Knowledge Based Improvement:

Simulation and Artificial Intelligence for understanding and improving decision making in an operations system

By

Thanos Alifantis BSc (Lancaster)

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Industrial and Business Studies

Operational Research and Information Systems Group
University of Warwick, Warwick Business School

August 2006

Contents

Figures and Tables

PAGE
MISSING
IN
ORIGINAL

Acknowledgements

I am indebted to a number of people who have offered help and encouragement during the course of this research.

Firstly, I would like to thank my supervisor Professor Stewart Robinson for his initial idea, his time, his encouragement and for offering me the opportunity to work as a research student in the EPSRC project GR/M72876.

Secondly, I would like to thank my second supervisor Dr Robert Hurrion for his comments during the early stages of this work and for inspiring me by being the originator of the Visual Interactive Simulation.

Thirdly, I would like to thank John Edwards (Aston University), John Ladbrook (Ford), Tony Waller (Lanner Group) and Peter Walton for their useful comments and for providing the information and the support that contributed to the validation of this research.

Also thanks go to the Group Leaders and production managers in the Ford engine plant at Bridgend for their involvement in this research and for contributing to its validation.

The research would not have been possible without financial support. The Engineering and Physical Research Council, Ford and the Lanner group provided the funding.

Finally, on a personal note, I would like to express my appreciation to my parents and to all my colleagues and consultants at AEA Technology Rail in London for their encouragement.

Abstract

The thesis investigates the possibility of using simulation for understanding and improving the design of decision making in a real context. The approach is based on the problem of representing decision making behaviour in Discrete Event Simulation.

An investigation of existing techniques led to the design of a methodology known as Knowledge Based Improvement (KBI). The KBI covers the key stages of the process of using simulation for understanding and improving the design of decision making. Using a research strategy that involves a case study in Ford, the research tests each stage of KBI.

The thesis explains how simulation can be used for understanding real decision making problems and for collecting the data required for modelling individual decision making strategies. The thesis demonstrates the possibility of a simulation based knowledge elicitation in a real context and it investigates the practical difficulties involved in this process.

The research tests the process of understanding decision making policies by modelling specific decision makers using Artificial Intelligence. It tests the use of simulation for assessing the decision making strategies and it shows that simulation can be used for identifying efficient strategies and for improving the design of decision making practices.

The thesis reports the degree of success of the approach in relation to the data that were collected and it describes the validation checks that were undertaken. In addition, it reports the lessons learned from the application of the KBI methodology, the overall success of the approach and the main limitations that were identified during the implementation.

Key Words: Simulation, Artificial Intelligence, Knowledge Elicitation, Decision Making, Manufacturing, Expert Systems, Decision Trees, ID3, Data Mining, Scheduling

Abbreviations

KBI	Knowledge Based Improvement
OR	Operational Research
Simulation	Discrete Event Simulation
VIS	Visual Interactive Simulation
VIM	Visual Interactive Modelling
AI	Artificial Intelligence
WES	Warwick Expert Simulator
ID3	Iterative Dichotomiser 3
UI	User interface
DM	Decision Maker
ERT	Estimated Repair Time

Chapter 1

Introduction

1.1 The foundations of simulation modelling

Discrete Event Simulation (Simulation) is one of the most commonly applied techniques in Operational Research (OR) (Jeffrey & Seaton 1995, Fildes & Ranyard 1997, Clark 1999, Robinson 2005). Simulation involves the development and analysis of computer based models that simulate systems which exist or are going to exist in the future. Amongst other things, those models can be used for cost effective and safe experimentation and as a tool for understanding the behaviour of real systems.

Simulation started to emerge in the early fifties, but it was formally introduced by Tocher (1963). Tocher's seminal book set the foundations which underlined the rationality of using simulation for solving real world problems. In his book, Tocher assumed that systems are collections of subsystems that interact with each other. This interaction generates an emerging dynamic behaviour which cannot be captured with mathematical representations. In an effort to predict the effect of a change in the configuration of the system, Tocher (1963) proposed to model systems by developing computer based simulations.

Almost 40 years after Tocher's original contribution, the use of simulation modelling has been expanded and now accounts for a major percentage of the OR applications (Fildes & Ranyard 1997). A significant contribution towards making Simulation a mainstream approach and for expanding its scope was Hurrion's (1976) development of Visual Interactive Simulation (VIS). The development of VIS made the use of simulation relevant for a wide variety of applications. Hurrion's contribution to the simulation community is summarised by Bell & O'Keefe (1994) who comment:

"Bob Hurrion at the University of Warwick in England first applied

> animation and run time interaction with a simulation model of an operational system in his 1976 PhD thesis and reported this in subsequent papers"

Hurrion (1978) & Hurrion (1980) report applications of running simulations of operating rules for a job shop scheduling problem. In these papers, which form the first journal-based published research in VIS, it was found that the human scheduler was using rules that were difficult to embed in the simulation model. To resolve this problem an interactive simulation model was developed and the control was passed to the human scheduler at various points in time. To provide the necessary information about the status of the system which the scheduler required in order to make a scheduling decision, an iconic visual display was used.

With this problem-solving approach, Hurrion created a new paradigm. The VIS paradigm which, soon after the appearance of the above publications, was generalised to produce a package for programming VIS called Vision, led to the development of See Why (Fiddy et al 1981) and made VIS commercially available. This was the beginning of a new era in simulation. VIS software became a simulation standard and attention in simulation research and practice was diverted towards a different methodological approach that involved the use of the new tool. In subsequent years, Hurrion's VIS approach became so popular that now, apart from its application in quantitative modelling, it is also used as a tool for interpretative modelling. The reason for this, as explained by Bell and O'Keefe (1994), is that it can be used for problem-structuring purposes since the visual model can act as problem-definer tool, a vehicle for the decision maker and the modeller to learn from the problem together.

1.2 Motivation and aim of the research

Today, with Hurrion's ideas to underline and shape the simulation modelling practice, simulation packages are available and are used for modelling virtually any type of operations systems. As described by Robinson (2004), an operations system is a configuration of resources combined for the provision of goods and services. Operations systems usually involve a significant element of human - machine interaction (Williams 1996) where human decision makers are involved in decisions that have a significant impact on the performance of the systems. The operations of manufacturing plants include production and maintenance scheduling that may involve human decision making (Proudlove et al 1998). In the service sector, service operations systems may include personnel who may be involved in a decision making process as they interact with the rest of the system. For example, in transportation systems such as railway networks and airports where simulation modelling is extensively applied (Robinson & Stanger 1998), decision makers such as signallers and air traffic controllers are usually involved in a number of scheduling activities which very often require decision making.

Although the modelling and representation of the human decision making process is essential for understanding and improving decision making using simulation, the latter usually overlooks the effect of human decision making or it adopts a very simplistic approach towards modelling it. This is probably because modelling human decision making in simulation presents a number of challenges. Most simulation tools do not provide a full set of functions which would be useful for modelling human decision making (Williams 1996). A more fundamental challenge is that of representing human decision makers with individual characteristics and behaviour. This would require an understanding of what the decision making process involves and how the decision makers reach decisions. In addition, if the purpose of the simulation exercise is to identify and assess good decision making practices, then this would require a methodology for assessing decision making strategies using simulation.

The extensive use and potential of simulation, combined with developments in related disciplines such as computer science and artificial intelligence (Doukidis 1987), have motivated a number of authors, including Flitman & Hurrion (1987), to propose approaches which attempt to identify and model complex entities with ill-defined behaviour representative of human decision makers. However, as will be explained in the next chapter, through a review of the literature, most of those approaches have not been tested or applied in industrial applications and, as standard practice in most of the applications in which human decision making entities are involved, the decision making strategies are either excluded or modelled simplistically. Motivated by the absence of substantive empirical evidence to justify and demonstrate the feasibility, the difficulties and the benefits of modelling human decision making in simulation, the aim of the research (section 3.1) is:

> *To develop and test the use of simulation for understanding and improving the design of decision making policies in a real context.*

1.3 Overview and outline of the research

The aim of the research (section 3.1) is derived through a review of the literature (Chapter 2). In order to achieve the research aim a number of objectives are identified. Based on these objectives the research forms and evaluates a methodology for capturing efficient decision making.

The research is applied and tested on an engine assembly line at a Ford Motor Company (Ford) plant in Wales. In the engine assembly, engines are passed through a series of automated and manual processes. From time to time the automated machines break down and require repair. It is the decisions surrounding what happens when a machine fails that are the focus of the research.

An existing simulation model of the engine assembly line is adopted to form a prototype that is used throughout the research. Originally developed for production planning and bottleneck detection, the existing model represents machine breakdowns and maintenance that involves immediate machine repairs. However, it does not reflect the decision making process that takes place in the real system.

Taking into account that the decisions surrounding what happens when a machine fails have a significant impact on production, the research develops and tests the use of simulation for understanding and improving the production decisions.

Chapter 2 describes the methodologies and the techniques that have been used in the past for modelling decision making in simulation. It explains how the current research is related to previous work and it concludes by identifying a number of potential research areas. Chapter 3 describes the research questions and the research design. It justifies the methodological approach and provides an overview of its limitations. Based on theoretical developments by previous authors, Chapter 4 forms a methodology for capturing efficient decision making using simulation. Chapters 5, 6, 7 and 8 test and evaluate the stages of this methodology. Chapter 9 discusses the findings of this research and concludes with the main lessons that were learned. Finally, in Chapter 10 the main conclusions and the contribution of the research are summarised and areas where further research is required are identified.

Chapter 2

Current approaches for modelling decision making in simulation

Introduced in Chapter 1 the motive of the research is to develop and test the use of simulation for understanding and improving decision making. To provide an overview of the relevant techniques and to formulate the research problem this chapter will address the following questions:

- What process is involved in the use of simulation for understanding and improving decision making and what general stages are involved in this process?
- What techniques have been used or proposed for each of the above stages and what techniques have been considered for related applications?
- What methodologies have been developed for the process involved in the use of simulation for understanding and improving decision making and in what context they have been applied?

The answers to these questions form the basis of the research aim, the research questions and the research strategy developed in Chapter 3. They also contribute to the development of the conceptual methodology discussed in Chapter 4.

2.1 The process involved in the use of simulation for improving decision making: an overview of the challenge

The use of simulation for understanding and improving decision making involves the modelling of decision making in simulation (Curram 1997, Williams 1996, Liang et al 1992, Robinson et al 1998, Flitman & Hurrion 1987, Mason & Moffat 2000). Introduced in Chapter 1 (section 1.2), the modelling of decision making is one of the main challenges in simulation. The challenge consists of understanding what the decision making process involves and how the decision makers reach decisions (Williams 1996). As has been shown by previous authors (Williams 1996, Liang et al 1992, Robinson et al 1998, Flitman & Hurrion 1987, Mason & Moffat 2000) the problem can be described as a process that involves three generic stages. The first stage, known as knowledge elicitation, includes all those tasks required for gathering the information and the data that are required for modelling decision making (Williams 1996, Robinson et al 1998 & 2003). In addition to the problem structuring process, this stage might involve data collection aimed at facilitating the development and calibration of decision making models (Checkland 1981, Liang et al 1992). The second stage, termed knowledge modelling, involves the development of quantitative or qualitative models for representing the decision making process (Curram 1997, Perry & Moffat 1997, Mason & Moffat 2000). By developing models of decision making, this stage focuses on the identification of the decision making strategies required to represent decision making. The third stage, known as knowledge representation, involves the process of linking the decision making models with the other entities of the simulation (Flitman & Hurrion 1987, O'Keefe 1986, Mason & Moffat 2000, Robinson et al 2003b). This is the stage where the simulation can be used in order to assess the impact of representing decision making (Flitman & Hurrion 1987).

In the past 25 years, a number of different approaches have been proposed for implementing each of the three stages of the process of modelling decision making in simulation (Hurrion 1978, Flitman & Hurrion 1987, Curram 1997, Perry & Moffat 1997, Mason & Moffat 2000). Among the elements which differentiate the approaches are the techniques which have been employed for each of the three stages above (Curram 1997). The approaches also differ on how they use the techniques, on what they try to achieve, on what they define as decision making and on the environment and the knowledge that they use to develop the models (Perry & Moffat 1997). Given this diversity, in order to identify the techniques which are available for modelling decision making in simulation, sections 2.2, 2.3 and 2.4 of this chapter will describe knowledge elicitation, knowledge modelling and knowledge representation techniques. The knowledge elicitation section will cover techniques that have been used as part of the process of eliciting knowledge for modelling human decision making in simulation. To provide a complete overview of the relevant data collection techniques, the knowledge elicitation section will also consider techniques which have been used to collect data to represent human elements in simulation as well as techniques that have been used for eliciting knowledge required for representing optimal decision making in simulation. The knowledge modelling section will cover AI and qualitative techniques which have been used or proposed as part of the process of modelling decision making in simulation. Finally, the knowledge representation section will cover techniques that have been used for implementing and linking decision making models with simulation.

