Trollhättan plant. Such shared facilities are referred to herein as common resources (CR).

2.1. FOCUSED FACTORY RESOURCES

The focused factory is comprised of over thirty workcentres. Most of these are unique, although there are a few identical fixtures and stations where operations may be conducted in parallel. The workcentres are organised into three areas; broadly speaking, these address machining, welding and assembly. There are also various CRs, outside the focused factory. Within VAC, these workcentres are identified by a four-digit number, but for the model we assigned each a one- or two-letter code.

In addition to the workcentres, there are operators represented within the model. In many cases, operators must be present at a workcentre for a process to be carried out, although at some it is necessary for an operator to be present only at key stages within the operation, perhaps to load the workpiece onto a CNC machine, and start it running, after which the operator is free to do some other job. Some operators are assigned to a wide range of jobs; others are responsible for a single process.

There are many shift patterns being operated within the focused factory. For some processes, a single shift is adequate. Elsewhere, a day and night shift, or three shifts can be seen. There is also an intensive working pattern borrowed from the oil industry, called the Norgeskift ("Norway shift"). The timetable is further complicated by the fact that some bottleneck machines can be loaded and then left to run untended for a few hours, until the next shift begins. In a boom period for aerospace manufacture, the high-value machine tools tend to be operated at as near to 100% of capacity as possible.

Some resources exist in sets, allowing several operations to be conducted at the same time. In a few cases, not all resources within a set are identical. For example, there is a large fixture that can be used to weld a PW2000 or the smaller V2500 TEC, and a fixture that can only be used for the V2500. In such circumstances, it is necessary to establish a fixture choice rule; ideally, the more flexible resources will be allocated last, to minimise the chance of a product having to queue for a the more scarce resource.

2.2. SCOPE OF THE INTERNAL LOGISTICS MODEL

The internal logistics model shows the activities of just one business unit, although it is influenced by events in the supply chain, insofar as it is dependent upon the flow of components from companies further upstream, and its activity is driven by demand from further downstream.

As Figure 2 shows, this means the focused factory model, although highly complex, can be thought of as a 'black box' that receives inputs of material and information, and processes these to produce a stream of goods for the next process in the supply chain. Ideally, the system will not just meet its performance targets in terms of annual production volumes, total value of WIP in the system, etc., but it will also be a 'good team player', achieving a high customer fill rate, even when this requires that variability at the input stage must be absorbed.
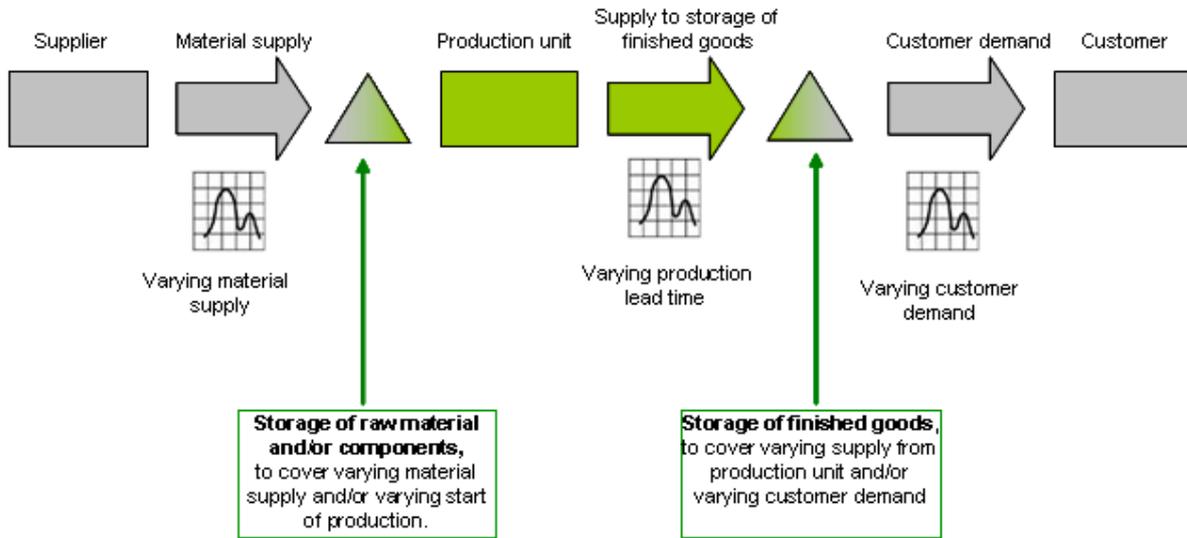
Figure 2: Scope of Internal Logistics Modelling activity

Each completed TEC has a large number of parts, but many of these are items of low value such as fasteners, and are therefore resupplied by a simple two-bin system. The availability of an adequate supply of each of these parts is assumed within the model. For each TEC there are also some major components that need to be represented in greater detail, either because of their high value, or because they have historically been the cause of a work stoppage.

For both TECs, the supply of outer rings and inner rings is modelled in detail. For the PW2000, the set of fifteen struts that are welded between the two rings are also represented in detail. For the V2500, struts had not been a problem, but a three-part housing needed to be studied closely. All other assembly operations are modelled as a process, requiring time and the presence of an operator as in reality, but the flow of small components is not represented. This has made the model more computationally efficient.

2.3. SEQUENCE OF MANUFACTURING OPERATIONS

Given the safety-critical nature of the product, the manufacturing operations are effectively fixed. They cannot be performed in a different sequence or a different way. For both types of TEC, manufacture takes place in a distinct sequence, as shown in Figure 3:
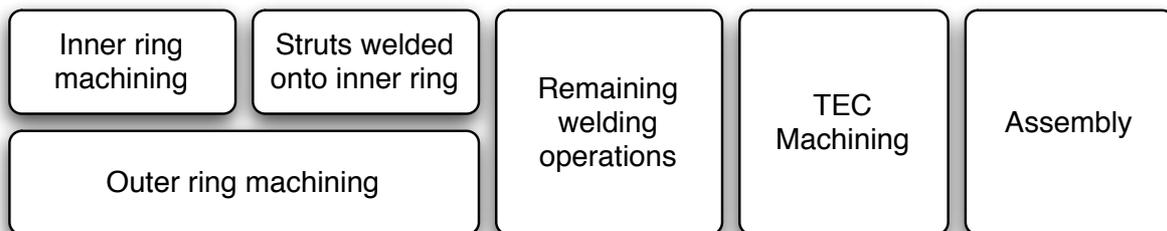
Figure 3: Basic operations sequence for TEC manufacture

The sequences do not correspond directly to the to breakdown of activities into the three areas mentioned in Section 3.1; nor does the physical layout of the focused factory match the progression of operations. (The production volume does not justify a dedicated flow line approach, and the time required to move a workpiece between operations seldom exceeds five minutes, so there is little need to arrange the facility in such a way.)

2.4. SET-UP OPERATIONS

Some of the operations, particularly those where large machine tools are used, require that the machine be set up differently when changing between V2500 and PW2000. Thus, the pattern in which jobs arrive at a workcentre will influence its effective capacity. This is one example of how a simple scheduling rule such as "allow a part to jump the queue if the machine is set up for it" could increase the amount of value adding time at a workcentre – although the consequences of this would need to be studied with care. This practice of temporarily grouping similar parts together is known as *informal batching*, and is a good example of how machine operators sometimes deviate from schedules. If such activities occur at VAC, they should be reflected within the model.

2.5. CONCLUSIONS

This chapter has shown the complexity of the focused factory where TECs are manufactured. As the next chapter shows, much of the complexity can be resolved within the software model, using a series of steps that are of no great complexity when considered individually – although it has also been necessary to make some simplifications and assumptions, each of which are explained in context.

3. MODEL IMPLEMENTATION

The VAC focused factory was modelled using Arena from Rockwell software. (The software selection process was detailed in the earlier project deliverable, D2.5.1_1 [Kim et al, 2004].)

Arena uses a graphical programming paradigm in which a manufacturing process resembles a flow chart. Figure 4 provides an introduction to some of Arena's basic modules, in the form of some slides from a VIVACE training event that took place in March 2006:

Create module

- Introduce new entities to the flow
  - Raw materials, or workpieces arriving from 'outside' the model
  - The arrival of orders
  - Populating the model with kanbans and WIP at 'time zero'
- Parameters include:
  - Entity (product) type
  - Arrival frequency
  - Time of first arrival
  - Number per arrival, and maximum number arriving

CREATE

Delay module

- Use for the simplest processes
  - When capacity is not a constraint (because no machines are needed?)
  - When using limited, observed data
  - Examples: cooling or drying time, delay while the part is sent to a facility that is not modelled, average transit time between sites
- Parameters include:
  - Means of specifying time distributions
  - Identification of the delay as value added, waiting, transfer, etc.

DELAY

Seize, delay, release module

- Use for processes where capacity is a constraint
  - Arriving parts join a queue, then seize the resources once they become available
  - Resources might be machines, fixtures, people... anything!
  - Useful for machining processes
- Parameters include:
  - Means of specifying time distributions
  - List of resources to seize

PROCESS

Decide module

- Use to establish branches in the operations flow
  - Use where the operations flow produces variants – some operations are added (or missed out) for some parts
  - Also good for causing machine tools to be set up, or for making some parts undergo rework
- Parameters include:
  - Number of alternative outcomes
  - Logical conditions for each outcome, or weightings for randomised outcomes

DECIDE Y N

Dispose module

- Causes entities to leave the model
  - Not just for scrap!
  - Finished goods are also 'disposed'
- Just one parameter:
  - Choice of recording the statistics of entities that leave the model, or not

DISPOSE

Other useful modules

- Assign
  - Apply values to global variables
  - Set attribute values that belong to the part
- Batch
  - Group parts together for processing
  - Assemble parts together
- Split
  - Split up a batch, or temporary assembly
  - Turn an 'order' into a stream of components

ASSIGN

BATCH

SPLIT

Figure 4: Introduction to some basic Arena modules

Given that the physical layout of the factory does not exert a significant constraint upon the performance of the manufacturing system (see Section 3.3), it is not entirely necessary to use a modelling paradigm that is based upon the plant layout. Indeed, the most informative representation of the manufacturing system, in the information provided by Volvo Aero, was the Operations Flow diagram. This is more useful than a plant layout in that it is easy to see the sequence of operations and the exact constituents of a queue, even when a part revisits a machine several times. The nodes on the operations flow are events rather than places. Figure 5 shows a small section of the Operations Flow diagram:
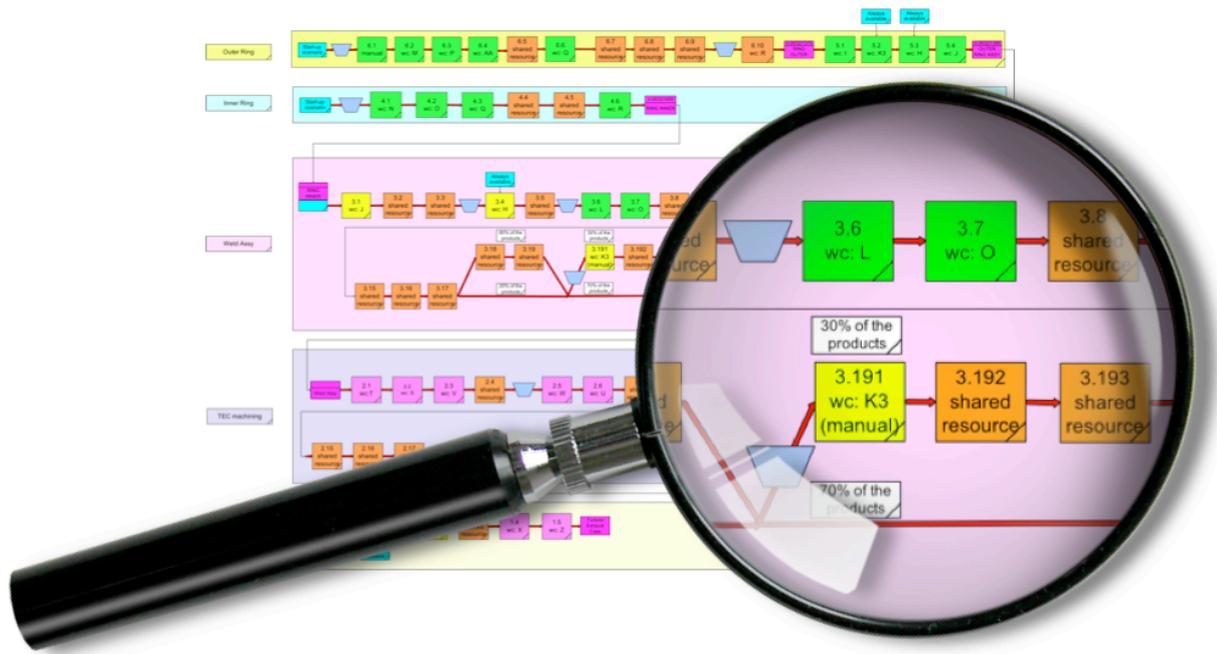
Figure 5: Partial Operations Flow for the PW2000

The full operations flow diagrams for each TEC is too large to be printed on a single page, but its nature should be clear from the sample shown. The diagram is supported by tables that describe the operations in more detail (giving information such as operation times) and tables describing the resources available and their shift patterns, etc. Confidentiality requirements mean that only a small sample of each can be shown here. Tables 1 and 2 provide an illustration of the kind of information that was provided:

Table 1: Extracts from the spreadsheet describing the operations sequence for the PW2000 TEC, including workcentre, transport and setup times, and time distributions for processing.
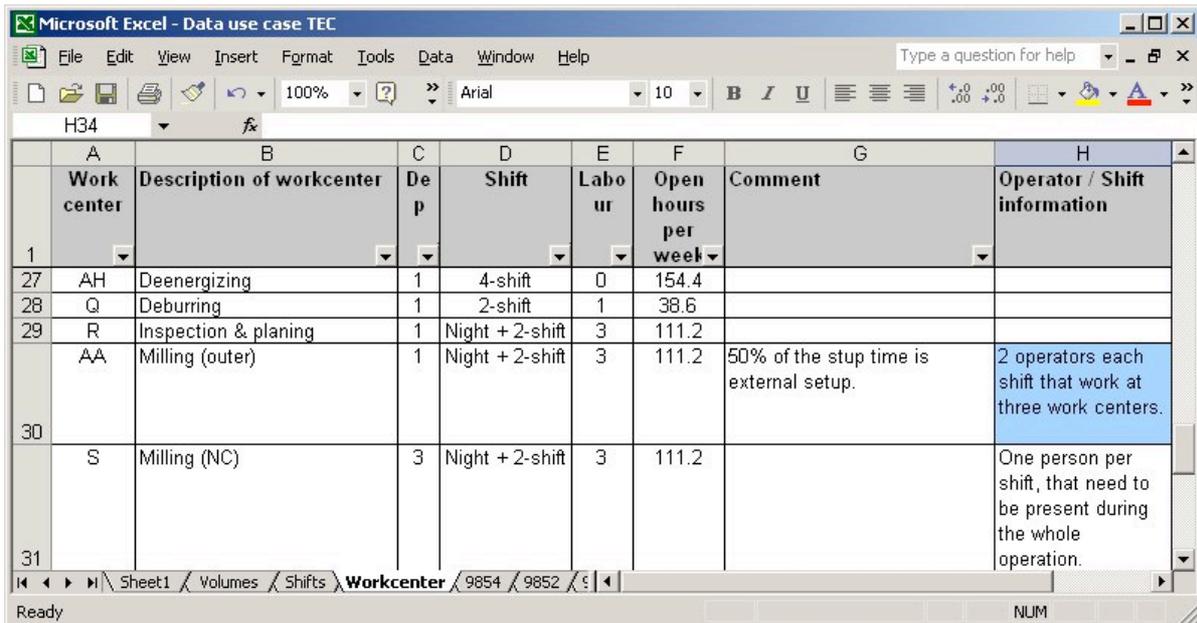
Table 2: Description of workcentres found within the focused factory, including the shift pattern operated

3.1. MODEL LAYOUT

It was aimed to copy the Operations Flow diagrams as closely as possible, within the Arena model, presenting VAC staff with a model that looked familiar, in the hope that this would make interpretation of the models easier, and 'buy-in' to the work more likely.

An initial model of the focused factory made use of a large number of Arena's 'seize delay release' modules, to represent the processes detailed in the operations flow.

These were laid out to match the diagram, and surrounded by coloured blocks meant to identify the department that owned the resource, conforming (as far as Arena's limited palette would allow) to the original VAC document.

This initial model proved the concept of the approach, but the basic 'seize delay release' module could not address the issue of machine set-up. As Section 3.4 explained, some machine tools within the focused factory incur downtime when changing between product types (and in a few cases where a product visits the same workcentre more than once, between operation numbers). For this reason, the basic 'seize delay release' modules were replaced with submodels. Building a submodel is simply a method for keeping a model tidy, by packaging a group of functions such that they appear as a single icon when seen from a higher level.

One necessary departure from the Operations Flow diagrams produced by VAC is that it is necessary to represent PW2000 and V2500 (and to a lesser extent, Trent 500) manufacturing operations within a single model. When an operation begins, the resources it seizes are unavailable not only for other products of the same type, but for the other product types as well. Thus, a single simulation must include the operations flow diagrams for every product. Figure 6 shows the top level of the model. Having the PW2000 and V2500 represented within a single screen means individual stages are shown at a very small size, but the user can choose to zoom in to study any area of interest, or look into a submodel.
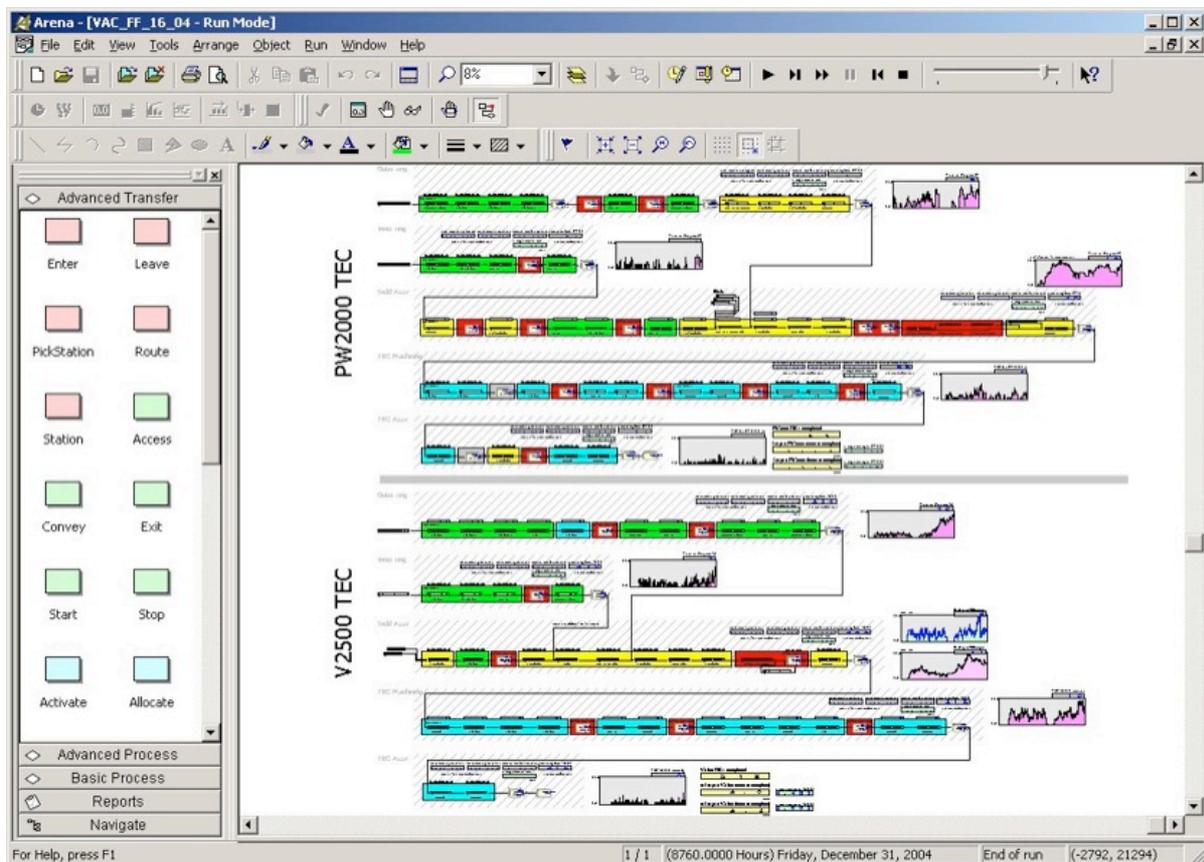
Figure 6: Representation of the VAC focused factory within Arena

3.2. MODULAR MODEL ARCHITECTURE

Organising the model into a series of submodels was necessary, in order to reduce its apparent complexity. The issue of machine set-up, for example, required that a 'decide' module be incorporated ahead of the delay. (Figure 7 shows a simplified submodel.) The diamond-shaped 'decide' module checks the value of a variable that records the current set-up of the relevant machine. This is held as a number that describes the part type and operational stage. If a match is found, the workpiece leaves the decide module at the exit labelled 'true' and moves along the upper path, incurring no additional delays. Otherwise, the workpiece takes the path labelled 'false', and moves through a delay (recorded as non-value added activity), and then reaches an 'assign' module to change the machine set-up variable. It then rejoins the main flow, and incurs a delay (recorded as value added) for the normal manufacturing operation.
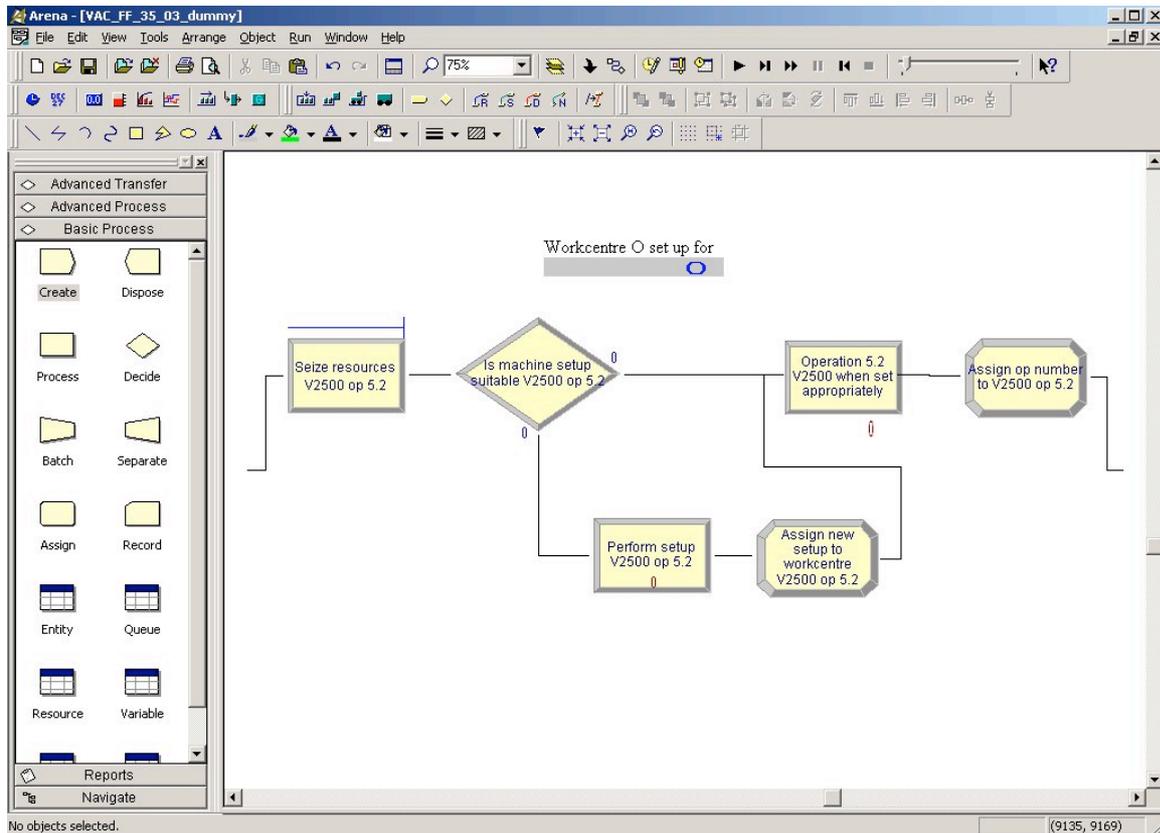
Figure 7: Sample submodel (simplified)

Many operations require no set-up time, in which case the delay for the set-up activity is simply set to zero. No other change is necessary, so a common submodel can be used for most machining operations.

Operations such as welding, where multiple components are joined together, are slightly more complicated, but again they have a common architecture. In addition to the modules shown in Figure 7, they have 'batch' and 'match' modules that hold arriving

parts until a full set of components in available, and then join them together into a single entity for the remainder of the simulation.

The modular approach has been extremely useful, since it has allowed the model to be built up with a reasonable degree of confidence that what had worked in one situation could be made to work elsewhere – with appropriate renaming of variables and modules.

3.3. MODEL SIMPLIFICATIONS

No simulation will ever encompass the full complexity of the real system. Indeed, if it did, it would be failing in one of the principal aims of simulation, which is to allow results to be obtained in a timely manner and at a reasonable cost. The subsections that follow describe the assumptions and simplifications that have been made during the project.

3.3.1. THE MODEL CALENDAR

A major difference between the real manufacturing system and the simulated one is the absence of holidays. In reality, work at VAC is interrupted by a long vacation in the summertime. While it would be relatively simple to represent this within Arena as an 'exception' time pattern, it was felt that such a feature would add no real insights to the model, and would in fact tend to make results more misleading. The time a component spends within a manufacturing sequence is recorded; if a holiday began after the component entered a sequence, its time in system would show a massive increase, despite the fact that no breakdown or quality problem had occurred.

For our purposes, it was judged best to assume that unless a specific scenario dictated otherwise, all simulated weeks were 'regular' ones, with no particular emphasis on meeting short-term targets or flushing the system in advance of a holiday. A consequence of this is that we have to expect a higher annual output from the simulation, since it operates 52 weeks per year. Holidays could be partially compensated for by running the simulation for just 45 – 48 weeks, but this would still not entirely reflect the reality of the situation at VAC, since the focused factory's loading varies throughout the year, in order to achieve a build-up of finished goods to keep the customers supplied during the summer vacation, for example. Thus, the 'real' calendar introduces a certain amount of complexity that could mask a number of minor influences. (A varying pattern of demand is one of the scenarios explored via experimentation, as described in Section 6.2).

3.3.2. REPRESENTATION OF COMMON RESOURCES

Within the focused factory, there are a number of CRs, where operations such as heat treatment and some inspection tasks are carried out. The machinery used for these tasks is not product-specific, nor would it be sufficiently heavily utilised to justify its inclusion within the focused factory. Instead, such workcentres are operated on a sharing basis. Scheduling is done on a simple 'first in, first out' basis, this being the fairest system for a workcentre receiving demands from many different business units.

These resources were a potential source of complexity in the modelling of the focused factory's operations, since modelling the 'seize' and 'release' of some finite capacity resource would have required the modelling of all product types that move through each CR. A brief investigation revealed that historical data of this kind were not available for the entire site, so CRs were represented using simple 'delay' modules, with the delay

being based upon the observed time that a product was away for, when sent to a CR. These delays have a relatively large degree of variance, when compared to the manufacturing operations that take place within the focused factory. This is a natural consequence of the unpredictable demands being made upon the CRs by other business units, parts being away over weekends, set-up operations and varying process times, etc. All of these are already accounted for by the use of historical data to provide a statistical distribution for the simulated delay.

3.3.3. REPRESENTATION OF MINOR LOSSES

In the real manufacturing system, operators take short breaks, and also give up some of their time for briefings, training, etc. These are not modelled precisely. Instead, such minor influences are absorbed into the processing time required for each operation. Likewise, minor problems with machines need not be modelled specifically. Since the duration of operations is based upon historical data, the statistical distribution already takes into account such sources of variability. If it was desired to explore the impact of a machine suffering a major breakdown that put it out of action for weeks, this would need to be modelled specifically, however. The basic 'steady state' version of the model does not allow for any such breakdowns to take place, although an experiment that involved a major machine breakdown was conducted (see Section 6.5).

3.4. MODEL INPUTS

While the earliest form of the model was a stand-alone Arena representation, this did not allow the construction of detailed manufacturing scenarios, such as changing patterns in the arrival of material, or in customer demand. The next step was to make use of a one-dimensional array for each component, detailing the number of parts arriving on each day. This proved functional, but was lacking in accessibility, since it required each pattern of arrivals to be pasted into place, within Arena's list of variables.

An improved user interface was developed, using a Microsoft Excel template, called VAC_inputs.xls, which features a number of named ranges where a user can view or modify the run conditions for the model. This use of a familiar application, available on most desktop computers, means that many people are able to define scenarios. As such, it achieves a democratisation of simulation that should increase the relevance of the work to decision-makers in the production facility.

One important element of each scenario concerned the initial conditions within the focused factory, since this manufacturing system will normally be seen to have a certain level of WIP. If every model run is started in a virgin state where there is no WIP, the model must be allowed to run for a considerable time, to allow it to accumulate material within the flow, and settle into a more realistic state. Only then can experimentation – or a verification against historical observations – take place.

To populate the model with WIP at the beginning of a run was simple enough, making use of additional 'create' modules that were set to operate only once, at the start of the model run. This initial solution was a 'hard-wired' one, where it was difficult to vary the location and quantity of WIP in order to explore scenarios. Any such changes required a working knowledge of the Arena software. As an alternative, the placement of initial WIP was defined within 'VAC_inputs.xls', these being read into the focused factory model and used to permit any initial WIP scenario to be explored. Figure 8 shows the new form for

each Arena submodel, now including a group of modules that check the external file to determine the level of WIP to be injected:
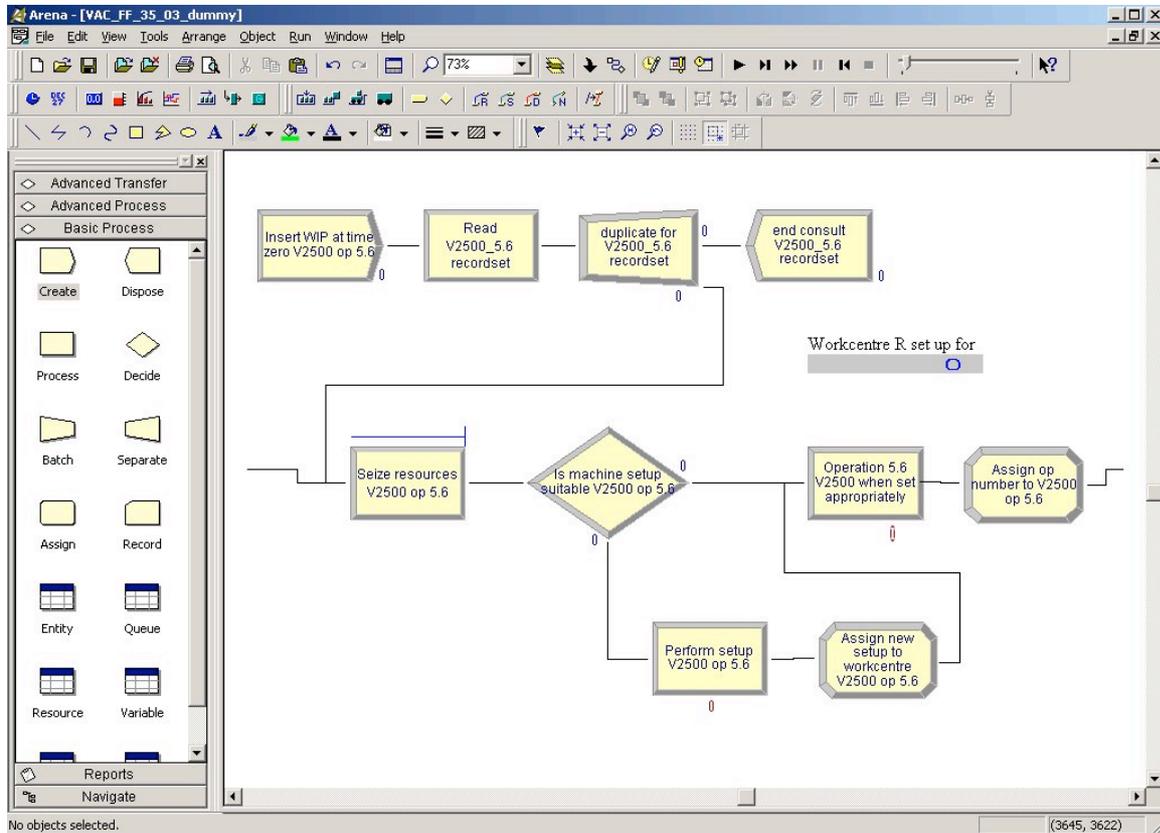
Figure 8: Focused factory submodel, including a means of populating the simulation with WIP at the start of the model run (the top four modules achieve this).

Little's law [1961] can be applied to determine a 'natural' level of WIP for the focused factory, based upon the desired output of the system, and the lead time that has been observed.

3.5. MODEL OUTPUTS

Getting a simulation model to the point where it contains all the resources, parameters and constraints of the real system does not indicate success. It remains necessary to verify that it is performing as designed. Figure 9, as used in the training events delivered on this task, shows the full sequence of activities to be performed during a well-executed simulation project:
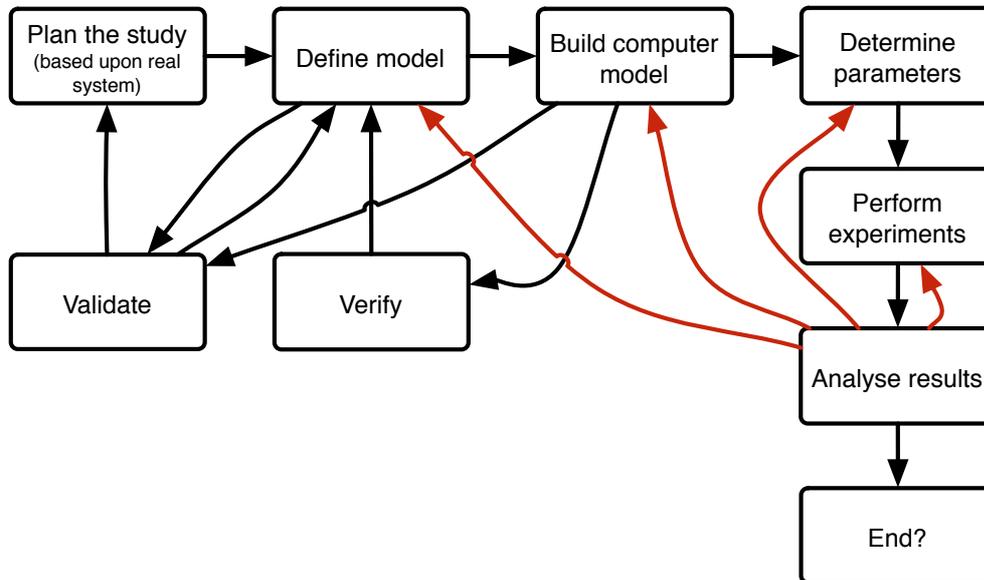
Figure 9: Schematic for model development and operation, including iterations

Validation and verification are addressed in the next section. However, the interpretation of a model at these stages requires that a number of things can be monitored during the model run. Thus, the model must yield results in a comprehensible format.

Arena's built-in reporting function proved to be somewhat misleading. For example, it was seen to describe workcentres as being utilised whenever a component was in the machine, even if the machine were in fact standing idle because the machinist's shift had ended. While it was possible to interpret these reports, with care, it was decided that an alternative method for collecting model statistics was desirable.

The possibility of integration with Microsoft Excel as an input device had already been explored. Now, it was to be used to receive the model output as well. This was achieved using Arena's 'write' module, with the target being specified as an Excel file, VAC_outputs.xls, featuring a number of named ranges where the data would appear.

Originally, the Arena model wrote out data to the spreadsheet each time a product left the line. The exchange of data between the two Windows applications proved to be too slow, however, yielding error messages when one application failed to keep up with the other. As an alternative, a Visual Basic script was added to the Arena model, to be executed at the *end* of a model run. Under the new approach, performance information was stored by Arena itself, in a number of arrays that were exported all at once, at the end of the model run. There was no longer a need to export the data to a specially formatted spreadsheet such as the old 'VAC_outputs.xls'; now, additional model runs would simply be stored as new Excel documents, defaulting to 'Book1.xls'.

This proved to be in the order of twenty times faster, allowing many more scenarios to be considered in a reasonable time frame, and perhaps paving the way for goal-seeking algorithms to be employed in the future.

