The Problem with Planning:
The Significance of Theories of Activity
for Operations Management

by
Robert Bruce Johnston
B. Sc.(Hons), Dip. Ed., M. Sc.

A Thesis Submitted to Monash University
for the Degree of Doctor of Philosophy

School of Business Systems

Monash University

March 1999

Table of Contents

Abstract

This thesis presents a new approach to understanding operations management in terms of the underlying theories of activity that inform various improvement initiatives and techniques. In particular, it is argued that the dominant plan-based approaches to operations management are informed by a common-sense, but flawed, conception of activity as the production and implementation plans derived from formal manipulations upon abstract models of the world. A detailed analysis of the conceptual basis and assumptions of this theory of activity demonstrates that the domain of applicability of such approaches is quite limited, and they should fail in characteristic ways already observed within other disciplines that have studied and experimented with the design of agents from this standpoint. The thesis also presents an alternative framework for the analysis and design of operations systems based upon a radical theory of the nature of intentional activity currently being articulated in diverse disciplines, but which has as yet had little impact on management science.

Management of operations in batch repetitive manufacturing is used as a case area to illustrate these ideas. A variety of pre-packaged management techniques have been aggressively marketed to this area. It is argued that these techniques are clustered into two competing theories, or paradigms, of operations management. This first, which may be labelled Computer Aided Production Management, stresses computer based planning and is clearly grounded on a planning theory of activity. The activity theory basis of the second, which is referred to as Lean Production, is in need of articulation. Case studies of companies who have had experience with both paradigms reveal that there are aspects of the management problem in this area for which the dominant plan-based approach gives no account, namely, the situatedness of manufacturing activity, and further, that this problem cannot be made to go away simply by making the planning approach more elaborate. These observations motivate a re-examination operations management in terms of the nature of activity.

The connection between management theories and theories of the nature of activity is made rigorous by noting that, at a certain level of description, all sufficiently complex systems, including management systems, can be viewed as intentional agents. At this same level of analysis, it is natural to formulate operations management as the design and control of intentional systems. When this level of description and what is meant by intentionality are carefully defined, it is possible to borrow fruitfully, from a variety of disciplinary areas, conceptions of the nature and design of intentional activity which have important implications for the design of operations systems and operations management.

With these new theoretical tools in place, the case area, and particularly the main case study, are revisited and reread. It is possible to examine the conceptual basis and likely usefulness of the plan-based approaches with new clarity. Using the new situational / interactional theories of activity, it is also possible to define and apply a new approach to operations management and to argue that the Lean Production paradigm of operations management can be understood as an example of this new approach.

Chapter 1

Introduction

When I took up my first serious computing position in manufacturing in 1982, I was led to the training room and shown a training video explaining the concept of Manufacturing Resource Planning, or MRP II. The company was implementing the IBM MAPICS manufacturing package which supported the MRP II concept. The video explained the idea that if a company has formalised what it wants to make, how its products are made, and how many of the parts involved are in stock, it is a simple information processing task to work out an action plan for the purchase of materials, the fabrication of parts and the assembly of finished items, and that this approach to manufacturing operations management promised immense reductions in the buffer stocks that are required when purchasing and manufacturing activities are not thus coordinated. I remember how compelling the argument seemed, and it must seem so for most people with something of an "engineering mind-set". I was not entirely naive about manufacturing having been a partner in a small manufacturing concern for the previous three years.

I spent the next thirteen years trying to implement this self-evident idea in about a half a dozen mid-size batch repetitive manufacturing companies, first as an employee and then as a consulting analyst / project manager. While most of these companies received substantial benefits from the implementation of the computerised infra-structure systems necessitated by the MRP II concept, all of them had difficulty in achieving the MRP II vision, or keeping it working once implemented. Along with a number of other practitioners and academics at the time (the early 1990s, see (Hopp and Spearman 1996, p173)), I started to wonder if there might not be a flaw in this seemingly self-evidently plausible conception of manufacturing operations management. At the same time a number of new ideas about manufacturing operations, going by the name of Just-In-Time or Lean Production were influential in my client companies. These ideas de-emphasised the role of computers and planning in improving operations and so seemed to be based on a counter-intuition to that of MRP II.

Then in 1992, I visited a client company for the first time since I had worked there five years earlier. During that period they had completed the MRP II implementation but had found that it had not solved their problems of huge inventories and long back-order list. Other initiatives such as quality circles, Just-In-Time, and self-managing production teams had also foundered until, in 1991, as a result of an Australian Best Practices Demonstration Grant, they totally revised the plant layout along Lean Production / Cellular Manufacturing lines. The transformation of the performance of the company's operations was dramatic, with a reduction in through-put time for a typical product from 5 weeks to 3 days, and the total elimination of back-orders. The managers I spoke to there insisted that the physical layout had been the crucial ingredient in the transformation.

I realised that this observation about the pivotal importance of the physical environment of work is a big problem for a management theory like MRP II, which sees the management problem as an essentially informational one. While the orderliness of the setting of work is certainly going to have an effect on productivity, within the information processing approach of MRP II this is seen as an *incidental* effect. On this account of operations management the setting is a passive backdrop against which work takes place. Yet this company's experience was indicating a different relation between operational activity and its setting: that activity is inherently "situated" in a way not easily taken into account within the information processing framework of MRP II, and that this "situatedness" is important to how operations should be managed.

This thesis reports the results of a series of investigations I made to clarify these issues. These investigations were:

1. An extensive review of the literature on the relationship between the various competing theories and techniques of manufacturing operations management.

2. An in-depth interview based case study of the company mentioned above, and two smaller case studies with companies that were similar to, but different in important ways, from the first company.

3. An extensive literature study of the nature of on-going, purposeful activity as conceived in multiple disciplinary areas, particularly those under the heading of cognitive science.

4. The development of a conceptual argument relating different conceptions of the nature of activity to appropriate management systems and interventions.

These studies were conducted over a six year period of part-time study from 1992 to 1997, initially for a Master of Business (by research) degree which was later converted to PhD candidature. The results have been published in three international journal papers, hereafter referred to as papers one, two and three (see appendix 1). This thesis presents a more fully worked out and rigorous presentation of the material in these papers.

I initially investigated the relationship between certain techniques of production management (particularly MRP II and Lean production) but found that the interesting issues were part of a larger problem of the relationship between theories of activity and theories of management, which has wider applicability than production management (as shown in paper two). The idea to which these investigations have led is that theories of operations management are always informed by conceptions of the nature of activity, although often implicitly so. Furthermore, making this connection explicit sheds useful light on the relationship between various management theories and their likely usefulness. My argument is as follows.

1. Different conceptions of the nature and workings of on-going, purposeful activity, which are often tacit or implicit, lead the problem of management at the operations level to be formulated in different ways, or conversely, that management theories are always implicitly informed by theories of activity.

2. For complex systems generally, the problem of explaining the nature of activity arises whenever one adopts a particular descriptive stance that views a system as goal-directed, or intentional. Thus, at a certain level of abstraction, operations systems are similar to all other intentional systems (animals, humans, robots, software agents) and multi-disciplinary sources can be usefully drawn upon to define candidate theories of activity.

3. The connection between activity theories and operations management can be made by noting that, at least at this level of abstraction, management can be viewed as the design and control of an intentional agent, namely, the operations systems. What is known from other disciplines about general approaches to the design of intentional agents can then be applied to define various approaches to operation management.

4. It is useful to explicate this connection between the nature of activity and the problem of management because:

   - by so doing the relationships between various proposed management theories, systems or interventions can be seen;

   - there are definite implications for the design of management systems, and the appropriate types of management interventions to be derived from studying the underlying theory of activity.

   - when exposed, many of the assumptions of the various theories of activity may be found to be untenable, causing the effectiveness of the management theory to be questioned;

5. In particular, there is a widely held folk theory of human activity that conceives it as being mediated by mental models, within the agent's head, of the world outside the agent, and effected by means of plans derived from these models. This theory has been uncritically imported to other disciplinary areas that study intentional systems, including management.

6. Adoption of this theory of activity, usually implicitly, leads to a certain conception of management as the automated production and implementation of plans, which I refer to as the management-as-planning approach.

7. Although this approach to management is ubiquitous and continually being re-invented, whenever it is seriously tested it is found wanting. The source of the problem can be traced to the untenability in many circumstances, of the underlying theory of activity.

8. In many disciplines, particularly those that make up the area of cognitive science (philosophy, psychology, neurophysiology, linguistics, artificial intelligence, robotics, computer science, ethology, sociology, and anthropology), a rather different conception of the nature of activity has recently been proposed that stresses the situatedness of activity, but which has as yet had little impact on operations management theorising.

9. This view of activity leads to a conception of operations management that stresses the role of the situation in which activity takes place, and consequently that the situation of work is a major point of leverage for managers upon practices.

10. Lean Production appears to be an example of this differently informed approach to improving operations. It promotes a raft of situational reforms which support simple decision systems, and down plays the role of planning in execution systems.

The specific aims of this thesis are therefore:

1. To show that it is fruitful to discuss management systems as intentional systems. This allows a level of description to be defined at which the issue of the underlying conception of activity can be defined. This level is common to all intentional systems so it allows for fruitful borrowing from multi-disciplinary sources.

2. To define and describe the main competing theories of activity and make explicit their assumptions and shortcomings.

3. To distinguish between several approaches to operations management in terms of their underlying theories of activity and to discuss their validity by reference to the validity of the theory of activity.

4. To give a framework in which the connection between theories of activity and management theories can be made.

5. To show that the theory of activity adopted by the designer of a management system has definite implications for the resulting design and the systems performance.

Although this argument has wide application, to make this discussion concrete I have chosen as a case area operations within batch repetitive manufacturing enterprises. Not only does this area present an extremely complex operational management problem, but it has been the target over several decades of various all encompassing, heavily hyped theories or techniques for managing operations, which serve my purpose well as illustrations of these ideas. Furthermore, the kind and scale of manufacturing which is the battleground for these management theories, is representative of a large proportion of medium sized manufacturing firms in Western countries, so the outcome is of some economic significance. In 1991, a column writer for Industrial Engineering estimated on then current figures that software supporting the MRP II concept would be a $2,300,000,000 industry in the US by 1996 (Anonymous 1991). Expenditure on even more expensive Enterprise Resources Management software, which may be viewed as a successor to MRP II, should continue this trend. The cost of an MRP II installation is on the order of millions of dollars (LaForge and Sturr 1986) and as we will see later, the success rate is not high. Any research which indicates that such levels of investment may be largely ineffective in solving production management problems is of potential economic significance.

The rest of this thesis is divided into three parts. The present chapter has defined the problem to be investigated, its context and the specific aims of the research. I have done so in a rather idiosyncratic way by relating my own experiences with the plan-based approach to operations management, so my

first task is to show that the problem also exists for other operations management academics and practitioners, in short that there is a genuine problem to be solved. I do that in the first part in three ways. Chapter 2 describes the competing theories and techniques of manufacturing operations management in an historical context which shows the changing attitudes to the problem and to the plan-based approach. That chapter also serves to define terms. Chapter 3 then reviews the academic and trade literature which has attempted to compare the competing theories and techniques. This reveals that there has been a good deal of difficulty in relating and comparing the two main types of techniques and the gradual recognition that actually two completely distinct theories or paradigms for manufacturing operations management have emerged. Chapter 4 then reports the case studies. The major case study is of the company mentioned earlier. It sets out in some detail the changing approaches to operations management in this company over a crucial five year period, the changes in attitudes of the managers to the nature of the problem, and their discovery of the crucial role of the physical environment for several of their initiatives. At the end of this section it is then possible to state with precision the problem to be addressed theoretically in the rest of the thesis.

The rest of the thesis presents my substantive theoretical proposal that operations management systems can be discussed in a new way, as intentional systems necessarily embodying a theory of activity. So Chapters 5 to 7, which make up the second part of the thesis, define my conceptual tools: intentional systems; the activity level of description; and theories of activity. Chapter 5 defines the notion of an intentional system and defines the level of description at which the nature of activity can be discussed. Chapter 6 then introduces and gives a formal definition of the most widely held theory of activity. This precision allows its assumptions to be clearly explicated and their validity assessed. Chapter 7 then presents a radical alternative theory of activity which is being articulated in a number of disciplines under the rubric of cognitive science, mainly in the last decade, but which has received little attention in management theory.

Part 3 of the thesis applies these tools to manufacturing operations management systems. In chapter 8, the proposal to discuss manufacturing operations management systems as intentional agents is discussed and justified. The idea of viewing complex systems as purposeful is not new and has been somewhat controversial. The approach, worked out in chapter 5, of viewing intentionality as the ascription of an outside observer and carefully accounting for levels of description overcomes many of these conceptual difficulties. When coupled with the notion that management can be viewed as the design and control of an intentional system, these ideas lead to a useful framework for making the connection between theories of activity and approaches to operations management. In chapter 9, it is argued that the folk theory of activity, the planning model, leads to a particular conception of the nature of management, of which MRP II is an example. The assumptions of the planning model of activity are examined in this setting and found wanting. Chapter 10, considers the consequences of adopting the radical alternative notion of the nature of activity for management practice. Lean Production appears to have many aspects in common with the resulting formulation of operations management. The experiences of the main case company are re-examined in the light of this alternative theory of activity.

The thesis is primarily theoretical and conceptual. The method is to draw upon a series of case studies of different kinds to make the central theoretical argument set out above. Chapters 2 and 3 describe the area of operations in batch repetitive manufacturing as a literature based case study of various approaches to operations management generally. The idea here is to give a concrete setting for the discussion of significance of theories of activity for operations management. It also shows that there are competing theories of operations management and that there is no wholly adequate explanation for their essential differences. Chapter 4 presents several interview based case studies of actual experiences of companies with the competing theories of operations management to further make the case that a new conceptual tool is needed to understand the relationship of these competing theories. They also point to the phenomenon that evades explanation in plan-based theories, namely the influence of the situation of activity on its outcome, and consequently indicate the direction for new explanations. The case studies of the first part are directed to articulating the problem which is the subject of the theoretical development of the rest of the thesis.

In chapter 5, I use the example of a hypothetical can-collecting robot, as a vehicle to introduce the notion of intentionality as an outsiders account of activity, and to define the level of description of complex systems which I want to use to describe the problem of operations management. Thus, the can-collecting robot is a didactic thought experiment upon which I draw repeatedly in making the argument about intentionality, observers, and theories of activity. Again in chapters 6 and 7, I use two literature based case studies to illustrate the competing theories of activity I describe. The first is the case area of autonomous robot and software agent design, and the second is that of theories of human intentional activity. The robot case study is again didactic in purpose. Using a mechanistic example here focusses attention on the pertinent question of the relationship between knowledge, sensing, acting, planning, and interaction, which is the issue that separates theories of activity, and therefore theories of management, at the level of description I adopt. With this case study I am not seeking to take a position on the design of robots, but rather to define the key issues of theories of activity in a setting where they are most easily appreciated. The actual details of the different approaches have been worked out most rigorously in areas where researchers have tried to build intentional agents and make them work. When I consider alternative theories of activity in chapter 7, the robotics and software agent case studies allow some ideas which could otherwise be dismissed as "humanistic", "psychological", "sociological" or even "mystical", to be seen to be also relevant in systems that are entirely mechanistic and deterministic by nature. My use of theories of human intentionality is slightly different: since I have argued that the planning approach to management is metaphorically derived from a folk psychological theory of human activity, it is important to demonstrate that there are other legitimate conceptions of human activity from which to draw metaphorical impetus. Again, I do not wish to take a position on the nature of *human* intentional action, but rather to use this example to motivate alternative designs for intentional systems in general.

Finally, in the last part of the thesis, I return to both the case area of manufacturing management in general, and the experiences of the main case company in particular, in order to present a rereading of these case studies in the light of the theory I have developed. I demonstrate the usefulness of the activity level of description in understanding the competing theories of operations management. It will then be apparent that all these case studies are about the same thing: the nature, design, and control of intentional systems.

Part 1

Case Area: Operations Management in Batch Repetitive Manufacturing

Chapter 2

Case Area: Manufacturing Operations Management

This chapter introduces the area of operations management that I wish to use as a case example, namely, operations management in batch repetitive manufacturing firms, and particularly production and inventory management. I first give an overview of the management problem in this area in order to define terminology, and then introduce the various competing operations management techniques with a brief historical review of their development and introduction.

2.1 The Problem

Manufacturing provides a good case area within which to discuss the management of on-going activity for two reasons: the activities are usually sufficiently complex to provide difficult organisational and management problems; and there are a number of general approaches to, or techniques for, the organisation of manufacturing operations currently being advocated by enthusiastic practitioners, academics, and consultants, and the defining differences between these, and the relative merits of these are not at all clear.

Manufacturing discrete items (as opposed to processing materials such as oil refining and chemical industries) involves the transformation of purchased materials (say, by a series of machining operations) into component parts and the assembly of components into more complex sub-assemblies and finished items. These fabrication and assembly processes are performed by machines with a restricted range of capabilities and operators with certain skill sets. Manufacturing operation types are generally distinguished by the product range, product complexity, and lifetime of the product being manufactured (Hopp and Spearman 1996, p9). It is normal to distinguish two extreme types of manufacturing. In **flow-shop** manufacturing, a small range of complex products is manufactured continuously usually on a dedicated production line consisting of single purpose machines. The large automotive manufacturers are a good example. At the other extreme is the **job-shop** where a large range of products of limited complexity are made in batches usually using fairly flexible machines. Although a batch may consist of a large number of (identical) units a given product may only be manufactured once. A repetition engineering works is a good example. On the middle ground is **batch-repetitive** manufacturing, where reasonably complex products from a large product range (typically 100s or 1000s of products) are made in batches whose frequency is determined by demand. However, the product designs have a considerable lifetime and many batches of a given product will be made. Batch repetitive manufacturing companies are probably the most common by number. Because of the large product range *and* product complexity, batch repetitive manufacturing presents very difficult management problems and has therefore been the target of many pre-packaged solutions to these problems. This combination of complexity and repetition makes it the ideal case area for this thesis.

Generally, subgroups of the products require specific machines types. Machines usually have to be given relatively fixed positions in the plant and there are two extreme approaches to this. In **process-oriented** or **functional** layouts, machines with similar functions are grouped together. Often these functional groups are arranged in the order in which they are used in the typical product, so that a product moves in a constant direction through the plant during production. At the other extreme is **product-oriented** or **cellular** layouts, where machines involved in a range of similar products are

grouped together. This idea is also known as **group technology** (Burbidge 1990). Individual machines or groups of identical machines are called **work centres**.

Typically, products are designed by **product engineers** and the manufacturing process is specified by **methods engineers**. The manufacturing method is normally specified in two documents or computer files: the **Bill of Material** (BOM) sets out the name and quantity of each part and sub-assembly that make up the finished item; the **Routeing** specifies the operations and machines involved in every fabrication process in the manufacture of the product.

Most companies carry some stock of the various items used in manufacture as a buffer against variation in supply or demand. Stocks of purchased materials are called **raw-material inventory**, stocks of items assembled for convenience or efficiency reasons to partial completion, are called **component inventory**, and stocks of the end-item are called **finished goods inventory**.

The problem of managing operations in a batch repetitive manufacturing environment can now be defined. According to Hopp and Spearman (Hopp and Spearman 1996, p4) "the term operations refers to the application of resources (capital materials, technology and human skill/knowledge) to the production of goods and services". On an on-going basis manufacturing operations management involves putting in place systems or procedures for deciding when to initiate purchase orders, works orders for finished items, works orders for any stocked sub-assemblies and short term changes in available labour (casual and overtime labour). Because of the dependencies between these various day-to-day decisions, dependencies which are expressed in the Bill of Material and Routeing, this can be an extremely complex decision scenario. Management must also make certain long term or policy decisions including the allowable stocks of purchased items, sub-assemblies and end-items, the quantities to be made in each production batch, layout of the plant, training of work force, delegation of control and responsibility, and so-forth.

Traditionally the objective of operations management has been formulated as the minimisation of the *cost* of producing goods. This is almost invariably the objective used for formal optimisation of operations decisions. However, in recent decades (Hayes and Wheelwright 1984) there has been an increasing emphasis on simultaneously increasing *responsiveness* by reducing through-put times, and minimising the *investment* required for a given production level by reducing the quantity and value of work-in-progress and reducing waste through better quality practices.

2.2 Beginnings

In the craft system of manufacture, which was widely practiced in the repetitive fabrication industries prior to this century (Womack, Jones and Roos 1990, p88), all work allocation and scheduling decisions were made by the foreman who was a craftsman of superior skill and experience (Zuboff 1988, p41). Finished items were normally made one at a time to completion by a group of craftsmen under his command. This was not only because, as a legacy of the guild system, the craftsmen guarded their expertise closely, but also because the lack of accuracy in machining parts and the consequent need for extensive fitting, made batch production difficult. Also in a time of unreliable communication, transport and banking, large inventories were often held as a way of storing wealth, so the efficient use of materials was not a major concern. Knowledge of the processes of manufacture was tacit, passed on by personal instruction, and seldom recorded (Zuboff 1988, p40). Operations management in this era was thus tacit, informal and not efficiency-oriented.

In the late nineteenth and early twentieth century, a number of innovations enabled a quantum leap in the scope and complexity of manufacturing, creating the need for much more formal methods manufacturing operations management. Although the idea of making products by the assembly of exchangeable parts has a long history, especially in the armaments industry (Hopp and Spearman 1996, p21), it was only achievable at a practical cost for complex precision items, such as automobiles, following the development of high speed steel and the scientific investigation of metal cutting by F. W.

Taylor in the late nineteenth century. Interchangeabilty of parts helped enable batch production and ultimately mass production of machinery. This made necessary the much greater formality and precision in the codification of manufacturing processes and parts specifications that was pioneered by F. W. Taylor (Wren 1994), and led to the discipline of methods engineering. The codification of manufacturing practices reduced the control of workers and foremen over the pace and organisation of work and gave the responsibility for scheduling the now more complex batch manufacturing to specialised production control departments. Formalised hierarchical management was also necessitated by the increasing size and vertical integration of manufacturing firms by the turn of the century (Hopp and Spearman 1996, p23). This pressure had also led to development of cost accounting within the U.S. railways by Andrew Carnegie which became standard practice in manufacturing companies (Gaines 1905; Eames and Eames 1990, 51) by the first decade of this century.

The scientific study of the decision aspect of operations management began with the development of the Economic Batch Quantity model by Harris (Harris 1913; Harris 1915) which attempted to answer the question: How many parts to make at once? This was later developed into a complete stock control system by Wilson (Wilson 1934). In the 1930s large US manufacturing organisations, such as General Motors and the Du Pont Powder Company, were using demand forecasting and inventory control systems, which were later widely imitated (Hopp and Spearman 1996, p38). The application of mathematical analysis to military logistics problems during World War Two, going under the name of Operations Research, led to the founding in peace time of the academic disciplines of Production and Inventory Management, and Production Control, with texts such as (Plossl and Wight 1967) and professional societies (Newell and Clark 1990) such as the American Production and Inventory Control Society, APICS, and its international off-shoots.

2.3 Early Computerisation

The commercial availability of electronic computers around the mid 1950s gave new hope that the areas of inventory control and production planning, which had become chronic and truly intractable problems for batch repetitive manufacturing, could be at last be solved. The earliest computer applications were aimed at recording inventory of parts and end-items and to using the emerging body of inventory control theory, going under the name Scientific Inventory Control, or Statistical Inventory Control (SIC), to aid the ordering decision for bought and manufactured parts. SIC used simple time series analysis and batch sizing theory to define optimal ordering strategies for independent demand items, primarily by the Order Point / Order Quantity (OPOQ) method. When applied to production inventories, the OPOQ method treats all sub-assembly and raw material demand as essentially independent of the end-item demand, and recommends the ordering of a predetermined order quantity when the stock of any item falls below a predetermined order point. The method generally results in large work-in-progress inventories (Orlicky 1975), but is easily implemented using only serial processing (and, indeed, even without computers). OPOQ and SIC are covered in all the early inventory management and production control texts, e.g., (Plossl and Wight 1967) and there are several good reviews (Aggarwal 1974; Fortuin 1977; Aksoy and Erenguc 1988).

However, when random access magnetic disk storage became available around 1960, the prospects for efficiently handling the complex dependency of demand implied in the Bill of Material of a manufactured product, were greatly enhanced. (The idea of exploding end-item demand through a Bill of Materials to find component demand had already been tackled using punch card computer technology in the 1950s (Orlicky 1975, p37).) When a company has computerised its BOMs, which is usually motivated by product costing, it is normally a short leap to using the BOM to enhance the OPOQ system. While components are ordered by the OPOQ method, BOM explosion is used to check their availability at order release time (Schonberger 1980; Schonberger 1983).

2.4 The Human Factor

After the rediscovery of the social aspects of work by the academic community, through the Hawthorne studies at Western Electric in Chicago carried out by Elton Mayo (Roethlisberger and Dickinson 1939), the Taylorist marginalisation of the human factor in manufacturing began to be replaced, in the new social science approach to work called the "human relations" movement, by a recognition of human individual and social needs in work. Trist (Trist 1980) argues that this attention to humanity was fostered by the full employment conditions that followed World War II, when the "stick" became unavailable. As a result, a change in the characterisation of workers from recalcitrant to a possible creative partner in management enterprise occurred, as expressed in McGregor's "Theory Y" (McGregor 1960).

Extensive action research field studies by the Tavistock institute in Britain (Trist 1982) led to the Socio-technical Systems perspective, which saw the operations system as explicitly having both social and technical components in interaction, both of which needed attention for successful reform. The idea was to optimise the socio-technical whole rather than each part individually. The Soft Systems Methodology for systems design (Checkland 1981) makes a similar break from solely technical optimisation. The Socio-technical Systems approach was not just a theoretical framework but was applied as a practical intervention methodology especially in Britain and Scandinavia, often in response to protracted industrial problems. Assisted by demands for a more democratic workplace in countries with socialist union movements, the concept of semi-autonomous, self-managing teams emerged as a component of this socio-technical optimisation. The notion that systems should be emancipatory as well as efficient has reached formal expression in Critical Systems Theory (Flood and Jackson 1991). While these ideas have remained influential with management theorists, their essentially psychological, sociological and emancipatory focus has led them to be seen by production and inventory management practitioners, particularly in the US, as apart form the main operations systems debate, especially as they did not lead to prepackaged "solutions" for the middle ground of batch repetitive manufacturing.

2.5 Material Requirements Planning

Material Requirements Planning (MRP) was developed in the USA in the 1960s. (Orlicky gives the first implementation as 1961 (Orlicky 1975, p119), Plossl mentions experiments in the late 1950s (Plossl 1980), and Anderson et. al. (Anderson et al. 1982) mention an MRP implementation dating from 1957). MRP takes a Master Production Schedule (MPS) and explodes this through a Bill of Materials augmented with product manufacturing lead times, to come up with time-phased material (raw materials and sub-assemblies) requirements which, when compared to the stock and order situation for these, yields an execution plan for production and purchasing (Orlicky 1975). The expected benefits are reduced work in progress and raw materials inventories and shorter production throughput times, by taking advantage of the dependency of the demand for subassemblies and purchased items upon end-item demand, to achieve a more nearly optimal timing of works order and purchase order release.

In 1972 the American Production and Inventory Control Society, APICS, launched a high profile MRP education and promotional program called the "MRP crusade" (Plossl 1980). At this time MRP was viewed merely as an order launching system (Wight 1981). Oliver Wight records two subsequent stages in the development of Material Requirements Planning. The first came with the realisation that in order for the Master Production Schedule to be realistic and maintained, the MRP system should receive feedback from the other operational systems such as capacity planning, shop floor control and purchasing, and this concept was known as Closed Loop MRP (it was first practised about 1969 (Wight 1981, p50)). A few years later Wight coined the term Manufacturing Resource Planning (MRP II) to denote a Closed Loop MRP system tied to other business functions such as marketing and finance in such a way that the Master Production Schedule becomes a major component of the overall business plan.

Wight (Wight 1981) proposed an ABCD classification of MRP II implementations, which has become widely used. A "D" class implementation is one in which MRP ideas are only influential in the computer department, whereas an "A" class site will be running the business via MRP II with the MPS being a true forecast of the business's intentions. Today, MRP II is generally implemented using proprietary integrated manufacturing software packages, most frequently on mid-range computers or networked PCs.

2.6 Doubts and Threats.

In the 1970s MRP and MRP II were widely promoted and tried in the US and elsewhere, but by 1980 concern about the lack of successful implementation was being expressed in trade and academic journals (Kochhar 1982; Voss 1986, p19). For example, Plossl wrote in 1980

> "Never ... has so much been proclaimed and expected and so little actually delivered." (Plossl 1980,p1).

and

> "Based on simple statistics, no sound management would authorise significant expenditures to develop MRP programs; the success ratio is still far too low." (Plossl 1980,p4).

Following the extraordinary success of Japanese manufacturing in the 1970's, MRP also had to face competition from the Japanese Just-In-Time (JIT) production methods during the 1980s. Also, other computer based production control methodologies, such as Optimised Production Technology (OPT), emerged in the same period. Nevertheless, a recent survey of manufacturers in Ohio USA (Newman and Sridharan 1992), revealed that MRP had achieved a sizeable penetration in the USA and the situation is probably comparable in Australia.

2.7 Optimised Production Technology.

The greatest technical weakness of MRP is that it takes no account of the finiteness of shop floor capacity (Hastings, Marshall and Willis 1982). The use of Rough Cut Capacity Planning (RCCP) (Orlicky 1975) can help insure that the production plan is feasible in an overall capacity sense, but the assumption of MRP that production lead times are independent of detailed work centre loading limits its effectiveness as a shop floor scheduling tool. The production of operational schedules within the constraints of finite capacity is known as Finite Capacity Scheduling (FCS). Attempts to use the new computing power against this computationally explosive problem are as old as MRP (Goldratt 1988) and resulted in some software packages such as the IBM CLASS, but never achieved the popular status of MRP, arguably due to the strong advocacy of MRP by the influential APICS organisation (Newell and Clark 1990).

In 1978 a new proprietary Finite Capacity Scheduling offering called Optimised Production Technology (OPT) was developed by Eli Goldratt (Goldratt 1988). Reports appeared in the literature from about 1983 (Fox 1983; Fox 1984; Lundrigan 1986; Fox 1987; Frizelle 1989; Wheatly 1989), but the product was initially difficult to evaluate independently because of the secrecy surrounding its methods (Jacobs 1983). Early versions of OPT were apparently based on familiar scheduling principles, but around 1982 Goldratt (Goldratt 1988) realised that a detailed schedule and detailed shop floor feedback were only required for the bottle-neck production processes and that non-bottleneck processes can and should be slaved to these. This idea became known as Drum-Buffer-Rope (DBR) production (Jones and Roberts 1990; Scgragenheim and Ronen 1990; Gardiner, Blackstone and Gardiner 1993). The book by Jones and Roberts is a good source on OPT (Jones and Roberts 1990) and the software is described in (Meleton 1986; Fry, Cox and Blackstone 1992). OPT has achieved far less penetration as a computerised production planning methodology than MRP II (Newman and Sridharan 1992), but has been implemented in some quite high profile companies, particularly in

process industries (Meleton 1986; Frizelle 1989; Jones and Roberts 1990), with reports of reduced work in progress inventories (Bylinski 1983), increased sales (Chantland 1993), and faster through-put (Manning 1983). Goldrat also advocated the adoption of new accounting measures of manufacturing performance and went on to develop a theory of operations management based on the recognition of constraints which was set out in the influential management novel, "The Goal" (Goldratt and Fox 1984).

During this time there must have been many systems developed by individuals, mostly for individual companies, that addressed the Finite Capacity Scheduling problem and the weaknesses of MRP (e.g. (Hastings, Marshall and Willis 1982; Hastings and Packham 1985; White 1985; Taal and Wortmann 1997) see also (Buxey 1989)) and a number of PC based Finite Capacity Scheduling packages are now available. But such was the influence of APICS and its British and Australian off-shoots, that leadtime-based MRP II has become virtually synonymous with computerised production management (Newell and Clark 1990; Newell, Swan and Clark 1993).

2.8 The Toyota Production System.

In the early 1980s, the western world started to become aware of a new group of production management techniques generally referred to as Just-In-Time manufacturing (JIT), or the Toyota Production System, which by then were well established in Japan. The Toyota Production System was developed by Taiichi Ohno and others at Toyota in the 1950s and 1960s (Monden 1983; Womack, Jones and Roos 1990). After the international oil crisis of 1973, the contribution that the lean methods at Toyota were making to its strong performance became apparent to other Japanese companies and the Toyota Production system was widely adopted there. By the late 1970s the competition from Japanese Automotive and Electronics industries was so strong that practitioners in the western world were forced to examine the Japanese production methods. The first academic article on the Toyota Production System appeared in 1977 (Sugimori et al. 1977), but the appearance of the book by Schonberger (Schonberger 1982) and the placing of JIT on the agenda of the 1984 APICS annual conference marked the beginnings of dissemination of these ideas in the western world (Newell, Swan and Clark 1993).

There were a number of unique conditions in the Japanese situation that led to the development of the Toyota Production System (Womack, Jones and Roos 1990).

1. The commitment that Toyota had made to lifelong employment of workers led to subsequent acceptance by the unions of multi-skilling.

2. The premium on space in Japanese plants meant that work-in-process inventory was viewed very unfavourably.

3. The dependence of Japan on external sources of raw materials.

4. Lack of access to capital in war ravaged Japan.

5. The local demand for a wide variety of car models from a low production industry.

These factors led to the development of a production system that stressed flexibility, elimination of waste, quality and worker involvement over the standard Western micro-economic concerns for economies of scale (Rice and Yoshikawa 1982).

2.9 Lean Production.

That these techniques constitute a new approach to production management is now clear, as a result of widespread use of the principles outside of Japan, initially through Japanese owned plants using local labour. Subsequently, the establishment of locally owned and staffed firms using the techniques in the US, Britain, Europe and Australia, (Voss and Clutterbuck 1989) made it clear that the ideas stand

independently of their Japanese cultural origins (cf. (Hieko 1989b)). This approach to manufacturing operations management has come to be known as Lean Production in the West, especially in the automotive industry (Womack, Jones and Roos 1990) where it is most influential. There is a vast literature on JIT / Lean Production (for reviews see (Sohal, Keller and Fouad 1989; Golhar and Stamm 1991; Goyal and Deshmukh 1992; Singh and Brar 1992; Ramarapu, Meha and Frolick 1995).

The essence of Lean Production (and hence the name) is the total elimination of waste. The two main notions for waste elimination are (Monden 1983,p2):

Just-In-Time, JIT, production.

> *"Just-In-Time basically means to produce the necessary units in the necessary quantities at the necessary time."* (Monden 1983,p2).

Autonomous Defect Control. This is the idea that the work force is responsible for, and empowered to, prevent any defect part travelling to the next process. In the West these ideas are advocated by the Total Quality Management (TQM) movement.

To support these two goals, a whole range of management, and process engineering methods, some of which were already familiar in the west, must be brought to bare on the production problem simultaneously. The following list has been extracted from various sources and uses Western terminology (Schonberger 1982; Hall 1983; Monden 1983; Hieko 1989a; Shingo 1989; Voss and Clutterbuck 1989; Womack, Jones and Roos 1990; Davy et al. 1992; Cowton and Vail 1994). The most detailed list can be found in (Gilbert 1990).

1. Changes to the physical environment: The production process is to be transformed as much as possible into flow production by the use of a cellular plant layout in which machine layout is more related to products than processes. Group Technology ideas (Burbidge 1990) are used to assign products to cells. Cellular layouts are adopted to simplify loading of production on machines, to allow visibility of the process to workers, and to reduce materials handling.

2. Product / Process changes: Products are designed for manufacturability, which involves promoting cooperation and communication between Product Engineering, Process Engineering, Production and Quality Control, and external parts supplier staff. Processes are standardised and simplified especially by the reduction of set-up times in support of the JIT objective.

3. Work-force skills: Workers have to be skilled in many operations to have the flexibility to meet the demands of Just-In-Time production which forbids the making of parts just to keep one specialised worker occupied. Workers move to different machines or cells according to demand for end-items.

4. Political factors: Organisation of workers into autonomous working groups and the flattening of management structures supports politically the flexibility of the work-force, and provides a structure for quality innovations by the work groups. The various skills such as Engineering, Quality Assurance and Preventative Maintenance, which are functionally separated in traditional plants, are brought together in the groups.

5. Cultural factors: Historical "laissez fair" attitudes to production, maintenance and quality are attacked through ideas such as Total Preventative Maintenance (TPM) and Total Quality Management (TQM) and good housekeeping.

6. Information Systems: On the shop floor, emphasis is on visual control and communication systems such as Kanban (a manual card-based production control system) and Andon (the display of a lighted symbol to indicate a breakdown in production).

7. Dealing with the outside world: Ideally the firm's customers and suppliers are brought into the Just-In-Time chain. On the supply side this means controlling the call-up of purchased parts using Kanban, and putting responsibility for purchase part quality on the supplier. On the demand side this means balancing and smoothing of customer requirements to protect the internal Just-In-Time systems from excessive fluctuation.

2.10 Two Paradigms for Production Management.

Many of the ideas listed above were already part of shop floor practice in the US and Europe. For instance, Burbidge had long been advocating levelled production, group technology, cellular layouts, and teams (Burbidge 1988; Burbidge 1990). The use in manufacturing of semi-autonomus production has a long history in Germany (Warnecke and Huser 1995) and Scandinavia (Schuring 1996). However, it took the success of Japanese manufacturing for these ideas to be viewed as parts of a consistent *theory* of manufacturing. Especially since the appearance of a world wide study of the Automobile Industry by Womack et. al. (Womack, Jones and Roos 1990), the influence of Lean Production can be seen as part of a shift away from the Fordist / Talorist mass-production methods of the first half of this century.

It is now clear that two distinct paradigms have emerged for the improvement of production management (Johnston 1995). The first, Computer Aided Production Management[1] (CAPM) sees work being mediated by a business wide computer based information system and reaches its ultimate expression in the Computer Integrated Manufacturing (CIM) notion (Ptak 1991). Production management techniques that are part of this paradigm include MRP, MRP II, OPT, and other Finite Capacity Scheduling systems. The other, Lean Production, advocates the use of simple visual, reactive control systems, operating in a structured physical, cultural and political work environment (Johnston 1995; Johnston and Brennan 1996). It emphasises the elimination of waste and de-emphasises computerised shop-floor systems. This paradigm is most closely associated in the West with JIT, TQM and Kanban.

This chapter has documented the rise of computerised plan-based approaches to manufacturing operations followed by the challenge from a newer paradigm de-emphasising, or at least redefining the role of planning in management. From the appearance of the first reports on the Japanese JIT methods a large literature has developed that attempts to explicate the relationships between the various production managements techniques introduced in this chapter. This literature is critically reviewed in the next chapter.

---

[1] The term Computer Aided Production Management, which has been used in the British literature, is used rather than Manufacturing Resource Planning, MRP II, since the latter is closely associated with Material Requirements Planning, MRP, which is only one method of implementing the ideas referred to. Similarly, the term Lean Production is used rather than the Toyota Production System in order to de-emphasise its Japanese origins and also to avoid the narrower focus of Just-In-Time. Thus MRP II is viewed as a particular instantiation of the CAPM approach and the Toyota Productions Systems is a particular instantiation, for a flow shop environment, of the Lean Production approach.

Chapter 3

Comparative Studies of the Production Management Techniques

In this chapter, the literature which attempts to compare the main advanced manufacturing operations management techniques described in the previous chapter is reviewed in terms of their underlying principles, applicability, and performance. In terms of my overall argument there are two objectives for this literature review. The first is to show that, especially in the academic community, there has been a great deal of difficulty in adequately dealing comparatively with the techniques, especially since JIT arrived in the West. I diagnose this problem as arising from the difficulty of evaluating the new Lean Production paradigm from the standpoint of the established one. I show that the recognition of the paradigmatic nature of the shift in thinking is quite recent and there is as yet no consensus on what characterises the difference between the paradigms. The second objective is to demonstrate from published accounts that, while the plan-based approaches have had considerable influence on business, there is a good deal of evidence that they have not delivered the expected benefits. Thus, in this chapter I use the literature to motivate the last two parts of the thesis by suggesting that the discrepancy between the evaluation and performance of the plan-based approach indicates that its underlying basis needs to be explicated and contrasted with the that of the emerging paradigm and, judged by the failure of the academic literature to adequately characterise these differences, this has not as yet been done. I will add to this motivation with the detailed case studies in the next chapter.

3.1 "Which is Best?": Competing Systems.

Confusion about the status of the techniques led to a rash of articles in the trade and academic press in the early 1980s and a number of threads to the subsequent literature can be discerned. I will argue that the debate over these techniques can be characterised as the asking of a series of ill-posed questions resulting from the initial failure to recognise the paradigmatic nature of the changes that were taking place. The earliest approach to comparing the techniques was to ask the question "which is best?" and attempt to answer it by determining which technique optimised some variable, usually related to cost, which was assumed to be the most appropriate measure of goodness. Even a quite recent paper (Newman and Sridharan 1992) posed the question: *"manufacturing planning and control: is there one definitive answer?"*. In this literature JIT is often simply identified with the Kanban control system.

The earliest papers (Fortuin 1977; Fortuin 1981; Ettienne 1983b; Ettienne 1983a) compared MRP and OPOQ in terms of cost of inventory holdings required for a given customer service level. Using highly simplified single product analytical models, these papers suggested that in all but very special cases OPOQ leads to larger inventories than MRP. A more recent single product simulation study (Jacobs and Whybark 1992) makes the same claim. However, the exact circumstances under which MRP is superior to OPOQ have proved difficult to determine (Miller, Berry and Lai 1976; Ritzman and Krajewski 1983; Bregman 1994). These papers can be viewed as the rigorous working through of the intuitive ideas first put forward by Orlicky (Orlicky 1975), in his ground breaking text, to justify the superiority of MRP over OPOQ, but even this most basic comparison is still an open question.

After 1982 the comparisons included JIT, and OPT was often included after 1985. Rice and Yoshikawa (Rice and Yoshikawa 1982) compared OPOQ, Kanban and MRP for repetitive manufacturing. They

concluded that while Kanban and lot-for-lot MRP[2] have the same JIT aims, traditional MRP is generally compromised in this respect, like OPOQ, by introducing larger lots of components in order to offset set-up costs. Schonberger (Schonberger 1983) compared the systems in terms of the amount of inventory carry they required, from OPOQ through MRP to Kanban with some intermediates described. He saw this as a kind of one dimensional scale of merit. Aggarwal and Aggarwal (Aggarwal 1985; Aggarwal and Aggarwal 1985), described and compared MRP, Kanban and OPT in detail, citing success cases for each. There was some mention of published misgivings for each system, but on the whole the reports were descriptive rather than critical. The treatment of OPT was based largely on the brief, early report of Jocobs (Jacobs 1983). There was no attempt to relate different production environments to the systems. Plenert and Best (Plenert and Best 1986) also compared MRP, JIT and OPT qualitatively. However, they did not consider the production environment and did not cite any actual experiences with the three systems. Gelders and Wassenhove (Gelders and Van Wassenhove 1985), compared MRP, JIT and OPT as to how well they deal with capacity constraints. They concluded that OPT combines aspects of JIT and MRP and dealt with capacity constraints better. Similar territory is covered by Lambrecht and Decaluwe (Lambrecht and Decaluwe 1988). Slack and Correa (Slack and Correa 1992) compared JIT and MRP in terms of flexibility and concluded that a simple answer was not possible because the approaches addressed different aspects of flexibility. Miltenburg (Miltenburg 1993) compared JIT and MRP for their usefulness in reducing wasteful processes in manufacturing.

Walbank (Walbank 1988), attempted to describe OPOQ, MRP, Kanban and OPT in terms of control system theory. While this seems to be a promising approach, the models became excessively formal, and the discussion had little connection with them. Boggs (Boggs 1988), compared MRP, JIT, OPT, within the framework of general systems theory, in particular in terms of the system "guarantor" in each case. In Churchman's (Churchman 1976) system theory, the guarantor is the sub-system which determines the truth of propositions generated by the system. Boggs used this idea to determine, for each of the production systems, which management level (strategic, tactical or operational) is responsible for insuring that the system is truly serving the goals of the organisation.

A number of authors have tried to compare the performance of OPOQ, MRP, OPT and Kanban systems by computer simulation methods. Sarker and Fitzsimmons (Sarker and Fitzsimmons 1989), used a very idealised linear production line model to compare "pull" (Kanban) systems to conventional "push" (MRP) systems for flow line production. As a work initiation discipline, "push" means to take action in anticipation of need (implying planning) while "pull" means to take action on request (or in response to the current circumstances (DeToni, Caputo and Vinelli 1988)). They found that the pull systems were much more vulnerable to operation time variations and machine breakdowns, as would be expected from the strong inter-stage coupling the Kanbans introduce. This work was later strongly and convincingly criticised by Baker et. al. (Baker, Powell and Pyke 1990). Lee (Lee 1989), also compared the performance of idealised pull and push systems, but with a model somewhat less biased towards flow line structure. He found that the pull system out-performed the push system on all measures, but since the randomly chosen set-up and run times had means of 8 and 72 minutes respectively this is not surprising. Siha and Davis (Siha and David 1990) have also investigate the performance of push and pull for topologies more complex than flow-line.

Motivated by recent claims that reactive systems such as Kanban handle demand uncertainty better than plan-based systems, Jacobs and Whybark (Jacobs and Whybark 1992) used a randomly generated and lumpy forecast demand and compared the performance of MRP with OPOQ when random deviations of actual orders from forecast occurred. The performance measure was inventory holding

---

[2] In standard MRP component and material batch sizes are chosen for efficiency independently of the end item batch size. In lot-for-lot MRP all materials are ordered in quantities only sufficient for a single batch of the end item. This approach sacrifices efficiency for reduced complexity and stock holdings.

cost at constant customer service rate. They found that MRP outperformed OPOQ unless the actual and forecast orders differed in a most extreme and contrived way. Ramsay et. al. (Ramsay, Brown and Tabibzedeh 1990), compared MRP, Kanban and OPT using a simulation of the production of two products with common components and a constraining resource. The three systems were implemented as different rules for authorising production at each machine. Because the simulation had no prioritising rule for a machine, it was incapable of simulating sustained production of even one of the finished products, limiting the value of the conclusions. An interesting simulation by Rees et. al. (Rees, Huang and Taylor III 1989), focuses on the relative performance of Kanban and lot-for-lot MRP, in what the author terms "an ill-structured production environment", that is one more like the job shop than the flow line, having multi-purpose machines, and lacking the physical layout and group technology improvements of the Lean Production philosophy. The simulation had two products with both assembly and process operations, and complex routeings between a small number of machines. They found that with Kanban implemented, considerable savings were gained by reducing lead, cycle, and set-up times. However, even greater savings were found when the same reductions were made under the lot-for-lot MRP system, which handled lumpy demand better.

3.2 "Horses for Courses?": Contingency Approaches

The next discernible thread supposes that the various systems should be chosen to suite the particular manufacturing environment of the firm. In the absence of empirical data on actual successes and failures in implementations of the various systems, one obvious approach is to relate an appropriate classification of environments to an appropriate classification of systems. A few authors have attempted this (Van Der Linden and Grunwald 1980; Schonberger 1983; DeToni, Caputo and Vinelli 1988; Grunwald, Striekwold and Weeda 1989). The idea is to classify the production environment along a number of axes (usually related to market conditions, product, and process) and then to determine from first, principles which production control systems are suitable for a given environment. This approach is reminiscent of contingency approaches to management theory (Hofstede 1981).

Typical aspects of the market environment that are thought relevant are: demand predictability, demand variability, and demand lumpiness (DeToni, Caputo and Vinelli 1988; Grunwald, Striekwold and Weeda 1989; Striekwold 1990). Typical product / process dimensions include: production lead time compared to customer order lead time, complexity of the bill of material, the distribution of value amongst the components, frequency of manufacturing various products, and repetitive vs. non-repetitive manufacture (DeToni, Caputo and Vinelli 1988; Grunwald, Striekwold and Weeda 1989), however there are several variations on how these are defined (Ettienne 1983a; Schmitt, Klastorin and Shtub 1985).

The simplest classification of the techniques is into "push" and "pull". Despite the simplicity of this idea, it has been employed inconsistently in the literature. All authors agree that Kanban is a pull system, but MRP is variously classified as pull (Rice and Yoshikawa 1982; DeToni, Caputo and Vinelli 1988), or as push (Rees, Huang and Taylor III 1989). The difficulty appears to be that MRP involves several decisions to which the push / pull classification might be applied (Pyke and Cohen 1990). Another theme is to classify the techniques by "leanness" (the amount of material and component stock they require for operation) (Schonberger 1983). More detailed classification schemes have been proposed by several authors (Grunwald, Striekwold and Weeda 1989; Striekwold 1990).

The next step in this kind of research is to try to find "dominance regions" (Grunwald, Striekwold and Weeda 1989) for each technique in the space of environmental possibilities by relating the capabilities of the techniques to the environmental requirements (see figure 3.1 for an example). Managers will then presumably identify their own market / product / process environment in this space of possibilities and choose a production management technique accordingly.
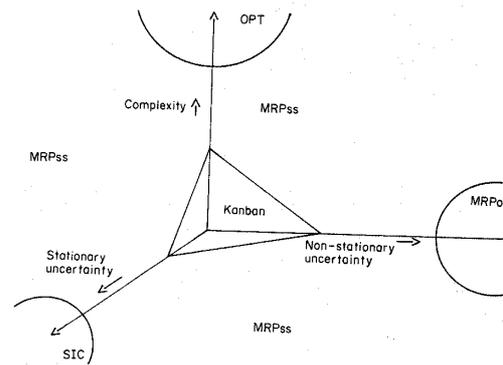
Figure 3.1: Dominance Regions of Production Management Techniques as a Function of Environmental Factors According to Grunwald (Grunwald, Striekwold and Weeda 1989, p291)

Another approach to evaluating the fit of techniques to environments is to use computer simulation. The simulations described in section 3.1 attempted to compare the performance of the various techniques in isolation. There have been a number of elaborate attempts to apply simulation within the contingency approach. Krajewski et. al. (Krajewski et al. 1987) used a complex parameterised simulation that was capable of being tailored to represent a wide range of manufacturing environments. The high and low values for the parameters were set by a panel of production managers, and the model was shown to be capable of simulating an actual plant. OPOQ, MRP and Kanban production control systems could be modelled. For MRP and Kanban, the parameters were varied systematically to determine which aspects of the manufacturing environment need to be shaped to maximise the efficiency of each system. These two systems could not be directly compared however, because different parameter values needed to be chosen for Kanban to make the environment look anything like a potentially JIT one. Another very detailed simulation is presented by Striekwold (Striekwold 1990). He makes the interesting point that if the comparisons are to be valid, the free parameters in the systems (e.g. safety stock levels) must be adjusted by experiment to optimal levels for each system before the comparison is made. However, because of the complexity of the model, he was only able to investigate one aspect of the production environment in detail, namely, convergent versus divergent product structure (i.e. process industry versus assembly). The basic conclusion was that MRP generally out performs OPOQ (JIT not included) but this depends strongly on the production environment.

3.3 "How do I Mix and Match?": Accommodations.

In the more recent literature a new perspective emerges that JIT, MRP, and OPT are not actually incompatible but can complement each other. This is closely related to the recognition that there are several production and inventory planning problems with different timescales and degrees of aggregation (Monniot et al. 1987; DeToni, Caputo and Vinelli 1988; Buxey 1989; Karmarker 1989; Bonney and Head 1993). Material requirements need to be controlled on the long term (strategic capacity and supply plans), the medium term (purchase and production planning) and the short term (order release and tracking), which may require different techniques. However, in much of this literature JIT is still being rather narrowly defined as Kanban.

This emerging opinion (Belt 1987; DeToni, Caputo and Vinelli 1988; Goodrich 1989; Karmarker 1989; Lee 1993) sees JIT / Kanban as useful at the operation level, but that computer based MRP type systems will be required for medium and long term planning and particularly for dealing with suppliers in environments where JIT purchasing is impractical. This rhetoric is now being incorporated into MRP package sales pitches (Webster 1991). The notion that MRP and JIT are compatible has been encouraged by APICS since 1984 (Newell, Swan and Clark 1993). Others include elements of OPT also in this synergy (Luscombe 1991; Maes and Wassenhove 1991; Ptak 1991). A number of papers deal with the mechanics of transforming MRP based operations to include Kanbans (Discenza and McFadden 1988; Ding and Yuen 1991; Flapper, Miltenburg and Wijngaard 1991; Miltenburg and

Wijngaard 1991). The approach is usually to remove corresponding levels from the parent bill of material as various sub-assemblies are put over to Kanban control.

3.4 "Which Philosophy to Adopt?": Manufacturing Paradigms

Most recently there has been a recognition that there are two distinct theories or paradigms represented by the techniques, along the lines I stated in the last chapter (Brodner 1990; Womack, Jones and Roos 1990; Alfred 1993; Cowton and Vail 1994; O'Neill and Sackett 1994; Buzacott 1995; Johnston 1995). This was clearly expressed in some of the early Japanese writings on the subject (Sugimori et al. 1977; Rice and Yoshikawa 1982; Monden 1983) if not stated in these terms, and was partly recognised in the western literature by saying JIT was a different "philosophy" of manufacturing. The important step in recognising this paradigm shift is to acknowledge that JIT or Lean Production is more than simply the Kanban "pull" control system, which will only work if the whole raft of interlocking environmental reforms are implemented. It then becomes apparent that Lean Production has a completely different analysis of how effective operations management can be achieved, to that of the computer based systems which rely on simulation of the future. The question then arises as to what actually characterises the differences in the paradigms. A number of suggestion have been put forward in the literature:

1. A number of authors see the paradigm shift as relating to the use of a different set of measures of manufacturing performance in opposition to the traditional micro-economic efficiency measure. These include: simplicity (Schonberger 1982), increased quality (Rice and Yoshikawa 1982), elimination of waste (Rice and Yoshikawa 1982; Cowton and Vail 1994), and cycle-time reduction (Stalk and Hout 1990). Cowton and Vail (Cowton and Vail 1994) have described the notion of Just-In-Time as an image of a "Weberian ideal state" which drives the whole Lean Production paradigm, and show how approaching this state through waste elimination leads in a logical way to the Lean Production techniques. The equivalent driving notion of the CAPM approach would be efficiency optimisation.

2. Other authors see the main difference in the degree of emphasis which is given to human skill and problem solving. Brodner (Brodner 1990) has called CAPM the "techo-centric" approach, and JIT the "anthropo-centric" approach to emphasise its connection to the support of human skill. Buzacott (Buzacott 1995) distinguished between "integrated manufacturing" which is driven by monolithic integrated information systems, and "cooperative manufacturing" which emphasises flexibility through human problem solving within teams. In a similar vein, Nohria and Berkley (Nohria and Berkley 1994) use the sociological notions of "design" and "action" to distinguish the two paradigms.

3. Yet another theme is to distinguish the paradigms in terms of the organisational structures they advocate: the old paradigm is associated with a monolithic and hierarchical structure, while the new paradigm is associated with a structure that is modular, or "fractal" (Warnecke and Huser 1993; Gardiner 1996), and "extended" (O'Neill and Sackett 1994) to include inter-organistional links, especially with customers and suppliers. Jenner (Jenner 1998) sees Lean Production as embodying the structural principles of self-organising dissipative dynamic systems far from equilibrium, as set forth by Prigogine (Prigogine 1980) and other chaos theorists. In this view, the older paradigm is informed by a static equilibrium model of organisation.

4. In papers one and two (Johnston 1995; Johnston and Brennan 1996), I advocated the view that CAPM and Lean Production can be distinguished by the underlying theories of activity they employ. Far from being at odds with other characterisations of the paradigm shift set out above, this proposal is consistent with all of them and, in fact, integrates them into a unified, coherent framework. Part three of this thesis argues this in detail.

3.5 Critique of the Comparison Literature.

The early "which is best" literature took an extremely simplistic approach to evaluating the competing techniques. We are acclimatised to the "consultant mentality" which formulates change as the implementation of pre-packaged systems, usually with three letter acronyms, which *they* sell. MRP was sold in that way:

> "Nowhere on the Proven Path does one see things like document the current system or design the new system ... the system is already designed; its called MRP II", (Wallace 1985, p33).

So it is not surprising that JIT, and specifically Kanban, should have been thought of as yet another production and inventory control package to be put up against the others. This is despite all the early Japanese literature (e.g. (Monden 1983, p4)) pointing out that Kanban and JIT were just part of a broader change in the way production is to be managed, and that the benefits of one portion of this change could not be achieved without adopting the whole approach. Thus the issue in these Japanese writings was between Fordism and Toyotaism, rather than MRP versus Kanban.

Although the "horses for course" approach tries to take account of the contextuality of the choice between the techniques, it too presents a simplistic view. While MRP II sits naturally upon the functionally organised, Taylorised firm of the mass production era, the fundamental idea of JIT / Lean Production is to *transform* this traditional production environment. One of the main ideas is that batch repetitive manufacturers should conceive themselves less in the job-shop mold and move in the direction of flow-shop manufacturing in order to improve throughput and responsiveness, and reduce buffer inventory. It also advocates changing product designs and manufacturing processes toward this aim. Therefore, dividing the techniques according to the environment / product / process classification of a firm's *current* operations is of limited usefulness.

Not only have the Lean Production advocates suggested that the production environment be changed but also that the measures of good manufacturing performance also be revised, favouring responsiveness, flexibility, and quality, over cost or efficiency. However, investigators within the operations research and management science disciplines have, by and large, uncritically continued to use cost minimisation as the principal means of evaluating any proposed management initiative. Therefore, it is not clear to what extent evaluations made on this basis will be a good indicator of business performance in a complex, competitive, and changing environment (Johnston 1996a; Johnston 1997b; Betts and Johnston 1998).

Another problem of the operations research approach, which applies equally to analytical and simulation work, is that many aspects of the Lean Production approach are difficult to effectively model. Models tend to adopt a reductionist view of the management problem, whereas Lean Production techniques promise improvement through a holistic approach (Hopp and Spearman 1996, p17). The result is that many of the simulations have modelled the Kanban system as if this adequately models the whole Lean Production system. Indeed, at this level of reduction OPOQ and Kanban are essentially the same: when OPOQ and Kanban were compared using control systems theory (Walbank 1988), no essential difference was found. The essence of the whole set of human-centred management changes that are required to support a working Kanban driven operation, compared with a loosely coupled operation using OPOQ, is not captured by this comparison. How would one take into account in a simulation the effects of the increased visibility to the operator of the whole process in a "U" shaped cell? This is not contained in the objective geometry of the cell but in the human experience of it. By contrast, the CAPM techniques which take management to be essentially an information processing task, do not present this problem for modelling.