Having reviewed the literature in terms of the techniques that are employed for each stage in section five, five characteristics which vary significantly across the proposed approaches are identified and are used in order to organise the research into categories with common characteristics. Based on these categorisations, the types of methodologies that are available for modelling decision making in simulation are identified. Section seven discusses potential research areas and section eight concludes by summarising the current research issues associated with the process of modelling decision making in simulation.

2.2 Knowledge elicitation techniques

Given that VIS (defined in Chapter 1) appears to be one of the most commonly applied quantitative knowledge elicitation techniques in the literature in modelling decision making (Liang et al 1992, Flitman & Hurrion 1987, Perry & Moffat 1997), the first part of this section describes various quantitative, VIS-based knowledge elicitation approaches. In the second part of this section, alternative knowledge elicitation techniques are described.

2.2.1 VIS for knowledge elicitation

VIS is a special case of Visual Interactive Modelling (VIM - Hurrion 1986). VIM is a well-established knowledge elicitation technique in the literature on modelling decision making in simulation (Hurrion 1986). The capabilities of VIM to establish a dialogue between the human expert and the computer system and the fact that it is usually combined with simulation models are two of the main reasons for the attention that it has received in the literature (Hurrion 1986).

Hurrion (1986) defines a VIM methodology as one where an OR analyst builds a model of some problem situation. The model has a graphics component, so that the user commissioning the study can observe in a suitable animated form the dynamics of the model. The user is then able to use his/her own expert knowledge and judgment of the original problem domain in order to interact with the model. In addition, Hurrion (1991) comments that one of the major modelling techniques used within the generic VIM framework is simulation.

VIS for knowledge elicitation involves the development of a simulation model where, initially, the representation of decision making rules requires the involvement of the user (Bell & O'Keefe 1994). This means that during the simulation when a decision point is reached, the simulation stops and the user, based on information about the status of the system reported through the visual display, must decide which of the decision options is the most appropriate. In some of the applications that have been reported (Liang et al 1992, Flitman & Hurrion 1987, Perry & Moffat 1997), the functionality of the computer model not only allows the user to input decisions which are taken into account in the next events of the simulation but also automatically creates a record of the decision and the attributes in a database.

Based on the above generic approach, a number of specific techniques have been developed and tested using full scale applications or simplified case studies that represent typical problems (such as the general job shop scheduling problem). According to O'Keefe & Pitt (1991), Perry & Moffat (1997), Hurrion (1986), Liang et al (1992) and Curram (1997), some of the key elements which differentiate the various VIS (for quantitative knowledge elicitation) approaches are the following:

- The type of visual display that is used.
- The number of decision situations that are presented to the user.
- The method that is used for generating the decision situations.

- The type of decision makers who are involved in the knowledge elicitation process and the level of their experience.

As is shown in table 2.1, based on these characteristics it is possible to identify at least three different VIS-based knowledge elicitation techniques (Flitman & Hurrion 1987, Liang et al 1992, Perry & Moffat 1997). Approaches I and II have been developed with the involvement of non-industrial experts (Curram 1997, Liang et al 1992, Flitman & Hurrion 1987). The main difference between the two lies in the type of visual display that is used. In the first approach, the representation of the system has the form of a diagram which represents the layout of the system (Curram 1997, Flitman & Hurrion 1987). In the second approach, the system attributes are visualised using logical diagrams such as charts and figures (Liang et al 1992). The third approach (Perry & Moffat 1997), shown in table 2.1, is quite different from the previous two. Apart from the involvement of real experts, a key difference between this approach and the previous two is the type of interaction. Unlike approaches I and II, in the third approach each time that a decision is required, besides the standard facilities that are provided to the user in order to input a decision, there is some extra functionality which allows the decision maker to request extra information, which is available only upon request. This creates a dynamic dialogue, which provides a more realistic simulation environment since it represents the fact that, in the real world, the user might need to find additional information before he makes a decision. Having identified three different variants of VIS for knowledge elicitation, the following paragraphs provide details about applications in which each VIS approach is employed.

Approach	Authors	Visual display	Type of interaction	Number of decision situations	Decision situations generation process	Type of problem	Decision makers' level of experience
I	Flitman & Hurrion (1987), Curram (1997)	Schematic	Static/ Passive	High	Sequential	Simplified example	Non-industrial
II	Liang et al (1992)	Logical	Static/ Passive	High	Sequential	Manufacturing example	Non-industrial (students)
III	Perry & Moffat (1997), Williams (1996) and Williams et al (1989)	Schematic	Dynamic	High	Randomised	Military (Real)	Experienced Commanders

Table 2.1: *Types of VIS for knowledge elicitation.*

Approach I: Visual schematic run time interactive simulation model with incomplete sequential dialogue and multi-participation of non-industrial decision makers

O'Keefe & Pitt (1991) and Flitman & Hurrion (1987) describe a visual interactive methodology for knowledge elicitation that is based on the use of a VIS. The simulations that model the operations of a simplified example of a coal depot (the same typical problem is analysed in both studies), are used to collect data on resource allocation decisions which depend on queue lengths. The models have been developed in Pascal and Prolog using the libraries Pascal_SIM (Davies & O'Keefe 1988) and WES (Warwick Expert Simulator, Hurrion 1991). During the data collection process (Flitman & Hurrion 1987) the model simulates the system, including the length of each queue. The appropriate facilities are provided to enable the human decision maker to interact at run time with the model and to alter the allocation of resource levels, taking into account the queue lengths that are simulated and dynamically displayed in schematic form in the computer model. Each time that a resource allocation decision is made, the decision and the corresponding attributes are recorded. The decision situations that are presented to the decision makers are generated sequentially (subsequent decision points: non-random sampling) from a short simulation run. As a consequence, the data set is not exhaustive and some decision situations are not presented to the decision makers, resulting in knowledge gaps. According to Hurrion's (1986) classification of interactive models, this dialogue is characterised as passive or incomplete due to the limited interaction with the model. In this knowledge elicitation experiment (Flitman & Hurrion 1987), five non-industrial decision makers participate during the knowledge elicitation process. In order to enable direct comparison of their efficiency, each of them is exposed to the same decision situations. Clearly, from the type of interaction and the type of visual display that is used, this work is an example of VIS for knowledge elicitation which represents Approach I (table 2.1).

A similar knowledge elicitation approach is also described by Curram (1997). In his research, a simulation that models the operations of a bank branch is used to collect data about 'queue joining' decisions, which depend on queue lengths and other characteristics of the system. Based on a pilot data collection that was conducted to investigate the range of possible situations and to address practical issues of the data collection process, a knowledge elicitation experiment was designed. During the design, the situations that were chosen to be presented to the decision maker were carefully specified to allow a structure to be imposed on the data collection. The aim was to achieve a good spread of situations given the limited number of scenarios. In view of the simplicity of the case study, no provision was needed to ensure that the scenarios generated did not fall into the same decision category and that extreme situations were presented. By simply applying common sense, it was easy to predict the extreme cases of a queue length in a bank branch and to develop a distribution of the frequency of each queue length. In order to collect the data a visual schematic display was used with passive interaction, where the user has to input a decision but is unable to interact dynamically with the system or request additional information about the status of the system.

Approach II: Visual logical run time interactive simulation model with incomplete sequential dialogue and multi-participation of non-industrial experts

Liang et al's (1992) work, which according to Bell and O'Keefe (1994) is one the most extensive investigations of intelligent VIS, describes a knowledge elicitation approach that is based on the use of a VIS model of a manufacturing facility. The model simulates a manufacturing process where materials have to be submerged in chemicals for specific times but go bad if left in a tank for too long. There is only one resource that may be used to move jobs from one tank of chemical to the next and the decision maker has to decide on the order in which material should be moved from one tank to the others. The aim of the scheduling decision making problem is to avoid under- or over-processing of each manufacturing item since this would spoil it. The knowledge elicitation process involves the use of that VIS which simulates the operations of the system and reports the value of various attributes of the system, using a bar chart display which is updated each time that the simulation reaches a decision point. When a decision is required, the model stops and the intervention of the user is required. Once a decision is taken it is recorded in a database for future reference. This knowledge elicitation application reports that a significant number of decisions (a total of 2125 decisions, 215 per person) can be collected using the VIS model but, as the authors comment, it raises issues associated with the fact that non-industrial experts were involved. It also raises issues associated with the fact that the decision situations that were selected might not be a representative sample, since no experimental design was undertaken. According to Hurrion's (1986) classification of visualisation, the visualisation approach that is used by Liang et al (1992) can be classified as logical, since charts and diagrams are used for displaying the information which the decision maker takes into account in order to make a decision. In addition, the type interactive dialogue can be classified as passive, since the decision makers could not define new decision actions or request information about additional attributes that were not reported in the logical display.

Approach III: Visual schematic run time interactive simulation with incomplete stratified scenarios and multi-participation of military decision makers.

Perry & Moffat (1997), Williams (1996) and Williams et al (1989) are perhaps the only full scale applications of VIS-based knowledge elicitation.

Williams (1996) and Williams et al (1989) describe a knowledge elicitation approach in which a VIS is used for collecting examples of scheduling decisions that involve prioritisation of replenishment. The author comments that, in their third VIS experiment, pre-prepared situations were presented to experienced commanders. Although specific details about the preparation of the decision situations are not given, it is clear that a stratified approach was employed in order to ensure that a representative data set was collected. In the methodology applied, apart from the decisions that were recorded automatically in the system, additional information was collected manually during the data collection. The situations were presented to the decision makers using electronic maps and therefore a schematic VIS approach is employed (Hurrion 1986). In addition, the fact that the decision situations are pre-prepared shows that, in this simulation-based data collection, the decision situations are not generated sequentially from the simulation and so a non-sequential dialogue is applied.

Perry & Moffat (1997) describe a VIS approach for collecting decision data about military decisions which depend on various attributes, formed from information that the real decision makers were able to request from the interactive model each time that a decision was required during the simulation run. The decision makers who participated in the study commissioned by Perry & Moffat (1997) were very senior acting or recently retired naval officers. The model is an interactive maritime stochastic event driven simulation at task force level, which can be interrupted by the player to probe for information or alter the simulation by inserting orders to the force. The decision situations and outcomes presented to the participants were generated by charting a reasonable course of events based on earlier decisions made and simulation adjudications. These events led to a single critical decision situation that was then presented to the user. Based on Hurrion's (1986) classification of interaction, the type of interaction used by Perry & Moffat (1997) is classified as active or dynamic, since the interactive dialogues which are used enable the user to retrieve information upon request and to input non-standard decisions that were not initially in the list of the anticipated decision options (this was achieved by proposing new emerging decision options that were later hard coded in the system). Moreover, the fact that additional information was available upon request supports the identification of emerging attributes, when initially it was not expected that they might support the decision making process. To make the experiment even more robust, in each of the ten experimental sessions that were conducted, a wide range of situations was presented to the decision makers with the aim of avoiding stereotyped responses. The decision option 'the situation is totally infeasible and there is no reasonable way forward' was also available to enable the identification of unrealistic scenarios. As the authors comment, the type of interaction that was employed in their research minimises the risk of producing a behavioural situation which is so constraining that the data set that is collected is purely a product of the experimental method and does not represent the situations that the decision makers actually face.

2.2.2 Alternative knowledge elicitation techniques

Apart from VIS for quantitative knowledge elicitation, a number of alternative knowledge elicitation techniques exist which, under certain circumstances, can be used for the first stage of the process of modelling decision making in simulation (Angelides & Paul 1999, Hurrion 1980, Carvalho et al 1998, Liang et al 1992, Pierreval & Ralambondrainy 1990). Some of those techniques have been proposed by previous authors and are explained in the following paragraphs.