The Visual Basic script executed at the end of the Arena model run grew in complexity, causing graphs to be displayed, and some basic analyses to be carried out. It was decided, however, that the most elegant solution was to have a 'master' document in Microsoft Excel that could import raw data from any Excel file produced by the Arena

model. As such, the 'master' document provides a dashboard by which the Arena model's results can be interpreted. (Chapter 6 describes the performance metrics against which each model would be compared.) Figure 10 shows the resulting system architecture for the logistic simulation system, as constructed:
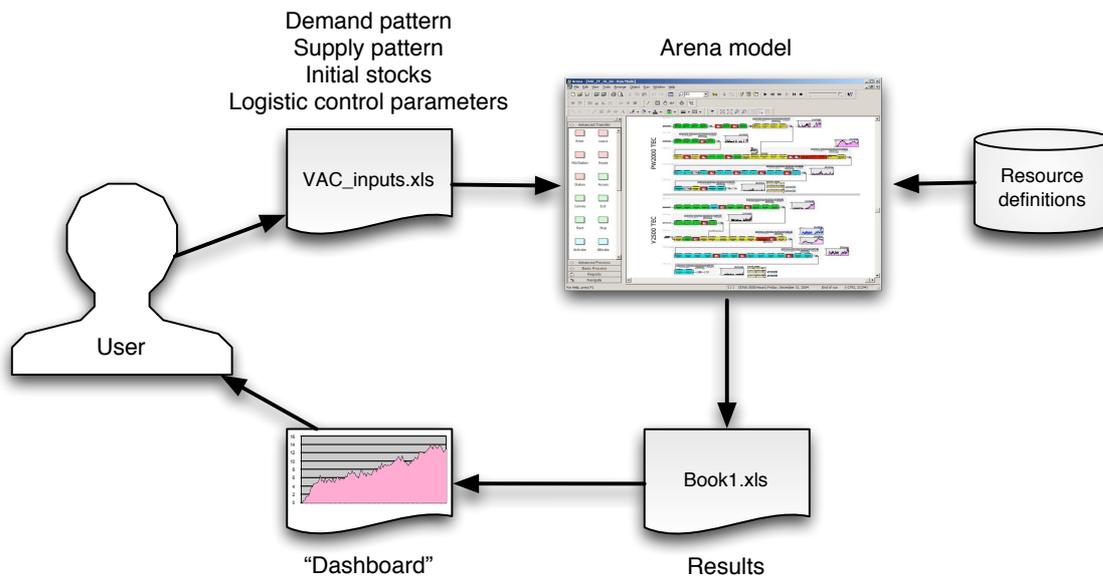
Figure 10: Logistic simulation system architecture

In Month 29 this approach was finalised, and the standard Excel 'analysis' document was developed by VAC staff. It could import data after any model run, and update a series of tables and graphs to showing the performance of the manufacturing system.

3.6. MODEL VALIDATION AND VERIFICATION

Initially, validation work was based upon historical data. Records were available detailing all operations within the focused factory since January 1st, 2004. However, these proved to be of only limited value because they did not record the underlying reason for a given pattern of operations. For example, was a workpiece launched into the focused factory late in response to a situation then occurring, or had there been a supply shortage?

A point was reached where it was decided that the experiments conducted for the purpose of validation would be better if they started from an artificial baseline condition, rather than beginning with the more turbulent reality. Therefore, experiments at this stage made use of an average figure for demand, repeated each week. The initial demand figure was for two PW2000 TECs, and seven V2500 TECs, per week. However, production volumes had increased since January 2004, and it was felt that experiments at the new level (two PW2000 and nine V2500 TECs) per week would be of particular value, in exploring the upper limits of the manufacturing system's capabilities. These figures were entered into the VAC_inputs.xls spreadsheet.

A sensitivity analysis showed that the ordering of work had a small impact upon the performance of the focused factory. If the two PW2000s were not launched together, they caused approximately twice as much time to be spent on machine set-up activities

at the early operations, since the arrival of a PW2000 TEC at a workcentre would usually necessitate a set-up, and and then another if it were followed by a V2500 type. If the two TECs were produced in more equal volumes, the number of set-ups caused could be even more significant. If this were the case, it might be wise to introduce a system whereby *informal batching* (see Section 3.4) is condoned within the model.

Validation involved running the model with the same shift patterns and resource availability as the real system, in the hope of demonstrating a similar level of finished goods output, and similar utilisation at key machines.

Four machine tools were selected for study in greater detail, their usage patterns being stored during the model run, and then written into the Microsoft Excel results file in detail. These were machines that had appeared to represent a bottleneck, at various times, although it was observed that the bottleneck could change, based upon the mix of products being demanded. The amount of usage made of various machines by the two types of TEC are not directly equivalent.

3.7. MODEL LOGISTICS

Within the real system, the scheduling of operations in the focused factory is driven by SAP, a piece of software for Enterprise Resource Planning (ERP). For scheduling in the present-day focused factory, planners use standard lead times (based upon historical data, plus a safety margin) to determine when materials should be launched into the focused factory.

This is a push system; parts are forced into the first stages of the manufacturing system, and advanced to the the next operation as soon as they have been processed, where they join a queue (if any). Since the average time that previous parts have taken to get through the whole manufacturing system is known, it is simple to determine a starting time for each component, such that it ought to be completed in time to be shipped to the customer.

As such, this *is* a system of logistic control – although once a component has been launched, little else can be done to influence its progress. Furthermore, planning that makes use of a standard lead time does not take into account the capacity of the facility, nor its other commitments.

Within the focused factory model, the task performed by the ERP system is represented within the 'VAC_inputs.xls' spreadsheet where the model run conditions are defined. The demand for completed TECs of each type is recorded in a column, to be assessed when considering the fill rate achieved. By default, the supply of raw materials uses the same demand data, offset by the standard lead time that has been determined for each component. Figure 11 shows an example:

Figure 11: Offset of delivery date and the time materials are introduced into the system. Values entered in the 'demand' columns are offset by a standard lead time, to determine the date when a component ought to enter the system.

This first system of logistic control provided the model with sufficient functionality to permit verification of the as-is system, but it remained necessary to investigate alternative systems of logistic control. Chapter 5 reviews these, and discusses their application within a software simulation.

3.8. CONCLUSIONS

This chapter has described how the focused factory was modelled using the Arena software, explaining some of the decisions taken and identifying key features of the resulting model. The process of its validation has also been described. With this tool in place, it was then necessary to determine exactly how it ought to be employed. In the next chapter, the alternative scheduling approaches that it was made to evaluate are discussed.

4. ALTERNATIVE LOGISTIC CONTROL METHODS

The simulation of the focused factory, as described in Chapter 4, was a necessary step in the evolution of a useful model, but it only yielded an as-is representation of operations at VAC. One of the principal advantages of having a simulation model is the ability to explore alternative options without disrupting the real system. A key aim, in the development of the simulation, was to investigate how alternative production scheduling and control strategies could affect the performance of the business unit.

At the earliest task-level meetings, in 2004, a matrix was drawn, showing the options for investigation, as Figure 12 shows. This presented a theoretical maximum of 25 different combinations of plant layout and method of logistic control. Of these, only two had been modelled; simplistic 'heap and hope' scheduling, in the earliest stages of the model, and the MRP-driven representation that had been produced by the time of validation.

	Simplistic scheduling	MRP, MRPII, ERP	JIT / Lean production	OPT	Period batch control
Job shop					
Continuous					
Flow shop					
Focused factory	present-day system				
Project					

Figure 12: Matrix of layouts and logistic control options

The limited time available within the Task meant that some combinations had to be eliminated, since it would not be possible to test every one of these manufacturing systems. Fortunately, some manufacturing layouts are not entirely compatible with the production of TECs. Similarly, some combinations of layout and logistic control disqualify themselves.

The 'job shop' can essentially be discounted, since it does not differ greatly from the VAC focused factory, except in terms of the focus. A job shop is essentially a functional production layout, grouping machines of a similar type together [DeGarmo et al, 1990], which is what we see within the use case facility. However, the job shops is characterised by the large variety of work undertaken. If it were decided to make a series of additional product types within the focused factory, it might lose its focus altogether, reverting to a regular job shop.

Continuous manufacture is seen in the production of commodities such as soap or sheet glass. This approach does not lend itself well to the manufacture of TECs, which require a series of discrete manufacturing operations of varying lengths, and therefore do not lend themselves to line balancing.

Similarly, the idea of reconfiguring the focused factory as a flow shop is unattractive, given that some components visit the same machine more than once, and that the sequence of operations is different for the V2500 and PW2000 product types. Finally, the machines in question are extremely large and expensive. The cost of relocating and duplicating machines in order to achieve a unidirectional flow would far outweigh the benefits in terms of reduced transit time between machines.

The project layout is not appropriate for these relatively small items, being more common for activities such as laying down the keel of a ship, or constructing a building. We are left, essentially, with the present-day focused factory, plus any variants that it is considered desirable to explore, such as the incorporation of a new machining centre.

In terms of the logistic control options, any control method might be applied within the focused factory. The logistic control of other production layouts such as continuous manufacturing is much less flexible, but a focused factory could be driven by a 'push' (MRP) system, a lean production 'pull' system such as Kanban, Optimised Production Technology, Period Batch Control or some hybrid of these.

When considering alternative logistic concepts in principle, it can be useful to remove some of the complexity, by imagining a fictional, simplified manufacturing system. Figure 13, shows an example in which raw materials enter a simulation, move through three processes in sequence, and then exit the model:
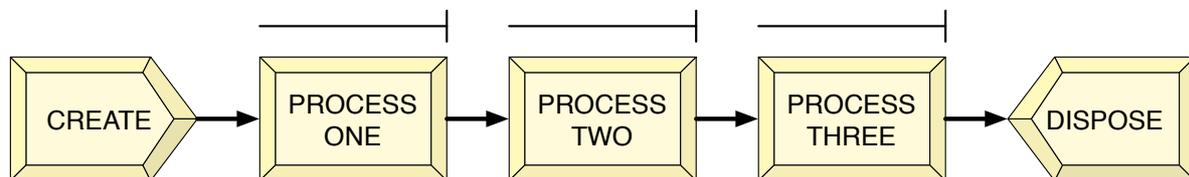
Figure 13: Simplistic arrangement of processes, as used for training in logistic concepts

The sections that follow present a range of alternative logistic concepts, making use of the simplified three-process manufacturing system to allow comparison, and discussing their appropriateness to the focused factory model.

4.1. PUSH SYSTEMS

With the complexities of shift patterns, quality problems and multiple product types, etc., removed for the moment, the role of logistic control can be communicated clearly. If a 'push' system is employed, the manufacturing system will run smoothly until the rate at which parts are created exceeds the rate at which one of the processes is cycled. If the flow of parts coming into the system (or being created) is too fast, the weakest point of the manufacturing system will be simple to identify, as a queue will form at that process. The machine where queue length tends towards infinity is the bottleneck.

If process one is the slowest, it will serve to keep the manufacturing system relatively uncluttered. It will throttle the whole system, since parts can only move on to process two

at the speed of process one. In a worse situation, the bottleneck process comes later in the flow. This allows WIP to build up within the system, and each product's overall makespan will tend to be longer than its predecessor. Some WIP can be useful, providing a smoothing effect or permitting economic batch sizes, but simply starting every operation as soon as possible is a dangerous strategy, achieving high utilisation at the cost of tied up capital, and a loss of flexibility.

Figure 14 shows the consequences of a lack of appropriate logistic control. The time a TEC spends in the manufacturing system can be seen increasing steadily throughout the model run. Parts are still leaving the end of the line, but there is a tendency for each to spend more and more time queuing within the system. This absence of logistic control represents a simplistic 'heap and hope' approach ; the capacity of the system is inadequate for the demands being made upon it, and launching additional material into the facility is the worst possible response to an inadequate output.
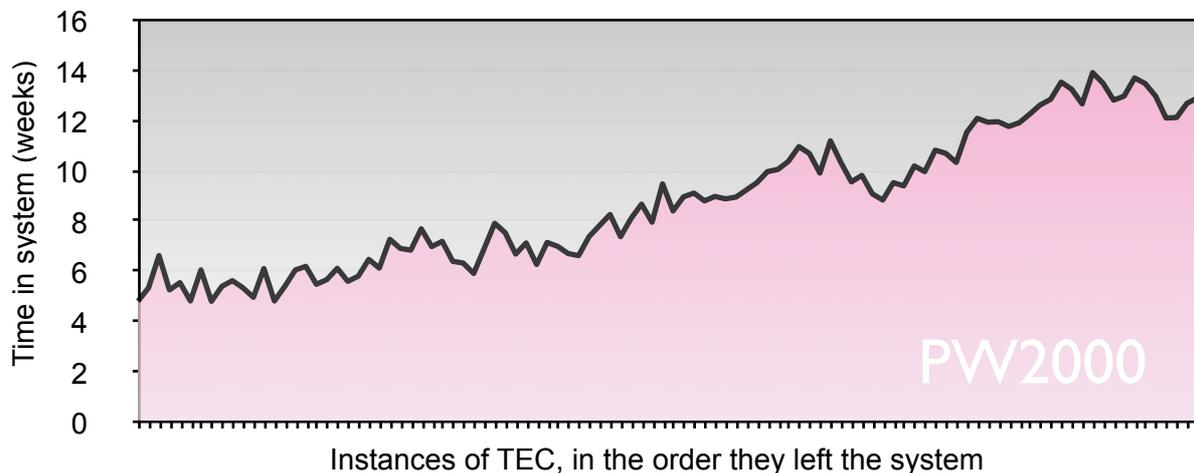
Figure 14: Steadily increasing 'time in system' is a symptom of a system with poor logistic control, where the rate at which products enter the business unit exceeds its capacity

Push systems can be more sophisticated than the most basic scenario described above. They can make use of techniques such as rough-cut capacity planning (RCCP), and can use observed results to determine the standard lead time for products. Even so, a push system cannot be said to be a key enabler for lean manufacturing; nor does the mechanism provide close logistic control within the business unit itself.

A small degree of responsiveness can be seen in some of the 'push' scenarios that were investigated, however; if there is a shortage of parts for the PW2000, for example, then production of PW2000s soon tails off, meaning that these TECs stop appearing in queues at machine tools. At such a time, the V2500 TECs tend to exhibit a reduced makespan, because they no longer have to compete for the capacity available. This reflects reality; if one product type cannot be made for whatever reason, it would be natural to concentrate upon the other product, to keep busy as much as possible, until the problem was resolved.

4.2. LEAN PRODUCTION

Lean production is an operational philosophy based upon targeting seven sources of waste, or 'muda':

- Over-production (making too much, or too early)
- Waiting (queuing products, or people waiting for a job to do)
- Transportation (moving products further than absolutely necessary)
- Processing (doing more than the customer values)
- Inventory (holding more stock than is necessary)
- Motion (people moving more than is necessary)
- Defects (money lost to identifying problems, rectifying them, or scrap)

It is an holistic approach, incorporating many initiatives such as the use of multifunctional teams, mistake-proofing ('poka yoke'), standardisation, ongoing process analysis ('kaizen') and most importantly for the simulation work described here, 'pull' production, achieved via kanbans (discussed in the subsections that follow). Under such a system, operations are driven by customer demand, rather than production expediencies.

Lean production principles can be hard for some managers to adopt. Trained to think in terms of high resource utilisation, the idea of eliminating surplus stocks and conducting operations only in response to demand seems bizarre. To go lean, however, one must accept that a production operation achieves nothing except a waste of energy and the spoilage of some raw material, unless it is conducted to achieve something that the customer values. The reduction of surplus WIP within the facility, for example, should save money and have the additional advantages of responsiveness (ease of change, since there are less products in the system at any given time) and visibility of problems (since inventory levels tend to mask problems, whereas lean production demands that any such problems are confronted).

Key texts on lean manufacturing include Ohno [1988] and Womack et al [1991]. Many of the components of the lean production system fall outside the boundaries of our simulation, since they involve human issues such as ergonomics, or design issues. VAC are not at liberty to vary the design of the TECs, nor to experiment with alternative production methods, due to certification requirements. Thus, our experiments can only address a subset of the lean 'toolkit'. Some of the seven wastes can still be measured, however, in terms of the impact of each alternative system of logistic control.

Two experimental 'pull' systems were created, and are described in the subsections that follow.

4.2.1. KANBAN 看板

Kanban is simply Japanese for 'card', but it also gives its name to a 'pull' system, described by Ohno [1988] as the means through which Just-in-Time (JIT) manufacture is achieved. Demand forecasts play no part in a kanban system. Instead, the interface with the customer is the *last* operation in the sequence. The customer's order causes a token to move *against* the direction of product flow, demanding activity from the preceding operation. Under a kanban system of logistic control, no job should be started unless a token has arrived. This ensures that any product created conforms to the customer's

expectations, and may allow a range of modular variants, without requiring a large stock of finished goods to be held.

Kanbans are not necessarily cards; in some cases it is the arrival of an empty container that triggers the manufacture or supply of a product (or a batch of products, to fill the container). Elsewhere, golf balls have been used as a durable form of kanban. There are also some computerised variants wherein the tokens become 'eBans' or 'faxBans'; a feature that may be useful for the virtual enterprise, where geographical issues make the passing of physical tokens impractical.

Ultimately, the usefulness of a kanban system depends upon the quality of forecast data. If customer requirements can be forecast with considerable accuracy, there are few reasons to adopt a 'pull' system. Where this is not the case, however, the responsive nature of the kanban system can be highly desirable.

Figure 15 shows a model built to demonstrate a how a kanban system could be implemented within Arena. It has three distinct elements; the three-process production facility itself, a separate flow to simulate the arrival and processing of orders (lower left), and an array of modules that populate the facility with a limited set of kanbans at the start of the simulation (lower right).

Figure 15: Kanban system, proof of concept

In the demonstrator, a product is temporarily 'assembled' with a kanban when it enters the workcentre. After the production operation, the kanban is 'disassembled' from the workpiece and returned to a pool from which it can be sent out again when another workpiece is needed. Thus, if no kanban is sent, the product cannot proceed.

This ensures that no work is started if (for whatever reason) there are already an excessive number of parts within an area. This makes Kanban an excellent method of control if, for example, a certain machine is prone to breakdowns or it exhibits considerable processing time variability, since the kanban mechanism stops the manufacturing system from drawing in additional materials until the problem has been resolved.

Kanban is thus an excellent control mechanism, although it is not a planning method as such, which complicates comparison with 'push' systems such as MRP. Kanban works well where a limited range of products are required. If the product or production facility undergoes change, it can impede production, since a key activity is the 'tuning' of the system to determine the best number of kanbans at each process step.

4.2.2. KANBAN SQUARES

The Kanban Squares logistic concept actually does away with kanbans! The tokens no longer exist, being replaced instead by designated areas where goods should be stored, called kanban squares. Typically a kanban square will be marked out on the factory floor using paint or tape, etc., or it might take the form of a shelving system. A key characteristic of this form of control is visibility. Operators must be able to see the input queue for all processes that depend upon them for replenishment. Whenever a part is taken from a square, the operator who can refill that square starts the necessary manufacturing operation. As that operator, in turn, takes material from his own input queue, he empties another kanban square, triggering a manufacturing or ordering activity further upstream, and so on.

Figure 16 shows the implementation of a kanban squares system within Arena. Of particular importance here are the three 'hold' modules that are located before each of the production operations. The final 'hold' module checks the kanban squares where a stock of finished goods is held. Whenever there is a space (stock of finished goods + number of parts in the turning operation < target stock of finished goods), it releases a part into the third operation. The other two 'hold' modules work in a similar fashion, each monitoring the kanban squares immediately downstream, and the WIP in the operation that will refill the square. A range of alternative strategies can be tried, simply by varying the number of kanban squares at each location. Thus, the lowest WIP value for a given level of customer service can be found.
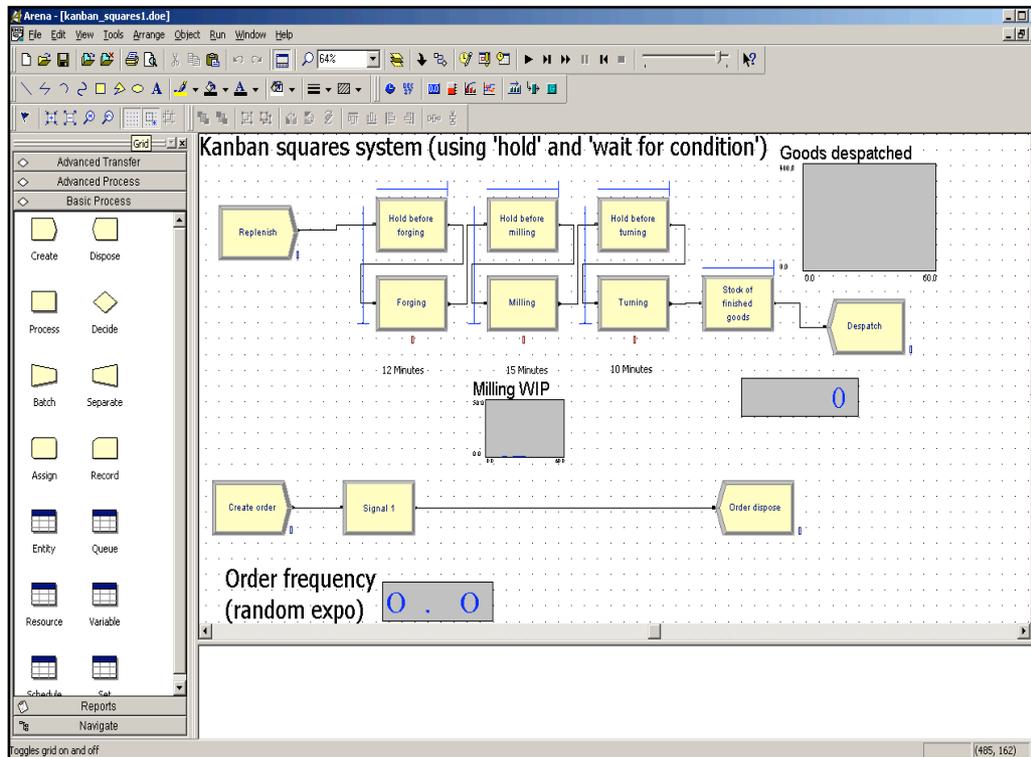
Figure 16: Kanban squares system, proof of concept

It was found to be simple to represent this approach within the arena software, and it was more computationally efficient than the endless recirculation of kanbans. As a system of logistic control for the VAC focused factory, however, it does not fit very well. In the real system it was not common for a planned buffer of any size to exist between manufacturing operations, so the size of most kanban squares would in fact be zero. Given the high value of even a partly-completed TEC, it was felt unlikely that a kanban squares system would be pursued.

4.3. TIME PERIOD SYSTEMS

A small experimental study was made into using Arena's 'hold' and 'signal' modules between operations, such that a workpiece would move between manufacturing operations only under strict control. Figure 17 shows a model in which the flow of manufacturing operations is joined by a parallel flow of information rather than products. This features three 'signal' modules. At intervals determined by a variable representing the Takt time (the frequency of drumbeats), an entity moves along this dummy flow, sending out three signals simultaneously and causing each 'hold' module to release one part.
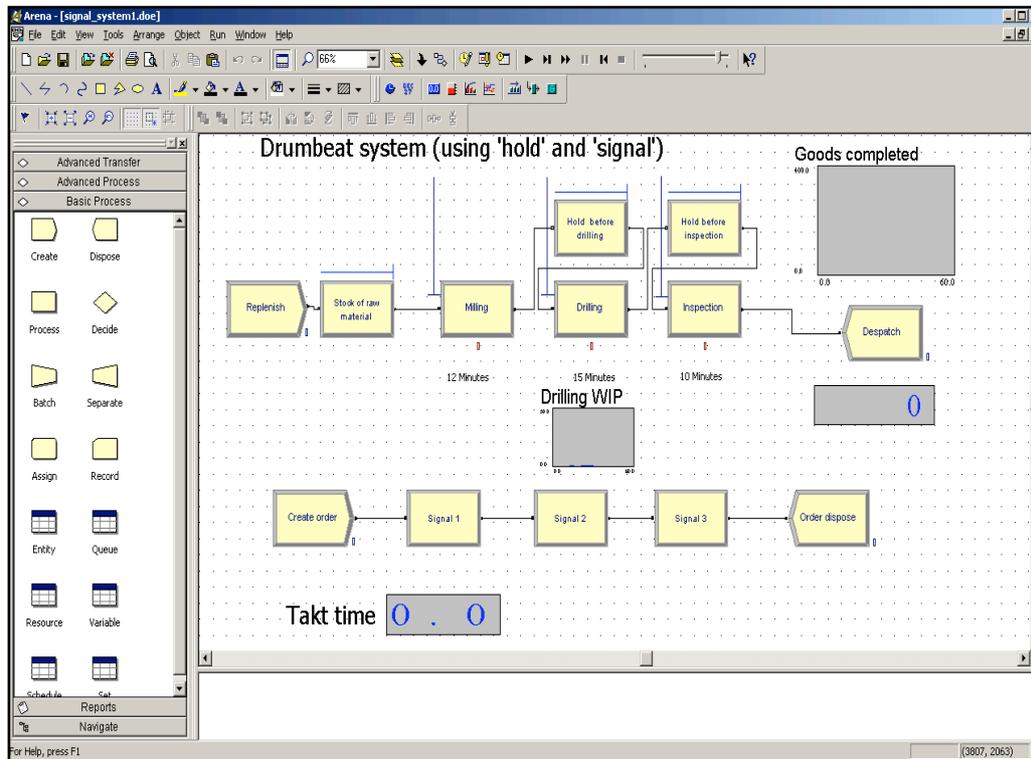
Figure 17: Simple Drumbeat system

In this way, the level of WIP within the facility is rigidly controlled. Providing the Takt time is greater than the duration of the slowest manufacturing process, an excess of inventory can never build up. (If this is not the case, the system will ultimately clog up with workpieces, producing results like those shown in earlier, in Figure 14.)

A drumbeat system of logistic control is best applied to a manufacturing system where each operation requires a similar amount of time, since this ensures that resources do not stand idle for too long, waiting for a drumbeat. The VAC focused factory is not entirely well-suited to a simple drumbeat system for the following reasons:

- It involves the manufacture of more than one kind of product
- Each product takes a different route through the factory
- The duration of manufacturing operations varies considerably
- Some workcentres are visited more than once

While these issues appear to make the a drumbeat system inappropriate, there remains the possibility of using a similar system of logistic control that operates at a higher level. In Period Batch Control (PBC), operations are grouped together, to produce sets of approximately equal duration]. For example, if all operations are broken down into sets that require one week, then manufacturing schedules can be planned on the assumption that a batch of parts will advance from one 'area' to the next, every week. PBC can also be achieved with grouped processes of different total duration, provided the time periods have a common denominator, so it is possible to have one-week and two-week periods within the same system if necessary [Burbidge, 1988].

PBC was originally included in the matrix of production layouts and logistic principles as a result of the literature survey for this work. The approach appears to have been overshadowed by the rise of computer-based simulations and planning tools, although it is still advocated by authors such as Benders and Riezebos [2002]. It is a very practical system, employed on an informal basis in many manufacturing systems, and bears a very close resemblance to the 'Planeringstavla' whiteboards that were found within the focused factory at VAC, where the operational sequence could be seen divided into sets, with magnetic symbols used to show the number of parts present in each.

While individual machines within the focused factory cannot be made to operate to a drumbeat, it was still possible that macro-level planning could be based upon a part moving from one sequences of operations to the next, at regular intervals. It was felt that this would merit closer investigation, since it would closely match the present-day planning system. This familiarity is a key point, when attempting to 'sell' the idea of a simulation to those who work with the real system.

4.4. OPTIMISED PRODUCTION TECHNOLOGY

Optimised Production Technology [Goldratt 1980, Goldratt and Cox 1987] is sometimes referred to as Optimal Production Timetabling, or the Drum-Buffer-Rope system. It is based upon the *theory of constraints;* that the capability of any bottlenecks should determine the logistic control applied. In essence, the productivity of the system will be maximised when the bottleneck machine is run at as near to its maximum capacity as possible. This demands that the bottleneck is never left waiting for material, and that its output is never blocked by an excess of goods. Thus, an OPT-based system of logistic control will include an element of 'pull' upstream, and 'push' downstream of the bottleneck. Since both these logistic control methods were to be evaluated, it seemed likely that implementing an OPT system would involve little more than reusing elements of the earlier models. In fact, as Section 5.6 describes, the implementation approach meant that conducting experiments with OPT was even simpler, being achieved by simply varying certain scenario parameters.

Experiments with the initial focused factory model described in Chapter 3 had already identified the potential bottlenecks within the manufacturing system. These can only really be described as potential bottlenecks, since the processes plans for the two kinds of TEC differ. Thus, a change in the relative volume of PW2000 and V2500 can move the bottleneck. Variability in yield, processing times, or the occasional demands made upon the focused factory by other business units could all move the bottleneck.

4.5. HYBRID SYSTEMS

In reality, few manufacturing systems demonstrate pure examples of logistic controls across a whole business unit. Few companies operate in an environment where demand is sufficiently constant to allow a production facility to evolve to the point where it can balance operation times closely, and introduce control mechanisms such as kanbans. These produce an enviable degree of control, and can reduce waste such as inventory holding costs, but they are disrupted by each shift in demand.

Instead, logistic controls are applied locally. Within a cell or a short sequence of machines, increased visibility and a better understanding of the local issues mean it is

relatively easy to implement an appropriate system of logistic control. Thus, one area might operate of a simple 'first in first out' basis while another seeks to schedule work to maximise the usage of a key asset, or minimise expensive inventory. The modular nature of the focused factory simulation, as developed, means that such strategies can be built into an experiment at any time. For now, the most likely hybrid systems are likely to be combinations of 'pull' and 'push' logistics in a single facility. In fact, Optimised Production Technology (Section 4.4) can be considered an example of a hybrid system.

4.6. IMPLEMENTATION

For a long time, all experiments conducted with the focused factory model concerned 'push' scheduling, this being achieved as described in Section 3.7. It was necessary to evaluate a variety of logistic controls, however, and attention then turned to 'pull' systems. (Some initial experiments into achieving this within Arena were described in Section 4.2.) This introduced a complication that would slow model development considerably, however; any adaptation of the existing focused factory to include a new system of logistic control would make it would be necessary to maintain an additional copy of the model.

Such an approach means that there would need to be a variant of the model for each system of logistic control, in much the same way that Figures 15 – 17 showed different models embodying the experimental approaches. This is a perfectly normal approach to simulation where alternative flows of information or alternative strategies are to be investigated, but it means that any information that is to be updated must be updated across all models. As such, it is a potential source of errors as well as being time-consuming. For these reasons, it was decided that if at all possible, a single model should be developed that could be instructed to apply a variety of logistic controls.

The Microsoft Excel file, 'VAC_inputs.xls' that had been used to define scenarios in terms of the demand patters, schedule for the delivery of raw materials and initial WIP was now expanded to include a means of specifying the strategy to be pursued. Visits to to the focused factory had revealed how its status was monitored via whiteboards where the operations flow was broken down into sequences, in a manner reminiscent of period batch control (see Section 4.3), and our own study of WIP levels within the system had made use of the same boundaries. (As section 2.3 explained, these sequences are not directly related to operation number classification, machine type or geographical location.)

Now, instead of merely monitoring the level of WIP within a sequence, a feature was added whereby a target value for the number of parts within a sequence could be specified. If, for example, we specified a target WIP level of 5 parts for V2500 TECs in final assembly, then a TEC of that type, having completed machining, would no longer advance into the assembly stage automatically. If there were already five parts in assembly (or more, in the event of an initial 'overstock' scenario) then it would be made to wait in an output buffer. During this wait, the machined TEC would count against the total number of parts within the machining sequence, with its own target WIP value, and might thus cause a welded TEC to be delayed in its own output buffer... and so on upstream, right back to the release of raw material into the facility. Varying the target WIP figures in each sequence, for each product type, provides a capability for a considerable degree of experimentation, *without the need to be an expert user*. A basic knowledge of Microsoft Excel is all that is required. Furthermore, by setting *very* high

target WIP levels, the user can open the 'flood gates' to allow material to move through the facility without waiting for a 'pull' – effectively switching between 'pull' and 'push' logistics within a single model, while making no changes to the model itself... and since target WIP levels can be set on a sequence-by-sequence basis, it is possible to investigate a combination of 'pull' and 'push' at this same time – achieving a form of OPT . The result is a robust model that can be configured by planners rather than simulation specialists. It allows a broad range of scenarios to be evaluated at speed, and the model is somewhat protected against changes in the operations to be conducted or the resources available, since it need only be changed once to allow a new set of experiments across the full range of logistic control scenarios.

Figure 18 shows the relevant section of the 'VAC_inputs.xls' spreadsheet, and Figure 19 shows the output buffers within the focused factory model, operating in response to the parameters entered in the spreadsheet.

Figure 18: View of 'VAC_inputs.xls' showing the means of specifying initial WIP, and the desired level of WIP for each sequence. To operate a 'push' system, desired WIP levels can simply be set to a very large number.
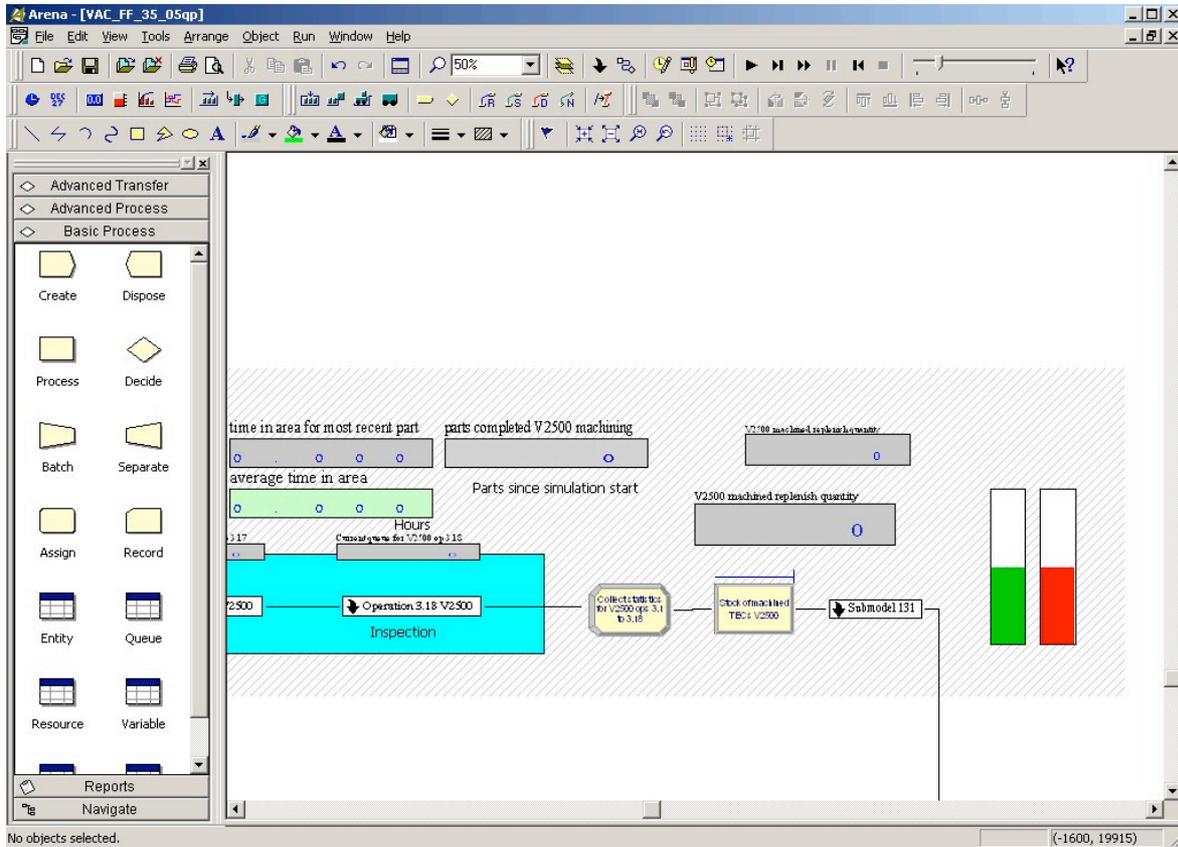
Figure 19: Output buffer at the end of a sequence of machining operations. The 'hold' module will retain any completed parts until the 'replenish quantity' (desired WIP level for the next sequence) is greater than zero.

4.7. CONCLUSIONS

This chapter has described a range of methods by which the flow of products within a manufacturing system can be controlled, and discussed their representation within an Arena simulation. The implementation of the most appropriate of these, within the VAC focused factory model, has been described.

At this stage in the project, a detailed model of the use case facility had been produced and validated, in accordance with the project objectives. It remained necessary to determine how the performance of the manufacturing system should be interpreted, before experimentation could begin in earnest. The evolution of a performance measurement approach is the subject of the next chapter.

5. SYSTEM PERFORMANCE METRICS

In order to perform experiments with the simulated system, it was necessary to decide what was to be measured. The sections that follow describe the selection of a set of metrics, evolved in collaboration with staff on Task 2.5.1 since a common system of metrics could then be employed at both the micro- and macro-level (e.g. for the logistics within a business unit, and those of a whole supply chain).

5.1. LEAD TIME

One of the principal measurements for the manufacturing system is the time that parts spend within it. For the purposes of the focused factory simulation, lead time is measured from the start of the first operation of the outer ring's machining until the completion of the last operation of the assembly. It was decided not to include the time the raw material spent in storage, since it might have been delivered early by the supplier, and this is something that the performance of the manufacturing system could not affect.

Lead times are recorded for every part produced, the times being stored in an array during the model run, and then exported to a Microsoft Excel spreadsheet as described in Section 4.5. These can then be shown in graph form.

When studying plotted lead times, it must be remembered that the x-axis is not time, as such. It shows the order in which the parts came out of the manufacturing system, certainly, but measures of lead time alone do not show delays. For example, the 61st and 62nd TEC to be produced might both have a lead time of six weeks... but they might have left the manufacturing system two months apart if some severe shortage of raw material had arisen. Thus, manufacturing lead times should not be considered in isolation.