One might feel that the "mix and match" approach is a satisfying conclusion to the techniques debate, but it is not. Patching up one manufacturing paradigm with isolated elements from the other might be attractive to academics or practitioners committed to one paradigm or the other, but, to the extent that this approach obscures the recognition that there are two extreme paradigmatic views of management here, it is a bad thing. We can learn much more of permanent value by carefully analysing the

fundamental positions about the nature of management that the *extreme* paradigms take, and by then learning to apply these lessons in an individualistic way.

The conclusion that should be drawn from this review of the practitioner and management science literature, is that great difficulty has been experienced by authors from these perspectives in defining the characteristics, applicability, and potential of the Lean Production group of techniques. Although a number of authors have recently suggested that CAPM and Lean Production are different paradigms or philosophies of manufacturing it has proved difficult to give a precise characterisation of how these philosophies differ. This has made it extremely difficult to adequately evaluate these technique or even to agree on the measures by which they should be evaluated. The proposal that I will work through in the remainder of this thesis is that what is needed is a different theoretical unit of analysis namely, "activity", and that these competing paradigms are distinguished by their attitude to the nature of activity.

3.6 Implementation Status of CAPM.

Comparative surveys indicate that MRP is the most widely implemented production management techniques in the US (Newman and Sridharan 1992), Britain (Little and Johnson 1990; Johnson and LaBarre 1991; Clode 1993), Italy (Bartezzaghi, Turco and Spina 1992), Finland (Matsuura, Kurosu and Lehtimarki 1995), Singapore (Sum and Yang 1993; Ang, Sum and Yang 1994), and Australia (Sohal and Chan 1991; Lowe and Sim 1993). JIT is the next most implemented technique in these surveys: OPT and other Finite Capacity Planning techniques are the least influential. Caution is needed in interpreting comparative implementation figures because elements of JIT are often implemented as an execution systems within and MRP II framework and often the degree of implementation is not considered in these studies. Similarly, in the supply chain of Japanese owned automotive companies where commitment to Lean Production is strongest, MRP is often used to explode customer schedules to produce vendor schedules, for informational purposes, without a commitment to the MRP II concept (Prajogo and Johnston 1998). However, it is clear from this literature that the CAPM planning-based approach to operations management has been extremely influential in western manufacturing countries.

A large industry has been built up around the MRP II idea, in terms of computer sales, package software providers, and consulting. According to one commentator (Anonymous 1991):

> "In 1989, MRP software accounted for almost one-third of the total market for computer services in the United States. At the current annual growth rate of 9.5%, this market is expected to reach $2.3 billion by 1996"

However, it is clear that the expectation of the "MRP Crusade" that computerised production management systems would transform the competitive position of the US, particularly against the Japanese, has not been realised (Hopp and Spearman 1996, p173). McLaughlin (McLaughlin, Vastag and Whybark 1994) found little reduction in inventory with MRP use. Groves (Groves 1990) points out that between 1948 and 1987 overall inventory turns in the US hardly changed with an average of about 7 whereas in 1981 the Japanese were averaging 50 with some companies achieving 200 and higher. The turn rate at Toyota in 1980 was 87 (Monden 1983, pvi). This finding has been confirmed by a detailed study of US census data (Hirsch 1996), which shows the only sustained reduction in material and work in process inventory for manufacturing in the period 1973 - 1993 was in the automotive industry after 1982 and is presumably the result of the influence of JIT ideas.

It is difficult from the existing literature to determine exactly how successful CAPM has been. Case studies of MRP and MRP II successes abound but failures are seldom considered material for publication. Often surveys sought highly subjective appraisals of the impact of MRP and they have often been performed on populations, such as members of APICS (Anderson et al. 1982), that would be skewed towards MRP usage, and completed by individuals with a vested interest in the MRP implementation. Even with more objective measures of performance, such as customer lead times or inventory levels, it is difficult to say how much improvement is due simply to computerisation *per se*,

and how much is generated by the adoption of the CAPM approach. In fact, one of the contributions of MRP II is that it has served as a pre-packaged strategy for computerisation of manufacturing firms in an era when system developments were often very *ad hoc*.

MRP II is only completely implemented when "class A" status is achieved. If the number of users reaching "class A" status is any indication, the MRP II concept is extremely difficult to implement. Up to 1993, over 2000 MRP II manufacturing packages had been sold in Australia (Homer 1993) but only about 10 companies claimed Class A status (according to the Victorian branch of APICS). A survey (Cerveny and Scott 1989) of manufacturing firms in the eastern US (N=433) found 60% claiming to be MRP users but only 9% of these classifying themselves as Class A. These figures are almost identical to those found in an earlier US survey (N = 679) of upper Midwest and Eastern US firms (Anderson et al. 1982). La Forge (LaForge and Sturr 1986) using a similar instrument (N = 107), found 25% were Class A. Another recent US study (Duchessi et al. 1988; Duchessi, Schaninger and Hobbs 1989) found 35% of their survey (N=247) of MRP users rated themselves higher than 8 on a 10 point Lickert scale for the success of their implementation. A recent survey (N = 128) of MRP practices in Singapore (Sum and Yang 1993) found that while top management support for MRP was higher than indicated by the US surveys, few MRP sites had attained Class A. A British interview based survey (Monniot et al. 1987) (N= 38) found that, while 66% of MRP users claimed to have a Master Production Schedule, it did not follow that it was negotiated between sales and operations managers in the way that "class A status" would require (see also (Burcher 1991; Clode 1993)).

It appears that contrary to the statement by Wight in 1981 to explain the results of his own survey (Wight 1981, p69), that low "class A" attainment was due to "learning curve" effects, low success rates in implementing the high level components of MRP II such Master Production Scheduling and Sales and Operations Planning have persisted to be a problem. On the other hand, there is good statistical evidence (Duchessi et al. 1988) that attainment of "class A" status is strongly correlated to receiving benefits from MRP. In addition, several studies (Anderson et al. 1982; LaForge and Sturr 1986; Lowe and Sim 1993) have shown that MRP II implementation is costly with costs on the order of millions of dollars in present day terms, and with a wide spread in reported costs. Project times are also long, on the order of several years (Hamid, Agus and Hassan 1991). The situation can be reasonably summarised by saying that while the MRP II idea has been widely disseminated and can yield substantial benefits, it has proved notoriously difficult and costly to implement fully.

Given the extent of the effort and money that has been invested in the CAPM plan-based approach, the apparent failure of the approach to be successfully implemented, and therefore, its failure to deliver the promised improvements over earlier loosely coordinated production managements methods, motivates a new look at the conceptual basis of the approach. This is undertaken in the last part of the thesis after suitable theoretical tool have been developed. This motivation is further developed in the next chapter, where three case studies of the actual experiences of companies with the two paradigms are presented. It will be demonstrated that there are aspects of the nature of batch repetitive manufacture for which the CAPM approach has no account, namely the situatedness of manufacturing activity, and further, that this problem is conceptual rather than technical.

Chapter 4

Case Studies

In this chapter, I present three case studies of batch repetitive manufacturers who have had experience with both paradigms of production management. The first is an in-depth case study with the company mentioned in the introduction, Bendix Mintex of Ballarat, Australia. This case study describes the events that occurred in the crucial period of change between 1988 and 1992, during which the company changed its commitment to the CAPM approach to a commitment to Lean Production. It appeared in truncated form in paper one (Johnston 1995). The other two cases, with Brake and Clutch Industries Australia, and Kodak Australasia, are less detailed and were undertaken mainly to triangulate the findings at Bendix Mintex. The experiences of these two companies were similar to those of Bendix Mintex, but also differed in interesting ways.

4.1 Methodological Considerations

4.1.1 Case Study Rationale

With any research method that attempts to elicit information about industry practices and concerns there is a trade-off between *access* and *generality* (Hakim 1987; Galliers 1992; Shanks, Rouse and Arnott 1993). To make generalisations about industry practices requires the application of standardised research instruments to a large sample of companies using surveys or highly structured case studies. While this approach has some advantages such as, uniformity, repeatability, and the feasibility of rigorous analysis (Hakim 1987, p47), it effectively limits the depth of enquiry that can be made because the questions that can be asked in this way have to be rather context independent and meaningful to all participants (Hakim 1987, p49). Only relatively objective, context-free data such as usage statistics, inventory levels and so forth can be validly collected in this way, although surveys are routinely applied to opinions as well. Little insight can be gained into the causes behind the phenomena under study (Galliers 1992, p154). Consequently, such research methods tend to be more useful in the theory confirmation stage of practitioner research projects, when the issues involved are well understood and objective data that would confirm a theory has been identified.

At the other end of this scale are deep access methods such as participatory observation (action research (Gummeson 1988; Galliers 1992)) and in-depth interview-based case studies (Yin 1989; Yin 1993). These methods aim at obtaining an in-depth knowledge of practices within *particular* organisations in their full individuality and contextuality (Yin 1989, p23). Unstructured or semi-structured interviews with practitioners can uncover information which was not anticipated at the design stage, interesting leads can be followed, and one can actively ensure that the questions are meaningful to the participant and within the organisational context (Yin 1989, p14). It is possible to gain information on the probable causes of the phenomena discussed with participants. This method is limited to knowledge that is already articulated in a firm and has the pitfall that the researcher may be receiving the "company line" which is often highly rehearsed in firms with a strong culture. Participatory research can overcome this observational limitation and can allow access additionally to the "tacit" or as-yet-unarticulated knowledge of the organisation (Gummeson 1988). This comes at the price of the study being practically limited to a single firm. These high access research methods are most appropriate in the theory building stage of research where deep access to the practices and

understandings of a small number of firms, while effectively prohibiting general statements, can allow the uncovering of important research question for further investigation.

It was in this theory building spirit that the case studies reported here were undertaken, quite early in the research project reported here. The major case study was quite extensive. Senior and middle managers from most functional areas of the company were interviewed. Project notes and official company publications concerning the company's initiatives were made available and examined. I had freedom to observe the operations on a number of occasions through plant tours. In addition, there was a participatory component because I had worked in a staff capacity as a systems analyst and project leader on the implementation of many of the foundation systems prior the period covered in the study. This allowed a detailed account of events and changes to be constructed over a five year period of great change in the company. The study cannot however be called truly longitudinal, since the practitioners' understanding of events was only sampled in detail in 1992 with a small follow-up in 1994. The depth of access to the real events and to the true opinions of the participants was also enhanced in this case, because most of the participants were known to me and could place a degree of trust in me, due to my knowledge of the culture and politics of the firm.

As a guard against the possibility that the experiences being observed at Bendix Mintex were idiosyncratic or contrived, I conducted two less in-depth case studies on companies that had similar characteristics to Bendix, in that they were both considered by the academic and business community at the time to have excellence in systems and management, and also they both had had experience with both Computer Aided Production Management techniques and Lean Production. They also differed in interesting ways from Bendix Mintex. Brake and Clutch Industries Australia was interesting because, when MRP failed to deliver performance improvement, they initially turned to a more complex computer based scheduling model, OPT, but later rejected that direction in favour of Lean Production. By contrast, Kodak (Australasia) had taken MRP II implementation to "Class A" certification and then started implementing Kanban and Just-In-Time ideas. They remain committed to the high level aspects of MRP II and in fact see no contradiction in using JIT with MRP II. While access to real company experiences was undoubtedly less in these two studies they served to confirm, or "triangulate", the messages extracted from the main case, because a number of similar themes were articulated by the interviewees there also.

Apart from two early interviews at Bendix Mintex, all interviews were recorded on audio tape and extensively reviewed afterwards. Participants were given confidentiality assurances. Wherever possible in the case descriptions that follow, opinions expressed by participants (as opposed to factual historical statements) are supported by direct quotes from the tapes. Questions were prepared before the interviews but were not common to all interviews, and conversational leads were freely followed. A write-up of the case study at Bendix Mintex was reviewed for accuracy by two participating senior managers before publication as paper one (Johnston 1995), and they confirmed it was a factual account of the events.

4.1.2 Case Study Materials

Bendix Mintex.

1. Conversation and site tour with R. Huchison (Manager Information Systems) 14/8/92, not recorded.

2. Semi-structured interview with Brian Walker (Systems Analyst on the MRP/MPS project) and Dorothy Shehan (Computer Operations Supervisor) 7/9/92, notes taken by hand.

3. Semi-structured recorded interviews each lasting about one hour with the following people:

   David Doyle (Production Manager) 9/9/92.

   Dale Guthridge (Demand Forecaster) 9/9/92.

Vera Pollard (Master Scheduler) 9/9/92.

Greg Simpson (Materials Manager) 14/9/92.

Ross Stoval (Director of Human Resources) 14/9/92.

David Doyle (Production Manager) 14/9/92.

Bob Jackman (Chief Executive Officer) 24/9/92.

Ross Stoval (Director of Human Resources) 24/9/92.

4. Project notes, memos, and minutes of the MRP II implementation project group.

5. Notes and overhead displays from a talk prepared by Ross Stoval, "Workplace Reform at Bendix Mintex", delivered to the Metal Trade Industries Association and to the Confederation of Australian Industry.

6. The "Leading Edge" booklet prepared by Bendix to explain their programs to new employees.

7. A follow up interview with the production manager (David Doyle) and plant tour was conducted in mid-1994.

8. In addition the author was employed by Bendix Mintex from 1983 to 1988 as a systems analyst and was deeply involved in the implementation of the foundation modules of the MRP II package, especially the manufacturing modules.

Brake and Clutch Industries Australia.

1. Plant tour of Kanban implementation on Mitsubishi Magna front calliper production line and accompanying handouts.

2. Semi-structured recorded interviews each lasting about one hour with the following people:

John McArthy (Materials Manager) 22/09/92.

Stewart Smith (Systems Manager) 22/10/92.

Peter Chiswell (Director of production) 22/10//92.

Kodak (Australasia).

1. Plant tour of Kanban implementation in the Plastic and Metal Products department.

2. Semi-structured recorded interviews each lasting about one hour with the following people:

Vivienne Triplett (MRP II Coordinator, Consumer Products Manufacturing) 15/12/92.

David Palmer (Planning and Procurement Manager, Consumer Products Manufacturing) 17/12/92.

David Mc Aloney (Production Supervisor, Plastic & Metal Products Manufacturing) 17/12/92.

3. Published description of the Kanban implementation for plastic components in the APICS Victorian chapter news letter, December 1992, by David Mc Aloney.

4.2 Bendix Mintex Case Study

This case study reports the changes in production systems that occurred at Bendix Mintex in the period 1988 to 1992. It records and analyses the swing that took place over that period from a commitment to MRP II over to complete commitment to a Lean Production and Cellular Manufacturing approach. A condensed version of this case study appears in paper one (Johnston 1995).

4.2.1 The Company

Bendix Mintex, formed in 1955, employs about 600 people and is the largest industrial employer in the provincial town of Ballarat, Australia. It manufactures friction products (brake pads, brake shoes, truck brake blocks and industrial products). Annual sales are about $60,000,000. Bendix Mintex supplies original equipment to the local car industry, but the largest proportion of sales is replacement parts both domestic and export. It has access to world class technology through its parent companies Allied Signal (US) and Pacific BBA (Australian). Its ability to manufacture a wide range of replacement parts (approximately 4000 line items) in small batches is its major competitive edge.

4.2.2 Bendix Mintex in 1988

Systems

Bendix had had minicomputer-based systems in the financial area since the mid 1970s. By the late 1970s these included on-line raw materials, work in progress, and finished goods inventory, bill of materials, and limited raw materials requirements explosion. These systems were all in-house developments. In mid 1987 an integrated manufacturing (MRP II) package was purchased as a result of a systems audit performed by external consultants. A full time Project Leader was appointed, a Management Steering Committee formed, and, using a user project group methodology, all existing systems were replaced and enhanced by the start of 1988. Thus by 1988, all the foundation systems for MRP II existed, namely, accurate computerised inventory, bills of material and routeings, purchase and works order tracking.

Production Planning and Control

At this time, production works orders and purchase orders were initiated by re-order report methods. There were about 1000 finished goods line items. A report of stock levels, customer back orders, and existing works orders was reviewed weekly by a member of the Sales department, and computer works orders were placed to replenish stocks using informal minimum and maximum stock levels. Computerised production batch cards showing the latest production method were printed by Production Control who processed transactions against the works orders that recorded progress, scrap and completion of the order.

The batch cards were released immediately to the production foremen who controlled detailed scheduling of production, with production for back orders having the highest priority. As there were approximately 500,000 units on back order at the time, the priority system had completely broken down with virtually every order having urgent status. There was a huge backlog of works orders and it was generally felt that the production foremen were producing what they wanted to:

> That was in the days when the production foreman decided what he was going to make. He used
> to pick out all the high volume products he liked making. (Chief Executive Officer)

Just prior to the period in question, a sales and production review process was put into place to ensure that the works orders raised in a given week did not exceed capacity.

Bought items and manufactured subassemblies (backing plates) were reviewed weekly, using a non-time-phased requirements explosion report, but were ordered in large batch quantities mainly by re-order point.

Culture and Management Style.

Prior to 1980, Bendix had been a classic "Industrial Era" company (Doll and Vonderembse 1991), being very hierarchical and functionally specialised, and the management style set by top management was autocratic. In the era up to 1988, top management and some middle management had adopted a more laissez-faire style but this had largely not filtered down to the supervisor/foreman level.

Manufacturing Processes

This case study concentrates on the disc brake pad manufacturing operation. A packaged disc brake pad set consists of typically four pads, not necessarily identical, each fabricated from friction material attached to a pressed steel backing plate. Other minor items of hardware may be attached or supplied loose. The two main manufactured sub-assemblies are the backing plates, which are punched from steel coil in up to three punching operations, and the friction material pre-mix, which is mixed in large quantities from a number of raw material and treated as floor stock in the moulding area. In the manufacture of a typical disc pad set the backing plates are pressed or removed from stock, the friction material is molded in a heated die onto the backing plate, various grinding and finishing operations bring it to size specifications, the various individual pads of the set (left-hand, right-hand, back and front) are shrink-wrapped together and placed in a brand name box with various accessories. Figures 1 and 2 show the Bill of Materials and Routeing for a typical disc pad set.
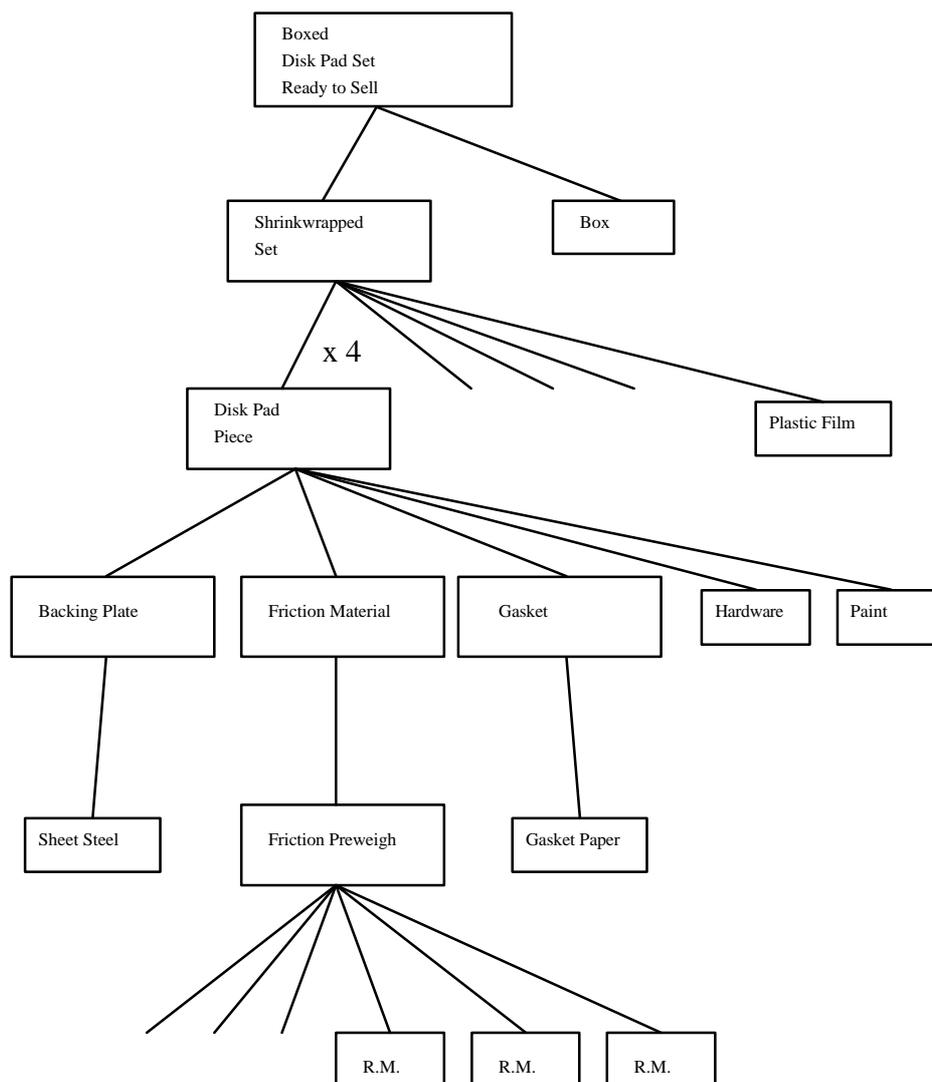
Figure 4.1: Bill of Materials for a Four Part Disk Pad Set

Human Resources

Prior to 1988, human resources issues had been receiving increasing attention, programs including middle and first line managements training, occupational health and safety, and Total Quality Control.

Step	Operation
10	Pierce, blank and shave backing plate
20	Inspect backing plates
30	Decrease and shot blast plates
40	Weigh count backing plates
60	Store in disc pad scheduling area
70	Mould disk pad
90	Blank and fit gasket to hot blank and clamp
100	Groove pad on grooving machine
110	Clean and paint black through wet paint line
120	Grind to thickness (non-magnetic)
130	Brand assembly
140	Fit brake accessory(s) hydraulic
150	Inspect to PCP
160	Release to "Kit Assembly Area"
170	Select parts from kit assembly area
180	Shrink wrap disk pad st components
190	Pack in wire basket
200	Inspect to standard specification
210	Release to distribution warehouse

Figure 4.2: Standard Routeing for a Simple Four Part Disk Pad Set.

Plant Layout

By 1988 some recent plant layout changes had moved all the machines for disc pad manufacture into a logical flow arrangement under two roofs. The minor product lines were less logically laid out. The machine layout however, was strictly by function and the positioning of the functional groups of machines reflected the manufacturing process. A $A1,000,000 warehouse (1988 dollars) had just been completed and the buildings had reached the boundaries of the site.

Production problems

1. There was a huge back order book (approximately 500,000 units) leading to poor customer perception of service and negating the official policy to supply all lines from stock with a high customer service level:

> We got the message loud and clear that our order fill rates were terrible. (Chief Executive Officer)

2. Large work-in-progress stocks were held, especially in the "kit assembly" area where the several pieces involved in one set were brought together and shrink wrapped. The large batch size (typically 3000) for the individual pieces and delays in routeing meant that the batches were seldom completed together. There were also large stocks of backing plates again due to large manufacturing batch sizes.

3. No formal master production scheduling system existed. For instance there was no formal system to allow large export orders with known future due dates to be made to schedule, and they were treated as though they could be supplied from stock.

4. There was a high reject rate in manufacturing.

5. Throughput time from works order initiation to warehouse was long (about 21 days). The main causes of this were the unreliability of backing plate stamping, difficulty of routeing the batches through the complex finishing processes, and difficulty in coordinating manufacture of the several parts of a given set which were made in separate large batches (typically 3000 pieces).

4.2.3 Changes at Bendix Mintex, 1988 - 1992.

Strategic Background

In the period just prior to the study period there were two events that drove rapid expansion in production at Bendix. Firstly, there was a period of low exchange rate against the US dollar that opened a window of export opportunity. Secondly, the company's major domestic competitor made the premature decision to market only asbestos-free products at higher prices. These two events led to sales orders far in excess of production and there was a high strategic priority to increase production capacity. The plant had reached its geographic limits of expansion, so there was a strong desire to increase productivity through better systems and practices. In addition there was a general background fear of potential competition from low labour-cost countries.

MRP II Implementation

With all foundation systems in place by 1988, in mid 1989 a project was initiated to implement the Material Requirements Planning, MRP, and Master Production Scheduling, MPS, modules of the MRP II package. By this stage a user-driven, management-steered project management methodology was well established in the company. A project group of eight people at middle management level representing all affected functional areas was formed and the project leader, the then Materials Manager, had had a great deal of project management experience from earlier software module implementations.

The project started with a 3 day course on the use of the MRP and MPS modules given by the software supplier to all project team members. This proved to be pitched too much at the system use level and it was decided to call on consultants to provide more conceptual education. The project group attended a 5 day off-site course run by David W. Buker Australasia Pty. Ltd., who also provided some consulting on implementation. This was followed up with a one day on-site course attended by all top management.

By August 1989 the group had a proposal and an implementation plan to have all product works orders initiated by MRP/MPS by 31/07/90.

The proposal was for:

• Forecasting of finished goods requirements from historical usage, using the provided forecasting module;

• Use of forecast and customer orders by the MPS module to create an MPS;

• Monthly Review of the MPS by the Production and Production Control managers for capacity feasibility;

• A monthly Sales, Inventory and Operations Group meeting where top management reviewed the plan for compatibility with sales and financial objectives;

- Weekly MRP explosion for disc pad manufacturing requirements, and purchased raw materials and hardware requirements;

- Production Control to release MRP planned order releases for production according to start date.

The plan was pursued vigorously and by March 1990 operator training had commenced. In mid-April 1990 a month long pilot run was made, with a single product line being controlled under the system. The cut-over to the MRP system for all products occurred in October 1990.

Immediately after the full product range came on-line, two design flaws became apparent. The first was that reviewing the forecast at a line item level was impractical with 4000 line items. This was eventually addressed by aggregating the forecast at a product class level within an ABC analysis classification. The aggregated forecast was then down-loaded to a PC spreadsheet, reviewed and altered by management and up-loaded back to the mini-computer system. The percentage changes at the aggregate level were then applied to the individual product forecasts in each product class.

The second flaw was that for products that were sold both to export and domestic markets, the export orders, which are infrequent and generally for large quantities, were disrupting the forecasting for these line items. An agreement was made with these customers that the products would be ordered with 5 week lead time and would always be made to order. A modification was made to the MPS program to remove this deterministic demand from the forecast. Previously, all products had been controlled by order points, largely because there was no formal system for separating deterministic demand. This had caused excess stock problems because the production planning person tended to keep at least part of the typical export order in stock, and also an excessive back order situation because export orders almost always went straight onto back order. Thus, the implementation of the MPS system allowed the real nature of this problem to be seen for the first time, and led to its solution.

More alarming though, was that despite encouraging results from the pilot run, it became apparent in the first year of operation that the MRP II system was not solving the finished goods and WIP inventory problem, because it was not insuring that the right products were being made, and it was not addressing the throughput problem. The first problem was that while the schedule was accurate at the aggregate level it was not accurate at the individual line item level, and it was only being reviewed at aggregate level. As the demand forecaster put it:

> If I had the time to go through each of those part numbers individually each month ... the inventory would improve. (Demand Forecaster)

The second problem was that there were only two control-points (points where actual production progress was measured) in the product routeings. So, while the system could recommend that the product be started on a certain date for completion by a certain due date, it gave little information to actually guide it through the plant:

> We were not improving lead times, we were not reducing WIP, in fact WIP was almost out of control ... we were not improving productivity - it was taking more people to produce that same amount of product. (Production Manager)

> The MRP did not provide any visibility whatsoever. (Production Manager)

Both these problem were viewed as being caused by the impracticality of completely fulfilling the logic of the MRP II model:

> If we put in {more} control points it would be too much paper work to control the system. (Production Manager)

> The only way you could get a lot of control over it {stock} would be if I had the time to do it {scheduling} by end item. (Demand Forecaster)

There were also difficulties in maintaining real involvement of the participants in the MPS review processes, which were stated by the master scheduler as:

- Lack of time to prepare.

- Those with the best market knowledge were too busy to contribute adequately.

- Loss from the company of the original coordinator of the Sales, Inventory and Operations meeting.

In mid 1992, a decision was made as part of the change to cellular manufacturing, that as the cells came on-line order release for the products made by the cells would revert to order point methods:

> The MPS is still needed, but we are not generating works orders because of the forecast. We generate orders now because of the demand. So it's a pull system instead of a push. (Materials Manager)

Thus, over a two year period, works order initiation had gone from order point methods, to MRP, and back to re-order point again, though as will be seen below, within an entirely changed environment. MRP/MPS was still used for purchased part ordering, although there was (in late 1992) a trial group of vendors supplying by Just-In-Time methods.

Human Resources

Prior to the period in question, Bendix had conducted a Total Quality Control (TQC) program (late 1985) aimed at harnessing work force problem solving abilities. There had also been an extensive training program aimed at multi-skilling, driven by the "second tier" wage negotiation of 1987 (this was a Labour government initiative to encourage unions to accept multi-skilling in exchange for award increases). The latter program was steered by a joint management / employee committee and run by departmental project teams.

About 1989 management started to agree that the reason these initiatives had not affected the bottom line was that the skills the *ad-hoc* teams of employees were able to demonstrate in the democratic TQC meetings and training environments, were not being carried over to the autocratic shop floor:

> We still had a very autocratic shop floor at that time and we were taking the people out of that environment and then putting them in a class room ... with engineers, management and some peers, and we were saying to them "we have a problem to solve, we have been set the task to solve that. We're all equal now {so} let's see if we can set about doing it". But they weren't all equal because as soon as they walked out of that class room, they were back on the shop floor getting the shit beat out of them if they did anything wrong ... they were never comfortable in that artificial environment, and being told they were equal. (Director of Human Resources)

They concluded that what was needed was a "socio-technical redesign" of the workplace. The objective was to reorganise the shop floor into self-managing teams. With the help of consultants a series of workshops were conducted, first with top management, then with first line supervisors and shop-floor representatives, respectively. From these the need for substantial training of supervisors in participatory management was identified, and a 9 month program including training, counselling, and on-the-job observation was conducted.

In July 1990 a "Future Directions Workshop" involving 45 employees, half from the shop-floor and the other half from a cross section of the rest of the organisation, was held off-site over three days, outside of working hours. The object was to produce a vision statement written and owned by the employees, to define key objectives in meeting it, and to develop action plans to achieve it. The remaining employees were briefed by workshop participants. Training programs on team concepts followed this initiative.

In Nov 1990 a "Social Redesign" program was initiated involving 20 employees for 10 days full-time work within working hours. This group came up with a new organisation structure which defined the role and responsibilities of the teams. The teams were supposed to take responsibility for certain product groups, indicated by colour coded batch cards, work on them selectively, and nurse them through the complex routeing between machines. The teams were also to act as quality and process improvement groups for these products. A joint management / employee Workplace Reform implementation group was formed. Formation and training of the teams commenced in April 1991, and much of 1991 was spent implementing the team concept.

Early attempts at Just-In-Time and Kanban

The earliest attempts at Just-In-Time production were before the period in question, and resulted largely from knowledge imported with the appointment from outside of two production management staff. The pre-formulated friction material mix used in production was "pulled" to production by a Kanban system by about 1986, and Kanban control of clips and hardware followed. These early applications Just-In-Time involved short lead-time, constant volume items.

An early attempt to control the production of disc pads by Kanban, around the start of 1991, was not successful. This consisted of a large Kanban display board in the production area which held the cards according to the next operation required. The operators were supposed to put the cards on the appropriate next operation after completing their work, but this did not happen reliably and this attempt at Kanban eventually failed. Reasons given for the failure were:

- Inability to implement small batch sizes in that production environment.

- The complexity of the product range and plant layout counteracted the simplicity and visibility of Kanbans:

> {The Kanban board} became as difficult to maintain as the data entry. (Production Manager)

- This attempted Kanban systems in the main plant addressed the lack of a scheduling system, rather than being an implementation of true Just-In-Time production and could easily be ignored by employees:

> They were not moving {the Kanban} along the process. They saw that as a time consuming waste of effort. (Production Manager)

A more successful Kanban system (to be described below) was introduced in the first disc-pad manufacturing cell in mid 1992.

Plant Layout

By 1989 the first round of plant layout changes, from a historically based "mess" of machines to a layout that rigorously reflected the production process from raw materials to finished goods warehouse, had been completed, and was responsible for a reduction in throughput time of production orders from 30 days to 21 days. A project to change the plant layout from one based on collecting machines by their process to one consisting of cells of machines which relate to groups of products, was driven by the recognition, in mid 1991, that the cultural change initiatives were being hampered by the unfavourable physical environment that the teams were working in:

> The teams were not working - these social teams were not working - we said "it's got to be the environment - it's got to be the physical environment". (Production Manager)

> We kept talking about socio-technical change and we said "we have done the socio-technical change" and then we said "no we haven't, we've done the social change; we haven't done anything to the technical system". (Production Manager)

The teams were supposed to take responsibility for specific groups of products which were identified by a colour code on the batch control cards:

> It was put on a batch card that this was the black team's pads so when you got it down the end {finishing} you made sure that you worked on a black team pad. So you had blacks and blues and greens and reds going through and what people said was "stuff it, I'm not working on that because that's a red one. I'm black". So they would go through all the WIP and pull out the basket that was black. All it created was argument, debates and fights. The environment just didn't suit. (Materials Manager)

What was needed was a physical focus for the teams:

> {They} needed the physical definition of the cells as well as the mental definition of people being in a team. (Chief Executive Officer)

> The members of the cells {teams} were spread from one end of the plant to the other and you wouldn't see someone for a week even though you were in the same cell {group} so there was no ownership of what you were doing with a product. (Materials Manager)

Reluctance to address this physical problem was strong because:

> It was too difficult, it would cost too much because you have to re-layout the plant. (Production Manager)

However, in 1992 a grant of $A500,000 (1982 dollars) became available for this task under the Australian Best Practices Demonstration Program. A tender for consultation was accepted from the UK-based Ingersol Engineering. These consultants redesigned the machine layout into five product based cells. These were based on the identification, using existing data on product routeings and Group Technology techniques, of five product groups with substantially similar operations and machine requirements. A major objective of the redesign was reduction in material handling. All machines and personnel involved in the manufacture of each group of a disc pad sets, from moulding to shrink wrapping, were to be assembled into "U" shaped cells with improved material handling systems and complete visibility of the whole process by all concerned. The cells thus fitted in neatly with the team concept. They also proposed a physical Kanban-based production control system to "pull" product to the finished goods warehouse, but this latter recommendation was felt to be unrealistic and a hybrid order point / Kanban system was adopted.

Under this system, a predetermined number of product Kanbans are held in the production control area. When the order point on a product is reached, as determined by a computer report, an appropriate number of Kanbans is released to production and a computerised works order is raised which results in a printed batch sheet. The purpose of the computer works order is to register on the computer that the potential shortage is being addressed, to cause the printing of the latest method, and to allow the valuation of scrap via the computerised routeing. Thus, the computerised works order system is serving some functions that a full JIT implementation would address, and reflects some old style concerns. The adoption of this hybrid system suggests that the company has not yet found an adequate solution to using Kanban in a lumpy demand situation (mainly caused by the deterministic export orders).

Kanbans for both disc pad and backing plate manufacture are placed on a Kanban board at the start of the cell line. The backing plate is manufactured first and the backing plate Kanban returns to the board. Then it travels through disc pad manufacture with the disc pad Kanban and batch sheet. The Kanbans are returned to production control on completion. When large export orders are to be made for parts that have their Kanban numbers set for the steady domestic market this Kanban cycle is too slow and it is necessary for Production Control to inject "special" Kanbans into the system temporarily.

At the time of the main part of this study, the first cell was operational and the second was just about to start operations. On a later follow-up visit in 1994 all five were running. The first cell manufactured the simplest disc pad sets, namely those with 4 identical pads. Nevertheless, the performance of the cell against previous methods was very impressive. The pads were made in lots of 300, being approximately one shift production, compared to lots of 3000 previously. Throughput time for a production order to reach the warehouse was now 2-3 days, with most of this time in raising the paperwork, against 21 days previously. Back orders for products made in this cell had been reduced to practically nil.

Disc pad backing plate manufacture was not included in the cell but they were attempting to manufacture in the small lots required (300 instead of up to 40,000 in previous times). Tool reliability and set-up time remained a continuing problem, and plans were under way to relocate the backing plate / tool making operation off-site as an independent Just-In-Time supplier to the disc pad operation.

Although the basic disc pad manufacturing process had not changed in the period in question, many gradual improvements in machinery and tooling had taken place, especially in the area of faster change

overs of disc pad moulding dies and standardisation of die fixtures. However, it is important to note that management decided to reduce batch sizes by a factor of 10 for disc pads, and up to 100 for backing plates, without these changes being accompanied by corresponding set-up time reductions. The reductions were entirely motivated by the presumed benefits of increased flexibility and response to customers, faster throughput, and simplification of synchronising the production of the various parts of a complete set.

4.2.4 Bendix Management thinking - late 1992

Status of MRP II

After initial enthusiasm, especially after the 5 day Buker course, and a long period of infra-structure preparation, it is fair to say that Bendix management had completely rejected MRP II as a production planning and control methodology:

> What we know now is that MRP II doesn't work with JIT production, which is what we're implementing here. (Director of Human Resources)

The reasons for this rejection were complex, not the least important being ideological. The CEO had early on expressed reservations stating:

> ... that you will install MRP over a period and then you will progressively dismantle it over a period ... as you changed {the factory} around and you went toward cells or dedicated manufacturing facilities, that you could then progressively take out that push type system and replace it with a flow or Kanban driven system. (Chief Executive Officer)

This may have biased some of the participants toward the "pull" direction, but on the other hand, there was no lack of top management support for the MRP II project itself.

But the main reasons for rejection of MRP II were, firstly, the amount of data processing required,

> MRP on the shop floor in my view is just too cumbersome given all the {data} processing. (Chief Executive Officer)

which was seen as

> Non-value-added cost associated with MRP. (Chief Executive Officer)

and secondly, its remoteness and complexity

> We've {the implementation team} gone through all the steps that many times that we know it's correct but if anyone else asks us, trying to explain your way through is very complicated. We trust the system but a lot of other people don't. (Demand Forecaster)

It was now viewed only as a system for dealing with the outside world, that is, planning customer's requirements and supplier orders:

> It's still good for forecasting future supply needs, purchase needs, the global things for the future, but when it comes down to running it {production} you need something that responds a whole lot quicker than MRP II will. (Director of Human Resources)

> "MRP was just a system that could be used up here for forecasting, ordering raw materials and perhaps creating batch cards" Production Manager.

for which it was working well.

> There is no doubt that this is a big improvement for getting things from outside. (Demand Forecaster)

However,

> If you have very short lead times ... then I don't believe you need MRP at all. (Chief Executive Officer)

On the other hand, Bendix accrued some substantial benefits from the MRP II program. Firstly, under the rubric of MRP II, a great number of infrastructure functions were systematised and automated, and vital product / process data was formalised.

> We are managing the whole process far better today than we ever have. And it's because there is a lot of information on the systems and it's good value information. You can get it and you can believe in it and you can work with it. (Materials Manager)

In particular, the existence of accurate routeing data was important in the production cell design, and recognition of the problem of attempting to supply the deterministic export demand from stock came out of better master scheduling methods. Secondly, the process of implementing the infrastructure systems for MRP II contributed to the development of an effective change methodology in the company which was repeatedly drawn upon in later projects. Because an MRP II system crosses functional boundaries, any company that is to successfully implement it must develop project management methods that allow people to organise for a more than usually political task.

The Importance of Self-managing Groups

A prominent theme of the case was the central importance for production improvement that Bendix management put on reorganising the work force into self-managing teams:

> {We} had to get people into small groups if we were going to bring in this continuous improvement culture - people need to see the start and the end of the process. They need to feel involved in the total process rather than just {be} a machine operator. (Chief Executive Officer)

Authoritarian management methods on the shop floor were to be rejected,

> We were really trying to flog this dead horse about whipping people to improve, you know, beat them into improvement... and then the penny dropped. (Production Manager)

and the functionalised organisation was to be broken down by a (physically) closer working partnership between production and production support personnel.

> In the past these guys {engineers} have sat in their office and they do their thing and they don't really know what in the hell is going on down in production and they don't really want to know. Now that they're down there, and they're out there working with them, they actually see the extent of the time it takes for some of the set-ups, they walk away and say "that's crazy", and so they're doing something about it. (Materials Manager)

Rejection of the traditional (Taylorist) work organisation was seen as a prerequisite for success with Lean Production:

> You can't take a traditional organisation and put in a pull system and expect it to work. (Director of Human Resources)

However, the implementation of the self-managing concept was not yet complete,

> {We} still have cultural issues - a lot of people have been here a long time. They don't want to rotate around different work processes, they don't want to come here and take some responsibility. (Production Manager)

nor without casualties.

> They {some foremen} wouldn't fit in with the new concepts and they were offered a package and terminated employment. (Materials Manager)

The Role of the Physical Environment.

Visiting the site after 5 years absence, to me the most striking thing was the pivotal influence that the changes in physical layout had had on several otherwise frustrated initiatives:

- The use of the team concept to improve quality;

- The efforts of the teams to improve throughput and reduce material handling;

- The use of visual pull systems for production control.

The cellular layout provided a focus for these initiatives, but also credibility:

> It's the first major move that's happened for the people on the shop floor to start realising "Hey, these guys (management) are fair dinkum finally". (Materials Manager)

Bendix management was seeing that improvement of production management was not just a matter of implementing isolated systems or ideas:

> You can't just attack the systems, you can't just attack the people {issues} or the physical environment. You have to do them all together. (Chief Executive Officer)

> {Lean production} is more than just the physical layout, it's more than just the social change, it's more than the paperwork system, it's all these things lumped together. (Production Manager)

The importance of the physical environment was stressed by the Human Resources Manager when he suggested that the logical order for the changes was: physical layout first; then teams; and only finally quality tools (TQC); which was the reverse of what they had done. There is no point giving people the tools for improvement until they have the physical environment, the political structure and the social environment for change.

4.2.5 Conclusion: Bendix Mintex Case Study

It is important to recognise that the failure of Bendix Mintex to adopt the MRP II model (for production control at least) cannot be described as a mere "implementation failure". All the factors that are generally recognised as critical to success in MRP II implementation were present: management support, as evidenced by financial commitment, aggressive project schedules and active participation in a steering committee; real user involvement; sound project management practices; and adequate education and training. (The MRP II implementation resembled in many respects the widely recommended "Proven Path" method (Wight 1981; Wallace 1985)). On the contrary, Bendix Mintex rejected the MRP II approach because of their gradual discovery of the importance of self managing teams and the nature of the physical environment in which the teams work to improving production performance, aspects of the problem not even considered in the MRP II framework.

This, the main observation drawn from the case, points to an important question, which is the subject of this thesis, namely, the relationship between the situation in which production takes place and attempts to improve production performance, and how the various paradigms of production management deal with this situatedness.

4.3 Brake and Clutch Industries Australia Case Study

This case study reports the adoption of Lean Production methods in another batch repetitive automotive equipment manufacturer which started out with a strong commitment to CAPM methods. In this case MRP II was initially replaced on the shop floor with a more complex computerised scheduling approach, OPT, but this was later rejected. Apart from confirming the observations at Bendix Mintex about the inadequacy of MRP II and the role of physical changes and teams in improvement, it also tends to indicate that simply removing some of the technical limitations of MRP does not solve the inherent problems with CAPM approaches.

4.3.1 The Company.

Brake and Clutch Industries Australia was formed in 1987 when the UK firm British Belting and Asbestos, BBA, acquired five Australian brake equipment manufacturing firms including Girlock, Repco, and Patons Brakes. They manufacture original equipment braking systems (approximately 200 lines) for all four automotive manufacturers situated in Australia, and support the after-market with both manufactured and imported parts (approximately 6000 lines). This study concentrated on the original equipment operation, which represents about 60% of business, because manufacturing of these lines has been controlled by MRP, OPT, and Lean Production at various times since 1987.

4.3.2 Performance

In 1987 the five companies had 2000 employees while in 1992 there were 1100. In the same period sales had only fallen 15% (mainly due to the down turn in car sales), while total inventory turns had increased from 4.5 to 7.5. Thus, considerable efficiencies accrued from the merger. These were attributed to:

- Better forecasting of finished goods requirements.

- Better product range management.

- Better systems, leading to less safety stock.

- In the case of WIP inventory, better plant layout and introduction of a single product database.

4.3.3 Strategic Focus

One of the strategic objectives of the BBA acquisition was to make the Australian operation, through its subsidiary, Pacific BBA, the seat of Pacific / Asia operations. Thus there was a strong desire that BCIA should have model production and management systems in place.

4.3.4 Material Requirements Planning, MRP

As part of the amalgamation, BCIA acquired the Repco Data Centre which was a large (approximately 20 staff), centralised department responsible for in-house computer system developments. They had been using an MRP system developed in-house around the IBM COPICS MRP module, running on an IBM mainframe computer, since the late 1960s. A number of Repco companies used the centralised system and some apparently had quite successful implementations of MRP as a production planning system.

After the takeover, a number of changes took place to the MRP system. Firstly, the multitude of inherited product databases had to be combined into the current two (manufacturing and purchasing). Secondly, the MRP system was tailored to be more appropriate for the work flow of the original equipment lines. MRP was run monthly, producing production schedules and suggested works orders. Although data was, and still is, collected on the progress of these works orders, the production schedules were being ignored by shop floor staff.

> It {the MRP schedule} was ignored by the plant. If you went out and asked them, they didn't use it. That was the last thing they looked at. It didn't mean anything. It worked in terms of the arithmetic but it wasn't reality, it wasn't valid. The only thing that they did take notice of, that was valid because it was a real demand, was what the customer order was. And what they did was go to the final assembly operation and ask him what he needed tomorrow and they went and made that. (Materials Manager)

As a result, they "ripped out" a lot of the MRP planning parameters, that is, they reduced many lead times to one or even zero days, reduced minimum order quantities, and so forth. This meant that the MRP system was suggesting a materials plan that was tightly coupled to the customer schedule on a lot-for-lot basis. It then became apparent that this tight coupling demanded that the finished product schedule be smooth, whereas the car company schedules were generally received as very discrete (lumpy) requirements. Therefore, around mid 1991, the master scheduling portion of the MRP system was modified to smooth these schedules to daily requirements before the MRP system was run. These changes were anticipating several features normally associated with "pull" systems, namely tighter coupling of materials requirements to finished goods requirements and smoothing of final assembly schedules.

Parallel with these changes, the company had been implementing the OPT finite capacity scheduling package. At this time, the idea was that the MRP system would deal only with long term scheduling and purchased materials planning, and OPT would fill the planning gap at the shop floor operations level.

4.3.5 Optimised Production Technology, OPT

The idea of using OPT for production scheduling came from the British parent company, BBA, who had then been using it for three or four years. In mid-1990, an executive project group visited the BBA site in England and also Bendix sites using OPT in France, and reported back favourably. The decision was made to install this system initially to address scheduling problems in the after-market product range, but later for some large set-up products in the original equipment range also. The OPT package was installed by a team from the vendor company over a six month period. They trained several BCIA people as "OPT analysts" to run the system. These were chosen from the industrial engineering area, rather than production, because they were felt to have the right "mind set".

OPT was run for about two years on selected product lines but failed to be adopted as the solution to the production scheduling problem, and in fact it had been "turned off" a few months before this study. As to its performance, the informants disagreed. One felt that it worked well for the products it was tried on, and that it produced technically valid schedules. Another stated that the schedules produced were clearly not in accord with sensible manufacturing practice, and the third felt that the system had not been given an adequate chance. This issue was controversial, fresh in recent memory and all informants were stakeholders. Nevertheless, they all agreed on certain problems that OPT had other than its technical validity:

- It required a very high level of expertise to operate, depending on highly trained "OPT analysts" (Manning 1983);

- Its output was not at all "user friendly";

- Its operation was foreign to anyone on the shop floor;

- It required a lot of data collection;

- It required the setting of an extremely complex set of parameters and there was no way to predict the answer it would give;

- It could not easily handle changing business conditions;

- It had to be trusted and followed blindly:

> He {the foreman} would say "why do I have to do this (set-up and tear-down)". And you would have to say "because OPT says so". (Materials Manager)

4.3.6 Lean Production

Since his appointment in 1990, the Manufacturing Manager had been reorganising the operations in a direction anticipating Lean Production. These changes included -

- Getting employees used to working in cross functional groups, initially through informal quality groups.

- Rationalising and flattening the plant management structure.

- Reorganising the plant physically to make space for a more cellular (product oriented) layout.

He stressed the importance of teams,

> Without teams you can forget the whole thing {Lean Production}. (Director of Production)

but rather than seeing the team concept as another "silver bullet" he saw it as an evolutionary process:

> First of all you get people working in teams to understand the concept ... the feeling of a group working together. Then you can formalise it, and say 'well that's actually a team, guys, now let's formalise it into work teams with team leaders' and you re-layout the plant in conjunction with those teams. (Director of Production)

In mid 1991, the Managing Director took part in a study tour, organised by the Federated Automobile Parts Manufacturers, to visit 13 plants in the US using Lean Production ideas and returned as a convert. Very soon after, he sent a group of 18 top managers from all functional areas to the US to study five of the sites visited earlier. These included US owned companies and Japanese implants, all using US workers. Several members of the party had been on study tours to Japan previously, but

> We all had this mind set that its Japanese culture and all that. (Materials Manager)

Thus they were impressed to find that some of these companies had made dramatic improvements in productivity in a Western environment using Western employees On the last day of this tour the group got together to discuss how Lean Production could be implemented in the company. They realised

> We had to do several things on the same front. We decided that it wasn't just printing Kanban cards when we came back, and hanging them on hooks. It wasn't just getting the people together and telling them they're empowered. We had to do all things on the same front. One of the core things we had to do ... we had to put the site into the physical shape for the system to work. (Materials Manager)

Returning to Australia, it was decided to adopt Lean Production for the original equipment product lines. One group worked to redesign the plant layout. At the time of the study they were in the process of laying out anew the entire original equipment operation along flow (cellular) lines. In another working group, five members of the group of 18 were taken out of their jobs, full time, to design the Kanban based production system. This group used physical simulation methods and PC simulation to test out their proposals. They took groups of shop floor people, from all shifts, ten at a time through these simulations. This was the first time shop floor personnel had been involved this directly with the design of a system that affected them. Furthermore, unlike OPT, where the foreman had to trust the calculations of the OPT system,

> They could see it on the table. They could see it working. They could actually see that when you got to here you stopped because there was no point in making any more. (Materials Manager)

4.3.7 The BCIA Situation, late-1992

At the time of this study the Kanban-based, Just-In-Time system was in operation for all processes in the manufacture of Mitsubishi Magna front brake callipers. The physical layout changes were not yet complete, and the processes took place in three separate plant locations each using internal Kanbans and supplying each other as Just-In-Time suppliers using "move" Kanbans.

The results were dramatic, not only exposing some vastly overstocked components, but reducing work in progress stock from 8 days to about 4 days supply. One sub-assembly that represented a large proportion of the material cost (clipped friction pads), had stocks reduced from six weeks to a few days. The stocks of these parts had got so large because:

> The foreman just went off and assembled pads whenever he felt like it. And we put the cards in and they hung on the hook for six weeks before the guy had to make one. (Materials Manager)

MRP remains as a forward scheduling and purchased item planning system, and for labour capacity planning, and all parties agreed that it would continue this role for some time. Production control would be via Kanban, within the Lean Production framework:

> We see the Kanban card as being that centre piece inside the MRP planning cycle. We see the skeleton of MRP in the way we have done it, remaining as it is. (Materials Manager)

All parties commented on increased customer focus.

> With a pull system, it certainly brings everybody's focus together on what you are doing in terms of servicing the customer. The other way, you are forever trying to find solutions to improve a number {such as} labour efficiency ... If we are servicing the customer and we have no inventory then there's not too much that can go wrong. (Materials Manager)

4.4 Kodak Case Study

Again this case study records the adoption of Just-In-Time methods within a batch repetitive manufacturing company that had formerly run their operations using the MRP II approach (they had been certified at "Class A" performance). In this case, the company did not have full commitment to the entire Lean Production approach (which at the time had greater influence in the automotive industry than elsewhere). However, we again see that when they wanted to improve responsiveness to customer requirements, they found the information intensive nature of the CAPM approach a barrier.

4.4.1 The Company

Kodak (Australasia) manufactures and distributes photographic film products to the Asia/Pacific market and is a subsidiary of the Eastman Kodak Company. It employs about 1500 people, 800 in manufacturing. This study concentrates on the consumer roll-film operation.

Kodak has been assessed as an excellent company. In 1992 they won the inaugural Australian Quality Prize, which is modelled after the Demming Prize in Japan and the Baldridge Award in the US, and is internationally audited. In addition they have implemented an MRP II approach to manufacturing that has been independently assessed at "Class A" level.

4.4.2 Strategy and Culture

Kodak competes in the Asian market, largely against other Kodak companies, on the basis of cost, but also on short supply time and minimal variation in supply time. Since 1989 shortening lead time has become a major focus and has affected manufacturing systems directions.

There is a strong culture of Total Quality Management and continuous improvement. There is a widely used and understood internal methodology for change management that emphasises involvement of shop floor personnel. Upward and downward communications are good, and natural work teams and flattening of management structures are influential concepts.

4.4.3 Manufacturing Resource Planning, MRP II

In 1984, the parent company, Eastman Kodak Company had initiated a world-wide commitment to MRP II using common computer systems. Implementation of MRP II in Kodak (Australasia) dates back to 1985 but in the early days MRP was seen as a computer system rather than a theory of manufacturing management. Senior management did not see that they had a role. The implementation group had two personnel with APICS certification, who were very knowledgeable on MRP II.

In 1989 Eastman Kodak threatened to close down the Australian operation. This event resulted in much greater commitment to, and involvement in MRP II by top management. Firstly, there was an urgent need to improve manufacturing performance, secondly, a committed and knowledgeable group in the MRP II camp now had an opportunity to be heard, and finally, the parent company announced that all operations were expected to provide a date by which they would achieve "Class A" MRP II status. The latter event meant that top management had a new motive for becoming familiar with MRP II principles.

Since 1989, Kodak has implemented MRP II through the roll film operation to "Class A" standard. The prime movers of this implementation have been the MRP II project team and the production planning department. The production departments, while not hostile to the changes, have seen it as a system belonging to production planning, and therefore have not had felt as much ownership of this system as with the later Kanban methods.

This MRP II implementation has placed a lot of importance on the Sales and Operations Planning (S&OP) and Master Production Scheduling (MPS) aspects of MRP II. The S&OP group consists of senior manufacturing and marketing management and middle production and planning management. Information on future orders is collated from customer and distributor sales forecast and is aggregated up to a product group level with inventory and production performance data. Each month they review the sales and orders at this aggregate level using a number of performance measures and key indicators, and prepare a plan for sales and inventory levels at the aggregate level to which both marketing and production are committed. This is disaggregated according to historical product mix and a Master Production Schedule is produced at the product level. Purchase orders and production schedules are produced from these.

Originally, the execution system was computer based Material Requirements Planning MRP, and after a move away from centralised mainframe computing, these systems where moved to a PC-based network with the MPS being down loaded to them from more centralised systems. The MRP based execution systems produced purchasing plans and detailed production schedules, and progress against these schedules was recorded as were movements of product between stores.

All the informants expressed complete commitment to the higher level functions of MRP II (S&OP and MPS), and believed that the benefits achieved had substantially accrued from the success of the implementation of these. They also believed that a really successful Kanban based execution system needed these functions to be in place.

> We do believe that we need to get control of the process by the implementation of the full MRP II process using shop floor documentation etc. first before we go look for opportunities to put in Kanban. (MRP II Coordinator)

As an indication of the performance improvement in the roll film area, part of which can be attributed to these systems, stock levels had gone from 13 weeks to 5 weeks in the period from early 1989 to late 1992, and in the same period the order turnaround time (time from placing an order to receipt of goods) for Asian customers had gone from 55 days to 20 days.

4.4.4 Just-In-Time Through Kanban

As noted earlier, a chief competitive focus of the company is reduction of lead time to the customer and reduction of the variability of this lead time. It is therefore natural that Just-In-Time ideas would be influential in the company. About two years before this study, the idea arose that Kanban based execution systems might be more appropriate than MRP in some areas where processes were sufficiently simple and reliable to support small batches of product required to increase responsiveness to customer requirements. Reduction of batch sizes within the existing MRP controlled process would mean an increase in shop-floor feedback transaction entry, and work associated with data entry errors, to unacceptably high levels.

A strong theme of the interviews was that the ground work laid in the implementation of the MRP systems, particularly gathering of process information and organisation of infrastructure, contributed to the success of later Kanban implementations. The S&OP process was seen as much as a prerequisite to getting the benefits from JIT efforts as from MRP, through the levelling of the production and setting of realistic Kanban numbers. Improvements in process reliability was also mentioned as a prerequisite.

4.4.5 Kanban in the Plastic and Metal Products Department

This area supplies all components except the film to the roll film processing and packing operation. These include high volume low cost components such as injection moulded plastic parts, metal spools and printed cartons.

The idea of using Kanban to control material flow for the plastic parts came up around 1990 in the discussions of a shop floor working group that was designing the shop floor procedures associated with the implementation of the PC based MRP system. These ideas came through the influence of APICS members in the organisation and Oliver Wight training tapes. However, the company at that time had a strong commitment to achieving "Class A" MRP II status and the MRP system was implemented. Under this system, computerised release schedules for the plant were generated and transactions were entered to record progress against the schedule.

The Kanban ideas surfaced again in mid 1991 when a working team, using TQM style methods was preparing a vision statement for materials flow in the area. A key motivation for this was the huge transaction demands of the, by then, implemented MRP system. (c.45000 transactions per annum). At the time the buffer stock between this department and its customer department was about 3-4 days stock and this necessitated and interim storage area. Movements in and out of this interim storage at a pallet level exacerbated the transaction problem which had first become apparent in the training sessions.