Historical data for quantitative knowledge elicitation

As is explained by Liang et al (1992), historical data of decisions and attributes can be a useful source of knowledge which can be used for modelling human decision making in simulations. It is an efficient knowledge elicitation technique, which does not require the development of knowledge elicitation systems or considerable involvement of human experts. In the methodologies that have been developed for modelling human decision making, historical data are not used as much as might be expected. One of the main reasons is the lack of historical data records, on which Liang et al (1992) comment that "most research is based on the assumption that training data are available and are error free. In some problem domains expert decisions are difficult to obtain or need to be collected on real time basis. When however historical data are available they can be used for developing models of decision making".

Mason & Moffat (2000) provide the most notable example where historical data on decisions are used to develop models. In one of their decision making representations they use a set of mentally stored references which have been accumulated by experience and training to model rapid planning decisions. Carvalho et al (1998) use historical data to model decision making on passenger choice between train or bus. Although their model is not specifically developed to fit in a simulation, it is developed using historical data on decisions. Similarly, Malhorta et al (1999) use historical data to develop decision making models that represent business decisions. As with Carvalho et al (1998), their model is not integrated with a simulation application, but the data set is derived from specifically historical decisions.

Questionnaires for knowledge about people

If the decision situations can be described in a questionnaire (Deslandres & Pierreval 1997), the use of questionnaires is a knowledge elicitation technique that can be used as part of the process of modelling decision making in simulation. Baines et al (2000) use questionnaires to collect data that are used in combination with data-driven modelling techniques to represent in simulation how individual characteristics and environmental conditions affect human performance in manufacturing environments. The questionnaires are designed to collect information about the personality of the individual and about how certain environmental conditions affect the time that they spend doing certain tasks or time that they spend on rework.

Optimal data sets derived from simulation experiments for quantitative knowledge elicitation

When the purpose of the modelling exercise is to represent optimal decision making, a technique that has been used in the past (Pierreval & Ralambondrainy 1990) involves the use of simulation for deriving a set of optimal decisions through simulation experimentation. For example, Pierreval & Ralambondrainy (1990) describe a technique in which simulation is used to generate a data set in which the decision instances are optimal and have been derived from exhaustive simulation optimisation experiments. In addition, Alifantis & Robinson (2002) propose the use of optimal decisions derived from exhaustive experimentation with simulation to develop a scheduling advisor.

Artificially generated data sets for quantitative knowledge elicitation

If the knowledge elicitation stage of the process of modelling decision making in simulation involves investigation about which technique is the most appropriate to be adopted, then artificially generated data sets can be used (Curram 1997, Carvalho et al 1998). Artificially generated data sets are used for testing the capabilities of methodologies and generic mathematical models. Although the data are not real, they behave as if they were collected from real world situations and sometimes they have properties that cannot be easily found in real world knowledge domains and data sets (Carvalho et al 1998). For example, artificially generated data sets are complete and follow predetermined statistical distributions which help to create controlled experiments by controlling assumptions that are required by certain techniques and models. Generally, artificially generated data sets are used as a preliminary stage in the process of investigating the performance and efficiency of various techniques under different conditions. In those cases, knowledge elicitation consists of finding which is the most appropriate technique to be used for modelling a specific situation. Curram (1997) describes a process that involves the use of artificially generated data to investigate the performance of stochastic neural networks models. Carvalho et al (1998) describe a process that involves the use of simulated data to undertake a comparative analysis of two data-driven modelling techniques (logistic regression and neural networks).

Visual interactive simulation and gaming for qualitative knowledge elicitation

Earlier in this section, VIS was described as a technique which can support the process of collecting quantitative data sets of decisions and attributes. However, when the purpose of knowledge elicitation is to collect qualitative information about the decision making process, VIS can also support qualitative knowledge elicitations. Hurrion & Secker (1978), Hurrion (1978) and Hurrion (1980) are the first published research studies (after Hurrion's PhD Thesis 1976) in which VIS is used as an aid to the decision maker. They are also examples of qualitative knowledge elicitation processes in which VIS is not used simply for collecting a data set of decision instances. As with recent investigations in gaming simulation environments (Angelides & Paul 1999), the purpose of those papers is to illustrate that if the user can see how the model is progressing, conditional on the decisions that are made, then he/she is in a knowledgeable position to learn from the progress of the model and from the interactive decisions that he/she makes. In these applications, the elicited knowledge is formed by understanding the responsiveness of the system in the decisions that are made during the interactive run and, therefore, VIS is used as a tool for knowledge elicitation.

Verbal protocol analysis for qualitative knowledge elicitation

A method that is usually applied for qualitative knowledge elicitation and sometimes supports data collection sessions is the observation of the decision maker in his working environment (Deslandres & Pierreval 1997). According to Deslandres & Pierreval (1997) this elicitation method, which is formally known as verbal protocol analysis, is usually applied in the problem-solving knowledge domain and is described as follows: "Verbal protocol analysis is a direct observation method of the expert's problem solving process. It consists of observing and noticing the experts' behaviour during their work, taking notes of everything that is done or said." The method is composed of two phases: syntactic analysis and semantic analysis. Syntactic analysis aims to identify operators, operands and the sequences of operations which can be observed during the verbal protocol of the experts. Semantic analysis involves exploring the reasoning by giving a succession of interpretations from the previous analysis. From the actions proposed by the expert the underlying purposes are identified, then an interpretation of the situation is given and, finally, the knowledge involved in this interpretation is characterised. According to Williams et al (1989), it is a time-consuming technique and the protocols collected do not always cover every situation.

Problem structuring as a qualitative knowledge elicitation technique

When optimisation and decision support tools are used during the decision making process then the knowledge elicitation process focuses on understanding how optimisation techniques can be used for representing the decision making process. For example, Mason & Moffat (2000) model decision making about deliberate planning using game theory. In this model it is assumed that when the commanders are not under pressure they use analytical techniques and not their experience to make rational decisions. In this case, the knowledge elicitation process focuses on understanding how game theory is used as a decision making tool. Emphasis is placed on the process of understanding what decision options are available, what are the enemies' decision options and what is the main objective. Clearly, the knowledge elicitation process is similar to the problem structuring process that would be applied for formulating an optimisation problem.

Structured interviews and knowledge acquisition grids for qualitative knowledge elicitation

Structured interviews supported by knowledge acquisition grids (questionnaires) are an alternative knowledge elicitation technique that is usually applied in qualitative knowledge acquisition. According to Deslandres & Pierreval (1997), in manufacturing applications structured interviews are appropriate for eliciting practical knowledge from foremen and operators, who are usually good at expressing individual considerations on the question of manufacturing constraints. In their research, knowledge acquisition grids are used. The purpose of the questionnaire given to the decision makers was to investigate the quality controls that certain manufacturing products need to pass. Williams et al (1989) comment that interviews as a technique are efficient in terms of time but require a generally accepted expert. In their research, interviews supported by VIS are employed in order to elicit quantitative decision making knowledge about replenishment decisions.

Action research

Action research can also be used for knowledge elicitation. When action research is used the modeller is involved in the decision making process. This enables him to understand what types of data are required and how the data can be collected (Hindle et al 1995).

2.2.3 Discussion: knowledge elicitation techniques

Having described a number of different techniques in the previous paragraphs of this section, the above review has shown that a number of different knowledge elicitation techniques exist and many of them have been used as part of the process of modelling decision making in simulation. In simulation literature VIS appears with great frequency, with a number of authors reporting experiments where VIS is used for quantitative and qualitative knowledge elicitation (Liang et al 1992, Flitman & Hurrion 1987, Perry & Moffat 1997, Curram 1997, Hurrion 1980). From the above review, it is clear that there are some areas related to the knowledge elicitation process that have not been fully investigated in the current literature. The current knowledge elicitation literature is based mainly on non-industrial applications, nor does it address the issues associated with the process of minimising the risk of collecting a data set which may be highly biased and essentially the product of the data collection process (Liang et al 1992, Curram 1997). This conclusion, which follows the general current trend in OR research (Pidd & Lewis 2001), highlights the lack of a detailed methodology for explaining how a representative set of decision situations should be generated from a VIS (Curram 1997). It also indicates that some of the early stages in the elicitation techniques, which involve the preparation of the data collection and perhaps some action research, have not been fully identified and addressed by the current literature (Curram 1997). Finally, it calls attention to the need for applied research and emphasises the importance and the benefits from research based on industrial applications (Curram 1997).

2.3 Knowledge modelling: techniques for modelling decision making using data collected from a simulation

Having reviewed the main approaches that have been proposed for collecting data for modelling decision making, the main quantitative and qualitative modelling techniques that have been used to develop models of decision making for simulations are now described.

2.3.1 Quantitative knowledge modelling techniques

The quantitative knowledge modelling techniques which are described first in the following paragraphs usually require the calibration of a mathematical model (Liang et al 1992, Curram 1997, O'Keefe 1986). The calibration involves optimising the ability of the model to predict the decisions that are included in the data set. Sometimes knowledge modelling techniques, instead of calibrating the parameters of an equation, involve the development of a database (Flitman & Hurrion 1987, Mason and Moffat 2000, Liao 2000). This database is used as a look-up table to predict decisions for a given set of attributes (Mason and Moffat 2000). If the attributes are not in the look-up table, an approximation algorithm is used to find the attributes in the data set which are closest to those in the new decision situation. The following paragraphs review specific quantitative knowledge modelling techniques which have been used to model decision making in simulation.

Pattern matching case-based reasoning

In pattern matching case-based reasoning, a database with historical decisions and attributes is used in order to determine what decision should be taken for a new situation (Flitman & Hurrion 1987, Mason and Moffat 2000, Liao 2000). When a decision is required, a query in the database is executed to find the decision whose attributes match with those in the current situation. Kolodner (1993) describes two possible case-based reasoning approaches. The first, known as pattern matching, involves the use of a database of stored decision and attribute pairs. This database is used to retrieve the most appropriate decision for a current situation. In order to find the most appropriate decision, a search takes place to match the current situation with the situations that are available in the database. Various algorithms have been applied for matching the current situation with those that are in the database (Flitman & Hurrion 1987, Mason and Moffat 2000).

Flitman & Hurrion (1987) and Hurrion (1991) describe a pattern matching implementation which is used to model human decision making in a simulation model. In one of the case studies (Flitman & Hurrion 1987) which they describe, they model decision making in a simulation by developing a database of decision and attributes instances. In the pattern matching implementation described in those papers, a separate database is used for each decision maker. The database is linked with the simulation model and, when a decision is required, a search algorithm is used to find the attribute pattern in the database which matches the attribute pattern in the current state of the simulation model. Once a pattern that is close enough to the one in the simulation is found, the database is queried for the decision that corresponds to the relevant pattern. The decision is then incorporated in the simulation, which continues the run until the next decision. Applying a similar matching approach but with a different pattern identification algorithm, Mason and Moffat (2000) model human decision makers in a military simulation model.

Adaptive case-based reasoning

The second case-based approach not only involves the use of an attributes-decision database but also an adaptation algorithm which is used to extend the knowledge base that is initially included in the database (Kolodner 1993, Liao 2000). This process represents the learning experience which takes place when individuals are involved in a decision making process. This approach is usually known as adaptive case-based reasoning or simply case-based reasoning. It is based on the use of various databases that come from different experts (Liao 2000). When a decision is required, all matched decisions that are in the database are collected to form a set of strategic propositions. As explained by Liao (2000), rules are used to form a set of detailed solutions which are then evaluated and assessed, based on certain criteria. The best of these forms the decision for the current situation. If the decision is very different from the decisions which have been stored in the database, then the new decision with the corresponding attributes is added in the database.

Liao (2000) describes a case-based decision support system which models military command and control decision making. It contains five independent modules: case search, heuristic solution search, decision making, execution and learning. A current situation is passed from the situation board to the case search and heuristic search which independently produce possible strategic decisions, though these are not directly applicable. The solution search module then adopts those general strategic solutions and produces detailed solutions which are assessed by the decision making module on the basis of time, risk and reward criteria. The best then forms the solution for the current situation. If this decision is very different from that which is already in the case database, the decision and the situation attributes are appended in the database.