In addition to the overall lead time from the start of the outer ring until finished assembly (selected because the time required for outer ring machining is greater than that for the inner ring), a detailed log of the time each product spent within each sequence of manufacturing operations is kept. Each entity is 'time stamped' as it passes through the key stages in question, these matching the boundaries to the sequences that are used for logistic control (Section 5.6). Table 3 shows the components of lead time that are measured:

Time spent undergoing outer ring machining operations	Time spent undergoing inner ring machining operations (and also some initial welding)
Dwell time in output buffer	Dwell time in output buffer
Time spent undergoing welding operations	
Dwell time in welding output buffer	
Time spent undergoing TEC machining operations	
Dwell time in TEC machining output buffer	
Time spent undergoing assembly operations	
Dwell time in assembly output (finished goods) buffer	

Table 3: Lead time components measured within the model

While early versions of the model yielded only an overall representation of the time each component spent within the manufacturing system, the greater volume of information now being exported made it possible to see where a product had spent time. Figure 20 shows an example, in which a problem can be seen to have a knock-on effect throughout the factory, requiring considerable time before conditions begin to resemble their earlier steady state. (Details of the experiments conducted appear in the next chapter.)
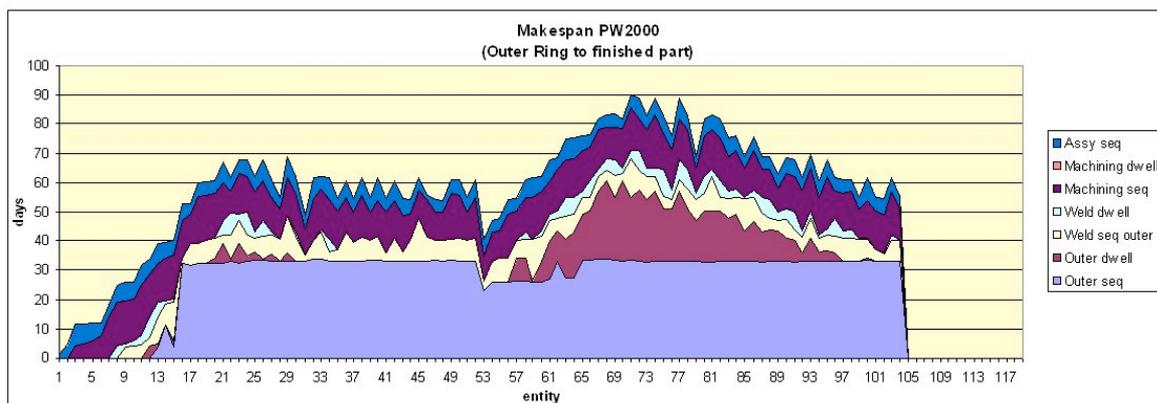
Figure 20: Stacked area graph showing lead time for the PW2000 TEC

It should be noted that dwell times in output buffers will always be zero where a 'push' system of logistic control has been selected. Furthermore, it is not possible to show both inner and outer rings in a single graph, because these lead times should not be stacked; the operations can take place concurrently.

5.2. CUSTOMER SERVICE

The ability of a manufacturing system to satisfy the demands of customers is somewhat at odds with other goals such as reducing the level of WIP in the system. It is, however, an important measure in that it makes a case for investment in safety stocks or surplus capacity when a pure 'optimisation' might produce a fragile system that is easily disrupted.

Within the model, customer service is measured in terms of the fill rate for the products ordered. (We rely upon the planner to refuse unreasonable orders, or at least bear them in mind when interpreting the results, since the capabilities of the manufacturing system are finite.) The fill rate is calculated on a weekly basis, and the number required each week includes any backlog of unfulfilled orders, carried over from the previous week. Thus, a fill rate of 100%, week after week, is the theoretical goal... although this may be achievable only at the expense of surplus capacity and excess WIP.

Figure 21 shows a graphical representation of the fill rate achieved during a model run. It allows the 'health' of the system to be seen at a glance, and its long-term viability can also be assessed, since a failing system's fill rate will decline steadily as backlogs build up. An average fill rate for the whole model run is also calculated.

Customer Service

Definition:
Customer Service is measured as an fill rate of planned demand per week.

	PW2000	V2500
Average fill rate	86.5%	
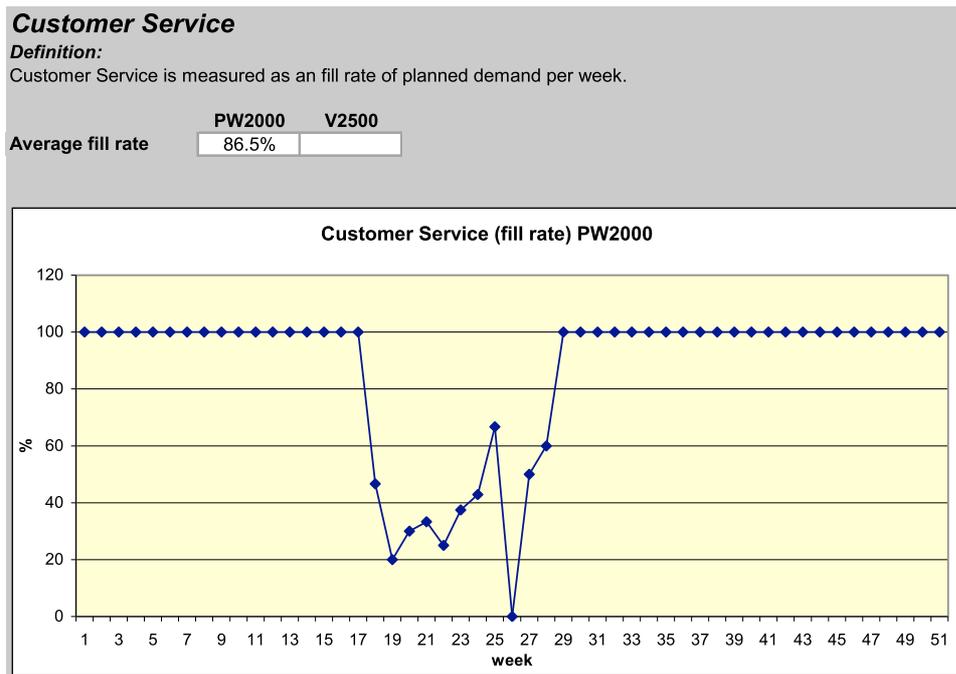
Figure 21: Fill rate graph, showing the degree of customer service achieved during a problem scenario

5.3. WORK-IN-PROGRESS

Work-In-Progress (WIP) is an important measure for the focused factory, since it shows how many of the high-value products are needed within the system, to maintain an

adequate customer service level and to provide a good flow of workpieces to bottleneck machines where high utilisation should be a goal.

Within the model, the level of WIP is sampled every 24 hours, for each manufacturing sequence and each output buffer. (A sensitivity analysis was conducted, varying the time of day when WIP levels were assessed, but it was found to have a minimal impact. This is perhaps unsurprising, given that TECs spend at least several days moving through a sequence.)

The accumulation of WIP is a clear symptom of inappropriate or inadequate logistic control, although problems may only appear when the model is subjected to a problem scenario. In an early stage of model development it was necessary to run the model in a 'slow' mode where animation could be seen. This allowed queue lengths to be studied, and graphs within the Arena model to be seen. In the newest version of the Arena model, WIP information is written to an Excel spreadsheet at the end of the model run, where it can be displayed graphically. Either way, bottlenecks can easily be identified. Figure 22 shows an output exploring the impact of a problem scenario:
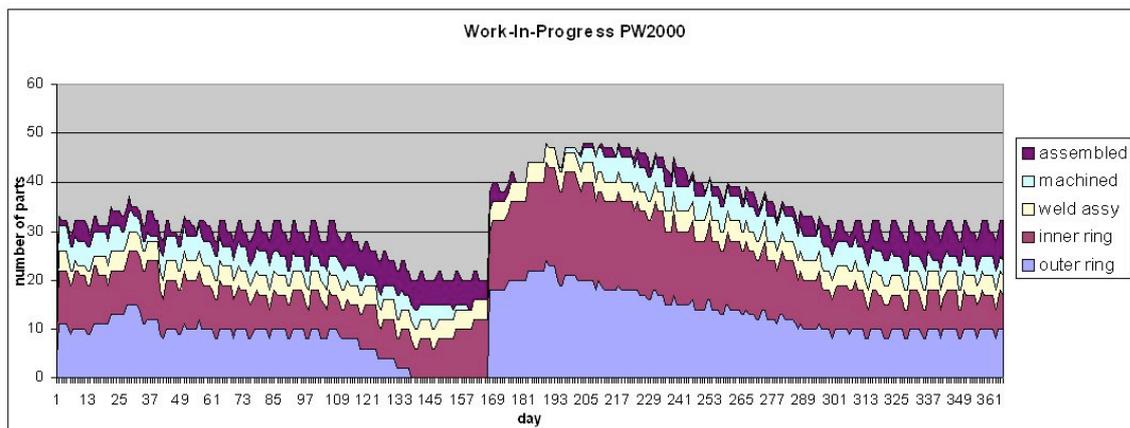
Figure 22: Sample output showing the level of WIP within the focused factory

5.4. TIED UP CAPITAL

Lean manufacturing demands that excess inventory be eliminated, since this is a source of waste. It is therefore necessary to have a method of determining the cost of the WIP, while exploring the impact of reducing the levels allowed to accumulate within the system.

Within the system described in this document, the value of tied up capital is measured as an average figure for the whole period. Standard costs have been determined for TECs and parts, depending upon the sequence in which they are found. We thus have approximate costs for raw materials, inner and outer rings, and TECs undergoing welding, machining and assembly, plus finished goods. Tied up capital in each sequence can then be calculated by multiplying the standard cost by the average WIP level observed.

Adding together the figures for tied up capital in each sequence for both TEC types yields a single figure that can be used to compare different systems of logistic control.

Figure 23 shows a conceptual aspiration for this work, allowing WIP and service levels to be compared under several different systems of logistic control.
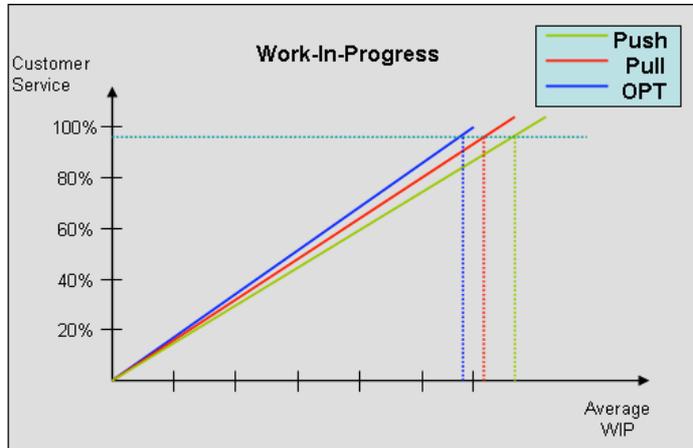
Figure 23: Proposed format for comparison of alternative logistic systems, based upon tied up capital and service level achieved.

5.5. RESILIENCE

Since the VAC focused factory is a complex system, and the demands placed upon it can vary, it would be naïve to seek an 'optimal' solution that assumes predictable demand, quality and machine reliability, etc. Instead, there may be some value in having a sub-optimal system that can recover well after a problem.

Resilience, then, is the ability of a system to return to its planned position after a disruption. This is measured as number of weeks that it takes until the backlog is back to zero after a disruption. The backlog is measured as number of parts that are late compared to the schedule defined in the 'VAC-inputs.xls' spreadsheet.

Figure 24 shows an example produced automatically by the simulation system, where the resilience of the system can be seen, and compared between different runs. Expressed as a single, numeric value for each TEC type, the resilience is the number of weeks during which there was any backlog.

Resilience

Definition:

Ability of a system to return to its planned position after a disruption.

Number of days that it takes until the backlog is back to zero after a disruption.

	PW2000	**V2500**
Resilience	11	1
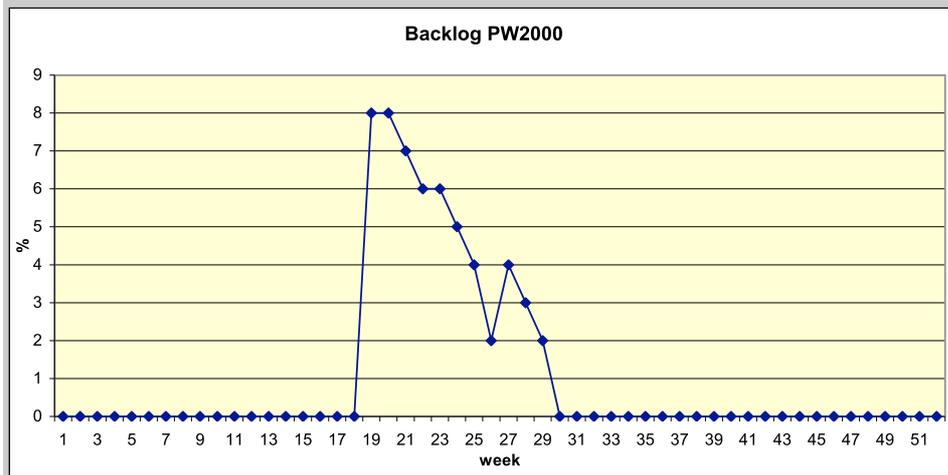
Figure 24: Resilience, expressed in terms of the time required to return to a normal condition after disruption.

5.6. ROBUSTNESS

Resilience (Section 5.5) is a measure of the ability of the system to recover from a disruption that causes a backlog, but there is also the possibility that a system can absorb a certain degree of variability, and still meet its commitments. A robust system will make use of surplus capacity, WIP or flexibility inherent in the system, and as such there should be a means of assessing the value of the system's robustness, since otherwise this might be traded away when trying to increase utilisation, or reduce WIP.

Customer service (Section 5.2) provides a measure of the success or failure of the manufacturing system, but this does not take into account the level of stress that has been placed upon the manufacturing system, in terms of variability in demand patterns, supply shortages, etc. Robustness is more difficult to measure, but an indication of the focused factory's sensitivity to the input variations defined for each scenario can be seen in the observed deviation of manufacturing lead time. The level of variation can be run with the model operating under a variety of systems of logistic control, and the results compared.

5.7. UTILISATION

Given the high value of some of the machine tools found within the focused factory, and the constraints they impose, it seems reasonable to assess the system of logistic control in terms of how much usage is made of key assets. High utilisation identifies a bottleneck, but low utilisation might indicate that the machine is being kept waiting, which

is disastrous if the machine in question is the bottleneck operation. This would mean that the system of logistic control being used had caused the manufacturing system to become *too lean*.

In a real manufacturing system, utilisation will seldom exceed 70%. Attempts to pursue a goal of 100% utilisation, where the machine cycle time is exactly that calculated in the process plan, will normally lead to targets being missed, simply due to misfortunes and the failure of machines, operators and workpieces to be ready at the same instant. Figure 25 identifies some common causes of losses, in approximate proportions:
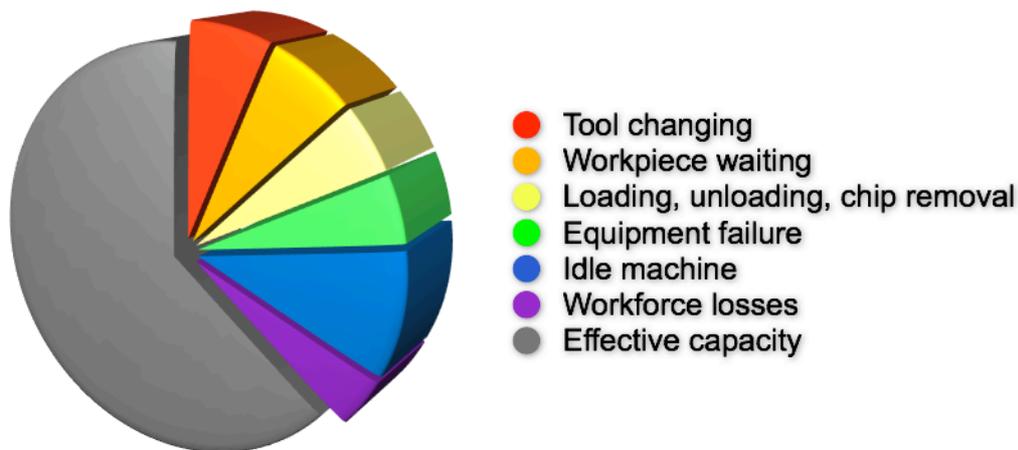
Figure 25: Losses in effective production capacity

The focused factory simulation was constructed using historical data, recording the time a workpiece arrived at a workcentre, and the time it left. Shift patterns were taken into account when calculating the total time a workpiece would spend at a machine, but minor delays such as operator breaks are absorbed into the total. Tool changing times were listed separately, since the order in which parts arrived at workcentres could influence the total amount of time spent setting machines (Section 2.6), but loading and unloading were included in the operation time, since they happen for every workpiece.

Equipment failures have not been programmed into the basic focused factory model; there are no random machine tool breakdowns, although it *is* possible for a specific scenario to investigate the effect of an equipment failure.

Four workcentres were identified as having acted as a bottleneck under certain circumstances, and the utilisation of each of these was recorded, on a weekly basis. Figure 26 shows the representation of machine utilisation generated:
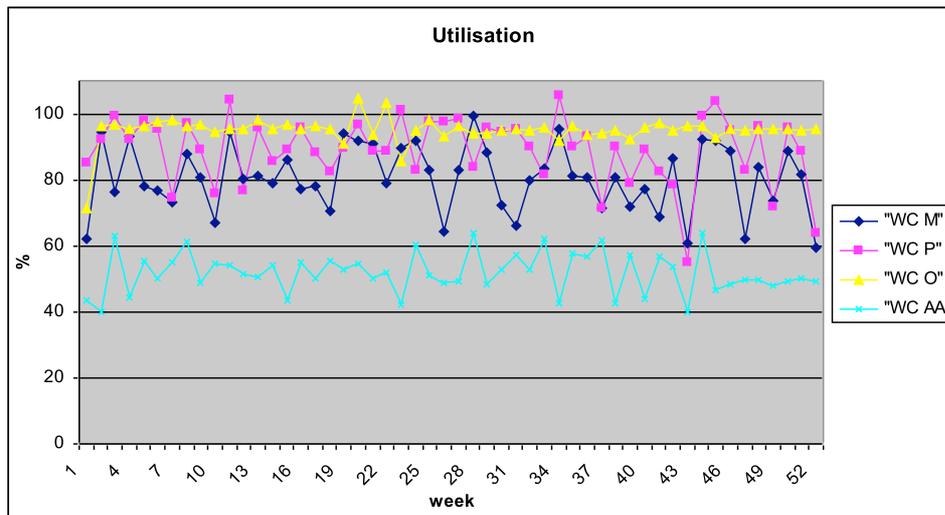
Figure 26: Sample workcentre utilisation graph

The utilisation statistics seen here can occasionally pass 100% due to the statistical collection method. A workpiece is time-stamped when arriving at one of these key machines, and again when leaving it. The difference between the two, which includes non-value-adding tasks such as machine set-up, is then added to a running total of hours used that week. Sometimes a part commences processing at a workcentre well before midnight on the 7th day of the week, but when it leaves, the processing time is assigned to the total for the following week. This effect was not judged to influence the results sufficiently to necessitate a more complicated modelling approach.

Utilisation figures from the model provide a useful indicator when selecting scenarios to explore. Figure 26 suggests that a breakdown at Workcentre 'O' would not form a useful experiment, unless capability of the same type existed elsewhere within Volvo Aero. If this machine suffered a breakdown, when already being operated on a multi-shift basis, it would be almost impossible to eliminate the backlog, making it impossible to assess the resilience of the business unit, no matter how it might be scheduled.

5.8. CONCLUSIONS

This chapter has described seven metrics by which the 'health' of the use case facility might be measured within a computerised 'dashboard' function, interpreting the results of each simulation run. The metrics available for study at the end of each simulation run are as follows:

- Lead time (for individual sequences as well as the whole product)
- Customer service (fill rate, taking into account any ongoing backlog)
- Work in progress (number of parts in each sequence)
- Tied up capital (total value of WIP found in the system)
- Resilience (time to get back on track after a given disruption scenario)
- Robustness (lead time variability after a given disruption scenario)
- Utilisation (percentage of available time, for a given shift pattern)

While it would be extremely unwise to maximise performance against just one or two of these metrics, the complex system of trade-offs between them becomes clear when alternative logistic control strategies are tried. In the chapter that follows, a number of scenarios are described, each being simulated with the model operating either 'push', 'pull' or OPT scheduling. The results can then be interpreted in terms of the the metrics presented in this chapter.

6. EXPERIMENTS CONDUCTED

The model described in the preceding chapters was run to assess the performance of various logistic control principles when confronted with a number of scenarios. Beginning with a basic scenario in which there are no 'surprises', the model was subjected to a series of deviations. Instead of attempting to optimise the model, to maximise its efficiency under ideal conditions, the experiments that were selected aimed to demonstrate how the present-day focused factory (or a similar system with another system of logistic control) might cope with problems.

The Prussian Field Marshal, Helmuth Graf von Moltke said that "No plan survives contact with the enemy" and although the expression is military in origin, it is no less true for manufacturing. The enemy, in this case, can be any of a number of unpredictable events, which we call *deviation scenarios*, to be compared with performance achieved under 'normal' conditions.

In the sections that follow, a 'push', 'pull' and OPT approach to logistic control are each applied to a scenario, beginning with normal conditions and then exploring the deviation scenarios.

6.1. EXPERIMENTATION UNDER 'NORMAL' CONDITIONS

Our earliest experiments were aimed at achieving a steady state with performance that resembled that of the real manufacturing system. For this, demand for each TEC was kept at a constant level, based on average customer demand. This equates to two PW2000 and nine V2500 TECs per week. This level is very close to the theoretical maximum output of the focused factory, as can be seen by a simple *static simulation* exercise; multiplying the number of hours required for each processes by the number of products passing through, and comparing this with the number of hours each machine can run per week.

Fluctuations in demand would make the results of other experiments less clear, so this steady level was used for every experiment except the demand fluctuation experiments. (It is simple to vary demand patters via the 'VAC_inpts.xls' spreadsheet, if it were desired to investigate 'combination' scenarios as well.)

The 'normal' conditions experiment may sound simple, but it was actually one of the most difficult to carry out satisfactorily because it represented a final verification stage for the model. Under a 'push' system of logistic control, it is a representation of the present-day system, and should therefore produce results that conform to the real system. For a long time the performance of the simulation was worse than that of the real system, but continued investigation revealed practices on the shop floor that saved a considerable amount of time, such as where an operator was able to leave an automated process running for a period, and do another job. Some of these practices had not been described in the information supplied, but with the information in place, the simulation's performance came to match that of the real system.

6.1.1. STEADY STATE, 'PUSH' SYSTEM

This is a representation of the 'as-is' system, and should therefore achieve a steady state, after a warmup period. (Adding WIP into the simulation at 'time zero' via the mechanism described in Section 3.4 had been found to shorten the model's warm-up time considerably, allowing simulations to settle down and begin producing useful data much more quickly.)

To conduct this experiment, the model was set to suffer no unusual breakdowns or quality problems, and the demand pattern was entered into the 'VAC_inputs.xls' spreadsheet, using regular quantities of two PW2000 TECs per week, and nine of the V2500 type. The model was then made to run for a year, and the results exported to a spreadsheet for analysis in the dashboard. Figures 27–29 show the results obtained.
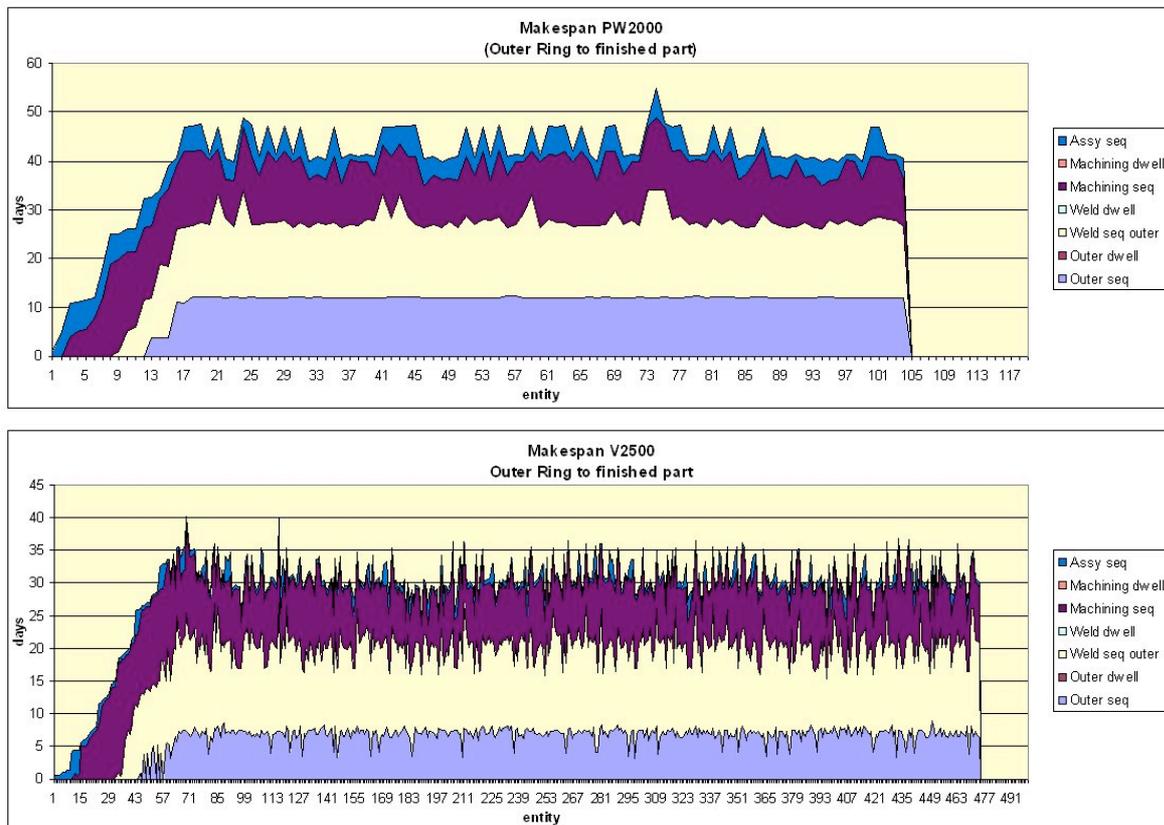

*Figure 27: Lead times observed, steady state 'push' system*
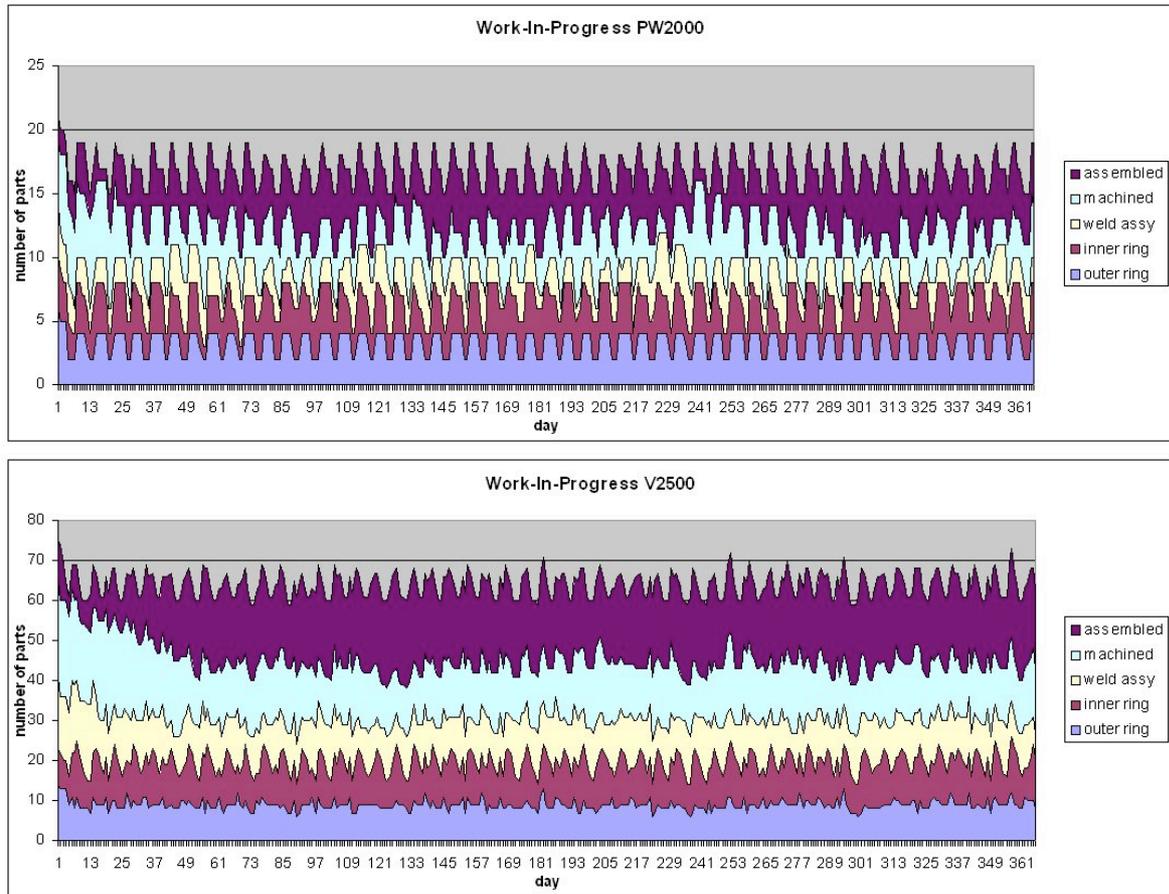
*Figure 28: WIP levels observed, steady state 'push' system*
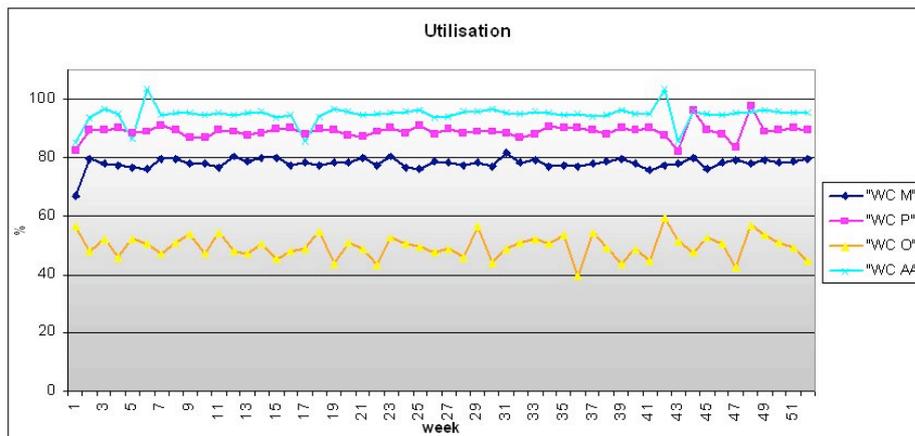
Figure 29: Observed utilisation of key workcentres, steady state 'push' system

The fill rate achieved in this scenario is, predictably, 100%. Measures of resilience and robustness are effectively meaningless, since the manufacturing system has not been stressed. These results provide a good baseline against which to compare the performance of alternative systems of logistic control, and to measure the impact of each deviation scenario. One such measure that will be important is the average capital tied up in WIP (Table 4):

Tied-up Capital

Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	1,632,982	1,210,881
Inner Ring	460,272	457,733
Weld Assy	6,165,457	4,398,544
TEC Assy (machining)	10,575,781	7,616,204
TEC Assy	15,324,577	11,434,221
Total	**34,159,069**	**25,117,582**

Table 4: Average tied up capital for the steady state 'push' scenario

A key issue with the 'push' system is the standard (or 'planned') lead time that is used to determine the launch date for each component (Figure 30). If the lead time allowed is too great, the system will have a needlessly high stock of finished goods. Padding out the standard lead time is effectively 'cheating' the MRP system. Since the fill rate achieved in this scenario is 100% for both products, it may be that the lead times used within the present-day focused factory are overly pessimistic... although the 'steady state' scenario is perhaps too good to be true. Sections 7.2 – 7.5 present disruption scenarios that may justify the standard lead time that has been selected.
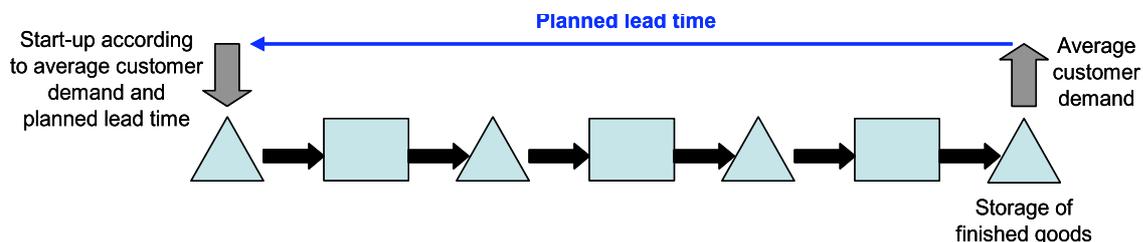
Figure 30: Application of a planned lead time in a 'push' scheduled system

6.1.2. STEADY STATE, 'PULL' SYSTEM

This experiment involves the same, uniform pattern of orders and components supply as described in Section 7.1.2, but with the application of a constraint to the amount of WIP

that would be permitted in each sequence. Section 5.6 described the means by which the 'pull' system was implemented. Figures 31–33 show the results obtained.
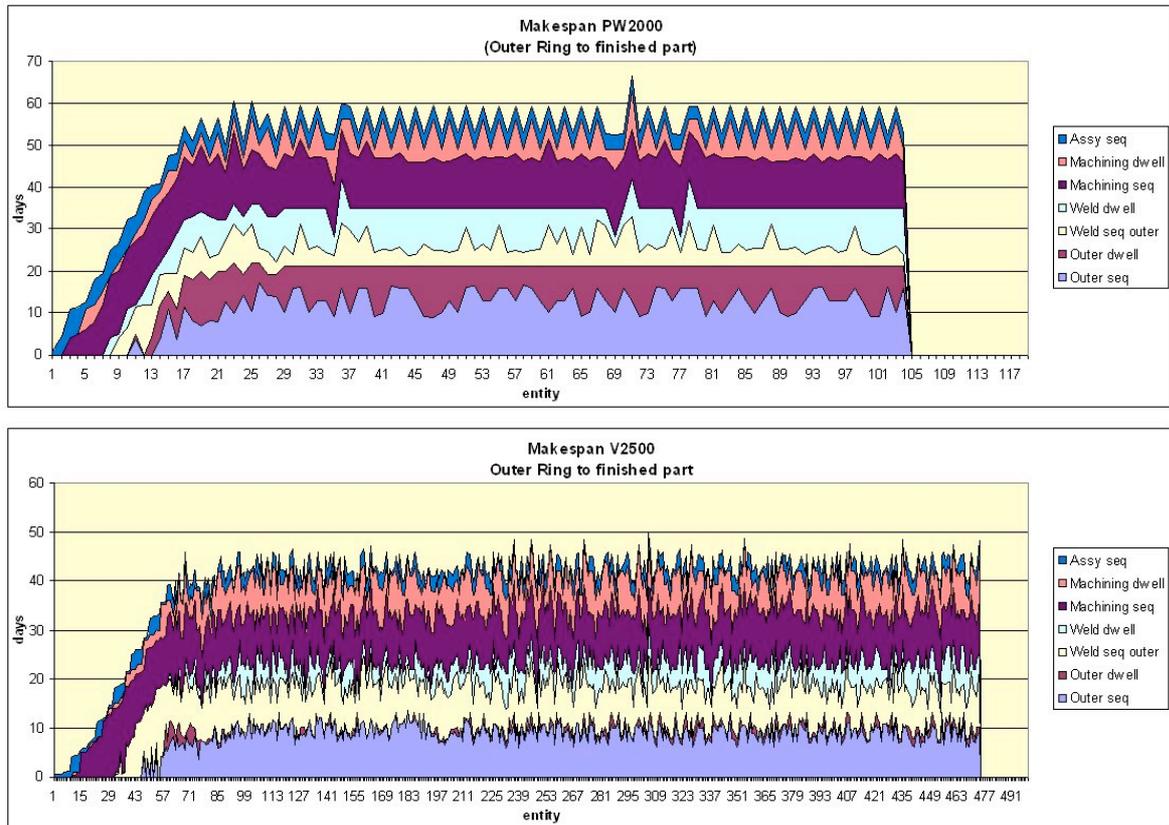
Figure 31: Lead times observed, steady state 'pull' system
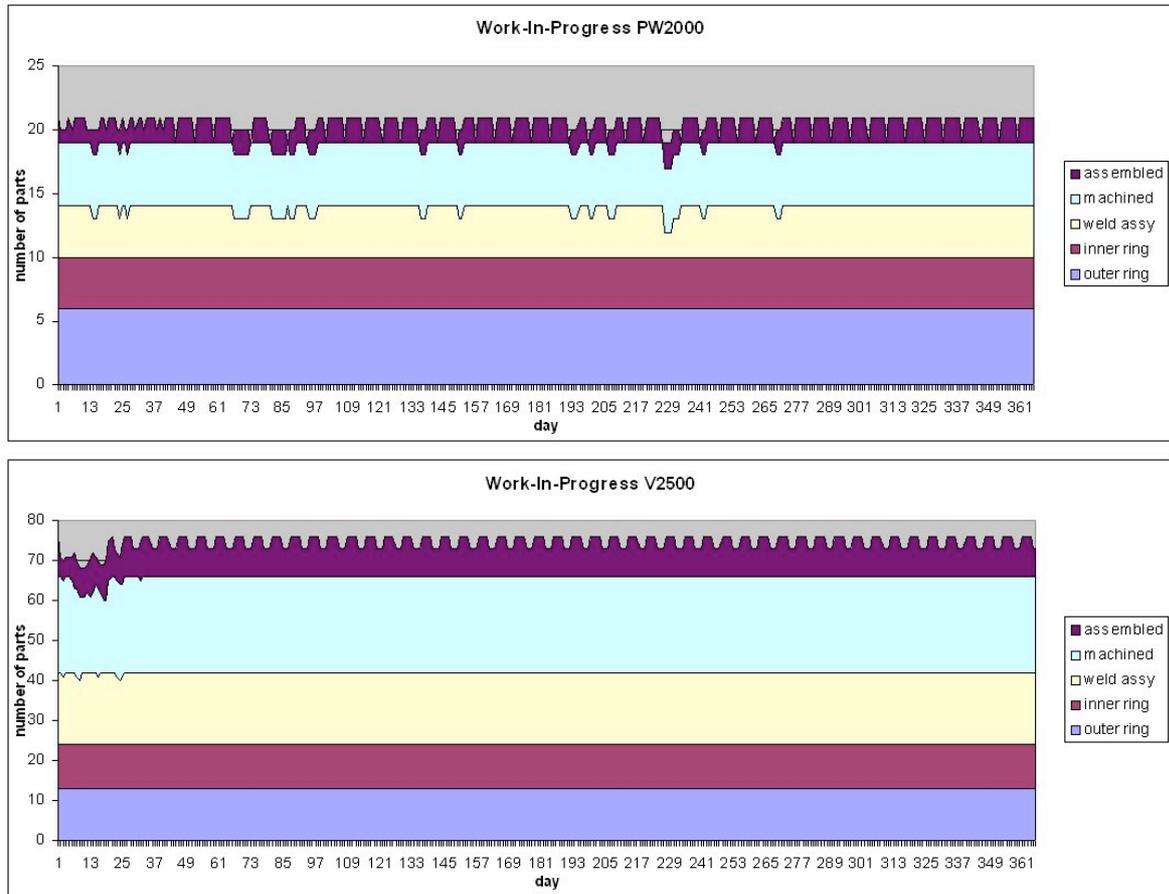
Figure 32: WIP levels observed, steady state 'pull' system
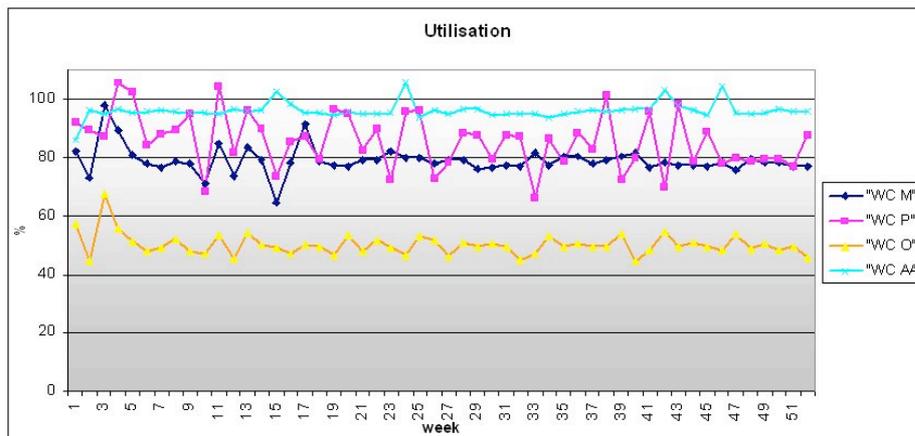
Figure 33: Observed utilisation of key workcentres, steady state 'pull' system

The value of WIP held within the focused factory can be seen in Table 5. This is a disappointing figure, being about 10% higher for the pull system than for the 'push'

system of logistic control. It must be noted, however, that this is when the manufacturing system is running at its best, with no perturbations. The experiments in Sections 6.2 – 6.5 may yet demonstrate some value for the 'pull' system.