Two project teams were set up. A management team looked at lot sizes, Kanban numbers, inventory and WIP valuation reporting, performance measures and other policy issues. A shop floor team looked at Kanban design, Kanban rules and preventative maintenance procedures. A single card Kanban system was designed to run through production to the supplier department and back. The Kanban cards start from a Kanban board that also states the Kanban rules simply and clearly. The level for these parts in the Bills of Material for the parent parts had to be removed, and changes made to the reporting of the picking of these parts in the down stream MRP system. Quite a deal of work was also done to increase process quality and reduce variability.

There was very strong ownership of the system by the production people, who viewed it as much simpler and understandable than the former MRP based system. Some production staff however, still had difficulties accepting the idea of stopping their machines in smaller runs. Also there were difficulties in negotiating with the unions involved the change from truck load deliveries of product to the customer department to more frequent fork-lift deliveries.

This area cut over to the Kanban system in late 1991 and it has resulted in reduction of buffer stocks to less than one day of stock, or about a fifth of its previous level, and has eliminated the interim store. In addition the elimination of non value added steps has reduced the manufacturing lead time of these products from 50 hours to 2 hours. Quality has also improved through quicker feedback about problems.

4.5 Conclusions Drawn from the Cases

4.5.1 Themes Common to the Cases

1. **MRP II did not solve the shop floor operations management problem:** In all cases the standard MRP II approach did not solve the manufacturing operations management problem. At Bendix this lead to almost complete rejection of the concept when other methods proved superior. At BCIA the problem was initially diagnosed as a technical one, as it often is (Hopp and Spearman 1996, p173), and the execution part of MRP II was replaced with a finite capacity scheduling system, OPT. However, the problems of abstractness, remoteness, lack of transparency, and cumbersomeness (in terms of transaction requirements), that had already been noted with MRP, only became worse, and this approach faltered. At Kodak, they were initially satisfied with MRP as the execution system (however the adoption of this process was under duress), but when, for strategic reasons, they needed to reduce cycle times the approach became too transaction intensive. All companies eventually implemented at least some of the components of Lean Production.

2. **The formidable information requirements of MRP II:** Virtually all the case study participants complained about the cumbersome and essentially non-value-adding information processes of MRP II. Bendix and Kodak used aggregate product forecasting to reduce the task of master production scheduling (BCIA received schedules from their customers). Bendix and BCIA limited the number of control points in the routeing to reduce transaction entry. These moves compromised the performance of the model but were required by practical considerations. Bendix and Kodak had to process inventory transactions for component and transfer stores which were only required to support the MRP II model. All this expensive information processing is required by the MRP II approach to management of activity via the intermediary of a model or simulation of reality.

3. **The transparency of JIT:** The positive effect that the simplicity, visibility and transparency of Kanban had upon user acceptance was commented on in the three cases.

4. **The importance of the environment:** From the Bendix case it was seen what a pivotal role the physical layout had on a number of change initiatives. In one sense this is obvious - if the physical environment is structured to suite the task at hand things proceed more easily. On the other hand, within the perspective from which MRP II, say, approaches the problem of improving manufacturing performance, the environment in which work takes place does not enter the formulation of the problem or its solution. This is because it sees the problem as a formal, one, that is one that can be discussed independently of the situation in which it arises. Thus on this account, managing manufacturing operations is an essentially information-processing problem. The BCIA case corroborated the importance of changes to the physical environment.

5. **The importance of teams, culture and management structure:** All cases indicated the importance of human factors in implementing change, and particularly the role of cross-functional working teams and flattening of management structures, that is, breaking down the Taylorist notion of separation of thinking and doing. Note that culture and political structure are part of the environment of work in the broader sense, so this finding supports the importance of the environment in effective operations management.

6. **The role of diversity and evolution:** All the cases illustrated the nonlinear nature of change. Many approaches are often pursued at once, even conflicting ones. Ideas previously rejected surface again later and succeed. Ideological forces are influential, particularly the ideas of important people. Unexpected events have decisive consequences. All this contrasts dramatically with the "silver bullet" approach to change often described in simplistic success stories and consultant's tales.

4.5.2 Differences Among the Cases

The main difference among the cases is the different roles the companies see for MRP II in their overall system. They all agreed that shop floor operations should be decoupled from high level planning. Bendix retained MRP explosion for purchase order initiation only, but believed that could ultimately be run by Kanban according the full Lean Production model. BCIA retained it for purchasing, capacity planning, and works order initiation for the final assembly purely for accounting purposes. The Sales and Operations Planning concept was rejected as impractical by Bendix. This was not an issue for BCIA with the operation studied being driven by car company schedules. The S&OP process was strongly advocated, even with JIT execution systems, by the Kodak participants who were active APICS members. My access in this case study was not sufficient to determine whether this was just "party-line" (this is the standard APICS position) or that this company was really able to make this cumbersome process work.

4.5.3 Problems Raised

The rest of this thesis focusses on explaining the main observations from these cases. Firstly, there are considerable difficulties in implementing the seemingly plausible MRP II concept in practice. Secondly, this seems to be connected, on the one hand, with its idea that purposeful activity requires the building and maintaining models of the world, and on the other hand, to the failure of the theory to take into account the pivotal importance of the environment in an operations management system, that is, its failure to take account of the situatedness of activity. To do this it is necessary to consider in detail the nature of on-going, purposeful activity, and how conceptions of its nature influence the design of management systems.

Part 2

Theoretical Tools: Intentional Systems, the Activity
Level of Description, Theories of Activity

Chapter 5

Activity

In this and the next two chapters, I introduce my substantive theoretical tool: the idea of "activity" and of theories of the nature of on-going, purposeful activity. My main proposal is that management systems can be usefully viewed, at a certain level of description, as intentional systems. At this level of description the nature of activity is a natural question, and we can draw on a wide literature on intentional systems to obtain possible answers to this question. Adopting this level of description of management systems allows an important issue, namely their stand on the nature of activity, to be singled out from many other confounding issues. This allows the relationship between theories of activity and the design of management systems to be explored in part three. Thus, first I must define this level of description and the notion of an intentional system.

5.1 Levels of Description

There are many levels of description of the world. They include physics, chemistry, biology, psychology, sociology, anthropology, politics, economics, and ethics. Each level has its own set of fundamental questions: the laws of motion, the interaction of atoms of various elements, the nature of life, the nature of society. Each level has its own units of discourse, or "atomic" ideas: particle, atom, molecule, cell, mind, organisation. These descriptive units represent a dissection of the world at some scale or "grain-size". Each level has its own mode of describing the dynamics of the world at its grain-size, using its own descriptive units, and possibly its own notions of causation.

We often think of the levels of description as forming an hierarchy. This is because in many cases the fundamental units of description can be viewed as having some deeper structure and explanation at some other descriptive level, which we consider therefore, to be more fundamental or *lower* in a descriptive sense. However, a very important point, which has been made forcefully by Polanyi (Polanyi 1962), is that although the explanation of the atomic ideas at one level in terms of those at a lower level is satisfying in terms of observing the unity of the world, it is descriptively irrelevant in the sense that a consistent and complete theory can be constructed at one level without any knowledge of the lower levels. The relationship between thermodynamics and statistical mechanics is a good example. Thermodynamics was formulated as a complete and consistent theory before it was understood that heat was the motion of atoms or, in fact, before it was certain that matter was composed of atoms. The idea is familiar to computer people also: you can be a perfectly good Cobol programmer without knowing anything about assembler, machine language or binary logic, because the Cobol level of description of the program is quite independent of the lower levels upon which it is implemented. The examples are, however, a little misleading because in both cases there is a very precise relationship between the high level description and the low level description. This need not always be the case as will be seen later.

What this all means is that a level of description of the world can be chosen which is appropriate to the grain-size at which the phenomena are to be viewed, appropriate to the kind of questions about the phenomena for which answers are sought, and which will give insights about the phenomena which will not be available using other descriptive levels. In this chapter, I define a level of description of systems that I call the "activity level". I will be seeking to define the questions that can be asked about systems at this level, to define the particular scale or grain-size of this level, to define the units of description, and their relationship to each other and the types of explanations of the dynamics of the world at other levels. I want to assert that this level of description is available for all kinds of systems - mechanical,

biological, social, socio-technical or virtual (informational) and is distinct from other levels of description such as mechanics, electronics, biology, sociology, psychology, politics etc. Finally, I want to begin to justify discussing the behaviour and design of systems at this level.

5.2 Intentional Description of Systems

In order to define the activity level of description, it is simplest to consider first a purely mechanical system. Take for example a robot which is observed to be wandering about an office building picking up empty soft drink cans and taking them to a recycling bin. Being entirely mechanical in nature, what is happening in this scene can in principle be given a complete description in terms of the known laws of physics. In fact the classical laws would suffice since, although any solid state electronic components could require a quantum mechanical description, it would be possible, in principle, to replace such components with purely mechanical (say, clockwork) mechanisms. Each component of the robot in its internal workings, and in its interaction with objects in the office, is merely acting in accord with the known laws of physics, and the evolution and dynamics of the whole scene can be thought of as governed by deterministic laws and by conventional physical notions of causality. Taking this descriptive stance is to adopt a certain level of description, arguably the most fundamental or lowest level. I will can call this a *physicalist* level of description to emphasise that it is a description in terms of the lowest conceivably pertinent level of phenomena or "physics" of the system. Actually, this description in terms of the physics of rigid bodies is already a fairly high level physical description, higher than the atomic level, for example. The term "physics" in this thesis is used to denote the lowest sensible level of description for any given system domain.

In this description it is not strictly valid to speak of a robot and a room - these entities have not yet been differentiated. There are simply objects in interaction. The point about this description is that understanding the evolution and dynamics is, at least in principle, unproblematic. What happens in the scene evolves in an inevitable way from the initial conditions. It just so happens that under these dynamics "cans" happen to end up in "the bin". No further explanation is needed for this observation than that it follows inevitably from interactions governed by the laws of physics. More precisely, we would say, in dynamical systems theoretical language, that can fetching is a "stable attractor" of the system dynamics (see (Beer 1995) for an application of this approach to insect walking gaits). Taking an extreme reductionist view, we could describe a human preforming the same can collecting job in the same low level way, but given the controversial and intractable nature of such a description, it is helpful to stick to purely mechanical devices for this discussion. However, it would be less far-fetched to apply this level of description (in the form of a molecular interaction description) to the single cell organism *Escherichia Coli* seeking nutrients by detecting chemical gradients in its surroundings (Johnson-Laird 1993, p23).

Now the observer of the can fetching robot may take a different view of what is happening in the scene and seek a different explanation of the dynamics occurring. He/she might describe the scene as we introduced it: there is an entity recognisable as a robot moving in an environment which includes, among other things, soft drink cans and this robot is engaged in a goal-directed activity, namely seeking out empty cans and taking them to the recycling bin. This is a totally different kind (and level) of description which invokes the notion of an *agent* acting in an *environment* in an *intentional* (that is goal-directed) way. Now the dynamics of the scene also requires a different kind of explanation. We perceive the robot as engaged in an *activity*. The activity is perceived to have a *goal*. The dynamics of the scene are perceived in terms of the *actions* of the robot and their relationship to *situations* the robot finds itself in. An understanding of the scene now would require an explanation of how the actions of the robot are chosen in such a way that the goal is repeatedly and accurately achieved.

It is worthwhile to analyse in detail the conceptual steps in this descriptive shift. This may appear somewhat pedantic for our present case where, by construction, the intentional description is highly plausible. However, we want to apply the intentional description and the activity level of description to

a wide range of systems and environments where the prerequisites and consequences of such a description may not be so obvious. There are four distinct conceptual steps in this descriptive shift:

1. The robot (focal system) is distinguished as an entity in its own right;

2. This entity is ascribed agency. It is assumed to be doing something - to be engaged in some activity;

3. This agent is assumed to be autonomous within its environment;

4. The agent is ascribed intentionality. What the agent is doing is assumed to be directed toward the achievement of some goal.

The first step is to draw a conceptual line around those parts of this mechanical scene which can sensibly be viewed as part of an entity which will be called "the robot". The rest will be called the robot's "environment". This move is the familiar systems theoretical decomposition of the universe of discourse into a focal system and its environment. As usual in defining a system, if this device is to be useful, the boundary must be placed in such a way that the interaction between the system and the environment is relatively weak. This is not true for an arbitrary decomposition. In addition, if we are going to name the bounded system a "robot", that is, an instance of some well defined and familiar category, then the interactions amongst its components must be sufficiently constrained, regular and stereotypical for the focal system to deserve the name "robot".

The second move is to conceive of the entity as an agent, that is, to view it as *doing something* or engaged in some *activity* within its environment. This involves describing the dynamics of the scene in terms of a series of *actions* on the part of the robot "agent" which take place within the environment. This is a matter of detecting a certain level of structure in the dynamics of the scene, a level of structure with a grain-size larger than the simple physicalist description would employ. The behaviour of the putative agent can be analysed in terms of recognisable "chunks" which recur frequently enough for them to be recognised as such, and named. We want to be able to say things like "the robot approaches a can", "the robot grasps the can", "the robot tests if the can is empty", etc. Approaching, grasping, testing, are *actions*. Actions are the new atomic units of description of the dynamics of the scene. They are a new descriptive unit which cannot be simply associated with, or reduced to, the low level interactions of the physicalist description, for several reasons. Firstly, a simple identification with precise movements of components of the robot or parts of the environment would destroy the atomic status of actions. We would not be able to identify the grasping of the can on the table and the grasping of the can on top of the filing cabinet as being to instances of the same atomic action "grasp", for the simple reason that the mechanical details would be different in the two instances. Therefore, a simple identification of actions with low level *movements* would lead to multiple splitting of the units of the new description. Secondly, actions have a certain directedness that interactions do not: they have the agent as the subject and the environment as the object. They are largely defined in terms of what they are intended to achieve for the agent, which is not true of motions, which are simply defined in terms of spatial movements of objects. This further blurs a simple association between action and movement because the intent of an action can often be achieved by a variety of movement sequences (Collins 1990, p30). We are now describing the scene in terms of an entity that does something. Therefore, we not only have to recognise some stereotypical entity to be the actor but also some stereotypical actions as units of this description. We are gradually building up a series of new descriptive units through which to describe and explain the dynamics of the scene. When we have defined all these descriptive units we will have defined the activity theory level of description.

Thirdly, if we are going to describe an entity as being engaged in some activity it must be plausible to assume that its behaviour is more or less autonomous from the environment, rather than a slave of a deterministic set of mutual interactions. Given that in this and many other cases we know that the latter is true in principle, for this autonomy assumption to be convincing the system viewed as a robot interacting with its environment must be capable of displaying sufficiently arbitrary behaviour and sufficiently unpredictable (in practice) dependence on initial conditions. If this condition is not satisfied then the description of the systems as an agent is rather lame. The autonomy assumption also demands

new descriptive terms. In order to conceive of an agent that takes appropriate action in an environment from which it is autonomous it is necessary to posit its co-ordination with that environment as being the result of a new (higher level) type of relation to the environment. The agent must *sense* its relation to the environment and acquire *knowledge* about it.

Finally, the fourth move is to describe the action of the agent (the robot) in its environment (the room) in terms of achievement of a goal. The activity of the agent has some purpose: it is *about* something. That is, we describe the actions of the agent as *intentional*. The reason this step is now necessary is, having ascribed autonomy to the agent, we need an alternative account of why things happen, in other words, an alternative explanation for the dynamics of the scene. In the former description the dynamics were explained as the deterministic result of physical laws and initial conditions. We are now adopting a descriptive stance which says that the things we see happening make sense because the robot is trying to achieve a goal, namely, clearing the room of empty soft drink cans. This is what Dennett (Dennett 1978, p4; Dennett 1996, p27), the controversial US philosopher of mind, has characterised as taking an "intentional stance" to description - the robot is viewed as an autonomous agent acting intentionally in an environment.

Throughout this thesis I use the term "intentional" in Dennett's sense of an observer's attribution of purpose, aboutness, or goal-directedness to an agent.

> A particular thing is an intentional system only in relation to the strategies of someone who is trying to explain and predict its behaviour. (Dennett 1978, p3)

> Intentionality in the philosophical sense is just *aboutness*. Something exhibits intentionality if its competence is in some way about something else. (Dennett 1996, p35)

I recognise that to some, intentionality may connote much more that simply having a goal, including consciousness, and being responsible for the consequences of achieving the goal, and so forth. These are notions of intentionality as an *inherent* property of humans, which is quite a different notion to the asciptive one than I use here.

5.3 The Activity Level and Theories of Activity

We can now define the *activity level* of description. It is concerned with describing and explaining the goal-directed or intentional activity of autonomous agent systems. Its units of description are agents, environments, situations, goals, agent actions, environment sensing, knowledge, etc. The term *activity* itself stands for goal directed or intentional action of an agent. The question that the activity level of abstraction begs is: How do agents engage in intentional actions in pursuit of goals? Various schemes that attempt to answer this question in terms of the constructs of the activity level of description, I will refer to, following Agre (Agre and Chapman 1987) as, *theories of activity*. The relation between the activity level of description and theories of activity is somewhat like the relationship between kinematics and dynamics in mechanics. The activity level of description provides the phenomena, the descriptive units and the kind of questions we might ask about intentional agents: theories of activity explore the possible answers to these questions. They are about different explanations of the dynamics of purposeful activity of complex systems - how things happen. As Dennett puts it:

> We do not just reason about what intentional systems will do, we reason about how they will reason. (Dennett 1978, p243)

The fundamental question of activity theory is to explain how an agent can select actions in the present that are directed towards, or at least consistent with the achievement, of goals in the future without resorting to final causes (that is, causation by future, non-proximal events) in the explanation. We will see that different theories of activity are distinguished primarily by the way in which they seek to answer this question. In answering this question the various activity theories will have to take a stand on other sub-questions such as what is the nature of the world in which the agent acts, what is the

nature of sensing, what is the nature of action, what is a plausible architecture for an agent, what is the nature and use of representations, rules, and plans.

5.4 The Usefulness of the Activity Level of Description

Central to my argument is the idea that the activity level of description can be separated from other levels of description in a useful way in a range of systems that admit an intentional description in terms of goal-directed behaviour. Wherever an explanation in terms of goal attainment seems compelling as a descriptive device, there will be a set of issues to be given an explanation in terms of a theory of activity, which are quite separate from issues arising at other descriptive levels. Intentional activity can be attributed to mechanical and organic systems, single and multi-agent situations. The contention is that in all these systems it is possible to define, and fruitful to employ, the activity theory level of abstraction. The basic descriptive constructs - agents, environments, situations, goals, sensing, actions - will be useful in each setting, and the theoretical question - an explanation of the nature of agency - will recur.

The range of systems admitting an activity level description includes the biological (human and other life-forms), mechanical (robots), software systems (complex programs and software agents), socio-techinical systems (manufacturing and other management systems) and societies (organisations, tribes). Each of these domains has its own low level or physicalist descriptions (molecules, cells and nervous systems; mechanics and electronics; routines and data structures; people and machines; groups networks and institutions, respectively. In each domain there will also be multiple higher level descriptive schemes. For instance, in an organisation consisting of people engaged in some task, the physicalist level of description would be in terms of the behaviour of individuals and their interaction, the psychological level. There are several other possible levels of description however: there is the organisational level of analysis which discusses the forces that encourage the formation of collectivities; there is the sociological level which discusses the interaction between the individuals and the group; there is the cultural level which discusses collective sense making in groups; and there is the political level which discusses the role of power and authority in explaining what happens. Often ideas from several descriptive levels are mixed and used very loosely in discussions of management systems. The proposal that I wish to work through is that there is also an activity level of description, which separates as far as possible the discussion of how goal directed behaviour can be explained and implemented in socio-technical systems from these many other levels.

What can we hope to get from attention to the activity theory level of description? There are three answers to this question.

1. Analytical Clarity. The activity level of description is the appropriate level of description for discussing questions of goal-oriented activity in a wide range of systems. Therefore, a description of these systems at this level adds to clarity. In particular, in operations management for example, there are issues about the nature of management and design of management systems which are strictly dependent on the implicit or explicit theory of activity employed. Often these issues are obscured in discussions using other descriptive levels such as the sociological, psychological, or political level. Bateson (Bateson 1979) argues for the explication of levels of analysis, and that many fruitless arguments arise from mixing ideas from different levels. The idea here is that the appropriate choice of descriptive level clarifies the description of certain phenomena and allows the commonality of certain issues across diverse settings to be seen. This leads to the fruitful borrowing of insights and results from diverse fields. Further, in complex systems of many interacting parts intentional description may be more parsimonious than physicalist descriptions.

2. Explication. Each theory of activity makes assumptions about the nature of the world, the nature of agents, the nature of sensing and acting, and so forth. In management theorising, many of these assumptions are not explicit, just as the underlying theories of activity are often implicit or taken for granted. Often these assumptions are not tenable in certain domains. Explicating these assumptions and

checking there validity is a contribution that the activity level of description can make to analysis of purposeful activity in various system domains.

3. Design guidance. Analysing systems is one thing but there is also the challenge of designing them. In the design of purposeful agents, whether robots, software agents or management systems, theories of activity play a crucial role: they determine the architecture and methodology chosen for the design. Often the activity theory is implicit and it is expressed indirectly through the system design. The point here is that although the physicalist level of description is often satisfying as an ultimate explanation of dynamics and causality, especially in mechanical systems, it is often not at all useful in the design of complex systems. It is almost always necessary to break systems down into relatively independent modules simply to ensure tractability of the specification and design process. Generally, the physicalist approach does not suggest any such decomposition but the activity level of description does. In fact we will find that different activity theories suggest different decomposition and therefore different architectures for complex goal directed systems. I intend, in the last part of the thesis, to apply this idea to operations management by treating management systems as designs for autonomous, purposeful agents and showing how the underlying theories of activity informing the designs affect the principles and architecture of the system.

5.6 Methods

The intention of the next two chapters is to outline two principle theories of activity, which recur in different guises in a number of disciplinary areas including psychology, robotics, computer science, linguistics, sociology, anthropology and management. I present a statement of the principle lines of these two activity theories removed from the disciplinary jargon to make their universal character clear. However, I have to draw on these disciplines for examples. I will emphasise examples of two kinds, for two different purposes. Firstly, as I have in this chapter, I will emphasise mechanical examples because it then becomes clear that the problems encountered using specific activity theories to describe and design intentional systems have nothing at all to do with the unique condition of being human. Mechanical settings offer conditions that often conform closely to the more reductionist and mechanistic assumptions of certain activity theories, and yet they fail in characteristic ways even in these settings. Secondly, I will use examples concerned with the nature of human activity because, while there exists the possibility of confusing the issue, these psychological ideas are very much the metaphorical basis of theories of activity. After all, intentional description is essentially an anthropomorphic device. Theories of activity generally have their greatest force as explanations of one's own behaviour, and can be made plausible by one's own introspection.

5.7 Literature

Ludwig von Bertalanffy (von Bertalanffy 1950) realised that the systems / environment split which was useful in the physical sciences, biology and many other areas could be studied as a topic in itself. He saw that a similar set of questions was raised whenever a collection of interacting parts, with properties greater than those of the sum of the parts, could be discerned. Thus, he proposed a common system theory level of description for phenomena in a diverse range of disciplines (von Bertalanffy 1968).

Around the same time Wiener (Wiener 1961, orig. 1948) and Ashby (Ashby 1956) founded the study of cybernetics. Weiner's approach focussed on the role of feedback in systems in a range of disciplines that could be seen as control systems. He realised that the feedback mechanism provided a model of goal directed behaviour that preserved physical causality (Rosenblueth, Wiener and Bigelow 1943; Gardner 1985, p20). Ashby's school emphasised the role of an agent's information about the environment in maintaining a goal state, and made use of Shannon's (Shannon 1948) recently devised method for quantifying information. This school considered both error correction and anticipatory modes of goal attainment (Conant 1969; Conant 1970). Thus the cyberneticians discerned a

rudimentary form of an activity level of description. However, their emphasis was on one aspect of activity, namely maintenance of control. More recently, a movement (Maruyama 1963; Maturana and Varela 1980; Varela 1984; von Foerster 1984; von Foerster 1986 ;Varela, Thompson and Rosch 1991; Maturana and Varela 1992) going by the name of second order cybernetics has tried to break away from the input-output and feedback descriptions of systems characterised by the early cybernetics. These discussions can be viewed as being at the activity level.

In philosophy the study of the relationship between intentionality and action has long been a well defined area. My treatment of intentionality is similar to that of Dennett:

> Intentional systems are, by definition, all and only those entities whose behaviour is predictable / explicable from the intentional stance. Self-replicating macro molecules, thermostats, amoebas, plants, rats, bats, people and chess-playing computers are all intentional systems - some much more interesting than others. (Dennett 1996, p34)

However, note that this chapter goes further than simply distinguishing intentional and physicalist descriptions, by also defining a particular activity level of description and posing the question as to the possible range of theories of activity. The distinction between activity and the physicalist notion of movement, and the consequences of this distinction is very clearly discussed by Collins (Collins 1990).

In psychology, the cognitive approach using the computer metaphor of thought (Newell and Simon 1972; Anderson 1983; Johnson-Laird 1993), to the extent that this is viewed as a high level description of thought and action separate from the neurophysiological one, recognises an activity level, although with a commitment to one particular activity theory. (Ackoff and Emery 1972) is an early attempt to put intentional notions on a firm theoretical footing in behavioural science. In the "action" approach to sociology (Silverman 1970) the distinction that is made between action and behaviour is similar to the one I have made between the physicalist and activity levels of description. However, this approach stresses the need for actions to be interpreted (by other humans) and this seems to me to place it at a sociological level of description which emphasises social sense making. Argyris and Schon (Argyris and Schon 1980) use the term "theory of action" in their treatment of organisational learning, but their use the term is quite different to my use of "theory of activity". They mean the theory that the actor (person of organization) has, explicitly or implicitly, about the nature of activity. I refer to the theory that *informs the design* of the agent, that is the theory that the *designer* has about the nature of activity. The theory of activity that informs these author's view of organisational intentional activity is a cybernetic one. In robotics and computer science, agent theories (Agre 1995a; Woolridge and Jennings 1995b) multi-agent systems (O'Hare and Jennings 1996) and agent oriented programming (Shoham 1993) have seen intense research activity in the past ten years.

The notion of considering various alternative theories of activity in particular areas of intentional behaviour dates from Dreyfus (Dreyfus and Dreyfus 1989; Dreyfus 1992, orig.1972) (philosophy of artificial intelligence), Suchman (Suchman 1987) (human-computer interaction), Brooks (Brooks 1986; Brooks 1991b; Brooks 1991a) (robotics) and Agre (Agre and Chapman 1987; Agre 1988 ; Agre 1997) (human and artificial intelligence). Bourdieu's (Bourdieu 1977; Bourdieu 1990) notion of "practice" is at the activity level of social systems and his discussion of various conceptions of the reproduction of practices are theories of activity. The idea that theories of activity are a proper area of study of management systems had not been stated prior to my paper two (Johnston and Brennan 1996).

Chapter 6

Information Processing / Symbolic Theories of Activity

A good place to start an exposition of theories of activity is with a particular model of activity, the *planning model of activity*, that has been described by a number of authors (Suchman 1987; Agre 1988; Brooks 1991a; Johnston 1995; Johnston and Brennan 1996). Although this is not the simplest theory of purposeful activity and is not the earliest to be exploited in system design, it makes a good point of departure because it will most likely be familiar and be almost common sense to the reader. This is because, as was argued in paper two (Johnston and Brennan 1996), this theory of activity is rather ubiquitous in system engineering and management thinking as well as being the folk theory of human activity. A thorough deconstruction of the assumptions of the planning model will reveal most of the issues that are controversial in accounting for activity and most of the dimensions along which alternative theories of activity can be constructed. A rather formal definition of the planning model will be given first. A formal approach will allow the assumptions underlying the planning model to be more easily and completely explicated in a following section.

6.1 The Planning Model of Activity

The planning model of activity takes the stand that purposeful activity is achieved through the automated creation and implementation of plans. It envisions an information processing agent that senses its current state in the world, devises plans that will transform its current state into a goal state through the execution of actions, and hands this plan to an executor function for implementation. To achieve this, it must build and maintain an internal model or representation of the world and its current state in that world. The agent deals with the world at arms length, as it were, via the intermediary of its representation of the world.

6.1.1 The Agent and the World

In the planning model of activity the agent is supposed to deal with the world using an internal, abstract, symbolic model of the world. When the word "world" is used here it is of course referring to the "universe of discourse" which is appropriate to the particular problem space to which the planning model is to be applied. For example if the planning model is being used to describe a human having dinner, the world in question is the every day world of tables, cutlery plates and food. If it is used to describe an automated manufacturing system then the world consists of products, processes, customers, vendors, work centres etc. The world model is expressed using some system of symbols and relations. In artificial intelligence applications the symbol system is often some form of predicate logic. In computer science a system of data structures, variables and pointers may be used. In operations research the formal system is often a form of graph or network theory. In other applications the symbolic system may be no more formal than ordinary language and everyday logic.

The *world model* consists of symbols which stand for objects in the world, relationships between symbols, and certain formal rules for manipulating the symbols that are intended to capture the properties of the relationships between these objects. For instance, if the world consisted of blocks, the world model would have a symbol class, or variable, for each type of block (CUBE, PRISM, PYRAMID), a set of possible relations which may obtain between individual members of each class of

block, (ON, NEXT-TO, NEAR FAR) and a set of constraints on possible relations (NOT CUBE ON PYRAMID). These restrictions represent the actual constraints that apply in the real world on possible relationships between objects represented. (The conceptual world consisting of an agent moving blocks is known as the "blocks world" and has been widely used in artificial intelligence and robotics as a restricted environment for testing and comparing planning theories. I will illustrate many concepts here using this conceptual world because of its simplicity, and because I intend to use robotic systems as convenient examples of activity concepts).

A particular *state* of the world is obtained from the world model by enumerating the instances of the object classes that are present in the real world (CUBE-1, CUBE-2, PARAMID-1) and the relations that obtain between the instances (CUBE-1 ON CUBE-2, CUBE-3 NEAR PRISM-2). This set of symbolic statements about the objects in a particular state of the world is a *representation* of the state of the world in the formal world model. It is a representation of a particular kind which associates symbols with objects, and I will refer to it as a *symbol-object representation.* (We will see later that there are other modes of representation). Defining the state of the world by specifying the particular objects and their relation is called in artificial intelligence *symbol binding* and in computer science *variable binding*.

When the word *object* is used in this discussion, it is not meant to refer simply to *things*. It refers to what would be called "entities" in systems analysis (Yourdon and Constantine 1979). An object in this sense is an instance of class of entity that can be defined by the possession by its instances of a fixed set of context-free attributes. This then includes entities which we normally associate with things, such as blocks, cars, humans, but also more abstract kinds of entities like products, customers, customer requirements, work centres, manufacturing processes and debts.

Actions in the real world are viewed as a changes from one state of the world into another. On-going *activity* is a trajectory of successive changes of state. Therefore, the formal system includes a specification of a range of allowable transformations of the symbolic equivalents of the states of the world (such as PUT BLOCK2 ON BLOCK 1) which are the formal counterparts of actions. The whole formal system consisting of the possible combinations of symbols and the possible transformations among them (the world model) can be thought of as a kind of model of the "physics" of the world. (The origin of this world model is not a concern of the theory, which is a theory of ongoing activity rather than of learning.). The world model, in a very strict sense, is assumed to stand in a relation of *isomorphism* to the real world. That is, the elements of the formal systems correspond one-to-one with the elements of the world and the relations between either set of elements are also in one-to-one correspondence. The essence of the planning model is that the agent can determine how to act in the world by performing formal symbolic manipulation upon the world model. Or put another way, the structure that we see in its actions which justifies the view that it is engaged in an activity, is due to that structure also being present in the plan it is pursuing. The relationship of the agent to the world is depicted in figure 6.1.

World Model - System of Symbols, relations, and transformations.

AGENT

WORLD

World "Physics" - Objects, Properties, Object Relations, Dynamical Laws of the World

Figure 6.1: The Agent and the World.

This form of description of the world borrows heavily from physics where the idea of a state space has been extremely effective as a ground against which to formulate its laws of motion. The state of the material world is completely defined by the position and momentum of all particles that make up the world. The state space idea is easily extended to include fields of force, and to indeterminate quantum systems. The evolution of the world is then a trajectory in this multi-dimensional state space. The planning model thus borrows a theory of motion as a basis for a theory of action.

Two important states of the world are its *current state* and the *goal state* intended by the agent. In this approach, intentional activity is viewed as a process of deducing, using the world model and its formal rules, a series of transformations, which are the formal counterparts of actions, that will take the world from its present state to the goal state. This series of actions is the *plan*. Again, the symbolically represented current and goal states and the plan are assumed to stand in one-to-one correspondence with their equivalents in the real world as shown in figure 6.2. The attraction of this approach is the possibility of exploring many possible trajectories of intermediate states and finding one that is optimal, before committing to action.

Figure 6.2: Correspondences between Entities in the World and their Symbolic Equivalents used by the Agent

Figure 6.3: The Agent's Sensor and Executor Systems.

In order to realise this approach the agent must be able periodically to acquire data concerning the state of the world. In addition it must be able to carry out the actions that transform the world toward the intended state. Therefore, it must have access to a sensing subsystem and an executing subsystem as
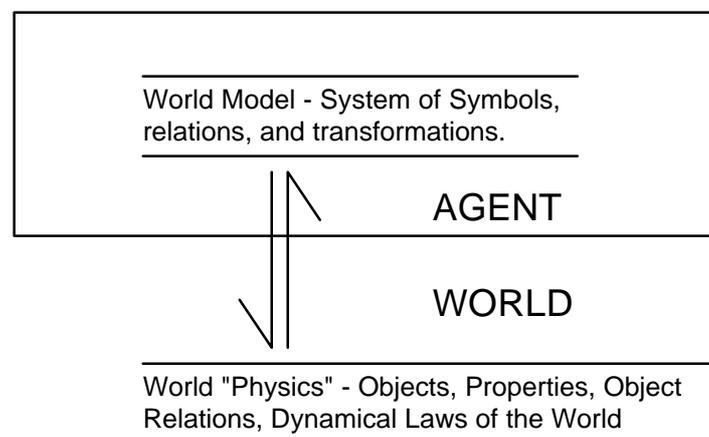
shown in figure 6.3. The sensing and executing systems are shown on the boundary of the agent to indicate that, although they have a symbol processing function which is internal to the agent, they also have a direct physical role in the world in which they act. They are in fact the agent's interfaces with the real world.

6.1.2 The Process of Planning

In paper two, I described the planning model using the Sense / Model / Plan / Act (SMPA) planning cycle (Brooks 1991a), familiar to management theorists, as well as those in artificial intelligence and other fields. The SMPA description concentrates on the four processes over which the agent has control. Here I include also the way in which the effects of the agent upon the world and the world upon the agent close the loop in the planning model of activity, so I will describe the dynamics of the planning model in terms of the processes Sense, Model, Plan, Execute, Effect, Respond. Intentional activity resulting in goal attainment is supposed, under the planning model, to be achieved through the sequential and repeated execution of this series of processes, each of which is conceived of as an input - output process.

Sense. The cycle begins with the agent receiving some sense data concerning the current state of the external world. The input to the Sense function is the physical effects of the world upon the agent and the output is sense data in the form of information about the current state of the world. What aspects of the world need to be sensed depends on the degree of detail (grain-size) of the world model.

Model. Since the world model is taken as given, this process is really one of using the sense data to determine the particular *current state* of the general world model. In other words, the objective is to determine from all the allowable or possible states of the world, which one the world is in at the present moment. This is the process of associating particular symbols with particular objects in the world and particular relationships which *may* hold among these symbols with the particular relations that *do* hold among these objects. Entities which in the world model refer to classes of objects (variables) are given a particular designation (bound) to particular instances present in the current world state. The inputs to the process are the world model and the current sense data and the output is the symbolic equivalent of the current world state.

Plan. Knowing the *current state* of the world, the desired *goal state*, and the allowable *transformations* of state that can be achieved by *actions*, it is now merely a formal problem to deduce a series of actions, the *plan*, that will take the world from its present state to the goal state with the greatest economy. In practice, the range of possibilities may be so great, and the formal manipulation so intricate, that compromises need to be made for the sake of tractability. There is a huge literature in artificial intelligence (Georgeff 1987a; Wilkins 1988), and operations research (Brandimarte and Villa 1995) on planning algorithms to perform this task. They amount to different ways of organising and controlling the search of the space of possible paths connecting the current state to the goal state. The inputs are the world model, the current state, the goal state and the output is the plan which is simply a list of actions, representing intended transformations of the world, in order of proposed executions.

Execute. The plan is now handed as input to an effector or executor function whose job is to execute it in sequence until such a time as it is handed a revised plan by the planner. Thus the plan is essentially a program of action like a computer program. The output of the execute process is a sequence of actions on the part of the agent. These actions are expected to transform the world, including the agent. into the desired goal state.

The Sense, Model and Plan processes described above could apply to any act of planning. In fact the description of the model to that point could apply to any form of analytical *reasoning* about action. What converts this description into a *theory of activity* is the assumption that implementation of the plan is sufficient to bring about action. Thus, the planning model of activity seeks to convert a theory of *analysing* activity into a theory of *performing* activity by automating the production and implementation of plans.

Effects: We can represent the effect that the agent's actions have on the world as an input-output process, but this time one which is a physical process (whereas, the previous processes were symbolic or information processing) and entirely out of the control of the agent. The inputs to this process are the actions of the agent and the given "physics" or laws of transformation of the world and the output is the resulting state of the world, which should be identical to the intended goal state. Due to unforeseen circumstances, inaccurate sensing, truncation of the planning process or unintended effects of action this, may not be the case. However this is catered for in the model by repeating the process as a loop.

Reaction: The change of state that the agent's action cause in the world in turn creates a response, or reaction, of the world upon the agent which can once again be viewed as an input - output process, but again a physical process not under the control of the agent. The inputs are the state of the world and the "physics" of the world and the outputs are the effects of physical interaction between the world, particularly upon the sensing components or functions of the agent.

It should be obvious that iteration of these steps will ensure the attainment of the goal by closing the gap between the goal state and the current state, that is, provided that a range of assumptions that are implicit in the model are valid in the application domain. The whole process including the inputs and outputs is shown in figure 6.4.
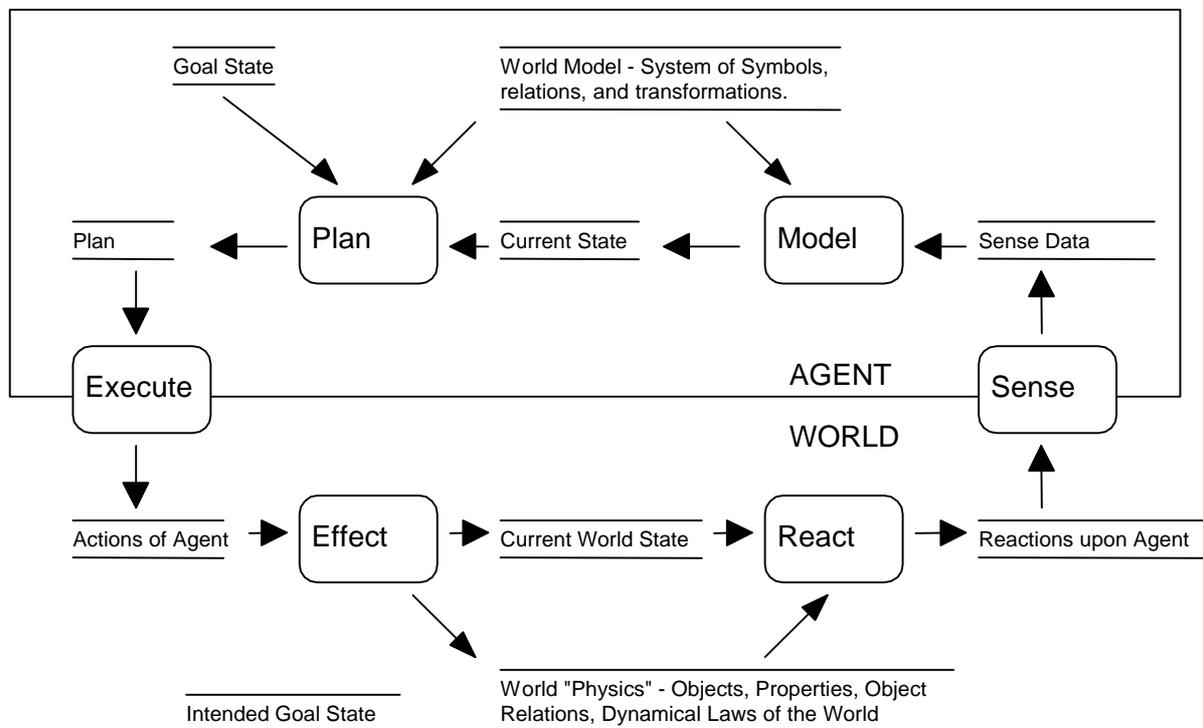
Figure 6.4: The Planning Model of Activity

6.1.3 Higher Functions of the Agent

To make this into a theory of *ongoing* activity*,* there needs to be a periodic revision of goals. Goals may be achieved or may be proved to be impossible and abandoned, so there needs to be a mechanism for generating new goals. These goals could be put in by hand by some means external to the agent or the process of goal generation may be automated by a higher level planning function that converts higher level persistent goals, often called "intentions" (Bratman 1987), into a series of new discrete goals as circumstances change. Examples of persistent goals in the business context would be "maximise return on investment" or "achieve at least a 90% customer service level". Clearly this only delays the question of where goals come from: in the planning model they are always ultimately external to the theory of activity. In principle, the planning process only needs to be repeated whenever

a new goal is formulated, but usually it is repeated more often than that to take account of change, inaccurate sensing and unintended effects.

6.2 Assumptions of the Planning Model of Activity

Having presented the planning model of activity in the previous section in a detailed and rather formal way, it is now a straight-forward matter to enunciate the underlying assumptions of the model. In this section, I outline these assumptions and postpone argument about their validity until we consider some examples of the model at work. This is convenient because it allows us to use the description of the examples to justify the correctness of the formal description of the model posed above and simultaneously to point out the appearance of the assumptions in the examples.

6.2.1 The Nature of the World

Objectivism. It is assumed that the world consists of objects and relations among them. For our purposes we can take an object to be an entity that maintains a constant identity with respect to a set of properties (Lakoff 1987). This assumption is necessitated by, and goes hand in hand with, the presumed feasibility of representing the world with a symbol - object representation scheme which is essential to the planning model.

Determinism. A high degree of continuity, predictability and causality is ascribed to the world as a result of the assumption that formal deduction can lead to a valid prediction of future world states. Actually the assumption is somewhat stronger than this. Deterministic systems can show chaotic behaviour which effectively disallows prediction of future states except over very short periods (Crutchfield et al. 1996). Such a situation would strongly limit the effectiveness of a simulation based theory of activity. So this assumption really is that the world is not only deterministic in principle but also predictable in practice.

Change. The Planning Model is intended to provide a theory of activity and therefore of change. However, the general world model itself is assumed to be unchanging. Change is modelled as a trajectory of different particular states that must be chosen from the set of possible states that comprises the world model. Therefore, the model presents a rather constrained treatment of change. In particular, the world model is not extended by knowledge gained through acting. It also assumes that the world state remains constant while each of the series of actions making up the plan is performed.

Boundedness. In modelling the world using a finite symbolic representation it is assumed that a neat boundary can be drawn in the real world such that things and events outside that boundary can be ignored at some definite level of approximation. The assumption is that we can work with a simplified model of the world, as demanded by practical or computational considerations, and that this simplification can still capture sufficient detail to provide a useful prescription for action. This assumption asserts the practical usefulness of cause-grained descriptions of the world. Although we regularly make use of coarse-grained description in *thinking about* the world, the planning model goes further by asserting that they are a reasonable basis for *acting* in the world.

6.2.2 The Relationship of the Agent to the World

Representation. It is assumed that action can, and must, be mediated via a centralised, symbol-object representation of the world. A symbol-object representation assigns to objects in the world, symbols having formal rules of manipulation in such a way that an isomorphism is created from the objects and their properties to the symbols and their rules of manipulation. Only in this way can the dynamics of the world be simulated in the planning process.

Unambiguous Representation. The model assumes that this representation is unique and unambiguous. In fact it is given from outside the model in an unspecified way. There is no room for multiple "world views".

Logicism. The image of action in this model is of the agent constantly solving the problem of acting from first principles. The model assumes that rational analysis is the appropriate foundation for a theory of action. It presumes that the world is itself structured logically and is thus amenable to logical manipulation of a deductive kind. Philosophical logic is the exemplary symbol manipulation language and predicate calculus is essentially a simulation language within logic. Consequently, many of the more formal instances of the planning model which appear in the robotics and the autonomous agent literature, use philosophical logic as their formal system, and practical implementations often use programming languages derived from and suitable to the implementation of formal logic rules of manipulation.

Causality: As stated in the previous chapter, different theories of activity take different approaches to the problem of how it is possible for an agent to act in the present, in such a way that current actions are determined by, or at least consistent with, future goals. The approach that the planning model takes is to invoke the internal symbolic representation as part of the causal chain of action. The represented current state and goal state together with the world model allow the deduction of the plan, which then becomes the *cause* of action. Each step in the plan production and execution is seen as an input-output process, as illustrated in figure 6.4, and therefore, the normal temporal order of cause and effect is preserved throughout the process. However, many of the inputs in this cycle are internal symbolic states of the agent, so if they are to be thought of as causative in the usual sense, they must also be considered to be *real*. Thus the planning model asserts the existence and necessity of real internal agent states representing its relationship to its environment as the causal force of action. Only in this way does the sequence of goal setting, sensing, state updating, planning and action preserve a normal forward-time causal chain.

In a robot or software agent, representations are ultimately expressed as bit positions in some set of registers. As a theory of human action the planning model asserts the reality of *mental* states. As a theory of organisational activity it asserts the necessity of publicly agreed-upon physical representation of the groups intentions and present state, which would often take the form of records in a computer system, jotting on a white board or more informal but still physically existent forms.

6.2.3 The Nature of Action

The Locus of Agency. In the planning model the entity from which purposeful activity springs, the proximate cause of activity, is the plan, which is assumed to contain and be the source of the order in the behaviour of the agent. The plan is derived from focal system's reasoning about the environment, goal states and current states, and is therefore the "intellectual property" of the focal system. It is this information processing ability that is the ultimate cause of intentional behaviour. Agency is thus located very much within the focal agent and is directed *from* the agent *toward* the environment. As we have seen in the previous chapter, intentional description attributes a high degree of autonomy to the agent, so the simple identification of agency with the focal system amounts to assuming a high degree of autonomy for the focal system over its environment. The planning agent achieves this autonomy through its assumed ability to simulate and influence its trajectory through the environment. Once underway, the actions of the plan depend (viewed at the physicalist level of description) on a strong interaction between the agent an its environment for their success. However, the states that arise through this interaction have already been *anticipated* in the course of planning. It is in this sense that the plan, rather that the agent's interaction with the environment, can be thought of as causing action. This is how, in this model, the autonomy of the agent (viewed at the activity level of description) is asserted. The *effect* of the action is already present in the *thought of the action* in this view, so *thinking about acting* and *acting* are essentially the same thing (Johnston 1996b; Johnston and Brennan 1996).

As Agre has pointed out (Agre 1997), once you adopt this planning model of action, and I argue that ubiquitously we tacitly do just that, it is hard to imaging the possibility of an action without a preceding plan.

This idea that intentional activity arises from the information processing of the agent leads to a certain ambiguity about the material status of the agent. To the extent that the agent as a part of the world figures in its own calculations about future states, the theory has a notion of the focal agent as a physical thing, made of the same stuff as the world. On the other hand as figure 6.4 depicts by drawing a boundary around what is internal to the agent and what is in the world, the theory demands that there is a "place" outside the world where these symbolic manipulations on representations of the world out there take place. If this interior of the agent were part of the world, it would have to figure in its own calculations leading to an homuncular infinite regress (von Foerster 1986). When the planning model describes human intentional action this internal place is the mind, and the assertion of the necessity of such a non-material construct as the locus of agency in human activity is simply the dualism of Descartes and Locke. The existence of a place outside the world where the "thinking" takes place is perhaps not so problematic in mechanical systems such as robots, where information processing is more easily viewed as an abstract process taking place on the concrete medium of binary switches. Nevertheless, the existence of this duality between the symbol processing domain and the actual world, has lead Agre (Agre 1997) to label all instances of this type of theory "mentalistic".

Role of the environment: Because interaction with the environment has been written out of the story by the agent's ability to simulate it, the environment has no role in this theory in the explanation of intentionality. The world gives no assistance to the actor: it is merely a passive backdrop against which activity is played out. Any structure the world has is already present in the plan. There is no recognition that certain types of structure in the world, possibly created by the past actions of this and/or other agents, can assist in purposeful action.

Transparency of Sensing. The planning model assumes that using sense data to determine the current state is a straight-forward matter. The model assumes that the process of gathering data about the state of the world is unproblematic. Data is always incomplete and noisy. Additionally, there is always a selection or interpretation involved in the choice of data determined by the design of the sensing instrument. But further to these well known data sensing problems, is the much more difficult problem that raw sense data is generally not in the appropriate *symbolic* form to be of direct use to the world model. The planning model does not recognise that there is a translation required at this level, one that is essentially knowledge dependant. The mapping from sense data to the world state is generally under-constrained and requires knowledge of what *might be* in the world to effect. In artificial intelligence, this problem is known as the "symbol binding" problem. The planning model assumes that the identification of (binding of) internal symbols (variables) with (to) external objects is a simple matter.

Transparency of Effects. Closely related to the assumption that sensing is a simple process is an assumption that translating a plan into action is a simple process. The model assumes that handing a plan consisting of a set of atomic action instructions to a "dumb" executor system is sufficient to ensure that the task will be accomplished. There is an implicit stance being taken here on the nature of action, namely, that action is essentially informational: once the symbolic plan is generated the job is essentially accomplished. There is no recognition of the role of sensing in acting, for instance. In practice, it is always necessary to confirm or disconfirm, by sensing, that an action has completed successfully, since initiating the action may produce effects not anticipated in the plan. Furthermore, the atomic actions generated by plan manipulation are coarse-grained descriptions and, as argued in the previous chapter, do not correspond in any simple way to low level interactions among components part of the agent and the world. They are, in essential ways, of a different logical kind to the motions required to effect them, and again a formidable translation must take place in the executor system, a translation that is again knowledge dependant and also probably sense dependent. An atomic action may be "move forward one metre" whereas to effect this action in a robot it must be translated to voltage - time patterns to be applied to various motors or in an organism, into nerve excitation patterns. Again, in the operations context an atomic instruction such as "weld part X to part Y" may draw on,

and in fact presuppose, a great deal of interpretation, improvisation, and tacit knowledge on the part of the operator. That this problem is not immediately apparent to proponents of this theory of activity is presumably partly because it is modelled on the physics theory of motion. Motion and action are units from quite different levels of description and a complex translation is necessary to go between the two. The translation from sense data to symbolic world state is one of *abstraction*, and that from symbolic actions to movements is one of *concretisation*.

6.2.4 Agent Architecture and Resources

Serialism. The last two points are particular cases of a deeper assumption that it is feasible to separate sensing, planning, and acting into discrete, sequential and linear tasks which are independent of the outcome of each other. Figure 6.4 presumes such a separation when it depicts activity as a repeating loop of these sub-processes. This assumption places quite severe restriction upon the possible architecture of the proposed agent. It *must* consist of modules which perform these task sequentially, and this implies a functional division of labour and serial processing.

Unlimited Resources. Finally there is an important assumption that the computational or planning resources available to the agent are unlimited. This is a significant assumption because there is a tendency to see increasing the complexity of the world model as a way around the limitations of many of the above assumptions. For instance, one might be tempted to reduce the grain-size of the world model to increase its fidelity and get round problems with the boundedness of the world model. In this way the previous assumptions are likely to be misinterpreted as mere practical limitations rather than difficulties in principle. However, it is known that a large class of planning problems require an exponential increase in computational effort as the level of detail of the model is increased, rendering them essentially intractable (Chapman 1988). Thus, unlimited resources is an illusory notion because a workable theory of activity must always guarantee that the agent can act *now*. Heuristic shortcuts generally do not solve this problem either, because it is quite possible that terminating the deductive process short of exhaustion might cause wildly inappropriate actions to be selected. A theory of on-going activity must be able to guarantee the generation of appropriate actions under extreme time constraint.

6.3 Case Studies of Implementations of the Planning Model

The approach of this thesis is to go on to criticise the planning model, as applied to management systems, on the grounds that the above assumptions are generally impossible to subscribe to, especially in complex environments. The reader will already see that some of these assumptions will be dubious in social system contexts. However, it might be reasonably expected that assumptions such as objectivism, determinism, logicism, unique representation, and so forth, would be good in purely mechanical systems. It may come as a surprise that a rigorous implementation of the planning model *even in mechanical systems* of moderate complexity is problematical. As Agre puts it:

> mentalism is not simply an inadequate description of people, but an untenable way of life for any creature in a world of any complexity. (Agre 1997, p151)

This is illustrated by the case study of Shakey the robot. This case was also used by Suchman (Suchman 1987) with similar intent, but the presentation given here undertakes a far more thorough deconstruction of the assumptions underlying Shakey's design. The purpose of this case study is firstly to show that our formal specification of the planning model indeed captures the essence of a certain extreme position on how autonomous, purposeful activity can and should be implemented in a system, a position that seems to be practically the default mode for the engineering mind set. Secondly, examining how these implementations fail, throws the validity of the assumptions into sharp relief and paves the way for a re-examination of these assumptions leading to alternative theories of activity.

The second case study describes how the planning model has been taken up as a model of human intentional activity. At least since the Enlightenment, the idea that mental representations of the world determine action has been the folk psychological theory of human activity. However, this idea has achieved scientific respectability since the computer has been available as a metaphor for the mind. I introduce this example here not because I want to discuss humans as intentional beings, but because this theory of human activity is the metaphorical source of the planning model: planning model theories of action in any domain can be thought of as deriving from this image of human *deliberative action.*

6.3.1 Shakey the Robot

Shakey was designed and enhanced over a number of years in the late 1960s at the Stanford Research Institute and was one of the first attempts to build an autonomous mobile robot that could perform specified tasks, using the planning model (Raphael 1976). Shakey sensed its world using a black and white television camera (plus a bump bar for collision detection) and moved in its world using motorised wheels. An off-board computer performed the modelling and planning tasks. Its world was static and specially contrived: it consisted of rooms with plane rectilinear walls and various blocks. Its tasks were mainly navigational, but it could also identify blocks and perform actions on them such as pushing them.

Shakey's general world model consisted of a list of possible propositions in first order predicate logic about rooms, walls, doors, objects, robots and their spatial relationships. States of the world were represented as sub-lists of propositions that held in that state. Actions were modelled as operations that changed the list of propositions for one state into the list for another state. Shakey's planning system (Fikes and Nilson 1971), the famous STRIPS (STanford Research Institute Problem Solver), was influenced by two preoccupations of early artificial intelligence research: automated theorem proving; and problem solving. Thus Shakey's goals were set up as a list of propositions describing a desired relation between the robot and its world. If these propositions could be proved as theorems from the propositions representing its present state, Shakey had achieved its goal. If not, the system sought and applied operators (actions) that changed the world state to another closer (in some well defined sense) to the goal. This new state was then a sub-goal of the planner and the process was repeated until a path from the present state to the goal state was found. At each stage, primitive actions such as MOVE and TURN were collected in a list that became an action plan that was handed to the motion sub-system. Figure 6.5 illustrates how Shakey's design is an implementation of the planning model.

Thus, Shakey's planner embodied the logicism assumption, that the world is a problem to be solved, in a very literal way. Many of the other assumptions about the world, and the relationship of the robot to it, are almost as explicit. The choice of a world consisting of blocks has a long history in artificial intelligence research and relies on the boundedness assumption to be taken seriously: it is assumed that a kind of self contained micro-world can be used for demonstration purposes and the results will be generalisable to the real world in a straight forward way. Even for this simple, formal world this assumption has proved hard to maintain (Brooks 1991a p2; Dreyfus 1992).

Shakey's main scientific contribution was to show that many tasks that are trivial for humans, such as perceiving, interpreting perceptions, deciding what to do next and executing planned actions, were extremely problematical to design into a machine. The vision sub-system had to take a scanned television image and hand to the planner as a set of propositions about objects in the world. The motor-control sub-system had to take high level atomic action commands and convert them into wheel motions which the planner assumed would be completed successfully. These are both translation problems that are now known to be extremely difficult. The former, due partly to the mapping from two dimensions to three dimensions being under-constrained (Marr 1982, p267), could only be accomplished in highly contrived environments and using a knowledge base that was put in by hand.
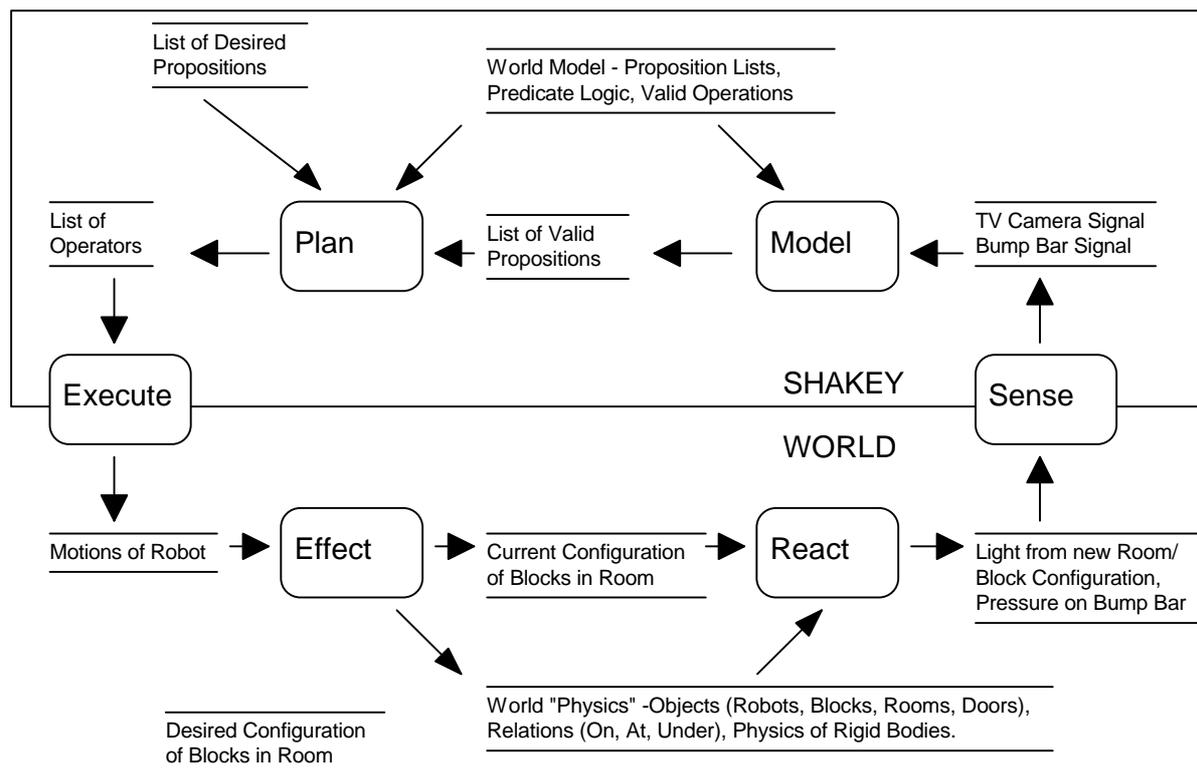
Figure 6.5: How Shakey Implements the Planning Model of Activity

Although the ability of Shakey to perceive objects, plan actions, and carry them out was considered impressive, its performance was also excruciatingly slow. Despite being connected to one of the most powerful computers of its time, it took hours to accomplish its tasks. It behaviour was brittle even in its contrived environment:

> An entire run of Shakey involved the robot going into a room, finding a block, being asked to move the block over to the top of the platform, pushing a wedge up against the platform, rolling up the ramp, and pushing the block up. Shakey never did that in one complete sequence. It did it in several independent attempts, each with a high probability of failure. You were able to put together a movie that had all the pieces in it, but it was really flaky. (Hans Moravec quoted in (Crevier 1993, p115))

Furthermore, it was completely unable to perform in ambient environments and, according to Brooks (Brooks 1991a p10), this has continued to be a problem for later generation autonomous robots designed on the planning model.