Neural Networks

Non-symbolic artificial intelligence in general and neural networks in particular have been used as a data driven technique for solving business problems and for modelling decision making in simulation (Wong 2000, Curram 1997). Unlike the pattern matching and case-based reasoning approaches, where the model is a database of decisions embedded in the simulation (Flitman & Hurrion 1987, Mason and Moffat 2000), neural networks are mathematical models which are calibrated using a database of attributes and decisions (Liang et al 1992, Curram 1997). Conceptually similar to regression models (Verbos 1994, Haykin 1999), the purpose of a neural network is to identify a causal relationship between input and output variables that can classify discrete or continuous output variables into distinct categories. For the classification, the inputs are used as predictors in the model. Technically, in a neural network model this is achieved using a heuristic such as Backpropagation (Verbos 1994), which iteratively adjusts the parameters of the model in order to minimise the total classification error between the model prediction and the actual observations. A distinct feature that has contributed to the increasing popularity of this technique is the multi-stage processing capability which enables the use of neural networks for universal approximation of any function (McClelland & Rumelhart 1988, Haykin 1999). The following paragraphs describe some examples in which neural networks have been used to model human decision making in simulation.

Liang et al (1992) describe a case study in which neural networks are employed to model scheduling decision making. Data sets with scheduling decisions and attributes (collected using a VIS - section 2.2) are used to calibrate the models. The efficiency of each decision making model is then assessed using simulation. The efficiency of the recorded decisions from each individual decision maker is also assessed using the same simulation model. Comparing the efficiency of the models with the actual decisions, the authors report that the neural network models are less efficient. However, when the neural networks are retrained by removing inconsistencies and inaccuracies from the dataset, the new models are reported to perform better than the actual human decision makers, whose performance was evaluated using the database with the recorded decisions.

Curram (1997) describes two examples for which neural networks are used to model human decision making. In the first instance, the capability of neural networks to model human decision making is tested using artificially generated data sets. In the second example, decision making is modelled in a simple simulation model. As has been explained in section 2.2, a local branch of a bank is used as a case study to test the feasibility of using neural network models to represent decision making in simulation. The neural network models that are developed in this case study represent customers' decisions about which queue to join when they arrive at the branch of the bank. The decisions about which queue the customers would join (if any) was modelled using seven attributes as predictors. The predictors represent the characteristics which determine customers' decisions about which queue to join. A number of stochastic and deterministic neural network models were developed and assessed. The evaluation of the models was based on the accuracy of the prediction and the internal consistency of the models. Validation using 'hold out' or 'one out' sample (UrbanHjorth 1994) was not undertaken given the limited data set, but qualitative checks were performed to make sure that the model was sensible.

Probability distributions conditional on attribute values

Probability distributions conditional on attribute values is perhaps the simplest and one of the most commonly applied techniques for controlling logic and for modelling decision making in simulation (O'Keefe 1986, Curram 1997, Pidd 1998, Robinson 2004). The technique involves the development of statistical or empirical distribution models which are calibrated using observed data sets.

The technique is usually applied when it is observed that a specific type of decision is almost always taken when a real valued attribute variable which follows a statistical distribution is less than a benchmark value (Perry & Moffat 1997). The implementation of this type of model in the simulation involves sampling a number from a statistical or empirical distribution that represents the current attribute value. The decision is then determined by comparing the sample with a benchmark value. The accuracy of this type of model depends on how consistently the rule is applied in the real system (Perry & Moffat 1997).

Perry & Moffat (1997) describe various experiments which involve representing decision making using statistical distributions. In their model, a statistical distribution is constructed about which type of information is more frequently required when a certain type of decision is taken. Using this concept, a number of experiments were undertaken in order to model decision making based on frequency of request for certain information. In their research, the authors report that the empirical data do not support the hypothesis that decisions are related to the information that the naval commanders request. As an extension of the above approach, Moffat & Witty (2002) present a preliminary investigation for modelling decision making using advanced conditional models which employ Bayesian statistics and the use of catastrophe theory.

Decision trees and rule induction

Symbolic artificial intelligence in general and decision trees in particular have also been used to model decision making knowledge in simulation (O'Keefe 1986, Flitman & Hurrion 1987). A decision tree, which can also be seen as a set of nested 'if-then' rules, is a logical model in which the attribute values are used sequentially to predict a decision. The most common decision tree derivation algorithm, known as Iterative Dichotomiser 3 (ID3 - Quinlan 1979, Quinlan 1983), is a data driven technique where the parameters for the conditional statements in the nested 'if –then' logical model are calibrated to fit the attributes and decisions in the dataset. The decision variable is divided into categories, each of which contains decisions of the same class. Each of the attributes is then divided into classes, based on their ability to predict the class of the decisions in the data set. The process is terminated when all the attributes have been used or all the decisions have been correctly classified. Each time that a new attribute is used, a new level in the tree is created. The attributes chosen first are those which can predict the class of the greatest number of decisions first. The following are some examples where decision trees are employed to model decision making knowledge.

O'Keefe (1986), in his taxonomy of expert systems, describes how the Quinlan ID3 (1979) algorithm could be used to generate rules from a data set to model decision making in simulations. Welbank (1983) appears less optimistic about the use of rule induction and reports that in his experiment it was not possible to generate a sufficiently complete case library of decisions for each possible decision situation.

Pierreval & Ralambondrainy (1990) describe a decision making problem of a simplified manufacturing environment, in which a decision is needed about which of the items queuing in front of two machines should be processed first on each occasion that a machine becomes available. An ID3 type decision tree is used to model the decision making process. The ID3 decision tree is calibrated using optimal decisions and observed attributes which are obtained from simulation experiments with a model of the manufacturing system. Robinson et al (1998) describe a model of a fictional lorry loading bay, where lorries arrive at the loading park on average every 10 minutes and require loads of between 5 and 20 items. The decision maker is a loading supervisor who has to make a decision about which bay should be allocated to a lorry which has just arrived. To illustrate the feasibility of representing decision making in simulation using decision trees, one of the authors in this research acted as a human expert. A data set of attributes and decisions which was collected from a VIS-based data collection session was used to develop an ID3 decision tree that was later linked with the simulation model. Mingers (1986) uses decision trees to model the decision making process that is involved during development of a regression equation. In this model, regression analysis output diagnostics are used as attributes to decide which independent variables must be included in the regression model. Williams et al (1989) use decision trees in order to model decision making about replenishment. In this case study, a decision making model is developed and is used in order to prioritise and schedule replenishment requests. In this model, decisions about what type of ship should be sent to implement the replenishment and when it should do so are predicted by taking into account who requests replenishment and where.

Fuzzy Logic

Fuzzy logic has also been proposed by a number of authors (Mundermann 1993 and Zadeh 1973) as an approach which can be used to model human decision making knowledge in simulation. Fuzzy logic can be used when the representation of uncertainty is required and when a number of conflicting rules need to be combined (Curram 1997). When fuzzy logic is employed the model, instead of predicting a specific decision for a set of attributes, predicts a fuzzy set for a number of fuzzy sets that represent the attributes. The fuzzy set that is predicted as an output of the model represents in a non-probabilistic way the likelihood of a decision to be of a specific type. Predicting fuzzy sets rather than specific decisions, the models allow for the representation of partial membership and statements with a linguistic nature such as maybe or probably. With this approach, fuzzy logic enables the representation of the fact that two individuals are tall but also to distinguish between one person who is taller than the other (MathWorks 1999). In fuzzy logic, fuzzy sets and rules are inputs in the modelling process. This allows the representation of uncertainty and the combination of conflicting rules. In fuzzy logic it is assumed that the decision makers are able to express the rules which they apply when they make decisions. As an example, Mundermann (1993) describes an experiment where fuzzy logic is used to represent agents of autonomous mobile vehicles in a highway simulation.

Problem Enumeration

Optimisation has been used to model human decision making when there is evidence that in the real system the decision makers apply optimisation techniques in order to make decisions (Curram 1997). Mason & Moffat (2000) use optimisation to represent human decision making. They use game theory to model planned decision making. They assume that, given sufficient time, the decision makers use all the available information to make rational decisions. Based on this assumption game theory is applied to represent the fact that all the possible decision options are considered and assessed in order to identify the decision option with the lowest risk.

2.3.2 Qualitative knowledge modelling

Analysis of Influence

Analysis of influence is one of the qualitative techniques which has been used to model decision making knowledge (Perry & Moffat 1997). It involves asking the decision makers to explain how each of the attributes affects each of their decisions (negatively or positively). With this process, the most influential attributes can be identified. Perry & Moffat (1997) use analysis of influence to model maritime decision making. The focus of the analysis in the second part of their work (in which analysis of influence is used) is on how the decisions are taken and how each decision maker is influenced by each attribute when he makes a decision. They are concerned about the individual human decision makers and how the attributes of the system are associated with their decisions. They also report that, in the specific application, the comparative analysis between analysis of influence and statistical distributions revealed that analysis of influence is more appropriate for modelling decision making. With the data set which they had available, they found that it was not possible to use statistical distributions to develop a model that can predict *what* decision must be taken given the attributes of the system. However, using analysis of influence (and influence diagrams) it was possible to represent how the decisions are taken by identifying how each attribute influences the decision maker (positively or negatively).

2.3.3 Discussion: techniques for modelling decision making strategies

Following the description of a series of different applications in the previous paragraphs of this section, it is clear that a number of different knowledge modelling techniques have been used as part of the process of modelling decision making in simulation (Perry & Moffat 1997, Flitman & Hurrion 1987, Liao 2000, Curram 1997). Pattern matching appears to be one of the most frequently applied, with a number of authors reporting full-scale applications where pattern matching is employed in knowledge modelling (Perry & Moffat 1997). From the above review, it can be concluded that in the literature the choice of the knowledge modelling technique is often arbitrary. This conclusion is based on the fact that there is not enough evidence to support the idea that the choice of the modelling technique is based on a comparative analysis of the performance of the alternative techniques. It indicates that, empirically, the pros and cons of the alternative knowledge modelling techniques have not been fully investigated by the current literature and further research is required. In addition, in the above review statistical modelling appears to be one of the less popular techniques, since there are very few examples where statistical modelling is used as part of the process of modelling decision making in simulation. This conclusion indicates that modelling decision making in simulation is a research area that is rather biased towards AI oriented techniques. This is consistent with what a number of authors (O'Keefe 1986, Doukidis 1987) have observed in the past and, although it is justified by the similarities that can be observed between AI and simulation, it indicates that further research is necessary in order to investigate the potential of statistical modelling.

2.4 Techniques for representing and linking decision making models with simulation

The final stage in the process of modelling decision making in simulation involves linking the decision making models with the simulation (Williams 1996, Flitman & Hurrion 1987, Robinson et al 2003b). The implementation of this stage depends on the degree of success of the decision making models and on the motivation of the modelling exercise (Perry & Moffat 1997 - quite often, when the motivation of modelling is to understand the decision making process, the third stage is not implemented since the objective has been achieved before the completion of the exercise).

In the literature, the approach which is frequently chosen to represent the model in the simulation is related to the nature of the technique that is chosen to model the decision making (Mason & Moffat 2000, Flitman & Hurrion 1987, Robinson et al 2003b). Depending on the development environment that is used to build the simulation, various techniques have been used to represent and link the decision making models with the simulation. The following paragraphs describe some of the most popular approaches that have been proposed.

2.4.1 Agent technology and object orientation

In the literature, agents are sometimes described as a knowledge modelling technique Schimidt (2000). Based on the agent definition given by Mason & Moffat (2000) and Moffat (2000), in this review agents are seen as a technique for representing and linking decision making models in simulation. Consistent with the Mason & Moffat (2000) approach, our decision to classify agent technology as a technique for knowledge representation has been taken on the grounds that an agent is a means of organising a decision making model. An agent can be seen as software architecture for representing a model of a human decision maker in simulation. It interacts with other entities of the simulation and with other agents if multi-level or multi-unit decision making is modelled.

Schimidt (2000) describes agents as virtual representatives of real human beings who have common properties such as individual worldview, autonomous behaviour, planning capability, social abilities, communication and co-operation capabilities.

Agents are usually implemented in object-oriented environments due to the benefits that occur from code reusability and flexibility (Moffat 2000, Pidd 1996). Although a full review of the approaches which have been proposed for implementing agents is beyond the scope of this thesis, the following paragraphs give an overview of the general object-orientated approach and describe some of the applications of modelling human decision making for which agents are used to link the decision making models with the simulation.