Tied-up Capital
Definition:
Standard cost of the average WIP per area.

|	PW2000	V2500
Outer Ring	2,376,790	1,762,426
Inner Ring	473,564	470,951
Weld Assy	10,349,723	7,383,671
TEC Assy (machining)	16,941,050	12,200,185
TEC Assy	7,037,113	5,250,645
Total	37,178,240	27,067,878

Table 5: Average tied up capital for the steady state 'pull' scenario

Since the 'pull' system is being subjected to a steady demand pattern, it is sacrificing one of the key advantages claimed for the approach; that of *responsiveness to problems*. In the basic 'pull' scenario under consideration (Figure 34) there were no fluctuations. It may be for this reason that the 'pull' system fails to outshine the more simplistic 'push' system of logistic control. The deviation scenarios in Sections 6.2 – 6.5 may demonstrate its strengths better.

Ingoing material is unlimited.
(Unless scenario "lack of material".)

Average customer demand
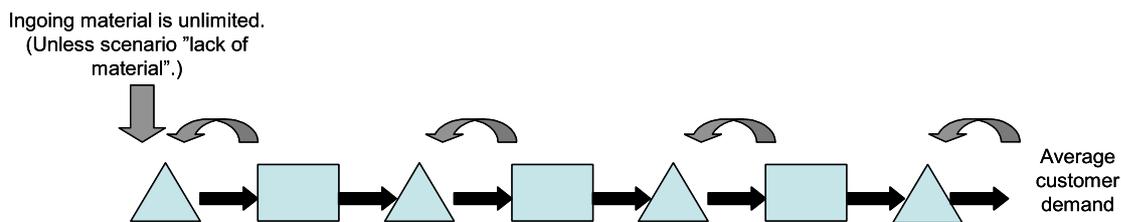
Figure 34: Application of the pull system concept

Another question that was to dog the 'pull' system implementation throughout these experiments concerned the target level of WIP to be employed within each sequence – effectively the number of kanbans to be used. Too many, and the system would not eliminate as much WIP as it could have done; too few and key machines within the manufacturing system might be starved of work. Effectively, there is no single 'pull' system of logistic control. Instead, there are a number of combinations that might be tried, striking a balance between tied up capital and fill rate risk. Beginning with sequence lead times derived from historical data, Little's law [1961] suggested an allowable level of WIP for each sequence. A little fine-tuning by trial and error produced a WIP strategy that could accommodate minor deviation scenarios, and this was used for all 'pull' system experiments described in this document.

The level of WIP within the system is still an experimental result, even though its upper level is constrained by the system of logistic control. Demand from 'downstream' can reduce the level of WIP in a sequence, and regular stock-outs may indicate an inappropriate set of logistic controls.

6.1.3. STEADY STATE, OPT SYSTEM

As Section 4.4 explained, a form of Optimised Production Technology is achieved within the model by setting the target WIP levels for sequences within the model differently. In some sequences a very high target WIP figure is specified, effectively freeing that sequence from any limitation. The logistic control in such areas functions as a 'push' system, while the constrained ones must 'pull' workpieces in. The experiments conducted here use a 'pull' system with the same parameters as Section 6.1.2 in the early stages, and a 'push' system like that of Section 6.1.1 at the end. In this, the system acts like Goldratt's [1987] "drum, buffer, rope" implementation of OPT. The bottleneck sequence (TEC machining) determines the pace at which the facility operates, and it is aimed to keep this facility supplied with workpieces. The operations downstream of the bottleneck are allowed to function on a 'push' system since they are unlikely to accumulate WIP; their cycle time is lower than that of the bottleneck. Figure 35 shows the lead time demonstrated by the focused factory under OPT logistic control. The makespan of both TEC types can be seen to be lower than for pure 'push' or 'pull' systems. This matches the claims made for OPT, in that it prioritises the right activities within the business unit. An hour lost anywhere else in the facility is actually a mirage; an hour lost in the bottleneck is an hour wasted for the whole facility. It should be noted, however, that the performance shown here was only achieved after a lengthy process of 'tuning' the logistic control, by varying the initial WIP and the target figure for each sequence. Care must be taken when configuring manufacturing systems to work in this way.
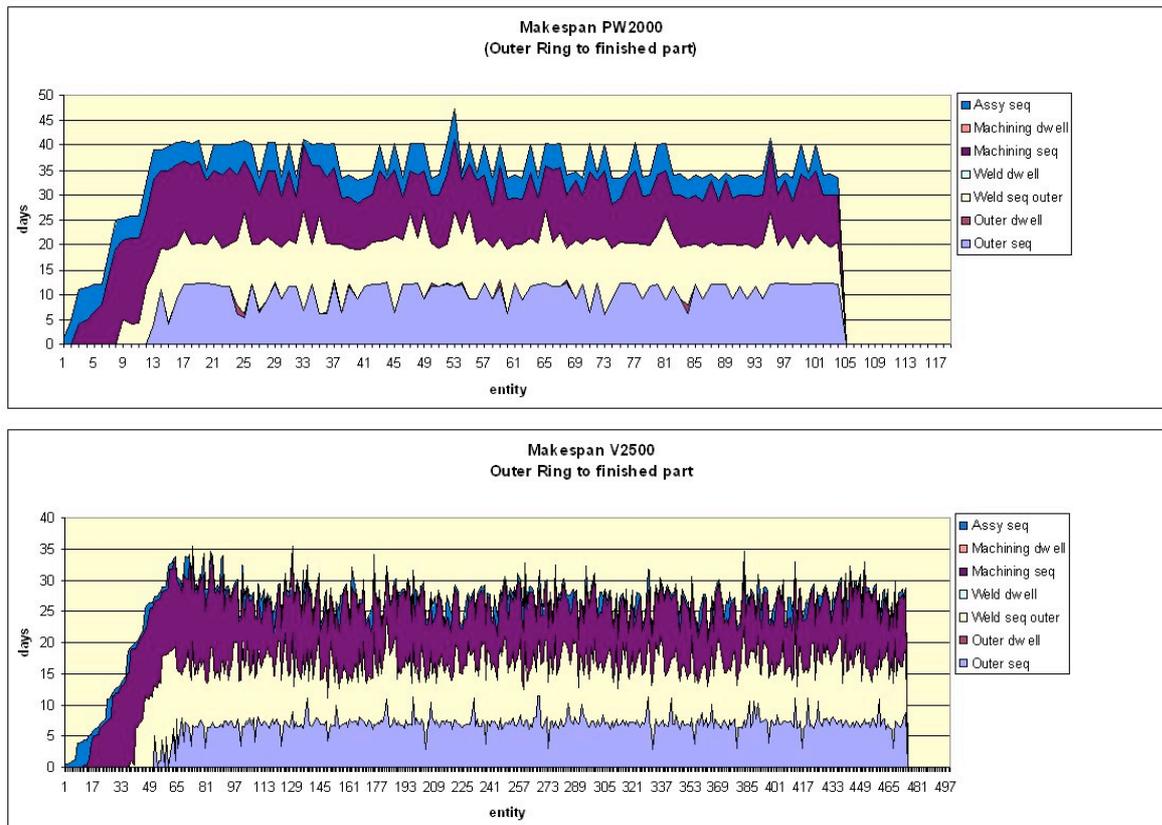
Figure 35: Lead times observed, steady state OPT system

Fill rate for both TEC types was 100% in this scenario. When we examine the level of WIP within the focused factory (figure 36), it appears that a considerable proportion of the WIP was to be found in the final assembly sequence. Since these operations are quite fast, the inventory in question was probably mostly finished goods, completed ahead of their delivery date. Further experiments showed it was possible for the system to be operated on a more lean basis, although this was catastrophic for the deviation scenarios. Thus, the system was permitted to operate with a more generous attitude to WIP than was strictly necessary.
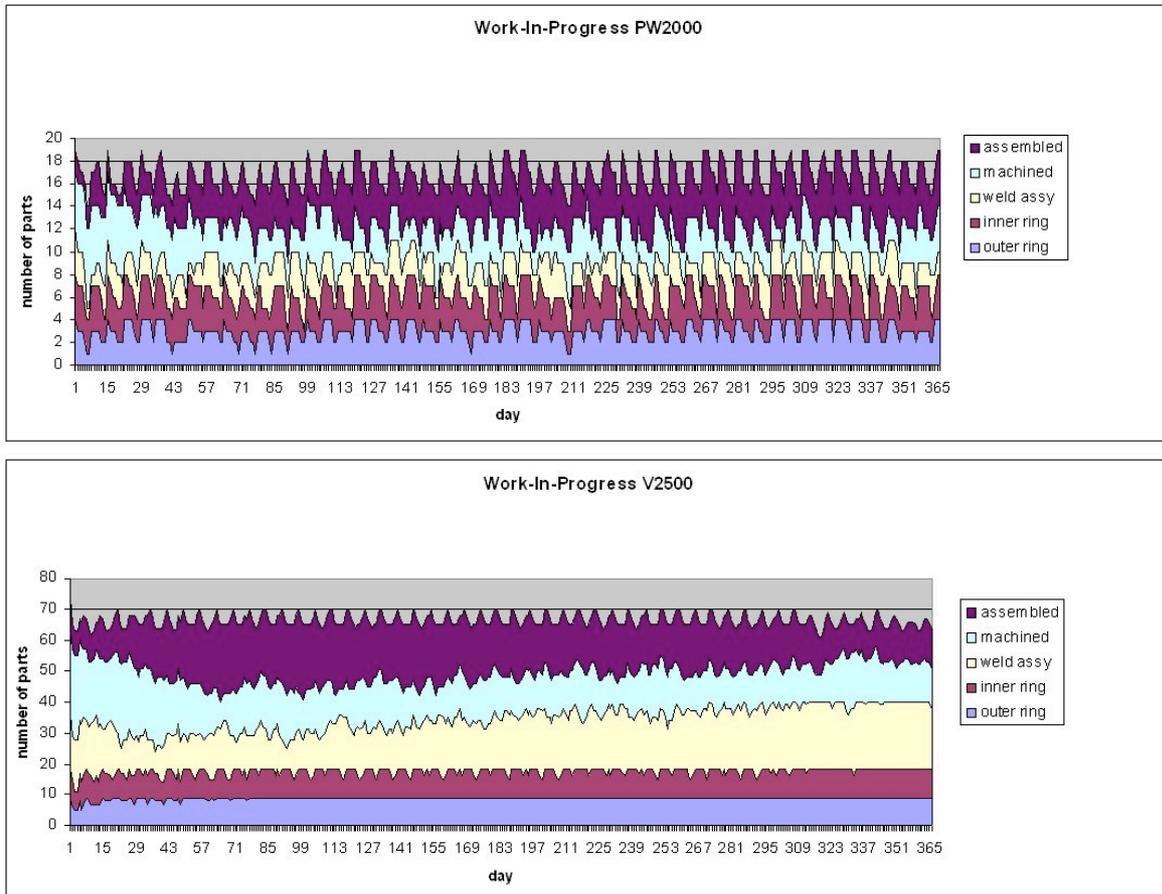
Figure 36: WIP levels observed, steady state OPT system

The average value of WIP within the system, under this scenario, can be seen in Table 6. The overall figure falls between that of the 'push' and 'pull' systems, suggesting it represents something of a compromise strategy.

Tied-up Capital
Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	1,618,495	1,200,139
Inner Ring	362,995	360,992
Weld Assy	9,871,341	7,042,385
TEC Assy (machining)	10,533,074	7,585,448
TEC Assy	13,954,086	10,411,649
Total	36,339,991	26,600,613

Table 6: Average tied up capital for the steady state OPT scenario
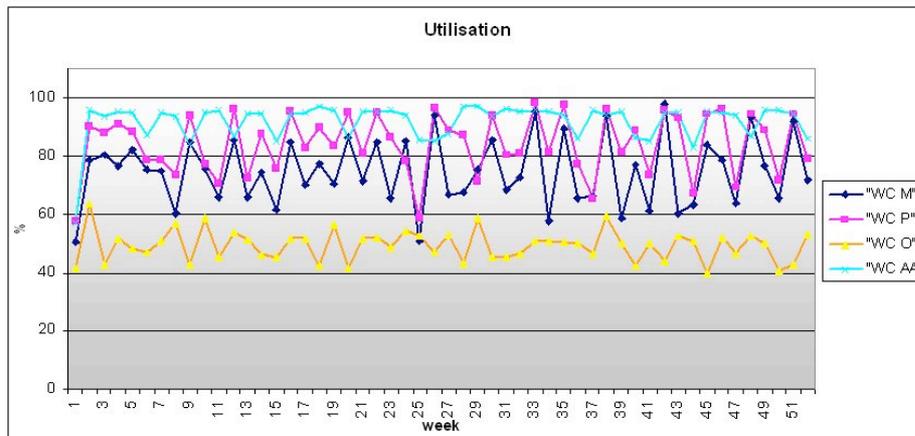
Figure 37: Observed utilisation of key workcentres, steady state OPT system

6.2. DEMAND PATTERN CHANGES

The demand for aero engines is not constant, although the industry is characterised by long lead times, and requires buyers to purchase options that constrain their buying behaviour somewhat. Some changes in the pattern of deliveries must be expected to occur, in response to long-term trends such as the economic cycle, emerging alternatives or sudden shifts such as the downturn in passenger flights that was seen in the aftermath of the attacks of September 11th 2001.

Since the VAC facility is a focused factory and not a continuous manufacturing system, it should be able to cope with changes to the level of demand, provided the logistic controls are appropriate.

In the subsections that follow, demand levels fluctuate by a small amount, as Figure 38 shows. (For the PW2000, the low volume being manufactured means that the variability is actually ±50%.) Both TECs undergo increases and decreases in demand at the same time, so the problem is one of overall capacity, not simply the product mix, but for each TEC the overall figure demanded during the year is the same as under the steady state scenarios presented in Section 6.1. Each of the three systems of logistic control was applied to the new demand pattern, and the results are recorded in the subsections that follow.
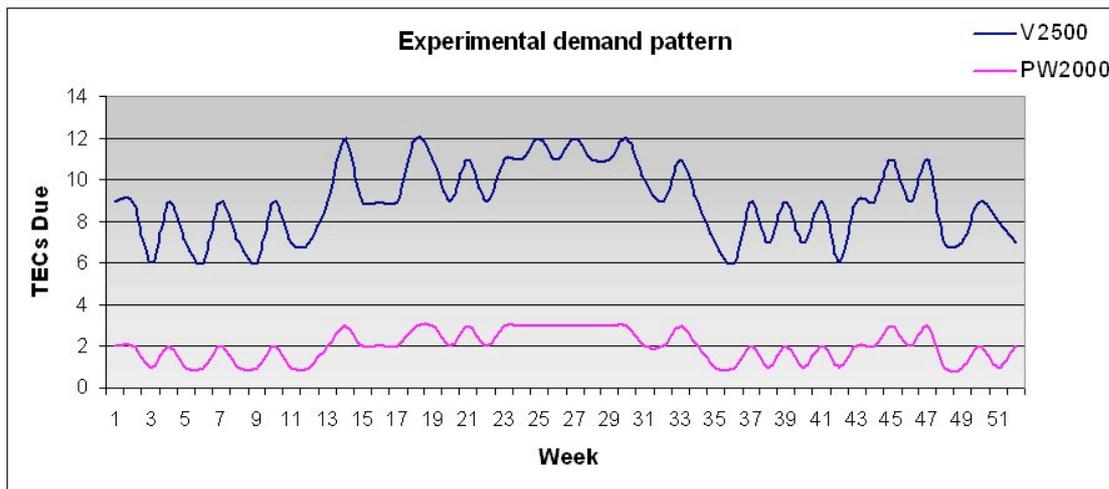
Figure 38: Varying demand pattern applied to the model

6.2.1. DEMAND PATTERN CHANGES AND 'PUSH' SCHEDULING

The push system's performance in a fluctuating demand scenario is very much dependent upon the quality of the forecasts that are employed. In the use case facility, however, we are not working to a forecast but to an agreed build schedule. Each TEC is a component part within an engine, and few will be required as service parts (the TEC can be damaged in rare circumstances, but it is not a consumable in the same way that turbine blades are). The manufacturing system under study is shielded from a good degree of variability as a result, although the variability that remains still has the potential to cause problems for the manufacturing system. Figure 39 shows shows lead times for two products; note how they vary much more than under the steady demand scenario in Figure 27.
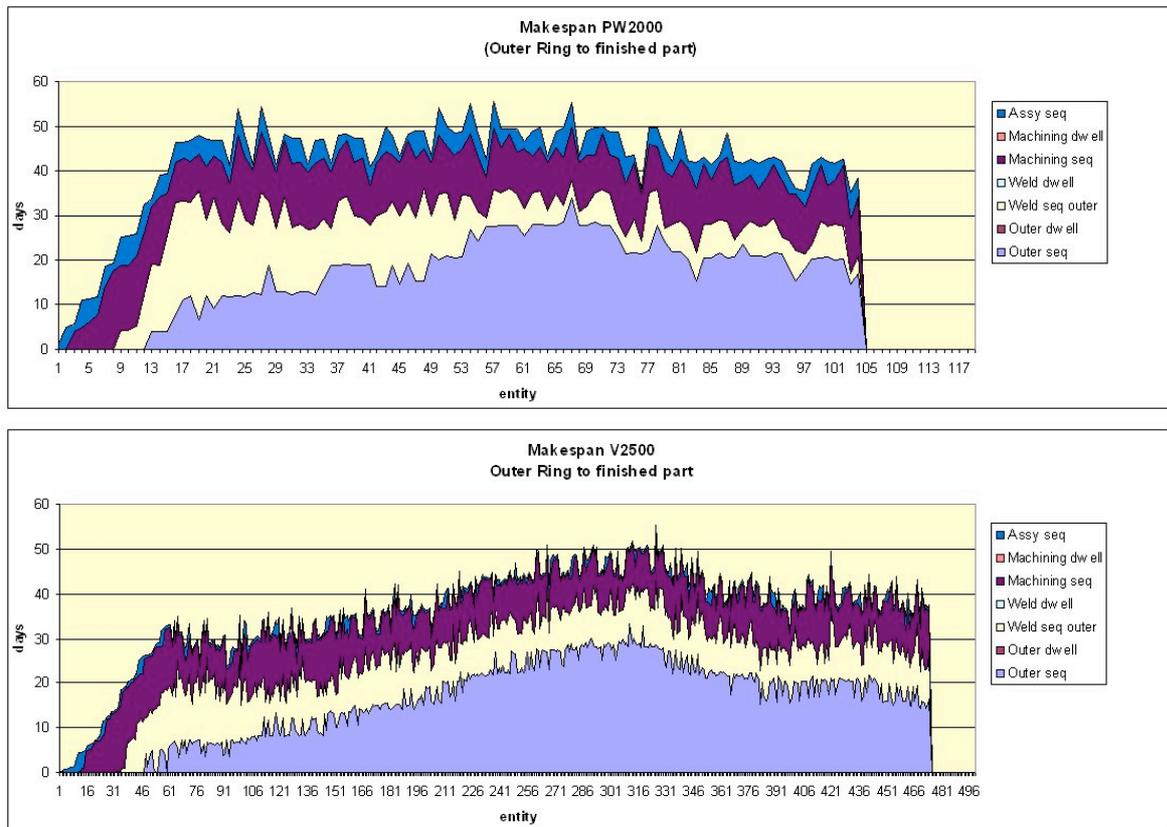
Figure 39: Lead times observed, with demand pattern changes under a 'push' system

The lumpy demand pattern also causes WIP to accumulate within the system (Figure 40). In effect, the focused factory is sometimes operating in arrears, using the weeks with slack demand to work against any backlog.
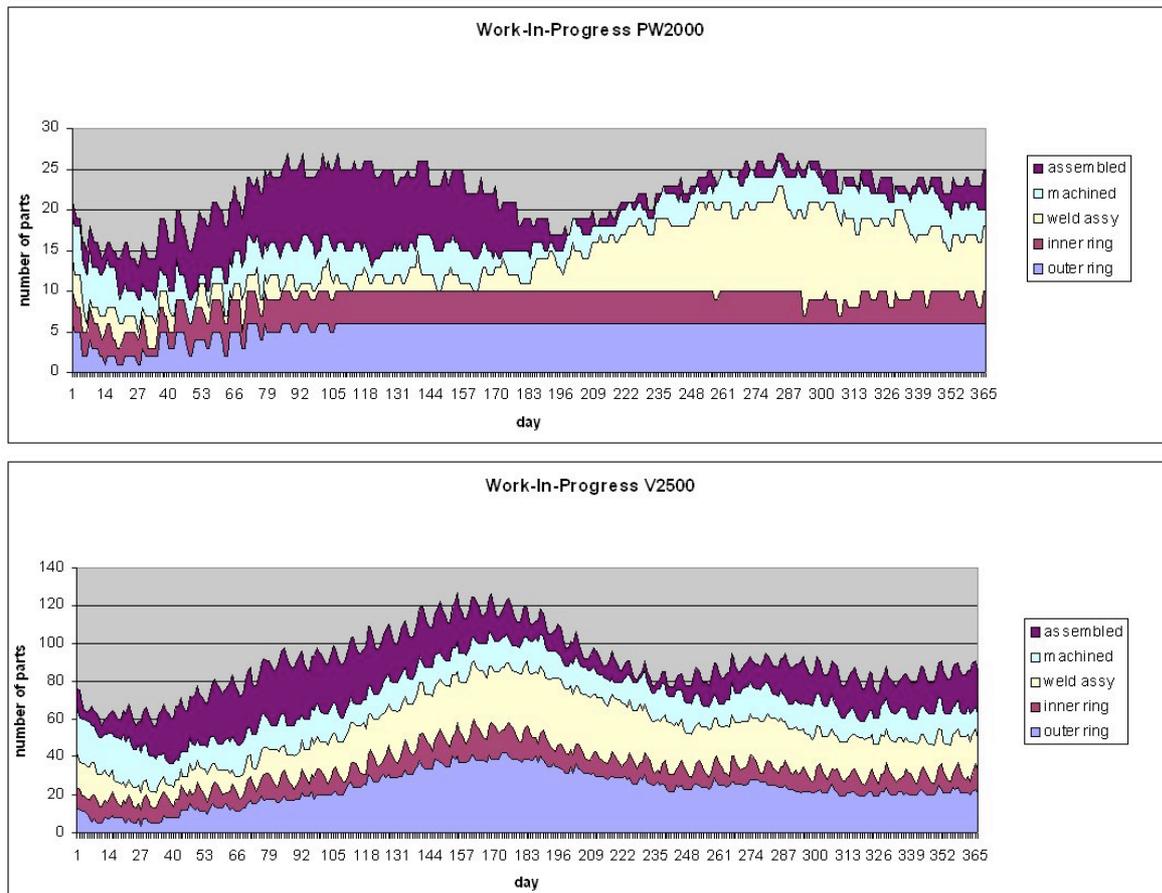
Figure 40: WIP levels observed, steady state 'push' system

The level of customer service achieved for the PW2000 suffered, although as Figure 41 shows, the supply of V2500 TECs continued to meet the demand pattern. Only the end of the high demand period allowed the PW2000 fill rate to return to 100%, so it appears that the present-day goal of manufacturing two of these TECs per week (along with an average of nine V2500s) is very close to the maximum that can be expected. As Figure 42 shows, the stock of finished goods for the V2500 comes *very* close to being exhausted at one point. The 100% fill rate was nearly spoiled, but then demand fell again, and the focused factory was able to rebuild its stock. (The generous standard lead time is a major contributor to the large stock of finished goods seen at most times.)

The machine utilisation pattern is interesting (Figure 43), in that activity at several of the key machines actually appears to dip not at the time of minimum demand, but somewhat out of phase with the demand pattern. This is because the workcentres being monitored work primarily on the earliest machining operations, so their activity is offset by a good part of the standard lead time. A better system of logistic control might introduce a buffer stock in the outer ring, to be depleted at times of peak demand and built up during reduced demand.
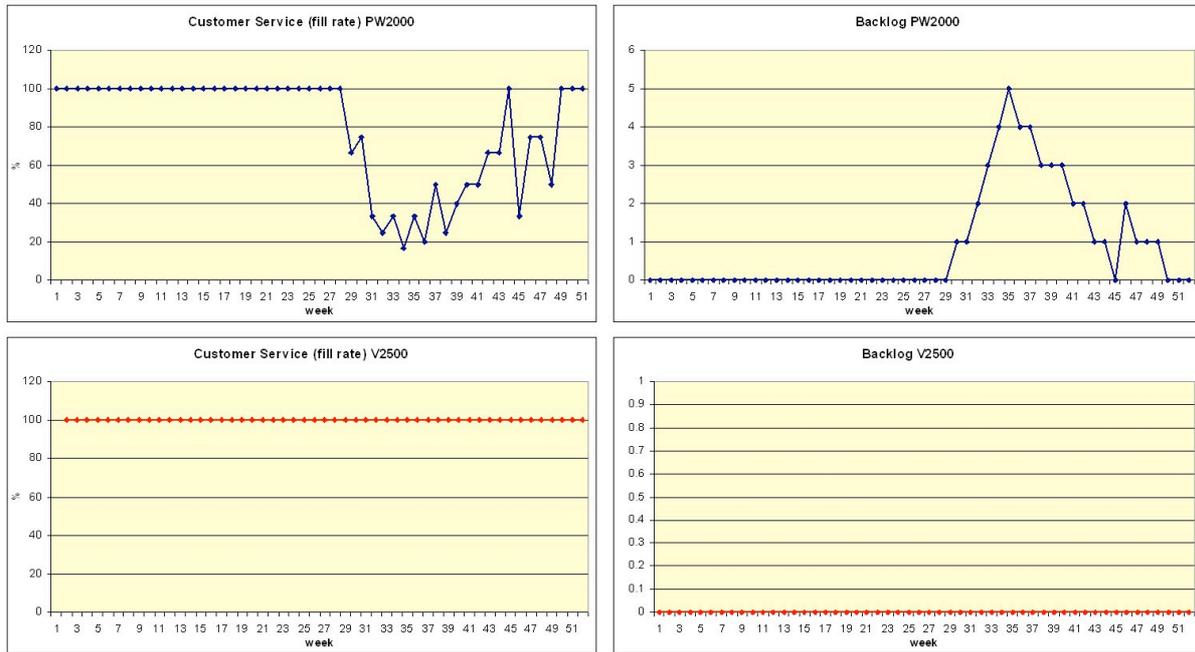
Figure 41: Fill rates achieved (and backlog), with the experimental demand pattern changes under a 'push' system
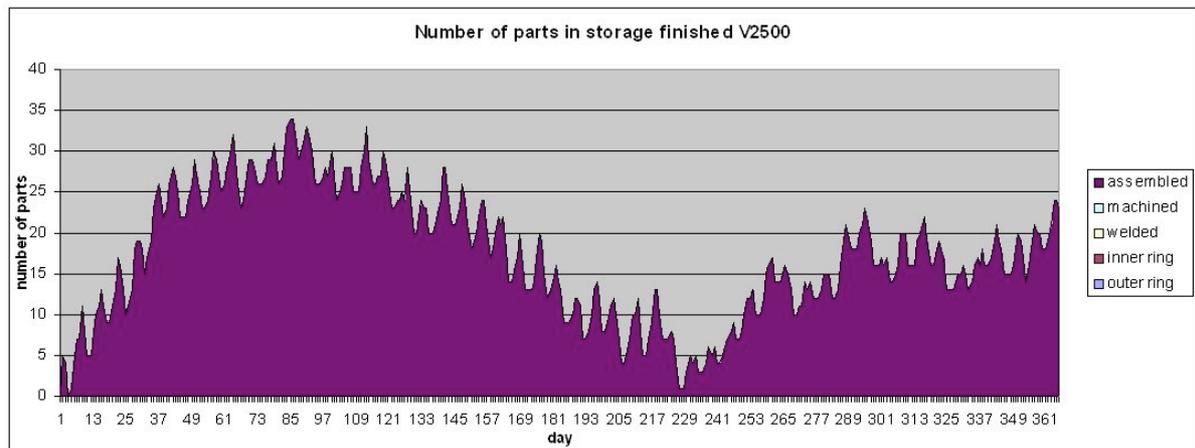

Figure 42: V2500 Stock of finished goods, with the experimental demand pattern changes under a 'push' system. (There is no stock held in the output buffers at the end of each sequence with this form of logistic control.)
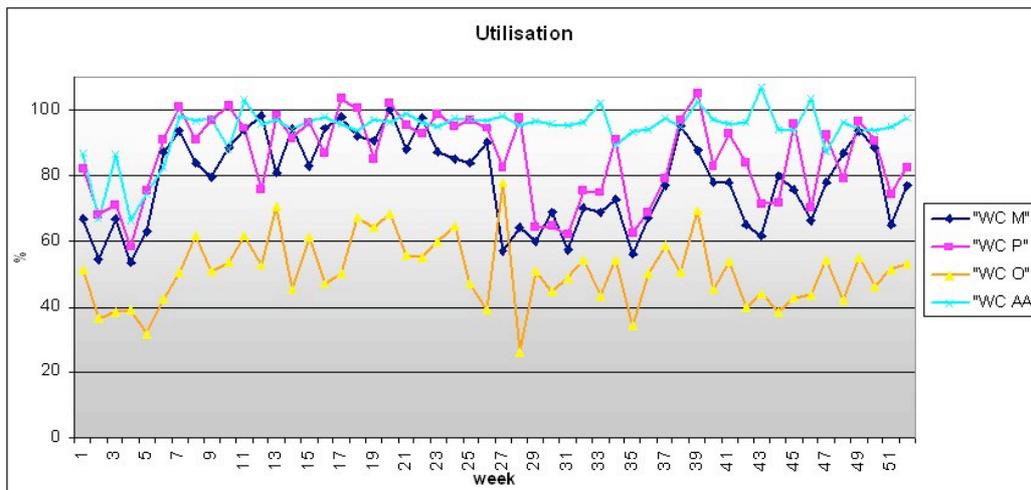
Figure 43: Observed utilisation of key workcentres, with demand pattern changes under a 'push' system

Finally, Table 7 shows the system's tied up capital under this system of logistic control. This can be compared with the alternatives presented within the chapter, to see which approach involves the greatest inventory holding cost:

Tied-up Capital
Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	4,271,528	3,167,404
Inner Ring	451,215	448,726
Weld Assy	12,422,185	8,862,201
TEC Assy (machining)	10,808,728	7,783,961
TEC Assy	15,084,296	11,254,939
Total	43,037,952	31,517,231

Table 7: Average tied up capital for the varying demand / 'push' scenario

6.2.2. DEMAND PATTERN CHANGES AND 'PULL' SCHEDULING

The pull system's performance in a fluctuating demand scenario depends upon the allowable levels of WIP within the system. If we select a generous target WIP level for the final assembly sequence, we could ensure 100% fill rate. Completed or near-completed TECs are the most expensive item of WIP, however, so it is debatable just how far this policy can be pursued. Figures 44–47 show the performance of the system when permitted the same target levels of WIP as were used within the 'steady state' scenario described in Section 6.1.2:
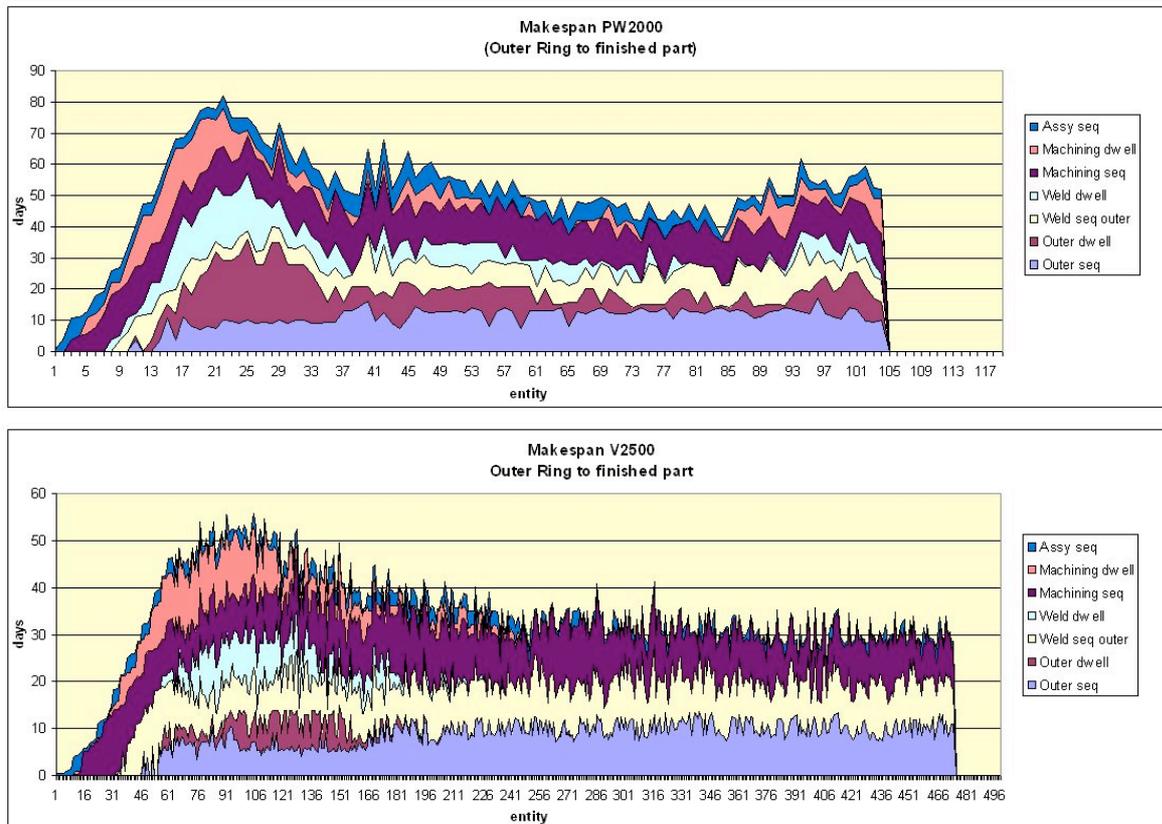
Figure 44: Lead times observed, with demand pattern changes under a 'pull' system

The lead times, when considered in isolation, look quite healthy. It must be remembered, however, that this shows only that the products were pulled through the facility at speed. It does not mean that the supply was sufficient to satisfy the demands being placed upon the system. To investigate this, we need to look at the fill rate achieved (Figure 45).
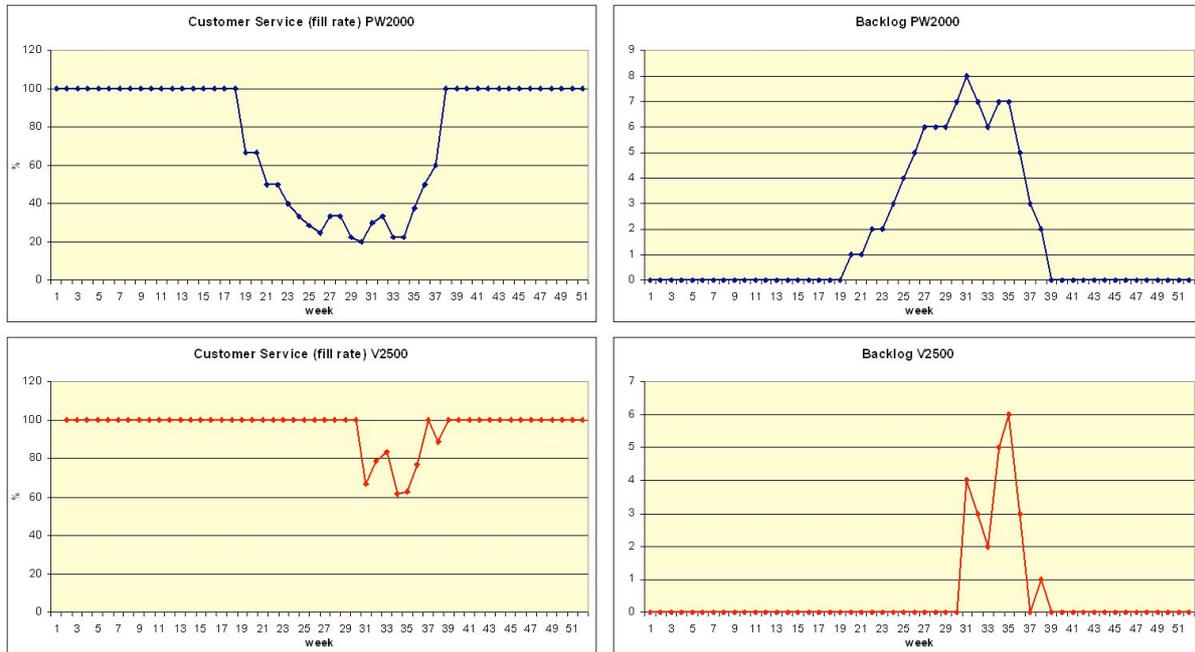
Figure 45: Fill rates achieved (and backlog), with the experimental demand pattern changes under a 'pull' system

The WIP levels within the system during this experiment show how the 'pull' system was mostly able to maintain a steady supply of components in the early stages (the target WIP level in the manufacture of outer and inner rings could actually be reduced) but the welding area was not *always* kept supplied with the materials it needed. This, in turn causes the situation in TEC machining to deteriorate, and at times of sustained high demand, there were not enough products in final assembly (Figure 46).