In no way do I want to appear to be critical of the quality of the research project that Shakey represents. Shakey's design is extremely ingenious, particularly the STRIPS planner and the PLANEX executor, which were the point of departure for much later work on planning and autonomous robots. On the contrary, the ingenuity and the apparent plausibility of Shakey's architecture when coupled with its virtual inability to act in any realistic kind of environment, represent an important part in my argument. The question that Shakey begs is: Why does such a plausible notion of activity fail in any moderately dynamic and complex world?

6.3.2 Human Activity

Shakey can be considered also as a "proof of concept" project for a parallel research effort to understand human intentional behaviour, which came to be known as "cognitive psychology" or more generally cognitive science (Gardner 1985; Johnson-Laird 1993). The key idea was that human cognition, and by implication human activity, could be viewed at a certain level of abstraction as

information processing or computation upon a symbolic mental representation of the world. This idea achieved its most formal expression in the influential Physical Symbol System Hypothesis of Newell and Simon (Newell and Simon 1981), which states that a system capable of manipulating physical symbols is a necessary and sufficient condition for intelligent behaviour. The underlying metaphor is the mind, located in the head, as a serial computer, manipulating symbols that stand for objects and relations in the world outside. It must be emphasise that this was proposed as a high level description of cognition. It was not intended to describe the neurological processes which underlie thought, but rather to assert that a consistent and complete description of the nature of thought with explanatory power can be set up using the methods and terms of computer science, which are quite natural to a description of symbol manipulations (note that the earliest code breaking computers manipulated lexical symbols rather than arithmetic ones). Once one assumes that the world can be represented in this symbol - object mode and that thought and action consist in symbol manipulation, this level of description emerges almost inevitably.

This view of human deliberative activity draws on the planning model of activity as figure 6.6 shows. Knowledge about the world is stored in so called long term memory. Goals are stored in short term or working memory. Sense data is received and stored briefly in specialised short term memory. Based on sense data, goals and world knowledge, reasoning takes place using short term memory to store intermediate sub-goals. In this version of general purpose deliberative activity, the reasoning is envisioned as a means-ends process (Newell and Simon 1963; Newell and Simon 1972) similar to that described in the discussion of Shakey. Reasoning results in the formulation of plans that are implemented by the various motor systems at a "sub-conscious" level.
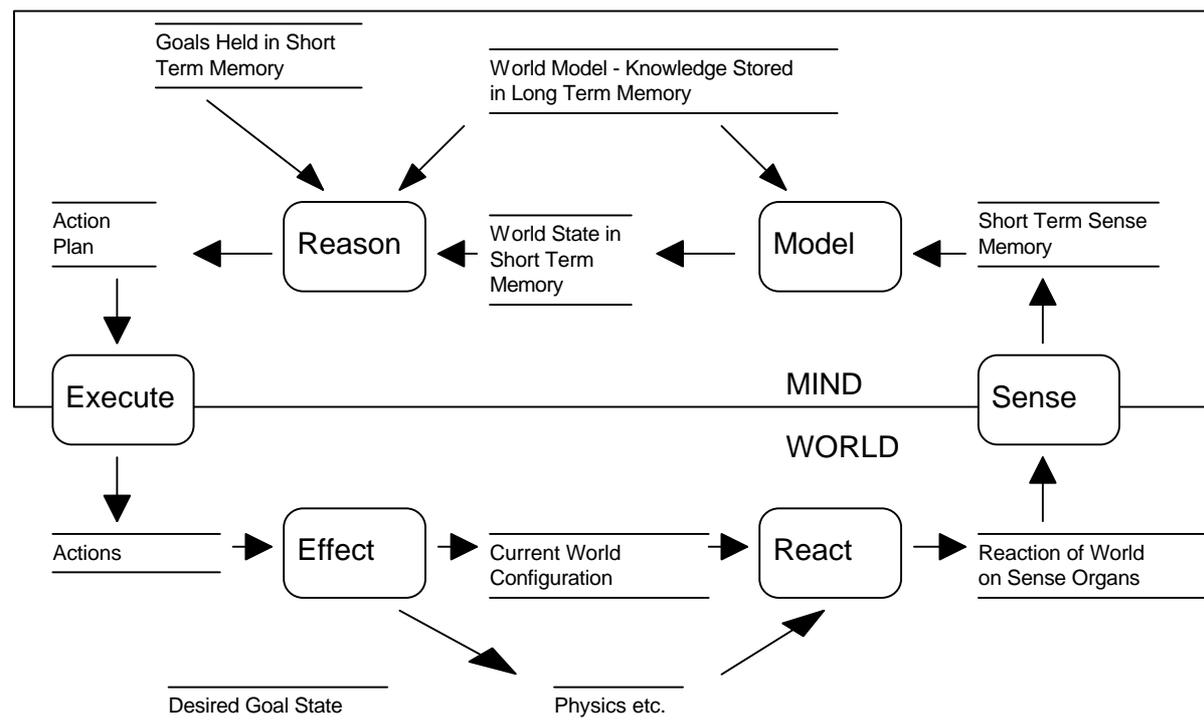
Figure 6.6: How the Cognitivist View of Human Activity Implements the Planning Model

Much of this literature concentrates on intellectual activities such as chess and other game playing (Chase and Simon 1973; Simon 1973), problem solving (Newell and Simon 1972), etc. It is implicitly assumed that these tasks are models for human activity in general, which is to take the stand that deliberative thought precedes and essentially *is* activity. It is this strong separation of thought and action together with this characteristic causal assumption which is the striking feature of the planning model of activity. With some modifications, this general model of the relationship of information processing and action is also assumed in the (dominant) cognitivist view of manual skills such as

sporting and artistic skills (Newell and Rosenbloom 1981; Ericsson and Smith 1991; Proctor and Dutta 1995; Ericsson 1996b).

In recent decades, this model of human activity has come under increasing criticism (Suchman 1987; Dreyfus 1992, orig. 1972; Agre 1997; Clancey 1997). These criticisms include the logical difficulties of grounding human intelligent behaviour on symbolic declarative and procedural knowledge structures, problems with the logicist approach such as the "frame problem" (the problem of bounding the deductive work required to predict the consequences of any action (Pylyshyn 1988)), the "qualification problem" (the difficulty of specifying all possible qualifications of a propositional statement (Crevier 1993, p120)), problems with accounting for the human ability to deal with ambiguity and relevance, and the difficulty of accounting for routines, tacit knowledge and style. I will not dwell on these criticisms: for my purposes it is sufficient to note that severe problems are already manifested in mechanical implementations of the model, such as Shakey, so it is important to check the validity of each assumption when this model is applied to management systems in chapter 9.

6.4 Variations on the Planning Model

There are two important variations on the planning model of activity. The first is simpler than the planning model and, as a formal theory, predates it. The second has been influential in accounts of skilled or expert activity.

6.4.1 Cybernetic Theories of Activity

This first variation assumes that goal directed behaviour can be accomplished through simple feedback mechanisms and it is actually a precursor to the symbolic information processing theories described above. We can call this approach the "cybernetic model of activity" and it arose out of second world war research on regulation devices such as automatic gun aiming devices, and on the role of information in regulation processes (Gardner 1985, p20). The cybernetic model of activity was the first proposal for an all purpose mechanism for goal-directed activity, being intended to apply equally to engineering systems such as regulators and servo-mechanisms, biological processes such as homoeostasis, management systems and some aspects of human cognitive activity.
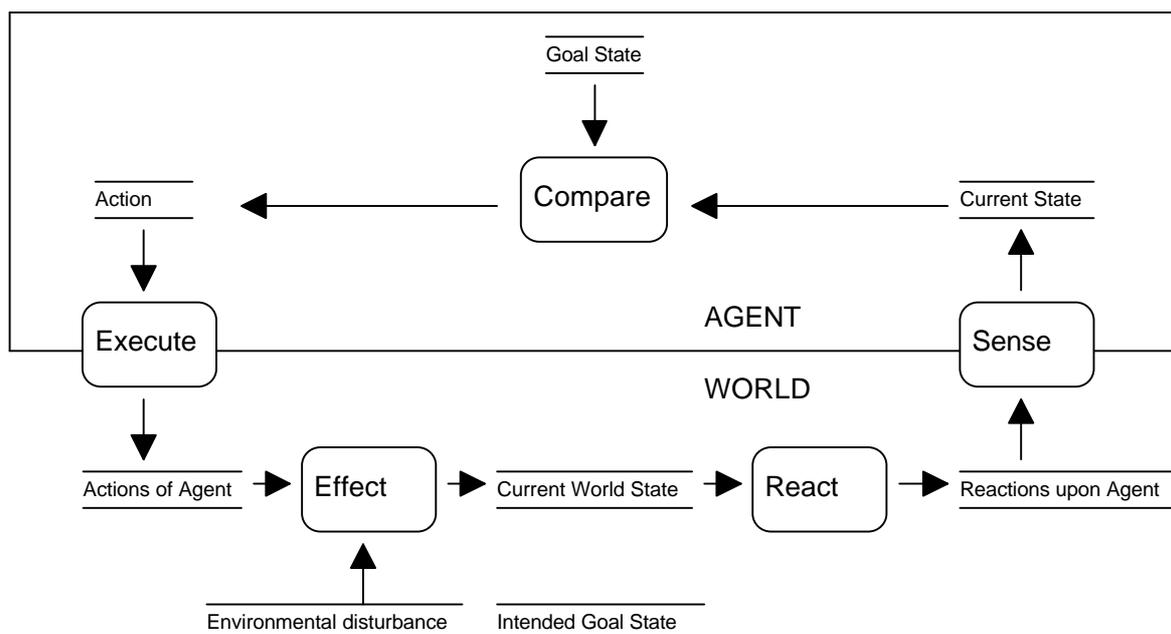

Figure 6.7: The Cybernetic Model of Activity as a Simplified Version of the Planning Model

The cybernetic model of activity differs from the planning model by assuming that both the current world state and the goal state can be represented by particular values of, in its simplest form, one real variable, or in its more advanced control theory elaboration, several possibly vector variables. The nature of the world is also seen in rather simplistic terms, being static except for unpredictable disturbances. The purpose of the agent is to take actions to counter the effects of these disturbances. An appropriate action, which is often conceived in binary terms such "switch on(off) the power to the motor" is determined by the difference between the goal state value and the current state value of this state variable, that is, an error signal. The cybernetic model is only applicable to the persistent type of goals. The planning model can be thought of as an extension of the cybernetic model to include discrete, future goals, and a complex deterministic world, and although the early cyberneticians foresaw this possibility in the form of "feed forward" systems (Ashby 1956; Conant 1969; Conant 1970), they lacked the full blown computational metaphor through which to implement this idea. Figure 6.7 illustrates the cybernetic model of activity. Specific implementation of this model include thermostats, the fly-ball governor, servo-mechanisms, periodic review inventory policies, and theories of homoeostasis in biology.

6.4.2 Rule-Based Theories of Activity

Another variation on the planning model sees the information processing part of the agent not as an all purpose reasoner, but as a domain specific rule-based reasoning system, or so-called "production system". A production systems consists of a set of rules of the type: IF <condition> THEN <action> which are held in a memory of some kind and can be retrieved efficiently according to their level of appropriateness to the situation prevailing. The conditions can refer to the current state, the goal state or to intermediate goals determined in the reasoning process, which are held in short term memory. The actions can be atomic actions which are again collected into a plan, or intermediate goals which are stored temporarily in working memory, which is generally conceived of as a stack. Given a goal state and the current state, the production systems will automatically "fire" rules, possibly producing and manipulating intermediate goals until a plan is built from the series of atomic output actions.
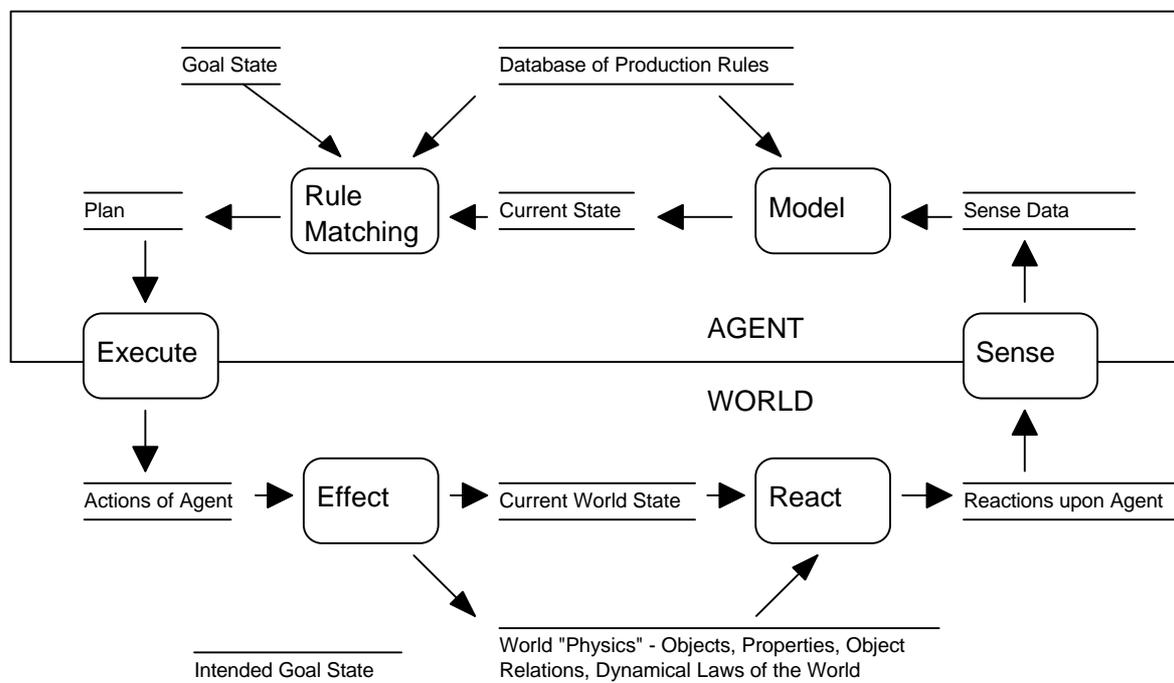

Figure 6.8: Production Systems as a Variation on the Planning Model

I will call this theory of action the "rule-based model" of action. This scheme is illustrated in figure 6.8. This model of action is very much a formalisation of the folk theory of procedural activity (as opposed

to deliberative activity) as the application of rules. Specific implementations of this theory are found in cognitive psychological theories of expertise and skill (Anderson 1983; Newell 1990), rule based expert systems (Crevier 1993. p153), and Chomsky's approach to grammatical sentence production (Crevier 1993. p151).The only difference between this theory of activity and the planning model is in the way the plan is produced from the goal state and the current state, and the computational metaphorical difference is akin to the notion of "compilation". The planning model assumes that knowledge is held in terms of generally applicable restrictions on possible world states (so called *declarative* knowledge (Anderson 1983)), that is, in the form of the "physics" of the world from which plans can be derived for any arbitrary goal and current situation, whereas, the rule-based model assumes that knowledge has been compiled into domain specific *procedural* form and under this model the agent cannot function outside the domain for which the knowledge has been collected in this form. The most sophisticated variants on the planning model assume the existence of both types of knowledge and reasoning (Anderson 1983; Newell 1990; Anderson 1993). There is the same commitment to the role of symbol-object representation of the world because, for the rule system to be finite, the conditions must refer to the same sort of symbolic *variables* which are given *particular* values by the sensing process.

6.4.3 Hybrid Planning Theories

Early work on planning systems concentrated on mental activities such as problem solving and game playing. The output from these programs were action plans which were not actually implemented in the real world. This was considered a trivial problem, and given that the task environments of these programs were so constrained as to be essentially symbolic in nature, this was probably true. However when these techniques were applied to mobile robots moving in even highly contrived environments, the error free implementation of the action plan was found to be far from trivial. Errors in sensing, truncation of the planning process, and failure of motor systems to reliably implement atomic actions in real situations, meant that the robot could rapidly find itself in a situation quite different from the one anticipated during plan production, and this could lead to totally inappropriate actions. It was quickly realised that performance would be greatly improved if the executor system had sufficient autonomy to at least detect failure of the plan, and preferably to do something about it. Shakey's STRIPS planner and PLANEX execution systems already anticipated this need (Fikes, Hart and Nilssen 1972; Agre 1997, p150). The STRIPS program not only handed the executor the plan, but also a summary of the preconditions for each action. PLANEX had sufficient autonomy and reasoning power to omit steps if the preconditions for a later move had already been accomplished, or to call for re-planning if the preconditions for no action obtained.

This theme of hybrid architectures with high level planning systems converting goals to plans using state space searching, coupled to low level execution systems with limited reasoning power or reactive behaviour handling contingencies has continued to be popular in autonomous agent research. This architecture, in keeping with the *modus operandi* of the planning model still asserts the primacy of reasoning about goals in determining action. To extend this approach to highly contingent environments one must either give the executor as much reasoning power and access to sense data as the planner, or re-plan after every action, both of which options are generally unfeasible (Agre 1997, p154).

6.5 Literature

The planning model of activity is the folk theory of human activity of most western people today and its fundamental premise, the separation of a machine like body and an ethereal mind (or in more contemporary terms an information processing mind), which drives the body by reasoning about an abstract representation of a world "out-there", is very much part of our taken-for-granted cultural background assumptions. In Western philosophy these ideas start to be clearly articulated by Descartes (see (Agre 1995b, p3)) and Locke (see (Gardner 1985, p54)). The intellectual status of this notion took a down turn at the beginning of this century at the hands of the behaviourists (Gardner 1985, p12;

Johnson-Laird 1993, p17). They argued in positivist vein, that since mental states were not observable, they were not suitable candidates for a causal explanation of activity in terms of goals. Consequently intentional explanations of activity relied on final causes, were unscientific, and to be avoided. They put forward what is essentially a physicalist alternative in terms of the response of an organism as a function of outside stimulus which made no reference to internal states. Complex behaviour resulted from complex chains of stimulus-response (S-R) rather than intentional activity. Essentially they saw an organism as a physical "particle" with no structure and with a rather odd, but nevertheless determinate, set of properties. This was not an intentional theory of agency and only had the weakest notion of an organism as an agent at all.

A way out of this causality bind was provided by the discovery of the feedback mechanism. In biology is was recognised that the maintenance of desirable states in organisms had a common explanation, homoeostasis (Cannon 1929; Cannon 1933), and about the same time, it was realised that in engineering the problem of control could be studied separately from what was being controlled (Hazen 1934). Weiner realised that these were examples of a common feedback mechanism and suggested the same mechanism could explain human intentional behaviour with respectable causality (Rosenblueth, Wiener and Bigelow 1943; Wiener 1961).

Gardner (Gardner 1985), argues that it was the appearance of the digital computer around the same time as a number of critiques of the explanatory power of behaviourism, notably by Lashley, that resulted in the reinstatement of the reality and causal status of mental states. The computer provided a metaphor for mental processing with physically real internal state, but also provide a new research method based on the simulation of mental capabilities. It is an interesting problem for the cognitive approach that nowadays neurophysiological discoveries have made it very difficult to conceive of the brain as a digital computer, and the kind of serial processing assumed in the planning model must be seen more as a high level description of the functioning of thought (ie, the mind rather than the brain). This status for mental states threatens to undermine their causal status once again, and we find cognitivists clinging to the assertion that the symbols of the representation correspond in some unknown way to actual physical (neuronal) states of the brain (Vera and Simon 1993, p9).

The mental part of the planning model is the basis for Newell and Simon's General Problem Solver program (Newell and Simon 1963), first run in 1957, for simulating problem solving using heuristic search, which influenced the development of the information processing view of the mind in psychology. Agre (Agre 1988) asserts planning as a description of mental activity dates from the book by Miller, Galanter and Pribram (Miller, Galanter and Pribram 1960). The idea that activity is the conversion of sense data to symbolic representation and then to action was first stated by Craik (Craik 1943; Craik 1948) in the context of human activity, and as a model of mechanical activity it is first fully worked in the Shakey project begun in the early 1960's (Raphael 1976). This is also the approximate date for the invention of Material Requirements Planning. Gardner (Gardner 1985) has argued that the idea of representation-mediated activity emerged as an important idea in all the cognitive sciences around this time. Thus the rise of the planning notion of activity is closely associated both metaphorically and directly with the invention of the digital computer, and is a remarkably recent idea despite its almost common sense status nowadays. Although the term planning occurs in management writing earlier than this, it is often used merely to distinguish *goal setting* from control which concerns *goal achievement* (Hofstede 1978).

Production systems were introduced as a model for intelligent behaviour by Newell and Simon (Newell and Simon 1972), and were strongly promoted for this purpose by Anderson (Anderson 1983; Anderson 1993) and Newell (Newell 1990).

The literature of post-Shakey planning research is too vast to review here (see (Wilkins 1988; Woolridge and Jennings 1995a)), but some examples with possible relevance to management systems are: the Phoenix forest fire planning simulation system (Cohen et al. 1989) which gave fire-fighting teams sufficient reactive autonomy to get out of unexpected danger, while still operating within an overall planning system; the experiments with Tileworld, an abstracted environment with tuneable

uncertainty which has been used to investigate the trade of between re-planning and execution time reasoning (Pollack and Ringuette 1990; Pollack 1992), Schoppers (Schoppers 1987) proposal to create, using traditional planning methods, highly conditional *universal plans* the appropriate portions of which can be selected reactively, and various proposals to use libraries of plans reactively (Firby 1987; Georgeff 1987b).

The planning model as a distinct theory of activity which could be compared to other theories was first described by (Suchman 1987), Brooks (Brooks 1986; Brooks 1991b; Brooks 1991a) and Agre (Agre and Chapman 1987; Agre 1988 ;Agre and Chapman 1990). As a theory of intelligence it had been articulated and strongly criticised by Dreyfus (Dreyfus and Dreyfus 1989; Dreyfus 1992, orig. 1972) and others (Winograd and Flores 1986; Clancey 1993). Similar, perhaps less explicitly articulated representational notions of activity were being criticised around the same time in a number of other fields: in biology (Maturana and Varela 1980; Varela 1984; von Foerster 1986; Varela, Thompson and Rosch 1991; Clancey 1993); in education (Lave 1988; Greeno 1989); in linguistics (Lakoff 1987); and anthropology (Bourdieu 1977; Bourdieu 1990). I criticised the planning model as a basis for operations management in papers one and two (Johnston 1995; Johnston and Brennan 1996).

Chapter 7

Situational / Interactional Theories of Activity

I now introduce an alternative approach to explaining intentional activity, which rejects the primary causative role of information processing on symbolic representations, that is planning, for action. Instead it proposes that intentional activity is primarily due to the interaction between the actions of agents in response to situations they experience, with the structure of the environment in which they act. This theory of activity has its routes in the European phenomenological philosophical and psychological tradition, but has only recently influenced workers involved in computational theories of agency, and seems to have had practically no influence in management theory.

7.1 The Situatedness of Agency

In chapter 5, I motivated the intentional view of the can-collecting scene by noting that an outside observer might give a more compelling account of what is happening than that provided by the physicalist view, by dissecting the scene into agent and environment, and attributing autonomy and intentions to the agent. In this section, I argue that the particular character of the planning model as a theory of activity stems from a failure to distinguish between what this outside observer can know about the world and what this hypothetical agent can know. The planning model of activity describes how an outside observer would design an agent based on his or her account of what is going on in the scene. However, the outside observer has a privileged position, what Agre calls the "aerial view" (Agre 1995a, p11), because he or she is not part of the scene being observed. This transcendental observer has ready access to information about the scene which is not readily available to the agent from its "ground view" by virtue of its being *situated* in the scene with which it wants to interact. There is a particular kind error involved here, which we might dub the "theoretician's error", which consists in assuming the that an observer's way of giving a *theoretical account* of the order in a scene, in terms of representations and plans, is also a valid account of how that order is *implemented* by the agent. Most of the problems with the planning model arise from attempting to design a (necessarily) embodied and situated agent as though it were able to make use of the transcendental point of view.

For a genuine outside observer of the can-collecting scene, it is a simple matter to work out what the robot should do next to achieve its goal. This is because the outsider can readily perceive the relations between various objects in the room and the robot, by virtue of his/her outside or aerial position. Furthermore, being acquainted (possibly in a tacit way) with the physics of the world he or she can see at a flash what the possible consequences of various moves might be, which things are clearly irrelevant, and can easily construct a plan of action, using this unobstructed view of the scene and her knowledge of the way things work in the world. And this can be done by the observer just as easily for any goal the robot is assumed to have, and any configuration the objects in the room assume. It is this wonderfully simple all-purpose plan building ability that the planning model assumes can be implemented in the robot.

Consider what it means to build an agent that simulates this transcendental outsider's view. Firstly, all the outsiders knowledge of the world has to be packed up into a neat symbolic form so that it can be transferred form the observer's head to the agent's memory banks, as it were. This necessitates gross simplifications, both in terms of the level of detail which can be effectively captured and in the kind of aspects of the observer's knowledge that are feasible to represent symbolically. It is this attempt to

model the world in a way independent of any particular agent's view, and independent of any particular agent's goals, which forces the objectivist world view on a designer implementing the planning model.

Secondly, consider the difficulty that the robot has in simulating the true outsiders view of the world. By virtue of is position in the room, that is, its being part of the scene it is trying to observe in a way that the true outside observer is not, the view the robot gets through any feasible array of sensing devices, is a very skewed and incomplete view of the scene from which to reconstruct an aerial view. Furthermore, its *embodiment* as a physical object, that is, its being of the same kind as the things beings observed, implies that there is certain information to which it cannot have access, but which is readily available to the true outside observer. If it is equipped with vision, it probably cannot see behind itself and certainly cannot see what is behind opaque objects. Even its scale is important: the everyday human world looks quite different to a two metre tall robot than it does to one a centimetre tall. These are not limitation for the transcendental outside observer. For the robot, the translation from the ground view to the aerial view is severely under-constrained, and so to complete its construction of the world around it, it must draw on its knowledge of what *might* be there, knowledge that suffers the limitation that it is incomplete and in a more abstracted form than available to the genuine outside observer. The reconstruction of the world state therefore involves considerable extrapolation and more importantly, interpretation or translation. The agent must translate sense data from its *situated* and *embodied* "ground" point of view to the "aerial" point of view, which becomes a knowledge dependent process, contrary to the assumption in the planning model that sensing is transparent. Furthermore, the robot's reasoning about what to do is much more laborious than that of the genuine outside observer because the robot cannot easily eliminate what is irrelevant. The outside observer's privileged aerial position allows him/her to see that most objects in the room are of no consequence to the can-collecting enterprise and can be ignored.

It is precisely this attempt to implement an outsiders approach to action in an embodied and situated agent that is the conceptual error of the planning model. This error has been pointed out by several authors, notably Clancey (Clancey 1991; Clancey 1993) Agre (Agre 1995a), Maturana and Varela (Maturana and Varela 1980) and Bourdieu (Bourdieu 1990), (see also (Slezac 1992)). The consequences of this confusion between agent-centred and transcendental points of view is easy to see for the case of a physical agent navigating a physical space as described above. But the same argument applies to all representational theories of activity. All agents are *situated* in their environment, whether this is an environment of space and time, or a social environment of a person within an organisation, or even an informational environment of a software agent, and this environment is always experienced from a point of view. An agent's *situatedness* always limits and skews its knowledge of the world. Similarly, all real agents are embodied. The limitations of having a body are rather obvious for individuals navigating a spatial environment, but again all real agents are in a similar sense limited by the fact that they are made of the same stuff as the world, are subject to the same laws of "physics", which limits their ability to have abstract knowledge of the world and to perform actions in the world.

One important bodily limitation of all real agents is bounded computational ability in the face of time constraint. The transcendental outside observer, not being material, does not have this constraint and so all manner of exhaustive search for plausible action seems possible. Besides, the outside observer is constructing an *account* of an agent's activity for the purposes of predicting or explaining its actions, and is not required to do this in "real-time". Real agents have to act *now*. The agent is "thrown" into the action rather, as a soldier who is thrown into battle and whose main concern is survival. Perhaps the only agents that do not suffer these problems of situatedness, embodiment, and thrownness are software simulations of agents, and it is precisely for this reason that Brooks (Brooks 1991b) has argued against simulations as a way of making theoretical progress on real agent design. This point should also be taken by those who use simulation to study real management systems.

I argue that this same theoretical error can also be detected in the widely held Cartesian notion that thinking and acting are much the same thing. Thinking about action is precisely the attempt, by mental simulation, to stand outside oneself and take a transcendental, totalized view, in order to contemplate and simulate possible courses of action, and humans do this, on occasions, with difficulty. But to do

this on a continual and repeated basis, in order to decide what to do next, is much too imprecise, complicated and time consuming to be a plausible explanation of on-going intentional activity. Mental models always fall short of capturing the complexity of the world (all blocks are either on top of other blocks or not on top). This is not a great problem when thinking about what to do when actions are simulated in an imaginary world of analysis. But *real* action has to be implemented in the *real* world where the details left out of the model, for analytical convenience, often play a crucial part in what actually happens (blocks teetering on the edge of another fall when bumped). *It is one thing to think about doing something and quite another to do it.*

Here are three ways of stating the shortcoming of the planning model as an explanation of on-going goal directed activity, which I believe to be entirely equivalent:

1. The planning model makes the "theoretician's error" of assuming an outsiders account of action yields a good design for an intentional agent;

2. The planning model fails to recognise the situatedness, embodiment, and thrownness ("situatedness" for short) of real agents.

3. The planning model takes a good theory of thinking about action and tries to use it as a theory of the on-going performance of action.

As I have argued in chapter 5, the notion of intentionality arises when an outside observer ascribes goal-directedness to the behaviour of a system, and it is useful in explaining, predicting and designing the actions of agents. So if we want to take an intentional stance on systems we must give some legitimacy to the outsiders view. The idea of goal-directedness is especially attractive in discussing management systems, so the question is: Can we create a theory of activity and a design for an agent which is motivated by intentionality but does not commit the error of ascribing to the agent the powers of the transcendent outside observer? Or another way: Can we develop a theory of intentionality for a situated agent? And while doing this we will have to avoid slipping back to a physicalist level of description which is inappropriate for a discussion of intentionality, and not particularly useful for guiding the design of complex systems.

7.2 Situated Activity

Let us return to the can-collecting scene, maintaining an intentional stance to its description, which means retaining the descriptive position that there is an agent which is effectively autonomous, engaged in an activity in an environment which it can sense, and the activity is directed toward the goal of moving empty cans to the recycling bin. These as before are an outside observer's ascriptions. Now, rather than basing an account of how it achieves its goals upon the assumption that it somehow imitates the way in which this outside observer might conceive of accomplishing it, we seek an alternative account in terms of what can be directly sensed and known by the robot. This is done by considering what the task might seem like as experienced by the robot. This obviously anthropomorphic device is purely rhetorical: in a moment it will become apparent that the mode of description developed here can be applied to non-human systems also.

In this new description, the agent encounters a succession of *situations* of which it is a part. It receives sense information which can be interpreted to some degree as indicating certain relationships between it and some of the objects in the room. However, only with great difficulty, if at all, can this sense data be used to define the position of *all* objects in the room including itself. Certain characteristic actions are performed when particular situations are detected and certain new situations ensue as a result of the combined effect of these actions and the "physics" of the world. Characteristic actions are again performed in the resulting situation and so on, until the situation in which the robot has just dropped a can in the recycling bin occurs.

So far, this is merely a description of the scene from an unusual point of view. There is no attempt in the above paragraph to explain *how* appropriate actions are selected in particular situations or to provide a reason (causal explanation of) why the final situation should occur at all. To provide those extra elements is to propose a *theory of activity*. But the new descriptive point of view does *suggest* a new theory by pointing to the resources that the agent can most reliably deploy in directing activity, namely, the relation of things to itself. The idea suggests itself that much of the conceptual difficulty of the planning model, which arises from the need for the agent deal symbolically with the world, could be avoided if the primary causal link in activity is taken to be between *situations* and *actions* rather than between *representations of transformations of the world* (that is, *plans*) and *actions*. The reason for this optimism is that the possibility exists for the agent to retrieve information about the situation more or less directly from sense data without having to perform the difficult transformation to the aerial point of view, using the complex symbol processing apparatus which this necessitates. Along with this change in the primary causal link, comes a change in the status of planning from central in the planning model to peripheral in this new theory of activity, which I will refer to, following Suchman (Suchman 1987), as the "situated action" model.

7.3 The Situated Action Model

In the situated action model, certain activity level units of description are retained: goals, sensing, actions, and knowledge. However, certain units which will now be seen as specific to the planning model theory of activity (rather than the activity level *per se*) are discarded, modified or played down. These include world models, representations (at least of the symbol - object kind), plans and reasoning. These are replaced by some new units of description and it is important to remember that these also are not easily associated directly with the physicalist level of description. These include *situations*, *attention* (relevance selection or focus) and *interaction* with the environment. Just as the planning model of activity was presented as an extreme position of reasoning-based action which has subsequently been modified in various ways to take account of the need for reactivity, so the model I now present is an extreme position that may need to be weakened (especially by the known relevance of planning and reasoning to human and possibly mechanical activity).

The situated activity model states that it is possible to implement intentional activity by having actions determined by situations the agent encounters. The basic idea is to implement a tight loop between sensing, acting and interaction with the world without the intervening steps of world modelling and reasoning. It is hard to give a formal statement of the situated action model, partly because there is no precise agreement between its authors (as explicitly stated by several authors (Agre 1993; Suchman 1993; Vera and Simon 1993) in the debate in the journal Cognitive Science on Situated Action), and also because the key concepts come from an unfamiliar way of thinking about the world. It is hard to draw a diagram of the processing envisage in the model because non-sequential and non-infomation processing approaches do not lend themselves to simple graphical representation. The situated action model does not give a precise recipe for implementing an agent design at the physicalist level, which indeed it should not as an activity level theory. Rather, it gives a high level specification of the kind of causal connections that need to implemented to avoid the difficulties of the planning model. My approach in this section is to first describe the key concepts of the model and then present an admittedly vague statement of the model.

7.3.1 Situations

To appreciate what is being claimed in the situated action model of activity, it is important to state what is meant in the theory by the term "situation". We use the word in everyday talk to mean "a state of affairs". More specifically, it describes a set of relations that obtain between a focal agent, usually ourself, and the world around it. The notion of a world state, as formulated in the planning model, purports to describe the relations that exist at a particular time among all the objects that make up the

world, including the agent. A situation differs in a number of ways from this notion of a world state. The first difference is the point of view. The world state notion is an aerial view of the state of affairs at some time, whereas the idea of a situation is a description of the state of affairs from the focal agent's *site*. An agent is *in* a situation, not simply *part of* it. Situations are reports on the world perceived from the agent's situated and embodied perspective. Secondly, a situation is not an inventory of everything in the world but a report on the relation to the agent of just those things or aspects of the world that can be directly perceived from the agent's situated point of view. However, we should not think of situations as simply snap-shots of the world taken from the agent's position. What counts as part of a situation also depends on the agent's goals or intentions. Aspects of the world which have no direct or potential relevance to the agent's goals are not considered to be part of the situation of the agent at this moment. If I am intending to take public transport to the city and I see a tram passing the end of my street then this observation is part of the situation which determines my next action: I head to the railway station instead of the tram stop. If my intention is to take a walk then trams are not relevant and I am unlikely to pay attention to or even notice the passing of a tram. The state of the weather (bright or overcast) on the other hand will be crucial to setting off on a walk, whereas for the trip to the city, unless it the weather is very bad, it is not likely to have any causal significance for my action. Thus a situation is neither totally "out there" or "in here". A situation is a function of what is happening in the world *and* what the agent is making of it. This goal dependency of situations means that in two universes two agents in the same world state can be in different situations. Conversely, two agents could be in the same situation in two different world states since objects and relations peripheral to the agent's goal may differ arbitrarily.

This is a difficult but fundamental point of difference in the way in which the world is described in these two theories of activity. The notion of a world state as used in the planning model is an attempt to describe the world in a way which is both independent of the goals of any particular agent inhabiting the world, and independent of any agent's situated and embodied point of view. It attempts to provide a description of the world equally valid for all agents with all possible goals. It is a notion of the world borrowed from Newtonian mechanics, where the concept of intentionality of the particles of the world does not arise. Having taken this position on the world, a *particular* agent in the planning model is now forced into performing complex symbolic manipulations upon this representation of the world in a way that reflects *its* goals, before the representation is useful for *its* actions. The whole symbol processing apparatus, and all its attendant problems, is forced on the theory by the adoption of an agent-independent, goal-free (that is, objective) representation of the world. By direct contrast, the notion of a situation is a description of the world in relation to an agent's intentions and situated point of view. In fact its intentions are part of what we term its situatedness. This mode of description of the world already has the intentionality of a particular agent built in to it, and therefore goals enter the causal picture at an early stage in this approach.

7.3.2 Relevance and Attention

If what counts as a situation depends on the agent's intentions, it follows that what things in the environment need to be considered in defining the current situation is very much less than all the objects in the world and all their relations, since only those that could potentially play a direct role in activity in the near future need be considered. Thus, the intentional sensitiveness of situations brings with it a notion of *relevance,* and this restriction of relevant relations of things to the agent brings with it a great reduction in the information processing overhead in relating action to situations, compared to the simulation based approach of the planning model. In the planning model, what is relevant is only bounded at the world modelling level, when the grain-size and boundaries of the universe of discourse are defined. Agre (Agre 1988) has argued that the computational explosion implied by the need, in the absence of an agent centred-notion of relevance, to bind variables for every object in the world is an overwhelming problem in implementing the planning model as a basis of activity in a time constrained world. Dreyfus (Dreyfus 1992, orig. 1972, p258), echoing Wittgenstein and others, has gone further

than this and argued that the lack of any concept of relevance in planning-model-type theories makes them entirely intractable because *anything* has the potential to be relevant. He argues that it is only because we are part of a situation, and have familiarity with situations that others might be part of, that we can deal efficiently and flexibly with questions of relevance.

The notion that what counts as a situation is a function of the agent's goals can be expressed another way: this kind of theory of activity must include a notion that the agent is paying *attention* to, or is *focusing* upon, certain features of the world, or more precisely, the sense data coming from its sensing system. The notion of attention or focus is a concept with no equivalent in the planning model. Because the world model is agent-independent and goal-free there can be no selection of relevance until the moment of computation when the world model is processed for the particular agent with particular goals, that is, well after perception.

7.3.3 Ontology

Finally, since it is not necessary in this stance to describe the things in the world in a way independent of, or equally valid to, all potential observers, it is not necessary to assume the world being experienced is made up only of objects with context-free properties. This position is forced on the planning model because only an object so defined can be represented by a symbol which has the same meaning to all possible agents. In the situated action view, there may be objects in the world which would be detected through the invariance of their properties over different situations and over time, and also possibly in the role they play in communicating and coordinating action with other agents (von Foerster 1984), which would be taken to be objects. But there could also be features of the situation, such as the threatening nature of a certain street corner, which have significance only for the focal agent and the given situation, but which may yet have a causal status for action. The situated action model is thus not committed to an objectivist ontology.

There is an important point to grasp here, however. The design of an intentional agent implies a choice of ontology to be used in building the agent, that is, a choice of the primitive notions that will be used *by the agent* to deal with *its* world. (For that matter it involves also a choice of epistemology - notions of how the agent can become aware of the world). This ontology does not have to correspond to the designer's own ontological assumptions about *the* world. You can hold the view that the world really consists of objective entities such as blocks and chairs, and still choose to design an agent that does not require that ontological assumption, because of the difficulty entailed in implementing that assumption in the agent. Many of the influential writers on situated action, such as Suchman and Dreyfus, were discussing the *essential* nature of *human* intentionality, and associated themselves to various degrees with an ontological position that asserts the fundamental socially constructed nature of the world (Garfinkel 1967; Collins 1981; von Glasersfeld 1984). This position is often unacceptable and alien to people from a scientific-engineering background, and tends to make these writing seem for them somewhat mystical (Hayes, Ford and Agnew 1994). When we are talking about the design of agents with the *ascriptive* type of intentionality used here, such a radical position is not necessary.

For Agre, the minimum ontology, which he terms a *deictic* ontology, required by a situated agent is one that assumes the existence of objects in the world only to the extent that they relate to the focal agent:

> If an agent has an intentional relationship to an *entity* then as far as the agent is concerned the latter is defined entirely in terms of the role it plays in the agent's activities. (Agre 1997, p243)

7.3.4 Representation

One of the key claims of the situated action model is that intentional activity need not be mediated by symbol-object type representation, with the agent building and revising a model of the world "out there" in its mind "in here". Brooks (Brooks 1991b; Brooks 1991a) has taken the strongest position on representation and claimed that intelligent behaviour, by which he means intentional activity, can be

exhibited by an agent which uses no representation at all, and encapsulated this notion in the slogan that "the world is its own best representation". Taking this hard line opens the way for philosophical debate which is not productive for my project (see (Slezac 1992)). It can be argued, for instance, that the curvature of the bi-metal element in a thermostat is a representation of the temperature of the room it is controlling. The intentional stance brings with it the presumption not only that the agent has goals ("desires") but also that agent "knows" in some sense its environment ("beliefs") (Dennett 1978, p3), and knowledge must be represented in some way. There are two point that can steer us away from this potential quagmire:

1. In putting up alternatives to the planning model we are seeking to avoid a causal role for *centralised*, that is all encompassing representations of the world. We are not concerned about fragmentary representations that can be said to exist in components of the agent's architecture.

2. It is the *type* of representation employed by the planning model, namely using symbols to stand for objects without reference to agents or agent goals, which leads to problems. I have called these symbol-object representations.

Agre has pointed out that representation need not be of the symbol-object type. An alternative is what he has termed at different times "deictic" (Agre 1997, p 241) or "indexical-functional" (Agre and Chapman 1987) representation. Indexicality is a term borrowed from linguistics. Indexical terms, such as "*here*", "*now*", "*my car*", "*the day after tomorrow*" are those that require a knowledge of the location and epoch of the speaker to allow their objective meaning to be determined. In other words they are descriptions of features of the world from the speaker's situated point of view. The "functional" aspect of this type of representation means that things in the world are represented in relation also to speaker's goals and projects, as in "*the cup I am drinking from now*". This expression represents a cup not only as a function of the location and time of my writing, but also in terms of my purpose for pointing it out as significant. Indexical-functional representations are agent-centred, fragmentary and purpose-laden, and are therefore, quite different to objective representations such as "cup-9 at 3 Pearson St. Richmond on the 24th of June 1998". They are ideal for representing situations or aspects of situations when necessary.

7.3.5 Sensing and Acting

In the situated action model, actions are supposed to be triggered as a direct response to situations. Many authors take this to mean that situated activity is purely reactive. A purely reactive system is one where actions are selected by a fixed mapping, or look-up table, from inputs to outputs. The agent is not suppose to posses internal state. This means that the entire system is not provided with any internal mechanism to hold a memory of previous states and therefore cannot determine actions as a function of a sequence of states. This strong position is often taken in order to set up the strongest possible opposition to the planning approach, since clearly, if there is no memory structure in the agent's architecture with which to store symbols representing objects in the outside world, the agent can hardly be conceived as employing a symbol-processing approach. However, equally it does not follow that the existence of internal state or memory implies the use of symbols in the sense that is meant in the planning model.

This direct action does not mean that sense data is not processed in any way in the situated action model. We will see that in the mechanistic implementations described below, that the designers have provided processes to select and filter sense data in situation-dependent ways. Implicitly, this pre-processing is also goal-dependent, although this point is often obscured by the fact that these systems are usually designed with a single goal in mind. This pre-processing generally takes the form of identifying features in the environment relating to external objects and taking actions that are a function of what features are detected. But in keeping with the stance of only using the resources directly available to the agent to drive action, objects in the world should only enter processing as relations to the focal agent, and through their direct significance to the agent rather than as abstract object classes.

These implementations of the model have included conflict resolution processes to deal with the conflicting action implication of various features detected. Thus, there is no simple look-up table or rule-like relation between situations and actions. The real point is not that there should be an unmediated link between the sense systems and the executor systems, but that any mediation should not take the form of symbol processing on a global, aerial view representation of the world and the agent's current and intended position in it. Rather, if aspects of the situation need to be represented for computational purposes the mode of representation should be deictic, as defined in the previous section. Thus, there is a certain degree of latitude for advocates of the situated action model can back away from this hard reactive position. The definition of situation used here is probably general enough to include at least some traces of past situations and goals as part of the current situation.

To implement a direct look-up table type relation between sense data and motor activity would be to implement a physicalist description of *movement,* rather than a high level intentional theory of *action*. It would be a return to the stimulus-response notion of activity which is a physicalist rather than intentional theory. Critics of situated action from the cognitivist tradition (Vera and Simon 1993; Hayes, Ford and Agnew 1994) have been very wary of a return to behaviourism, which is not surprising given their struggle to overthrow that approach. However, it is a mistake to identify situated action and S-R theories, because in situated action, sense data is assumed to be processed in a situation-dependent, and therefore goal-dependent way, and the notion of a situation that is causal to action is a high level abstraction at the activity level, unlike stimulus, which is at the physicalist level.

It is also a mistake to identify situated action with production systems as did Vera and Simon (Vera and Simon 1993). Production systems inherit the symbolic representation position and the aerial view of the planning model and simply provide a different account of reasoning as taking place through the application of rules. In order that they be sufficiently general, the rules have as arguments, variables representing possible goal and world states. These variables have to be bound to real objects just as in the planning model. Thus, the processing of sense data is assumed to be of the same symbolic representational kind as in the planning model. It is tempting to think of situated action theories as implementing rules between features, detected in pre-processing of sense data, and actions, but this does not do justice to the complexity of the sense-action relationships that is found in the examples described below.

7.3.6 Agent Architecture

In the planning model of activity, because the link between sensing and acting is mediated by deductive processes on a centralised representation of the world, a decomposition of the agent into sensing, planning and acting models, with serial processing is necessary. If a tight, direct connection between sensing and acting is sought, and if, at most, fragmentary representations should mediate this connection, then a parallel and distributed agent architecture becomes a possibility.

7.3.7 Interaction

In the planning model the cause of action is the plan and thus, what happens is considered to be the result of the mental work of the agent. By contrast, the situated action model stresses interaction between an agent of limited reasoning power and a structured environment. This position is demanded by the nature of situations and their putative causal role in actions. To see this, let us once again view activity from the agent's perspective. The agent is in a situation and takes an action determined by that situation. The action is the agent's response to a situation. But the next situation faced by the agent depends not only on the previous action but on how the world is after that action (as seen from the agent's point of view). This depends very much on the structure of the environment. In fact the mapping from agent's actions to the next situation can be viewed as the *behaviour of the environment* under the agent's action. The size of the repertoire of responses of the environment to agent actions is a measure

of the complexity of the structure that the environment possesses, just as the size of the agent's repertoire of appropriate actions for situations is a measure of the complexity of the structure of the agent's behaviour. Therefore, the outcome of the sequence of situation, action, situation, action, . . . (to the extent that it can be viewed as a discrete sequence) is just as much determined by the complexity or simplicity of the environment as the complexity or simplicity of the agent's response to situations, and whether the goal situation appears in this sequence is also a function of both these things. If we try to construct a causal chain from a certain action of the agent to the goal state, the response of the environment to the agent's actions is part of that causal chain.

This view of the environment's response to the agent's actions is rather like a computer "adventure game" where the agent (from whose point of view the action is usually represented) performs action which transform the world, usually resulting in access to new rooms, and so forth, whose relationships (the structure of the environment) are described by a set of rules embedded in the game program. This idea can be made quite formal by treating both the agent and the environment as finite state automata (see (Horswill 1995) and references therein). This leads to the view of activity under this model, as a "dance" between the focal system possessing "agency" and an environments also with an equally valid claim to "agency" (Chapman 1991).

7.3.8 Locus of Agency

The series of events, encountered by our can collecting robot, would appear exactly similar to an agent using the planning approach, except that the appearance of each situation would have been anticipated in advance (implicitly at least) in the planning process, and would not have a causal status for the next action (except perhaps if the situation was unanticipated in which case re-planning would be called for). It is because the planning approach has *anticipated* the future through its symbol manipulations on a model of the world, that we attribute the dynamics of the scene entirely to the agent's actions and we do not think of the environment as having a direct causal influence at all. The structure of the environment "caused" the plan to be produced, but between this and action stands the plan which is thus the efficient (or proximal) cause of action. We therefore attribute the complexity of the agent's behaviour to the structure of the plan, rather than the structure of the environment. But because the situated action model assumes no simulation of the future, we are *forced* to include the environment's response in any causal analysis of the events that take place. The environment is a co-agent in the description of the scene, and the locus of agency in the Situated Action model is not simply the focal agent, but the *interaction* of the focal agent with its environment.

It is important to note here that both the planning model and the situated action model depend equally on the world being structured. In the planning model, if the world had no structure there could be no world "physics" to exploit in the planning process (no predictable effects of actions as state transformation, for example), and in the situated action model there could be no stable interaction between the agent and the world. Therefore, the existence of structure in the environment is not an assumption of the situated action model *per se*: it simply enters the causal picture at a different time. However, we will see later that there are certain *types* of structure that are particularly useful to situated agents.

7.3.9 Goal Attainment

In the situated action model, goals are attained through the interaction of an agent, which has a limited repertoire of actions for stereotypical situations, with a structured environment. Goal attainment is thus *emergent* from this interaction. It is pretty clear that we cannot expect situated action to produce an *optimal* trajectory in state space from an arbitrary initial state to a given goal state. As argued in the previous chapter, this attraction of the planning approach is probably an illusion in real environments anyway. What we hope to get instead is a behaviour that is *robust* in complex and relatively

unpredictable environments, by avoiding the brittle symbolic translation process inherent in the planning model. But why should we expect that it will be possible to design agent behaviours that will reliably achieve intended goals? There are several reasons:

1. Environments have a lot of structure;

2. Interaction is ubiquitous;

3. Certain environmental structures are particularly conducive to situated activity;

4. In many cases, goals, environments and situated agents have co-evolved, in others they are situation-specifically designed.

The structure of an environment is the degree of constraint that it places on possible actions. Taking human activity as an example, it is easy to underestimate, or fail to notice completely, the amount of structure in our everyday environment and its role in constraining possible actions and thereby relieving us of some of the cognitive burden of acting. The reason we fail to notice this structure is that we routinely attribute all the structure of our actions to the plans that we assume precede action. As you enter a house the mere existence of, and the connectivity of, the rooms dividing up the space, constrains enormously the possible next situation you will encounter. Within rooms the very existence of, and the stereotypical nature of, artefacts such as furniture that you encounter, further constrains the possibility of certain actions and enables others. When you are in a kitchen and you need a knife, the very existence of kitchens, dressers, cutlery draws, and knife compartments, means that you do not have to engage in an exhaustive search of the entire house to find a knife. It is just this type of exhaustive search of possibilities in a problem space that the planning model takes as the basis of intentional activity. Going in the opposite direction, we can detect a hierarchy of compartmentalisation constraining activity in useful ways. Starting with the possibilities offered and denied by our own body parts (shape of hand, symmetry etc.), we move to clothing (pockets, etc), to temporary work-spaces (used to collect together parts of a projects), to private work places (making tools permanently available), to shared work places, to public spaces (Agre and Horswill 1997, p139). These examples deal with artifacts and spaces. The same constraining and enabling character is true of other aspects of the human environment. The legal and political environment constrains what you can do without sanction. The cultural environment constrains what ideas and questions are considered meaningful.

Interaction is ubiquitous for embodied agents because they are part of the world. Again, in human intentional activity we fail to notice this because, by placing the locus of agency in our plans, we assert a strong form of autonomy from our environment, based on our ability to predict it. If we back away from this position, we can see that interaction is a natural outcome of our physicality, our embodiment. It is easy to set up reactive behaviour in robots that results in wall following because walls place such strong constraints on a wandering robot - a mobile robot cannot fail to interact with walls. Without wishing to retreat to a physicalist explanation of movement, the situated activity model recognises the connectedness of parts of the agent with parts of the environment by placing the interaction of agent actions with environmental structure to the fore.

Certain kinds of structure in environments are particularly useful to situated agents. These are structures that constrain the agent's possible space of actions in such a way that very simple strategies can be used by the agent, strategies that require little remembering of past actions, and use a simple "hill-climbing" or "greedy" approach to closing on the goal. An example is the property of a class of mazes that allows them to be traversed with the simple rule "keep your left hand in the left wall at all times". This is the property of being "singly-connected" (which means there are no closed circuits in the maze) and is true of many real mazes. Jigs used in manufacturing are another extreme example. As part of the environment of work, they constrain action so severely that they remove the need to plan or decide completely. Agre and Horswill have developed a formal theory of agents and environment which allows the enabling properties of certain environmental structures to be precisely defined (Agre and Horswill 1992; Horswill 1995; Agre and Horswill 1997). For instance, using the simple repetitive activity of making breakfast as an example, they have shown how the particular properties of tools

which are present as part of the task environment, allow the task to be segmented into episodes that require only situation-reponse chains of action, with little or no problem space searching.

Spatial compartmentalisation, jigs, and tools are simple examples of how the environment takes part of the cognitive burden for a situated agent. But there are also ways that environments can be actively used by agents to simplify the possibility of situated action. One is using the environment to take the burden off short-term memory by, for example, placing things near the door so they will be remembered on leaving home. Another is the process of "stabilising" environments (Hammond, Converse and Grass 1995) (Agre 1995a). This refers to actions that agents take to maintain the cognitively useful properties of their environment, such as returning tools to certain places, and dedicating spaces to certain projects. To the extent that the environment plays a crucial role in making intelligent action possible, "it is almost as if these surroundings were an extension of one's mind" (Agre and Horswill 1997, p139).

In many cases, environments and situated action, and indeed goals, have co-evolved. Tools, house layouts, customs, myths are all cultural artifacts, whose continued existence as part of the environment is ensured by their usefulness to activity. The activity and the tool may have, in many cases, been invented simultaneously, and at the same time given rise to previously un-thought-of intentions. Once the utility of a tool is realised, the knowledge of the use of the tool and the potential intention to use it are passed on simply by passing on the tool. Tools are compiled plans. The tool, its use, and its purpose are all packaged together and are passed on as cultural artifacts. The same is true of many other aspects of environments, physical and cultural.

It is difficult in many cases to distinguish between the environment structure as the cause of a certain activity and the activity as the cause of the environment structure. If we observe a teacher in a classroom with all the desks arranged in neat rows, can we tell whether the teacher adopts a teacher-centred approach to education *because* of the layout of the classroom, or if the classroom is laid out that way *because* she want to use a teacher-centred approach?. This is the type of mutual causality that results when the notion of strong separation of system and environment, inherent in an input-output systems analysis, is replaced by the notion that the behaviour and structure of systems is *closed* through the response of the environment (Maruyama 1963; Maturana and Varela 1980).

In the case of designed agents, many aspects of the environment and the agent's goals are generally given. Rather than evolution shaping the agent's response repertoire, the agent has to be designed to show the correct goal attaining behaviour in this given environment. Agre (Agre 1997, p62) recommends that the designer should analyse the structure of the environment and its potential for interaction, design a simple agent that exploits this potential, test, and tweak. Brook's (Brooks 1986) adds the extra possibility of building layers of more complex behaviour upon simpler ones whose robustness has already been tested. These processes once again highlight the different role of planning and deliberation in the planning and situated action models. In the planning model, the designer builds declarative knowledge and the ability to perform all-purpose planning *into* the agent as a *modus operand*i: in the situated action model the *designer* uses these planning abilities to construct a parsimonious agent which then does not have to use these unwieldy abilities in ongoing activity.

This discussion does not conclusively demonstrate the possibility of designing agents based on the situated action model that reliably achieve goals. The degree to which the reader believes this is possible will depend on the degree to which he/she believes that:

1. Complex systems consisting of agents with limited behavioural repertoires interacting with complex environments are conducive to the appearance of stable patterns of interaction;

2. Such stable interactions can be "tweaked" either by small changes to the agent behaviour, or the environment structure, or the way the agent uses and interacts with the environment, to ensure that the goal state is visited in the interaction.

3. Stable interactions can form a basis for building further stable interactions of greater scope allowing the bottom-up construction of intentional behaviour of considerable complexity.

These are empirical questions that are just beginning to be investigated empirically (Agre 1995a).

7.3.10 The Model

The process that the situated action model sees occurring in intentional action can now be stated.