In an object-oriented environment, an agent is implemented as an object with properties and methods (Moffat 2000). Usually, the number of objects depends on the number of the decision makers who are involved in the decision making process. If there is only one level of decision making but many decision makers, then all the objects that represent agents are defined from the same class. If multi-level decision making is modelled, different classes are used to define the objects at each decision making level (Moffat 2000). Normally, object methods are used to represent different decision actions and object properties are used to check the availability of the decision makers and to store the knowledge which comes from the decision making model. In a pattern matching model, for example, a property of type 'list' can be used to store the database of the decisions and attributes (Moffat 2000). Properties can also be used to model its position and condition. They might be used, for example, to check whether the agent has to decide under pressure or is in a situation where a decision is not required urgently, in which case planning skills and knowledge can be applied in order to make a rational decision (Moffat 2000). Depending on the level of detail, different classes might be used to define the objects. This is usually the case when multi-level decision making structures are modelled and so various types of agent need to be represented.

Mason & Moffat (2000) and Moffat (2000) describe an agent architecture that captures the key command and control processes. These processes include intelligence activities of data fusion, recognised picture compilation, decision making and planning. The purpose of representing these processes was to gain a perception of the outside world, to understand what is going on and finally to decide what to do next by formulating a plan of how to achieve this. The agents consist of the following objects: a 'com' object, a collector, a planner, a promulgator, the recognised picture and the plan.

The 'com' object in their research allows the agent to interact with other agents. The collector is used to request data from the simulation and to apply pattern matching techniques to see what should be decided in the short term. The collector is also responsible for informing the planner component of what is happening. The planner, whose role is to decide what to do in the long term, controls the collector by requesting it to collect and analyse information. The promulgator is used in the simulation logic to make sure that the decisions of the agent are executed. The recognised picture is a database which contains all the decisions that the agent has taken about the current situation in each geographical zone. The plan is also a database which contains the current plan that the agent has decided to apply.

2.4.2 Remote expert controller

The remote expert controller was an early attempt to separate the logic from the rest of the simulation model (Flitman & Hurrion 1987 and Hurrion 1991). The technique involves the use of a separate machine that contains the logic of the simulation and the implementation of the decision making model (Flitman & Hurrion 1987). The machine which is used for the other components of the simulation interacts with the machine that contains the decision making strategies. It does so by requesting decisions for specific situations which are described through a set of attribute values. Very often in the literature the motivation for proposing this type of technique is based on the benefits which can occur when specialised development environments (such as LISP logic language: Jackson 1998, Edwards 1991) are used to develop, maintain and operate complex decision making rules used by the simulation.

An implementation of a remote expert controller is described by Flitman & Hurrion (1987) and Hurrion (1991). Both papers describe a model where the expert system function was clearly separated from the simulation. The two parts of the system were placed on separate micro-computers, which interact with each other in real time during the simulation run through an RS232 interface. As the authors explain, by separating the components of the system it was possible to make full use of the capabilities that the expert systems can offer to control the logic of the simulation model.

2.4.3 Integration with third party applications

The development of decision making representations as third party applications, which interact with the simulation model in real time during the run using .com objects and OLE windows functionality, is perhaps the most recent approach for implementing and linking decision making models with simulation applications (Robinson et al 2003b). In addition, it is the most commercialised approach (Lanner Group 2003, Attar software 2000) and the one which, according to various authors (Robinson et al 2003b and Standridge & Steward 2000), is most likely to be adopted in industrial applications since it is compatible with commercial simulators and AI packages.

Robinson et al (2003b) describe a simplified example of a simulation model that is linked with an expert system software. The Witness (Lanner Group 2003) simulation model runs until a decision point is reached. When a decision point is reached, the simulation halts and the expert system becomes the active application. The software then retrieves information about the attributes of the simulation model and sends a decision to the simulation model, which once more becomes the active application and continues the simulation process until the next decision point.

Standridge and Steward (2000) describe a simulation model which is linked with an expert system to represent decision making about patient appointment scheduling. The simulation model was written in Slam system. The expert system was implemented in C language and the link in Fortran.

2.4.4 Discussion: representing and linking decision making models with simulation

It is concluded, from the above discussion, that there are at least three types of technologies which have been proposed for implementing and linking decision making applications with simulation. Agent technology, which has been very popular in military applications, involves the development of a decision making representation inside the simulation model by using object-oriented structures, which allow the agents to act independently and to interact with each other and with the rest of the simulation (Moffat 2000). Remote expert controllers involve the implementation of the decision making representation in separate software in separate machines (Flitman & Hurrion 1987). This allows a more efficient development, maintenance and operation of the decision models in specialised environments (Hurrion 1991). Finally, a more industrial-oriented approach, which does not involve the use of separate hardware but enables the separation of the decision making model from the rest of the simulation, is the use of third party applications and OLE connections to develop and link the decision making models with the simulation (Robinson et al 2003b).

2.5 Current methodologies for modelling human decision making in simulation

Having described the techniques that have been used as part of the process of modelling decision making in simulation, it is clear that that a number of quite different approaches have been developed to model decision making in simulation. From the above review it is clear that, although most of the methodologies involve three stages (elicitation, modelling and representation), the emphasis which is given in each of these three stages of the process and the type of problem that each methodology attempts to solve vary significantly across the proposed approaches (Perry & Moffat 1997, Curram 1997). Clearly, the techniques which are employed by each methodology are not the only element that differentiates them. The methodologies that have been proposed also differ on how they use the techniques, on what they try to achieve, on what they define as decision making, on the environment and the knowledge that they use to develop the models and on the focus of representation of the decision making process (Pierreval & Ralambondrainy 1990, Flitman & Hurrion 1987, Curram 1997). Taking into account these levels of differentiation, it is possible to identify at least the following five characteristics which vary significantly across the methodologies.

- The nature of decision making. This is a measure of the extent to which the decision making model must attempt to represent specific human decision makers or an efficient decision making practice (Curram 1997).
- The motivation of the modelling exercise. This is author's motivation and objective for modelling decision making in a simulation model (Curram 1997, Flitman & Hurrion 1987).
- The psychological content. This shows to what extent the model should attempt to explain how the decisions are taken (Perry & Moffat 1997).
- The contextual content of the methodology. This measures how generic the decision making model must be (Perry & Moffat 1997).
- The nature of knowledge that needs to be acquired for developing the models. The nature of knowledge reveals whether the decision making process is based on explicit or tacit knowledge and, as has been seen in the above review, it determines the type of elicitation technique which must be used in the methodology (Cornnell et al 2003).

In this section, in order to identify the types of approaches that are available to model decision making in simulation, the above five characteristics are used to categorise and group the approaches which have been proposed in the past. Having identified the types of approaches that are available for modelling decision making, this chapter concludes by highlighting the areas which have not been addressed in the past and by identifying some of the current research issues.

2.5.1 Nature of decision making

Taking into account that the nature of decision making is one of the methodological characteristics that vary significantly across those methodologies which the various authors have proposed, this is the first attribute that can be used to categorise the research in modelling decision making in simulation. As is clear from the previous three sections, the methodologies proposed vary between those which have been developed for representing human decision making (Curram 1997, Perry & Moffat 1997) and those which have been developed for modelling optimal decisions (Mason & Moffat 2000).

Modelling human decision making in simulation involves modelling specific human decision makers who are part of the system that the simulation model represents (Curram 1997). As has been shown in the previous sections of this chapter, representing human decision making might involve modelling inefficient decision making practices if this is what is applied by the decision maker in the system (Curram 1997). Models of human decision making are usually a conceptual representation of the fact that the efficiency of the decisions varies across the decision makers who are involved in the decision making process (Flitman & Hurrion 1987). In addition, it is a conceptual representation of the fact that different people have different experiences and, therefore, might take different decisions for the same decision situation (Flitman & Hurrion 1987, Curram 1997).

On the contrary, as has been shown by methodologies such as those proposed by Mason & Moffat (2000), Moffat (2000) and Pierreval & Ralambondrainy (1990), modelling optimal decision making is usually impersonal. The models do not represent a specific decision maker but an optimal decision making process or a hypothetical efficient decision maker (Curram 1997). Optimal decision making is normally a conceptual model for representing the fact that human decision makers sometimes have the time to use analytical techniques, algorithms and decision support systems in order to make a decision under conditions of perfect information (Moffat 2000). Taking into account that methodologies such as those proposed by Hurrion (1980) are also associated with optimality, it is clear that modelling optimal decision making in simulation is a theme that has received significant research attention. One of the main reasons for this attention is the fact that very often clients of simulation projects believe that the most appropriate decision option is chosen when a decision is required.

2.5.2 Motivation of modelling

Using as an attribute the motivation of modelling, which is the second methodological characteristic that has been identified above as one that varies significantly across the proposed methodologies, it is possible to categorise the current research into three relatively distinct research streams (Mason & Moffat 2000, Flitman & Hurrion 1987, Hurrion & Secker 1978).

The first group includes methodologies such as those of Perry & Moffat (1997) and Mason & Moffat (2000) that have been proposed for situations for which a model of decision making is needed to improve the specification, the accuracy and the planning capabilities of the simulation. Mason & Moffat (2000), in their research on representing the C2 process in simulation, justify the implementation of the decision making models in military simulations by explaining that the decision making process needs to be represented in models of conflict in order to simulate realistic force behaviour and effectiveness. Perry and Moffat (1997) also conclude that simulation models for which decision making is subsumed in a command and control module would benefit more from the results of external analysis, such as the statistical representation of decisions. In their research, the benefit in the simulation is translated into better specification and, therefore, more realistic military behaviour. Williams et al (1989) also comment that a model in which the decision making was represented proved effective as a planning aid since, having modelled the decision making process, the simulation model could be used to predict the effect of various policies associated with stock levels.

The second group which can be formed when the motivation of modelling is used for classifying the methodologies, includes the approaches that have been proposed for situations for which a number of decision making models are developed to assess the performance of alternative decision makers using simulation (Flitman & Hurrion 1987, Liang et al 1992). This is the motivation of a more recent theme in research into modelling decision making in simulation. The aim is to identify improved decision making practices by modelling and comparing the performance of individual human decision makers. Flitman & Hurrion (1987) and Hurrion (1991) are representative of this research theme and they propose the use of this approach to develop simulation-based expert systems for improving decisions. As has been instanced in sections 2.2.1 & 2.3.1, the above authors model the behaviour of five decision makers in a simulation and use the model to evaluate the efficiency of each decision maker. To identify better decision making practices, the strength of the five separate decision making models are combined to develop a model which could perform better than each individual decision maker. The simulation assessment of the combined model confirmed that the combination of the strength of each decision maker results in superior performance. Liang et al (1992) is another research study which falls into this group. As has also been seen in sections 2.2.1& 2.3.1, they describe the development of an improved practice decision making model. In order to improve the decision making practice, the data sets are processed using the Markov process to remove inconsistencies and obvious incorrect decisions. The processed data sets are used to calibrate a neural network scheduling advisor. In their research, they report that the simulation assessment of the model showed that it performs better than the human decision making models.

There is a third research stream which can be formed when the motivation of modelling is used to classify methodologies. This stream includes all those approaches that have been proposed for situations where representations of decision making are developed within the simulation for deriving optimal decision rules through simulation experimentation (Pierreval & Ralambondrainy 1990, Hurrion & Secker 1978). This is perhaps one of the most common motivations in the research in modelling decision making in simulation. As evidenced during the review of the elicitation techniques (in section 2.2), methodologies such as those proposed by Hurrion & Secker (1978) show that VIS enables the user to watch the simulation process and to alter the decision making strategies during the simulation run. This helps in the understanding of how the decision making strategies affect the performance of the system and enables the user to identify the most efficient decision making strategy.

2.5.3 Psychological content

The psychological content is the third characteristic that was identified at the beginning of this section as one which varies significantly across the methodologies that have been proposed and, therefore, it is the third attribute which serves to discriminate among different types of research in modelling decision making in simulation. The psychological content (Perry & Moffat 1997) is a measure of the extent to which the analysis focuses on the understanding of the mental process involved in the decision making. Using the psychological content as a continuum, it is possible to locate at least three types of methodologies for modelling human decision making in simulation.

At one end of the continuum there are external methodologies, such as those proposed by Curram (1997), Liang et al (1992), Perry & Moffat (1997) and Perry & Moffat (2000 first case study). Their principal focus is on developing models for predicting decisions. From the review of the modelling techniques in section 2.3, it is clear that very often the authors use non-symbolic AI, statistical modelling and probability distributions. In all these cases the modelling process normally focuses on representing *what* decisions must be taken rather than representing *how* those decisions are taken. This is clearly an external approach for modelling decision making since it does not involve understanding the mental process which takes place when a decision is taken.