Workcentre utilisation looks as might be expected; it corresponds roughly to the demand pattern to which the model was subjected, although of course it is constrained at or about 100%. The backlog means that high utilisation continues for a while after the demand level has returned to a more normal level (Figure 47).
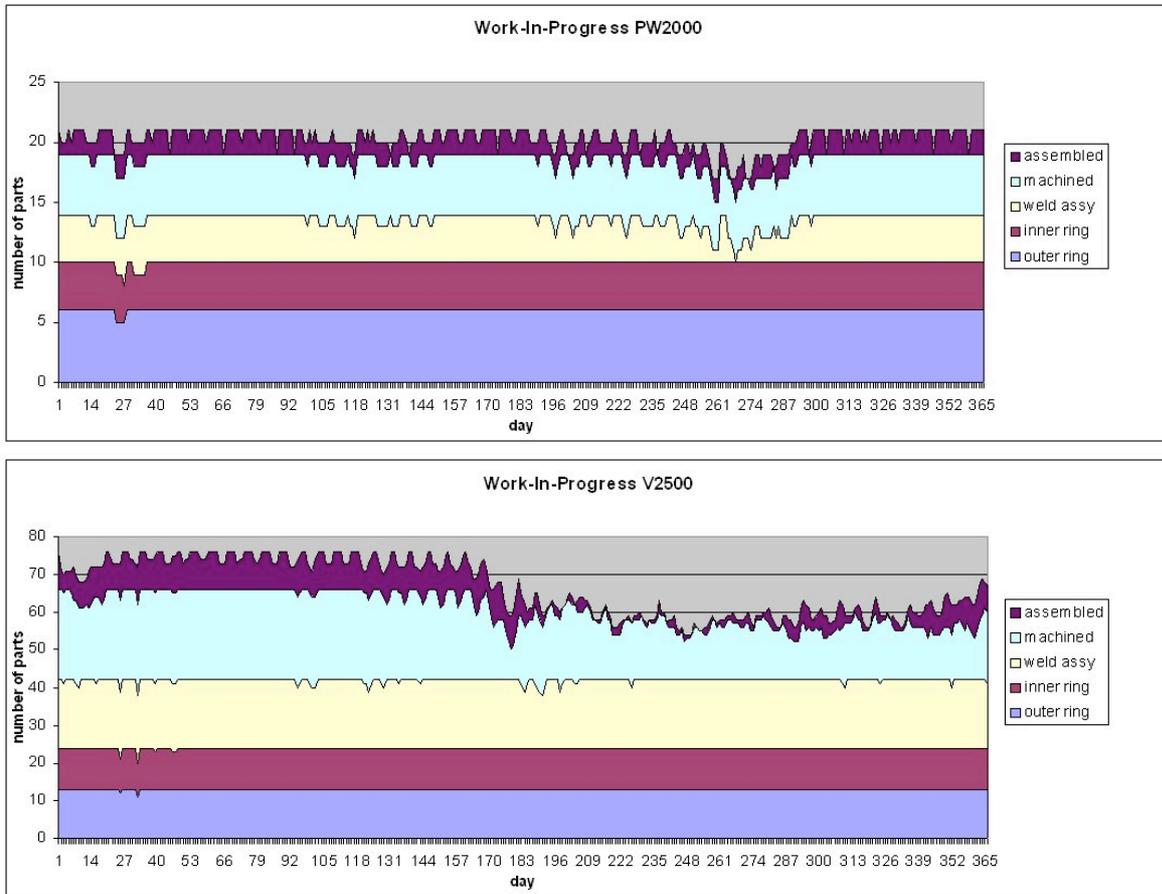
Figure 46: WIP levels observed, with demand pattern changes under a 'pull' system
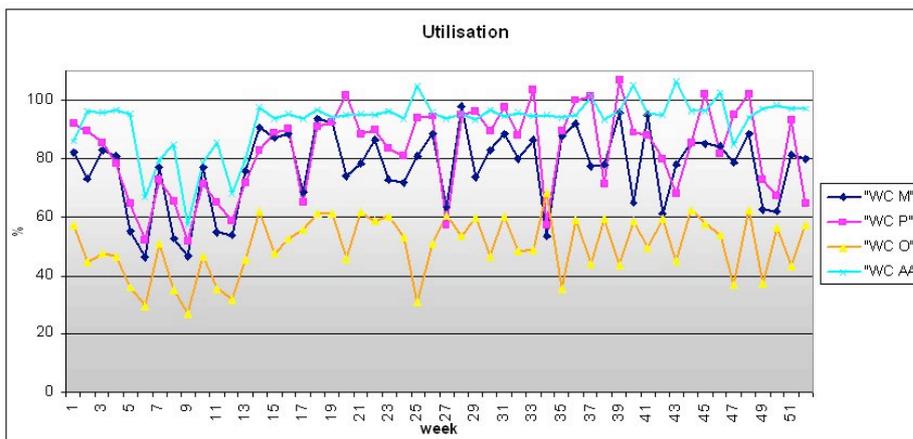
Figure 47: Observed utilisation of key workcentres, with demand pattern changes under a 'pull' system

Finally, Table 8 shows the tied up capital for the scenario. The average value of WIP held within the manufacturing system is substantially lower than that for the 'push' system

described in Section 6.2.1. This represents a change from the 'steady state' scenario, and demonstrates how a system of logistic control that appears sub-optimal when measured under ideal conditions can prove useful in a more complex context.

	PW2000	V2500
Outer Ring	2,375,291	1,761,315
Inner Ring	472,740	470,132
Weld Assy	10,278,910	7,333,152
TEC Assy (machining)	13,365,318	9,625,103
TEC Assy	4,652,104	3,471,103
Total	31,144,363	22,660,805

Table 8: Average tied up capital for the varying demand / 'pull' scenario

6.2.3. DEMAND PATTERN CHANGES AND OPT SCHEDULING

Appropriately configured, OPT should combine the better points of the 'push' and 'pull' systems, getting the most out of the bottleneck component of the system. We have something of a complication with the focused factory, in that the bottleneck machine may differ, depending upon the relative proportion of PW2000 and V2500 TECs that flow through the system. Even with the demand pattern for both engine types following a similar trend, the kind of work to be done at a workcentre in any given week may differ significantly.

As Figure 48 shows, the OPT system continued to exhibit short makespans for most products, as it had under a steady demand scenario. The rate at which products were passing through the system was not *quite* good enough, however. The V2500 fill rate dips halfway through the year – failing to keep up with peak demand – and although some moves towards eliminating the backlog are made, the system continues to operate in arrears for the remainder of the simulation. The backlog is roughly one week's production (Figure 49).
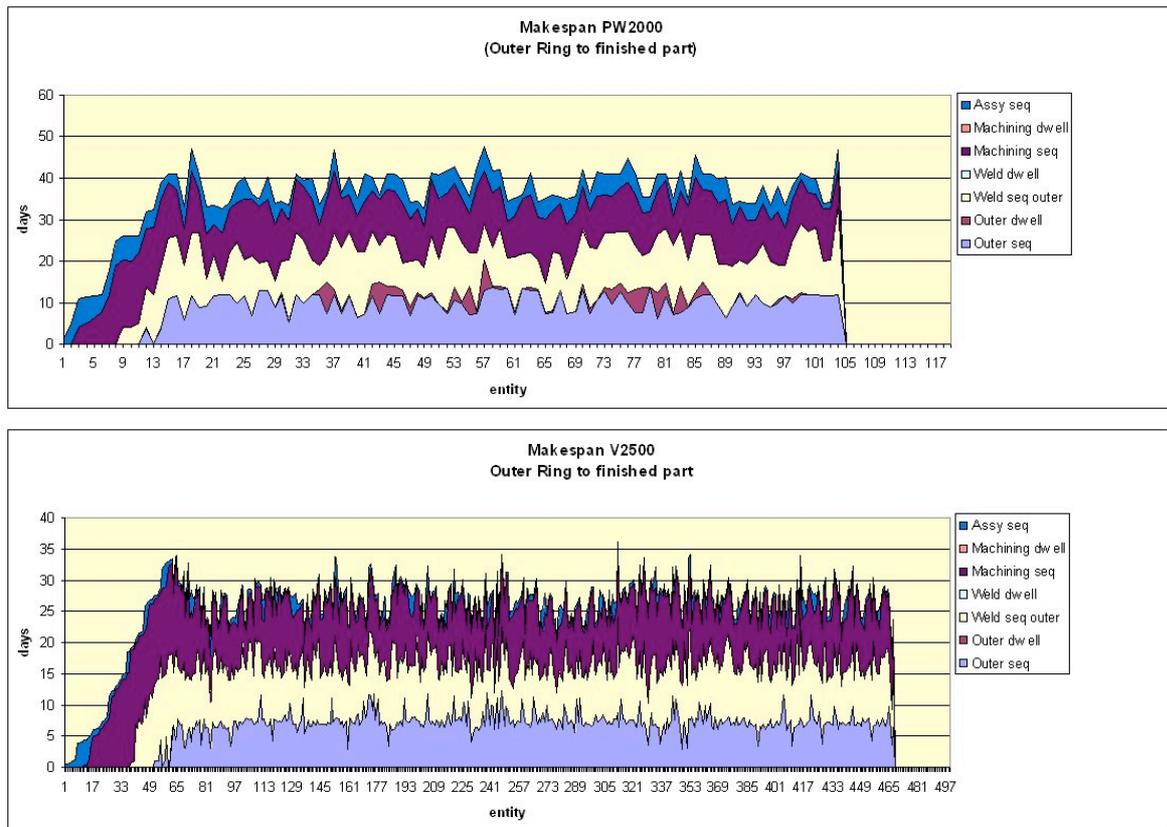
*Figure 48: Lead times observed, with demand pattern changes under OPT*
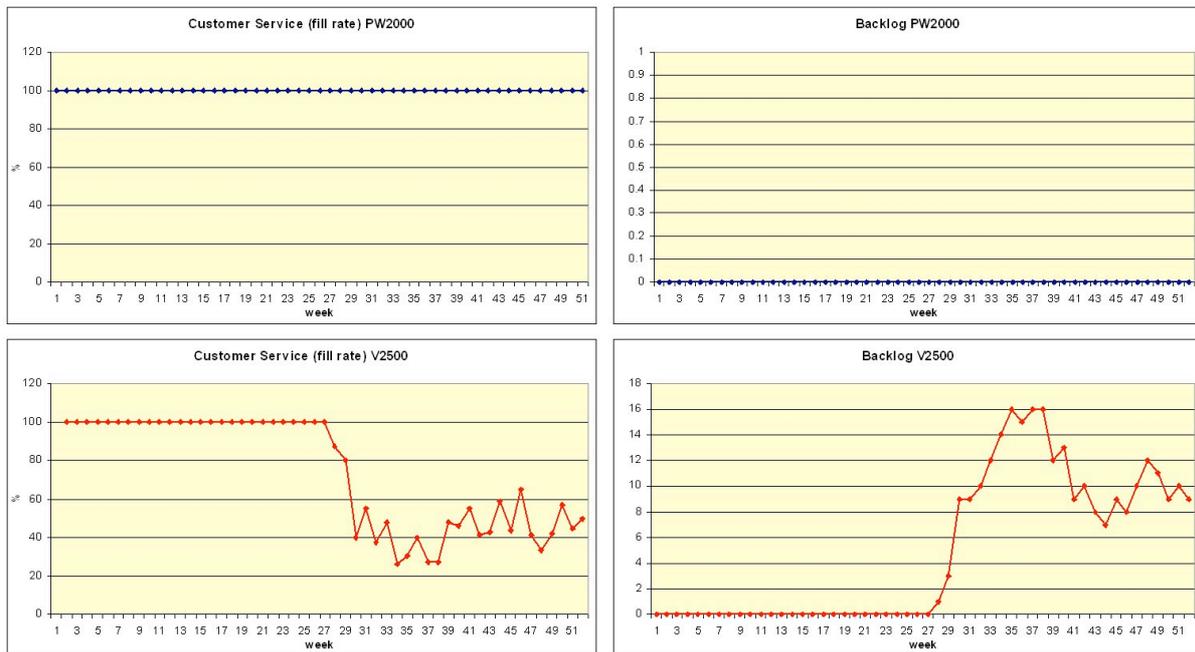
Figure 49: Fill rates and backlogs, with demand pattern changes under OPT

The WIP help by the focused factory under this scenario is quite well controlled in the early stages, as Figure 50 shows. Outer ring, inner ring and welding stages are all constrained by an upper limit that appears to work well, although the material in TEC machining fluctuates, implying that this sequence is not being utilised to the maximum. This is confirmed by the fluctuations in machine tool utilisation (Figure 51). The 'push' system of logistic control employed to flush the operations downstream of the bottleneck is effective, but causes a considerable stock of finished goods to be accumulated during the period of low demand. At peak demand this is eroded, until a point is reached where the fill rate suffers. The minimal WIP found in the assembly sequence for V2500 in the second half of the simulation run illustrates the problem, and makes a strong case for a buffer stock.

The better performance of the PW2000 in this scenario goes against what might have been expected, since the lower volume of this TEC's manufacture tends to reduce cushioning effects. (In some scenarios, a part might suffer a delay, but there are plenty of others that can overtake it and satisfy the customer order.) The utility of even a single extra product in a buffer is much greater for the lower-volume product, it seems.

This serves to demonstrate the value of WIP, which must be borne in mind when pursuing any optimisation strategy based upon inventory holding cost. (Table 9 shows the average WIP cost for this experiment.)
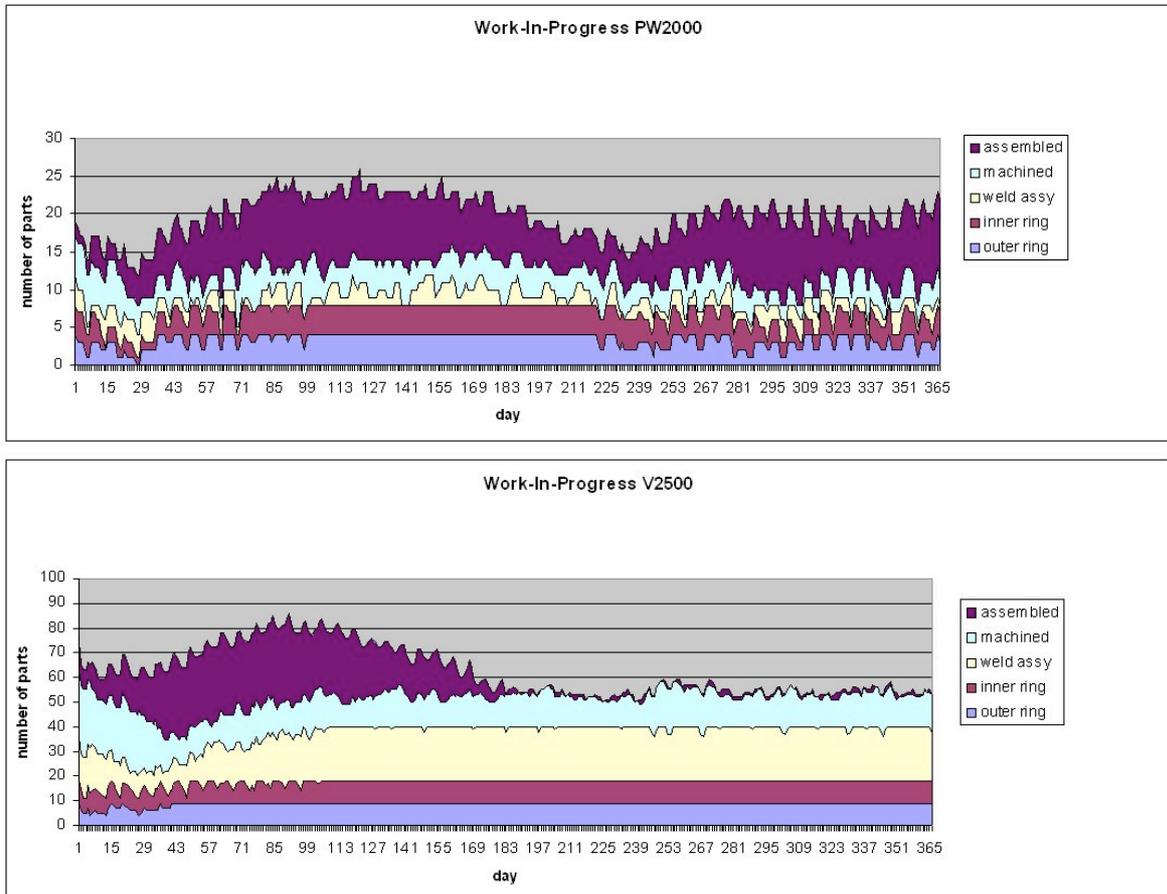
Figure 50: WIP levels observed, with demand pattern changes under OPT
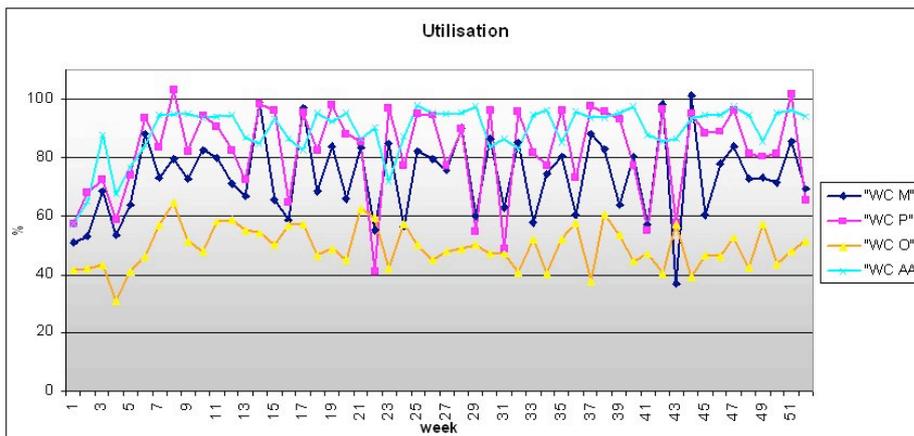

Figure 51: Observed utilisation of key workcentres, with demand pattern changes under OPT

Tied-up Capital

Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	1,597,015	1,184,211
Inner Ring	369,112	367,075
Weld Assy	11,410,345	8,140,337
TEC Assy (machining)	10,205,008	7,349,189
TEC Assy	8,756,901	6,533,841
Total	**32,338,380**	**23,574,653**

Table 9: Average tied up capital for the varying demand / OPT scenario

6.3. QUALITY PROBLEMS

The most minor quality problems are already built into the model, in that the historical data used to evolve the statistical distributions used for processing times includes some basic remedial actions. Within the operations flow for each TEC, there is also an inspection stage where it is quite common that the product should require extra testing, and in some cases additional welding. Consequently, a rework loop was included in the operations flow for all experiments because the problem is a common occurrence.

This deviation scenario investigates the effect of unanticipated, more severe problems of the kind that arise from time to time. It is relatively simple to represent a quality problem within the model, although at this time it *does* require some knowledge of Arena. Theoretically, an input file could be used to define the yield at each operation, in the same way that the model is populated with WIP when it is initialised, but for now a quality problem is introduced by adding a 'decide' module to the operations flow that sends some percentage of the products onwards, while others are delayed and/or routed through additional processing stages.

The high value of TECs means that quality problems do not automatically result in scrap. An investigation is started and if at all possible, a rework strategy is developed, working in collaboration with the customer. This means that there will be a substantial delay – in a major rework scenario, a damaged workpiece could be awaiting agreement on its rework strategy for as much as six months – and some demands may be made upon the resources within the focused factory before the part can rejoin the operations flow.

All such issues are relatively simple to model within Arena. The only problem is that each defect would tend to be unique, requiring a different rework strategy. As such, any simulated quality problems must be seen as representative, rather than encompassing the full range of potential defects.

The steady demand pattern of two PW2000 and nine V2500 TECs was used, not the variable demand pattern from Section 6.2. The representative quality problem that was selected is one that occurs after an inspection operation in the TEC machining sequence. For the purposes of the current experiment it was decided that 8% of PW2000 TECs being inspected should fail, along with 12% of the V2500 type, with

further investigation and repair being necessary and the part being delayed by 2–6 weeks. This means that the part leaves the operational flow; an issue that affects 'push' and 'pull' logistic controls differently, as the subsections that follow explain.

6.3.1. QUALITY PROBLEMS AND 'PUSH' SCHEDULING

The 'push' system is effectively unable to respond to variability that occurs after products have been launched. It simply uses a safety margin in the standard lead time to provide time for problems to be absorbed. When a quality problem occurs, a new workpiece should be allowed to enter the affected system, but because the basic 'push' system has no buffers within the operational flow, any replacement must start at the raw materials stage. Even then, it may be that ordering lead times for raw material (not included within this simulation) play a part, leading to an even longer delay before a replacement can be produced. For the selected quality problem with its 2–6 weeks delay, any such timescale would mean the replacement part fails to overtake the part being reworked.

Figure 52 shows the lead time performance of the manufacturing system with 'push' scheduling and a quality problem scenario in each TEC's manufacturing operations. The lead time plots each show 'spikes' in the TEC machining element of the lead time; these are where the quality problems that took place. The larger manufacturing volume for the V2500 makes the problem look much more severe.
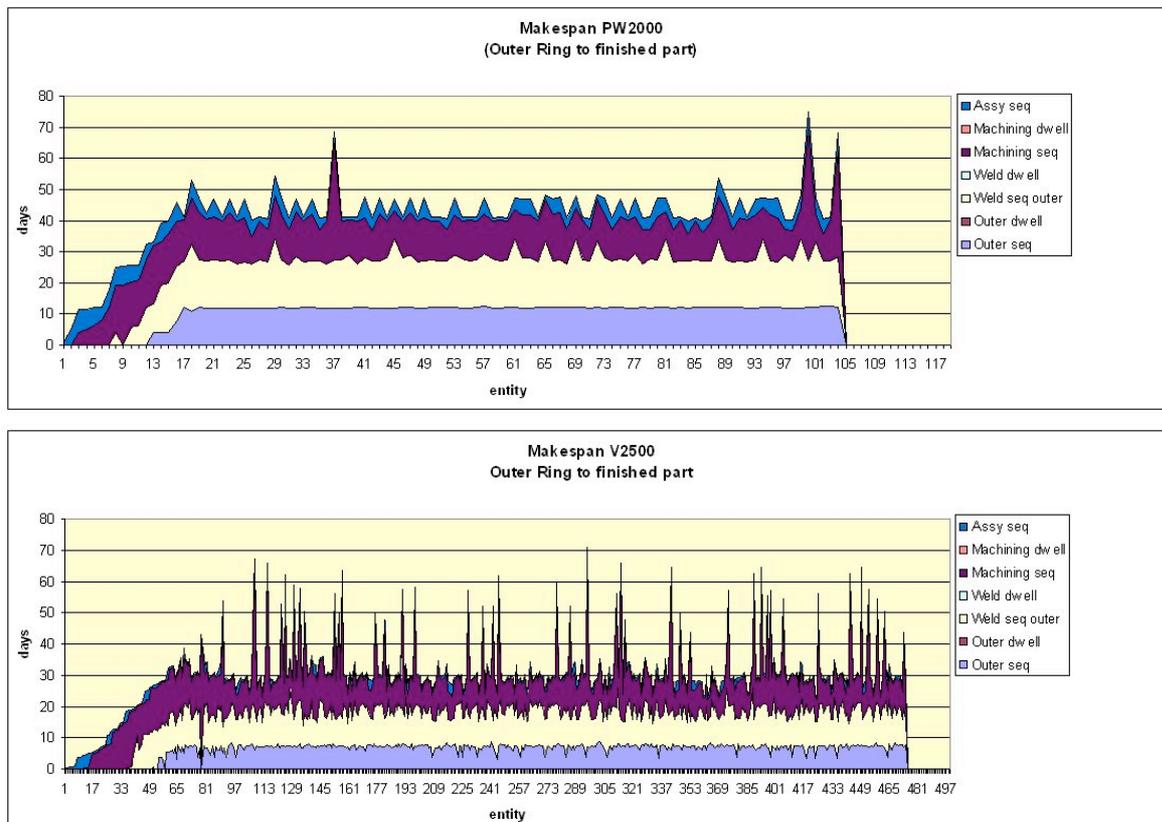

Figure 52: Lead times observed, with quality problems in a 'push' system

The level of customer service achieved is 100% for both TEC types, so the fill rate/ backlog diagrams are not shown. Utilisation of key workcentres system was highly uniform, as Figure 53 shows. It appears that the factory maintains a steady rhythm despite the occurrence of quality problems. It may be significant that the remedial action carried out in response to a quality problem is modelled as being carried out by a common resource, as opposed to requiring the involvement of the focused factory's own resources. (When time permits, an additional set of experiments could be conducted, in which a bottleneck or near-bottleneck resource must be used to return a defective part to the operations flow.)
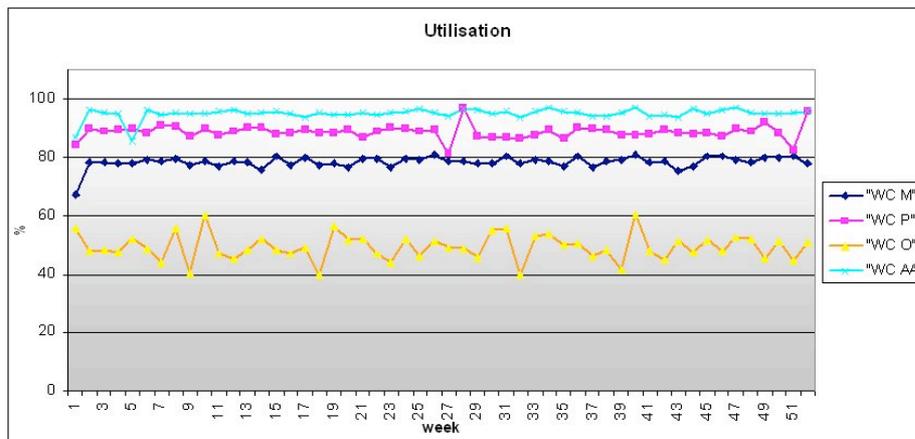
Figure 53: Observed utilisation of key workcentres, with quality problems in a 'push' system

The level of WIP within the manufacturing system is generally steady, if a little high. Figure 54 shows where WIP was to be found during the model run, and Table 10 shows the value of that material.

Figure 54: WIP levels observed, with quality problems in a 'push' system

Table 10: Average tied up capital with quality problems under a 'push' scenario

6.3.2. QUALITY PROBLEMS AND 'PULL' SCHEDULING

There is a danger, with the 'pull' system, that quality problems will disrupt the carefully-configured lean production system. Indeed, one of the advantages often ascribed to such systems is that problems cannot "hide" within the factory. The reduced inventory level highlights sources of variability, which can then be designed out.

The VAC facility is not entirely suited to a lean production approach, however. Sources of variability such as the time a part spends at a common resource cannot be controlled, so this must be absorbed by some level of surplus capacity, or surplus WIP. The target WIP figures specified for all the 'pull' system experiments are generous enough to cushion the system against these minor variations, although it remained necessary to see how a quality problem within the focused factory might affect its performance.

Figure 55 shows the lead times achieved for a 'pull' system in the quality problems scenario. Despite the ability to react when problems occur, the best lead times demonstrated by the 'pull' system are worse than those of the 'push' system. This makes sense, because this system of logistic control requires parts to dwell in output buffers between sequences until they are needed. (These dwell times can be seen in the lead time graphs.)
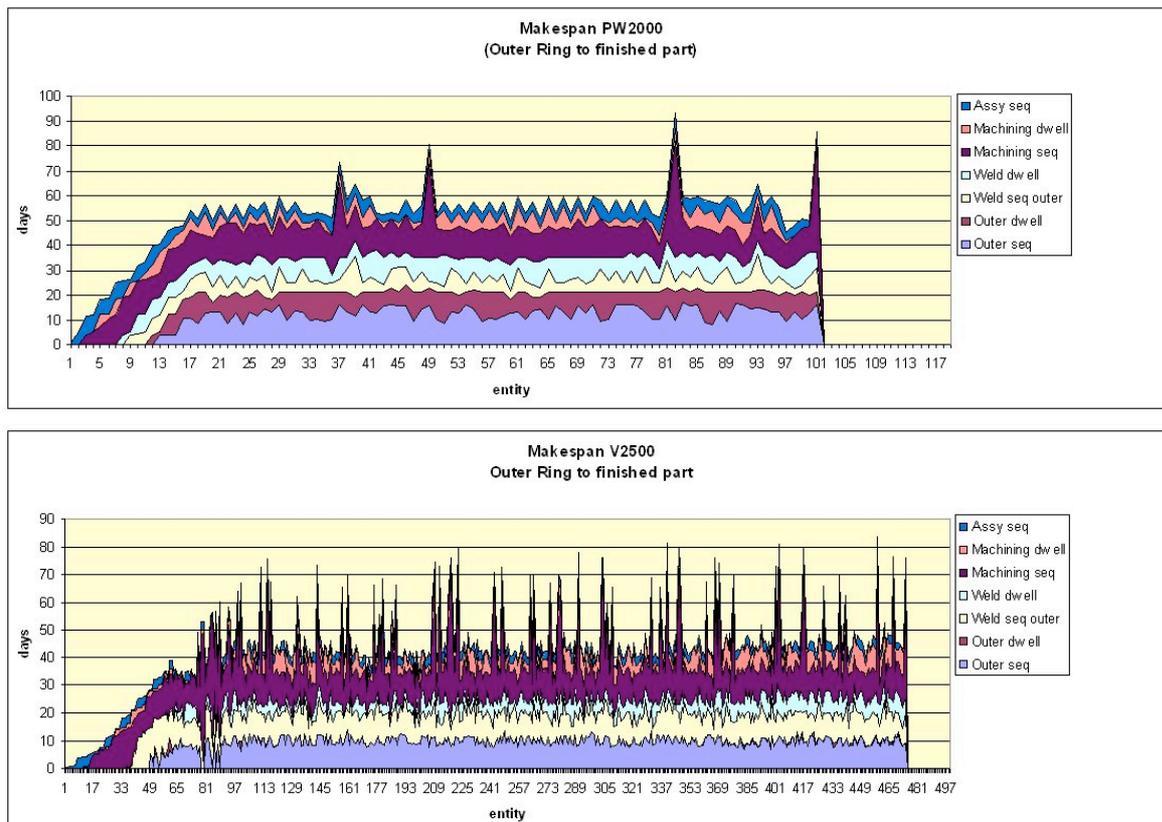
Figure 55: Lead times observed, with quality problems in a 'pull' system

The fill rate achieved with 'pull' logistics is near-perfect (Figure 56). A slight problem can be seen at the end of the year, but it should be remembered that for the scarcer PW2000 TEC, a fill rate of 50% indicates a shortfall of just one component. (It might be wise to

allow this experiment to run for a longer duration, to investigate the nature of the year-end problem more closely, but for the purposes of this document, all experiments took place under the same conditions.)
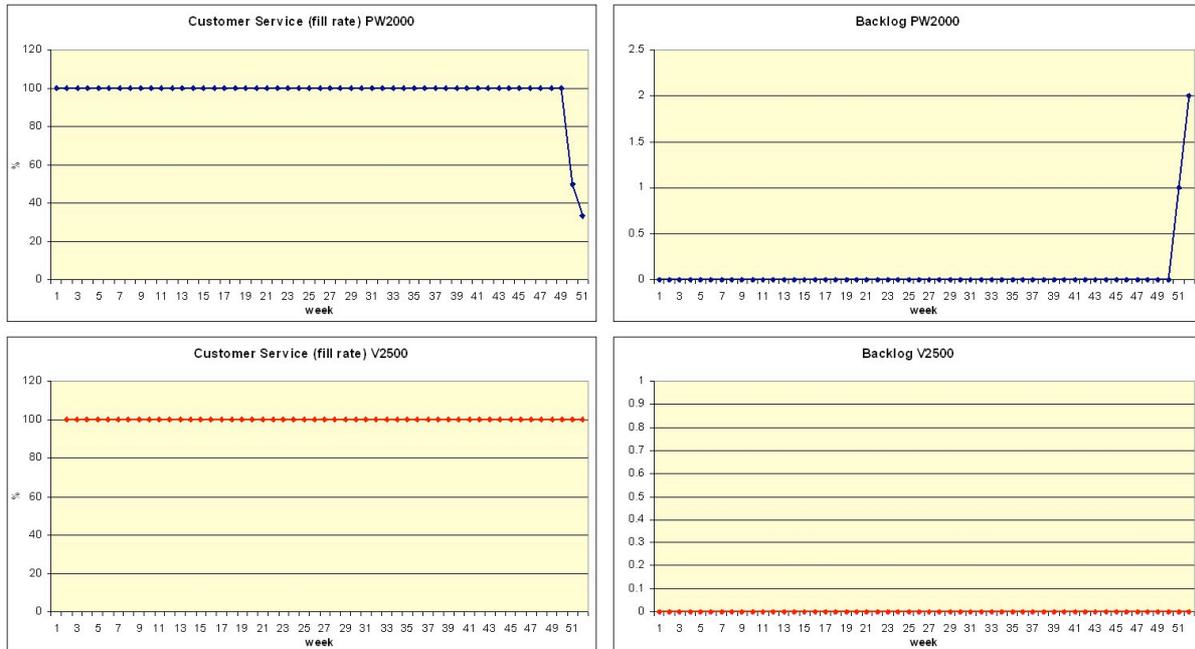
Figure 56: Fill rates and backlogs, with quality problems in a 'pull' system