1. Actions are responses to situations which occur as the agent interacts with its environment.

2. A situation is the focal agent's situated, embodied, and goal-directed perception of what is happening in the world.

3. There is a tight (that is, essentially direct) connection between sense data and actions of the executor systems. This process does not make use of symbolic representations of the goal state, the current state of the world, simulation of future states, or reasoning about proposed actions (plans). However, the link between sense data and action may be mediated or filtered using processes that implicitly depend on the intended goal situation. This will be some form of goal-directed *attention* to sense data.

4. Because situations are functions of both what is going on in the world *and* the agent's perception of what is going on, the outcome of an agent's actions, and in particular the achievement of goals, is as much due to the structure and dynamics of the world as to the agent's actions.

7.4 Case Studies of Applications of the Situated Action Model.

To give some concrete examples of the situated action position at work, I draw once again upon examples theorising about the nature of human activity, and mechanistic examples of attempts to construct physical or virtual agents using the approach described above. In chapter 6, I used the example of Shakey to show that the problems of the planning model occur also in purely mechanical systems, and therefore derive from the underlying position on the nature of activity, rather than having any essential connection with the human limitations. I now use mechanistic implementations of situated action ideas to dispel the idea that they are purely humanistic, or even mystical, an impression that one can easily gain from phenomenological writing on human action.

7.4.1 Mechanistic Implementations of Situated Action

Brook's "Creatures"

In the seminal paper "A Robust Layered Control System For A Mobile Robot" Brooks (Brooks 1986) described a new architecture for an autonomous goal directed mobile robot. Brooks argued that, while some kind of decomposition of a design project like building a robot was necessary to make it tractable, the traditional functional decomposition into sensing, planning, and execution was the source of much of the behavioural fragility of earlier approaches (for example, Shakey). He proposed an alternative decomposition into modules based on the behaviours they performed, and argued that this architecture would lead to more robust behaviour in ambient environments such as office settings. The architecture called for a number of interacting sub-systems each of which was functionally complete in that each had the full complement of sensing, acting (or at least they had direct access to shared sensors and effectors) and a small amount of computing power. Each sub-system was to carry out a certain distinct behaviour which was designed on the basis of an analysis of the behaviour, hard-wired into a dedicated controller, and "tweaked" for robustness by trial and error testing in ambient environment. The systems could interrupt each other if predetermined environmental conditions occurred. There was to be no central symbol-object representation of the world: the sub-systems interacted directly with the sensed

world and even communicated with each other to some extent through the world (or in Brook's words, used the world as its own representation).

The first implementation of this idea, the robot "Alan", had only three behaviour layers. The first handled avoidance of objects both stationary and moving. The second made the robot wander randomly when no other behaviour was active. The third layer sought out distant objects and caused the robot to head towards them. While Alan could hardly be described as intelligent or versatile, being largely a demonstration of principle, it nevertheless displayed purposeful behaviour - at least as ascribed by an observer - which was largely emergent, rather than programmed in, and also highly robust in ambient environments such as office buildings. Brook's group (Brooks 1989; Brooks and Flynn 1989; Brooks 1990; Brooks 1991c; Brooks 1991b; Mataric 1992) has gone on to design a series of robots of increasing complexity and with more useful behaviours, including insect-like creatures that can learn to walk, and a can-collecting robot (which inspired the thought experiment used throughout this thesis), using this "subsumption architecture". Recently Brooks has initiated a project to build a humanoid robot which learns by interacting with the human world based on these principles (Brooks and Stein 1993).

Sensors ⟶ | perception | modeling | planning | task execution | motor control | Actuators ⟶
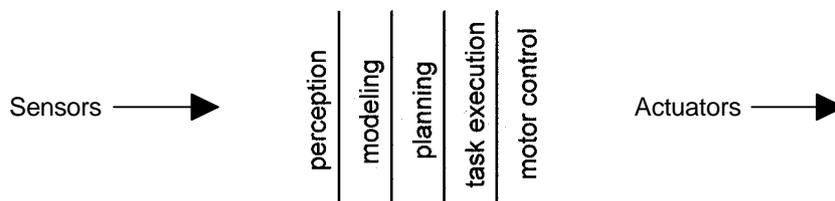
Traditional decomposition of a mobile robot control system into functional modules.

reason about behavior of objects
plan changes to the world
identify objects
Sensors ⟶ monitor changes ⟶ Actuators ⟶
build maps
explore
wander
avoid objects
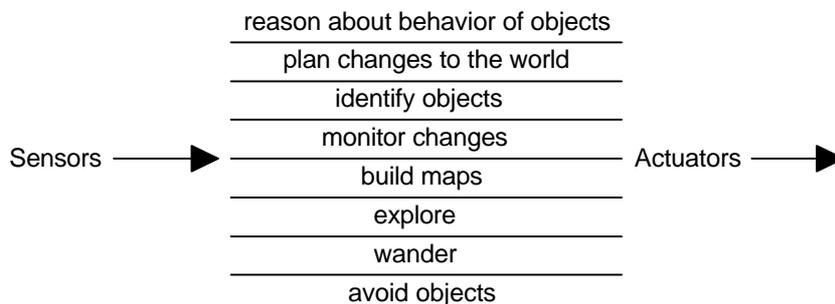
Decomposition of a mobile robot control system based on task-achieving behaviors.

Figure 7.1: The Alternate Modular Decomposition Proposed by Brooks. After (Brooks 1986, p14)

Brook's creatures implement a number of aspects of the situated action model of activity. Brooks has consistently maintained the position that centralised symbol-object representation in not required for goal-driven behaviour. In fact he argues that his robots do not use representations at all (Brooks 1991b). By the use of a parallel, distributed modular design, his subsumption architecture is able to implement a tight coupling between sensing, computation and action. These behaviour modules act directly upon the limited data available to the situated and embodied robot, and while some modules could be thought of as employing some sort of representation (such as records of sonar responses) these are fragmentary and agent-centred. Brooks has argued that this tight coupling greatly assists action based on incomplete and noisy sense data and unpredictable motor systems. Although each behaviour module can be said to have its own goals (Brooks 1986), which conflict with the goals of other modules, the goal achievement of the robot agent as a whole is emergent from the interaction of the low level behaviours with the environment, and with each other. The interrupt mechanism can be thought of as implementing the concept of attention inherent in the situational approach. Although the intended behaviour of each module is based on sense data in an essentially reactive rule like fashion, the effect of

the interrupt mechanism is that the behaviour of the agent as a whole is more complex than simply a rule-like connection between situations and actions.

The robot "Toto" designed by Maja Mataric (Mataric 1992; Hendriks-Jansen 1996, p141) is particularly interesting in relation to the status of representation in activity. This robot, using subsumption architecture principles, builds maps of its environment and uses them to navigate. But rather than starting from an architecture where low level reactive behaviours are incorporated within a representational planning system, which has been the usual approach of developments of the planning model, Toto builds representations bottom-up based on data extracted from the low level reactive behaviours. Toto has a lower level which uses four behaviour modules ("stroll", "avoid", "align", "correct") which can interrupt each other to establish an emergent and robust wall-following behaviour. Even though this lower level uses no map of the rooms in which it achieve this wall following behaviour, and indeed uses no concept of "wall" at all, the sequence of data from its sensors and the interrupts that occur between its low level behaviours, contain substantial information about the geography of the environment. The high level behaviour modules of the robot recognise recurring regularities in this data stream which are interpreted as "features" of the environment, and are used to build a "map" of the environment as the robot explores it. This map consists of a linear list of features and the order of the list implicitly defines the topology of the environment. The map uses no aerial view coordinate system, only regularities in its directly available sense data. Now the high level layer can use this map to navigate. Because the landmarks are ordered in the map, it can be used to direct the robot from its current location to a goal location (specified by an outside operator) by a simple hill-climbing algorithm. Recognition of the intermediate features is accomplished by recognition of the sense data patterns that define these "features", and what to do to alter the behaviour of the low level wall following layer is implicit in the description of the features stored in the map. If the robot loses track of its position it will *automatically* wander until its sense data match a feature description and it can then recover automatically. Because the high level behaviour was built over a robust low level behaviour, the low level behaviour is all the while handling the contingencies of the environment, such as obstacle avoidance, without any reference to higher modules.

This example not only illustrates how increasingly complex behaviours can be built bottom-up upon robust low level behaviours, but also how behaviours depending on representation such as planned navigation could conceivably emerge from reactive low-level behaviours. The point is that low level behaviours that combine to produce a stable patterns of interaction with features of the environment generate, by virtue of this stability of interaction, data which can be used by further layers to produce behaviour correlated to more detailed and global features of the environment, and so on. Hendriks-Jansen has recently argued in detail that this bottom-up approach to representational behaviour provides the only plausible explanation of how intelligent behaviour could have evolved in animals and humans (Hendriks-Jansen 1996).

Pengi

While Brooks has repeatedly insisted that his work is motivated by engineering considerations, Agre and Chapman (Agre and Chapman 1987; Agre 1988; Chapman 1991; Agre 1997) wrote the program "Pengi" specifically to demonstrate the feasibility of an intentional system that did not use symbol-object representation or planning. Pengi is a computer simulation of an agent that plays an arcade game called "Pengo" (which is also simulated). In Pengo, a penguin moves in response to commands from a button and a joy-stick, in a two dimensional video screen world of ice-blocks which effectively form a maze. The world is also inhabited by bees whose behaviour is almost entirely random. If a bee gets too close to the penguin, the penguin dies, and the game ends with a loss. Both the penguin and the bees can kick the ice-cubes and they slide in the kicked direction until they reach an obstacle. If the obstacle struck is a bee it dies, and if all bees die the game ends with a win. If the obstacle is the penguin it dies, and the game ends with a loss. The environment is thus fairly complex, dynamic, and has a high degree of unpredictability. Pengi's behaviour, in controlling the actions of the penguin, is clearly goal directed.

Note that Pengi is a simulation of the game player's actions and point of view, not the penguin's, although the player can only control the penguin's movements.

Pengi does not use any planning: it is merely equipped with a set of appropriate responses to each type of situation, which have been put in by hand by the designers on the basis of an analysis of the dynamics of the game. In fact Pengi's action selection is done using a network of binary logic gates which have no internal state. These generate a series of simulated joy-stick (left, right, up, down) and button (press) movements which are passed to the game simulator. Thus, Pengi's executor system is essential trivial. However the visual system is quite sophisticated. It is designed to detect characteristic features of the situation on the game board which are relevant to Pengi's goals of avoiding defeat and winning. While the vision system has access to the data structure representing the positions of the various objects on the game board, this symbolic representation is not made available to the simulated agent. Instead, this data structure is used like a retinal image and various features in it are extracted by "visual routines" which are used to pick out objects of relevance, such as bees threatening the penguin, ice-cubes that can be kicked, and other spatial relationships. The visual system can also place a small number of markers on objects to keep track of them as a simulation of visual attention. Most of the visual operators ask binary questions about the relationships between markers the answers to which are used as input to the central logic system.

Although the visual system individuates objects, it does not build a symbolic representation of them. Instead of labelling them BEE-34, ICE-CUBE-62, that is, binding variables to particular instances, it labels them *indexically* by there relation to the penguin (*the-wall-I-am-lurking-behind*) and *functionally* in terms of their significance to Pengi's (i.e. the player's) endeavour (*the-bee-I-am-attacking*). In contrast to a symbol-object representation which would set up a permanent binding of names to particular objects, these names, which actually refer to markers, will represent different objects on the board as the need arises. For instance, when a bee is killed or gets into a position where it cannot be attacked by a particular bee, the name *the-bee-I-am-attacking* can be applied to another bee. It is this indexical and functional labelling of objects, together with the reactive approach to action selection, which accounts for Pengi's ability to act quickly and deal with a constantly changing environment. Rather than having to track all objects on the game board, which would involve a heavy overhead in constantly binding variables, and reasoning about what to do about them, it simply responds in a variety of stereotypical ways to situations consisting of stereotypical relationships between generic bees and ice-cubes. Some actions involve decisive moves such as killing a bee by sliding an ice-cube, and occur when the opportunity arises. Others, such as pushing ice-cubes to more advantageous locations are performed because they tend to create future opportunities. However, they are not steps in a plan but merely a suitable (hopefully the most suitable) response to the immediate situation. If they lead to killing a bee or winning the game at a future time step this is as much attributable to the dynamics of the environment as the performance of the act. According to Agre:

> Pengi plays a pretty decent game of Pengo. In its present state it is a little better than I am, which is to say it wins from time to time and puts up a good fight. (Agre 1997, p265)

Pengi provides an existence proof that at least some aspects of the situated action model of activity can be used in the design of mechanistic intentional systems. This should be an antidote to the somewhat mystical feel of the phenomenological description of human intentionality. The simulation shows clearly how indexical, functional representation of an agent's situation, coupled with largely reactive action selection, can keep computational explosion in check and allow for rapid adaption to changing circumstances. It also demonstrates how goal attainment can emerge from interaction between situated action and a structured environment. It is easy to see, using a physicalist description, that the separate simulations of the agent and of the environment together form a closed system. At the physicalist level of description, the ultimate clearing of the video screen of bees is a result of the detailed deterministic evolution of this whole system from initial conditions (even the pseudo-random number generators used to control the bees are deterministic). The view of Pengi taking goal-directed situated actions is the activity level description of the same phenomenon.

The concrete example of Pengi provides an opportunity to clarify the relationship of situated activity model to the stimulus-response models of activity. Although Pengi acts to a large degree reactively it cannot be viewed as a reversion to a Stimulus-Response approach to activity for two reasons. Firstly, the complexity of the visual processing unit guarantees that there is not simple transfer function relation between Pengi's input (the game board) and its output (joy-stick and kick-button actions) as envisioned in S-R theories. And this is despite there being only limited internal state (the visual system has a one time step memory of past events). Secondly, although goal attainment is largely emergent, Pengi's operation involves intentionality in the way that situational features are processed. The visual systems implement ideas of attention and relevance as a function of Pengi's goal. Computational simulations that involve a single goal tend to obscure this difference between situated action and stimulus-response theories. If Pengi had the goal of committing suicide as quickly as possible or of stretching out the game as long as possible it would not only have to use a different logic structure to select action from information provided by the visual routines, but also it would certainly need to use the visual routines in a different way and may need different routines to achieve these goals efficiently. The point is that what counts as significant in the situation for action is goal dependent in situated action theories, which are therefore thoroughly intentional theories in a way that S-R theories are not.

Neither can Pengi be viewed as using a production system. Although Pengi's central logic system could be viewed as implementing a set of rules (on account to the universal computing power of production systems), the referents of the rules envisioned in the production system account of skilled action are goals, sub-goals, and variables which need to be bound to particular objects in the current world state. The referents of Pengi's rules are features of the situation which have already been selected as salient to Pengi's goals. This may seem like hair splitting but it points to the essential difference in how intentionality is treated in planning and situational models of activity. Production system models envision the application of situational variants of goal specific rules - they are still within the planning model framework trying to implement action as reasoning about goals and "situations" (in the sense of objective features of the world) - whereas, in the situated action model, goals enter the theory at the level of what constitutes situations (in the sense of agent-centred and goal- directed, history dependent views of the world) which may then be viewed as determining action.

7.4.2 Human Activity

The situated action model, applied to human activity (Suchman 1987), holds that normally, humans simply do what is appropriate to the situation in which they find themselves, and reasoning and planning of the type envisioned by the Planning model plays a minimal role as a cause of action. On this account, activity is not the result of plans which are transformed into actions by motor systems, but is the direct response, routine or improvised, to situations or some salient features of situations. Plans are viewed as either a method of analytically dissecting proposed action *before* it takes place in order to prepare oneself for acting in unfamiliar circumstances, or alternatively, a way of organising retrospectively an account of what was done, usually for accountability purposes. Plans have an immediate causal relation to action only in the most deliberative and faltering kinds of activity characteristic of a learner. In fact, according to Heidegger (Heidegger 1961, orig. 1927), the need to dissect the world into objects and relations and form mental representations, arises only when this normal transparent coping with the world encounters a *breakdown*, either a totally unfamiliar situation or a totally unexpected response to an action. Even then, representations and the formal manipulation of them with respects to goals and current states, only support thought *about* action rather than action itself. In this respect plans are one of many *resources* for human action rather than the *cause* of action, and plan making is just another type of situated activity elicited either by unfamiliar situations or situations where there is a need to give a formal account of action.

An example from Suchman (Suchman 1987) may clarify this image of activity and the status of plans in it.

..... in planning to run a series of rapids in a canoe, one is very likely to sit for a while above the falls and plan one's descent. The plan might go something like "I'll get as far over to the left as possible, try to make it between those two rocks, then backferry hard to the right to make it around the next bunch." A great deal of deliberation, discussion, simulation, and reconstruction may go into such a plan. But, however detailed, the plan stops short of the actual business of getting your canoe through the falls. When it really comes down to the details of responding to currents and handling a canoe, you effectively abandon the plan and fall back on whatever embodied skills are available to you. The purpose of the plan in this case is not to get your canoe through the rapids, but rather to orient you in such a way that you can obtain the best possible position from which to use those embodied skills on which, in the final analysis, your success depends. p52

In terms of the account of situated activity given previously, I would translate this account thus. The canoeist, when poised at the top of the rapids, simulates an aerial view of the scene in order to reason about his/her relation to various objects in the path and to simulate possible future relations. But when he/she pushes off, the aerial view becomes too cognitively burdensome to maintain as a principle for action and he/she is forced to rely on a ground view and respond to situations in real-time by reflex or procedural knowledge. For an observer to assume that h/she continues to make use of the aerial view during the action is to make the "theoretician's error".

The situated action theory of human activity thus reverses the status of reasoning and representation on the one hand, and improvisation and reactivity of the other. In the planning model representing and planning were viewed as the central phenomena of activity, and reactivity and improvisation were introduced to overcome some technical difficulties in accounting for our ability to cope with a complex, dynamic world. In the situated activity model reactive coping unmediated by reasoning and representation is the normal state, and deliberative planning and modeling the world is a last resort.

According to this new view the need to build objective representation of the world around us arises only when a breakdown in this coping state requires us to plan. If things go hopelessly wrong for the canoeist in his/her descent, he/she may try to prop against a rock or the bank to get out of the action long enough to reconstruct the aerial view. Many authors from the phenomenological tradition (Heidegger 1961, orig. 1927; Merleau-Ponty 1962, orig. 1945; Polanyi 1966) have found evidence for this statement is the way in which we relate to equipment in situations of transparent coping, compared to situations of breakdown. When we are using equipment in a fluent, skilled way, as when a trades-person uses a tool, it is as though the equipment is an extension of our body and we are not aware of it. In fact the equipment extends our perception, rather than it being perceived. Heidegger calls equipment in this state ready-at-hand (available) (Dreyfus 1991, pxi). However, if the equipment becomes involved in a breakdown of transparent coping, say if it is found to be inappropriate, or an accident occurs, then it becomes the object of contemplation and it suddenly appears as an object mediating activity. His famous example is a carpenter's hammer which is normally not experienced as intervening between the carpenter and the act of hammering unless he/she slips or the hammer is found to be the wrong type or weight. Heidegger calls such equipment present-at-hand.(occurrent).

> Consider the example (used by Wittgenstein, Polanyi, and Merleau-Ponty) of the blind man's cane. We hand the blind man a cane and ask him to tell us what properties it has. After hefting it and feeling it, he tells us that it is light, smooth, about three feet long and so on ....But when the man starts to manipulate the cane, he looses his awareness of the cane itself; he is aware only of the curb (or whatever object the cane touches); or, if all is going well, he is not even aware of that but of his freedom to walk, or perhaps only what he is talking about with a friend. Precisely when it is most genuinely appropriated equipment becomes transparent. (Dreyfus 1991, p65).

Notice again the role of the aerial and ground views. When contemplating equipment from the detached aerial view, either due to a breakdown, or though meeting it in an unfamiliar circumstance, the full symbolic representation apparatus is brought to bare on it. It is perceived as objectively exsistent apart from any agent's situation or purposes. When it is in fluent use, viewed from the ground, it only needs to be represented, if at all, to the extent that is part of the agent's goal directed interaction with the

environment of the task at hand. As Agre (Agre 1997, p332) notes, deictic (indexical-functional) representation is the appropriate way of representation ready-at-hand objects.

The situated action model posits that most activity occurs in a world where objects are ready-at-hand. The planning model is an image of action in a world where all objects are always present-at-hand: they always need to be explicitly represented, contemplated and reasoned about. Our experience tells us, so the argument goes, that this is only ever the case in the most hostile and unfamiliar environments, so to use this as the underlying metaphor for human action is a fallacy. Because the planning model is an agent-independent, purpose-free, aerial descriptive framework it has no way of making this kind of experiential distinction.

7.4.3 Relation to Skill

It is tempting to make an association between the type of human activity portrayed in the situated action model, and skilled human activity. Skill is characterised by "smooth, automated, and highly integrated patterns of behaviour" (Rasmussen 1983). The planning model of activity would then be a description of the faltering, jerky, deliberative activity of a novice working out what to do from first principles. There are a number of good reasons for such an association.

1. It is known that highly skilled persons, or experts, deal with situations in the direct, immediate sort of way envisioned by the situated action model. This was found for instance in the earliest study of chess expertise by de Groot, and has been confirmed by many subsequent studies (Chase and Simon 1973; Simon 1973).

   > Research has shown that chess masters rapidly perceive the game configuration as a whole, with the best move effortlessly retrieved as part of this process. (de Groot 1978, orig. 1946, p234)

   Chess playing has been a favourite case area for the study of skill because there is a clear cut and objective measure of skill (international rating) and fragments of behaviour that capture the essence of the skill can be studied under laboratory conditions (Ericsson 1996a).

2. Expertise appears to be related less to deliberative skills and more to pattern recognition. de Groot was expecting that chess masters would gain their advantage by being able to mentally examine more future moves (that is greater deliberative ability) than simply good players, but found that there was little difference in search depth between the two. His research indicated that the key difference was their greater ability to recognise characteristic positions. Evidence of this is found from laboratory tests of chess master's abilities to remember and recall chess positions (Chase and Simon 1973). Performance of chess masters is not significantly hampered by the simultaneous performance of mental arithmetic (Dreyfus and Dreyfus 1996), and this has been confirmed for sporting skill (Leavitt 1979).

3. Skill is sentient rather than intellectual and it is difficult to separate the *knowing* from to *doing*. It is often difficult for experts to explain what they are doing and skill is often best communicated by demonstration (Zuboff 1988). In addition the accuracy and reaction time of sporting skills often deteriorate greatly when actions are performed in response to videotaped situations, compared to when they are performed under natural conditions (Shea and Paul 1996), indicating their inherent situatedness.

4. There is quite a deal of evidence that the progression from the faltering, slow, and brittle novice performance, to the efficient, fluent performance of an expert, is accompanied by one or more *qualitative* changes in the cognitive apparatus involved. This was noted in the earliest work on word list memory (Ebbinghaus 1964, orig. 1885) and telegraphic operator skill (Bryan and Hartner 1897), and the classic laboratory experiment is (Fleishman and Rich 1963). Fitts (Fitts 1964) saw three stages: a cognitive stage in which actions are consciously monitored; an associative stage during which direct associations between inputs and actions are developed; and an autonomous stage where no conscious control is required. Anderson (Anderson 1983) and Rasmussen (Rasmussen 1983) have given similar stage models. Dreyfus and Dreyfus have given a detailed phenomenological analysis of skill acquisition

(Dreyfus and Dreyfus 1989; Dreyfus and Dreyfus 1996) which has been used to explain observations on acquisition of nursing skill (Benner 1984). They recognise five stages: deliberate application of non-contextual rules; recognition of contextual variants of rules; deliberate decision making based on recognised features of situations; intuitive response to features of situations interspersed with deliberative decision making; and finally, absence of deliberation except when complete breakdown occurs. Accompanying this transition from deliberation to direct response to situations, is an increased involvement of the actor with the actions and increased responsibility taking. The whole process of skill acquisition for these authors is very much a transition from the detached "aerial" view of one's actions to the involved "ground" view.

As a corollary, Zuboff describes the confusion and loss of control that skilled workers experienced at a paper mill, when they were required to perform their usual skilled work in deliberative fashion, through the mediation of a remote computer system.

> The chlorine has overflowed and .... its all over the floor, but you can't see it. You have to remember how to get into the system to do something about it. Before you could see it and you knew what was happening - you just knew. (Worker quoted in (Zuboff 1988, p62))

Sensing had been redefined as reading data from their screens which, rather than immediately suggesting action, required the deliberate bringing to mind and applying rules and procedures.

> Before computers we didn't have to think as much, just react. You just knew what to do because it was physically there. Now, the most important thing to learn is to think before you do something, to think about what you are planning to do. (Worker quoted in (Zuboff 1988, p74)).

They were reduced to novices in their own industry and often expressed a deep loss as a result. In fact the central control room with its information displays, stylized graphical representation of the process and indirect, non-corpereal responses, in which these workers were now reduced to novice performance, is a perfect metaphor for the homuncular vision of activity recommended by the planning theory of activity.

5. As Dreyfus has pointed out (Dreyfus 1992, pxxi), deliberative theories of activity scale badly with increased knowledge of situations. In deliberative theories, based on problem space search, greater knowledge of situations should result in slower performance whereas we know from experience that it results in faster, more fluent performance. Associating skilled performance with the use of increasingly fine-grained world models also threatens to overload the well known information processing limits placed on the brain imposed by the "rule of seven" (Miller 1956) and the "single channel" theory (Broadbent 1958). It appears that a mechanism other than deliberation is required to explain skilled performance.

6. Simon and Chase (Simon 1973) made the observation that ten years practice was required to achieve expertise in chess playing and this "ten year rule" has received verification in other areas also (Ericsson, Krampe and Tesch-Romer 1993). So on this account, adult humans would be expert in a great proportion of everyday activity, and we would expect everyday activity to bare more resemblance to skilled performance than to the deliberative performance that the planning model suggests.

While this association between situated action and fluent, skilled activity on the one hand, and the planning model approach with deliberative, novice action in unfamiliar circumstances on the other, is very much the position of the phenomenological philosophical and psychological camp, it is not generally accepted by the mainstream cognitivist position in the psychology of skill (Ericsson and Smith 1991; Proctor and Dutta 1995; Ericsson 1996b). According to their position, the symbolic, information processing view of activity can be extended to the domain of skilled activity by assuming the skill acquisition is the process of compiling declarative knowledge down to procedural production rules (Anderson 1983; Newell 1990) and "chunking" the units of perception (Chase and Simon 1973; Newell and Rosenbloom 1981). In the complilation process, the recognition of recurrent rule like relationships

between world states and action, together with rule generalisation, gradually replaces deliberative (search-based) reasoning with procedural (rule-firing) reasoning. Chunking is the process of using larger elements of experience and situations as the units manipulated by perception, memory and rule-firing. The necessity of these processes is supposed to account for the long required duration, and diminishing returns, of practice (Simon 1973; Newell and Rosenbloom 1981).

This is an ingenious theory but, by retaining a commitment to the whole symbol processing apparatus, it is open to all the criticisms generally levelled by proponents of situated action at the planning model of activity. In addition, it is also subject to well known criticisms of rule systems as a basis for intelligent action (Wittgenstein 1972, orig. 1953; Dreyfus 1992). Arguably, the application of rules, even situation specific ones, as conceived in this rule-based model, is typical of a novice approach to action at best. Fortunately, I do not have to resolve this debate on the nature of human skilled action. My project is to explore the consequences of adopting the alternative situated action perspective for the design of management systems. The robust, fluent, efficient, direct response to situations of skilled activity can serve as a guiding metaphor without making a commitment to the particular processes that explain *human* skill.

7.5 Theories of Activity Summarised

By way of closing this discussion of theories of activity, this section summarises the two basic approaches in table 7.1 and compares their stand on key issues for the explanation agency and design of agents in table 7.2.

7.6 Literature

Situational / interactional theories of activity have their roots in the phenomenological tradition of philosophy and psychology of Heidegger (Heidegger 1961, orig. 1927), Merleau-Ponty (Merleau-Ponty 1962, orig. 1945), the later Wittgenstein (Wittgenstein 1972, orig. 1953) and Polanyi (Polanyi 1966). Accessible accounts of Heidegger are found in (Winograd and Flores 1986; Dreyfus 1991), and of Merleau-Ponty in (Dreyfus 1992). The idea of phenomenology was to found philosophy and psychology on a detailed study of the nature of human experience of the world from the "ground view". Heidegger's elaboration of *Dasein*, or human "being", can be read as a description of the transparent coping of a situated agent.

While the phenomenological tradition has influence European scholarship on agency (eg. Bourdieu) it has only influenced computational theories of agency relatively recently. The connection seems to be Dreyfus's critique of artificial intelligence (Dreyfus 1992). Dreyus is a US authority on phenomenological philosophy. As a theory of agent activity the Situated Action perspective has been explored in many of the disciplines that make up cognitive science: in robotics (Brooks 1986; Brooks 1991b; Brooks 1991a); in human-computer interaction studies (Winograd and Flores 1986; Suchman 1987); in artificial and human intelligence (Agre and Chapman 1987; Agre 1988; Chapman 1991; Agre 1997); in human cognitive function (Varela, Thompson and Rosch 1991; Clancey 1997); in neurophysiology (Maturana and Varela 1980; Varela 1984; Maturana and Varela 1992; Clancey 1993); in education (Lave 1988; Greeno 1989; Lave and Wegner 1991); in anthropology (Bourdieu 1977; Bourdieu 1990); in linguistics (Lakoff 1987); and in ethology (Hendriks-Jansen 1996). There is no canonical model espoused by these authors, but I believe they would all agree on limiting the significance of representation, and increasing the significance of the interaction of agents of limited cognitive facilities with structured environments, in an explanation of intentional activity.

Information Processing / Symbolic Approaches to Intentional Activity	Situational / Interactional Approaches to Intentional Activity
The agent senses the world. The agent uses its world model (declarative knowledge employing symbol-object representations) to determine the current state of the world. It deduces a set of transformations that will change the current state into its intended goal state. These transformations are collected into a plan which is used by its execution systems to implement the changes of world state. Action thus consists in the information processing that is required to translate from the real world to the symbolic internal model of the world.	The agent senses the world. The agent uses goal-dependent, attention-dependent feature extraction processes to resolve these direct sensations into situations. Actions are implemented by executor systems in direct response to situations. using non-deliberative processing, and procedural knowledge represented in indexical-functional form. Sensing, processing, and acting may not be serial but may be distributed and parallel. Goal states are attained as a consequence of the agent's actions and the structure of its environment. Intentional activity consists in the interaction of situated agents of limited cognitive abilities with structured environments.

Table 7.1: Summary of the Two Main Approaches to Intentional Activity.

An extensive, and formal study of the role played by the environment in reducing the cognitive burden of human and machine intentional activity has been made by Agre and Horswill (Agre 1988; Agre and Horswill 1992; Agre 1995a; Agre and Horswill 1997) and it has been followed up by other authors (Hammond, Converse and Grass 1995; Kirsh 1995). Agre emphasises the importance of *routines* in the explanation of everyday purposeful activity:

> A routine is a frequently repeated pattern of interaction between and agent and its familiar environment (Agre 1997, p108)

They are neither obligatory (law-like) or consciously maintained. They depend in essential ways on the structure of the environment in which they occur for their maintenance, and on improvisation and opportunistic evolution, for their development. In a very similar vein, but with a larger unit of analysis, Bourdieu (Bourdieu 1990) argues that cultural practices are shaped by the environment in which they occur and are reproduced by their shaping of the environment. Von Foerster's notion of "eigen-behaviours" is similar (von Foerster 1984). The idea that much of intelligent behaviour can be attributed to structure in the environment was early on articulated by artificial intelligence pioneer Herbert Simon (Simon 1969, p4) and then largely ignored by mainstream research in that field.

The embedding of human intentional activity in the social environment is also explored by the Russian "activity theory" school of psychology (Vygotsky 1978, orig. 1934; Nardi 1996) and social constructivist sociology (Collins 1981; Collins 1990). The former emphasises the importance of social interaction and cultural artifacts, such as tools, in learning and cognition. The latter argue, essentially, that if intentionality is an outsiders ascription, then what constitutes a correct and competent performance is a social construct, rather than intrinsic. The idea that the environment, and indeed individual entities not necessarily human in the environment, can be viewed as agents has been explored in the "actor network theory" of sociology (Callon 1986) and technology studies (Pickering 1995).

This different view of the relationship of system and environment and causality has also been articulated by the second order cybernetics school (Maruyama 1963; von Foerster 1984), particularly in the biological writings of Maturana and Varela (Maturana and Varela 1992). It involves a change in view from the traditional input / output view of the open system, to one that is "operationally closed" by its interaction with the environment ("structural coupling"), at least as far as causality is concerned. Their notion of systems that reproduce their own structure ("autopoiesis" (Mingers 1995)) is also relevant. Agre (Agre and Horswill 1997) has coined the term cognitive autopoiesis to describe the relation between individual human cognition and the enabling, reproductive effect of structured

environments and cultural artifacts. The situational theories of activity can be viewed as an attempt to define the place of intentionality within this new systems perspective.

Issue	Information Processing / Symbolic Stance	Situational / Interactional Stance
Ontology	Objectivist	Not necessarily objectivist
Representation	Symbol-object Centralised Unique	Indexical-functional Fragmentary
Nature of the world	Deterministic Predictable	Not necessarily deterministic
Agent's relation to the world	Detached Autonomous	Situated Embodied Thrown
World-view implemented	Transcendental "aerial"	Agent-centred "ground"
Agents knowledge of the world	Declarative "know-that"	Procedural "know-how"
Cause of action	The plan	The agent's situation
Locus of agency	The focal agent	Interaction between agent and environment
Relation between sensing and acting	Mediated by deliberation upon representations	Mediated by direct interaction
Agent architecture	Serial Functional	Modular Parallel
Resources	Unbounded	Bounded Time-contrained
Causality	Simple proximal cause-effect chain	Mutual
Systems perspective	Open (input-output) Strong separation of system and environment	Operationally closed Weak separation of agent and environment
Underlying metaphor for activity	The deliberative thinker	The skilled artisan
Promise	Optimal and reliable goal attainment in a wide range of environments	Robust, fluent behaviour in complex dynamic environments under severe resource constraint
Potential downfall	The complexity and brittleness of the symbolic translation process	No guarantees of goal attainment Activity may be too domain specific

Table 7.2: The Differing Stances Taken by the Planning and Situated Action Models on Various Issues.

The idea that the representation approach to intentional activity arises from the mistaken application of the view of a theoretician observer has been strongly argued by Clancey (Clancey 1991; Clancey 1993) (see also (Slezac 1992)), Agre (Agre 1995a), Maturana and Varela (Maturana and Varela 1980) and von Foerster (von Foerster 1984).

The notion that purposeful behaviour might be emergent from the interaction of multiple agents of limited cognitive powers, and possibly conflicting goals, has been explored in a number of disciplines: in robotics (Brooks 1986; Maes 1990); in the field of distributed artificial intelligence (O'Hare and Jennings 1996); in Minsky's 'society of mind" notion (Minsky 1985); in artificial life (Resnick 1997); complex adaptive systems (Holland 1995); and in ethology (Hendriks-Jansen 1996).

Part 3

The Significance of Theories of Activity for Operations Management

Chapter 8

Theories of Activity and Operations Management

In this last part of the thesis I argue that operations management can be viewed as the design of an intentional agent - the operations management system. At first sight this may not appear to say much, only that management is about goal achievement, which is reasonably uncontroversial. But coupled with the argument that I have made in part two, that at the activity level of abstraction, the design of intentional systems raises a common set of issues in *all* intentional systems, it has considerably more power. It means that we can borrow fruitfully from other disciplinary fields where the nature of intentional activity has been discussed with more precision, and where the process of intentional agent design has been more explicitly specified. As I have argued extensively in part two, designs for intentional systems are informed by theories about the nature of on-going, goal-directed activity. We would expect this to be so also for management systems. However, there is practically no tradition of discussing management systems at the activity theory level of abstraction. It is my purpose in this part of the thesis to undertake that project.

8.1 Operations Management Systems as Intentional Agents

I have already argued in chapter 5 that any complex system can be viewed as an intentional agent engaged in some goal directed activity. This is an observer's attribution and is to be judged on whether it adds to clarity, whether it simplifies describing and predicting what the system will do, and most importantly, if it suggest principles for the design of such a system. The new element encountered in treating *management* systems as intentional agents is that they include people as part of the system, and people are themselves intentional agents. There has been a certain degree of controversy about attributing goals to systems consisting of people who themselves have goals. I will argue that the strict adherence to the idea that activity discussions about systems are a particular level of description, not to be confused with other possible levels, together with the approach to intentionality expounded in part 2, helps avoid a lot of these problems.

The first point is that the notion of a management systems as an intentional agent is meant as an abstraction, formulated at the activity level of description, which is adopted for a specific analytical purpose, namely, to enable a discussion of how managements initiatives and systems designs are informed by theories of activity. It is not intended to provide an account of *all* issues of management, or even all aspects of organisational goal setting, but only issues about how systems can be designed to reliably achieve goals through organised on-going activity. There will be many other important topics that will not be touched by such a discussion, such organisational politics, culture or collective sense making, the connection between individual and collective aims, organisational values, ethics and responsibility. The separation of a multitude of other management issues from those of intentional systems design, I believe, is a virtue of the approach. The activity level of abstraction, like any abstraction allows simplifications to be made in an explicit and analytically controlled way. Fruitless discussion, at cross purposes, can be avoided by carefully accounting for the levels of abstraction involved. From the point of view of these other perspectives, treating the whole socio-technical management system as a single agent might appear to be treating the humans in the system in an impoverished way. However, I will show in this and the subsequent chapters that this abstraction is quite adequate for my purposes. The frequently voiced objection to treating organisations as systems,

that the approach writes human factors out of the story, will turn out to be only true for representational approaches to activity: within situational / interactional approaches to activity, the human factors re-enter the picture as aspects of the systems environment, or what constitutes the situation, as will be seen in chapter 10. The second point is that the present approach to intentionality as an observers ascription greatly clarifies a discussion of goals in socio-technical systems.

I will briefly justify the intentional systems approach to manufacturing operations management systems, by showing that the four steps in taking this descriptive stance, set out in section 5.2, are reasonable in this case.

8.1.1 The System / Environment Split

My proposal is to treat the management system that controls the flow and timing of materials through a batch repetitive manufacturing organisation as the focal system. This system is a socio-technical system which includes, all the staff and line workers who are involved in deciding what to make, when, and how, in scheduling and sequencing the work, in actually making it, and in organising the replenishments of materials. It also includes all the technical resources used in these processes, such as procedure manuals, computers, computer files, tools and machines. This is the operations management system. The environment of this systems includes: the customers (or more specifically their requirements); the commitments of the suppliers of materials; the products that are being made; the method of manufacture; the physical layout of the place of work; and the social, cultural and political setting in which the operatives work. Certain high level policy parameters, such as levels of acceptable inventory holding, batch sizing policies, and manpower capacity may also be thought of as part of the environment.

To justify treating these things as a *system* we merely have to note that this collection of people resources and skills has been collected together because they can do something that the uncoordinated actions of the individuals or machines cannot. This indicates that the collection is greater than simply the sum of its parts. The system is an *open system* because it exchanges material, energy, and information with its environment. Finally, the focal entity just described occurs in similar form in a wide varieties of firms, which justifies its naming as a *thing* for general study.

The split between system and environment may be a bit more controversial. In discussions of this kind, the whole organisation is often taken to be the focal system and the environment is taken to be everything extra-organisational. For my purposes, I want to discuss what might be loosely termed the work system, so I want the environment of work to include the place of work, the things being made, the power structure in which work is done, and the cultural environment of shared meanings. The design of the products and the way they are made are viewed as environmental constraints on the work system because they are not routinely changed in the cause of operations. For this focal system viewed as an agent, the sensing systems are those that acquire information about this environment, including customer requirements, availability of materials, and the current state of materials conversion process such as inventories, and work in progress. The executor sub-systems are the people and machines that make the products. As with other agent theories, the sensing and execution systems are both part of the agent and part of the environment. I want to consider the machines as used by the workers to make product as part of the system, but to also consider them, to the extent to which they constrain activity, to be part of the environment. For instance, the *physical layout* of machines is part of the environment.

8.1.2 Activity

The second step to an activity level description is to be able to conceive of the focal system as doing something, as being engaged in an activity, whose description does not simply reduce to a description of the motions of its components. This is clearly the case here: the activity is making a range of brand products whose nature is relatively stable despite many of the details of the market environment, the

supply environment, and the processes performed differing on each occasion of the activity. We routinely speak about, and wish to theorise about, the making of products at this high level of description: "How do we get this product out tomorrow?" does not ask for a specification of the movements of humans and material objects, but for action on such high level abstractions as "manpower", and "resource allocation".

8.1.3 Autonomy

The ascription of intentional agency also implies a degree of autonomy for the focal system: there is no point intending something if one is not free to do it. As we have seen in part two, this autonomy is expected to be only apparent at the activity level of description, as we know that at the physicalist level it is not literally true. It amounts to the notion that the choice of actions is sufficiently arbitrary that they may be viewed as arising from the agent's sensing of the environment, and as being directed toward goals. For a management system, this amounts to the assumption that activity within the firm at this high level of description, can be controlled, directed, and reconfigured towards certain ends, rather than just being a passive response to environmental interaction. This assumption is central to any kind of management reform initiative. Many would argue that this over estimates the amount of influence that managers actual have on what happens in firms. I think this is beside the point, and is a confusion between descriptive levels. This autonomy assumption is made at the activity level of description: it is at the stage of testing and evaluating management theories that follow from different theories of activity that their ability to achieve the difficult translation form the activity level to actual human and machine behaviour will be tested. In other words, management autonomy is an abstraction used to speak about management theories and system designs, rather than a description of the state of affairs at the physicalist level of description. While not exactly the same point, we also want operations systems to be largely automatic and not require constant intervention from management.

8.1.4 Goal - Directedness

The final move in attributing intentionality to systems is to assert that the focal system, at a certain level of description, can be thought of as an agent pursuing some goal. We must be careful about what is meant by at "some level of description". The "agency" being referred to is the agency of the *collectivity* which is the system as a whole. This level of description is to be distinguished from the level where individual workers or perhaps even machines are viewed as agents. Some sociologists (eg. (Silverman 1970)) have argued that it is an invalid move (a so called "reification") to treat an organisation as having goals, and that goals where they exist are always those of individuals. This amounts to denying that organisations can validly be viewed as intentional systems. This view, by placing a high weight on the actions of individuals as the explanatory unit in organisations, amounts to a physicalist or low level descriptive stance. This kind of argument is usually made as an argument against the systems view of organisations, on the grounds that it tends to obscure the essentially political and possibly coercive nature of organisations. A discussion of the issues of ascribing to organisations goals is found in Hall (Hall 1980), Argyris and Schon (Argyris and Schon 1978), and in the context of organisational responsibility by (Weaver 1998). Similarly, some authors question under what circumstances organisations can be said to perform actions. Are they the actions of individuals and do individuals have to act in the organisation's name for them to be viewed as arising from the organisation? Does an organisation have to have a formal decision making process to count as intentional? Is a riot an intentional agent?

The ascriptive approach to intentionality I have taken in this thesis, and a rigorous accounting of descriptive levels, sidesteps a number of these thorny issues. Firstly, the proposal to treat management systems as intentional is an abstraction, and it is not supposed to explain everything. It does not explain how goals are negotiated in organisations, or how dysfunctional actions of individuals might arise, and be dealt with. Secondly, many of these conceptual problems arise from the assumption that

intentionality has to be an intrinsic property of the entity in question, in the way that human intentionality seems to be. If we take the view that intentionality is a descriptive stance of an observer of the entity, many of these problems do not arise. A riot *is* an intentional agent when viewed from a police helicopter. Looked at from outside, it is difficult to reject the view that manufacturing systems are "about" something, and that an explanation of how they do it would be appealing. The goal is to bring materials together to form products that customers will buy, and it seems rather natural to adopt the standard measures of goal attainment: cost minimisation, timeliness, quality and efficient use of resources (Hayes and Wheelwright 1984). Without this ascription it is hard to account for why the firm exists at all.

This does not involve saying that the entity in question *has* goals, merely that its behaviour can be explained and predicted in a parsimonious way by viewing it as pursuing goals. On this view, the goal it is pursuing may not even correspond to the goals of any individual or group in the organisation. If one is discussing politics, say, this may not be a useful descriptive stance and indeed, may lead to some important issues being hidden, because the important topics of discussion are probably best framed at the lower descriptive level which emphasises the agency of individuals. My argument is that it is quite a different theoretical project to ask: "Given that management systems appear to behave like intentional agents, what are the implications of this for their design?" to: "What is the process by which the goals of the manufacturing company are negotiated?" and their separation, although it can never be completely clean, will clarify many important issues about system design that might otherwise be clouded by conflating these two questions.

Most of the confusion surrounding the validity of intentional descriptions of organisations result from conflation of levels of abstraction. Actions of the manufacturing system are an observer's classification of high level recurring patterns of interaction with the environment, such as producing products, fulfilling customer orders, and replenishing supplies, and are not simply associated at this descriptive level with individuals movements, or even with what would count as individual actions in an intentional description of individuals. Design, and by implication management, is conducted at this activity level of abstraction. As with physical robots, making the connection between this high level conception of actions, and the physicalist level behaviour of individuals, will become a problem at implementation time, and different theories of activity will give more or less compelling accounts of how this can be done.

8.2 Management and the Manager

We saw in part two, that certain things about intentional systems are inherently outside the descriptive framework, namely, the goals themselves and the theory of activity itself (which is embodied in the agent design). In the view I have been espousing, these are the ascription of an outside observer. So, in the spirit of treating operations managements systems as intentional agents, I now associate the *manager* with this outside observer, and *management* with goal setting and agent design. Again, this is intended as an abstraction. Interpreted literally, that is, by incorrectly applying it to an individual in the organisation, it appears to fall victim to some obvious management myths, namely that the manager (management team) is essentially *external* to the system, and the manager (management team) can be viewed as a single external agent with a single intention. Neither of these assumptions is literally correct. The first would mean that the manager is some kind of ethereal being not part of the system or the environment, without embodiment, environmentally shaped values or personal goals, and the second would mean that there is no managerial politics. But questioning these assumptions is the topic of other projects, such as the study of organisations power (Pfeffer 1981) and critical systems theory (Flood and Jackson 1991).

My argument is that this abstraction is natural and good enough for the purposes of this thesis. I want to discuss the influence of theories of activity on approaches to management and management systems design. Facing the problem of designing a complex system, such a hypothetical manager must make

important decisions, such as the way the system will be decomposed into parts and the way control systems will be implemented, which are determined by the theory of activity adopted. It really does not matter for this project where goals come from, or indeed, what they are. It is not important who designed the system, or indeed, whether is was actually designed by one person. So it is adequate to lump these things under a single heading "manager". Again, when it comes to looking at how various management theories stand up to being implemented, the inevitable need to reconcile the abstraction and the physicalist description will arises, and need to be considered. We will find that situational activity theories have something extra to offer here, because they allow for the possibility that an intentional agent might consist of sub-agents with different and even conflicting goals.

8.3 Aims and Methods

My proposal then, is to consider operations management as the design and control of intentional systems, and to see what light theories of activity can shed on the design process and the nature of control, and thus on the likelihood of success of various operations management approaches. These are the tools that can be brought to bare on such a project:

1. We can examine various well defined approaches to operations management and try to determine which of the theories of activity laid out in part two appears to inform then. Because approaches to activity follow rather directly from assumptions the analyst or designer makes about the world of activity, the nature of agents, and the relation between the two, one approach is to examine management systems in terms of the assumptions they appear to make. These are revealed in the mechanics of the approaches, and in the writings of their advocates. We can also look at the *process* of activity proposed, that is, the way the approach proposes to relate goal directed actions to sensing, knowledge, and goals. Agent architecture will be an important clue here.

2. We can start from first principles and try to determine how management of a socio-technical system of production would be conceived on the account of each particular theory of activity. This would involve determining an appropriate role for managers, and for operatives, and determining the points of leverage upon on-going activity that the theories suggest for management. These points of leverage will be determined by what the designer considers to be the primary *cause* of actions.

3. We can look to other areas, such as robotics and software science where the building of intentional systems is a routine affair, and try to make use of their accounts of the design process.

In the next two chapters, which deal with specific conceptions of operations management, I will apply each of these approaches.

Chapter 9

Representational Approaches to Management

In this chapter I argue four propositions more or less concurrently:

1. That adoption, explicit or implicit, of the planning model of activity leads to a particular formulation of the problem of operations management, and its solution as the automated generation, implementation, and enforcement of plans. In paper two (Johnston and Brennan 1996), Brennan and I called this management approach *management-as-planning*.

2. Manufacturing management techniques from the Computer Aided Production Management paradigm, MRP II in particular, are examples of the management-as-planning approach.

3. The assumptions about the world, the feasibility of representation, about the nature of agency, and so forth, that are imported to management-as-planning from the planning model, while apparently plausible, do not stand up under closer examination.

4. The invalidity of these assumptions in any reasonably complex dynamic environment, and therefore the inappropriateness of the planning model of activity as a basis for a theory of operations management, is a likely explanation for the difficulty companies have had making MRP II work.

My method is to first show how the planning model of activity leads to this particular approach to management and then to characterise that approach. I then recap the process involved in MRP II and show that it is a very literal translation of the planning model described in chapter 6. I then show the attitudes to the formulation and solution of the management problem, that I have defined as management-as-planning, are present in the writing of Oliver Wight, a leading advocate and namer of the MRP II concept, in his text (Wight 1981). I then examine each of the assumptions of the planning model of activity to show that they are present in MRP II, and where necessary, consider other CAPM techniques. I also present a critique of the validity of each assumption in the batch repetitive manufacturing domain, and to save repetition, I do that also in the same section. The sources I draw on for this last exercise are the case studies presented in chapter 4, the text of (Wight 1981), other published research papers, and to some degree, my own consulting experience with MRP II.

My aim is not to prove that MRP II is "wrong", and indeed, the data I use would not be sufficiently generalisable to do that, if it were my intention. Rather, in keeping with the aims of the thesis, I want to show that the planning model of activity is not a plausible basis for a theory of operations management in a commercial world of any degree of complexity and uncertainty, by showing how the various assumptions break down on closer inspection. I am using MRP II as a case study of management-as-planning and I would expect these assumptions to fail in similar ways in other examples (as indeed Brennan showed for the case of a public policy initiative in education, based on the approach, in paper two). In any case, MRP II is not "right" or "wrong": it is merely "consistent" or "inconsistent" and years of scrutiny testify to its consistency (its weakest area is the infinite capacity assumption which is addressed by other CAPM techniques). Like any other consistent theory, it has a domain of applicability defined by the conditions under which its assumptions prove tenable. My assertion is that this domain is not nearly as wide as is generally imagined, and in particular, does not include the centre ground of manufacturing - batch repetitive manufacturing - to which MRP II has been aggressively marketed.

9.1 Planning

This chapter is going to define and criticise an approach to management based on planning. Since planning is an almost sacred concept for management, to avoid being accused of heresy, I must first define what sense of planning I am attacking. There are at least two distinct senses in which the notion of planning is used in everyday life, as well as in management. Planning is used in analysis and simulation to examine possible future worlds, plans are used to decide what to do next in unfamiliar circumstances, and plans are used as frameworks in which improvisation takes place. In all these cases, plans are used as *resources* for action (Suchman 1987). An example is the plan for a domestic house construction. It represents a good deal of formalised thinking about the most suitable layout, materials, and methods of construction, and is intended to be a resource for the builders to use in construction. This kind of plan is a representation of the world at some grain-size constructed for the purpose of enabling analytical thought about the world and possible future states.

However, there is another conception of planning which is the idea that plans can be used as a *control structure* for action, in the sense that the plan serves as a *program for action*, somewhat as a computer program serves to control the execution of a machine process. Agre and Chapman (Agre and Chapman 1990) have called this notion of planning the *plans-as-programs* view and the former usage as the *plans-as-resources* view. This is a different sense of the word "plan", where the plan is a representation of a series of actions to be taken in a programmatic way to achieve a goal. The management-as-planning approach asserts that plans of the first type can be used for the automated generation of plans of the second type, and that this is essentially what operations managements is about. It is as if the plan for a house were used to generate a schedule and sequence of actions for the builders, and this was taken to solve the problem of building a house. Although the house plan might contain some information on the steps to be taken in construction it does not specify a complete set of *procedures* for building it. It seems particularly absurd to suggest doing this for a domestic house construction, yet the feasibility of doing this is routinely asserted by the management-as-planning approach to project management and operations management. I want to concentrate on operations management because its repetitive nature lends itself to the consideration of alternative theories of activity.

In paper two, we argued that management-as-planning is based on a misconception of the nature of on-going activity, that is already present in its informing theory of activity, that *thinking about acting* is essentially the same as *action*. Plans as resources are an important part of management activity that is directed towards thinking about the future, seeing opportunities, and positioning against threats. I am not here, and would not want to, argue that plans are not valuable *resources* for management. But to take planning to be the basis for managing action, through the automation of plan generation and implementation is quite another step. This amounts to taking what managers do best, which is thinking about what to do, to be a plausible basis for the way the organisation itself should do things. This error is exactly the "theoretician's error" which arises from a confusion between the aerial and ground views as a plausible basis for activity.

9.2 Management-as-Planning

Whatever criticisms we may raise against the planning model of activity, it does lead to an obvious formulation of how to manage on-going, purposeful activity. In the planning model all orderliness of activity is contained in the plan, so if managers are responsible for imposing order, then under the planning model they are also responsible for *producing* the plan, or providing the tools and systems to do so, which closes the gap between the current state and the goal state, and for *enforcing* the plan. This then is the image of management if one assumes the planning model is a correct explanation of how on-going, purposeful activity is produced:

The problem: In the absence of management, the efforts of operatives to deal with a changing, complex environment is disorganised, chaotic and inefficient.

The solution: The efforts of operatives should be coordinated and directed toward a common goal. This is achieved by means of a plan to which all must adhere and which provides the information for each operative to know what to do.

What managers do: Managers set goals and produce plans that logically lead from the current state to the goals state. Managers measure adherence to the plan, and define the criteria for measuring adherence. Managers enforce operator adherence to the plan by rewards and sanctions.

What operatives do: Operatives execute the plan. Their actions are determined by the plan rather than their situated experience. Thus they are not viewed as agents themselves, or more precisely, their individual intentionality is not a part of an explanation of the overall system's intentionality. Operatives are accountable for there adherence to the plan. The performance of operatives is measured in relation to the plan since the plan defines the *correct* (logical) series of action to go from the current state to the goal state.

Role of the environment: None. All pertinent information about the behaviour of the operational system and its environment is captured in the plan, which is executed against the passive backdrop of the environment. Any effect that the environment may have on activity is *incidental*.

The rigorous separation of planning by management, and execution by operatives, is of course characteristic of Taylorist management theory. It is also exactly the same relationship that Descartes saw between the mind and the body. I argue that the connection is their common acceptance of the planning model of activity, and its aerial view. Managers are the "mind" of the organisation, and their decisions are implemented by a mechanical, robotic organisation. The attraction of this conception of management is that it seems to hold out the hope of a kind of all-purpose intelligent behaviour of the organisation which, in particular, promises optimal or efficient use of resources. Thus, the metaphor of organisation behind this conception is that of the deliberative behaviour of a thinking person ("organisations as brains" (Morgan 1986, p77)).

9.3 CAPM as an Example of Management-as-Planning

9.3.1 The MRP II and Planning Model Processes Compared

There are a number of computerised planning techniques that have been developed to tackle the manufacturing operations managements problem which I have grouped under the rubric of Computer Aided Production Management, CAPM. Examples are Manufacturing Resource Planning, MRP II (Wight 1981), Optimised Production Technology, OPT (Jones and Roberts 1990), and Finite Capacity Scheduling (Goldratt 1988; Taal and Wortmann 1997) or Schedule Based MRP (Hastings, Marshall and Willis 1982; Hastings and Packham 1985; Yeh 1997). All the techniques adopt a common planning based approach to the management of manufacturing operations, differing mainly in the detail of their models of the production process and in the complexity of their planning algorithm. In the rest of this section I will mainly use MRP II as an example of the approach, but everything that is said here applies equally to the other CAPM techniques, which I will only mention explicitly when their differences are pertinent.

In MRP II, the manufacturing processes for the products made by the firm are modelled in a set of computer files which can be called the Product Structure database. This usually consists of the bills of material and routeing, but these may be combined into a single file (Goldratt 1988; Hastings and Yeh 1992). These files describe, in terms of network or tree data structures, the relationships between finished goods and the subassemblies, purchased items and manufacturing processes that go into making them, including required quantities and lead-times. The production goals of the company are

formalised into a Master Production Schedule (MPS). This involves a high level task, generally referred to as Sales & Operations planning, involving production, marketing, engineering, and finance management, and is often supported by additional computerised planning systems. The MRP II literature (Wight 1981) stresses the need for all levels of management to commit to the MPS as a statement of the company's goals. The Master Production Schedule is revised periodically, usually monthly.