At the other end of the continuum, there are internal methodologies such as those proposed by Perry & Moffat (1997 second case study). These concentrate mainly on understanding the decision making process using techniques such as influence diagrams, cognitive maps and VIS. The aim of the research which focuses on modelling human decision making internally is to resolve the problems that exist when external approaches are applied. Modelling decision making internally involves understanding and representing the mental process that the decision markers undertake in order to make decisions. Clearly, when internal methodologies are applied, the focus of representation is on how the decision makers decide rather than on what they decide.

Between these two are located methodologies such as those proposed by O'Keefe (1986), Robinson et al (1998) and Pierreval & Ralambondrainy (1990). These methodologies, using rule induction for developing decision trees, combine understanding of the decision making process with decision predictability. With this kind of technique, as has been explained in section 2.3, external data such as pairs of observed decisions and attributes are used to develop models which are useful for understanding the decision making process and which can also be used to predict decisions.

2.5.4 Contextual Content

The contextual content is the fourth attribute which can serve to discriminate among different types of research in modelling decision making, put forward once again by Perry & Moffat (1997). It measures the extent to which the decision maker in the analysis operates within a context similar to the one in which real life decisions are made. It also measures the extent to which the methodology has been developed as part of the process of solving a real world problem. Using the contextual content as a continuum, it is possible to locate at least three types of methodologies for modelling human decision making in simulation.

Firstly, at one end of the continuum are the methodologies which are based on industrial or military applications, where the decision makers operate within a context similar to the one in which real world decisions are made (context full methodologies: Mason & Moffat 2000). Secondly, at the other end of the continuum, are the methodologies which have been developed mainly on theoretical grounds and which address issues related to the technical end of the process of modelling decision making in simulation (context free methodologies: Doukidis 1987, Doukidis & Paul 1990, Bell & O'Keefe 1994). In those approaches the decision makers, when and if they are involved, operate in an environment that is different from the one in which real decisions are taken.

Thirdly, in the middle of the continuum, are located the methodologies which have been developed and tested involving decision makers who operate within an environment that is slightly different from the one in which real world decisions are made. The decision maker might be hypothetical, such as students (Liang et al 1992, Angelides & Paul 1999, O'Keefe & Pitt 1991, Bell & O'Keefe 1995) or the authors of the papers (Robinson et al 1998) or the decision making environment might be a simplified example of a real world decision making environment (Flitman & Hurrion 1987). Simplified examples of real world decision making environments are normally used as a preliminary stage for illustrating concepts, theorems and for testing algorithms. The general job shop scheduling problem (Hurrion 1978) and the general travelling salesman problem (Reeves

1995) are two examples of research that has been undertaken using simplified examples of real world decision making environments.

Military simulations, such as those developed by Mason & Moffat (2000), Perry & Moffat (1997) and Williams et al (1989), are perhaps the most common examples of research in modelling human decision making in simulation which is based on real world problems and which involves the use of situation specific simulations with real world decision makers. Standridge and Steward's (2000) research in modelling patient appointment scheduling in a simulation model can also be classified as context full research since it is applied in a specific medical clinic.

Methodologies such as those proposed by Pierreval & Ralambondrainy (1990), Hurrion (1991), Flitman & Hurrion (1987) are only some of many examples of research which is located in the middle of the context free - context full continuum. As has been noted, most of this research describes approaches for modelling decision making which were developed based on simplified generic scheduling decision making environments. Liang et al (1992) is a slightly different example of research that is located in the middle of the context free - context full continuum. It is located in the middle of the continuum because the data sets that are used are collected from students. As the authors comment, the fact that non-industrial decision makers are used might reduce the applicability of the findings, since it could be argued that real decision makers would have challenged the decision making process, the attributes and the options available in the simulation model.

Research investigations, such as those carried out by Doukidis (1987), Doukidis & Paul (1990), Bell & O'Keefe (1994) and Paul et al (1997), provide reviews of progress in the integration of AI with simulation and, as such, they are classified as context free research which provides useful insights and motivation for combining AI with simulation (Perry & Moffat 1997).

2.5.5 Type of knowledge

The fifth attribute that can be used to locate in a continuum the methodologies that have been proposed for modelling decision making is the assumption about the nature of knowledge which the decision makers use in order to take decisions (Cornell et al 2003). The knowledge required for making decisions can be located in a continuum that is defined by two points. Around the first point are concentrated all those methodologies which have been developed on the assumption that the decision makers base their decisions on predetermined hard rules, given to them in the form of a user guide or other type of documentation. According to Cornell et al (2003) this knowledge, termed explicit, consists essentially of concepts, information and insights that are specifiable and can be formalised in rules and procedures. Around the second point of the continuum are concentrated all those methodologies which have been developed on the assumption that the decision makers make decisions based on knowledge which has been acquired through learning and experience. This type of knowledge, which is known as tacit, according to Cornell et al (2003) involves less specifiable insights and more skills embedded in individual or organisational contexts. Usually decision making is based on tacit knowledge when the problem is too complex or when it involves too many decision options (Cornell et al 2003).

Most of the approaches that have been described in the previous sections of this chapter (Perry & Moffat 1997, Mason & Moffat 2000, Pierreval & Ralambondrainy 1990, Hurrion 1978, Williams et al 1989 and Flitman & Hurrion 1987) describe decision making environments with many decision variables and many decision options. For all these cases there are no written rules to describe explicitly how decisions must be taken. They are, therefore, decision making problems that are based on tacit knowledge.

2.5.6 Summary of methodologies

Having categorised the research activity in modelling decision making in simulation using the five attributes described at the beginning of this section, it is clear that, as a result of the research activity in modelling decision making in simulation, a number of different approaches have been developed (Flitman & Hurrion 1987, Perry & Moffat 1997, Mason

& Moffat 2000, Hurrion & Secker 1978). The above discussion has shown that the diversity of approaches is not limited only to the techniques which are employed. Throughout the years, technological advances in computational power and the emergence of artificial intelligence have enabled a broader diversification, which is reflected in the psychological content, the motivation, the context and the nature of the decision making models that have been developed. Based on previous work, it is possible to identify four existing approaches for modelling decision making in simulation.

- Modelling human decision making externally to compare and assess the performance of alternative human decision makers in simplified examples of real world problems (Liang et al 1992, Flitman & Hurrion 1987).
- Modelling optimal decision making to assess the performance and effectiveness of alternative heuristic rules in simplified examples of real world problems (Hurrion & Secker 1978).
- Modelling human decision making internally or externally to improve the specification and accuracy in military simulation models (Perry & Moffat 1997).
- Modelling optimal decision making to improve the specification and accuracy in military simulation models (Mason & Moffat 2000).

The above approaches and the methodological ideas which have been described in the previous pages of this section can be summarised as in table 2.2, which also outlines the current trends in the literature.

	Approach 1	**Approach 2**	**Approach 3**	**Approach 4**
Authors	Liang et al (1992), Flitman & Hurrion (1987)	Hurrion (1980), Hurrion & Secker (1978), Hurrion (1978)	Perry & Moffat (1997), Mason & Moffat (2000- First case)	Mason & Moffat (2000 – Second case)
Motivation	Assessing performance of human decision makers	Assessing performance of heuristic rules	Improving accuracy and specification of simulations	Improving accuracy and specification of simulations
Nature of decision making	Human	Optimal	Human	Optimal
Psychological content	External	External	Internal or External	External
Contextual content	Simplified examples of real world problems	Simplified examples of real world problems	Military simulations	Military simulations
Knowledge	Tacit	Tacit	Tacit	Tacit

Table 2.2: *Current approaches for modelling decision making in simulations.*

Table 2.2 shows that in non-military simulations the motivation for modelling decision making is often to derive optimal decision strategies through experimentation with alternative heuristic decision making representations (Approach 2: Hurrion 1980, Hurrion & Secker 1978, Hurrion 1978). On the basis of this conclusion, it is evident that optimisation is one of the main themes which has motivated many of the authors who have proposed methodologies for modelling decision making in non-military simulations (Hurrion 1980, Hurrion & Secker 1978, Hurrion 1978). A relatively recent research stream involves the use of simulation and artificial intelligence to assess and compare the performance of alternative decision makers (Approach 1: Liang et al 1992, Flitman & Hurrion 1987, Hurrion 1991). Owing to the novelty of this kind of approach, most of the proposed methodologies have been developed and tested using simplified examples of real world problems (Flitman & Hurrion 1987, Liang et al 1992, Hurrion 1991, Bell & O'Keefe 1994). Based on these experiments, a number of authors have concluded that simulation and AI can be used for the development of expert systems for improving decisions, but the implementation would require a methodology for identifying and assessing decision making strategies using simulation (Curram 1997). The military applications which have been described show that the motivation for modelling the decision makers in military simulations is slightly different. This conclusion is drawn from the fact that in the majority of the military applications which have been described, decision making is modelled mainly in order to improve the accuracy and the specification of the simulation model. Comparing the motivation and the type of decision making which is modelled in the various papers that have been described, the conclusion is reached that there is not a clear relationship between these two. For example, in military simulations where the motivation is to improve the specification and the accuracy of the simulation models, optimal and realistic decision making models have been developed for the same motivation (Perry & Moffat 1997, Mason & Moffat 2000).

2.6 Potential research

From the above categorisation it is clear that, although many types of approaches have been proposed for modelling decision making in simulation, only the military models have been developed and implemented in tackling full scale problems (Perry & Moffat

1997, Mason & Moffat 2000). Having been unable to identify a significant industrial application which demonstrates the feasibility and the benefits from modelling decision making in simulation, it is clear that there is lack of empirical evidence in the research in modelling decision making in simulation (Curram 1997). This review shows that there are many issues for potential research, some of which are as follows:

- As was found in section 2.5.6, the use of simulation for improving the performance of alternative decision makers is an area which has only been addressed in the past on theoretical grounds (Approach 1: Flitman & Hurrion 1987, Liang et al 1992). Most of the approaches that have been proposed in this area conclude by suggesting that the use of simulation for understanding and improving the design of decision making would require a modelling methodology (Curram 1997). This indicates that the methodological issues which are associated with the above process have not been fully addressed and the development and testing of the use of simulation for understanding and improving the design of decision making is a potential research area.

- From the discussion of the knowledge elicitation techniques in section 2.2, it is clear that from the current literature it is not known to what extent it is possible to use VIS-based knowledge elicitation in a real context. Very few of the approaches which have been proposed address issues associated with the process of identifying the decision situations that must be presented to the decision maker (Perry & Moffat 1997, Curram 1997). In addition, none of the approaches proposed in the past address issues associated with the process of understanding the elements of the decision making problem. On this basis, it is clear there is lack of a methodology and software (Williams 1996) for simulation based knowledge elicitation and this must be identified as another area where further research is required.

- In section 2.3 it was found that in most of the non-military applications, techniques such as neural networks and pattern matching have been proposed or used for modelling decision making in simulation (Flitman & Hurrion 1987). In section 2.5.3, using as an attribute the psychological content, a distinction was made between internal and external techniques and it was found that both neural networks and pattern matching are classified as external, since they do not focus on understanding and representing the decision making process but only on the prediction of decisions. Given this distinction, it is evident that the effectiveness of internal techniques (such as influence analysis or decision trees) for modelling decision making in non-military applications has not been fully investigated and, therefore, it is an additional research issue which can be identified from the previous categorisation.

- Finally, the impact on the prediction of the model and the benefits which can occur as a result of modelling decision making have only been investigated in military models (Moffat et al 2004). One of the reasons for the lack of research in this area could be that various authors in the past (Curram 1997) have questioned the benefits and the need to improve the accuracy of simulation models. Leaving aside debate as to what extent this is true, it is clear that research is required in order to investigate the impact that modelling decision making has on the prediction of simulation models.

2.7 Conclusion

In this chapter, the main stages of the process of modelling decision making in simulation have been explained and the techniques proposed for implementing each stage of the process have been discussed. In addition, five characteristics which vary significantly across the proposed approaches have been identified and have been used to provide categorisations of the research. From these categorisations, the main types of approaches that are available for modelling decision making in simulation have been identified and a number of potential research issues have been discussed.