The WIP levels in Figure 57 demonstrate once again the strict control that the 'pull' system has over inventory. It barely fluctuates from the target figures specified in the 'VAC_inputs.xls' spreadsheet. This success is achieved at the expense of machine utilisation (Figure 58), which varies considerably as operations are timed to maintain the steady WIP figure.
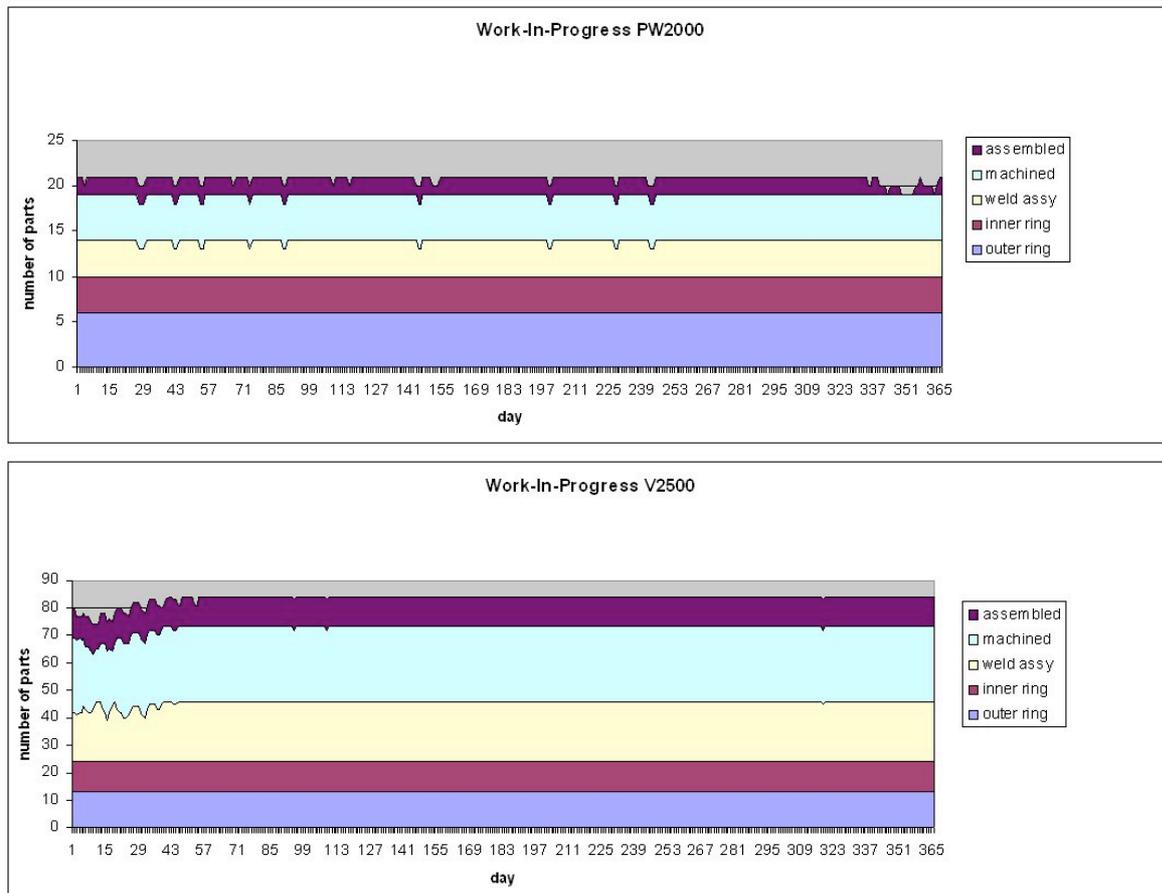
Figure 57: WIP levels observed, with quality problems under a 'pull' system

There is an opportunity to make the 'pull' system still more reactive to issues on the shop floor. If a quality problem occurs, it may be able to respond by drawing a replacement workpiece from the preceding output buffer. Attempting to do this involves exceeding the target WIP figure for the sequence temporarily. In effect, such a system of logistic control recognises that a part undergoing a lengthy quality problem need not hold a kanban while it is undergoing investigation and repair. This strategy would allow a replacement part to enter the sequence (if one is available). When the damaged workpiece returns from its inspection and repair, the sequence will be overstocked with WIP (the number of parts in the sequence is greater than the target WIP value). This would be permitted, but the *next* part to leave the sequence would not trigger the release of a workpiece from the preceding output buffer, and thus normal practices would be resumed. This form of logistic control could be the subject of a future experiment.
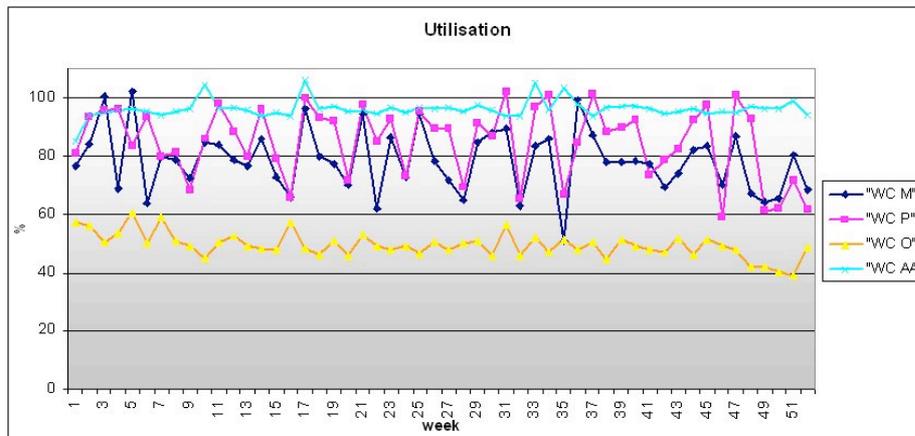
Figure 58: Observed utilisation of key workcentres, with demand pattern changes under a 'pull' system

Table 11 shows the average value of inventory for this scenario:

Tied-up Capital

Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	2,376,790	1,762,426
Inner Ring	473,564	470,951
Weld Assy	12,475,688	8,900,371
TEC Assy (machining)	19,041,453	13,712,801
TEC Assy	8,881,491	6,626,802
Total	43,248,986	31,473,352

Table 11: Tied up capital within the 'pull' system, with quality problems

6.3.3. QUALITY PROBLEMS AND OPT SCHEDULING

Appropriately configured, OPT should combine the better points of the 'push' and 'pull' systems. The OPT strategy employed in the 'steady state' and 'demand fluctuation' scenarios was reused, this time being applied to the modified model, featuring the quality problem described in Section 6.3. Figure 59 shows the lead times achieved:
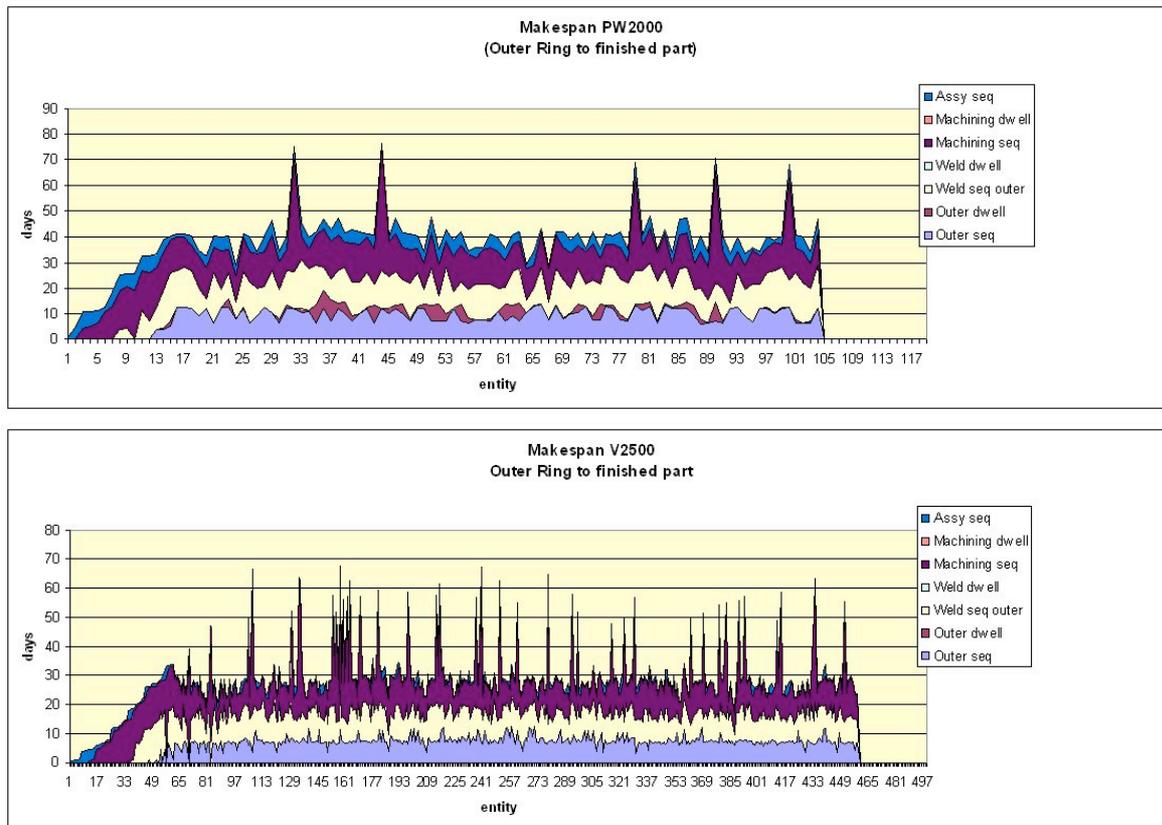
Figure 59: Lead times observed, with quality problems under OPT

As in Sections 6.1 and 6.2, the lead times exhibited by the OPT system are the lowest. Given such short lead times, the delay to a part suffering the quality problem is relatively great. The fill rate achieved in this scenario is perfect for the PW2000, but disappointing for the V2500, for which a problem is seen after week 26 (Figure 60). This leads to a four-week period in which the fill rate declines, after which is stabilises at around 30%. Effectively, the system is operating in arrears, with a backlog of around 17 parts. It does not manage to significantly erode the backlog during the remainder of the year, and the WIP for the V2500 remains at a very low level which suggests that TEC machining is unable to keep it supplied (Figure 61). Table 12 details the average level of capital tied up in WIP.
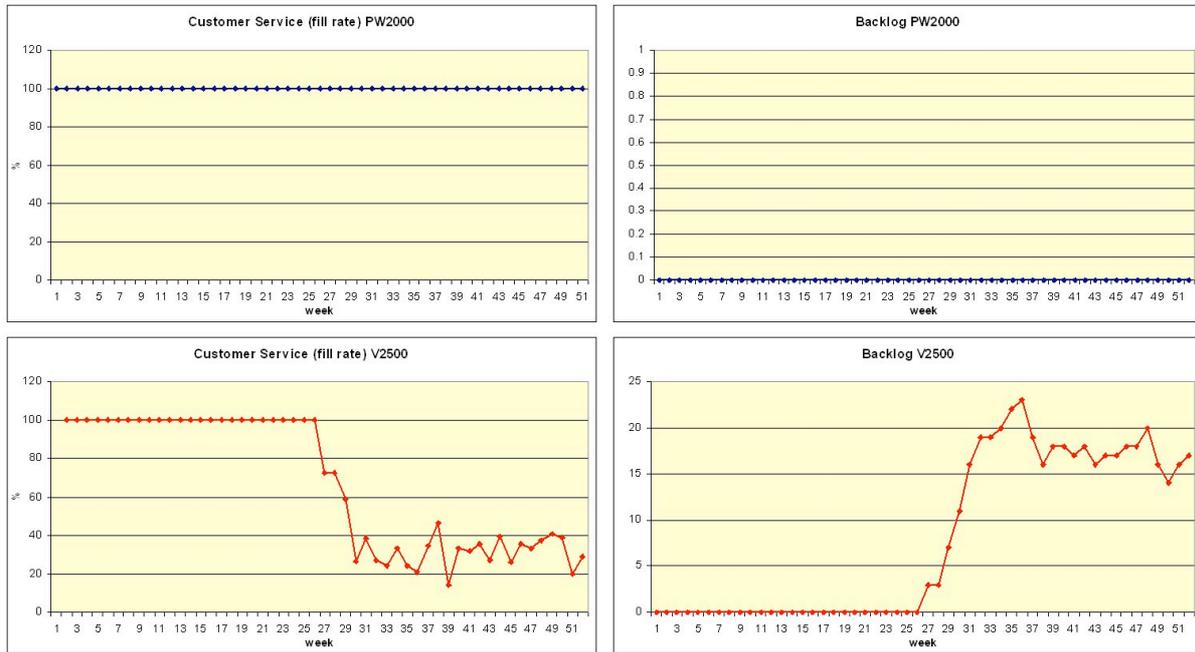
Figure 60: Fill rates and backlogs, with quality problems under OPT
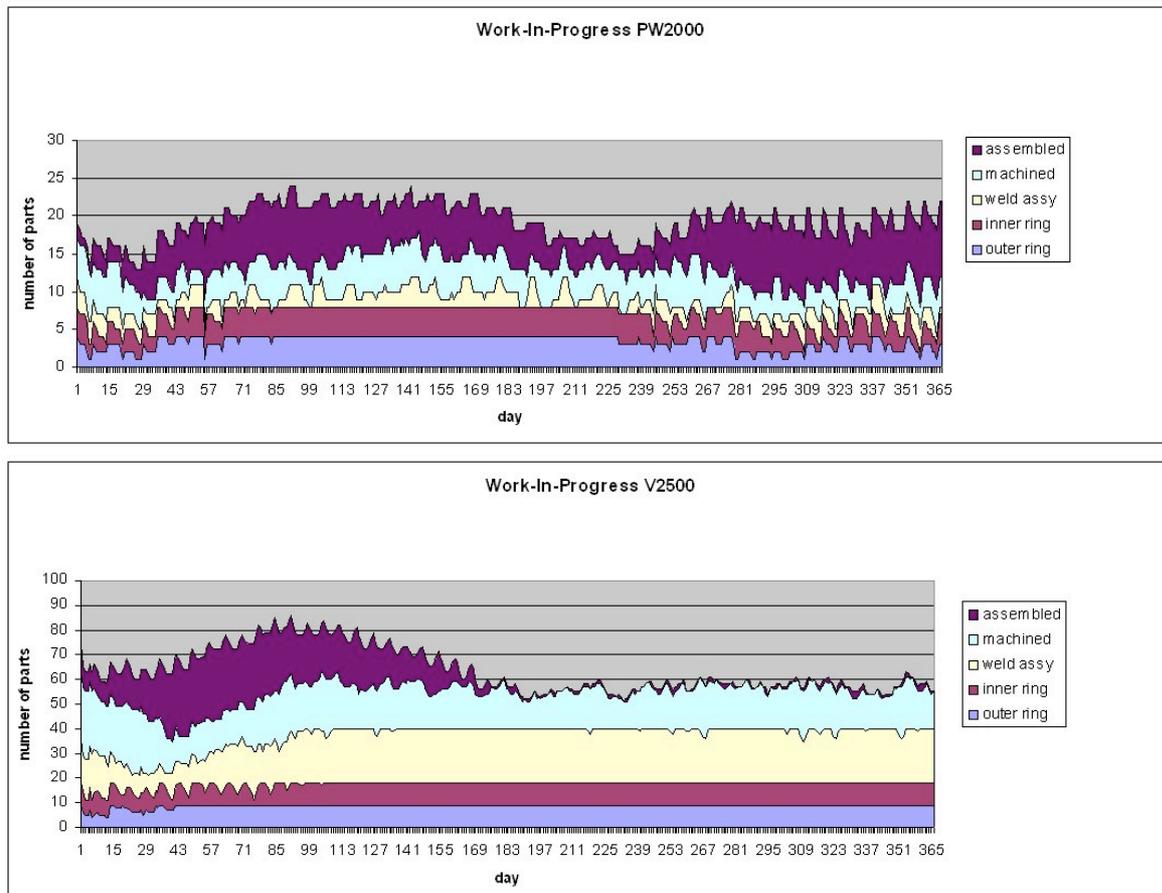
Figure 61: WIP levels observed, with quality problems under OPT

In truth, this is something of an unfair test for an OPT system, since the "drum, buffer, rope" concept demands that any troublesome process upstream of the bottleneck should be followed by a buffer large enough to mask its variability. In this programme of experiments it was desired to apply the same system of logistic control to each deviation scenario, so there is no such buffer, other than the usual sequence output buffer. With a failure rate of 12% causing a delay of 2–6 weeks, an average of four parts can be expected to be undergoing additional inspection and remediation at any given time, although a run of bad luck could increase this figure. The size of the buffer would need to be determined by experimentation, having selected a desired fill rate.

Tied-up Capital
Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	1,599,013	1,185,693
Inner Ring	369,112	367,075
Weld Assy	11,347,400	8,095,431
TEC Assy	12,241,350	8,815,672
TEC Assy	7,533,249	5,620,829
Total	**33,090,124**	**24,084,700**

Table 12: Tied up capital within the OPT system, with quality problems

Figure 6.2 shows machine utilisation for the scenario. As with the demand pattern changes in Section 6.2.3, machine utilisation fluctuates considerably. The machinists who serve workcentres 'O', 'P' and 'AA' do so interchangeably, so the wildly changing utilisation is less inefficient (in terms of personnel) than it might seem.
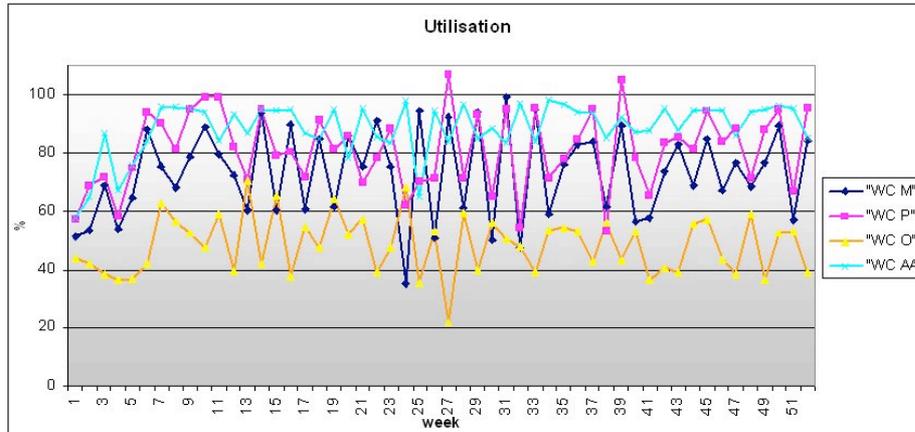
Figure 62: Observed utilisation of key workcentres, with quality problems under OPT

6.4. LACK OF KEY MATERIALS

While poor performance in the event of a raw material shortage is to some degree a reflection upon the quality of the supply chain rather than the manufacturing system itself, a programme of experiments was tried with the aim of showing if one or another of the logistic control mechanisms was better able to tolerate input variability, absorbing it and still meeting the delivery schedule.

The scenario selected concerns the supply out outer rings for the PW2000, although the shortage affects both TEC types' manufacture, as the subsections that follow describe.

The supply of PW2000 outer rings, normally reaching the focused factory at a rate of two per week, are constrained from week 17. For eight weeks, none arrive, and then the whole backlog is cleared by the arrival of a batch of eighteen. It is clear that this will impact upon the fill rate for the PW2000, but the time to recover and the effect upon V2500 production vary, depending upon the system of logistic control employed.

6.4.1. LACK OF MATERIALS AND PUSH SCHEDULING

Using standard lead times to determine the launch date for each component, the push system runs the risk of failing to respond adequately when there is a shortfall in the supply of one component. The other elements of the TEC's bill of materials will still be launched, and their production operations begun, even when it is impossible to complete the product.

The lead time for PW2000 TECs appears to get better when the shortage occurs, but it must be borne in mind that the x-axis of the lead time graphs (Figure 63) only shows the order in which products left the factory; it is not a timescale as such. The TECs appear to have been processed quickly, but this is because the facility has been waiting for outer rings.
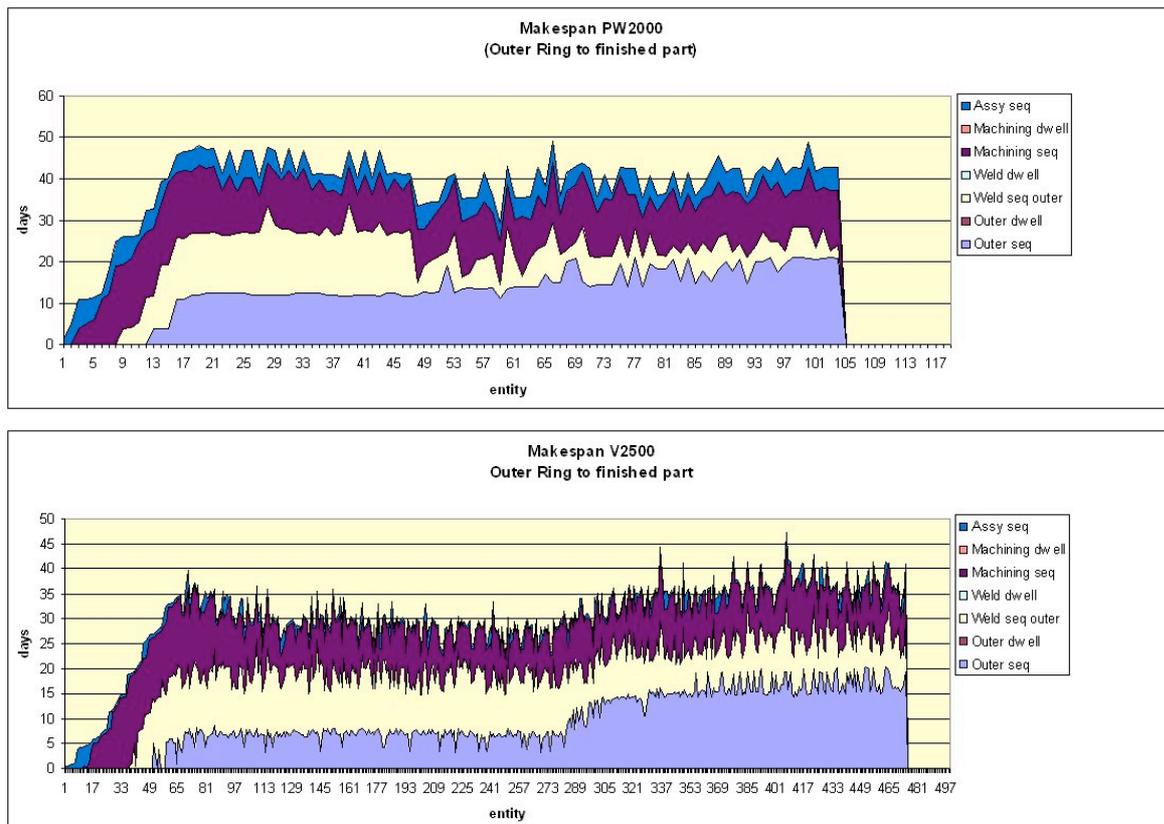
Figure 63: Lead times observed, with a supply shortage under 'pull' logistic control

The fill rate achieved is good, considering the 8-week stoppage on arrivals of outer rings (Figure 64). The material within the facility is depleted, but this 'flushing' allows the fill

rate to be maintained, to some degree. (The maximum backlog of nine TECs is only 4½ weeks' production, despite the eight week shortage.)
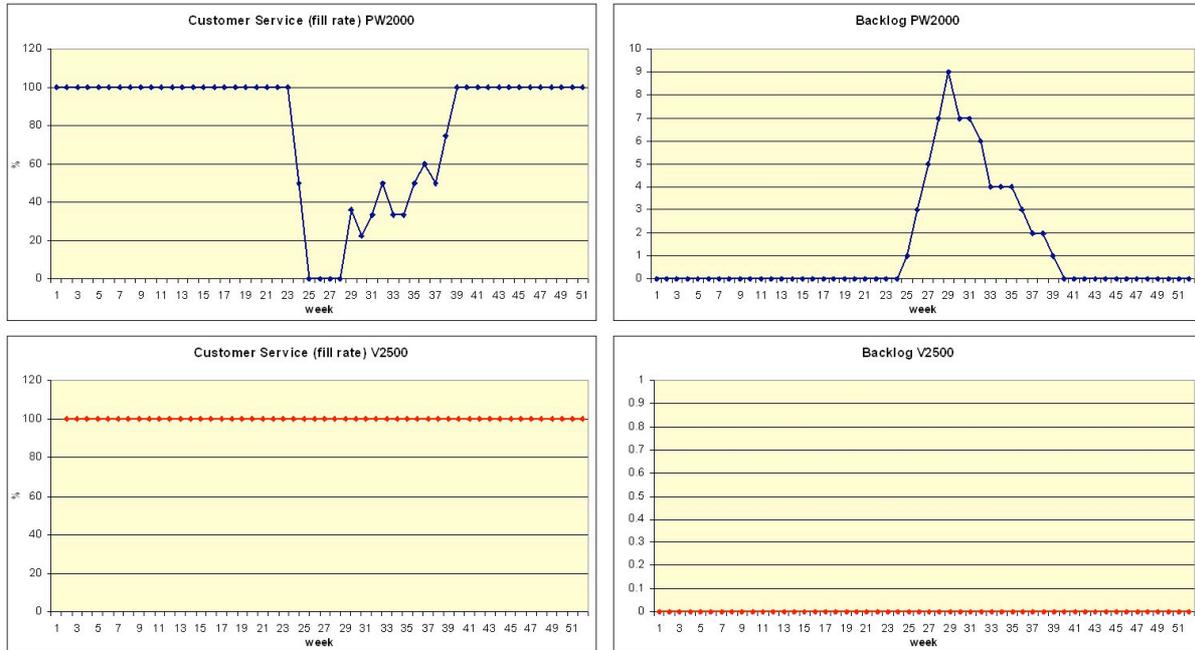
Figure 64: Fill rates and backlogs, with a supply shortage under 'pull' logistic control

The level of WIP within the system is highly problematic in this scenario. The focused factory continues to consume raw materials such as inner rings when no outer ring is available (Figure 65):
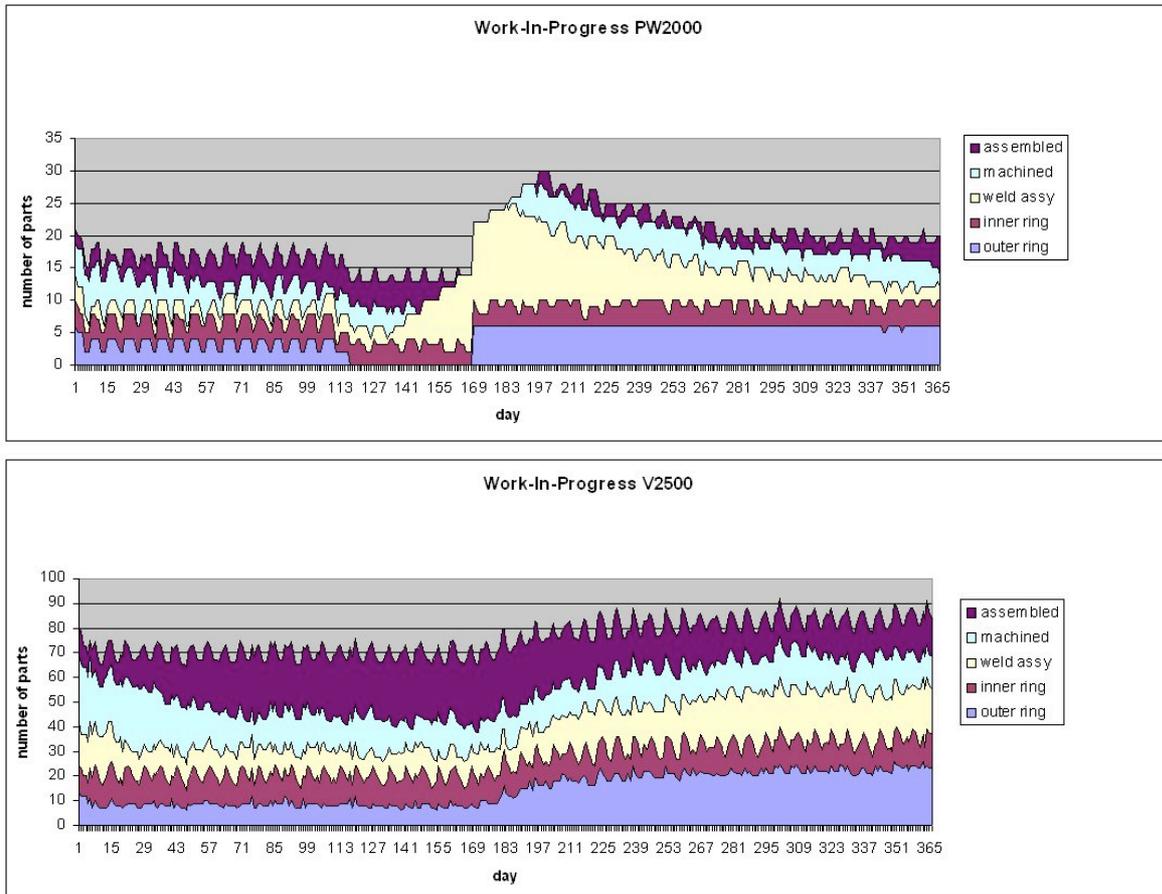
Figure 65 WIP levels observed, with a supply shortage under 'pull' logistic control

The utilisation of machines shows a distinct slump when the PW2000 outer rings stop being delivered (Figure 66). When the backlog of outer rings is supplied, a period of high utilisation begins, and the PW2000 backlog is cleared.
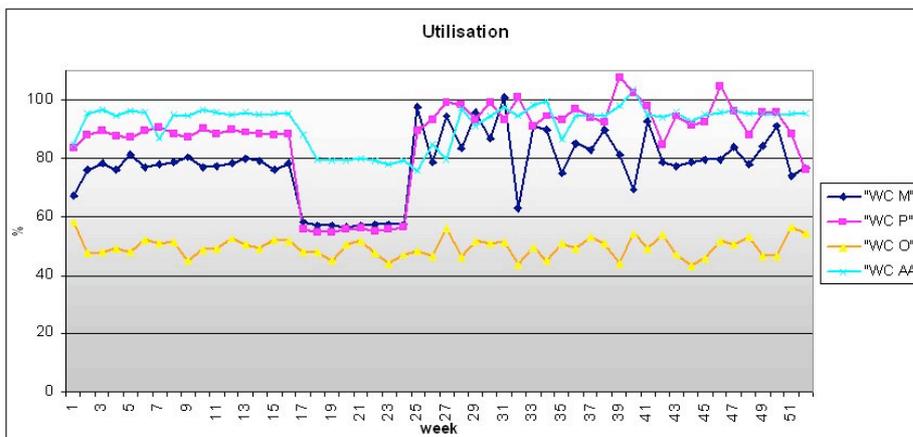
Figure 66: Observed utilisation of key workcentres, with a supply shortage under 'pull' logistic control

Table 13 shows the average amount of capital tied up in this scenario. It is substantially higher than for the same system of logistic control under steady supply conditions, due to the continued accumulation of inner rings.

Tied-up Capital
*Definition:*
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	2,650,036	1,965,042
Inner Ring	502,735	499,962
Weld Assy	8,415,739	6,003,934
TEC Assy (machining)	10,595,193	7,630,183
TEC Assy	16,216,731	12,099,889
Total	**38,380,434**	**28,199,010**

Table 13: Tied up capital within the 'pull' system, with a supply shortage

6.4.2. LACK OF MATERIALS AND PULL SCHEDULING

The 'pull' scheduling approach responds better to the supply shortage, reacting to the shortage of outer rings by ceasing to process inner rings once the output buffer from the inner ring sequence is full. As in the basic scenario, lead times are longer than with 'push', but here they are more stable – particularly the tangentially affected V2500 – and recover to a state that more closely resembles that before the shortage, as Figure 67 shows:
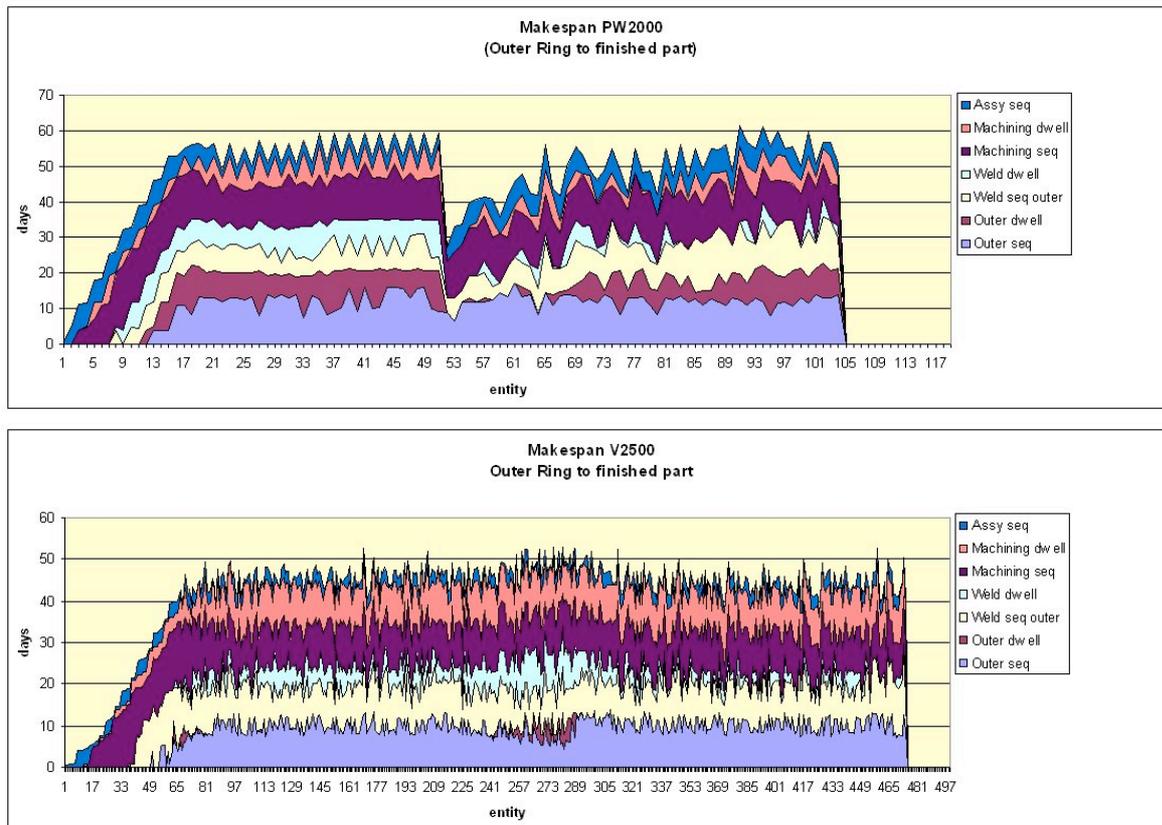
Figure 67: Lead times observed, with a supply shortage under 'pull' logistic control