On a more frequent basis, say weekly, the Master Production Schedule is processed against the Product Structure database to produce a time phased operational level schedule for the initiation of purchase orders and production orders. In MRP II, this process is relatively simple but time consuming: knowing what materials go to produce any assembly, the quantities required and the time it takes to manufacture or purchase the items, it is possible to work back from the end item requirement to a schedule for purchases of bought items and manufacturing of intermediate sub-assemblies. These schedules are then distributed to purchasing and shop floor personnel who are expected to work to them. In the process of producing these schedules it is necessary to net the requirement against the current inventory of purchased items, assemblies, and finished goods found on the inventory file, and existing purchase, production and customer orders stored on appropriate open order files. In principle, this data only needs to be updated just before the planning run. However, there are generally computer systems devoted to maintaining records of these data on at least a daily basis for other control purposes.
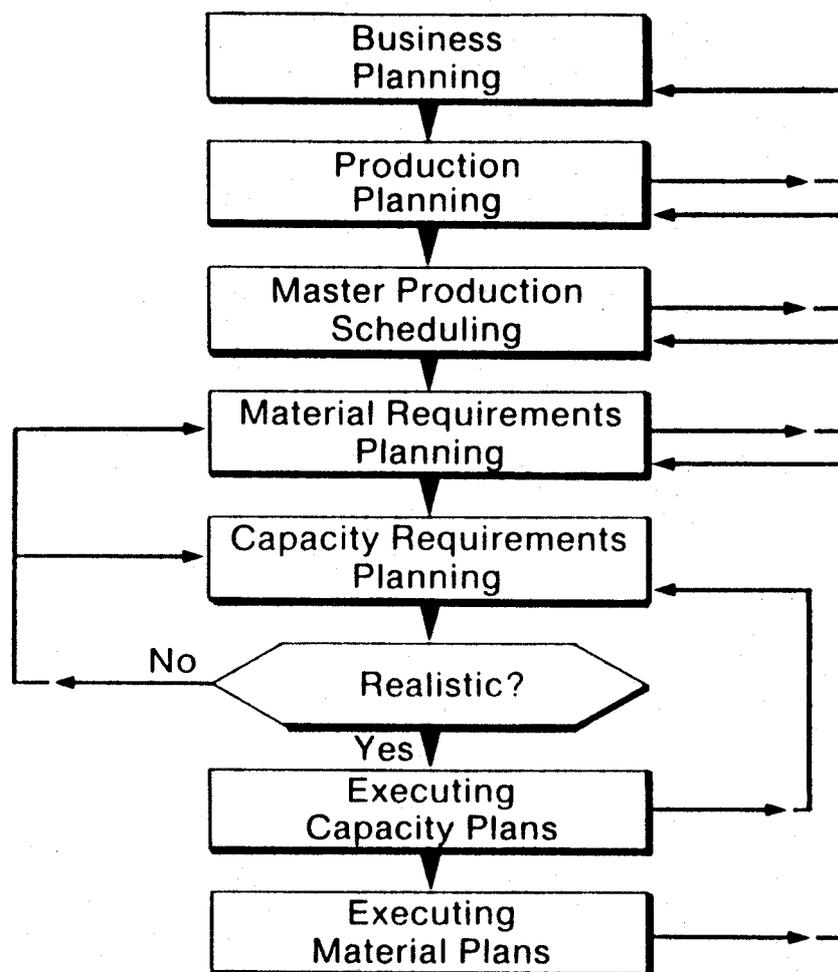

Figure 9.1: MRP II according to Wight (Wight 1981, p54).

In MRP II, the explosion process takes no account of the finiteness of production capacity and the feasibility of the operating schedules has to be checked, using Rough-Cut Capacity Planning, after it is produced, possibly resulting in modification of the MPS. Most of the other CAPM techniques have

been devised to tackle this constraint of finite capacity (both overall and at each machine) of production. When this is taken into account the explosion becomes extremely complex and heuristic methods are invariably required.

Figure 9.1 shows this process according to Wight (Wight 1981, p54). It can be seen to be a top-down process which converts high level management goals into executable plans. Wight envisions feedback at all levels to ensure the feasibility of the plans. It is actually a double loop process. In the first loop, high level management policy (persistent goals) and market reconnaissance is used to derive discrete goals expressed in the MPS. In the second loop, the MPS is processed against existing inventory and open orders to arrive at executable plans
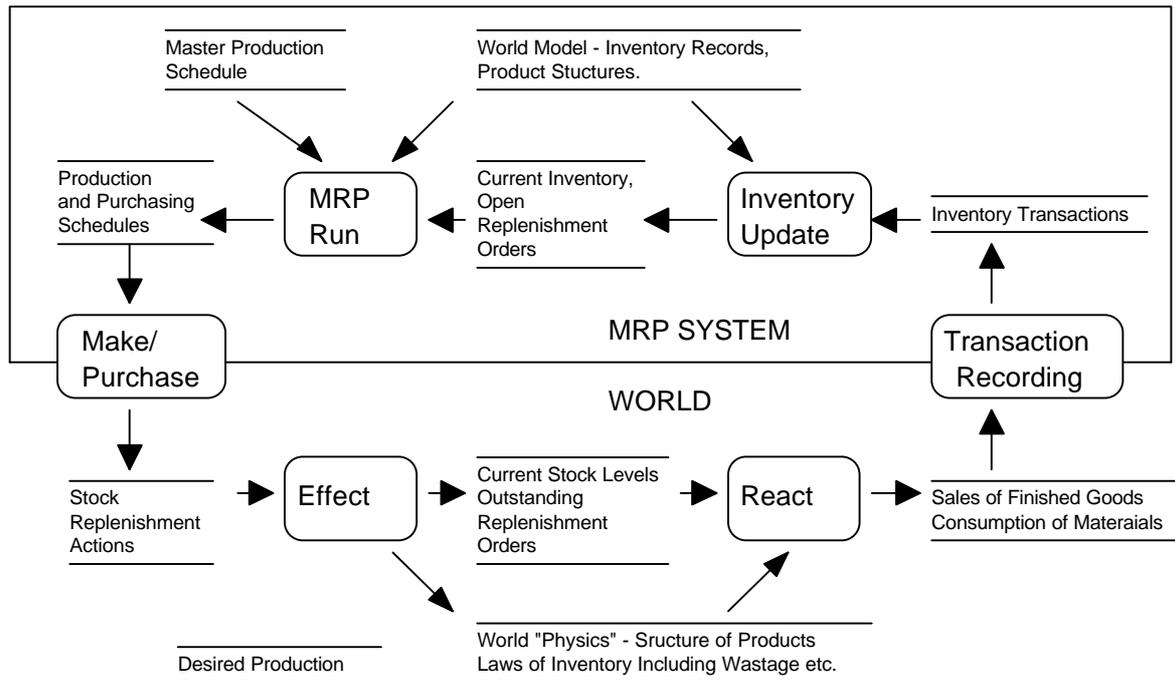
Figure 9.2: The MRP II Process as an Example of the Planning Model of Activity
For simplicity, this diagram is for the case of a make to stock situation, so that sensing only involves recording of sales and inventory transactions

Figure 9.2 expresses the second loop in the form used to describe the planning model of activity in chapter 6, for the somewhat simplified case of a make-to-stock company. With the following identifications, MRP II can be seen to be a literal translation of the planning model:

Persistent goals   Management policies

Discrete goals   Master Production Schedule

World model   Implicit arithmetic of material combinations and inventory
                           embodied in the data structures of the
                               Product Structure database
                               Inventory files
                               Open order files
                           and the algorithms of their processing routines.
                           (Wight refers to this as the "fundamental manufacturing
                           equation" (Wight 1981, p56).

Representations   Data records of
                               Product Structure database

Inventory files
Open order files

Sense data   High level market reconnaissance
                          Inventory and other transactions

Plan   Purchase and production schedules

Execution   Execution of schedules by operatives

9.3.2 The Image of Operations Management under MRP II

The problem: Wight sees the manufacturing operations problem as operatives reacting to immediate crises, rather than engaging in coordinated behaviour. This leads to loss of efficiency (productivity). (In this and later sections all emphasis in quotations is from the original, and all quotation without a citation are from (Wight 1981)).

> One of the problems is the lack of a company game plan that everyone can work to. As a result, a great deal of effort is expended, but it isn't directed to well defined goals. This means that the productivity of the entire organization is effected. p17

Indications of this chaotic activity are constant expediting of orders;

> Purchasing says, "we bring material in when somebody expedites it. If we actually worked to the dates on the purchase orders and brought in all of the past due material tomorrow, you wouldn't know where to store it." p15

working to a shortage list rather than a schedule;

> The shop foremen say, "the schedule is a joke - we work to the shortage list; and anyway we don't have the material we need from purchasing." p15

and particularly, the appearance of multiple informal systems which actually run operations.

> A company cannot measure performance against the numbers from the *formal* system when the really run with an *informal* system. p31.

The solution: Wight clearly believes that the solution to the problem is to coordinate the actions of operatives through a plan or schedule:

> ... the first problem that must be addressed is scheduling, because in a manufacturing company: *Scheduling is fundamental.* p19

> ... the problem is lack of a valid schedule, and for that matter, the lack of the tools to make a valid schedule. p38

The way to achieve coordination is to simulate reality:

> Since a good system is basically a simulation of reality, it can be used to simulate what would happen if various policy decisions were implemented. p53

The plan is produced by management and executed by the operatives:

> MRP II makes it possible to translate overall management plans as expressed in the business plan and the production plan down to the detail level plans that everyone else will be working to execute. p87

When a state of order is achieved there will be no reacting to crises:

> If MRP is being managed properly, *there should not be a shortage list*. p131

What managers do: Wight states explicitly the functions of management:

      1. Defining Objectives
      2. Assigning Accountability
      3. Developing Understanding
      4. Providing the Tools
      5. Measuring Performance
      6. Providing Incentives. p107

Management is responsible for creating and enforcing a formal system of work, embodied in the plan.

> The CEO should set a very high priority on using the formal system to manage. This means keeping the numbers right, making valid plans, ... p131

> The CEO has a stewardship, a responsibility for keeping the formal system operating properly. p132

Management sets the goals of the organisations.

> One of management's prime responsibilities is to assign proper objectives to each function in the organisation so that each function is working to the overall objectives of the organization. p108

and holds operatives accountable to it.

> The function of the planning department is: Making valid plans that other people will be held accountable for executing. p195

The fundamental tool for holding operatives accountable for their actions is the plan.

> Accountability implies measurement, and a prerequisite to measurement is a valid plan. p111

What operatives do: Operatives execute the plan and a good operative is one that executes the plan exactly as prescribed.

> Undoubtedly, more and more emphasis will be placed on *execution* in the future because that, in the final analysis, will make the difference between the good users and the excellent users.

They do exercise judgement when their knowledge indicates a flaw in the plan,

> There is nothing wrong with having the computer generate a vendor schedule - but the planner should be held accountable for approving the schedule. p100

and they are expected to provide the information to correct it,

> Another concept that is vital to the "closed loop" approach is the feedback to indicate when there are significant deviations from the plan p385

but their individual intentionality is driven by the goals of the organisation as set by management and transmitted by the plan.

> All of these *objectives of the business are* the *objectives of all of the employees* of the business *at all times*. p147

The role of the environment: Judging by Wight's text, the nature of the situation of work has no influence on the outcome of the operations system. It is nowhere mentioned in the book. Although he does distinguish different types of manufacturing firms (Wight 1981, p18), nowhere is plant layout mentioned. Similarly, other aspects of the work environment - political structures, culture, worker skill - are not mentioned, and all we can presume is that the work situation has no causal effect on the outcome of operational activity. I argue that this is because he sees the operations management problem as purely informational, that is, as the *coordination* of action through the dissemination of *information*. It could be argued that he does not include these issues because he sees them as part of another topic, but that is exactly my argument. The only way the situation of work can be excluded from a discussion of operations management systems is if the formal models (bill of material, routeing, etc.) are assumed to contain all the relevant structure of the environment in informational form, and consequently any interaction of the focal systems with its environment is already accounted for in the plan.

9.3.3 Assumptions of the Planning Model Illustrated in MRP II

If the assumptions of the planning model of activity are checked against the MRP II model, as represented in MRP II practice and writings, it is found that they are all present and it is possible to see limitations to their validity. For each assumption, I will demonstrate its presence in MRP II, and then indicate ways in which it is likely to fail. Again these criticism apply generally to CAPM techniques, but I will refer to MRP II specifically in most cases. I engage in this rather lengthy discussion here because the common sense nature of these assumptions requires something of a "deconstruction" if the limits of the management-as-planning approach are to be appreciated. I would expect these assumptions to fail in other applications of the approach in similar ways, as indicated in paper two.

Nature of the world

Objectivism: MRP II assumes an objectivist ontology: the entities that make up its world of action - customers, suppliers, products, inventories, customer requirements, product structures - are assumed to be definable in terms of fixed sets of context free attributes. This is forced on the approach by the choice of representing them as records in the various associated computer files. The very process of Structured Data Analysis (Yourdon and Constantine 1979) presupposes such as objectivist ontology (Klien and Hirscheim 1987). This assumption would appear to be safe in the world of products and inventories. There are actually a remarkable number of areas where this objectivity is difficult to maintain even in assembling goods to customer orders. These problems are well known to MRP practitioners, but are usually regarded as annoying details rather than problems of principle.

1. **Features and options:** Business is about making the products customers want rather than the ones that are convenient to make. The manufacturer thinks of the product range as a certain finite set of generic products which are given distinct product codes. To the production manager these are the objects to be made, and for production planning these are the products to be scheduled. But the customer very frequently wants the addition or substitution of various components or materials to these standard products. These are known as "features and options" and include such things as colour choice, optional extras, custom features such as the customer's own logo, a choice of handles on a generic tap. So the database designer faces a dilemma. All the possible variants so implied can be treated as separate products and given a separate product code and product structure. This approach has severe disadvantages. Firstly, only a small subset of all possible combinations may ever actually be ordered but in this approach, they would all have to be included in the databases to allow for the possibility that they may. Secondly, treating these variant as fundamentally different products results in splitting demand, making forecasting and planning more difficult. It makes more sense internally to have demand collected for the generic product types than split over all possible minor variations, because all internal decisions, such as capacity, tooling, routeings, forecasting, etc., are more closely associated with the generic products. On the other hand, if the structures do not reflect the exact product, invalid recommendations and commitments will be made.

2. **Substitutes:** It is frequently the case that there are several grades or suppliers of materials available, which are all suitable for the product in question, but which differ in ways that are significant enough (cost, quality inspection criteria, etc.) to demand that they be treated as different materials and allocated different product codes. The buyer wants the flexibility to use a substitute material if this is necessitated by temporary changes in supply. Should the product structure database contain a structure for every affected end-product using each of the different substitutes? If so, they will all require different product codes and be different products, yet they are interchangeable from a selling point of view. Again this will split demand for forecasting, complicate finished goods despatching, and so forth. If we use the standard structure, the MRP logic will request the wrong material to be used in production, and if, as is common, automatic relieving of component inventory is used, this will lead to stock errors.

3. **Re-works and conversions:** Frequently a batch, or part of a batch of product is found to be sufficiently out of quality specification that it cannot be sold as, or continue production through to, the intended finished product. When the process is stopped, what has been made may have considerable recoverable value and can be re-worked into another defined product. However, since it failed quality control testing at the last operation attempted, it has no name - there is no product number for these parts that are to be used in the re-work process. Also there will be no bill of material or routeing to make the down graded product from this unnamed assembly (this may involve non-standard materials and operations). We can put up a special product structure to make the down graded product from the defective assembly, which is time consuming and has undesirable effects such as demand splitting, or we can use the standard product structure for the finished good that is to be made, in which case the systems will plan for the replenishment of materials which are already in the partially made assemblies being re-worked. Similar problems arise when a customer order for non-standard product can be fulfilled more easily by altering an existing product from stock than by making it as a one-off. This is usually referred to as a "conversion" and the same issues arise with regard to the product structure to be used.

Most MRP packages include elaborate software features that address all, or some, of these problems (feature and option bills, planning bills, soft bills). All these approaches involve using a different product structure for planning purposes to that used after the production plan is firmed-up. This approach of having the nature of the planned products change over time is not entirely satisfactory, and results in problems with aggregate planning, forecasting, and variance cost accounting, and also detracts from the simplicity of operation of the vanilla-flavoured MRP approach. All the problems mentioned above have the same ultimate source, which is that what is considered by the MRP logic to be the objective products of the company, actually differ according to whose point of view of them is taken and for whose purposes: the production planner; the customer; the material buyers; method; production; or accountants.

Determinism: This assumption means that causes are supposed to have unique and predictable effects. This assumption is necessary if MRP is to allow management by simulating the future.

> MRP II is a simulator. It simulates material requirements far enough in advance so that shortages can be predicted and prevented, rather than fixed after they happen. p56

Wight sees the dynamics of a manufacturing system in rather physics-like terms: he repeatedly refers to an underlying "manufacturing equation",

> MRP, by definition, is a simulation of the fundamental manufacturing equation. p56

and running a business depends fundamentally on a valid mathematical model of activity.

> MRP II results in management finally having *the numbers to run the business*. p54

MRP is a strictly deterministic approach to management. The logic has no place for stochastic or fuzzy variables. When stochastic demand forecasts are used, they must be converted into "planned orders" for definite products with definite quantities and definite due-dates. Except in a pure make-to-order environment, these crisp "objects" do not correspond to anything really "out there" and even in make-to-order environments their certainty is not absolute. Even in internal operations there is great variability in processes that are assumed by the model to be precise. This process of modelling a largely stochastic environment with a deterministic model, which treats all variables to be equally reliable, leads to a phenomenon known as "MRP nervousness" (Hopp and Spearman 1996, p131). When the effects of variability are reported to the system it can cause massive rescheduling, often of no practical consequence, in response to insignificant variations in the environment. The standard method of combating this is to introduce extra buffers of stock or lead-time, which has the effect of reducing to usefulness of the information produced.

Change: According to Wight:

The manufacturing environment is one of constant change. There is no way for a manual system to cope with the constant changes. The scheduling systems must be able to answer the question, "When do we *really need* that material?" it must be able to establish and maintain valid due dates. p23.

Change is modelled in MRP II through the timing relationships (due date, lead-times) built into its database. Between planning episodes, the world is assumed to evolve as predicted in the schedule outputs. This assumption is clearly going to fail if the re-planning period is too long. Wight sees computing power as the solution to this problem.

Before the computer was available, companies took six to thirteen weeks to calculate the requirements manually ... {t}hey typically "ordered" once every thirteen weeks (this was called the "quarterly ordering system"). This didn't work a lot better, in many applications, than the order points. The computer made it possible to calculate requirements over a weekend! p44

Companies frequently do master schedule planning monthly and MRP explosion weekly (these were The settings used by Bendix Mintex). Many MRP installations can not schedule MRP explosion on mainframe computers more than once a week because it ties up a dedicated computer for 8 - 12 hours. This is a problem that will potentially be solved with technology, but there are other factors limiting the response time of MRP. Firstly, the responsiveness of the planning systems depends not only on the speed of rescheduling but also on how frequently the master production schedule can be revised. Given the high time cost of the management involvement in this process, there will always be a limit to how frequently it is practical to do this. The second limitation is that throughput time for the typical batch of product sets the time scale at which it is worth replanning (Ashton, Johnson and Cook 1990). If batch sizes are reduced in order to increase the responsiveness of the whole manufacturing operation, as Kodak found, the data collection required in initiating, tracking, and closing works orders, becomes prohibitive.

Boundedness: The boundedness assumption asserts that a finite and course grain model of the world can produce an action plan that will be useful in an infinitely interconnected and fine-grained real world. This is an assumption which is imported from the "thinking about action" approach of the planning model. In thinking about acting this assumption is routine and benign, because the simulated action are not required to be performed. The capacity independence of lead-times assumed in standard MRP is an extreme example of such a simplifying assumption which ignores the interconnectivity of lead-time and machine loading. The independence of optimal batch size from loading decisions is another. Are plans developed from these dubious assumptions of any use? In the case of the first assumption many would say "no" (Hopp and Spearman 1996, p175), and as mentioned earlier many of the variations on the MRP II theme within the CAPM framework are dedicated removing the infinite capacity assumption. However, there is always going to be a point where the detail missing from what is specified on the schedule has to be filled in by the tacit knowledge of those trying to implement it. To use Bateson's favourite expression, the map is not the territory (Bateson 1979, p110).

The Relationship of the Agent to the World

Representation: The planning model asserts that it is necessary and feasible to manage operations through the intermediary of a model or representation of the world.

The system is just a representation of the way it really happens - a simulation. p105

For MRP II this representation is the records of the product structure database, the inventory file, and the various open order files of product structures, material and finished goods inventories, forecast and actual customer orders, works orders and replenishments orders. These all have to be maintained at a high level of accuracy if the system is to work (Breznier 1981; Fisher 1981). (Figures normally quoted are 98% accuracy for bills of material, and 95% for routeings and inventory records (Wight 1981, p356)). The databases have different degrees of volatility, inventory files being the highest and product structure information being the lowest.

It is not a trivial matter to track inventory entering and leaving the plant at an accuracy of 95% or better. However, there are considerable benefits to doing so, other than merely to support the MRP II process, such as the elimination of frequent physical checking, and the ability to promise products on-line. So these infrastructure benefits can be leveraged to achieve this accuracy. Also the updating of this data is usually performed at traditional accounting control points by a small number of people with a strong grasp of the consequences of error.

Although the recording of inventory entering and leaving a plant is a pretty standard accounting requirement, the MRP II model requires much more than this. For the MRP II process to provide useful schedules for work release in the factory, the bill of materials must be structured as an hierarchy of sub-assemblies which are (at least in principle) manufactured and placed into stock (usually called component inventory), and withdrawn as materials for the next stage. In MRP it is common to make components in lot sizes chosen for efficiency, which may not be the same as that of the end products, so these bookings into intermediate storage are often real. In this case there are real costs involved in viewing the end-product as being made of various objective sub-assemblies, both due to the cost of recording the inventory transactions, and due to the cost of providing the stores infrastructure. Because these transactions are completely internal it is probably harder to justify and enforce accurate compliance. Figure 10.2 (next chapter) shows the subassemblies that were defined in a typical disk-pad set manufactured by Bendix Mintex around 1988. Inventory movements were recorded for each of the shaded material and assemblies, which were also physically moved to, and withdrawn from, actual stores.

The standard MRP II recommendation is that *all* inventory stores should have high levels of security and control (the "locked stock-room" concept).

> Limited access stores is a prerequisite to accountability. p201

> The cost of a stock-room is not something a company can avoid. They will always pay for stores control. The only option is whether they get control or pay for not having it. p204

In addition to sub-assembly detail in the bill of material, to provide useful progress information for trouble shooting and expediting, there must be sufficient control points defined in the routeing. A control point is an operation whose completion must be reported to the system. If it is considered necessary to be able to see the size of the queue ahead of a work centre, then "move" transactions must be processed as well as "operation completion" transactions.

Data recording associated with subassemblies and operation completion is necessitated by the representational approach of MRP II, and is essentially a *non-value-adding overhead* of the approach. The greater the fidelity of the representation to the process required, the greater this overhead of data collection. So again there is a representational dilemma here. Bendix found that they got no useful shop floor scheduling information from their MRP II system because they could not justify the cost of maintaining many control points. They also found a similar problem with the fidelity of their aggregated representation of customer requirements. It was not feasible to have the Sales and Operations team consisting of operations, financial and marketing managers, to consider the product range line-by-line, yet this was what was required if the MRP system was to work correctly. Similarly, Kodak found that when they wanted to reduce their cycle time to increase customer responsiveness, the consequent paper work of the MRP progress reporting would be overwhelming.

Unique representation: Wight insists that there must be one sanctioned system in use and informal systems must be eliminated

> A common bill of material for all functions in the company is essential for MRP II. .... A common bill of material data base means one uniform set of product structure information used by all functions and, consequently, drawing upon one set of *correct* information. p261

Informal systems to accommodate varying perspectives are to be stamped out because of their erosive effect on the accuracy of the official database.

> If the master schedule, in particular, is managed properly, there can be one system to answer the question, "What do we really need when?", rather than many conflicting systems. p51

In my experience with MRP II, maintaining a highly accurate centralised product database is one of the most problematic areas of implementing the concept. The packages usually provide storage facilities in the database only for that information vital to the operation of the MRP II system, whereas, there is usually all manner of product data and product structure data distributed throughout the organisation, which may include unstructured text information. One company to which I consulted had separate product structure recording systems for costing, material issuing, engineering change control, and quality testing. These separate databases have historical origins but also reflect different data needs and emphasis of the various functions. This is quite a typical situation and is being exacerbated by the proliferation of personal computers in companies. Therefore, rather than supporting the MRP II package's product structures as an *extra* database, it seems a good policy to modify it to allow it to *replace* all the others. This is not usually a difficult programming task but can often be an extremely difficult implementation task, involving many functional areas, and personnel with very little incentive to cooperate with such a cumbersome centralised approach to data management. Wight recognises the problem:

> When the system is not a valid simulation of reality the users will not keep the information correct. p105

but his solution is unconvincing, relying as it does on enforcing accountability.

> The inventory transaction form for an unplanned withdrawal should have a place to note whether this was because of a bill of material error or some other reason. p272

This issue of multiple databases and informal systems is a well known problem; I only mention it here to point out that what the planning model sees as a simple informational issue - defining how the product is made - is actually embedded in a lot of difficult issues of how to ensure that it is maintained on an on-going basis. Again this can be viewed as an incidental complication, possibly to be dealt with by coercion (as Wight sees it), or as a deeper insight that different participants in the system see the product in different ways, contrary to the assumption of unique representation. As an example, in the company mentioned above, I never managed to persuade the raw materials quality manager to transfer his data to the central database because the mainframe (an IBM AS400) could not print his quality testing requirements for the suppliers with "sexy" fonts. This is a nice example, because the MRP II package, treating the product data as pure information, has no concern for data presentation, but that was an important issue for this user of the system in his situated use of the data, because it impacted the company image.

The World as a Problem: MRP II views operations management as a problem solving task: we know what we want to make when, we know how to make it, so its a purely deductive task to schedule material acquisition and production. Operations management is thus an information processing task.

> MRP is a simulation of a manufacturing business. The shortage list is made by asking, "What are we going to make? What does it take to make it? What do we have? What do we need to get?" That is exactly the logic of MRP. p123

The **logicism** assumption of the planning model, which asserts the usefulness of formal deductive manipulations on abstract representations in capturing the dynamic nature of the real world, is clearly present.

> The functions in the closed loop MRP systems are just a logical way to run a manufacturing business. p67

This approach completely abstracts production management from the environment in which it occurs. The production management problem is entirely about the logic of conversion of materials, the "fundamental manufacturing equation". It does not depend on the particular circumstances of the company, the company's own particular problems, or ways of doing things. This attitude is neatly

captured in the following sequence of quotes which sums up Wight's quaint theory, expressed throughout the book, to explain why previous systems did not work:

> In the early years of the computer age, the users tended to think that their problems were unique and the systems people were taught that their job was to "design" a system to fit the unique "needs of the business." p57

> But, in manufacturing, the systems never did work. To go to the user and ask him what he wants will not generate useful answers p308

> The reason that many people think that systems have to be designed uniquely for each business is that the systems that they had in the past didn't work, and thus they all looked different. The systems that do work all look the same, since MRP II is a simulation of the universal manufacturing equation. p124

> If, however, the objective is to have a system that actually works, it will have to simulate the universal manufacturing equation. p58

> When the systems is challenged, however, the best possible response is, "This is MRP II, the same thing everyone else uses. Let's stop blaming the system and figure out what we need to do to manage better". p58

Of course this could be just "consultant talk", but to say this sort of thing you need to have a pretty extreme view of the relationship between logic and the situated activities of particular people, in particular companies, in particular circumstances.

Causality: Wight clearly sees the simple linear causal chain of the planning model linking manager's intentions to operative action, via a physically existent representation (the MRP II computer files) and the plan:

> MRP II makes it possible to translate overall management plans as expressed in the business plan and the production plan down to the detail level plans that everyone else will be working to execute. p87

Nature of Action

Locus of Agency: The locus of agency in the planning model is the plan. In MRP II, the source of action and the source of all order is the plan. It tells people what to do (with certain caveats), it specifies the "correct" thing to do, and it is the standard against which all behaviour is to be judged. Management is consequently the production and enforcement of plans. I have already cited examples of these attitudes in Wight's text.

As I have noted in the discussion of activity theories, if you locate activity and order in the plan, it reduces the environment to a passive backdrop for action. All the structure of the environment is already captured in its representation, and all interaction between the focal agent and the environment is extracted from the world model ("physics") by the deductive process of producing the plan. For MRP, all necessary information about the environment of work is captured in the product structure, inventory, and open order files, and the dynamics or interaction of the operation of the manufacturing system with its environment is captured in the implicit world model ("fundamental manufacturing equation") of the system programs. Therefore, it is not necessary to consider the environment further in the theory. The arrangement of the machines in the factory is not included in the theory, for instance, because it is not part of the formal relationships between end-items, component assemblies, and materials which are the essential relations that determine the planning of work. The layout of the plant is external, *incidental*, to the problem of operations management conceived of as planning. One could object, as did a critic of paper 1 (Lawrie 1996), that companies using the MRP II approach may also have other improvement initiatives in place such as Total Quality Management which might address these environmental issues. As I argued in my reply (paper three), this is not the point. The point is what is *central* to the management theory and what is *marginal*. MRP II and management-as-planning in general see

planning as central, the environment of work as marginal. Reforming the environment is seen as something like "good house-keeping" which is just "common sense", and is marginalised in the theory in just the same way the skill and interpretive abilities of the operators is marginalised. The structure of the environment is not seen to have any primary causal effect on the outcome of work, which is not already captured in the representations and, consequently, the plan.

Transparency of Sensing: Sense data is all the information coming from the environment that indicates changes of the state of the world, and is used to update the representation of the current state. The current state is defined by the stock figures on the various inventory files, and data recording the status of various promises such as replenishment orders, works orders and customer orders. The sense data are the records of transactions which record changes to inventory and the status of various orders, and also, for the higher level planning task of setting the MPS, informal information about changing market conditions. With the caveats set out in the section on objectivism above, the conversion of materials within the plant lends itself reasonably well to representation in terms of objective entities and their conversion. So the sensing problems in relation to those parts of the environment within the plant walls have more to do with problems of incomplete or noisy data within the various inventory recording systems. It is stated in all the texts that an accurate inventory system is essential for MRP II.

However, when we come to sensing customer requirements we have in addition, problems of the *translational* type discussed in relation to the planning model of activity in chapter 6. The MRP II logic requires that we must ultimately define the production goals in terms of discrete orders, for particular products, with definite quantities, with particular due dates, possibly tied to particular customers codes. This process is only translation-free if all requirements come from standard products, made purely to customer order, with due-date further in the future than the planning horizon of the system. This is a highly exceptional situation, and most companies have to do some form of demand forecasting. The inputs to this forecasting process are the previous history of sales of products, estimates of trends in the market, knowledge of competitor strategies based on market reconnaissance, knowledge of internal initiatives (such as promotions and new product launches) and their likely effects, and estimates of trends in the economy as a whole. These data, particularly those that are fuzzy hunches, are clearly not of the same *kind* as the necessary discrete orders for particular products with definite due-dates, so considerable translation is required. The knowledge needed for this translation, apart from the historical transaction data, is held by sales and marketing people, product developers, and economy watchers. It is often highly qualitative, fuzzy, intuitive, and its validity is only guaranteed by human experience and tacit knowledge.

The approach of MRP II to this problem is characteristically utopian. The people with this knowledge are supposed to meet periodically, say monthly, and translate this knowledge into a master production schedule of discrete production orders. Often forecasting packages are used to set a base-line from past history, and the marketing input is often at an aggregate level (with typically thousand of products to forecast).

> The CEO should get his {sic} top marketing, manufacturing, engineering, and financial managers together at least once a month to make sales & operations plans. In this meeting, a course should be established for the company and conflicts resolved. The sales & operations plans drive the entire MRP II system. p130

> At this meeting (chaired by the CEO), marketing, manufacturing, finance, and engineering should discuss the problems, discuss the alternatives, and often make hard unpleasant choices. p142

This process which drives the whole MRP II approach is quite unrealistic, and is bound to fail for a number of reasons:

1. From an organisational behaviour standpoint this process is unlikely to succeed. Even if it proves possible to translate the intuitive knowledge of the sales, marketing, and finance people into hard estimates in such a meeting, it is unlikely that this process could be maintained as a periodical process for any length of time. This is exactly what Bendix found. The people with the best input knowledge are

the people who are most indispensable elsewhere. In practice, they are usually accountable in ways other than through the performance of the manufacturing system (at least as measured by accuracy of the requirements schedule) so, although they may intellectually understand its importance, when the chips are down, it is not going to be number one priority. Bendix found that after the champion of the sales and operation process left the company, that is, the only one on this committee directly accountable for the accuracy of the schedule, it became increasingly difficult to organise these meetings with the most valuable people present. Like many companies, they resorted to master production scheduling based entirely on forecasts, which means on the basis of past events. This is another example of thinking about acting, and acting, being confused. It is quite feasible to get these people and their knowledge together for a brain-storming exercise on a topical problem. It is quite another matter to automate this process as a driver for on-going activity.

2. Bendix found that to have any manual input into the master production schedule, it was necessary to aggregate the forecast to a small number of product classes. After manual adjustment they were disaggregated by adjusting the line-item forecasts according to the class aggregate changes. This is the process advocated by Wight (Wight 1981, p145). As a consequence of this process, they found that the products the system recommended making, frequently did not correspond even to those that had immediate shortages, and they eventually returned to order point methods. They may not have performed this aggregation - disaggregation process optimally and there is considerable theory on this problem (for example (Brandimarte and Villa 1995, p53). Nevertheless, it is obviously an information losing process.

3. Except in make-to-order environments, the forecast of demand is subject to a great deal of uncertainty and there are large possible deviations between the demand that is planned for, and what eventuates as real orders. It is just possible that this uncertainty (or rather the cost of reducing it) outweighs any stock reduction benefit that accrues from the time-phasing of production and replenishment using the MRP explosion. This question has been investigated by a number of authors (Miller, Berry and Lai 1976; Ritzman and Krajewski 1983; Bregman 1994) but it has proved difficult to give a definitive answer.

These three sense data translations together bring into question the assumption of MRP II that it is possible to represent the world of future requirements in the crisp objective way required by the planning model of activity. Since the ability to do this is a major driver for the whole approach, we begin to see the precise, deductive, optimum seeking approach of MRP II and CAPM as a fantasy of control applied to a situation where limited predictability is really present.

Transparency of Effects: In the planning model, action essentially ends at the production of the plan, the implementation of which is assumed to be trivial. Management-as-planning inherits this view holding that plan production and enforcement are the central management activities. In practice converting the plan into production output may be far from trivial, depending in essential ways upon the interpretive abilities of operatives, their explicit and tacit knowledge, and on their ability to improvise in the face of unexpected events. Just as with plan based robots, the problem is that the execution systems, which receive clean, abstract plans as inputs, are also intimately implicated in the messiness and uncertainty of the real world by virtue of their physicality. It is not just that the world is likely to change in unpredicted ways during the planning period. It is also that the world is far more complex than the plan can ever capture. The unintended effects of actions need to be dealt with in a more immediate way than simply re-planning, which is the only option that the official systems allows. Much real management effort at the low levels is directed to this complex translational task, but this is not part of the official management model. Again it is a matter of central and marginal concepts. The central concern in management-as-planning is the order and efficiency offered by the central plan, whereas the type of situational management of day-to-day contingencies that allows this fantasy of order to be translated into real production is viewed as marginal, a mere implementation detail. In practice it may be a large part of total management activity (Webster 1991).

Wight concedes that there must be interpretation on the part of the operatives,

> MRP is not a computer decision model. It is a simulation that tells people "what will happen if." The results will depend upon how people execute these plans. p385

but their ability to do this is seen not as a potential *cause* of the orderliness of the systems but as a *repair measure* when the orderliness which is suppose to reside in, and be guaranteed by, the plan, fails to be realised in practice.

> Another concept that is vital to the "closed loop" approach is the feedback to indicate when there are significant deviations from the plan p385

In fact, despite Wight claiming that

> The greatest value of the computer and computer techniques is in helping us to use our human resources more productively. p98

the book leaves the reader with impression that managing manufacturing operations has nothing to do with the *skill* of the operatives, the management abilities of foremen and middle management who direct them, let alone to the situation in which they conduct work.

Agent Architecture

Serialism: MRP II, as conceived by Wight, implies an hierarchical and functionally divided management structure which he took for granted. The former is necessitated by the use of a centralised world model which requires centralised planning by management an hierarchical transmission of instructions to operatives. The latter is required by the serial nature of the planning approach. Marketing gathers information about customers, Stores about inventory holdings, Methods and Design Engineering take responsibility for maintaining accurate product structure data, top management sets goals, Information Systems runs the MRP II systems and schedules are distributed to foremen. Although one can now conceive now of distributed MRP systems being used in production cells, this certainly was not the original MRP II conception, which depended heavily on a single central world model which would be used by all parties in decision making:

> MRP II results in management finally having *the numbers to run the business*. One set of numbers, valid numbers, and everybody using the same set of numbers. p54

This was necessary if top management goals were to be unambiguously translated into operative actions through the plan to which they could be held accountable.

Unlimited Resources: It is tempting to attribute failures of the planning approach to the lack of sophistication that is used in the world model. The solution then is to use a finer grained model and apply more computing power to the problem, with the implicit assumption that with enough resources the problems of the approach will go away. Later developments of the CAPM approach such as Optimised Production Technology and Schedule-Based MRP have taken this direction. With MRP II, the weakest assumption is that work centre capacity is infinite. If the finiteness of work centre capacity is included the computational complexity the optimal scheduling task becomes exponential as a function of the number of work centres (Brandimarte and Villa 1995, p144). This means that it is intractable in practice for realistic problems using any real computer. However, this computing problem can be solved. For instance, Hastings *et al* (Hastings, Marshall and Willis 1982; Hastings and Packham 1985) have shown that a very simple heuristic with tractable computational complexity can produce a feasible schedule and emphasises that it is more important that the planners and the operators can understand the schedule for it to be optimal. In addition, there are powerful graphical methods available to present the schedule to users in a graphical way, and an number of quite sophisticated PC based finite capacity scheduling packages are now available. A planner can load the bills of material, routeings, and work-centre capacities onto such a system and run planned production against it to find feasible promise dates and identify bottle-neck resources. However, this is merely thinking about activity. The problem with using finite capacity planning to control activity at the work-center level is that to realistically plan future production requires shop floor feedback on the actual present load on every-work centre, and

additionally every time there is a discrepancy between actual an planned production the schedule must be regenerated. All processes are subject to variability and breakdowns. There is no point trying to use a scheduling approach to controlling work at the work-centres if the schedule information does not take into account that an upstream machine just broke down and the work is not going to arrive.

Thus the main concern is not the exponential complexity scaling of the planning generation process, but the rapid scaling of the apparatus associated with the translation between reality and the abstract model. This is the fundamental planning dilemma: if the model is to be fine-grained enough to actually direct activity then the translation problem becomes intractable, but if it is course-grained enough to be implementable it is not precise enough to solve the problem. Bendix Mintex discovered exactly this problem. As a result of this rapid complexity scaling, eventually the amount of effort involved in dealing with the world at arms length becomes greater than the task itself, and arguably this happens before any significant impact is made on the problem of creating orderly activity. BCIA tried to overcome the deficiencies in the MRP model by employing the OPT system. Not only did they experience this increased model maintenance effect, necessitating specialist positions to be created simply to run the system, but also loss of system transparency and a consequent loss of faith in its recommendations. There seems to be an almost irresistible temptation for technical people to diagnose the problems of the planning model as the need for more epicycles, made plausible by the assumption of boundless resources, rather than as the inappropriateness of the processing of abstract information to the task of creating order in the real world.

9.4 Other Representational Approaches to Management

In the discussion of representational theories of activity I noted three variations on the planning model, each of which fell short of adopting the full simulation approach of the planning model.

9.4.1 Management-as-Control

The application of periodic review and correction approaches to manufacturing inventories probably dates back to antiquity, and was quite formalised in some large US companies by the early twentieth century. Similarly, the fly-ball governor and float-ball valve are feedback control mechanisms of some antiquity. However, it was not until the advent of control theory and cybernetics that the feedback mechanism was seen as an issue separate from the machines that utilised it, in fact a theory of activity. In a similar way it was the infiltration of general systems theory and systems thinking into management theory which led to review and correction approaches like the Order Point - Order Quantity method to be seen as an example of a more widely applicable control theory of operations management.

The *management-as-control* approach is based on the cybernetic model of activity which states that purposeful activity is achieved by monitoring and closing the gap between the goal state and the current state. As with management-as-planning, managers set goals but not as discrete future states to be attained, but as an ideal state to be continually *maintained*, which can be represented by a small number of quantifiable, and observable variables. The current state is measures by the current values of the same variables, and organisational activity is managed by holding operative accountable for reduction of the deviation between the goal and current values (using rewards and sanctions). The reorder point approach to controlling manufacturing operations is an example. All assemblies and purchase materials are replenished by periodic review using order points and order quantities preferably determined by Statistical Inventory Control methods. The goal state is thus a certain average inventory level for each materials. End-items are either ordered for stock by the same methods, or assembled to customer order. Management by Objectives (MBO) (Drucker 1954) and Planning, Programming, Budgeting Systems (PPBS) (Wildavsky 1975) are examples from general management (note that "planning" here refers to setting goals).

Management-as-control shares the same general framework as management-as-planning particularly the strong causal connection it asserts between management goal setting and operational activity. However it depends entirely on feedback rather than simulation of the future (feed-forward) making it subject to the problems of all "hill-climbing" approaches, and inappropriate to achieving discrete goals. Its simplified notion of a "state" and how it can be quantified, makes the approach simpler to implement and with less ambitious sensing requirements, probably more robust than planning approaches (Johnston and Betts 1996b; Johnston and Betts 1996a; Johnston and Betts 1998). As a manufacturing operations management system, it results in inefficient use of financial resources due to excessive buffer inventories, but as mentioned earlier, there is no general agreement upon the circumstances under which this inefficiency is balanced by avoidance of the complexity of the management-as-planning approach.

Hofstede (Hofstede 1978) has given an interesting analysis of the limitation of management-as-control which uses a similar approach to the one we directed at management-as-planning in paper 2. (we become aware of this paper in the reviewing process). He works out the assumptions of the cybernetic model and uses the conditions under which they will fail to determine the applicability of "management control theory". He conclude that management-as-control will be ineffective when: "objectives are missing, unclear, or shifting"; "accomplishment is not measurable"; and "feedback information is not usable" (Hofstede 1978, p455), and gives many example of where these conditions are not met, particularly in public sector management. However his alternative was framed in terms of older organismic metaphors of organisation.

9.4.2 Management-by-Procedures

If I am correct in asserting that rule-based theories of intentional action described in chapter 6, such as production system theories of the mind, are simply a variation on the planning model of activity where the atomic actions of the plan are collected from the application of rules rather than being deduced from declarative knowledge, then this theory of activity would lead to a position on operations management, which shares most of the characteristics of management-as-planning, but differs from it by the use of procedures to determine actions from management goals. We could term this *management-by-procedures*.

I am not aware that anyone has proposed that a manufacturing production schedule could literally be prepared in this way although heuristic scheduling approaches have some of this flavour. Nevertheless, this image of management arguably informs management initiatives that try to achieve order through the explicit documentation of procedures to be applied in objectively defined situations, to which operatives are held accountable. The ISO9000 approach to quality assurance is an example. Contingency approaches to management (Hofstede 1981), which try to match management approaches to types of environmental situations, also have the flavour of this approach. All the arguments against a rule-based approach to intentionality also apply to such approaches, in addition to the above criticisms of the management-as-planning approach.

9.4.3 Hybrid Theories

Just as with robotics systems designed on the planning model (section 6.4.3), hybrid approaches to operations management have been advocated that attempt to introduce reactivity at the execution level within a planning framework. There has been something of a theme in the academic and practitioner literature in recent years that advocates Kanban as a control systems on the shop floor within an MRP based system that handles finished item assembly and purchase order initiation (see 3.3). This was the position that was advocated in the Kodak cases study. Optimised Production Technology (OPT) also implements a hybrid approach since material requirements at non-bottleneck work centres are slaved to requirements at bottle-neck resources in an essentially "pull" fashion, while the full apparatus of

"push" scheduling and optimal batch sizing is applied to requirements at bottle-necks. In OPT bottle-necks are seen as sites of permanent breakdown (to use Heideggerian language) to be dealt with by planning and in the "JIT within MRP" approach the world outside the plant is seen that way.

However, both these approaches maintain an overall commitment to the symbolic representation / planning-based approach to activity at the base level, and inherit all the problems of that approach. This is because reactivity is implemented *within* the planning framework. Rather than the planning being simply a resource of management thinking, it is still acting as a control structure, and therefore, the plan has to be continually kept up to date. In the JIT within MRP approach there is still the burden of recording the completion of works orders and purchase orders if the on-going recommendations are to be useful, which requires the traditional top-down MRP II systems to be kept in place, although with simpler shop floor scheduling output. Within the OPT framework, what constitutes a bottle-neck is determined from an elaborate global modelling process which has to be repeated as circumstances change (as Brake and Clutch Industries found). Consequently these hybrid approaches, as in the robotics case, maintain a commitment to centralised top-down symbolic world modelling. I will argue in the next chapter that another hybrid relation between planning and reactivity is possible, where planning is introduced *on top of* an essentially reactive base more akin to that used in the robot Toto (section 7.4.1).

9.5 Evaluation of Representational Approaches to Management

Representational approaches to management assert the feasibility of averting the chaos of uncoordinated activity by the automated production, implementation, and enforcement of plans. This approach necessitates the building and maintenance of an elaborate representation of the world of activity through which managers are supposed to deal with the world "at arms length", as it were. The feasibility of building and maintaining this representation at a sufficient level of detail that it may be truly causal to action, and the feasibility of continually translating between the real world and the representation, are subject to a set of assumptions whose validity is highly suspect in any complex and dynamic world. Evidence from the lack of success of MRP II, which is probably the most elaborate attempted implementation of the idea, in the most conducive problem space of material conversion, suggests that management-as-planning is not a widely useful idea. Yet the idea keeps being promoted by consultants professional organisations, and keeps being reinvented in more grandiose forms, such as MRP III (Louis 1991), and Enterprise Resource Planning.

There are grave methodological difficulties in establishing that MRP II, say, does not work. Who do you ask: the operator who has to translate its recommendations; the manager who claims the system is correct but the operators won't follow it; or the champion who got promoted on the back of it? How do you measure its contribution to profit, or better still return on investment? Most of the investment in money and effort in an MRP II implementation goes into implementing infrastructure systems, such as inventory control, which have their own benefits. Often benefits accrue simply from getting people organised. Often firms that can implement a system as complex and all encompassing as MRP II are already excellent companies, like the case companies, with an excellent chance of survival. They often have many improvement initiatives running at the same time as implementing MRP II. An implementation failure may not be interpreted as a failure of the principle, especially from the management-as-planning standpoint, which tends to see the system and its implementation in two separate categories. An implementation failure can be put down to the world model being too approximate, or to the recalcitrance of the operatives (Wacker and Hill 1977; Safizadeh and Raafat 1986). Implementation failures are seldom reported anyway.

Besides, like all consistent theories CAPM does work, in some sense. It just may work only in a domain too atypical to be useful. What we would then expect to see when it is applied outside of this domain, is the endless addition of epicycles to repair the theory. The concepts and the software would grow more

complex to cover the inadequacies of the original clean conception. I believe there is ample evidence that this has happened with MRP II. Wight makes a big point that computer systems should be transparent to, and understandable by, the user:

> The planner can see *through* the system to the simple logic being used. p101

This is scarcely true of a modern MRP II package. As the demand forecaster at Bendix Mintex put it:

> "We've (the implementation team) gone through all the steps that many times that we know it's correct but if anyone else asks us, trying to explain your way through is very complicated. We trust the system but a lot of other people don't."

I have spent hundreds of hours helping several companies implement the MRP and MPS modules of the software package referred to, and I still do not know the effect of all the system control options available. One can view "features and options", "soft-bills", "planning bills", "phantom assemblies", system control options, and so forth as added features, or as epicycles. They are software complications that are required to allow the concept to be implemented in a variety of complex manufacturing situations, contrary the claim that MRP II was the single right answer. For BCIA, the situation became even worse when they went to the more sophisticated OPT package, the operation of the system being so obscure and counter-intuitive as to require specialised OPT analysts, and gave rise to even greater mistrust by the users.

One batch repetitive manufacturing environment where most of the assumptions of the planning model are close to valid is component manufacturers in the automotive manufacturing supply chain. They typically get detailed requirements schedules from the car company customers, the products are highly standardised and the suppliers also accept MRP produced schedules. Interestingly this is also the environment where Lean Production has had a huge impact (Prajogo and Johnston 1998). This is explained by the observation that *both* the planning model and the situated action model depend on there being order in the world. If the former can utilise it so can the latter, and I will argue in the next chapter, better.

The continual reinvention of management-as-planning can be partly attributed to the widespread blindness to the implausibility of the planning model of activity, and the Cartesian tendency to conflate thinking and doing. However there is another possible explanation for its continued popularity, namely its *convenience*. It places the manager in a central position and, like the "mind" in Cartesian dualism, in a different category of "stuff" to the executors. This is convenient for the ego and political ambitions of managers. The strong and direct causal connection between managers actions and organisational outcomes that it asserts, is convenient to a management fantasy of control. At the same time the axiomatic correctness of the plan conveniently allows managers to distance themselves from poor outcomes which are "implementation failures". If operatives are accountable for putting the plan into practice, it follows that managers are only accountable for producing a valid plan, so it is possible to do a good job of planning even if the plan fails. Furthermore, it is a convenient excuse for inaction because, due to the widespread conflation of thinking and doing, having *made a plan* is widely misinterpreted as having *done something*.

If we now shift the point of view from that of the manager to that of the *customer* of the system we are likely to find a different evaluation. From the users' point of view, systems designed under the planning model of activity are likely to be viewed as inflexible, unresponsive, and down right rude. For instance, in the mechanical domain, Suchman (Suchman 1987) documents in detail the frustrations that users, going about their situated activities, experience when dealing with photocopiers designed using a planning conception of interaction. Similarly, the movement that emphasises customer focus or customer responsiveness in manufacturing (Schonberger 1986; Womack, Jones and Roos 1990), is motivated by a desire to improve users' perception of the manufacturing system along the dimensions of flexibility and responsiveness. A CAPM based system will be perceived by the user as attempting to place their unique requirements into constrained categories that reflect to designers perception of their problem rather than their own. There is a clear need for managers to trade their own comfort level for that of their customers in the design of systems.

Chapter 10

Situational Approaches to Management:

In this chapter I argue that accepting situated / interactional theories of activity as the basis for intentional systems design leads to a radically different conception of the nature of operations management. This conception stresses management as the design of systems of limited complexity interacting with suitably structured environments to achieve robust goal-attaining routines of activity, with minimum, representation, planning and simulation. The manager is the *organiser* of the conditions for intentional activity as an *emergent* phenomenon. Firstly, I sketch this conception of managing, and apply it to the Bendix Mintex case study. Then, I deduce form the situated action model what a management system thus informed would look like and compare it with Lean Production to show that the latter can be viewed, at least in its approach to intentional activity, as an example of this theory of management.

10.1 Management-as-Organising

On the view I am taking in this part of the thesis, managers as goal setters and designers of systems are outsiders to the system. Therefore, there is a potential commit to the "theoretician's error". I believe this is exactly the fallacy of the management-as-planning approach. Planning, model building and simulation are all ways of analysing and thinking about the world and future worlds (plans-as-resources). This kind of planning is a primary task of the manager. The error is to see this outsider's view of the system's operation as a basis for the operation of the system itself. This is to take the outsider's way of dealing with the environment, the managers natural *modus operandi* which is planning, to be a viable basis of the *modus operandi* of the system itself. This is what the management-as-planning approach does when it takes the position that management is the design of systems which automate the planning process as the cause of goal directed activity.

We can introduce an alternative conception of management by inverting this role for planning in management, and at the same time recognising the situatedness of activity. On this alternative view, managers analyse and plan *in order to* design operational systems, rather than design operational systems which plan. If we accept the situated action account of intentional activity, then the locus of activity is the *interaction* between the focal system and the structure of the environment, and this interaction is therefore the manager's main point of leverage. Consequently, the task of the designer / manager is to analyses the potential for constructive (that is goal-attaining) interaction between simple systems and properly structured, or restructured, environments as a way of designing operational systems. This view, which we have called *management-as-organising* in paper two (Johnston and Brennan 1996), therefore, explicitly recognises the situated nature of activity, by recognising that activity must be analysed *in relation* to the environment is which it is expected to occur.

In previous chapters, I have argued in detail that it is possible, at a certain level of abstraction, to isolate the issues of explaining and designing intentional agents from their material implementation, so we can look to disciplinary areas where the processes of designing situated agents has been made most explicit for guidance here. The work of Brooks and his co-workers, and of Agre, Chapman, and Horswill will be drawn upon here. Agre (Agre 1997, p62) has been the most explicit about this design process, and the following account paraphrases his recipe:

1. The environment is analysed for its potential to encourage interactions which will simplify the decision tasks of the focal system. If these potentials are not present then those parts of the environment that are under the designers control will need to be modified until they are.

2. The focal system is designed with the simplest decision structures that can take advantage of the structure of the environment to consistently achieve goals through this interaction.

3. The environment is changed, the focal system is built, and the interaction is examined for the expected behaviour. The environment or the decision mechanisms are "tweaked" until the desired interaction is stable and reliable. This can be done by trial and error, or by simulation: mental; physical; or computerised. Note here again, the proper role of simulation is in design, not in driving activity itself.

This method has something in common with the socio-technical systems approach of Emery and Trist (Trist 1982). They advocate an analysis of the interaction of the technical and social systems components in a joint optimisation procedure. However their units of analysis are rather different from mine. They saw the technology and social systems as sub-systems and the environment as extra-organisational. Also they were interested in psychological, sociological, and emancipatory issues rather than on-going operational activity. Checkland's Soft Systems Methodology (Checkland 1981), which stresses negotiation of multiple world-views of system user's and designer's, and their consequences, is also relevant, but once again with different units of analysis.

An important part of this view of management is that operational activity should be largely routine. This would be a platitude if it weren't for the fact that the planning model of activity has no way of giving an account of routine activity. All it has to offer is the idea that the symbol processing of deliberative activity may be compiled to "subconscious" rule following. Apart from the objection that there is no equivalent of subconsciousness in the production management system, this notion of a routine is far too deterministic: it does no allow for improvisation. On Agre's analysis of routines (Agre 1988; Agre 1997), or Bourdieu's analysis of practices (Bourdieu 1990), routines are stable patterns of interaction between agents and environments that occur because agents, with particular goals, encounter similar situations in their pursuit. Without interference routines will continue by virtue of the orderliness of the environment in which they occur, but they are not compulsory. Small adjustment to *either* the familiar situation-action mapping *or* the structure of the environment may alter a routine. As a result of such an adjustment, the routine might perform its function better, or it may evolve into a different routine for this, or another unanticipated, purpose. Or it may breakdown and create an occasion for symbolic analysis and re-design. Routines cannot be thought of as essentially rule-driven because, although the initial situated action that sets off the chain of events might be viewed that way, it does not explicitly control the resulting chain of events that constitute the routine: that is done by the physical determinism of the interaction. Additionally, being the result of, or a pattern of, situated actions, they do not involve mediation by symbol-object representation or translation.

We want operational activity to be routine precisely because routines are fluent, relatively reliable, efficient in the sense of not requiring the representational overhead, and stable platforms on which a safe form of improvisation can take place. In paper one, I struggled to express this desirable state of fluent unrepresented activity with the notion of "tacit activity" and "tacit manufacturing practices" which I defined more explicitly in paper three. Like Polany's notion of tacit knowledge as "what we know but cannot say" (that is, linguistically represent) (Polanyi 1966), it meant doing things without resource to the symbol-object representational apparatus. I drew on Bourdieu's conception of the reproduction of cultural practices, but his unit of analysis is probably too large for my purpose. Agre's notion of a routine is a more appropriate unit, and I will use it below.

10.2 Bendix Mintex Revisited: The Situatedness of Manufacturing.

Let us return to the Bendix Mintex case study and apply these principles of situational management, and the situated action perspective generally, to the events that occurred. I do this in order to show that this perspective makes sense of what happened. However, this is with the benefit of hindsight, and indeed, with the benefit of a position outside the action from which to observe. As one of the participants (the Information Systems Manager) on reading paper one commented, while he agreed with the analysis of events given therein, the management and staff at all times during the case period were doing what needed to be done from their understanding of the problems faced by the company, that is they were "situated" and "thrown" into the action I described. My account is not therefore to be taken to be in any sense a criticism of management action in this case. On the contrary, it shows how critical reflection by management on the interactions between situations and environments, especially where a number of conflicting, competing approaches were sanctioned, led to the evolution of a fine outcome.

From the start of the period studied there were two key ideas that influenced management thinking:

1. That the abilities of the operators organised into teams, to solve production quality and production problems, should be harnessed.

2. That the potential of simple, visual, reactive systems such as Kanban to simplify the detailed operational decision processes, should be harnessed.

I will argue later that these are the insights that also motivate Lean Production but initially they were not connected. The enthusiasm for Kanban systems came from the newly appointed Production Manager (Chief Executive Officer by the time of the interviews) who was familiar with Just-In-Time notions as they were perceived in the mid-1980s, and also had experienced the simplicity of flow oriented production in his previous production experiences. The team ideas came from the personnel manager and others, and consultants they employed, who were influenced by the Total Quality Control movement (as it was then called) and the Socio-technical Systems perspective of Emery and Trist (Trist 1982). There was also the (with hindsight) incompatible notion that operations should be controlled by scheduling, which came from various quarters including finance, computing, engineering and production, but I will leave that for the moment.

Both of these goals were pursued tentatively at first, followed by attempts at full scale implementation, which in each case foundered because the implementation was of a formal nature, not supported by an appropriately structured environment. There was initial success with taking *ad hoc* teams away from the work environment to brain-storm production and quality problems. This is the kind of *thinking-about-acting* activity which is easy to organise with a high profile campaign but difficult to maintain as a permanent way of *acting*. Permanent teams were then named and given portions of the product range to monitor for potential quality improvements, and to nurse them through the complex routeing between functionally organised machines. As collectives these teams were entirely formal entities, essentially names on someone's white-board. The product batch cards were coded with coloured dots to indicate which team was responsible for them. Again the association between the products and the teams was of a formal nature. Spread through the cluttered, functionally organised, authoritarian work space, with other teams and other teams' products, these teams had no spatial, social or political identity. This approach to team based production failed as documented in the case study.

Similarly, following some easy wins with implementing the reactive Kanban systems with components which were already organised into a smooth flow, a visual production control system was tried in the much more intractable main production plant. Again this initiative had a distinctly formal, symbolic flavour. Although the central Kanban board made the process more visible, it required a conscious symbol processing effort to make use of the information it contained. There was a formal relationship only, between what was represented on the board and what was supposed to happen on the shop floor, a relationship which required a difficult, conscious, symbolic translation to make use of. Apart from a

token physicality of the representation it was not different in essential ways from the computer representation in use at that time, whose abstractness it was supposed to remedy.