In section 2.6 it became clear that modelling decision making in simulation is an area with significant research activity, which in the past has focused on the investigation and theoretical development of techniques and approaches for modelling decision making in simulation. In this chapter, it also became clear that modelling decision making in simulation is an area with many unanswered research questions. The lack of empirical evidence regarding the feasibility of the approaches for industrial applications has been identified as one potential research issue. This has raised a number of specific research questions associated with the industrial applicability of the early stages of the approaches.

In section 2.5, it was concluded that the use of simulation for understanding and improving the performance of the alternative human decision makers is one of the most recent themes of the research in modelling human decision making in simulation. Taking into account that most of the research in this stream is based on non-industrial - usually scaled-down problems - it is concluded that there is an empirical research gap, since fundamental research questions related to the difficulties and limitations of the use of simulation for understanding and improving the design of decision making have not been addressed.

Taking the above conclusion as its basis, the next chapter sets out the objective of this study by outlining the research questions and describing the research approach that will be followed in order to answer these questions.

Chapter 3

Research programme design

Chapter 2 outlined the current techniques and methodologies for modelling decision making in simulation. The review indicated a number of theoretical and empirical research shortcomings that provide the motivation for this research. Based on the research shortcomings identified in Chapter 2, in this chapter the research questions and the objectives of the research are described and the method of approach employed is explained. Taking into account the conclusion from the review of the literature, section one outlines the research objectives and the research questions. Section two describes the research programme. It explains the research method which is employed in this research and it provides an overview of the research design. Section three describes the limitations of the research programme and section four concludes by highlighting the main elements of that programme.

3.1 Aim, objectives and research questions

From the literature it was found that the use of simulation for improving decisions would require a modelling methodology (section 2.6). Previous authors have shown that this is a process which involves three phases: simulation-based knowledge elicitation, knowledge modelling and knowledge representation. Having found that there is little evidence of applied work which brings together the above three phases, the aim of this research is:

> *To develop and test in a real context the process involved in the use of simulation for understanding and improving the design of decision making practices.*

In order to work towards this aim, a number of objectives and research questions have been set. These arose out of the literature reviewed and investigate three research issues that have not been addressed by previous research.

In Chapter 2 it was concluded that previous work in the area is focused on the theoretical and technical aspects of the process of modelling decision making in simulation. The previous research has shown that VIS and AI can support the process of modelling decision making but the findings have not been tested in industrial applications (Curram 1997). Several authors (Curram 1997, Liang et al 1992) have observed that, as an implication of the lack of industrial applications, the presence of decision making in simulation presents a number of problems to simulation modellers. Firstly, it is not known to what extent it is possible to use VIS-based knowledge elicitation in an industrial application (Curram 1997). Secondly, it is not known to what extent AI can be used to determine decision making strategies employed by real decision makers (Liang et al 1992). Thirdly, the potential benefits (in identifying efficient decision making strategies) and the impact of the presence of decision making in simulation are not known, since previous research investigations are based on hypothetical problems (Curram 1997, Flitman & Hurrion 1987).

Clearly, the above problems discourage the presence of human decision making in simulation. This is noted by Williams (1996) and Curram (1997) who comment that, due to the lack of a working methodology, the practical difficulties of modelling decision making in simulation have not been addressed. In attempting to resolve the above problems, the research will work towards the following research objectives:

- **Objective 1**: *Form a conceptual methodology to capture efficient decision making.*

- **Objective 2**: *Investigate the feasibility of the conceptual methodology by addressing the following research questions:*

    1. Is it possible to use VIS within an industrial environment to elicit the knowledge and collect the data required for modelling individual human decision makers using AI?

2. What are the practical difficulties of the above process and what are the methodological implications for the OR analyst?

3. Is it possible to identify, within an industrial environment, decision making strategies of individuals (that are appropriate for representing the individuals in simulation) by modelling a sample of their decisions using AI?

4. Is it possible to identify efficient decision making strategies by representing and assessing individual decision making in a simulation?

- **Objective 3:** *On the basis of the testing, refine the methodology and discuss the lessons that have been learned.*

3.2 Research programme

The research questions having thus been identified, this section provides an overview of the research programme. It describes the research methods which will be used to address the research questions and it explains the research design that will be applied.

3.2.1 Research methods

From the previous section, it is clear that the research is focused on what can be learned from the industrial application of a methodology for modelling decision making. All the research questions are derivatives of 'what' type questions about a contemporary set of events over which the investigator has little control. Their purpose is to bring together and test a set of specific techniques which have been proposed by previous authors. They all aim to cover contextual conditions and all attempt to address a situation which benefits from the prior development of a theoretical proposition that can guide data collection and analysis.

Keeping in consideration that the research questions are exploratory and that they focus on the contextual conditions which are missing from previous theoretical developments in the area, a deductive case study approach is used in this research (Bryman & Bell 2003).

The case study approach was chosen from a set of alternative approaches because, according to Yin (2003), it can cope with exploratory what questions. It can facilitate a process for covering contextual conditions and it can be used for situations which benefit from the prior development of theoretical propositions.

3.2.2 Research strategy

As part of the objectives of the research a methodology will be formed and tested using the case study described in section 1.3. The methodology will cover data collection, modelling and simulation-based model assessment. Using theoretical developments from previous authors, the methodology will provide a set of stages for implementing the process which the research questions aim to test. The implementation of each stage will provide the empirical evidence required for addressing the research questions.

As will be explained in detail in the next chapter, the methodology will involve the development of a case specific model. Owing to model development time constraints, a single case study strategy will be used. This limits the deductive power of the conclusions but provides in-depth empirical evidence about the feasibility of the approach. Similar case study-based deductive strategies have been used by previous authors who have contributed in this research area (Curram 1997, Flitman 1986).

The case study at the Ford engine assembly plant, introduced in Chapter 1, was seen as appropriate. The principal reason was that the simulation model that was available (section 3.1) could be adopted in order to:

- Support the knowledge elicitation process.
- Determine and compare the impact of representing individual characteristics of human decision makers in the model.
- Compare the performance of the alternative decision making strategies.
- Validate the findings since it represents realistically part of an existing production system.

3.2.3 Validation strategy

In order to investigate the feasibility and the limitations of the use of VIS as a knowledge elicitation technique (research questions 1 & 2), industrial experts will be involved in knowledge elicitation experiments. These experiments will be implemented using a number of alternative VIS prototypes. The prototypes will be developed based on the pre-existing simulation model of the case study and they will be used to capture the required knowledge. In order to validate the collected data and to test the use of AI for representing decision making strategies (research question 3), a set of AI models will be developed and the conclusions will be discussed with the decision makers. Finally, in order to test the use of simulation to identify and assess the decision making policies (research question 4), the modelled strategies will be assessed using the simulation model.

From the above it is clear that, in order to validate the findings, an amalgam of quantitative and qualitative validation approaches will be used (Checkland 1981). To minimise the risk of failing validation checks, an incremental validation approach is used. The implementation of each stage of the methodology is validated soon after its completion and not once the whole methodology has been implemented.

3.3 Research boundaries and limitations

In the light of the method used to address the research questions, it is clear that the approach has certain limitations which set the boundaries in the scope of the research.

As already pointed out, using a case study the conclusions of the research are relevant only for a certain type of application. More specifically, applying the research programme in a production facility of an automotive engine assembly plant, the validity of the conclusions regarding the applicability and the benefits of the research programme in industry are perhaps limited to only these types of operations.

Although the research programme involves the development of AI models which represent and replicate decision makers, it is not the intention of the research to investigate how these systems could replace human decision makers in operations. This is because the models developed during the course of this research are aimed at representing and assessing decision making strategies using a simulation model. As a result, the level of detail which the models contain is sufficient for assessing decision making strategies though perhaps not enough for the development of operational objects that can replace human beings.

The purpose of the research is not to derive optimal strategies but to compare and possibly improve current decision making practices. Only a subset of all the possible techniques is employed to identify the strategies adopted by the decision makers. Therefore, the conclusions on the feasibility of the approach are mainly relevant for those particular selected techniques.

The research programme described in the previous section, aimed at overcoming problems with conventional data collection approaches, proposes to use VIS to collect the required data. However, the use of VIS as a data collection method is not free from problems and it is the purpose of this research to investigate how significant are those problems (section 3.1- research question 2). It is anticipated that four specific difficulties might arise. Firstly, the model needs to contain and report all the key attributes in the decision making process. This probably requires a very detailed model, which could be time-consuming to develop and would have onerous data requirements. A second problem is the need to involve the human decision maker in entering decisions in the model. A very large number of example decisions may be required to obtain a full set of data, which in itself could also be time-consuming. A third problem, known as the gaming effect, as highlighted by O'Keefe & Pitt (1991) and Wielinga & Breuker (1988), is whether the human decision makers are likely to take realistic decisions in a simulated environment. It is quite likely that they will take greater risks, as there are no real consequences resulting from their decisions. Fourthly, the statistical analysis may present difficulties due to the limited data that can be collected using a simulation model.

Finally, given the scale of the production facility, as will be explained later during the implementation of the research, in order to keep the amount of work manageable, the research focuses only on a specific segment of the production line and only on unplanned maintenance activities.

3.4 Conclusion

In the previous sections of this chapter, consequent upon the literature review, the aim of the research has been refined and the research objectives and research questions have been determined. From the research questions, it has been concluded that the purpose of the research is to contribute to the provision of empirical evidence for evaluating a methodology for capturing efficient decision making using simulation and AI. On this basis, it was decided that the research programme, as part of its first research objective, should involve the formation of a methodology which will be tested using a case study. The following chapter describes the formation of this methodology.

Chapter 4

Formation of the Knowledge Based Improvement methodology

As part of the first objective of the research (section 3.1) this chapter describes the methodology and how it was formed. Section one describes the relevant theoretical underpinnings for the methodology. This same section explains how the previous research, extensively discussed in Chapter 2, has influenced the development of the methodology. In section two, the stages required in a methodology for understanding and improving decision making are outlined, based on theoretical developments by previous authors. Having identified the required stages, the methodology known in this research as Knowledge Based Improvement (KBI) is formulated in section three. The chapter concludes in section four by summarising the main benefits and the key aspects of the methodology.

4.1 The theoretical foundations of KBI

Previous research has contributed to specific stages of the process required for understanding and improving decision making using simulation, but it has not formed a complete modelling methodology (Curram 1997). Despite the lack of a step-by-step approach, the research undertaken by previous authors has provided most of the theoretical insights required to support the formation of a practical methodology. Hurrion (1980) has demonstrated the power of VIS as a gaming tool and he has recommended the use of VIS for problem understanding. Liang et al (1992), Curram (1997) and Flitman & Hurrion (1987) have shown that VIS can be used as a tool for data collection. As explained in section 2.2.1, the authors propose the use of VIS as an approach for resolving problems with conventional data collection techniques. Curram (1997) has shown that representing individual behaviour in simulation requires modelling individual decision making. Reference was made in section 2.3 to the modelling stage in which he proposes the use of AI. Flitman & Hurrion (1987) have shown how decision making strategies can be assessed using simulation. Having provided a hypothetical example of assessing decision makers using simulation, the authors explain how the most efficient decision maker can be identified.

4.2 Required stages

On the basis of the above theoretical recommendations, the main stages of a methodology for understanding and improving decision making can be derived by considering the potential challenges involved in the implementation of the above ideas in a practical problem. As explained in section 2.1, the process of modelling decision making in simulation involves three phases. From the previous theoretical developments it is anticipated that the implementation of some of these phases would require several stages.

In the knowledge elicitation phase, previous research has shown that VIS can be used to collect examples of decisions. This conclusion is based on experiments for which the process started from a model (section 2.2 VIS Approach I &II). In an industrial application, the process would start from a problem and, as anticipated by Checkland (1981) and Liang et al (1992), this means that prior to the data collection a stage would be required for understanding the problem and the data that should be collected. The above indicates that the knowledge elicitation phase in a practical methodology would be implemented in two stages. The first stage would involve a problem-understanding process and the second stage would involve data collection.

In the knowledge modelling phase, previous research (section 2.3) has shown that the process would involve the use of AI as a tool to determine the strategy employed by each decision maker. This conclusion is based on non-industrial applications. Due to the technical nature of the process (the derivation of a model or a set of rules from a data set of example decisions), the implementation is not expected to diverge significantly from the way it has already been applied in the non-industrial applications (Curram 1997). This indicates that the knowledge modelling phase can be implemented in one stage and this would involve the process of determining the strategies employed by each decision maker.