The fill rate for the PW2000 suffers, with the backlog lasting some thirteen weeks, but recovery does not spoil the V2500 fill rate (Figure 68). In demanding exactly the workpieces that are needed, when then can be used, the 'pull' system does not choke the manufacturing system when the material shortage comes to an end; Figure 69 shows the tight control that is maintained over WIP levels.
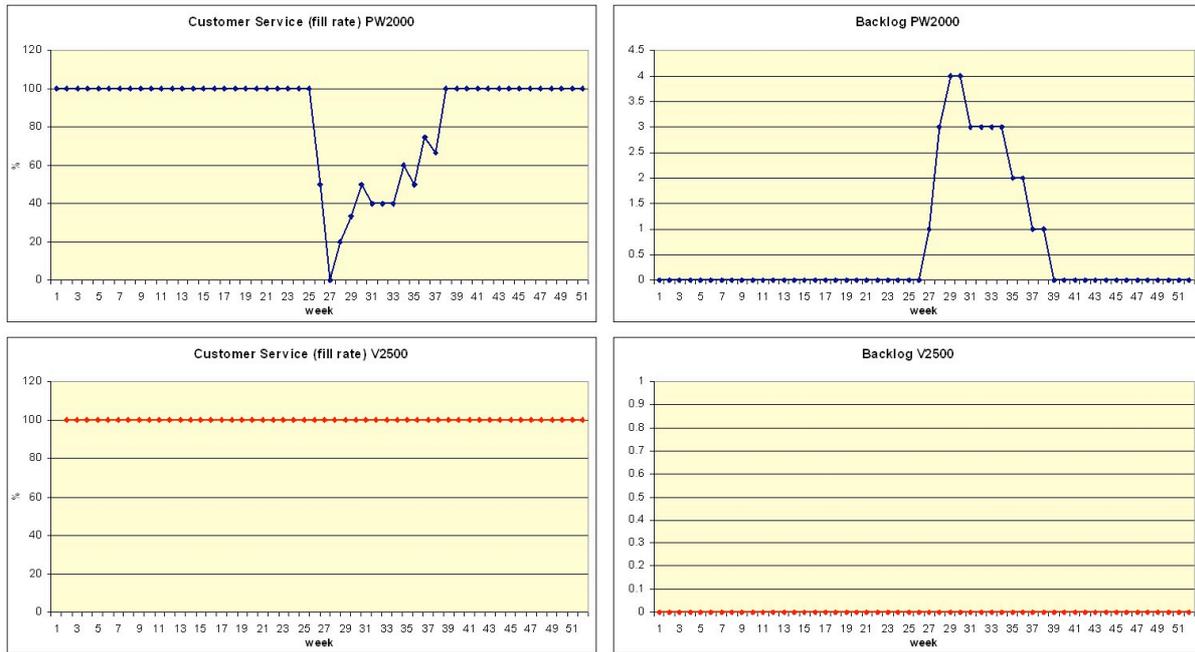
Figure 68: Fill rates and backlogs, with a supply shortage under 'pull' logistic control
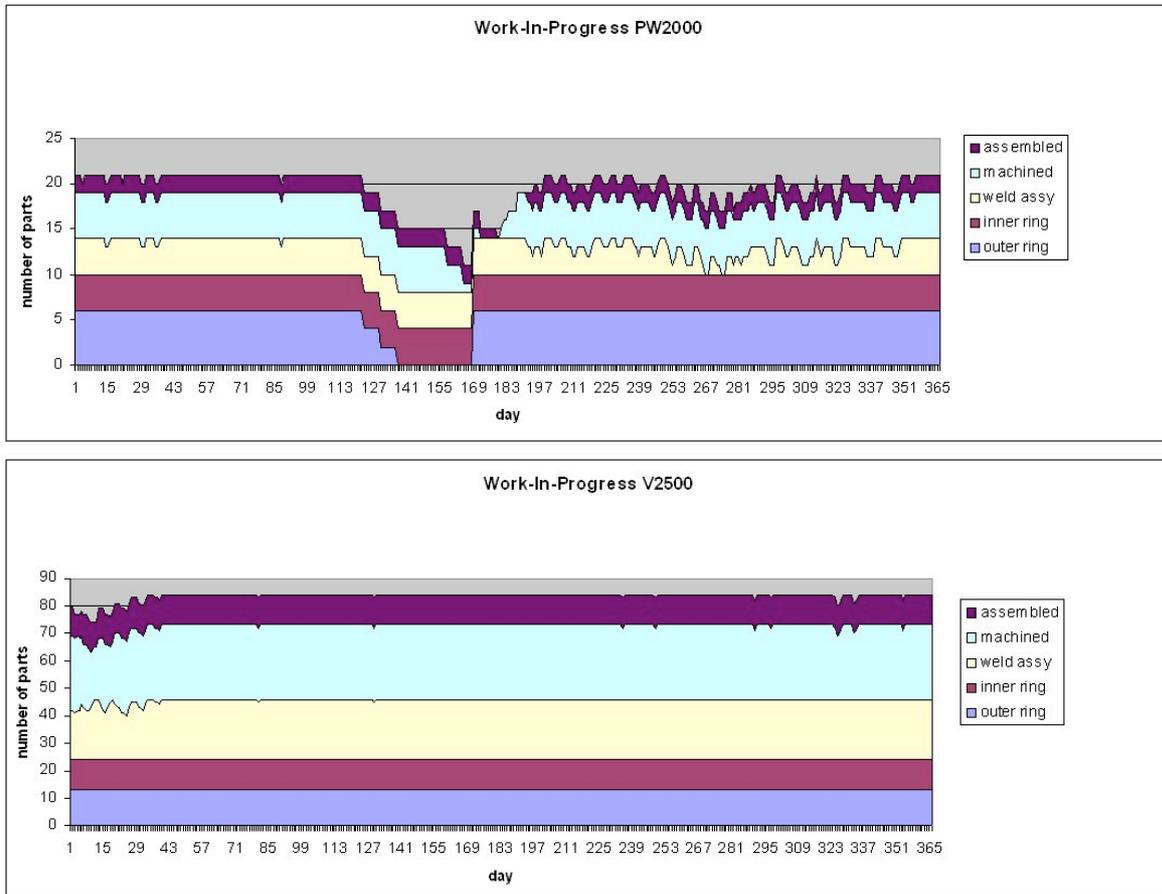
Figure 69: WIP levels observed, with a supply shortage under 'pull' logistic control

The period when the shortage of material most affects the factory can be seen in Figure 70, where the utilisation of workcentres 'M' and 'P' drops for a period. The utilisation of workcentres 'O' and 'AA' drop as the recovery begins; this makes sense since machinists are shared between 'O', 'P' and 'AA'.
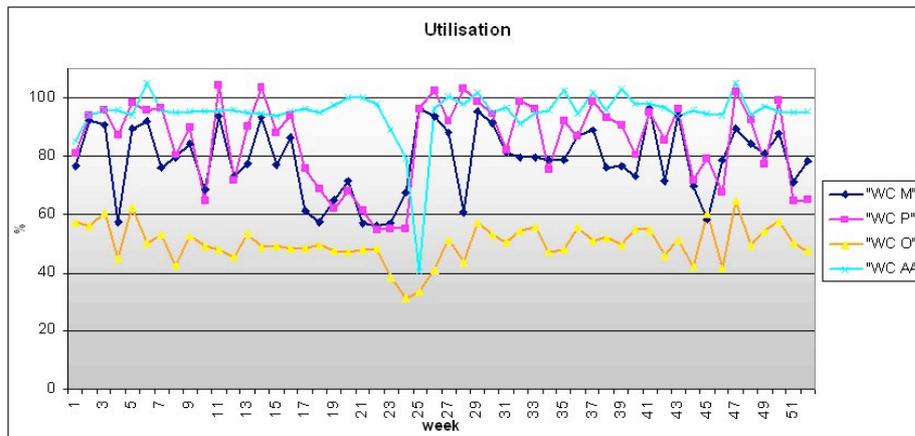
Figure 70: Observed utilisation of key workcentres, with a supply shortage under 'pull' logistic control

Table 14 shows the tied-up capital within the 'pull' system. Disappointingly, it is higher than for the 'push' system, which indicates that the target WIP levels used in the experiment are too high:

Tied-up Capital
Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	2,376,790	1,762,426
Inner Ring	473,564	470,951
Weld Assy	12,515,029	8,928,438
TEC Assy (machining)	19,012,335	13,691,831
TEC Assy	8,928,212	6,661,662
Total	**43,305,930**	**31,515,309**

Table 14: Tied up capital within the 'pull' system, with a supply shortage

6.4.3. LACK OF MATERIALS AND OPT SCHEDULING

Once again, the combination of 'push' and 'pull' logistics produces the shortest lead times, as Figure 71 shows. Unfortunately, the fill rate suffers; this system is less able to cope with the supply shortage, and does not entirely recover by the end of the simulation period (Figure 72).
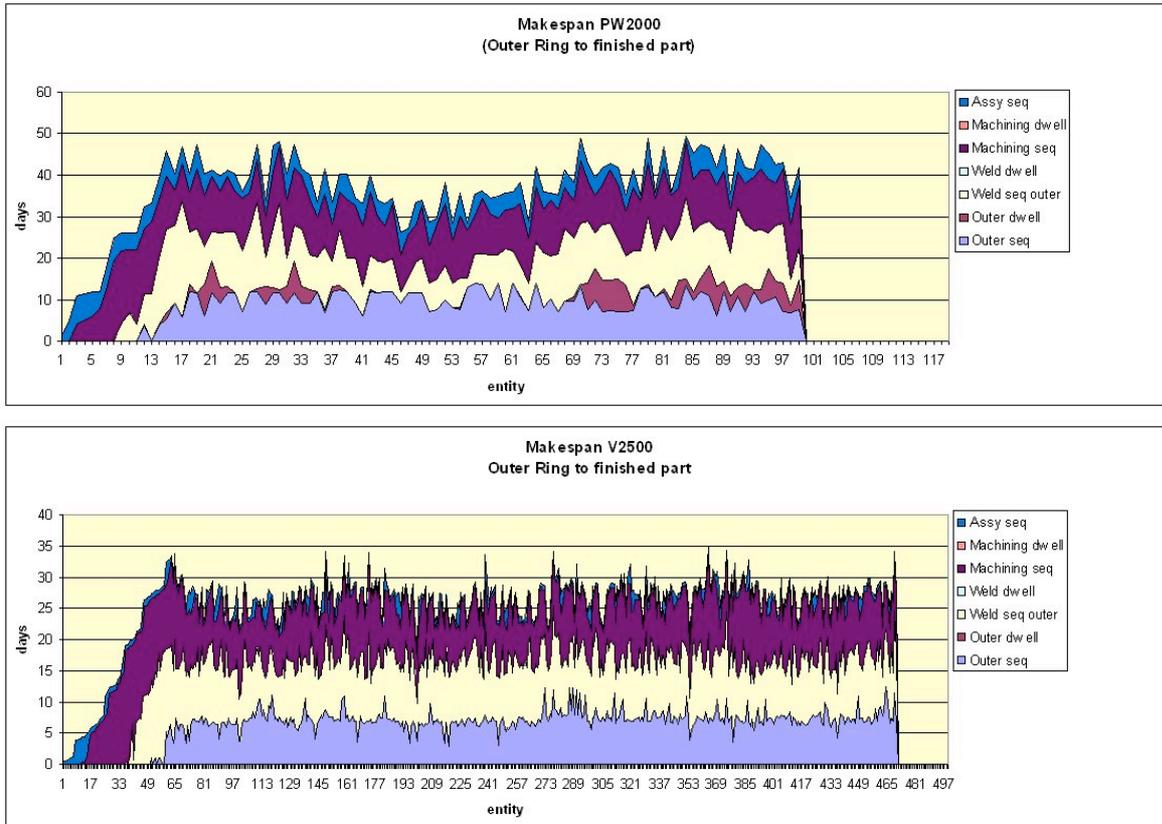
Figure 71: Lead times observed, with a supply shortage under OPT
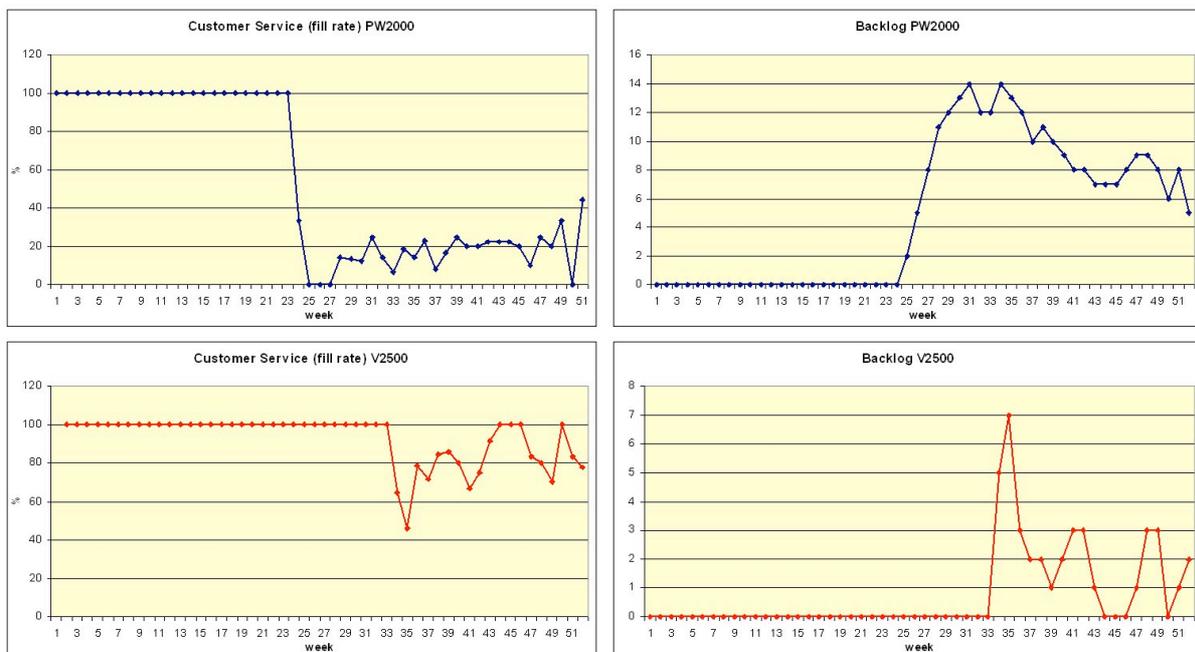
Figure 72: Fill rates and backlogs, with a supply shortage under OPT

The WIP levels observed within the system clearly show the 'pull' scheduled sequences, as bold, plain stripes where the inventory is held constant. The 'push' scheduling downstream of the bottleneck causes an overstock of finished goods (shown as being in the final sequence) but this is eliminated by the shortage, and thereafter the operations downstream of the bottleneck are very lean (Figure 73). Utilisation figures fluctuate a great deal as the system works to maintain target WIP levels (Figure 74).
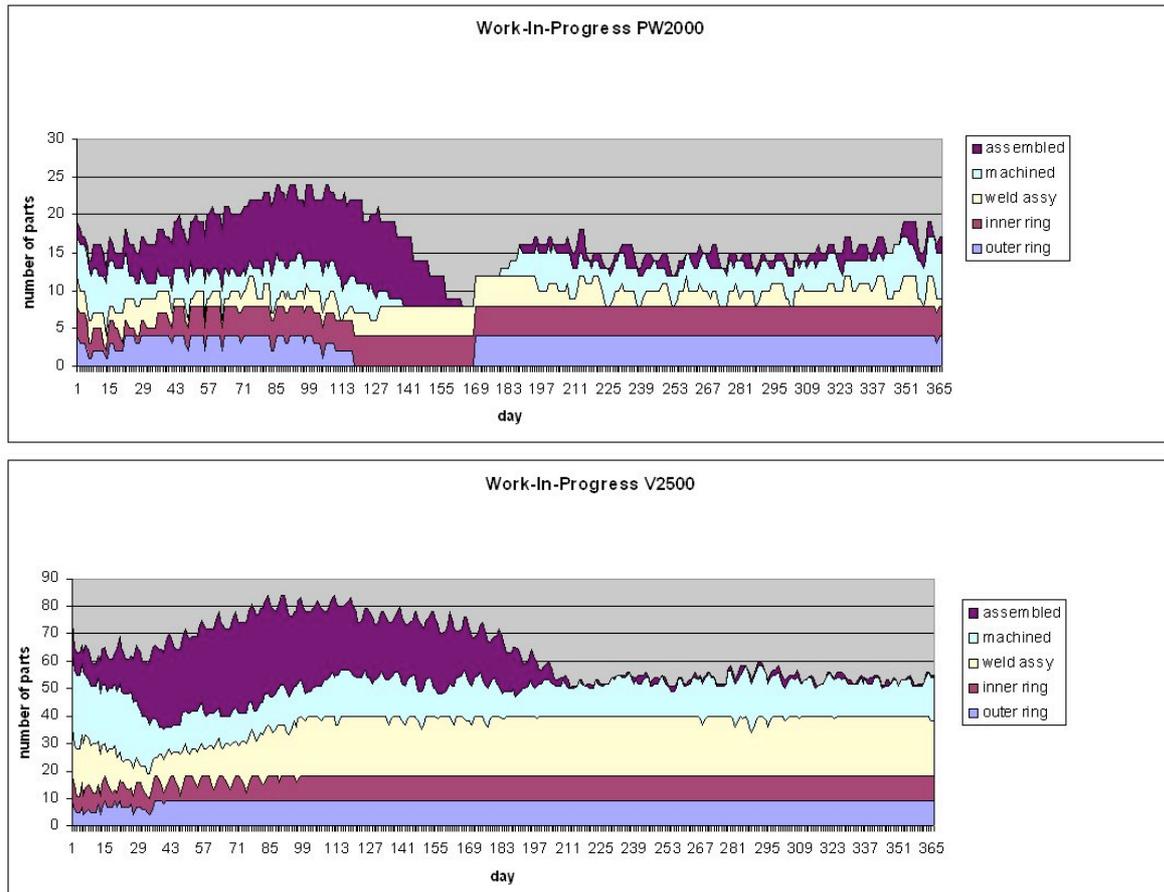
Figure 73: WIP levels observed, with a supply shortage under OPT
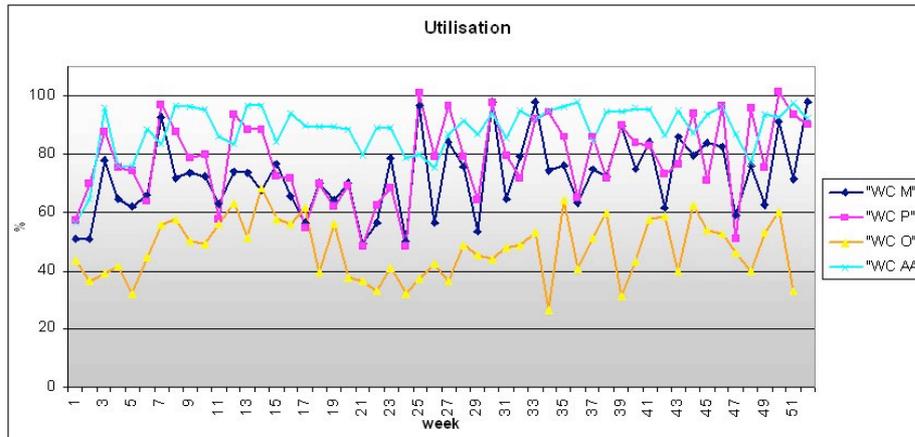
Figure 74: Observed utilisation of key workcentres, with a supply shortage under OPT

Table 15 shows the amount of capital tied up in inventory under this scenario. It produces the lowest WIP cost of the three systems of logistic control:

Tied-up Capital

Definition:
Standard cost of the average WIP per area.

|	PW2000	V2500
Outer Ring	1,599,013	1,185,693
Inner Ring	371,111	369,064
Weld Assy	11,218,363	8,003,374
TEC Assy (machining)	9,810,939	7,065,399
TEC Assy	10,581,255	7,895,057
Total	33,580,682	24,518,587

Table 15: Tied up capital within the 'pull' system, with a supply shortage

6.5. MACHINE BREAKDOWN

The machine breakdown scenario is a particularly fraught one for for focused factory, because it is operating very close to its maximum capacity. This can be seen from the processing times at some operations, given the volume of products that must flow through them.

It was decided to investigate a breakdown in Workcentre 'P', something that had occurred in the past. The breakdown was not left to chance, since this might bias the results and make comparison between the three systems of logistic control difficult. Instead, it was determined that a breakdown would take place ninety-five days into the simulation run, and that the machine should be out of commission for eighteen days. The

timing of the breakdown allows the model to settle into a steady state before the disruption begins, but allows enough time to see if it can recover.

6.5.1. MACHINE BREAKDOWN AND PUSH SCHEDULING

Under a 'push' system, the model continues to launch raw material, despite the fact that it will stall at Workcentre 'P'. Both TEC types visit this resource, so both are affected. A backlog builds up at the workcentre, and since it is operating very close to its maximum output, that backlog can clear only slowly. Progress against this is shown as a slight reduction in the outer sequence lead times, as Figure 75 shows:
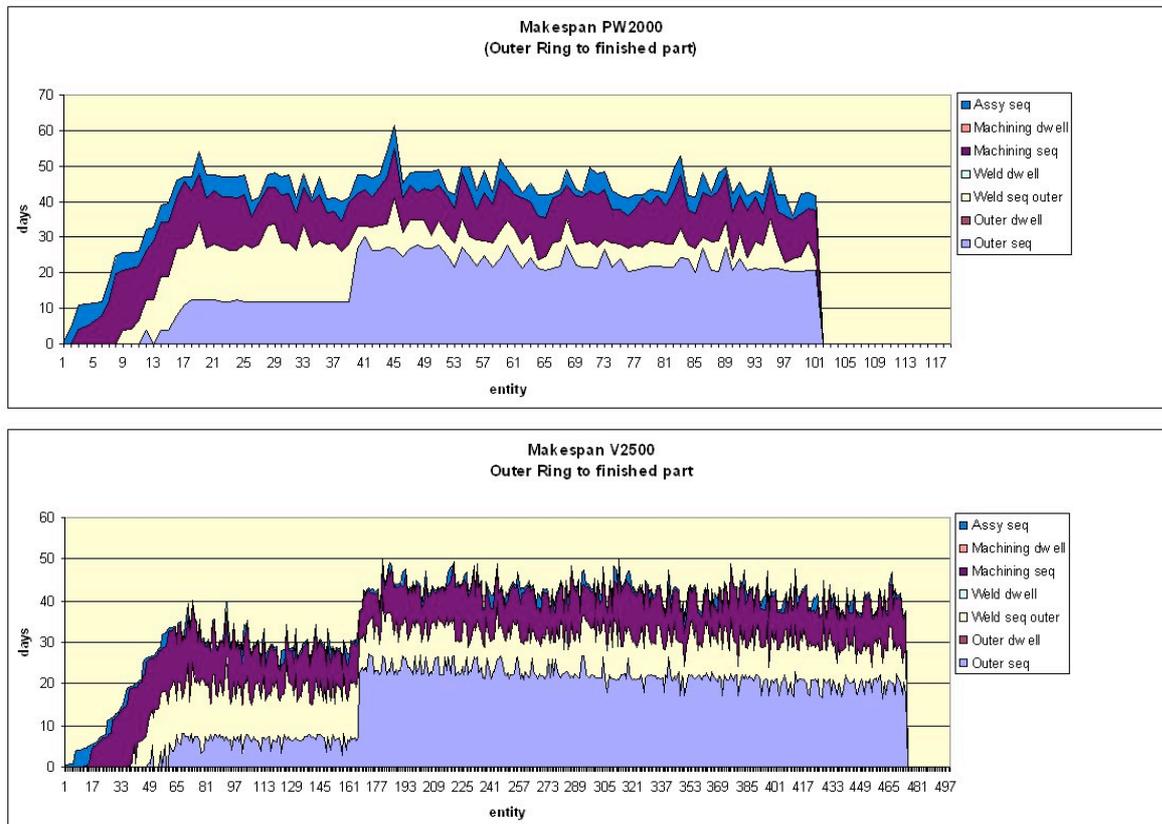
Figure 75: Lead times observed, with a machine breakdown under a 'push' system

The fill rate under this scenario suffers for the PW2000, although it remains perfect for the more numerous V2500. The backlog for the PW2000 is of up to three parts, or 1½ weeks' production (Figure 76). The V2500 comes perilously close to suffering a drop in the fill rate, as the stock of V2500 finished goods shows (Figure 77). These finished goods gradually increase in number again, although that is something of a shame, since the resources involved in their manufacture would have been better utilised in clearing the backlog of PW2000 orders. Under this relatively simplistic system of logistic control there is no facility for examining the system's obligations, however.
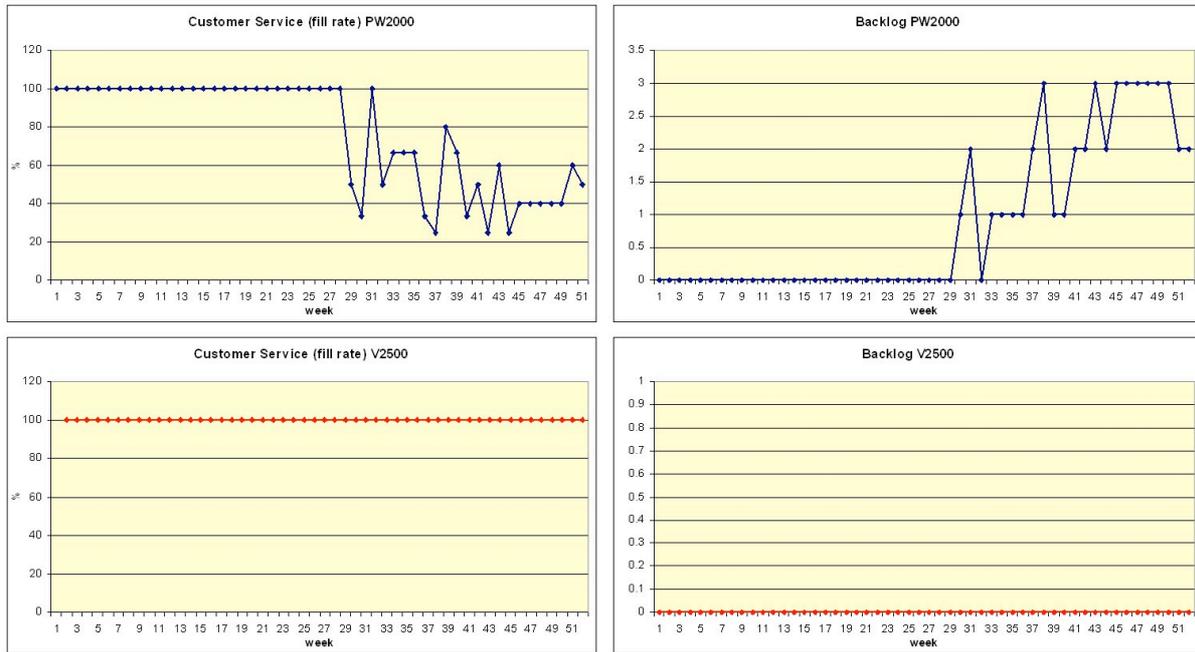
Figure 76: Fill rates and backlogs, with a machine breakdown under a 'push' system
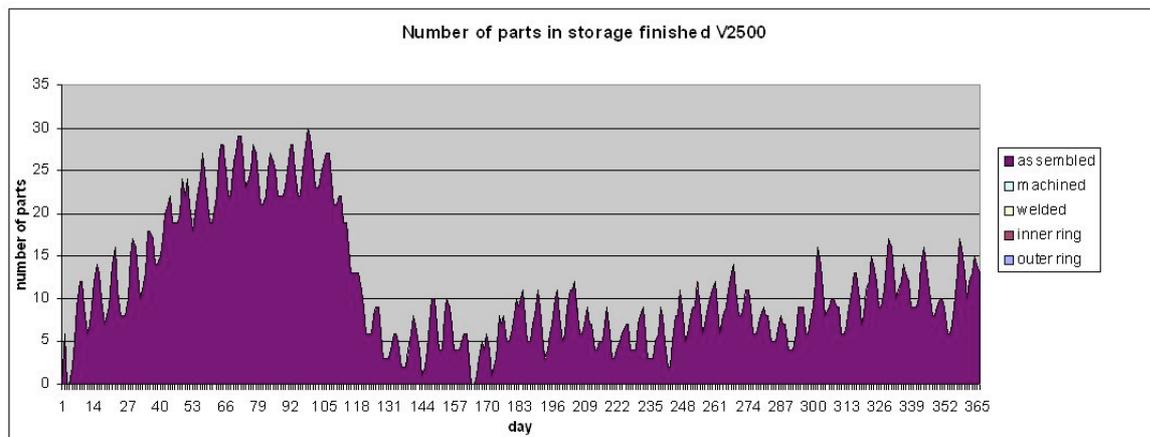
Figure 77: Stock of V2500 TECs, with a machine breakdown under a 'push' system

In Figure 78, the lasting effect of the disruption to Workcentre 'P' can be seen, in the form of an excess of work within the welding sequence. The backlog of parts arriving late from outer ring machining cannot be accommodated. Without intervention, it seems unlikely that the system will return to its previous state, despite the fact that utilisation of Workcentre 'P' falls back to a normal level after about thirteen weeks, suggesting that it has dealt with the backlog caused by its breakdown (Figure 79).
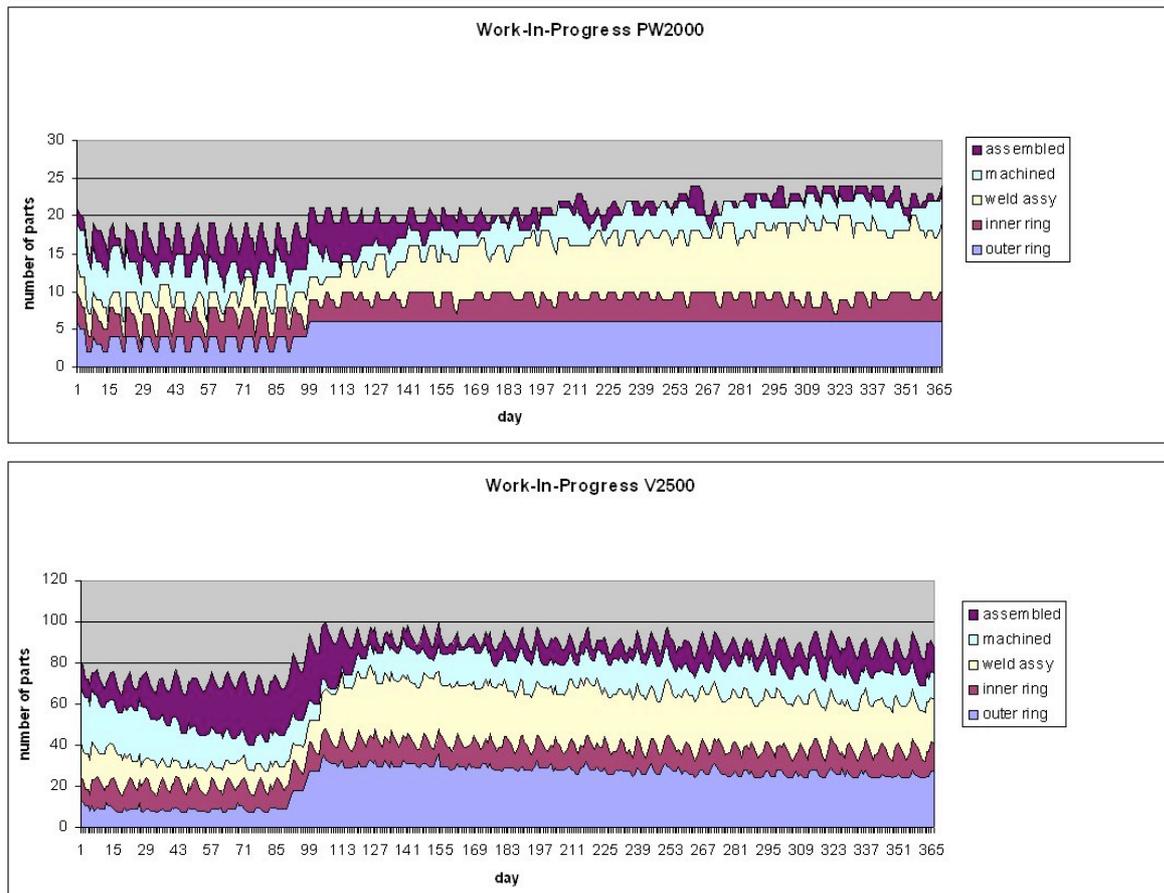
Figure 78: WIP levels observed, with a machine breakdown under a 'push' system
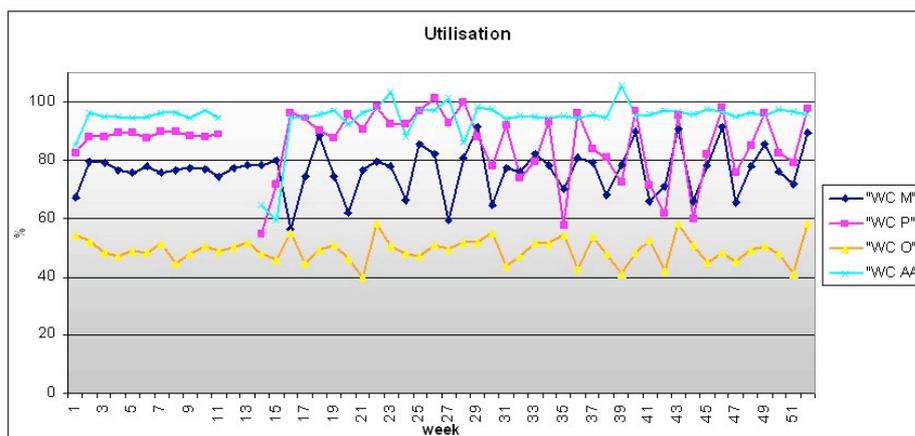
Figure 79: Observed utilisation of key workcentres, with a machine breakdown under a 'push' system

Table 16 shows the average value of WIP within the manufacturing system under this scenario:

Tied-up Capital
Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	4,198,596	3,113,323
Inner Ring	503,794	501,014
Weld Assy	12,876,962	9,186,647
TEC Assy (machining)	10,698,078	7,704,276
TEC Assy	10,316,501	7,697,515
Total	38,593,932	28,202,776

Table 16: Tied up capital within the 'push' system, with a machine breakdown

### 6.5.2. MACHINE BREAKDOWN AND PULL SCHEDULING

The machine breakdown scenario is one to which a 'pull' system should be well-suited, since it involves variability that cannot be included in a forecast. The existence of buffers should allow the disruption caused by the breakdown to be controlled, to some degree, while the failure of outer rings to complete their machining sequence should prevent additional WIP being drawn into the system.

Figure 80 shows the lead times observed during this experiment. For both TEC types, the increased time spent in outer ring machining sequence is clearly visible, but in both cases the impact is absorbed by a reduced time in processing and/or buffers downstream. The PW2000 lead times return to normal quite quickly, while the V2500 is not the same after the breakdown incident, although it adopts a new steady state. In this, the stock of finished goods after machining has been eliminated; pulled onwards to final assembly and then delivery to the customer. As Figure 81 shows, the WIP in the later stages of V2500 manufacture is below the target level, although welding has a full amount of WIP. This suggests that welding has become the bottleneck under these conditions, for some reason. In a real system, managers would intervene, putting a recovery strategy in place as soon as Workcentre 'P' became usable again. This might be achieved by changing shift patterns (if possible), or obtaining the use of a resource from outside the business unit, in order to clear the backlog. Changing the queues at machine tools from their present-day 'first in first out' rule might also be of some use.