We can now understand with some clarity why these initiatives did not work. To make these formal representations the basis of on-going activity requires the full symbol processing framework of the planning model of activity. There needs to be a constant translation between reality and the representation, followed by complex symbol processing to determine what to do next, and finally another difficult translation back to the field of action. For all the reason so far expounded this is unlikely to work. But worse than this, the particularly human capabilities that these initiatives were trying to exploit - visual pattern recognition, ambiguity tolerance, relevance sensitivity, improvisation, creativity - were not the ones required to perform the essentially symbolic, information processing task required to make the system work. In fact, that task could be better done by a central computer. The nature of these early initiatives indicates how we routinely take activity to be essentially informational, and propose information processing solutions to it.

I can now take a privileged position, and examine how these initiatives might have benefited from the use of the management-as-organising approach suggested here. The idea is to examine what aspects of the work environment would need to be present to allow teams and reactive systems to work with the minimum of deliberation, that is, by taking maximum advantage of interaction with the environment, in order that desired practices would become routine.

Team ideas depend on cooperation. Therefore, there is a certain social requirement of the environment: the team members have to know each other and trust each other if brain-storming and cooperation is going to work. There are some physical requirements of a team environment: team members must encounter each other on a regular basis so they can get to know each other and develop some *esprit de corps*. There is no club without a clubhouse. There is a certain cultural dimension required of the environment: it has to be permissible to question existing methods of doing things. There is a political dimension required: management has to be genuinely prepared to delegate authority to team leaders if they are to make management decisions, with financial consequences, independently. There are certain perceptual or information processing requirements of a team environment: it must be possible to easily identify the products for which the team is supposed to take responsibility. This is a visual problem: the environment must be clear of visual clutter. There is also a problem of movement: if teams are to take varying roles as required, to expedite production, they must be free to move around the work area. They also need the necessary variety of skills. We would therefore expect that if the environment lacked these kinds of structure, it would be difficult to maintain stable goal-attaining routine activity envisioned by the situated action model.

For reactive control systems to operate in a routine way there are also prerequisite environmental conditions. Reactive systems won't work if the decision scenario is too complex. They will only work when they control a more or less fixed sequence of activities. This is because reaction, being a "hill-climbing" or "greedy" approach, cannot be used as a basis for complex scheduling and routeing decisions which require full problem space search, using backtracking, branch-bounding, and so forth. In other words, reactive systems require the environment to take much of the cognitive burden away from the limited decision algorithm. Small production batches are also required for reactive systems if sufficient timeliness of production is to be achieved reactively. If the reactive system is to be a visual physical control system, there are certain visual and physical aspects of the environment which will facilitate their use: the environment must be visually uncluttered, and should have a physical layout that closely reflects the routeing, and leads to a simple movement of the Kanbans.

What management at Bendix Mintex eventually realised, having modified many aspects of the work environment, such as skills, political structures, cultural aspects, was the deep implication of the physical aspects of the environment in all these requirements. The solution was to place the machines that were required for the product groups that the teams were meant to control, into separate "U" shaped cells. This move contributed to satisfying almost every environmental requirement. It compiled a great amount of the production routeing decision into the environmental layout. The routeing through

each cell was now essentially fixed with perhaps various operations being skipped if not required. This rendered a suitable decision environment for a reactive control system. It also provided a suitable visual environment for both the visual control cards and for simplified control and expediting of the production process by the teams. It meant that the teams were only physically and visually confronted by the products and processes for which they were responsible. It meant that the team members worked and socialised together. The cell was a focus for the social, cultural and political dimensions of the teams. The management team at Bendix Mintex discovered the situatedness of manufacturing operations.

One point here is worth exploring in more detail since it gives a perfect example of how environments may lessen the cognitive burden of systems. Job shops generally adopt a functional layout of machines - similar machines are clustered together. This is because they are all purpose manufacturing environments, so there is no typical sequence on which to base a product-focused layout. When the possible routeings of jobs through this environment is unlimited, the scheduling problem is extremely complex, in fact NP-hard (Brandimarte and Villa 1995, p144). However, in repetitive batch manufacturing, the possible machine routeings are very much more limited because, although there may be a great range of products, these products are largely variations of several stereotypical ones. They typically require the same general sequence of operation, for instance. A technique known as Group Technology (Burbidge 1988) can exploit this hidden regularity. The products can be grouped, by the similarity of the operation routeing required, into a small number of generic groups. Now, if it is feasible to dedicated the machines required for these operations to the production of these groups only, then a great simplification of the scheduling problem results. Each group of dedicated machines becomes a flow line, conceptually. Provided also that batch sizes are small, jobs can be released onto the first operation in sequence and largely kept in sequence through the operation, or they can be backward scheduled from due date. In any case the simplification is enormous.

So far this simplification is purely formal, a scheduling technique. But we can go further and place the machines dedicated to the groups in dedicated physical cells. This effectively *compiles* this simplification into the environment. We then only need a simple visual control system such as Kanban to initiate production and the subsequent routeing to successive machines, which is now completely obvious (to situated humans at least), requires no planning. Any minor contention problems can be sorted out in real-time because the machine loading is immediately apparent through visual inspection. This example shows how a structured environment coupled with simple reactive systems can lead to an interaction that will nearly guarantee the correct outcome. This particular example is very close in spirit to Agre's analysis of tools in the cooking task (Agre and Horswill 1992; Agre and Horswill 1997). Just like implements, the cells partition the environment so that many tasks can be carried on simultaneously with little central simulation of potential conflicts. The shop floor is "jigged" for production, but not so tightly that a range of products including unusual variants cannot be made: improvisations is still possible within routines. It is likely that this approach will work in many batch repetitive companies. Batch repetitive manufacturers often think of their process as more similar to jobbing than flow production (as did Bendix Mintex at the start of the case period) because they are more attentive to the complexity of the product range than its simplicity. Besides, from the information processing perspective, the plant layout is arbitrary, so the apparent orderliness of a functional arrangement will do.

This example is easy to follow because the interaction is between an *informational* system and a *physical* environment, both of which are easy to visualise. But all the environmental structurings I have described above, work in much the same way: the environment compiles some aspect of the problem of goal-attainment, and thus relieves the necessity for informational control in the focal system. The groups are bound to socialise by the cell environment, a degree of political autonomy avoids hierarchical communication and authorisation, the inclusion of design and methods engineers in the cell ensures the products and processes are designed to be tractable with the resources of the cell, and so forth.

The example of the cell formation is also a perfect example of the proposed new relationship between planning, representation, and action. The product structure and routeing data from the MRP II system was used by the Ingersol consultants to define the product groups, and thus the cells, that would effectively *eliminate* the day-to-day planning problem. The same kind of explicit, symbolic reasoning about products and processes that was previously supposed to form the *basis* of the on-going activity, was used instead to design the system and compile intended action into the interaction between the system and its structured environment.

In hindsight, many aspects of this analysis of environmental requirements are obvious. However, from the dominant information view of activity and management, they take on the status of *mere* practical wisdom, *details* of implementation. If the situatedness of activity were explicitly and routinely recognised in manager's thinking about activity, it would be possible to anticipate these connections. It would be possible to sit down and analyse the simplest systems that would implement the team and reactivity ideas, and specify by analysis, the precise environmental structure and interaction that would be necessary to make the desired performance routine. Various methods of formal analysis would be provided by management science, rather than managers relying on hard-earned practical experience. The interaction between system and environment could be simulated with various degrees of formality, and the design "tweaked". In fact, the Bendix Mintex solution actually *evolved* (with the principled intervention of management). This also is a recommendation for the situational approach. Just as Hendriks-Jansen (Hendriks-Jansen 1996) sees it as the mechanism of evolution of human and animal intelligence, and Brooks recommends it for the development of robots, so it is the building "bottom up" upon stable patterns of interaction that enables the incremental improvement of systems without necessitating continual total redesign. This has considerable implementation advantages for management systems compared to "slash and burn" methods like Business Process Re-engineering and MRP II.

Finally, the concurrent MRP II initiative foundered at least partly because processes such as the Sales & Operations review meetings could not be made routine. In other environments it is the maintenance of valid product structures or inventory records, that can not be made routine. Of course, the information processing framework of MRP II does not even see this as a potential problem. It is possible that certain environmental structures could even make these abstract processes automatic, a kind of second order routine of explicit representation. It would require a culture of discipline, attention to detail, and comfort in dealing with the world on an abstract level. Zuboff (Zuboff 1988, p90) for instance noted that, unlike the traditional skilled workers of the manual process, a later generation of computer literate control operators at her paper mill case company, were able to treat interaction with the computerised control system as though it were a sentient experience. For them the computer representation became ready-at-hand. The difference, of course, is that the paper mill was totally automated, whereas in traditional manufacturing it seems more desirable, flexible, and economical, to make use of the situated abilities of the work force at large.

10.3 Lean Production as Situational Management

I now argue that the Lean Production manufacturing paradigm is an example of the management-as-organising approach which is informed by the situated action theory of activity. The key goals of Lean Production are:

1. Just-In-Time production and supply chain management through small lot-sizes and simple reactive (pull) systems.

2. Quality at the source. Defects should be detected and corrected at the point of production.

3. Teams. The hierarchical and functionally departmentalised organisation should be replaced by functionally complete teams to increase communication between functions and harness human potential.

These are a more or less literal translation of the goals of the Toyota Production Systems (Monden 1983) which can be taken as an instance the of Lean Production paradigm, with perhaps a slight bias to flow line production. All the specific techniques of the Toyota Production System are directed to these goals. We can be reasonably sure that its designers were not influenced by European phenomenology, although Hieko (Hieko 1989b) finds some resonances between Lean Production and Japanese cultural notions. It is more likely that it evolved as a solution to practical problems, much as did the system at Bendix. According to one of the originators of the Toyota Production System:

> "The technique of Kanban as a means of Just-in-time production, the idea and method of production smoothing, and Autonomation (Jidoka), etc., have all been created from such trial-and-error processes in the manufacturing sites" (from Taiichi Ohno's foreword to (Monden 1983, pi)).

The guiding principle was the "notion of eliminating all kinds of unnecessary functions in the factory" and most authors who have written on Lean Production as a paradigm have used waste elimination and simplification as the organising principle (Schonberger 1982; Monden 1983; Womack, Jones and Roos 1990; Cowton and Vail 1994). As I will detail below, these sources of waste are also sources of breakdown of stable routine activity.

Cowton and Vail have pointed out (Cowton and Vail 1994) that the notion of Just-In-Time production is a vision of an "ideal state" of activity, rather than a "system" or "philosophy". I equate this state to the transparent coping of a situated agent in a routine interaction with its environment. The metaphor is of a skilled worker engaged in fluent, efficient action largely unmediated by explicit knowledge, representation or reasoning. I will argue that the way this flavour of intentional activity is implemented in Lean Production is precisely in accord with the situated action model. The simplest way to do this is to first define what a system of manufacturing operations based on the situated action model should look like, and then to show that Lean Production conforms to this image. From the discussion of the situated action model in chapter 7, it would be expected that management based on this notion of activity would have the following features:

1. **Planning:** While strategic and tactical planning systems would inevitably be used by managers in the firm, the operations systems would be completely decoupled from these. Such planning activity would be used to shape the environment, encountered by the focal system, rather that implementing a causal link between planning and operational activity. When planning and model building is present at the operational level, it would be built over the normal routine situated activity, as necessitated by breakdowns in that routine activity.

2. **Representation:** It would de-emphasise the role of centralised symbol-object representations and worlds models as mediators of operational activity. Representations when they are present would be fragmentary and indexical-functional.

3. **Action selection:** Actions would be driven by response to features immediately present in the agent's current situation. In manufacturing systems this corresponds to the use of "pull" systems.

4. **Architecture:** In order to implement a more direct coupling between sensing and acting, we would expect to see a modular rather than functional decomposition of the operations system.

5. **Role of operatives:** Recognising explicitly the translation, interpretation, and relevance selection that is involved in sensing and acting, we would expect the system to be organised to make the maximum use of the ability of humans to excel in these areas.

6. **Communication and Control:** Without the need to update and make use of a centralised representation, we would expect hierarchical communication and authorisation to be minimal. In other words we would expect the modules to be semi-autonomous and fully functional, interacting through conflict resolution or focus mechanisms.

The following sections will demonstrate that these characteristics are present in the Lean Production paradigm, drawing on the Toyota Production System as a well documented prototype (Sugimori et al.

1977; Monden 1983; Shingo 1989), other World Class Manufacturing (Schonberger 1982; Schonberger 1986) and Lean Production (Womack, Jones and Roos 1990; Jenner 1998) manifestos, my own observation of Lean Production at Toyota Australia (Johnston and Lee 1997), and the Bendix case study. Many of the Lean Production ideas such as production levelling, teams, cells and group technology are present also in the writings of Burbidge (Burbidge 1988; Burbidge 1990), the origins of which predate Lean Production.

10.3.1 Planning

The way in which operational activity is decoupled from high level planning is well illustrated by the material replenishment systems used by automotive manufacturers using Lean Production compared to the traditional US "push" approach (Monden 1983; Womack, Jones and Roos 1990; Johnston and Lee 1997). In the traditional approach, car requirements are ascertained from dealer orders and forecasts, and a detailed car assembly schedule is produced. An MRP explosion is used to convert this into time phased parts requirements schedules which are sent to each part supplier, usually monthly. The first one or two months will be firm orders, which the car company is committed to receive. The remaining months data are informational. These schedules specify product, quantity and dispatch date, at daily or weekly intervals. When these dispatch dates fall due, the supplier is supposed to deliver the parts, which in theory will be just about to be consumed. Due to inevitable fluctuations in supply and demand, the parts are usually placed in a buffer stock store for somewhat delayed consumption. This approach is driven by a signal which was generated in anticipation of future demand, and which is activated on the supply side of the exchange. These two characteristics define it as a "push" system in operations management parlance. Its operation depends on direct coupling between the final assembly plan and replenishment of parts on the operational level.

In companies that use the Toyota Production System the relationship of replenishment to production plans is different. As before, a production schedule is prepared from dealer and forecast information. The production plan is again used to sequence final assembly. Additionally, a process called "production levelling" is used to create the most uniform sequence of daily assembly by uniform mixing of car models (quite contrary to traditional "economy of scale" batching). This makes the consumption of parts on the shop floor very uniform over time, but necessitates the quick model changeovers for which Toyota is famous (Shingo 1989). This production schedule may be exploded by an MRP-like process and time phased parts requirements sent to the suppliers monthly. However these are purely informational, and are used by suppliers for their own high-level production planning. At the car company, parts are only ever stored at the work station where they are used, and only in small quantities (say, enough for one shift). They are held in appropriately sized uniform containers, to each of which is attached a removable Kanban card of some durable material, which "represents" this temporary quantum of parts. The number of Kanbans that are allowed to exist in circulation is fixed by management policy. When a part container is emptied at the production work station, its Kanban is freed and is taken by material handling staff to a highly visible series of hooks or pigeon holes, organised by supplier name, in the goods receiving area. These free Kanbans now become both a signal and an authorisation for the supplier to deliver the associated quantity of parts. The suppliers deliver parts frequently and regularly (often several times per day, and within fixed time windows). They pick up free Kanbans (and empty containers) on one trip and return the parts with Kanbans on the next. The parts are taken directly to the production work centre where they are to be used. Thus, there is no causal connection between the forecast requirements derived from the production plan and the actual delivery of parts. The replenishment of parts occurs through the operation of a routine, one which automatically adjusts to changes in the production volume and mix. Note that the Kanban signal is both a decision mechanism and a stock control mechanism. This example describes the replenishment of parts from outside, but the same method is also used to control material movements to downstream "customer" production cells.

This example illustrates a number of the themes of situational management. Plans are used in the two examples in different ways. In Lean Production, the plan is used to shape the environment in a number of ways. The planning involved in production levelling only serves to protect the focal system from fluctuations in demand: in the "push" method it drives replenishment directly, and any reactivity is within this planning framework. Other management actions in this example also can be viewed as shaping the environment of the system. Supplier compliance is important for operation of the routine and may be involuntary, but it may also be shaped by establishing, closer business ties and sharing information (which is derived from planning). Management control of the total number of Kanbans is also a tuning parameter. It must not be so low that shortages regularly occur, and not so high that accumulation of goods at the work centre, or loss of tight coupling with the supplier, allows the routine to destabilise.

The relationship of planning to breakdown is also illustrated in the Toyota Production System. When a production problem occurs which is likely to pass a defective assembly down stream, a production worker is empowered to stop the production line. This activates a visual signal called an "Andon", a clearly visible coloured light, which summons all cell workers with appropriate skills to fix the problem as quickly as possible. Team problem solving skill are brought to bare on the problem. We might say, in Heideggerian vein, that the production facility, the product, and its manufacturing process, become present-at-hand in the breakdown and become the objects of explicit representations and reasoning, whereas, in smooth operation they are ready-at-hand as implements of transparent coping. This process may not be just one of repair, but also of evolution: solutions may result in permanent changes to the production facility, product, or process, and as a result a new slightly altered interaction between the system and the environment is established. This technique is explicitly employed when managers use the trick of reducing the number of Kanbans in a controlled way to expose inefficiencies in the replenishment systems. The important point is that this kind of planning is built over the ability for the system to run in a routine way without planning, and provided that routine is well established, should create less of a fire fighting-situation than if operators are constantly translating largely invalid control plans through their problem solving abilities into workable actions.

In implementations of cellular manufacturing in complex batch repetitive manufacturing enterprises, it may well be that certain products never entirely fit within the capabilities of any one cell. There is no reason why these should not be handled through an explicitly represented and planned routeing between cells, provided they are not too great in number. As this becomes more burdensome, with careful management, the machine composition may well be "tweaked" and evolved to more nearly accommodate these products in the core system. Again the feasibility of this incremental approach to change depends on the existence of a stable routine to manage the residual complexity. This is the basis of the Japanese concept of "Kaizan", gradual improvement. In a different context, Hutchins (Hutchins 1995) has shown in detail how routine situated activities of individual members of a team can be re-jigged to solve problems of considerable complexity in an emergency.

10.3.2 Representation

The plan-based CAPM approach requires that customer requirements, planned production, supplier commitments and inventory stocks be represented in a discrete symbol-object manner, *in order that* a production and purchasing schedule can be deduced from them. In Lean Production customer requirements are represented in this way to some extend to facilitate long term planning and model mix smoothing. There will most likely also be other representations of this kind, particularly for accounting purposes, but these representations are not implicated in a symbol processing deduction of what to do next. Operational activity is entirely decoupled from planning, inventory, and financial systems in complete contrast to what Oliver Wight (Wight 1981, p51) proclaimed as the definitive step in completing the MRP II vision. The conflicting representational requirements for operational and financial management under Lean Production, were commented on a number of interviewees in the case studies.

In the Lean Production operational system, representation only arises through the need to convey certain information between the environment and the production system to control their interaction. This role is played by the Kanban cards. However, Kanban cards do not represent anything as discrete or objective as the requirements of a certain customer on a certain date, just the immediate downstream shortages of the particular generic product. Nor do they represent particular works orders with particular due dates even though they control production. Nor do they directly represent the amount of stock of a particular product in a particular place even though they control inventory levels. Each free Kanban says: "There is a downstream (customer) shortage which is not currently covered by an available part". Its referents, as a representation, are neither entirely in the environment nor the focal system: they are in both. They are an *indexical* representation of the necessity for a certain type of interaction between the environment and the focal system. They are also *functional* in the sense that they are a purpose-specific representation, unlike records in the product master file of an MRP II computer system, which provide an all encompassing description of products for all systems, for all purposes.

Like all indexical-functional representations they are efficient (Agre 1997, p230), because they do not have to name every object in the world, and to the extent that they refer indirectly to individual downstream requirements, this pointing is rapidly transferable to the next individual downstream requirement as the agent does something about the first. The limitation of the number of Kanbans is effectively a limitation of the number of features of this interaction that can be simultaneously dealt with by the system, just like Pengi's visual markers, and serves to limit the potential "cognitive" (combinatorial) complexity of the system. Although it can probably be looked at in other ways, Agre's account of the efficiency of indexical-functional representation accounts for the efficiency of this type of physical token control systems.

This lack of object-orientation can be seen in the way Lean Production views the conversion of materials in the process of manufacture. In conventional component batch manufacture for which MRP II was primarily designed, manufacturing is seen as the repeated production of objects which in turn become inputs to other production episodes at later times. These objects are completely tangible: they have to be put somewhere, and they have to be accounted for until they are used at some possibly quite remote time. Transactions need to be recorded to maintain a representation of their quantity. Thus, they must be represented objectively as something that exist for a time independently from the interaction which is the active conversion of materials. The more of these sub-assemblies that exist for economy of scale or control reasons, the more complicated the planning task is.

Lean production (by virtue of its small batch sizes) views manufacturing as a continuous process of conversion of supplier sourced material to fill customer requirements. Ideally none of these materials ever stops in one place or steps outside this interaction long enough to require being treated as an independently existent stocked component requiring objective representation and accounting. This is, of course, from the operation management perspective: finance will require accounting of flow of value in and out of the plant, but a great deal of accounting control can be relaxed in side the plant due to the small value of work-in-progress in Lean Production firms. Viewed from the Lean Production perspective then, the component stocks of MRP II are a symptom of a *breakdown*, due to the inability of the systems to maintain a continuous conversion of material. They are set aside from the production activity and pop into existence as *something*, some thing which needs to be represented and accounted for. They become present-at-hand. While the very same configuration of material comes into existence momentarily during lean production, its relation to the focal system is ready-at-hand. As Winograd remarks (Winograd and Flores 1986, p69):

> It is often remarked that the Eskimos have a large number of distinctions for forms of snow. This is not just because they see a lot of snow (we see many things we don't bother talking about), but precisely because there are recurrent in activities with spaces of potential breakdown for which the distinctions are relevant.

Figures 10.1 to 10.6 illustrate this point using data from the Bendix Mintex case. They trace the way the bills of material were structured, that is, the way the production of the end product was thought to involve the separate production of various intermediary objects (all of which were actually stocked, existed for a considerable duration, and were accounted for) through various phase of the changes that occurred in the case study period. Note that the actual *method* of manufacture hardly changed during that time.

In 1986 the company was practicing classic component manufacturing where at various stages from raw materials to finished products things where put in and out of inventory stores and converted into each other (figure 10.1). These inventories where controlled by reorder point methods. This approach involves a great deal of unnecessary material handling; for instance, the shrink wrapped sets are shelved in the warehouse and then picked later for boxing. The need for all these inventoried items was related to several seemingly intractable problems, or breakdowns:

1. There was a long standing belief that large stocks of backing plates were required due to unreliability of the stamping tools. There were 1.8 million backing plates in stock at this time, being about 8 weeks stock.

2. Large stocks of unmatched disc pad pieces (a set often requires 4 different pieces) existed in the "kit assembly area" because of the near impossibility of making all pieces of a set at once given the batch sizes that were then thought to be normal (3000+) and the systems in place. Note that neither of these items has any saleable value, yet they were viewed as inventory rather than work-in-process.

3. There were large stocks of shrink wrapped sets in the finished goods warehouse because it was not viewed as possible to commit them to final customer form until dispatch.

4. There were large stocks of purchased items due to the loose arrangements with suppliers typical of that era.

All this was normal thinking for that period. After an new production manager came in bringing JIT ideas, Kanban was used to control the friction mix production (figure 10.2). This was an easy win because informal pull systems already existed. As part of the first round of plant layout changes around 1988, manufacturing accepted for the first time that the inventory of loose disc pad pieces was work-in-progress and their number was an adverse measure of performance (figure 10.3). This was a significant event because previously manufacturing had viewed disc pad pieces as the object being manufactured (eg. production was measured in pieces) despite the fact that they could not be sold if there where no matching parts to make a set. Similarly, loose backing plates were accepted as work-in-progress, and early attempts were made to manufacture them lot-for-lot with disc pad production. This was a major culture change forced by new management. It was found that there was some truth in the tool breakage "mythology", but mainly in limited ranges of complex pressings. Backing plate stocks were reduced to about 1 weeks supply.

In 1992, when the cellular arrangement overcame the long standing problem of introducing Kanban effectively in the main disc pad production area, the significance of individual disc pad pieces as objects was further eroded, with the cells producing shrink wrapped sets straight to the warehouse (figure 10.4). In years to come it will probably seem ridiculous that the individual pieces were ever thought of as the thing being manufactured. In addition many suppliers were being brought into schedule driven JIT supply agreements. In 1994 the opinion was that backing plate production may well be a permanent site of breakdown and the plan was to move backing plate manufacture off site and treat it as an autonomous JIT supplier.
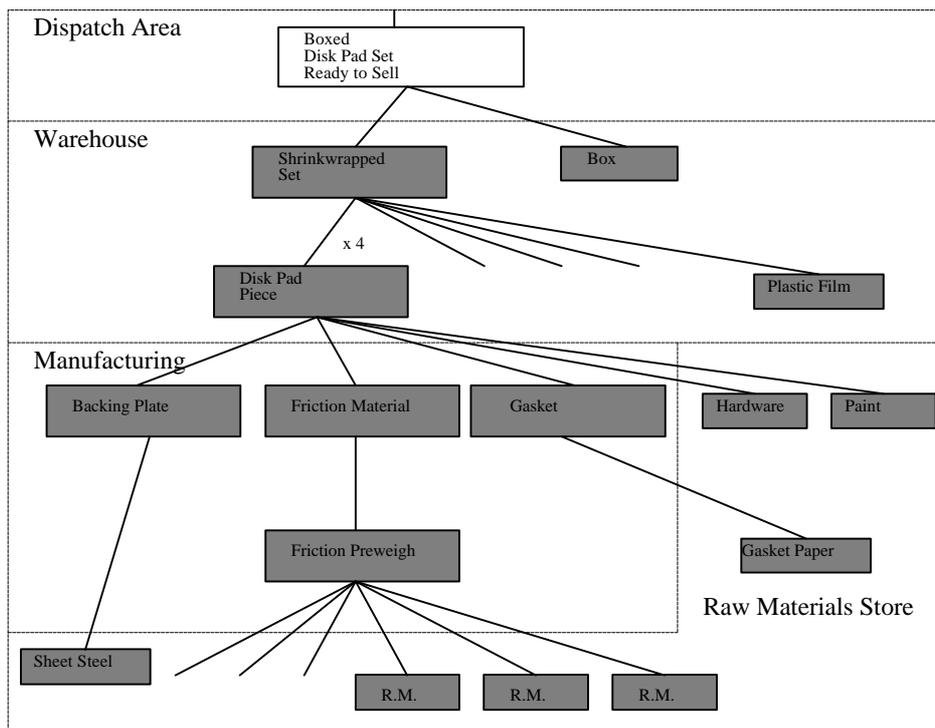
Figure 10.1: Disk Pad Product Structure Around 1986. Shading shows components that had sizable stock. Dotted lines indicate the store area in which they reside. In 1986 essentially all components and subassemblies where held in stock. The only exception was the boxed set which was packed by dispatch to customer needs. Manufacturing was seen as the successive conversion of inventoried items into each other.
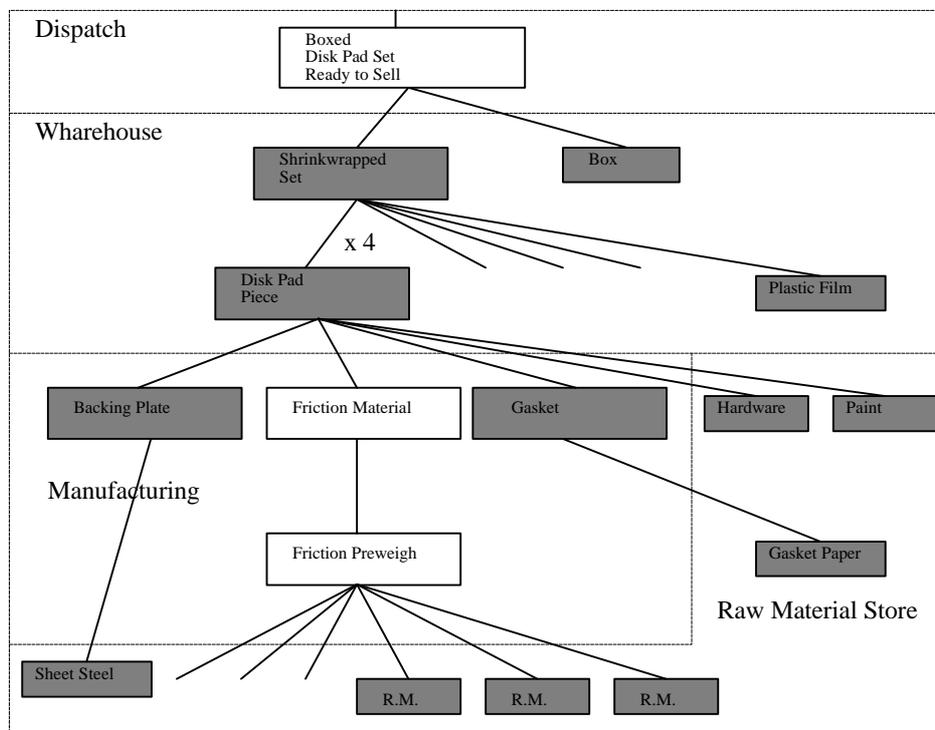

Figure 10.2: Disk Pad Product Structure after the First Round of Kanban Initiatives (1988). In the first round of Kanban implementations pre-weighing, mixing and delivery to the shop floor of friction material mix was controlled by Kanban cards, and computerised inventory figures for these items were no longer used for ordering.
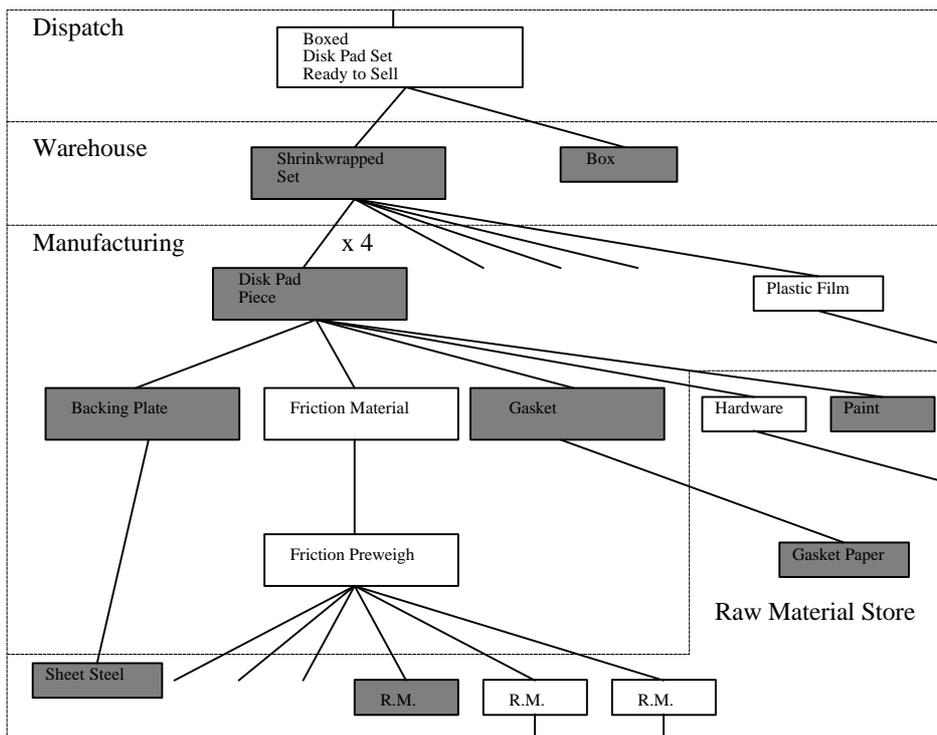
Figure 10.3: Disk Pad Product Structure after Establishing JIT Arrangements with Some Suppliers. Over a long period JIT supply arrangements were negotiated with suppliers of some locally sourced materials. Some other materials with long lead-times or small value were still under order point control. A major achievement of this period was to get manufacturing to accept miss-matched disc pad pieces as work in progress and the use of MRP to coordinate their ordering.
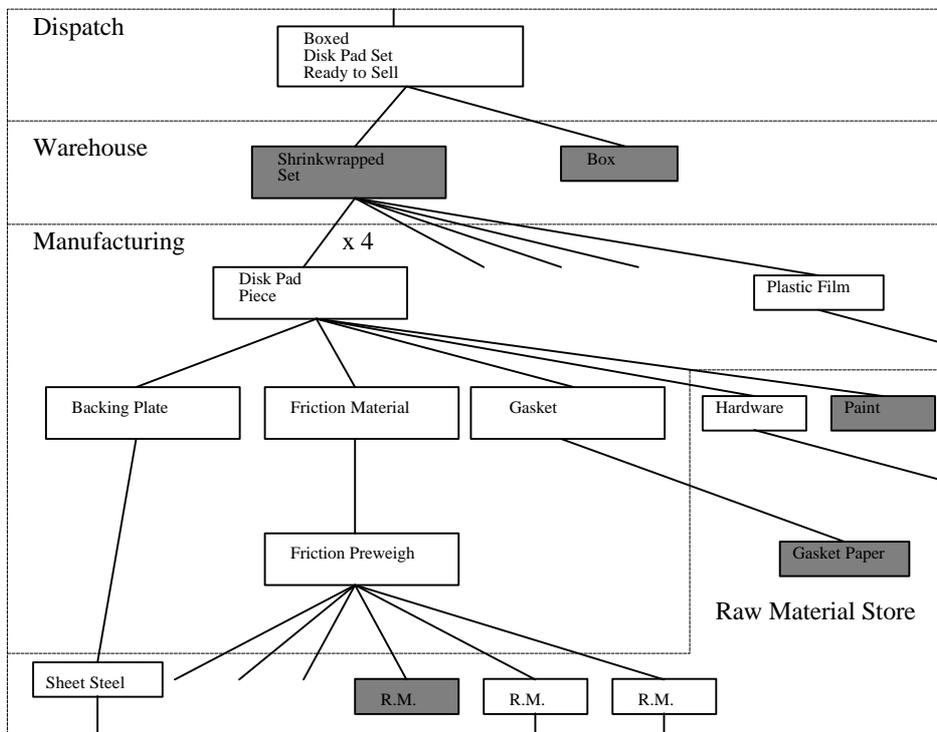

Figure 10.4: Disk Pad Product Structure after the Introduction of the First Cell. After a number of unsuccessful attempts to control backing plate production and disk pad production by Kanban, by 1992 the simplest disk pad sets were being made in the first cell by JIT methods. Shrink wrapped sets were still being stocked in the finished goods warehouse for later packing in customer boxes, but backing plates, gaskets and disk pads were being made in lots of 300 being one shift production.
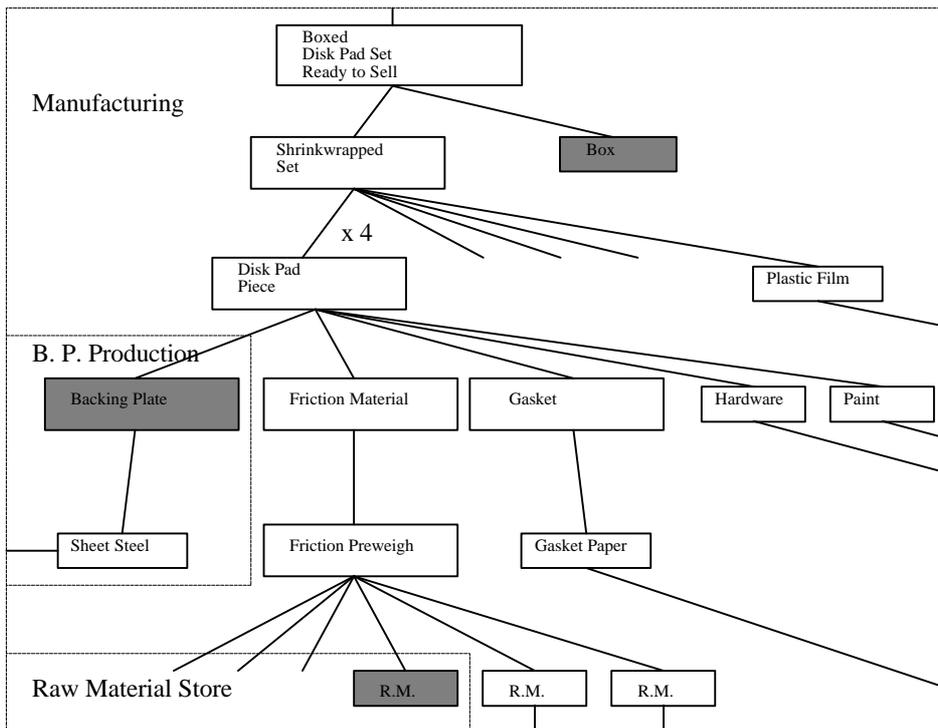
Figure 10.5: Image of the Ideal Solution (1994). Difficulties with small volume backing plate production continued and when the cells for sets with different parts came on line it became apparent that some stocks of backing plates would be needed to avoid stocking odd disk pad pieces in the cells. Backing plate production would be set up as an autonomous JIT supplier in another location. The idea that individual customer orders could be made and boxed in the cell with 24 hour turn around, effectively eliminating storage of shrink wrapped sets, was being expressed for the first time.
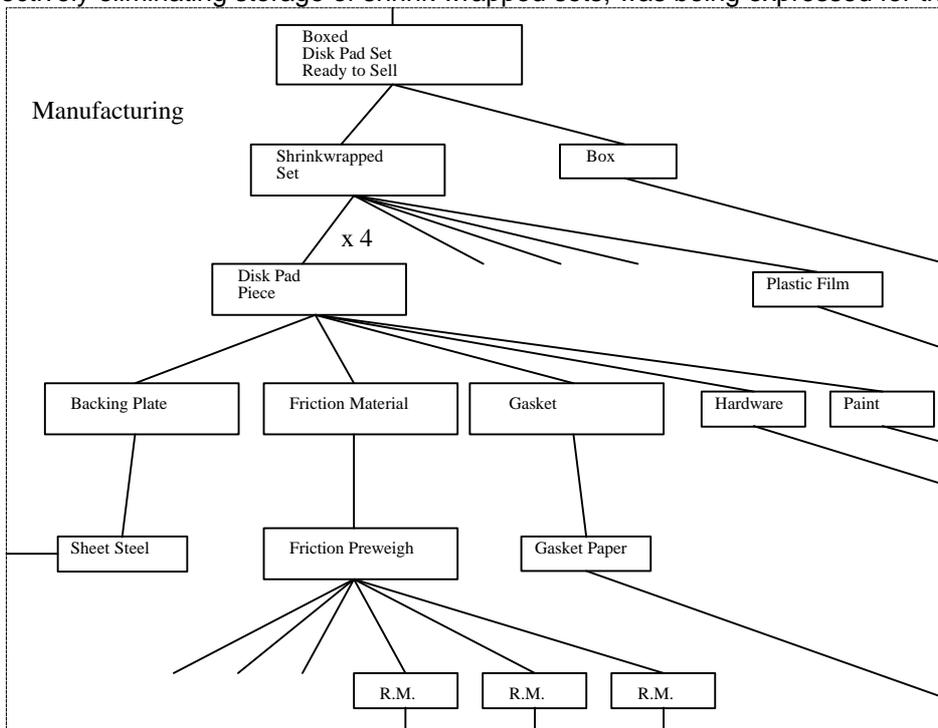

Figure 10.6: The Lean Production Ideal. In the Lean Production ideal, all barriers to JIT production are eliminated. set-up and tooling problems by process development, supplier lead times by encouraging suppliers to move closer or carry stocks. Manufacturing is now seen as provision of a service to customers by adding value to materials from suppliers rather than as making objects for sale.

On one visit to Bendix Mintex the Master Scheduler spoke of the vision of disc pad manufacture shown in figure 10.5. A customer order is received, the disc pads are manufactured in a single shift and are immediately boxed in the appropriate customer carton and sent to dispatch for delivery the next day. Now even the shrink wrapped set would loose its status as an object. Although this ideal had not been fully implemented by 1994, the important point for this discussion is that *this idea was absolutely un-sayable in 1986.* In the lean production ideal (figure 10.6) all such breakdowns that would prevent the transparent coping embodied in the JIT notion are to be resolved, and manufacturing is viewed as the adding of value to materials to customers requirements in an unbroken chain of suppliers and customers. In this, possibly unrealisable ideal, manufacturing ceases to be making objects and becomes the provision of a service to customer through adding value to materials.

10.3.3 Architecture: Autonomous Teams in Cells.

Batch repetitive manufacturing facilities using Lean production are generally organised into cells to simplify routeings and transform the processing into a flow line within the cell (Schonberger 1982; Womack, Jones and Roos 1990). The cells are operated by teams which should ideally include design, methods, maintenance and quality personnel, so that these functions which were formally functionally separated are now dispersed though the plant in close visual and practical contact with the processes they control. The teams are supposed to be self-managing and semi-autonomous, to reduce the need for hierarchical communication. In Burbidge's system (Burbidge 1990), production levelling, group technology, and cells was specifically motivated by the desire to make the production scheduling problem simple enough that it could be done by the cell foreman without a computer system. Monden (Monden 1983) sees the whole structure of the Toyota Production Systems as built in a bottom-up way upon the team concept (figure 10.7). This diagram also nicely depicts how management goals of increased productivity and profit *emerge* from reforms at the grass-route level, rather than the reverse relation depicted by Wight (figure 9.1).

How the Lean Production components, viewed as environmental choices, support the human problems solving potential of teams has been described in connection with Bendix Mintex above and to avoid repetition they are just summarised in table 10.1 below. Here I merely note that the approach adopts a modular architecture in order to enable a direct connection between sensing production requirements and acting on them, both through the Kanban systems that the cellular design enables, and by making use of the human cognitive abilities instead of symbolic information processing.

10.3.4 Action Selection

Within the plant, the activities within the cells are determined by Kanbans that have been brought to the cell from downstream cells to indicate that they require replenishments of certain assemblies. They are collected on a Kanban board at the start of the "U" shaped production line. There will be Kanbans for a number of different products that can be made by the cell. The Kanban board often makes use of visual methods to indicate the priority of the Kanbans, and there are usually "Kanban rules" printed on the board to indicate how to resolve any ambiguities. I have seen companies in the automotive supply chain who use the Kanban board at the final assembly cell to perform production smoothing manually. In many companies, if the demand for production in a cell (indicated by number by the Kanbans on the board), becomes low the workforce is meant to relocate to some other cell where they are needed. This is only possible if the operatives are multi-skilled, which is therefore part of the Lean production approach. The Kanbans and the JIT philosophy mean that they cannot overcome a demand shortage by making unwanted products and stocking piling them. I have seen the ultimate of this system in one company to which I consulted that made a relatively small range of speedometer cables. Production cells were set up one for each product as long parallel benches. At the head of each was a Kanban board holding the Kanbans for that product which were received from their car manufacturer customer. A small number of teams moved from one cell to another as a critical number of Kanbans appeared at

the head of each product cell. The routines that the operatives develop for each product within a given cell, or between different cells, are analogous to the "behaviours" of Brook's creatures. Action selection is the swapping between these behaviours based in cues relating to immediate situation of shortages.
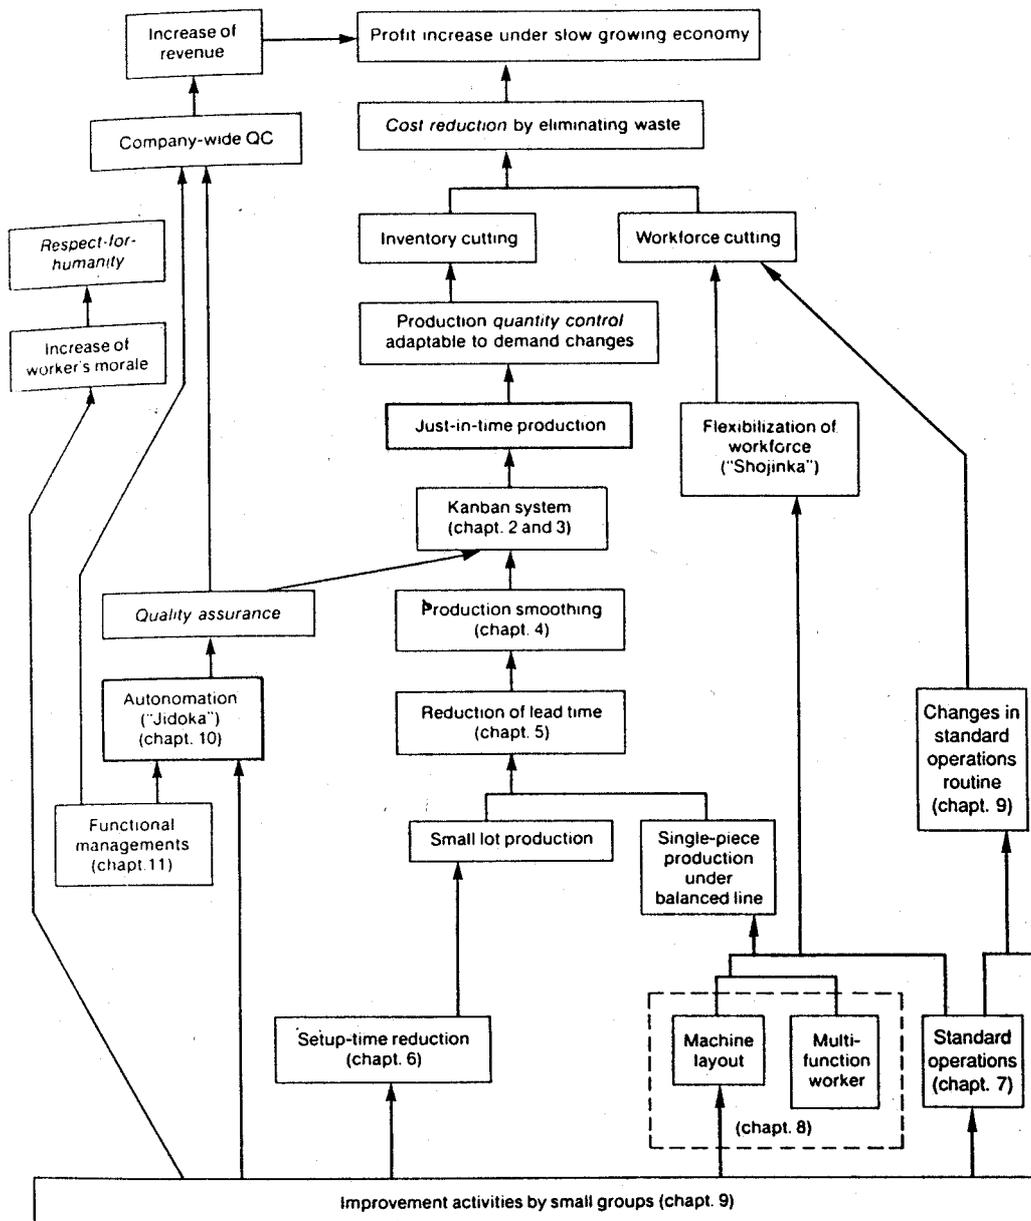

Figure 10.7: The Toyota Production System According to Monden (Monden 1983, p3)

10.3.5 Improvisation

One of the apparent attractions of the planning model is its general purpose reasoning facility that seems to hold out the hope of dealing with unusual circumstances as easily as familiar ones. Yet as I demonstrated in chapter 9, the representational approach of MRP II becomes extremely burdensome when customer variations, re-works, conversions and one-off products occur. MRP II therefore, fails to give a convincing account of how a manufacturing system can improvise on the standard manufacturing theme. However, within Lean production, provided that these manufacturing episodes are not too frequent or different from the standard ones for which the cells are designed, they can be easily accommodated *on top of* the established routine. This is because all processes are controlled by self-

adjusting reactive systems that are not sensitive to the representational differences in the variants. Raw material replenishments of a standard item will adjust to its temporary replacement by a substitute, production will adjust to re-work being sent back to a previous operation. Customer variations are easily handled at the output end of cells because the end-items are not stocked (for example, Bendix Mintex boxing sets at the end of the production cell which eliminates the need to stock generic shrink-wrapped, un-boxed sets). And again, the possibility exist for gradual evolution of the cell design if certain variants come to dominate. This alternative shows clearly that the complexifying elaborations on the standard bill-of-materials, that these variants entail for MRP II, are *problems* not *features* of that approach.

10.4 Evaluation of Managing-as-Organising

I have argued in this chapter that adopting the situated action view of activity leads to a new conception of the management of on-going, goal-directed activity of people, and a different conception of the place and role of the manager. The manager is no longer a plan producer and enforcer, but a designer and organiser. Appropriate operative action is no longer connected to management action in a direct causal chain, but is emergent from the interactions that are shaped by management organising.

This will be a difficult idea to sell. The conveniences of the planning approach set out in the last chapter are no longer present. Management becomes real work, involves actually dealing with the individual situation and uniqueness of the firm, its employees and its environment. This approach is unlikely to be catered for by turn-key solutions offered by consultants.

The potential advantages of this approach include:

1. Routines of production using simple largely manual systems stabilised by structured environments, are potentially very efficient in a total system sense, even though isolated component may not be optimised (for examples, batch sizes may be much smaller that recommended by traditional optimisation theory (Johnston 1996a; Johnston 1997a; Betts and Johnston 1998) - at Bendix Mintex both disc pad and backing plate batch quantities were reduced without an accompanying set-up cost reduction program, contrary to traditional wisdom). The efficiency offered by component oriented MRP II approaches may well be illusory when the enormous over-head of the symbol processing approach is taken into account.

2. The leanness of the approach allows for greater flexibility, and hence responsiveness to customer requirements, which is increasingly being recognised as an important whole system performance measure (Johnston 1997b; Johnston 1997a). When systems performance measures other that traditional micro-economic efficiency are used, reacting is found to be a more attractive strategy than the forward commitment of planning approaches, when facing complex, unpredictable customer environments (Johnston and Betts 1996a; Johnston and Betts 1998).

3. Management-as-organising shows the way that the potential of humans individuals and teams can be incorporated in a theoretical and practical account of operations management. The symbolic, representational approach gives no account of the inherent and vital situatedness of human activity, the utility of stable routines, the ability of humans (on the shop floor as well as in the board room) to improvise, and to see opportunities for new practices. From a theoretical point of view, the approach points the way to a satisfactory resolution of the old problem of reconciling individual, team, and system intentionality, since the approach is compatible with a modular multi-agent design for a system.

4. The approach gives a satisfactory account of how an operations system might evolve and learn as a result of the existence of robust routines as protective fall-backs. It gives a more satisfactory account of how systems can be designed and implemented incrementally. Neither of these issues is easily addressed in the symbol processing framework, because the source of declarative knowledge is external to the intentional agent. It also shows how system knowledge can become compiled into the physical and social environment of the operations system - the system can have *tacit* knowledge. This fact is ignored in accounts of radical change that focus on processes from a purely informational point of view, such

as Business Process Re-engineering (Hammer 1990). Not only is what works compiled into the environment, but what does not work (which has been tried and rejected) is also implicitly present.

Aspect of focal system	Structuring of Environment	Aspect of Environment
Reactive control system	Visual Kanban system implements reactive control through a physical system	Physical / visual
	Production smoothing protects reactive system demand fluctuation	Customer requirements
	Group Technology and cellular layout simplifies scheduling by creating many small flow lines with relatively fixed routeings	Physical / Problem structure
	Information sharing with vendors and other supply chain cooperative measures protect Kanban from fluctuations in supply	Vendor reliability
	Product redesign makes group technology and fast set-up possible	Product design
	Set-up and lead-time reduction, and other process changes allow small lot production and thus reactive production initiation	Process design
Small group production management	Cellular layout creates production environment which models production steps	Physical
	"U" shaped cell layout provides visibility of process allowing a direct link between sensing and acting	Visual / Cognitive
Small group improvement Activities	Cells provide social contact for teams	Social
	Total Quality Management initiatives provide a culture of problem solving	Culture
	Flattening of management structures provides political legitimacy to teams	Political
Worker initiated quality assurance	Authorised line stop prevents defects from being passed on	Political
	Andon lights provide visual indication of production problems	Visual
	Cell and team structure provides environment for quick response to problems	Physical
	Political changes empower workers to prevent defects	Political
Flexible deployment of labour	Multi-skilling of workers allows mobility of workers	Technical
	Cells provide environment for mobility of multi-skilled labour	Physical

Table 10.1: How the Various Components of the Lean Production System Structure the Environment to Support the Focal System.

The main disadvantage of the approach would seem to be the possibility of reducing flexibility by depending too heavily on the physical environment for "jigging". I have shown above that this jigging need not be too purpose-specific, because the interpretive abilities of teams can fill the gaps when the

environment is sufficiently organised. In this regard, the flexibility of the symbol processing approach may be an illusion born of confusing thinking and acting. Metaphorically the choice is between a smart dilettante who relies on explicit reasoning, and a dumb artisan who works more fluently but has more trouble adapting to changing circumstances. The relative flexibility and efficiency of systems designed under the two approaches is a theoretical and empirical question for the future.

Chapter 11

Conclusion

11.1 Summary

This thesis has presented a new approach to understanding operations management in terms of the underlying theories of activity that inform various improvement initiatives and techniques. In particular it has argued that the dominant plan-based approaches to operations management are informed by a common-sense, but flawed, conception of activity as the production and implementation of plans derived from formal manipulations upon abstract models of the world. A detailed analysis of the conceptual basis and assumptions of this theory of activity demonstrates that the domain of applicability of such approaches is quite limited, and they should fail in characteristic ways already observed within other disciplines that have studied and experimented with the design of intentional agents from this standpoint. The thesis also presented an alternative framework for the design of operations systems based upon a radical theory of the nature of intentional activity currently being articulated in diverse disciplines, but which has as yet had little impact on management science.

The method used in the research has mainly been the presentation and analysis of a series of case studies of different kinds, and serving different purposes in the theoretical argument. In the first part of the thesis, the main case area of operations management in batch repetitive manufacturing was presented. Using the literature I discussed the range of techniques that are currently competing in this case area, and proposed that they can be viewed as forming two distinct and differently informed theories, or paradigms, of operations management. I demonstrated that existing comparisons of these techniques were inadequate, mainly because they were motivated by ill-posed question asked from the positions of the dominant paradigm, and that a new analysis of the paradigmatic differences was required. I continued this theme with actual company experience from three practitioner case studies, which present strong evidence that there are aspects of the nature of batch repetitive manufacture for which the dominant plan-based approach gives no account, namely the situatedness of manufacturing activity, and further, that this problem cannot be made to go away simply by making the planning approach more elaborate.

The second part presented the theoretical apparatus for a re-examination of operations management in terms of the design and control of intentional systems. This began with the careful definition of a level of abstraction at which all suitably complex systems can be described as intentional, and at which the observer or designer must make a commitment to a theory of the nature of intentional activity. Two competing theories of activity were then presented using literature based case studies from robotics, computer science, and human cognitive science, areas where these theories of activity have been most fully worked out.

These tools where then applied to the area of operations management. To do this, it was first necessary to make precise the notion that management theories are informed by theories of activity, and this was done by viewing operations management systems as intentional agents with the manager as the designer and controller of the intentional system. The plausibility of this move depended strongly of the notion, developed in chapter 5, of an activity level of abstraction at which discussion of the nature of activity can be separated from other current management debates. The plan-based CAPM approach could then

be associated with plan-based theories of activity and its assumptions and likely usefulness examined with a new level of clarity. An alternative proposal to view operations management from the point of view of situated action theories of activity was then developed. It was possible to reread the practitioner case studies in terms of this new approach to management, and make sense of the observations that came from these case studies. Finally it was argued that the Lean Production paradigm of operations management could be understood as an example of this new situational / interactional approach.

11.2 Contributions of the Research

The main contribution of this research is the rigorous working though of the notion that operations management theories and interventions are always informed by theories of the nature of on-going purposeful activity, and that the recognition of this relation has important implications for both the understanding of, and the practice of, operations management. These ideas were elaborated in three publications which appeared in 1995 and 1996, and this notion had not been explicitly stated prior to this. These publications already contained the following contributions:

- The demonstration of the situatedness of production work and its importance for operations management through the analysis of a case study of the experiences of a batch repetitive manufacturing firm over a five year period of reorganisation of its operations systems;

- The suggestion that the making of work practices routine, or *tacit*, is an important aspect of intervention in operations management, which cannot adequately be dealt with by the dominant representational theories of activity and management;

- A detailed explication of the assumptions that underlie the planning conception of activity and an examination of their validity using case studies from manufacturing operations management and from cognitive science disciplines, where the design of intentional systems has been developed in the most detailed and experimental way;

- The suggestion that radical alternative theories of activity and intentionality being articulated in many branches of cognitive science might form the basis for re-thinking operations management in a way that includes the situatedness of work activity;

- The articulation of a new understanding of the nature of operations management, management-as-organising, which requires that managers pay attention to structuring the environment of the operations system as well as to the design of focal system itself, in order that management goals emerge through the robust interaction of simple reactive systems with suitably structured environments. This proposal contrasts dramatically with the dominant position which recurs in many operations management settings and asserts that desired system outcomes are achieved through the automated production and implementation of plans or schedules derived by formal manipulations upon abstract representations of the world of activity;

- An argument to the effect that the paradigm shift in manufacturing operations management accompanying the influence of Lean Production in Western countries should be understood in terms of a change in the underlying conception of the nature intentional activity, and in particular, that Lean Production can be understood exemplifying this new situational or management-as-organising approach.

In addition to developing these themes in greater detail, the present work adds to these insights by providing a unified conceptual framework within which the connection of management theories to theories of activity can be made rigorous and explicit. This framework consists of:

- The recognition and definition of a distinct level of description or abstraction at which it is appropriate to discuss the activity of complex systems independent of other aspects of their nature, behaviour and realisation;

- The proposal that at this level of abstraction all sufficiently complex systems, including management systems, can be viewed as intentional agents. Although this idea is not new, the recognition that it allows management science to draw fruitfully upon other disciplinary areas that deal with intentional systems, such as artificial intelligence, robotic, computer science, sociology, anthropology, and ethology, is new. In addition the consistent adoption of the position that intentionality is a descriptive and predictive ascription of an observer overcomes much previous confusion in treating socio-technical systems as purposeful;

- A detailed explication of the way in which the lines of the main competing theories of activity follow from the way in which the outside observer conceives that the focal agent might make use of its knowledge and sense data to act in its environment. Planning or representational theories of activity write the agent's environment out of the story by assuming that the agent can build and maintain a symbolic model of the world and that its relation to it, which may be processed by formal means to determine a series of goal-directed actions. Situational or interactional theories of activity imagine an agent which only makes use of sense data and knowledge to which it has direct access, and this data and knowledge is therefore essentially agent-centred and limited by the agents situatedness, embodiment and thrownness. Action is a more or less direct response to the situation of the agent and therefore depends in essential ways upon the structure of the environment. The pieces of this picture are spread over a number of previous publications by other authors so the claim for a contribution here is mainly for a unified and clear explication rather than novelty;

- A method of connecting previous research on intentional systems in various disciplines with the nature and practice of operations management through the proposal that, at the activity level of abstraction, it is natural to view management as the design and control of an intentional agent, namely, the operations system. The viability of this proposal depends on the apparatus of the activity level of description and observer ascription of intentionality that was been built up. Rather than being a naive proposal about management in practice, it should be viewed as a useful abstraction (at the activity level of description) which allows the design and architecture of management systems, and the way in which a causal connection between management action and operative action is conceived, to be seen as a separate topic of study distinct from the many other ways in which management can be theorised. It amounts to associating the manager with the outside observer who attributes intentionality to the system. This avoids the confusion that can occur between the intentionality of agents that may be part of the system and intentionality of the system as a whole. Although this could be thought to be a weak form of intentionality, it nevertheless is sufficient to make a rigorous connection between theories of activity and management. This is because, facing the problem of designing a complex system, such a hypothetical manager must make important decisions, such as the way the system will be decomposed into parts and the way control systems will be implemented, which are determined by the theory of activity adopted. This proposal also overcomes the frequently voiced objection to viewing socio-technical systems as goal-directed, that such an approach writes human intentionality out of the story. Under the approach presented here, it can be seen that this is only true of representational theories of activity, since in situational theories of activity individual intentionality reappears as relevant to the human (political, cultural, cognitive, skill-based) aspects of the environment of the system;

- The thesis includes a discussion of theories of activity which are modifications or hybrids of the two main theories dealt with in the papers. Although they complicate the discussion somewhat, it is necessary to include them to negate such incorrect claims as that situated action theories are nothing but stimulus response theory or production systems revisited.

Together these additional developments constitute a rigorous articulation of what it means for a theory of activity to inform management, which was not present in the papers.

11.3 Implications and Future Directions

Inappropriate, expensive and disruptive interventions in all sectors of working life are continually reinvented based on naive and dangerous fantasies of control and order. Many of those who work with or under these systems, or who are involved in attempting to implement them, are dubious of their worth, but often argument against them is characterised as "user resistance", that is recalcitrance, rather than the articulation of real flaws in their underlying plan-based approach. Frequently problems with these systems are diagnosed as "implementation failures", or due to inadequate scoping of the system, rather than as due to any problem of principle. The research reported here provides a theoretical position from which the excesses of the dominant representational and plan-based theories of management can be attacked.

On the other hand, it provides a theoretical and practical basis for an alternative way of conceiving and implementing change in operations systems. On this view, management is a design process rather than a control process. The process is to analyse the environment in which the system is to act, then design of the simplest system which will achieve the desired goals in interaction with a properly structured or restructured environment, followed by testing and tweaking. This is probably the way most successful systems are actually designed in practice, but the lack of theoretical respectability for this approach can lead to it being derided as "muddling through". Such bottom-up system design is expected to be robust by virtue of the existence of stable low-level interactions as fall-backs, and is entirely compatible with and an evolutionary approach to change. This contrasts with the expensive, disruptive and risky "slash and burn" attitude associated with the implementation of increasingly ambitious "silver bullet" systems such as Enterprise Resource Planning and radical restructuring methodologies such as Business Process Re-engineering.

The extension of the work presented here should be directed toward making more explicit the principles of designing systems for particular environments and of the restructuring of environments in such a way that the environment shares the cognitive burden with the focal system, the outlines of which have been given in this thesis. This amounts to a call for the development of, in Agre's words, "principled characterisations of interaction between agents and their environments to guide explanation and design" (Agre 1995a). Agre and co-workers have made some progress on this project within the computational agent domain, and methods such as Group Technology represent a start within operations management. But compared to elaborate methodologies such as Structured Systems Analysis and Design which are available for the design of systems from the representational standpoint, this project is in its infancy, arguably due to the lack of theoretical legitimacy. The development of such a set of principles or tools would provide not only the basis for practical intervention and action research within particular companies attempting change, but also new tools for the analysis of success and failure cases with greater precision than provided by existing "success factors" approaches.

An important issue in such a development will be the role that computers will play in operations management. The answer was clear within the plan-based approach: computers are the logic engine which enables the automation and rapid regeneration of plans. The vision 18 years ago, when I entered the operation management area, was of robust computer work stations everywhere on the shop floor receiving data from operatives concerning the progress of work, in order to maintain a near-real-time model of the situation for the benefit of planners and accountants. This was a vision of computerised Taylorism. As production scheduling is simplified and made more tacit by the adoption of cellular layouts and self-managing teams, and the value of work-in-progress is rendered negligible by the adoption of lean methods, this activity directing and measuring role for the computer is diminished, although computerised planning still will be necessary to provide a resource for management thinking. This diminished role of computers in the control of shop floor activity is strikingly apparent when one visits companies that have implemented Lean Production ideas.

Just at the time when the significance of computers as logic engines is being challenged, there is emerging a recognition of a new role of computers and computer networks as *media*. It is likely that the new role for computers in operations management will be in the support of cooperative work through

computerised communication. This is entirely consistent with the proposal here that the true role of representations in activity is as resources for thinking about activity, rather than as control structures within which it is done. The framework described here could therefore be used to theorise computer mediated cooperative work.

References

Ackoff, R.L. and Emery, F.E. (1972). *On Purposeful Systems*. London, Tavistock Publications.

Aggarwal, S.C. (1974). "A Review of Current Inventory Theory and its Applications." *International Journal of Production Research* **12**(4): 443-482.

Aggarwal, S.C. (1985). "MRP, JIT, OPT, FMS? Making Sense of Production Operations Systems." *Harvard Business Review* **63**(5): 8-16.

Aggarwal, S.C. and Aggarwal, S. (1985). "The Management of Manufacturing Operations: An Appraisal of Recent Developments." *International Journal of Operations and Production Management* **5**(3): 21-38.

Agre, P.E. (1988) *The Dynamic Structure of Every Day Life*. Technical Report No. 1085. MIT Artificial Intelligence Laboratory, Cambridge, MA.

Agre, P.E. (1993). "The Symbolic World View: Reply to Vera and Simon." *Cognitive Science* **17**(1): 61-69.

Agre, P.E. (1995a). "Computational Research on Interaction and Agency." *Artificial Intelligence* **72**(1): 1-52.

Agre, P.E. (1995b). "The Soul Gained and Lost: Artificial Intelligence as a Philosophical Project." *Stanford Humanities Review* **4**(2): 1-19.

Agre, P.E. (1997). *Computation and Human Experience*. New York, Cambridge University Press.

Agre, P.E. and Chapman, D. (1987). "Pengi: An Implementation of a Theory of Activity." *Proceedings of AAAI-87. The Sixth National Conference on Artificial Intelligence*, Menlo Park, CA., Morgan Kaufman, Los Altos, CA.: 268-272.

Agre, P.E. and Chapman, D. (1990). "What are Plans for?" *Robotics and Autonomous Systems* **6**: 17-34.

Agre, P.E. and Horswill, I. (1992). "Cultural Support for Improvisation." *Proceedings of AAAI-92. The Tenth National Conference on Artificial Intelligence.*, Menlo Park, CA., AAAI Press: 363-368.

Agre, P.E. and Horswill, I. (1997). "Lifeworld Analysis." *Journal of Artificial Intelligence Research* **6**: 111-145.

Aksoy, Y. and Erenguc, S.S. (1988). "Mult-Item Inventory Models with Co-Ordinated Replenishments: A Survey." *International Journal of Operations and Production Management* **8**(1): 63-73.

Alfred, J.K. (1993). "Changing the Manufacturing Paradigm: A Blueprint for U.S. Industrial Competitiveness." *Industrial Engineering* **25**(8): 16-17.

Anderson, J.C., Schroder, R.C., Tupy, S.E. and White, E.M. (1982). "Material Requirements Planning Systems: The State of the Art." *Production and Inventory Management* **23**(4): 51-67.

Anderson, J.R. (1983). *The Architecture of Cognition*. Cambridge, MA, Harvard University Press.

Anderson, J.R. (1993). *Rules of the Mind*. Hillside, NJ, Lawrence Erlbaum Associates.

Ang, J.S.K., Sum, C.C. and Yang, K.K. (1994). "MRP II Company Profile and Implementation Problems: A Singapore Experience." *International Journal of Production Economics* **34**(1): 35-45.

Anonymous (1991). "Competition in Manufacturing Leads to MRP II." *Industrial Engineering* **23**(7): 10.

Argyris, C. and Schon, D. (1980). "What is an Organization that it may Learn." *Organizations as Systems*. Lockett, M. and Spear, R. (Eds.). Milton Keynes, The Open University Press: 128-137.

Argyris, C. and Schon, D.A. (1978). *Organizational Learning: A Theory of Action Perspective*. Reading, MA, Addison-Wesley.

Ashby, W.R. (1956). *An Introduction to Cybernetics*. London, Chapman and Hall.

Ashton, J.E., Johnson, M.D. and Cook, F.X. (1990). "Shop Floor Control in a System Job Shop: Definitely Not MRP." *Production and Inventory Management Journal* **31**(2): 27-32.

Baker, K.R., Powell, S.G. and Pyke, D.F. (1990). "The performance of push and pull systems: a corrected analysis." *International Journal of Production Research* **28**(9): 1731-1736.

Bartezzaghi, E., Turco, F. and Spina, G. (1992). "The Impact of the Just-In-Time Approach on Production Systems Performance: A Survey of Italian Industry." *International Journal of Operations and Production Management* **12**(1): 5-17.

Bateson, G. (1979). *Mind and Nature: A Necessary Unity*. New York, E. P. Dutton.

Beer, R.D. (1995). "A Dynamical Systems Perspective on Agent Environment Interaction." *Artificial Intelligence* **72**(1-2): 173-215.

Belt, B. (1987). "MRP and Kanban - A Possible Synergy?" *Production and Inventory Management* **28**(1): 71-80.

Benner, P. (1984). *From Novice to Expert: Excellence and Power in Clinical Nursing Practice*. Reading, MA, Addison-Wesley.

Betts, J.M. and Johnston, R.B. (1998). "Efficiency vs. Flexibility: A Risk vs. Return Analysis of Batch Sizing Decisions." *Proceedings of The 3rd International Conference on Management (ICM-98)*, Shanghai, Republic of China, Springer Verlag: R268.

Boggs, R.A. (1988). "Implication of Formal Production Information Systems for Production Management." *International Journal of Operations and Production Management* **8**(4): 22-33.

Bonney, M.C. and Head, M.A. (1993). "Progress on Developing a Framework for Computer Aided Production Management." *International Journal of Production Economics* **30,31**: 345-353.

Bourdieu, P. (1977). *Outline of a Theory of Practice*. Cambridge, Cambridge University Press.

Bourdieu, P. (1990). *The Logic of Practice*. Stanford, Stanford University Press.

Brandimarte, P. and Villa, A. (1995). *Advanced Models for Manufacturing Management Systems Management*. Boca Raton, FL, CRC Press.

Bratman, M. (1987). *Intentions, Plans and Practical Reasoning*. Cambridge, MA., Harvard University Press.

Bregman, R.L. (1994). "An Analytical Framework for Comparing Material Requirements Planning to Reorder Point Systems." *European Journal of Operations Research* **75**(1): 74-88.

Breznier, N.W. (1981). "The Odyssey of Inventory Management." *Production and Inventory Management* **22**: 22-36.

Broadbent, D.E. (1958). *Perception and Communication*. Oxford, Pergamon.

Brodner, P. (1990). "Technocentric-Anthropocentric Approaches: Towards Skill-Based Manufacturing." *New Technology and Manufacturing Management*. Warner, M., Wobbe, W. and Brodner, P. (Eds.). Chichester, UK., John Wliey & Sons: 101-111.

Brooks, R.A. (1986). "A Robust Layered Control System For A Mobile Robot." *IEEE Journal of Robotics and Automation* **2**(1): 14-23.

Brooks, R.A. (1989). "A Robot that Walks: Emergent Behaviors from a Carefully Evolved Network." *Neural Computation* **1**(2): 253-262.

Brooks, R.A. (1990). "Elephants Don't Play Chess." *Robotics and Autonomous Systems* **6**(1-2): 3-15.

Brooks, R.A. (1991a) *Intelligence Without Reason*. A.I. Memo No. 1293. MIT Artificial Intelligence Laboratory,

Brooks, R.A. (1991b). "Intelligence Without Representation." *Artificial Intelligence* **47**(1-3): 139-159.

Brooks, R.A. (1991c). "New Approaches to Robotics." *Science* **253**(9): 1227-1232.

Brooks, R.A. and Flynn, A.M. (1989). "Fast, Cheap and Out of Control: A Robot Invasion of the Solar System." *Journal of the British Interplanetary Society* **42**(10): 478-485.

Brooks, R.A. and Stein, L.A. (1993) *Building Brains for Bodies*. Technical Report No. 1439. MIT Artificial Intelligence Laboratory, Cambridge, MA

Bryan, W.L. and Hartner, N. (1897). "Studies in the Physiology and Psychology of the Telegraphic Language." *Psychological Review* **4**: 27-53.

Burbidge, J.L. (1988). *Group Technology, The State of the Art*. London, Kogan Page.

Burbidge, J.L. (1990). "Production Control: A Universal Conceptual Framework." *Production Planning and Control* **1**(1): 3-16.

Burcher, P. (1991). "Closing the Loop in Manufacturing Resource Planning Systems?" *Control*(August/September): 35-39.

Buxey, G. (1989). "Production Scheduling: Practice and Theory." *European Journal of Operations Research* **39**(1): 17-31.

Buzacott, J.A. (1995). "A New Perspective on Paradigms in Manufacturing." *Journal of Manufacturing Systems* **14**(2): 118-125.

Bylinski, G. (1983). "An Efficiency Guru with a Brown Box." *Fortune* **108**(5): 120.

Callon, M. (1986). "Some Elements of a Sociology of Translation: Domestication of the Scallops and the Fishermen of St Brieu Bay." *Power, Action and Belief: A New Sociology af Knowledge?* Law, J. (Ed.). London, Routledge and Kegan Paul: 196-233.

Cannon, W.B. (1929). "Organization for Physiological Homoeostasis." *Physiological Reviews* **9**: 399-431.

Cannon, W.B. (1933). *The Wisdom of the Body*. New York, Norton.

Cerveny, R.P. and Scott, L.W. (1989). "A Survey of MRP Implementation." *Production and Inventory Management Journal* **30**(3): 31-34.

Chantland, R. (1993). "Strategic Subcontracting and the OPT System." *Proceedings of APICS 26th Annual International Conference*: 601-602.

Chapman, D. (1988). "Planning for Conjunctive Goals." *Artificial Intelligence* **32**(3): 333-377.

Chapman, D. (1991). *Vision, Instruction, and Action*. Cambridge, MA., The MIT Press.

Chase, W.G. and Simon, H.A. (1973). "Perception in Chess." *Cognitive Psychology* **4**: 55-81.

Checkland, P.B. (1981). *Systems Thinking, Systems Practice*. Chichester, John Wiley.

Churchman, C.W. (1976). *The Design of Inquiring Systems: Basic Concepts of Systems and Organisations*, Basic Books Inc.

Clancey, W.J. (1991). "Book Review of Israel Rosenfield, The Invention of Memory: A New View of the Brain." *Artificial Intelligence* **50**: 241-284.

Clancey, W.J. (1993). "Situated Action: A Neuropsychological Interpretation. Response to Vera and Simon." *Cognitive Science* **17**(1): 87-116.

Clancey, W.J. (1997). *Situated Cognition: On Human Knowledge and Computer Representations*. Cambridge, Cambridge University Press.

Clode, D.M. (1993). "A survey of U.K. Manufacturing Control Over the Past Ten Years." *Production and Inventory Management Journal* **34**(2): 53-56.

Cohen, P.R., Greenberg, M.L., Hart, D.M. and Low, A.E. (1989). "Trial by Fire: Understanding the Design Requirements for Agents in Complex Environments." *AI Magazine* **10**(Fall): 32-48.

Collins, H.M. (1981). "Stages in the Empirical Programme of Relativism." *Social Studies of Science* **11**: 3-10.

Collins, H.M. (1990). *Artificial Experts: Social Knowledge and Intelligent Machines*. Cambridge, MA, The MIT Press.

Conant, R.C. (1969). "The Information Transfer Required in Regulatory Processes." *IEEE Transaction on Systems Science and Cybernetics* **5**(4): 334-338.

Conant, R.C. (1970). "Every Good Regulator of a System Must be a Model of that System." *International Journal of Systems Science* **1**(2): 89-97.

Cowton, C.J. and Vail, R.L. (1994). "Making Sense of Just-In-Time Production: A Resource-based Perspective." *OMEGA International Journal of Management Science* **22**(5): 427-441.

Craik, K.J.W. (1943). *The Nature of Explanation*. Cambridge, Cambridge University Press.

Craik, K.J.W. (1948). "Theory of the Human Operator in Control Systems: II. Man as an Element in a Control System." *British Journal of Psychology* **38**: 142-148.

Crevier, D. (1993). *AI: The Tumultuous History of the Search for Artificial Intelligence*. New York, Basic Books.

Crutchfield, J.P., Farmer, J.D., Packard, N.H. and Shaw, R.S. (1996). "Chaos." *Scientific American* **255**(6): 38-49.

Davy, J.A., White, R.E., Merritt, N.J. and Gritzmacher, K. (1992). "A Derivation of the Underlying Constructs of Just-In-Time Management Systems." *Academy of Management Journal* **35**(3): 653-670.

de Groot, A. (1978). *Thought and Choice in Chess*. The Hague, Mouton.

Dennett, D.C. (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*. Brighton, Harvester.

Dennett, D.C. (1996). *Kinds of Mind: Toward an Understanding of Consciousness*. New York, Basic books.

DeToni, A., Caputo, M. and Vinelli, A. (1988). "Production Management Techniques: Push-Pull Classification and Application Conditions." *International Journal of Operations and Production Management* **8**(2): 35-51.

Ding, F.Y. and Yuen, M.N. (1991). "A Modified MRP for a Production Systems with the Cooextistance of MRP and Kanbans." *Journal of Operations Management* **10**(2): 267-277.

Discenza, R. and McFadden, F.R. (1988). "The Integration of MRP II and JIT Through Software Unification." *Production and Inventory Management Journal* **29**(3): 49-53.

Doll, W.J. and Vonderembse, M.A. (1991). "The Evolution of Manufacturing Systems: Towards the Post-Industrial Enterprise." *MEGA International Journal of Management Science* **19**(5): 401-411.

Dreyfus, H.L. (1991). *Being-In-The-World: A Commentary on Heidegger's Being and Time, Division 1*. Cambridge Massachusetts, The MIT Press.

Dreyfus, H.L. (1992). *What Computers Still Can't Do*. Cambridge, Massachusetts, The MIT Press.

Dreyfus, H.L. and Dreyfus, S. (1989). "Why Computers May Never Think Like People." *Computers in the Human Context*. Forester, T. (Ed.). Oxford, Basil Blackwell: 125-143.

Dreyfus, H.L. and Dreyfus, S.E. (1996). *Mind over Machine: The Power of Human Intuition and Experience in the Era of the Computer*. New York, The Free Press.

Drucker, P.F. (1954). *The Practice of Management*. New York, Harper and Row.

Duchessi, P., Schaninger, C.M. and Hobbs, D.R. (1989). "Implementing a Manufacturing Planning and Control Information System." *California Management Review* **31**(3): 75-90.

Duchessi, P., Schaninger, C.M., Hobbs, D.R. and Pentak, L.P. (1988). "Determinants of Success in Implementing Material Requirements Planning (MRP)." *Journal of Manufacturing and Operations Management* **1**(3): 263-304.

Eames, C. and Eames, R. (1990). *A Computer Perspective: Background to the Computer Age*. Cambridge, MA, Harvard University Press.

Ebbinghaus, H. (1964). *Memory: A Contribution to Experimental Psychology*. New York, Dover.

Ericsson, K.A. (1996a). "The Acquisition of Expert Performance: An Introduction to Some of the Issues." *The Road to Excellence: The Road to Expert Performance in the Arts and Sciences, Sports, and Games*. Ericsson, K.A. (Ed.). Mahwah, NJ, Lawrence Erlbaum: 1-50.

Ericsson, K.A., Ed. (1996b). *The Road to Excellence: The Road to Expert Performance in the Arts and Sciences, Sports, and Games*. Mahwah, NJ, Lawrence Erlbaum.

Ericsson, K.A., Krampe, R.T. and Tesch-Romer, C. (1993). "The Role of Deliberate Practice in the Acquisition of Expert Performance." *Psychological Review* **100**: 363-406.

Ericsson, K.A. and Smith, J., Eds. (1991). *Toward a General Theory of Expertise: Prospects and Limits*. Cambridge, Cambridge University Press.

Ettienne, E.C. (1983a). "Comparative Behavior of Statistical Inventory control and Materials Requirements Planning Under Varying Industrial Conditions: Presentation and Evaluation of a Model." *International Journal of Operations and Production Management* **3**(2): 3-17.

Ettienne, E.C. (1983b). "MRP May Not Be Right For You: At Least Not Yet." *Production and Inventory Management* **24**(3): 33-45.

Fikes, R.E., Hart, P.E. and Nilssen, N.J. (1972). "Learning and Executing Generalized Robot Plans." *Artificial Intelligence* **3**(4): 251-288.

Fikes, R.E. and Nilson, N.J. (1971). "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving." *Artificial Intelligence* **2**(3): 189-208.

Firby, J.A. (1987). "An Investigation into Reactive Planning in Complex Environments." *Proceedings of the Tenth Joint International Conference on Artificial Intelligence*, Milan, Italy: 202-206.

Fisher, K. (1981). "How to Implement MRP Successfully." *Production and Inventory Management* **22**(4): 36-54.

Fitts, P.M. (1964). "Perceptual-Motor Skill Learning." *Categories of Human Learning*. Melton, A.W. (Ed.). New York, Academic Press: 243-285.

Flapper, S.D.P., Miltenburg, G.J. and Wijngaard, J. (1991). "Embedding JIT into MRP." *International Journal of Production Research* **29**(2): 329-341.

Fleishman, E.A. and Rich, S. (1963). "Role of Kinesthetic and Spacial-Visual Abilities in Perceptual-Motor Learning." *Journal of Experimental Psychology* **66**: 6-11.

Flood, R.I. and Jackson, M.C. (1991). *Critical Systems Thinking*. Chichester, John Wiley & Sons.

Fortuin, L. (1977). "A Survey of Literature on Reordering of Stock Items for Production Inventories." *International Journal of Production Research* **15**(1): 87-105.

Fortuin, L. (1981). "A Comparison of SIC and MRP, Two Methods for Material Procurement in Industry." *European Journal of Operations Research* **6**(4): 386-392.

Fox, R.E. (1983). "OPT vs. MRP: Throughware vs. Software. Part I." *Inventories & Production Magazine*(November-December): 23-33.

Fox, R.E. (1984). "OPT vs. MRP: Throughware vs. Software. Part II." *Inventories & Production Magazine*(January-February): 33-35.

Fox, R.E. (1987). "OPT: Leapfrogging the Japanese." *Just-In-Time Manufacture*. Voss, C.A. (Ed.). London, IFS Ltd. UK: 107-122.

Frizelle, G.D.M. (1989). "OPT in Perspective." *Advanced Manufacturing Engineering* **1**(1): 74-80.

Fry, T.D., Cox, J.F. and Blackstone, J.H. (1992). "An Analysis and Discussion of the Optimized Production Technology Software and its Use." *Production and Operations Management* **1**(2): 229-242.

Gaines, M.W. (1905). "Tabulating-Machine Cost-Accounting for Factories of Diversified Product." *The Engineering Magazine* **37**(12): 364-373.

Galliers, R.D. (1992). "Choosing Information Systems Research Approaches." *Information Systems Research: Issues, Methods, and Practical Guidlines*. Galliers, R.D. (Ed.). Oxford, Blackwell Scientific: 144-162.

Gardiner, K.M. (1996). "An Integrated Design Strategy for Future Manufacturing Systems." *Journal of Manufacturing Systems* **15**(1): 52-61.

Gardiner, S.C., Blackstone, J.H. and Gardiner, L.R. (1993). "Drum-Buffer-Rope and Buffer Management: Impact on Production Management Study and Practices." *International Journal of Operations and Production Management* **13**(6): 68-78.

Gardner, H. (1985). *The Mind's New Science*. New York, Basic Books, Inc.

Garfinkel, H. (1967). *Studies in Ethnomethodology*. Englewood Cliffs, NJ, Prentice Hall.

Gelders, L.F. and Van Wassenhove, L.N. (1985). "Capacity Planning in MRP, JIT and OPT: A Critique." *Engineering Costs and Production Economics* **9**(1-3): 201-209.

Georgeff, M.P. (1987a). "Planning." *Annual Review of Computer Science* **2**: 359-400.

Georgeff, M.P. (1987b). "Reactive Reasoning and Planning." *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Seattle, WA: 677-682.

Gilbert, J.P. (1990). "The State of JIT Implementation and Development in the USA." *International Journal of Production Research* **28**(6): 1099-1109.

Goldratt, E.M. (1988). "Computerised Shop Floor Scheduling." *International Journal of Production Research* **26**(3): 443-455.

Goldratt, E.M. and Fox, J. (1984). *The Goal*. Aldershot, Gower.

Golhar, D.Y. and Stamm, C.L. (1991). "The Just-In-Time Philosophy: A Literature Review." *International Journal of Production Research* **29**(4): 657-676.

Goodrich, T. (1989). "JIT & MRP Can Work Together." *Production Engineering* **36**(4): 46-48.

Goyal, S.K. and Deshmukh, S.G. (1992). "A Critique Of Literature On Just-In-Time Manufacture." *International Journal of Operations and Production Management* **12**(1): 18-28.

Greeno, J.G. (1989). "Situations, Mental Models and Generative Knowledge." *Complex Information Processing: The Impact of Herbert A. Simon*. Klar, D. and Kotovsky, K. (Eds.). Hillsdale, NJ, USA., Erlbaum.

Groves, B.R. (1990). "The Missing Element in MRP - People." *Production and Inventory Management Journal* **31**(4): 60-64.

Grunwald, H., Striekwold, P.E.T. and Weeda, P.J. (1989). "A framework for quantitative comparison of production control concepts." *International Journal of Production Research* **27**(2): 281-292.

Gummeson, E. (1988). *Qualitative Methods in Management Research*. Lund, Chartwell-Bratt.

Hakim, C. (1987). *Research Design: Strategies and Choices in Social Science Research*. London, Allen & Unwin.

Hall, R.H. (1980). "Organizational Goals." *Organizations as Systems*. Locket, M. and Spear, R. (Eds.). Milton Keynes, The Open University Press: 88-96.

Hall, R.W. (1983). *Zero Inventories*. Homewood, IL., Dow Jones Irwin.

Hamid, A.A., Agus, A. and Hassan, M.E.M. (1991). "Computerised Materials Requirement Planning in Manufacturing Companies in Malaysia." *International Journal of Production Economics* **25**: 73-79.

Hammer, M. (1990). "Reengineering Work: Don't Automate, Obliterate." *Harvard Business Review* **68**(4): 104-112.

Hammond, K.J., Converse, T.M. and Grass, J.W. (1995). "The Stabilization of Environments." *Artificial Intelligence* **72**(1): 305-328.

Harris, F.W. (1913). "How Many Parts to Make at Once." *Factory, The Magazine of Management* **10**(2): 135-136, 152.

Harris, F.W. (1915). "What Quantity to Make at Once." *The Library of Factory Management, Vol V. Operation and Costs*. Chicago, A. W. Shaw Company: 47-52.

Hastings, N.A.J., Marshall, P.H. and Willis, R.J. (1982). "Schedule Based M.R.P.: An Integrated Approach to Scheduling and Material Requirements Planning." *Journal of the Operations Research Society* **33**(11): 1021-1029.

Hastings, N.A.J. and Packham, B.J. (1985). "Developments in Schedule Based Manufacturing Resource Planning." *Proceedings of 7th National Conference of the Australian Society for Operations Research*, Adelaide, Australia.

Hastings, N.A.J. and Yeh, C.H. (1992). "Bill of Manufacture." *Production and Inventory Management Journal* **33**(4): 27-31.

Hayes, P.J., Ford, K.M. and Agnew, N. (1994). "On Babies and Bathwater: A Cautionary Tale." *AI Magazine* **15**(Fall): 15-26.

Hayes, R.H. and Wheelwright, S. (1984). *Restoring our Competitive Edge: Competing Through Manufacturing*. New York, John Wiley & Sons.

Hazen, H.L. (1934). "The Theory of Servo-Mechanisms." *Journal of the Franklin Institute* **218**: 279-331.

Heidegger, M. (1961). *Being and Time*. New York, Harper & Row.

Hendriks-Jansen, H. (1996). *Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought*. Cambridge, MA, The MIT Press.

Hieko, L. (1989a). "A Simple Framework for Understanding JIT." *Production and Inventory Management Journal* **30**(4): 61-63.

Hieko, L. (1989b). "Some Relationships Between Japanese Culture and Just-In-Time." *Academy of Management Executive* **111**(4): 318-321.

Hirsch, A.A. (1996). "Has Inventory Management in the US Become More Efficient and Flexible? A Macroeconomic Perspective." *International Journal of Production Economics* **45**(1-3): 37-46.

Hofstede, G. (1978). "The Poverty of Management Control Philosophy." *Academy of Management Journal* **3**(3): 450-461.

Hofstede, G. (1981). "Management Control of Public and Not-For-Profit Activities." *Accounting, Organizations and Society* **6**(3): 193-211.

Holland, J.H. (1995). *Hidden Order: How Adaptation Builds Complexity*. Reading, MA, Addison-Wesley.

Homer, G. (1993). Personal Communication. Glyn Homer keeps an index of MRP system usage in Australia.

Hopp, W.J. and Spearman, M.L. (1996). *Factory Physics: Foundations of Manufacturing Management*. Chicago, IL, Irwin.

Horswill, I. (1995). "Analysis of Adaption and Environment." *Artificial Intelligence* **73**(1-2): 1-30.

Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA, The MIT Press.

Jacobs, F.R. (1983). "The OPT Scheduling System: A Review Of A New Production Scheduling System." *Production and Inventory Management* **24**(3): 47-51.

Jacobs, F.R. and Whybark, D.C. (1992). "A Comparison of Reorder Point and Material Requirements Planning Inventory Control Logic." *Decision Sciences* **23**: 332-342.

Jenner, R.A. (1998). "Dissipative Enterprise, Chaos, and the Principles of Lean Production." *OMEGA International Journal of Management Science* **26**(3): 397-401.

Johnson, D. and LaBarre, J. (1991). "Inventory Control: A Target for MIS Support." *Production and Inventory Management Journal* **32**(1): 62-66.

Johnson-Laird, P. (1993). *The Computer and the Mind: an Introduction to the Cognitive Sciences*. London, Fontana Press.

Johnston, R.B. (1995). "Making Manufacturing Practices Tacit: A Case Study of Computer Aided Production Management and Lean Production." *Journal of the Operational Research Society* **46**(10): 1174 - 1183.

Johnston, R.B. (1996a). "From Efficiency to Flexibility: Entropic Measures of Market Complexity and Production Flexibility." *Proceedings of Complex 96: From Local Interactions to Global Phenomena*, Charles Sturt University, Albury, NSW, IOS Press, Amsterdam: 130 - 140.

Johnston, R.B. (1996b). "Thinking about Acting and Acting are Different - Reply to the Comment on my Paper." *Journal of the Operational Research Society* **47**(8): 1080-1082.

Johnston, R.B. (1997a). "Reconciling the Economic Batch Size Model with Current Manufacturing Rhetoric." *Proceedings of The Fourth Conference of the Association Asian-Pacific Operational Research Societies*, Melbourne: MB6.4.

Johnston, R.B. (1997b). "The Trade-off Between Efficiency and Flexibility in Operations Management." *Proceedings of The Fourth Conference of the Association Asian-Pacific Operational Research Societies*, Melbourne: WA2.1.

Johnston, R.B. and Betts, J.M. (1996a). "A Complex Systems Simulation of Planning versus Reacting for Inventory Replenishment." *Proceedings of Complex 96: From Local Interactions to Global Phenomena*, Charles Sturt University, Albury, NSW, IOS Press, Amsterdam: 339 - 350.

Johnston, R.B. and Betts, J.M. (1996b). "Evaluating Planning vs. Reactive Trading Strategies using a Complex Systems Simulation." *Proceedings of Commerce, Complexity and Evolution*, University of New South Wales, University of NSW.

Johnston, R.B. and Betts, J.M. (1998). "Evaluating Planning vs. Reactive Trading Strategies using a Complex Systems Simulation." *International Symposia in Economic Theory and Econometrics, Vol 11*. Barnett, W. (Ed.). Cambridge, Cambridge University Press: (Forthcoming).

Johnston, R.B. and Brennan, M. (1996). "Planning or Organising: The Significance of Theories of Activity for the Management of Operations." *OMEGA, International Journal of Management Science* **24**(4): 367-384.

Johnston, R.B. and Lee, R.P.W. (1997). "The Role of Electronic Commerce Technologies in Just-In-Time Replenishment." *Proceedings of Hawaii International Conference on Systems Science*, Hawaii, IEEE Computer Society Press: 439-448.

Jones, G. and Roberts, M. (1990). *Optimised Production Technology (OPT)*. Bedford, UK., IFS Publications.

Karmarker, U. (1989). "Getting Control of Just-In-Time." *Harvard Business Review* **67**(5): 122-131.

Kirsh, D. (1995). "The Intelligent Use of Space." *Artificial Intelligence* **73**(1): 31-68.

Klien, H.K. and Hirscheim, R.A. (1987). "A Comparative Framework of Data Modelling Paradigms and Approaches." *The Computer Journal* **30**(1): 8-15.

Kochhar, A.K. (1982). "Why Computerised Production Control Systems Continue to Disappoint." *The Production Engineer* **61**(2): 37-38.

Krajewski, L.J., King, B.E., Ritzman, L.P. and Wong, D.S. (1987). "Kanban, MRP, and Shaping the Manufacturing Environment." *Management Science* **33**(1): 39-57.

LaForge, R.L. and Sturr, V.l. (1986). "MRP Practices in a Random Sample of Manufacturing Firms." *Production and Inventory Management* **27**(3): 129-136.

Lakoff, G. (1987). *Women, Fire and Dangerous Things: What Categories Reveal about the Mind.* Chicago, The University of Chicago Press.

Lambrecht, M.R. and Decaluwe, L. (1988). "JIT and Constraint Theory: The Issue of Bottleneck Management." *Production and Inventory Management Journal* **29**(3): 61-65.

Lave, J. (1988). *Cognition in Practice: Mind, Mathematics and Culture in Everyday Life.* New York, Cambridge University Press.

Lave, J. and Wegner, E. (1991). *Situated Learning: Legitimate Peripheral Participation.* Cambridge, Cambridge University Press.

Lawrie, N.L. (1996). "Making manufacturing Practices Tacit - A Comment on Johnston's Paper." *Journal of the Operational Research Society* **47**(8): 1078-1080.

Leavitt, J. (1979). "Cognitive Demands of Skating and Stick Handling in Ice Hockey." *Canadian Journal of Applied Sport Sciences* **4**: 46-55.

Lee, C.Y. (1993). "A Recent Development of the Integrated Manufacturing System: A Hybrid of MRP and JIT." *International Journal of Operations and Production Management* **13**(4): 3-17.

Lee, L.C. (1989). "A Comparative Study of the Push and Pull Production Systems." *International Journal of Operations and Production Management* **9**(4): 5-19.

Little, D. and Johnson, S.B. (1990). "Survey of UK Manufacturing Control Practice." *Control*(April/May): 31-32.

Louis, R.S. (1991). "MRP III, Material Acquisition Systems." *Production and Inventory Management* **32**(July): 26-28.

Lowe, J. and Sim, A.B. (1993). "The Diffusion of a Manufacturing Innovation: The Case of JIT and MRPII." *International Journal of Technology Management* **8**(3-5): 244-258.

Lundrigan, R. (1986). "What is this Thing Called OPT?" *Production and Inventory Management* **27**(2): 2-11.

Luscombe, M. (1991). "Design and Implementation of Integrated Production Control Systems." *Integrated Manufacturing Systems* **2**(4): 4-8.

Maes, J. and Wassenhove, L.N.V. (1991). "Functionalities of Production-Inventory Control Systems." *Production Planning and Control* **2**(3): 219-227.

Maes, P. (1990). "Situated Agents Can have Goals." *Designing Autonomous Agents*. Maes, P. (Ed.). Cambridge, MA., MIT Press: 49-70.

Manning, E.J. (1983). "OPT. A Production Technique Case Study." *Proceedings of APICS 26th Annual International Conference*: 596-598.

Marr, D. (1982). *Vision*. San Fransisco, CA., Freeman.

Maruyama, M. (1963). "The Second Cybernetics: Deviation Amplifying Mutual Causal Processes." *American Scientist* **51**: 164-179.

Mataric, M.J. (1992). "Integration of Representation into Goal-Driven Behavior-Based Robots." *IEEE Journal of Robotics and Automation* **8**(3): 304-312.

Matsuura, H., Kurosu, S. and Lehtimarki, A. (1995). "Concepts, Practices and Expectation of MRP, JIT, and OPT in Finland and Japan." *International Journal of Production Economics* **41**(1): 267-272.

Maturana, H. and Varela, F. (1980). *Autopoiesis and Cognition*. Boston, D. Reidel.

Maturana, H.R. and Varela, F.J. (1992). *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston, Shambala Publications, Inc.

McGregor, D. (1960). *The Human Side of Enterprise*. New York, McGraw-Hill.

McLaughlin, C.P., Vastag, G. and Whybark, D.C. (1994). "Statistical Inventory Control in Theory and Practice." *International Journal of Production Economics* **35**(1-3): 161-169.

Meleton, M.P. (1986). "OPT-Fantasy or Breakthrough." *Production and Inventory Management* **27**(2): 13-21.

Merleau-Ponty, M. (1962). *Phenomenology of Perception*. London, Routledge & Kegan Paul.

Miller, G.A. (1956). "The Magic Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information." *Psychology Review* **63**: 81-97.

Miller, G.A., Galanter, E. and Pribram, K.H. (1960). *Plans and the Structure of Behaviour*, Holt, Rinehart and Winston Inc.

Miller, J.G., Berry, W.L. and Lai, C.F. (1976). "A Comparison of Alternative Forecasting Strategies for Multi-Stage Production-Inventory Control Systems." *Decision Sciences* **7**(4): 714-724.

Miltenburg, J. (1993). "On the Equivalence of JIT and MRP as Technologies for Reducing Wastes in Manufacturing." *Naval Research Logistics* **40**: 905-924.

Miltenburg, J. and Wijngaard, J. (1991). "Designing and Phasing in Just-In-Time Production Systems." *International Journal of Production Research* **29**(1): 115-131.

Mingers, J. (1995). *Self-Producing Systems: Implication and Application of Autopoiesis*. New York, Plenum Press.

Minsky (1985). *The Society of Mind*. New York, Simon and Schuster.

Monden, Y. (1983). *Toyota Production System*. Norcross, Institute of Industrial Engineers.

Monniot, J.P., Rhodes, D.J., Towill, D.R. and Waterlow, J.G. (1987). "Report of a Study of Computer Aided Production Management in UK Batch Manufacturing." *International Journal of Operations and Production Management* **7**(2): 7-42.

Morgan, G. (1986). *Images of Organization*. Newbury Park, CA, Sage.

Nardi, B., Ed. (1996). *Context and Consciousness: Activity Theory and Human-Computer Interaction*. Cambridge, MA, MIT Press.

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA, Harvard University Press.

Newell, A. and Rosenbloom, P.S. (1981). "Mechanisms of Skill Acquisition and the Law of Practice." *Cognitive Skills and their Acquisition*. Anderson, J.R. (Ed.). Hilside, NJ, Lawrence Erlbaum: 1-55.

Newell, A. and Simon, H.A. (1963). "GPS: A Program that Simulates Human Thought." *Computer and Thought*. Feigenbaum, E.A. and Feldman, J. (Eds.). New York, McGraw-Hill: 279-293.

Newell, A. and Simon, H.A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ, Prentice-Hall.

Newell, A. and Simon, H.A. (1981). "Computer Science as Empirical Inquiry: Symbols and Search." *Mind Design: Philosophy, Psychology, Artificial Intelligence*. Haugerland, J. (Ed.). Cambridge, Mass., MIT Press: 35-66.

Newell, S. and Clark, P. (1990). "The Importance of Extra-Organizational Networks in the Diffusion and Appropriation of New Technologies." *Knowledge: Creation, Diffusion, Utilization.* **12**(2): 199-212.

Newell, S., Swan, J. and Clark, P. (1993). "The Importance of User Design in the Adoption of New Information Technologies: The Example of Production and Inventory Control Systems (PICS)." *International Journal of Operations and Production Management* **13**(2): 4-22.

Newman, W. and Sridharan, V. (1992). "Manufacturing Planning and Control: Is there One Definitive Answer?" *Production and Inventory Management Journal* **33**(1): 50-53.

Nohria, N. and Berkley, J.D. (1994). "An Action Perspective: The Crux of the New Management." *California Management Review* **36**(4): 70-92.

O'Hare, G.M.P. and Jennings, N.R., Eds. (1996). *Foundations of Distributed Artificial Intelligence*. New York, John Wiley and Sons.

O'Neill, H. and Sackett, P. (1994). "The Extended Manufacturing Enterprise Paradigm." *Management Decision* **32**(8): 42-49.

Orlicky, J. (1975). *Material Requirements Planning*. New York, Mc Graw-Hill Book Company.

Pfeffer, J. (1981). *Power in Organizations*. Marshfield, MA, Pitman.

Pickering, A. (1995). *The Mangle of Practice*. Chicago, IL, University of Chicago Press.

Plenert, G. and Best, T.D. (1986). "MRP, JIT, and OPT What's "Best"." *Production and Inventory Management* **27**(2): 22-28.

Plossl, G.W. (1980). "MRP Yesterday, Today and Tomorrow." *Production and Inventory Management* **21**(3): 1-9.

Plossl, G.W. and Wight, O.W. (1967). *Production and Inventory Control*, Prentice Hall.

Polanyi, M. (1962). "Tacit Knowing: Its Bearing on Some Problems of Philosophy." *Reviews of Modern Physics* **34**(1): 601-616.

Polanyi, M. (1966). *The Tacit Dimension*. New York, Anchor Books.

Pollack, M.E. (1992). "The Use of Plans." *Artificial Intelligence* **57**(1): 43-68.

Pollack, M.E. and Ringuette, M. (1990). "Introducing the Tileworld: Experimentally Evaluating Agent Architectures." *Proceedings of AAAI-90*, Boston, MA, AAAI Press: 183-189.

Prajogo, N.H. and Johnston, R.B. (1998). "A Barriers Framework for Just-In-Time Implementation in Small Manufacturing Enterprises." *International Journal of Technology Management*: *(Forthcoming)*.

Prigogine, I. (1980). *From Being to Becoming*. San Francisco, W. H. Freeman.

Proctor, R.W. and Dutta, A. (1995). *Skill Acquisition and Human Performance*. Thousand Oaks, CA, Sage Publications Inc.

Ptak, C.A. (1991). "MRP, MRP II, OPT, JIT and CIM - Succession, Evolution, or Necessary Combination." *Production and Inventory Management Journal* **32**(2): 7-11.

Pyke, D.F. and Cohen, M.A. (1990). "Push and Pull in Manufacturing and Distribution Systems." *Journal of Operations Management* **9**(1): 24-43.

Pylyshyn, Z.W. (1988). *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*. Norwood, NJ, Ablex Publishing Corporation.

Ramarapu, N., Meha, S. and Frolick, M.N. (1995). "A Comparative Analysis and Review of JIT "Implementation" Research." *International Journal of Operations and Production Management* **15**(1): 38-49.

Ramsay, M.L., Brown, S. and Tabibzedeh, K. (1990). "Push, Pull and Squeeze. Shop Floor Control With Computer Simulation." *Industrial Engineering* **22**(2): 39-45.

Raphael, B. (1976). "The Robot "Shakey" and "his" Successors." *Computers and People* **25**(10): 7.

Rasmussen, J. (1983). "Skills, Rules, and Knowledge: Signals, Signs and Symbols, and other Distinctions in Human Performance Models." *IEEE Transaction on Systems, Man, and Cybernetics* **SMC-13**: 257-266.

Rees, L.P., Huang, P.Y. and Taylor III, B.W. (1989). "A Comparative Analysis of an MRP Lot-for-Lot System and a Kanban System for a Multistage Production System." *International Journal of Production Research* **27**(8): 1427-1443.

Resnick, M. (1997). *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds (Complex Adaptive Systems)*. Cambridge, MA, MIT Press.

Rice, J.W. and Yoshikawa, T. (1982). "A Comparison of Kanban and MRP for the Control of Repetitive Manufacturing Systems." *Production and Inventory Management* **23**(1): 1-13.

Ritzman, L.P. and Krajewski, L.J. (1983). "Comparison of Material Requirements Planning and Reorder Point Systems." *Simulation and Inventory Control*. Bekrioglu, H. (Ed.), Society for Computer Simulation.

Roethlisberger, F.J. and Dickinson, W.J. (1939). *Management and the Worker*. Cambridge, MA, Harvard University Press.

Rosenblueth, A., Wiener, N. and Bigelow, J. (1943). "Behavior, Teleology, and Purpose." *Philosophy of Science* **10**: 18-24.

Safizadeh, M.H. and Raafat, F. (1986). "Formal/Informal Systems and MRP Implementation." *Production and Inventory Management* **27**(1): 115-120.

Sarker, B.R. and Fitzsimmons, J.A. (1989). "The Performance of Push and Pull Systems: a Simulation and Comparative Study." *International Journal of Production Research* **27**(10): 1715-1731.

Scgragenheim, E. and Ronen, B. (1990). "Drum-Buffer-Rope Shop Floor Control." *Production and Inventory Management Journal* **31**(3): 18-22.

Schmitt, T.G., Klastorin, T. and Shtub, A. (1985). "Production Classification System: Concepts, Models and Strategies." *International Journal of Production Research* **23**(3): 563-578.

Schonberger, R.J. (1980). "The ROP Shortage List System." *Production and Inventory Management Journal* **21**(3): 106.

Schonberger, R.J. (1982). *Japanese Manufacturing Techniques: Nine Hidden Lessons in Simplicity.* New York., The Free Press.

Schonberger, R.J. (1983). "Selecting the Right Manufacturing Inventory System: Western and Japanese Approaches." *Production and Inventory Management* **24**(2): 33-44.

Schonberger, R.J. (1986). *World Class Manufacturing*. New York, The Free Press.

Schoppers, M.J. (1987). "Universal Plans for Reactive Robots in Unpredictable Environments." *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan Italy: 1039-1046.

Schuring, R.W. (1996). "Operational Autonomy Explains the Value of Group Work in both Lean and Reflective Production." *International Journal of Operations and Production Management* **16**(2): 171-182.

Shanks, D., Rouse, A. and Arnott, D. (1993). "A Review of Approaches to Research and Scholarship in Information Systems." *Proceedings of The Fourth Australian Conference on Information Systems*, Brisbane Australia, Department of Commerce, The University of Queensland.

Shannon, C.E. (1948). "A Mathematical Theory of Communication." *Bell Systems Technical Journal* **27**(3): 379-423.

Shea, J.B. and Paul, G. (1996). "Capturing Expertise in Sports." *The Road to Excellence: The Road to Expert Performance in the Arts and Sciences, Sports, and Games*. Ericsson, K.A. (Ed.). Mahwah, NJ, Lawrence Erlbaum: 321-335.

Shingo, S. (1989). *A Study of the Toyota Production System From an Industrial Engineering Viewpoint*. Cambridge Massachusetts, Productivity Press.

Shoham, Y. (1993). "Agent Oriented Programming." *Artificial Intelligence* **60**(1): 51-92.

Siha, S.M. and David, H.T. (1990). "Pull and Push Low-Volume Production Systems." *People and Product Management in Manufacturing*. Edosomwan, J.A. (Ed.). Amsterdam, Elsevia: 99-119.

Silverman, D. (1970). *The Theory of Organisation*. London, Heinemann.

Simon, H.A. (1969). *The Sciences of the Artificial*. Cambridge, MA, MIT Press.

Simon, H.A. (1973). "Skill in Chess." *American Scientist* **61**: 394-403.

Singh, N. and Brar, J.K. (1992). "Modelling and Analysis of Just-In-Time Manufacturing Systems: A Review." *International Journal of Operations and Production Management* **12**(2): 3-14.

Slack, N. and Correa, H. (1992). "The Flexibilies of Push and Pull." *International Journal of Operations and Production Management* **12**(4): 82-92.

Slezac, P. (1992). "Situated Cognition: Minds in Machines or Friendly Photocopiers?" P*roceedings of The McDonnell Foundation Conference on the Science of Cognition.*, Santa Fe, NM.

Sohal, A. and Chan, J.S. (1991). "Implementation of MRP II in a Large Australian Company." *Proceedings of Proceeding of the 1st International Meeting of the Decision Science Institute.*, Brussels, Belgium: 181-183.

Sohal, A.N., Keller, A.Z. and Fouad, R.H. (1989). "A Review of Literature Relating to JIT." *International Journal of Operations and Production Management* **9**(3): 15-25.

Stalk, G. and Hout, T.M. (1990). *Competing Against Time*. New York, NY., Free Press.

Striekwold, P.E.T. (1990). "A Qualitative Comparison of Production and Inventory Control Concepts." *International Journal of Production Research* **28**(11): 1921-1937.

Suchman, L. (1993). "Response to Vera and Simon's Situated Action: A Symbolic Interpretation." *Cognitive Science* **17**(1): 71-75.

Suchman, L.A. (1987). *Plans and Situated Actions*. Cambridge, Cambridge University Press.

Sugimori, Y., Kusunoki, K., Cho, F. and Uchikawa, S. (1977). "Toyota Production System and Kanban System: Materialisation of Just-In-Time and Respect-For-Human System." *International Journal of Production Research* **15**(6): 553-564.

Sum, C.C. and Yang, K.K. (1993). "A Study on Manufacturing Resource Planning (MRP II) Practices in Singapore." *OMEGA International Journal of Management Science* **21**(2): 187-197.

Taal, M. and Wortmann, J.C. (1997). "Integrating MRP and Finite Capacity Planning." *Production Planning and Control* **8**(3): 245-254.

Trist, E. (1980). "A Socio-Technocal Critique of Scientific Management." *Organizationa as Systems*. Lockett, M. and Spear, R. (Eds.). Milton Keynes, Open University Press: 56-65.

Trist, E.L. (1982). "The Evolution of Sociotechnical Systems as a Conceptual Framework and as an Action Research Program." *Perspectives on Organization Design and Behaviour*. Van deVen, A.H. and Joyce, W.F. (Eds.). New York, John Wiley: 19-75.

Van Der Linden, P.M.J. and Grunwald, H. (1980). "On the Choice of Production Control System." *International Journal of Production Research* **18**(2): 273-279.

Varela, F. (1984). "Two Principles of Self-Organisation." *Self-Organisation and Management of Social Systems*. Ulrich, H. and Probst, G.J.B. (Eds.). Berlin, Springer-Verlag: 25-32.

Varela, F.J., Thompson, E. and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA, The MIT Press.

Vera, A.H. and Simon, H.A. (1993). "Situated Action: A Symbolic Interpretation." *Cognitive Science* **17**(1): 7-48.

von Bertalanffy, L. (1950). "The Theory of Open Systems in Physics and Biology." *Science* **3**: 23-29.

von Bertalanffy, L. (1968). *General Systems Theory: Foundations, Development, Applications*. New York, Braziller.

von Foerster, H. (1984). "On Constructing a Reality." *The Invented Reality*. Watzlawick, P. (Ed.). New York, W. W. Norton & Company: 41-61.

von Foerster, H. (1986). Seeing with the Third Eye: A Night and a Day with Heinz Von Foerster (Video of Lecture), 25 - 26/6/86. Feltonchrist school, California.

von Glasersfeld, E. (1984). "An Introduction to Radical Constructivism." *The Invented Reality*. Watzlawick, P. (Ed.). New York, W. W. Norton & Company: 17-40.

Voss, C. and Clutterbuck, D. (1989). *Just-In-Time: A Global Status Report*. Kempston, Bedford, UK, IFS Publications.

Voss, C.A. (1986). "Implementing Manufacturing Technology: A Manufacturing Strategy Approach." *International Journal of Operations and Production Management* **6**(4): 17-26.

Vygotsky, L.S. (1978). "Mind in Society: The Development of Higher Psychological Processes. Cole, M.et al (Eds.). Cambridge, MA, Harvard University Press.

Wacker, J. and Hill, F.S. (1977). "The Key to Success or Failure of MRP: Overcoming Human Resistance." *Production and Inventory Management* **18**(4): 7-16.

Walbank, M. (1988). "Information Flows - The Foundation of Manufacturing and Stock Management." *Engineering Costs and Production Economics* **15**: 245-260.

Wallace, F.T. (1985). *MRP II: Making it Happen. The Implementers Guide to Success with MRP.* Essex Junction, VT., Oliver Wight Publications.

Warnecke, H.J. and Huser, M. (1995). "Lean production." *International Journal of Production Economics* **41**(1-3): 37-43.

Warnecke, H.-J. and Huser, M. (1993). *The Fractal Factory: Revolution in Corporate Culture*. New York, Springer Verlag.

Weaver, W.G. (1998). "Corporations as Intentional Systems." *Journal of Business Ethics* **17**(1): 87-97.

Webster, J. (1991). "Advanced Manufacturing Technologies: Work Organisation and Social Relations Crystalised." *A Sociology of Monsters*. Law, J. (Ed.). London, Routledge: 192-222.

Wheatly, M. (1989). "OPTimising Production's Potential." *International Journal of Operations and Production Management* **9**(2): 38-44.

White, C. (1985). "Production Scheduling with Applications in the Printing Industry." *European Journal of Operations Research* **22**(3): 304-309.

Wiener, N. (1961). *Cybernetics, or Control and Communication in Human and Machine*. Cambridge, MA, MIT Press.

Wight, O.W. (1981). *Manufacturing Resource Planning: MRP II*. Essex Junction, VT., Oliver Wight Publications, Inc.

Wildavsky, A. (1975). *Budgeting: A Comparative Theory of the Budgetary Process*. Boston, Little, Brown & Co.

Wilkins, D. (1988). *Practical Planning: Extending the Classical AI Planning Paradigm*. Los Altos, CA., Morgan Kaufman.

Wilson, R.H. (1934). "A Scientific Routine for Stock Control." *Harvard Business Review* **13**: 116-128.

Winograd, T. and Flores, F. (1986). *Understanding Computers and Cognition*. New Jersey, Ablex Publishing Corporation.

Wittgenstein, L. (1972). *Philosophical Investigations*. Oxford, Blackwell.

Womack, J.P., Jones, D.T. and Roos, D. (1990). *The Machine That Changed The World*. New York, Rawson Associates.

Woolridge, M.J. and Jennings, N.R. (1995a). "Agent Theories, Architectures, and Languages: A Survey." *Intelligent Agents*. Wooldridge, M.J. and Jennings, N.R. (Eds.). Berlin, Springer-Verlag: 1-39.

Woolridge, M.J. and Jennings, N.R., Eds. (1995b). *Intelligent Agents*. Berlin, Springer-Verlag.

Wren, D.A. (1994). *The Evolution of Management Thought*. New York, John Wiley & Sons.

Yeh, C.H. (1997). "Schedule Based Production." *International Journal of Production Economics* **51**(3): 235-242.

Yin, R.K. (1989). *Case Study Research: Design and Methods*. Thousand Oaks, CA, Sage Publications, Inc.

Yin, R.K. (1993). *Applications of Case Study Research*. Newbury Park, CA, Sage Publications, Inc.

Yourdon, E. and Constantine, L.L. (1979). *Structured Design*. New Jersey, Prentice-Hall.

Zuboff, S. (1988). *In the Age of the Smart Machine: The Future of Work and Power.* New York, Basic Books.

Appendix 1

Relevant Publications by the Author

Publications on Which this Thesis is Based

Paper one

Johnston, R.B. (1995). "Making Manufacturing Practices Tacit: A Case Study of Computer Aided Production Management and Lean Production." *Journal of the Operational Research Society* **46**(10): 1174 - 1183.

Paper two

Johnston, R.B. and Brennan, M. (1996). "Planning or Organising: The Significance of Theories of Activity for the Management of Operations." *OMEGA International Journal of Management Science* **24**(4): 367-384.

Paper three

Johnston, R.B. (1996). "Thinking about Acting and Acting are Different - Reply to the Comment on my Paper." *Journal of the Operational Research Society* **47**(8): 1080-1082.

Publications Which Extend the Work in this Thesis

Johnston, R.B. and Betts, J.M. (1996). "A Complex Systems Simulation of Planning versus Reacting for Inventory Replenishment." In *Proceedings of Complex 96: From Local Interactions to Global Phenomena*, Charles Sturt University, Albury, NSW, IOS Press, Amsterdam: 339 - 350.

Johnston, R.B. (1996). "From Efficiency to Flexibility: Entropic Measures of Market Complexity and Production Flexibility." In *Proceedings of Complex 96: From Local Interactions to Global Phenomena*, Charles Sturt University, Albury, NSW, IOS Press, Amsterdam: 130 - 140.

Johnston, R.B. (1997). "The Trade-off Between Efficiency and Flexibility in Operations Management." In *Proceedings of The Fourth Conference of the Association Asian-Pacific Operational Research Societies*, Melbourne: WA2.1.

Johnston, R.B. (1997). "Reconciling the Economic Batch Size Model with Current Manufacturing Rhetoric." In *Proceedings of The Fourth Conference of the Association Asian-Pacific Operational Research Societies*, Melbourne: MB6.4.

Johnston, R.B. and Betts, J.M. (1998). "Evaluating Planning vs. Reactive Trading Strategies using a Complex Systems Simulation." *International Symposia in Economic Theory and Econometrics, Vol 11*. Barnett, W. (Ed.). Cambridge, Cambridge University Press: (Forthcoming).