The fill rate for both TEC types was a uniform 100%, so no fill rate / backlog plots are shown. Utilisation of the key workcentres is shown in Figure 82; several of them can be seen coming to a halt as a result of the breakdown. Only buffers *between machines* could prevent this – with consequences for level of capital tied up in the system.
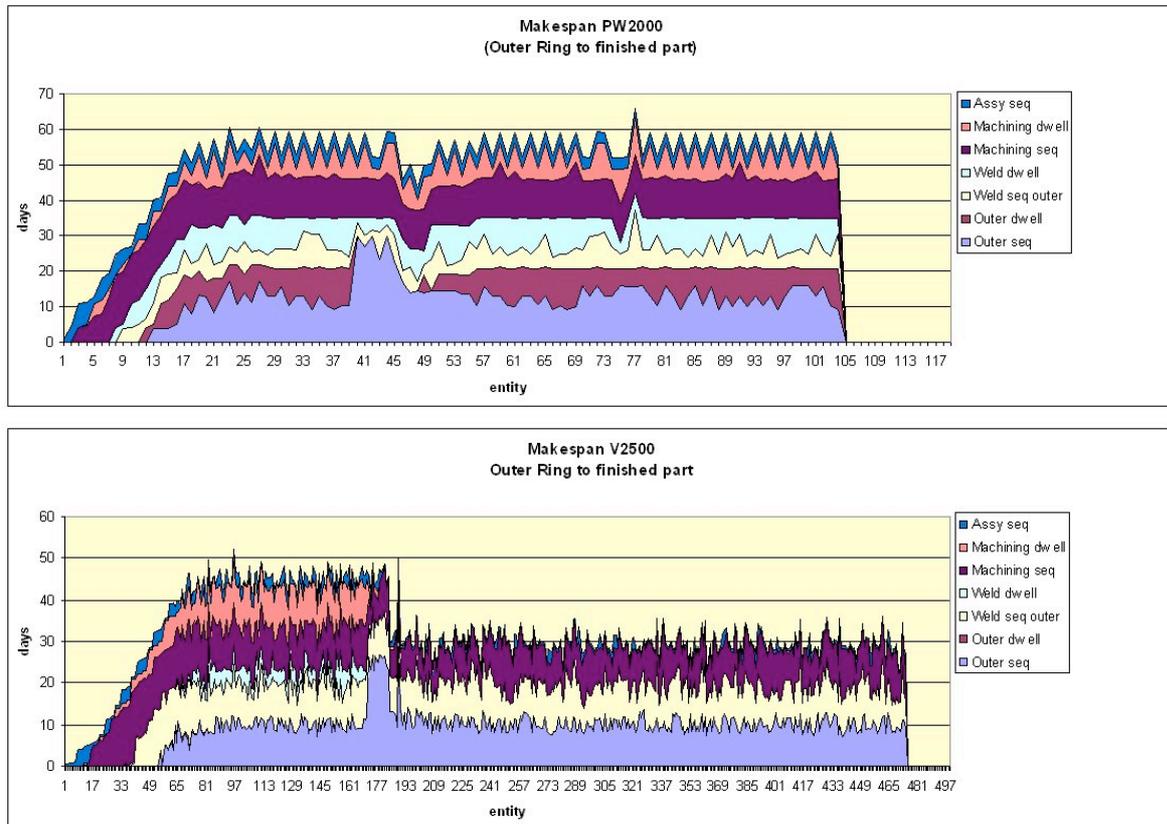
Figure 80: Lead times observed, with a machine breakdown under a 'pull' system
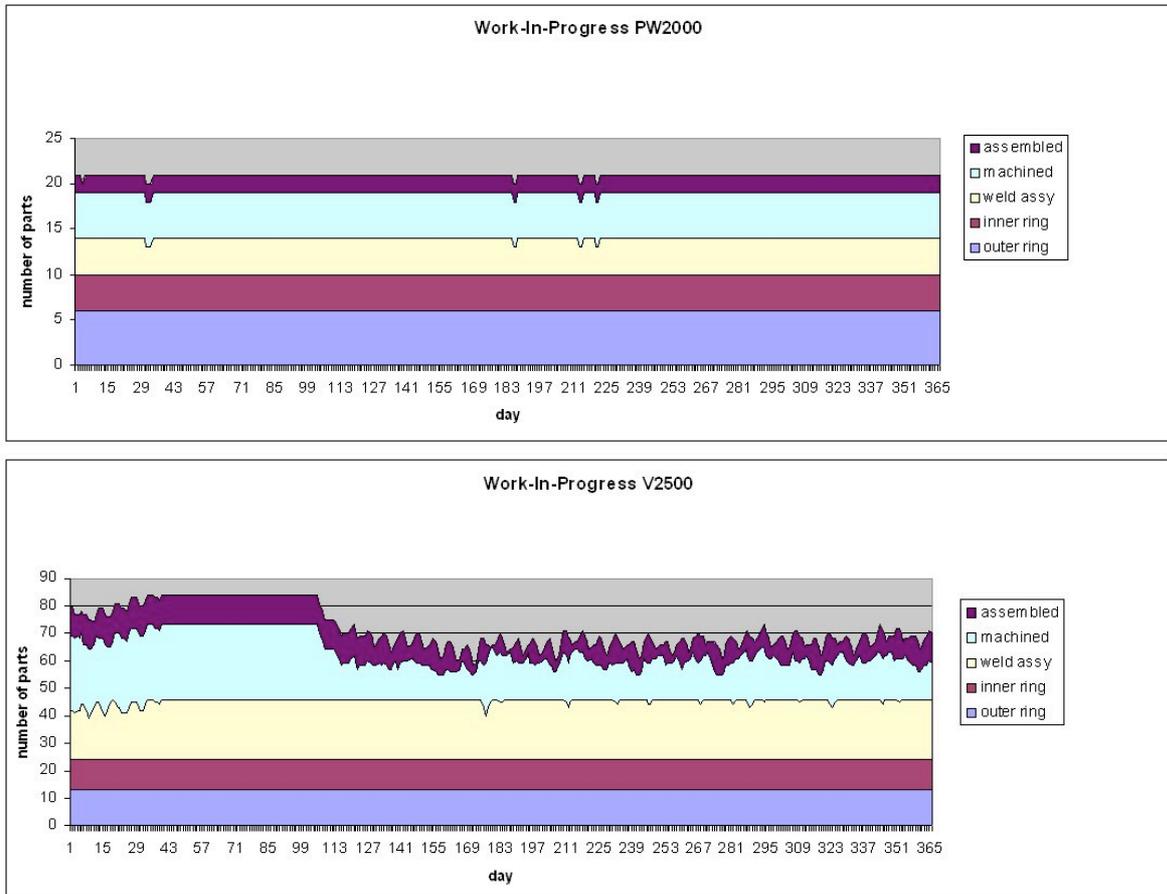
Figure 81: WIP levels observed, with a machine breakdown under a 'pull' system
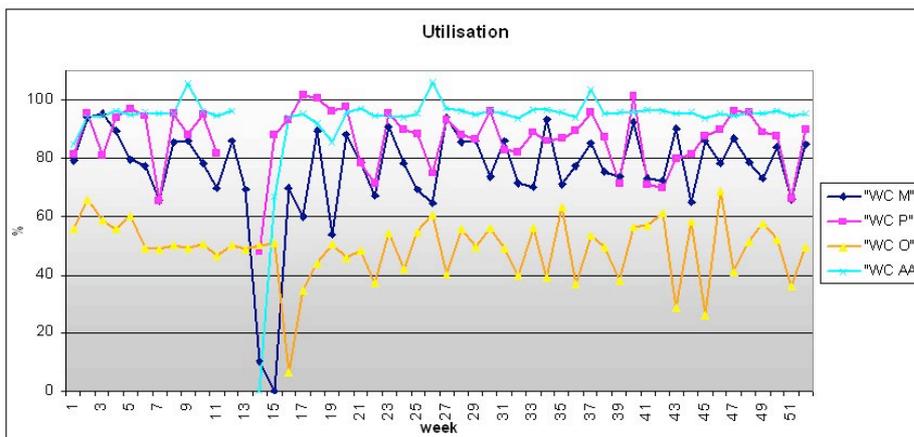
Figure 82: Observed utilisation of key workcentres, with a machine breakdown under a 'pull' system

Table 17 shows the average WIP cost for the scenario:

Tied-up Capital

Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	2,376,790	1,762,426
Inner Ring	473,564	470,951
Weld Assy	12,423,759	8,863,324
TEC Assy (machining)	12,452,943	8,968,052
TEC Assy	6,894,725	5,144,404
Total	34,621,780	25,209,157

Table 17: Tied up capital within the 'pull' system, with a machine breakdown

6.5.3. MACHINE BREAKDOWN AND OPT SCHEDULING

Again, the OPT system of logistic control displays short lead times (Figure 83), but the impact of the machine breakdown can be seen as a sudden increase in the time some products spent undergoing the outer ring machining sequence. For both products, the system then returns rapidly to the 'steady state' lead times, and functions as designed for the remainder of the year. The fill rate for the V2500 suffers (Figure 84), falling sharply and at its worst having a backlog of 38 parts, although the system is clearly making inroads into this total by then end of the simulation run.
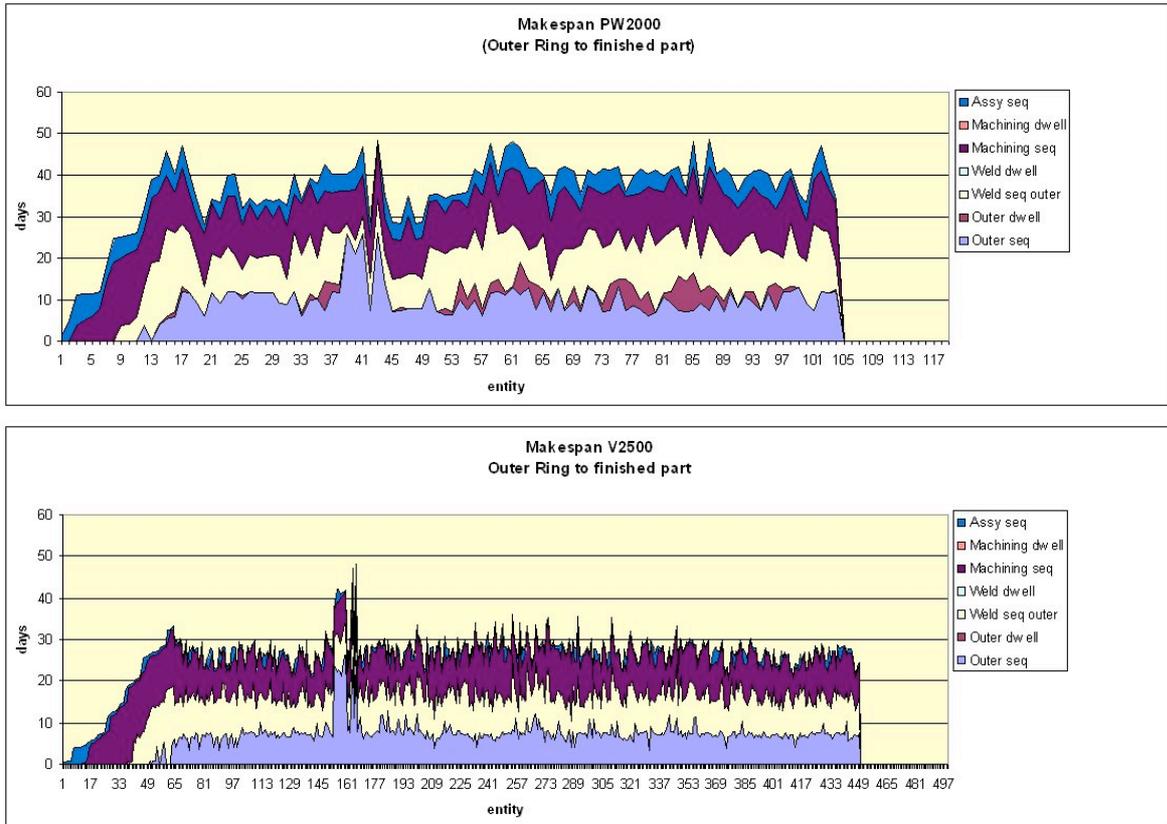
Figure 83: Lead times observed, with a machine breakdown under OPT
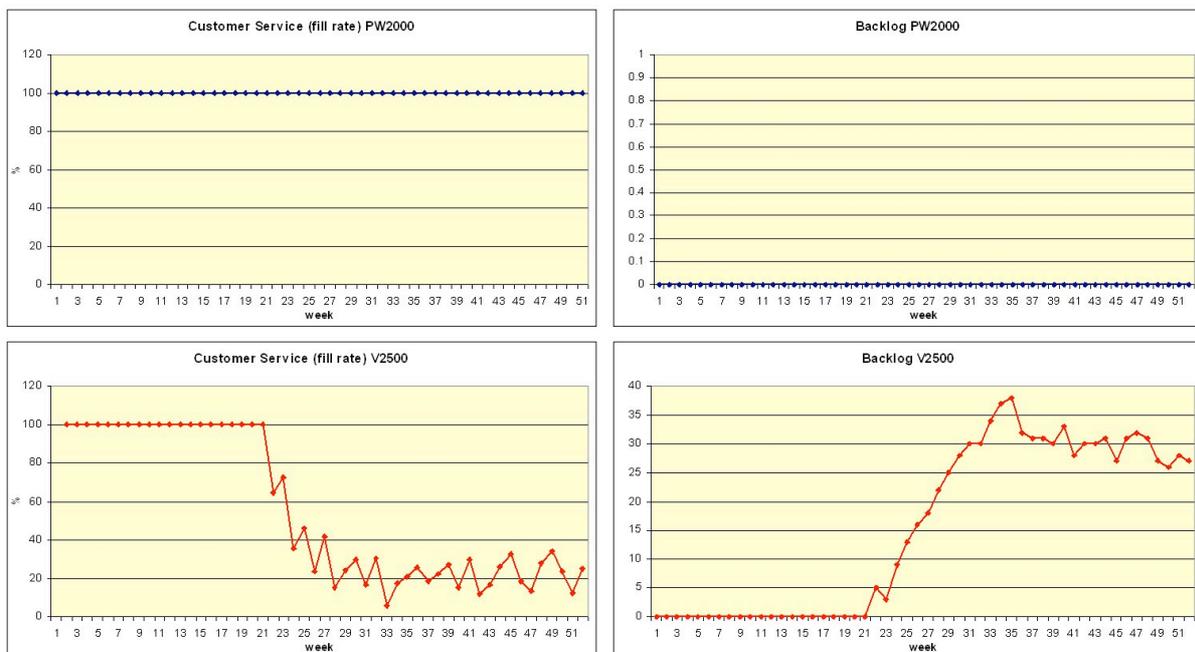
Figure 84: Fill rates and backlogs, with a machine breakdown under OPT

The figures returned for WIP found within the focused factory (Figure 85) show that the 'pull' scheduled early operations enjoy a steady level of WIP, once the manufacturing system has settled down. The number of components in the later manufacturing stages continues to fluctuate, however, and this may be a cause of the fill rate problems shows in Figure 84.
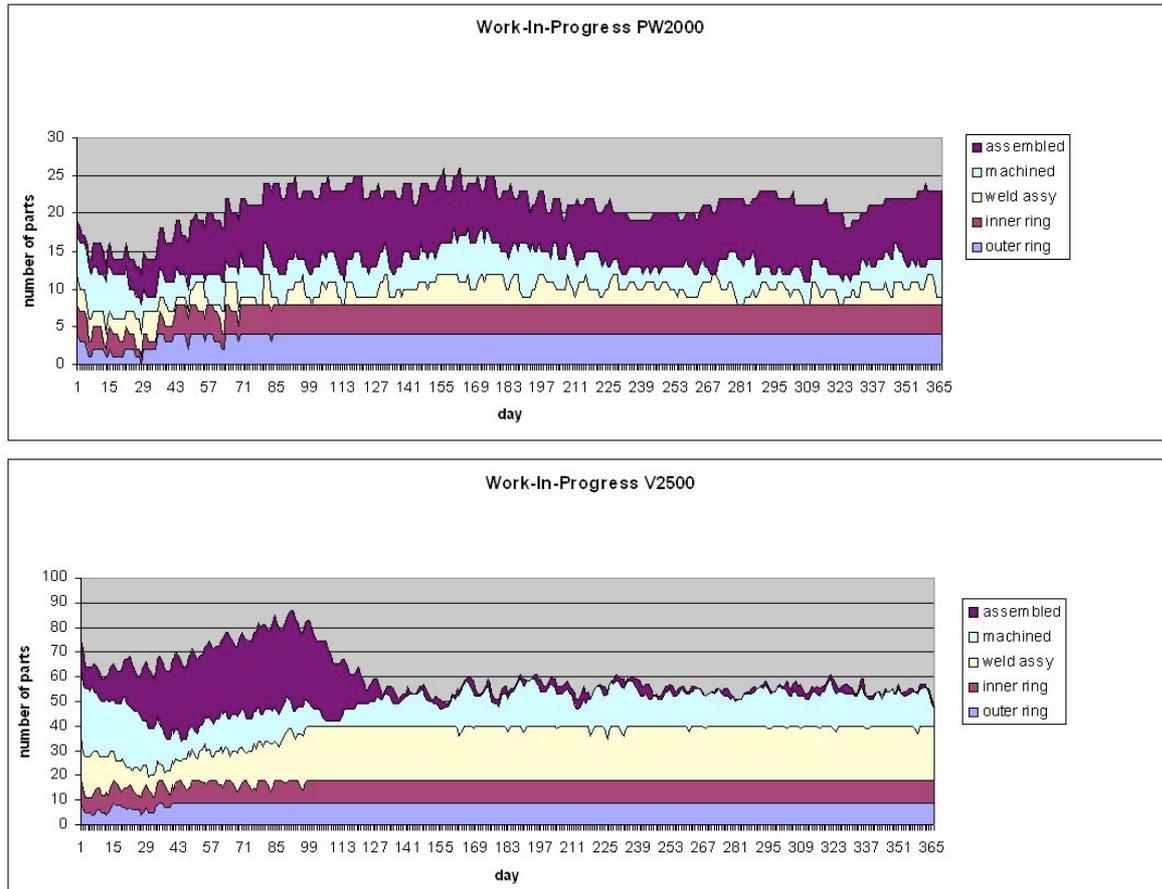
Figure 85: WIP levels observed, with a machine breakdown under OPT

The utilisation levels presented in Figure 86 again show the 'chaotic' usage pattern that appears to be a characteristic of this form of OPT; something that would need to be borne in mind if the system were to be implemented in reality.
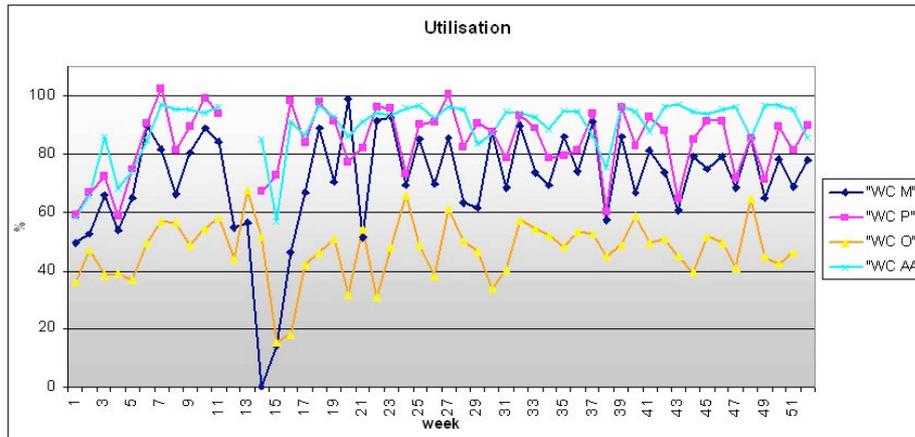
Figure 86: Observed utilisation of key workcentres, with a machine breakdown under OPT

Table 18 shows the amount of tied-up capital under this scenario:

Tied-up Capital
Definition:
Standard cost of the average WIP per area.

	PW2000	V2500
Outer Ring	1,592,519	1,180,877
Inner Ring	372,758	370,702
Weld Assy	11,281,308	8,048,280
TEC Assy (machining)	9,657,583	6,954,958
TEC Assy	7,884,771	5,883,113
Total	**30,788,938**	**22,437,930**

Table 18: Tied up capital within the OPT system, with a machine breakdown

7. DISCUSSION

The experiments detailed in Chapter 6 yielded some interesting results, although little that is at odds with existing theory. A 'push' system, for example, is known to create individual parts in a shorter time than a 'pull' system, *provided nothing occurs to prevent the system operating as designed.* The time products spend in buffers, waiting to be drawn into a downstream operation seems like a waste, although it is only truly a waste if the delay means that a bottleneck machine is left idle.

Lead time is just one measure of the effectiveness of the manufacturing system; level of WIP another. Chapter 5 described an array of performance metrics, any one of which might be selected as the subject of further investigation. The wide range of experiments carried out, such as the fifteen described in the preceding chapter, have illustrated the folly of attempting to 'optimise' such a complex manufacturing system. Being a real-world manufacturing system, it is subject to the vagaries of the marketplace, and the supply chain. Despite the high cost of retaining a part-built (or event completed) TECs within the factory, this 'waste' can be extremely useful when faced with a disruption. It is probably impossible to calculate the exact value of having a 'surplus' piece of inventory at a given place within the factory, since that would require a detailed understanding of the perturbations that might be expected, the penalty clauses that might be invoked when a component is not supplied, and less tangible costs such as the loss of reputation.

In one sense, performance measurement was actually made easier by VAC's relationship to the prime contractor. Because the focused factory produces each TEC for just one customer and with no variants, decisions relating to the sequencing of jobs are relatively simple. When there is a demand for a V2500 TEC, the next TEC to leave final assembly will fulfil that requirement. It doesn't have to be the TEC built from the outer ring and inner ring that were launched in response to the original order. (Rework activities and workcentres where parallel processing in possible mean that it is possible for one TEC to overtake another.) In a system where each product was differentiated to some degree, such as if the facility made cars with different engine size and body colour choices, it would be much more complicated to schedule jobs.

7.1. CHOICE OF EXPERIMENTS

The experiments that were detailed in the previous chapter do not present an exhaustive range of the scenarios that might challenge the focused factory. They do provide enough variety to show that no one system of logistic control can be wholly endorsed – or dismissed. The experiments were not entirely 'fair', although every effort was made to achieve consistency when testing each scenario. To compare the MRP-driven 'push' system of logistic control with a kanban-based 'pull' one is not fair; if the company is in possession of forecast information that is good enough to allow an MRP-driven approach, why ignore this information and require the system to respond directly to customers' demands at the *end* of the manufacturing sequence? Also, can it be reasonable to operate a 'pull' system that has not been permitted to benefit from the other initiatives that comprise lean manufacture, such as empowerment, quality initiatives and reduced changeover times?

7.2. SIMULATION PARAMETERS

One real difficulty in running the experiments detailed in the previous chapter concerned the parameters that should be used. With all systems, there was the question of initial WIP levels, although these were intended only to settle the model rapidly into a steady state such that experiments could take place. Another issue is the spacing of deliveries; there are many different delivery patterns that all equate to 'nine a week'. Their timing has a small but measurable effect.

More problematic among the parameters were the standard lead times to be used as an offset, determining the delivery date for raw materials based upon the due date for the finished product. For the 'push' system, these delivery dates can influence the resulting fill rate tremendously; padding out the standard lead time is a well-known way of improving the on-time delivery performance – at the cost of having additional WIP. A test was made in which the observed lead time from a 'steady state' model run was used as a new, reduced figure for the standard lead time; this was found to have the effect of *increasing* the lead time. In other words, small changes to the inputs to this complex system can have unpredictable outcomes. It is necessary to assign a standard lead time that is slightly pessimistic, in order to accommodate the inevitable variability that comes from sources such as the triangular distributions for processing time at each operation... but this, in turn, biases any judgement of fill rate.

For 'pull' systems, the key parameter is the target level of WIP within each sequence – effectively the number of kanbans in that part of the production system. Too few, and the system will run too lean, with customers being let down and machines standing idle. Too many, and the system operates with an excessive level of WIP, increasing costs and lead times.

For the OPT system of logistic control, both the standard lead time and (where a 'pull' system is operated) the target WIP value are relevant parameters. Such a system requires careful 'tuning' if it is to operate efficiently.

7.3. THE ROLE OF THE PLANNER

In a real system, a human planner would react to a problem scenario as it arose. Where the model has only its logistic control 'rules' to determine what work should or should not be started, some of the excesses displayed by the simulations should not happen in reality. A machine breakdown, for example, might be addressed by reallocating staff to other jobs, or obtaining the use of a substitute. The simulation can do neither of these things. Indeed, under a 'push' system, it goes on accumulating raw material, even though it will be impossible to complete any products.

From what has been observed during this study, it is clear that the staff in the use case facility are skilled, intelligent and well-motivated. It may be that a full implementation of a lean manufacturing approach does not suit this company. The impersonal control provided by kanbans may be valuable in highlighting problems in a culture where staff fear they might lose face if *they* called a halt to production, but such a culture does not appear to exist within VAC. The need for this aspect of the 'lean toolkit' is reduced, therefore.

7.4. THE BEST SYSTEM?

It is impossible to say which system of logistic control is the most appropriate for the focused factory at this time, given the multiple criteria against which each system must be judged. Also, as Section 7.2 explained, the parameters to which each system is operated means that there is no single candidate representing each logistic control approach. Instead, there is a near-infinite combination of control parameters, and deviation scenarios against which they might be tested.

Furthermore, the importance of the use case facility – and the fact that it is currently operating at a production rate that is very close to the theoretical maximum – suggest that it should not be tampered with. While it is true that the present-day system is operating in arrears, several of the experiments in the previous chapter yielded similar results, where deliveries were being made but little could be done to clear the backlog of orders. Given that the focused factory has suffered problems that cannot be prevented, such as shortages of raw materials, a backlog is almost inevitable.

If we return to the concept of plotting fill rate against the average value of WIP held by the system, as presented in Section 5.4 / Figure 23, it is now possible to use the experimental results obtained, to produce diagrams of this kind. (We have only single points of data for each scenario at this time; to plot lines showing performance against various WIP levels would require a substantial number of additional simulation runs.) The fill rates achieved came nowhere near 100% in some cases, as a result of the deviation scenario to which they were subjected, but some comparisons can be made, as Figures 87–91 show:
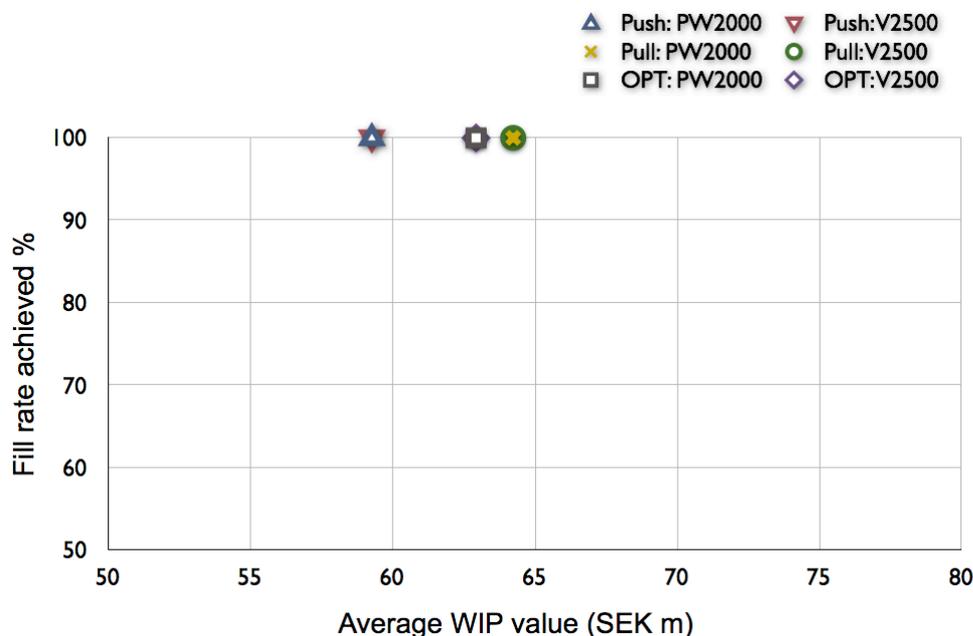

Figure 87: Comparison of logistic control systems under 'steady state' conditions
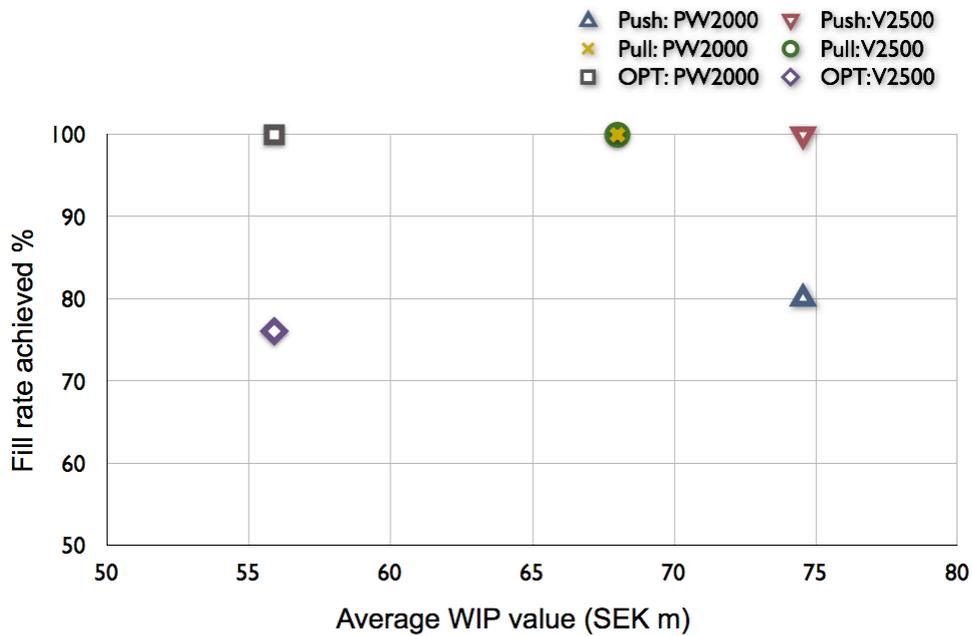
Figure 88: Comparison of logistic control systems under varying demand
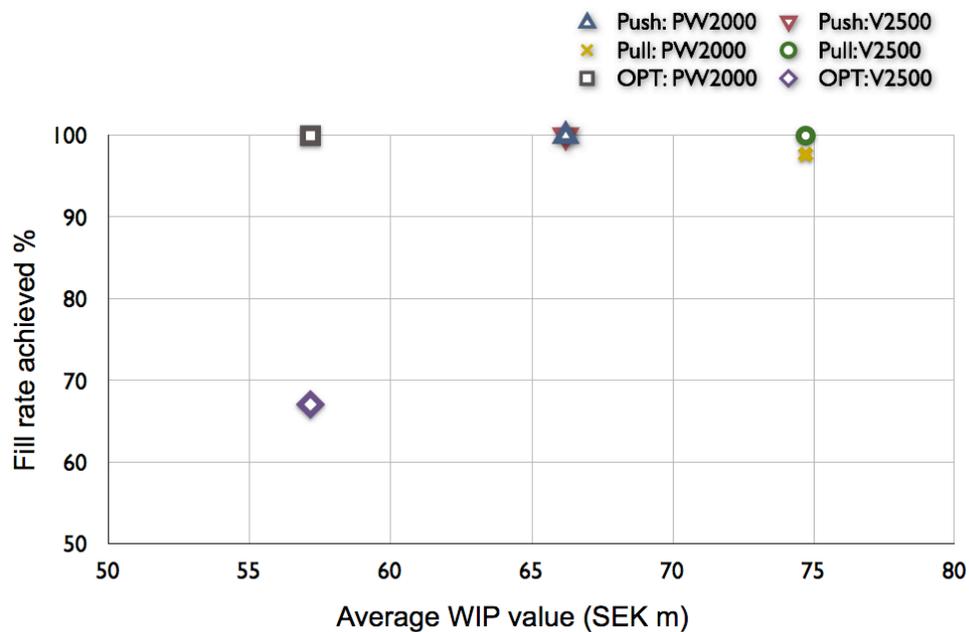
Figure 89: Comparison of logistic control systems with variable quality
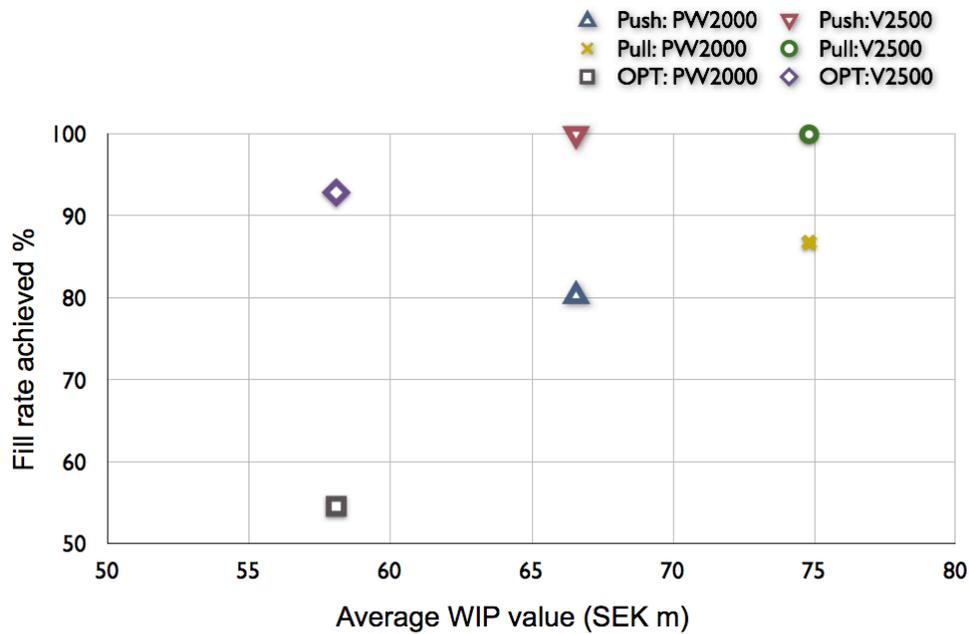
Figure 90: Comparison of logistic control systems with a supply shortage
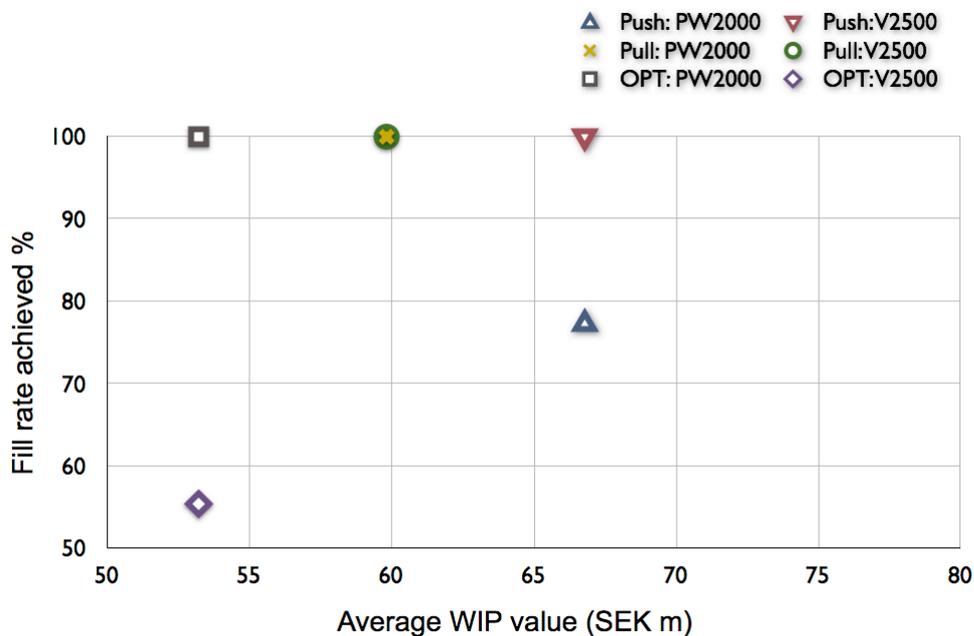
Figure 91: Comparison of logistic control systems with a machine breakdown

The system of logistic control selected would depend upon the planner's attitude to risk, and the likelihood of each deviation scenario coming about. Clearly, the OPT system minimises the tied-up capital, but the customer service level for the V2500 suffers under any conditions other than the steady state. Conversely, the 'pull' system produces the

best fill rates, at the cost of much higher WIP levels. Experiments with a very lean pull system had yielded terrible fill rates, but some further experimentation might be conducted, to produce a new, hybrid system of logistic control, paring away at the inventory where it is felt that a deviation scenario is unlikely, or that the resulting delays are not catastrophic.

8. CONCLUSIONS

This document has described the design and implementation of a complex simulation, representing a focused factory at Volvo Aero Corporation, where two kinds of Turbine Exhaust Casing are manufactured. The work has demonstrated how simulation can be employed to explore alternative production planning and control strategies, without risking disruption in a focused factory that is operating very close to its limits, and which could therefore fail if it were mistreated.

The approach to model construction was novel, allowing a variety of logistic control strategies to be selected, although only a single model needs to be maintained. This saved a considerable amount of time, and yielded greater confidence in the quality of the resulting simulations. It was also computationally efficient, allowing logistic control to take place as if kanbans were circulating within a sequence of operations, but without the need to represent them explicitly.

The use of Microsoft Excel for both the input of scenario details and the output of results provides a familiar, easy-to-use interface to the more complex model itself. As such, it makes the model available to a wider potential audience, while the use of an input file to populate the model with WIP at the beginning of a model run speeded the system's settling down into a steady state. It also raises the possibility of using the model to investigate scenarios in the very short term; a planner could use shop floor data to see the condition of the factory at that instant, and investigate the consequences of a schedule change, for example.

Finally, the Excel-based user interface means the model may be of use in the training and dissemination activities that form a part of this task. The modular nature of the simulation means that it could be scaled down to aid understanding during training. Just such an approach has been agreed with MTU.

There is no clear answer as to which system of logistic control would be best for the use case facility, given the considerable variety of deviation scenarios that might be inflicted upon it. Nonetheless, with a further programme of experiments, designed to represent plausible patterns of events, it might be possible to use the simulation to 'fine tune' the three logistic control mechanisms described in this document, to achieve a comfortable balance between the seven performance metrics described in Chapter 5. Only then can it be argued whether the system should become leaner, or if the value of the inventory (in terms of utilisation and customer service level) exceeds its cost.

To summarise:

- Discrete event simulation, based upon existing manufacturing system data, can provide insights into how the system would perform under alternative control strategies
- A novel method of simulation construction was developed, whereby elements within a single model can be controlled as if they were being controlled via 'push' or 'pull' scheduling
- No single system of logistic control can be said to be 'the best' for the use case facility; this depends upon the challenges it faces in terms of demand variability, supply shortages, machine breakdowns, etc.

- It was confirmed that the present-day manufacturing system is operating at a level very close to the maximum possible, given the resources currently available, and the operations it must perform
- The value of the contribution made to the real system by human planners should not be overlooked, but it cannot easily be represented within the model. (Lacking this, the model tends to simply highlight problems; in reality a manager would solve them.)
- Further experimentation is possible, perhaps making a case for the inclusion of additional buffers within the operational flow. (The simulation could be used to determine the ideal size and location of these.)
- The model may also have some utility in a training and dissemination role

8.1. RECOMMENDATIONS FOR FURTHER WORK

Experimentation proved to be a more involved process than had been anticipated, due to the almost infinite number of scenarios that could be explored by making small changes to the level of WIP permitted at in each manufacturing sequence. (The 'push' scheduled system could also be altered, by changing the standard lead times.) Further experiments may be desired, therefore, to reduce lead time and WIP while keeping the service level above a target figure.

With the model in place, further experiments can be conducted, perhaps investigating combinations of the scenarios presented in Chapter 6, working towards a 'most likely' future scenario that can be used with a weighted, multi-criteria system of performance measurement.

Future logistic control within the focused factory might be attempted at a higher resolution than that described in this report; instead of scheduling on a sequence-by-sequence basis, buffers and target WIP values might be established for smaller groups of operations – even singe, critical operations. The establishment of such a close system of logistic control within the simulation would be a mammoth task (the simulation already contains several thousand Arena modules) but with the approach proven, the implementation is merely a task involving repetition. As such, it might be achieved with the data-driven model construction method developed under Task 2.5.1 (described in [Kim et al, 2004]).

9. REFERENCES

Banks, J. (1998) "Handbook of Simulation", John Wiley and Sons, USA

Benders, J. and Riezebos, J. (2002) "Period batch control: classic, not outdated", Production Planning & Control Vol. 13, No. 6, pp497–506

Burbidge, J.L. (1988) "Period Batch Control", Computer-Aided Production Management, Ifip State Of The Art Reports pp71–75

DeGarmo, E.P., Black, J.T. and Kohser, R.A. (1990) "Materials and Processes in Manufacturing", Macmillan, New York

Kim, C-S., Tannock, J.D., Byrne, M., Farr, R., Cao, B. and Er, M. (2004) "State-of-the-Art Review: Techniques to Model the Supply Chain in an Extended Enterprise", VIVACE deliverable D2.5.1_1

Little, J.D.C. (1961) "A Proof of the Queuing Formula $L = \lambda W$", Operations Research, 9, pp383–387

Ohno, T. (1988) "Toyota Production System: Beyond Large-Scale Production", Productivity Press, New York

Goldratt, E.M. (1980) "Optimized Production Timetable: A Revolutionary Program for Industry, Proceedings of the 23rd Annual APICS Conference, Falls Church VA, pp172–176

Goldratt, E.M. and Cox, J. (1987) "The Goal: A Process of Ongoing Improvement" (revised edition), North River Press, New York

Skinner, W. (1974) "The Focused Factory" Harvard Business Review, 52, pp113–121

Womack, J.P., Jones, D.T., and Roos, D. (1991) "The Machine That Changed the World: The Story of Lean Production", Harper Perennial, New York