In the knowledge representation phase, previous research has shown that this involves a process of linking the decision making models with the simulation (section 2.4). Flitman & Hurrion (1987) have also shown that this phase involves the process of determining the effect of each decision making strategy. Clearly, in an industrial application, this phase would also require a process of recommending improvements based on the comparison of the effect of each decision making strategy (Hurrion 1991). This indicates that in a practical methodology the industrial implementation of the knowledge representation phase would require at least two stages. One stage would require the determination of the effect of each decision making strategy and the other would involve the recommendation of improvements.

4.3 Outline of the KBI methodology

Taking into consideration the above discussion, the KBI methodology outlined in this section has been formed in this research and it is based on the theoretical foundations which were discussed in the previous two sections of this chapter (Flitman & Hurrion 1987, Williams 1996, Robinson et al 1998). It has already been disseminated in a number of articles (Robinson et al 2005, Edwards et al 2004, Alifantis et al 2001) and it consists of the following five key stages:

Stage 1: Understanding the decision making process
Stage 2: Data collection
Stage 3: Determining the experts' decision making strategies
Stage 4: Determining the consequences of the decision making strategies
Stage 5: Seeking improvements

As shown in figure 4.1, in the KBI process (Robinson et al 2005) a VIS is a core element of the methodology. First of all, it supports the problem understanding stage by facilitating interviews and pilot data collections which contribute to the understanding of the problem. It supports the data collection stage by generating the scenarios required to elicit decisions from the human decision makers. It also contributes to the assessment of the human decision making strategies by simulating the system under different policies. Finally, it supports the process of recommending improvements by assessing and scoring new decision making policies. AI and statistical models are also essential parts of the methodology since they can be used to identify the existing human decision making strategies currently applied in the real system. A summary of each of the stages which are shown in figure 4.1 is described in the sections which follow.

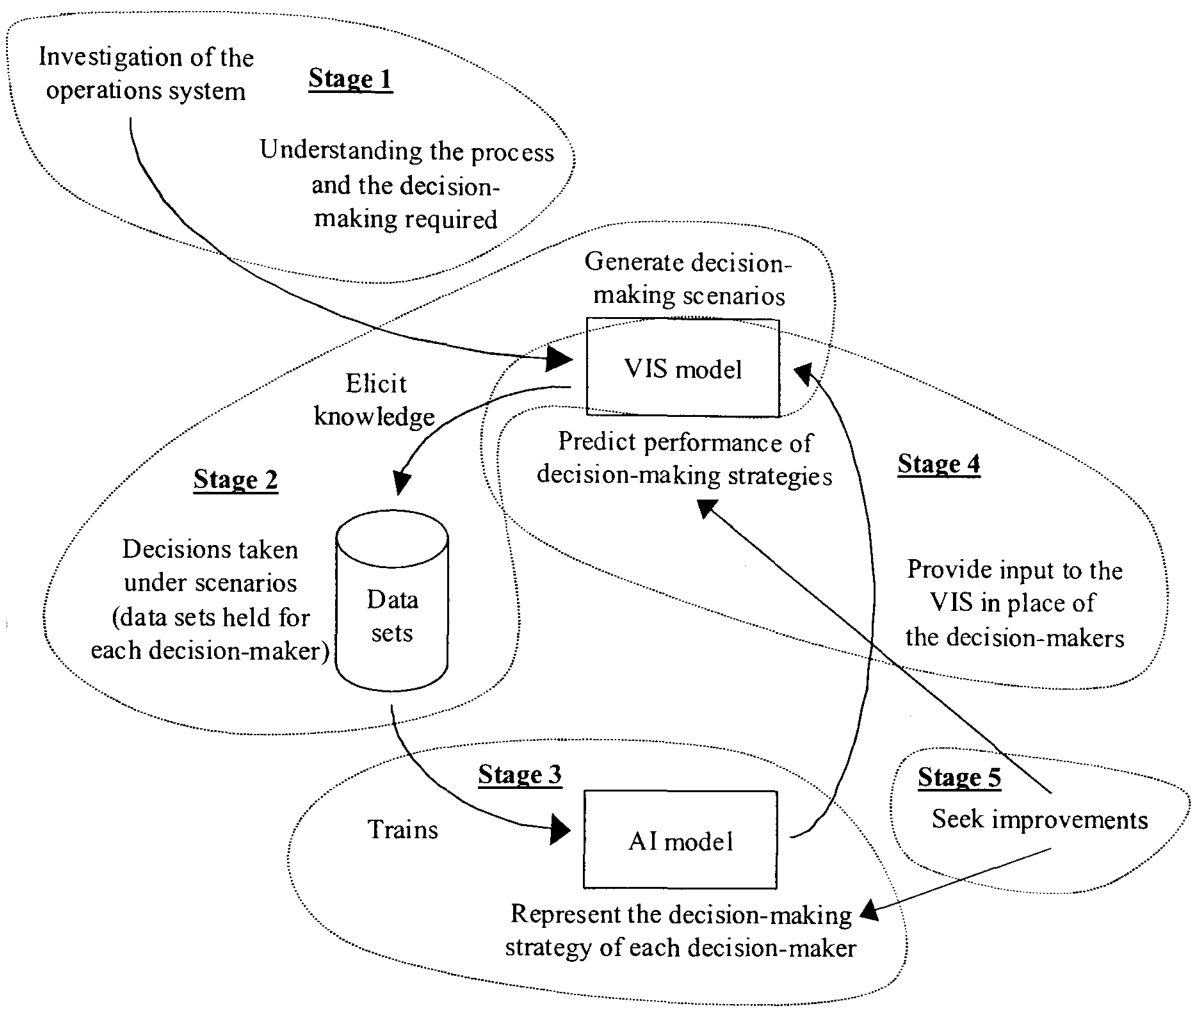

Figure 4.1: KBI stages and process - Robinson et al 2005

4.3.1 Stage 1: Understanding the decision making process

The first step in determining the experts' decision making strategies is to identify the component parts of the decision making process: objective, decision variables, decision options, attribute variables and attribute levels. For instance, in a simple maintenance scheduling problem, where the objective is to minimise the delays caused by machine breakdowns through taking various forms of action, if there are two actions that are not mutually exclusive and two engineers who can be asked to act if they are available then, as such, there are four *decision variables*. The first two variables correspond to the actions and the other two to the engineers. Each of them has two alternative *decision options*: the action can either be taken (denoted 1) or not taken (denoted 0); the engineer can either be asked to act (denoted 1) or not asked to act (denoted 0). Assume, for the simplicity of the example, that the decisions are determined taking into account an estimate of the repair time and the type of fault. In this case, it is clear that there are two *attribute variables* in the decision making process. The first attribute can take the value of any real number that represents the estimated repair time. The second attribute (type of fault) can take values that represent the code of any particular fault. The range of estimated repair times and number of fault codes define the *attribute levels*.

Although conventional interviews and discussion with the decision makers can reveal some information about the decision making process, as demonstrated by Lightfoot (1999), usually the human expert cannot explicitly identify and list the decision making components. In order to do so, the modeller should observe the human experts as they take decisions. In addition, to build a complete model of the decision making process, the modeller may need to make assumptions by considering other rational decisions which may be taken by the decision maker. These assumptions can be tested using VIS to facilitate discussions and small pilot data collection, where the decision makers might be asked to determine whether specific decision situations are realistic.

A decision making process can be represented as two row vectors $[\boldsymbol{\delta}_{i,j}, \boldsymbol{\alpha}_i]$. The first vector $\boldsymbol{\delta}_{i,j}$ corresponds to a decision taken at time i, with each element representing a decision variable δ. The second vector $\boldsymbol{\alpha}_i$ corresponds to the attributes of the decision at time i, with each element representing an attribute variable α. In the context of the simple maintenance scheduling example described above, the decision making process can be represented as follows:

$$\boldsymbol{\delta}_{i,j} = \Phi_j(\boldsymbol{\alpha}_i) \qquad i=1,2,3,\ldots\nu \text{ and } j=1,2,\ldots\mu \qquad (4.1)$$

Where:

$$\begin{aligned}\boldsymbol{\delta}_{i,j} &= [\delta_{1,j} \quad \delta_{2,j}]\\ \boldsymbol{\alpha}_i &= [\alpha_1 \quad \alpha_2]\end{aligned} \qquad (4.2)$$

The subscript i indicates the time at which the decision was taken and the subscript j indicates the human expert who took the decision. The function Φ_j represents the decision making strategy of the individual expert. It represents the concept of a complete correspondence that resides in the mind of the individual decision maker and associates every possible decision situation which can occur in the system with a specific set of actions – decisions. For the derivation of this correspondence, a complete list of decision situations which can occur is required along with the associated decisions that the decision maker would take. In practice, such a relationship cannot be obtained, since the population of decision situations is usually infinite (Johnston & DiNardo 1997). As a consequence, the attribute - decisions relationship which represents each individual's decision making strategy can only be estimated using a sample of attributes and decisions.

To mark the fact that the decision making strategy which can be identified is only an approximation of the actual relationship that links the attributes with the decisions, the decision model of an individual decision maker, in the simple decision making process outlined above, can be represented as in expression 4.3. Matrices **A** and $\mathbf{D}_j$ contain n records of vectors $\mathbf{A}_i$, $\mathbf{D}_{i,j}$ and represent the sample of attributes and decisions pairs which are used to estimate the relationship Φ_j. The function f_j represents the estimated relationship which is derived by calibrating quantitative models using the sample **A**, $\mathbf{D}_j$.

$$\mathbf{D}_{i,j} = f_j(\mathbf{A}_i) \qquad i=1,2,3,\ldots n \text{ and } j=1,2,\ldots m \qquad (4.3)$$

Where:

$$\begin{aligned}\mathbf{D}_{i,j} &= \begin{bmatrix} d_{1,j} & d_{2,j} \end{bmatrix} \\ \mathbf{A}_i &= \begin{bmatrix} a_1 & a_2 \end{bmatrix}\end{aligned} \qquad (4.4)$$

4.3.2 Stage 2: Data collection

Having identified the decision components, as explained earlier in this section, as a next step in determining the decision making strategies KBI proposes a data collection of examples of decisions from each expert j. Each example i in the data set should include the value of each decision and attribute variable. The data set should have the form of the two matrices described above: $\mathbf{D}_j$ and $\mathbf{A}$. $\mathbf{D}_j$ represents the decisions made by decision maker j under specific attribute values (identified in $\mathbf{A}$). Each row of the matrix $\mathbf{D}_j$ corresponds to the row vector $\mathbf{D}_{i,j}$, that is, the decisions taken at time i. Each column in the matrix $\mathbf{D}_j$ corresponds to a decision variable. Each row of the matrix $\mathbf{A}$ includes the attribute values at a particular decision point i. Each column corresponds to an attribute variable. For example, in the simple decision making process outlined above, the data set to be used in determining the decision making strategy of expert j should have the following form:

$$\mathbf{D}_j = \begin{bmatrix} d_{1,1} & d_{1,2} \\ . & . \\ d_{i,1} & d_{i,2} \\ . & . \\ d_{n,1} & d_{n,2} \end{bmatrix} \qquad \mathbf{A} = \begin{bmatrix} a_{1,1} & a_{1,2} \\ . & . \\ a_{i,1} & a_{i,2} \\ . & . \\ a_{n,1} & a_{n,2} \end{bmatrix} \qquad (4.5)$$

One method of collecting these data would be through observation of the experts at work. This, however, would be extremely time-consuming, particularly if the elapsed time between decision points is long. It would also be difficult to record the full set of many attribute values at a specific moment in time and, because the values are likely to change continuously, inaccuracies would occur if there were any delay. As a result, the methodology preliminary proposes the use of a VIS (Liang et al 1992, Flitman & Hurrion 1987). Adopting this approach, different decision makers can be presented with the same series of decision situations. The expert interacts with a visual simulation of the system in question. The simulation model stops at a decision point and reports the values of the attribute variables. The expert is then prompted to enter his/her decision into the model. The model records the value of each decision and attribute to a data file. As a result, a set of values for the matrices $\mathbf{D}_j$ and $\mathbf{A}$ are collected.

4.3.3 Stage 3: Determining the experts' decision making strategies

Having collected a series of examples using the VIS, the next step that KBI proposes is to use the data in the matrices $\mathbf{D}_j$ and $\mathbf{A}$ to determine the decision making strategies f_j of the individual experts. Amongst other approaches, a decision making strategy can be represented by the use of a decision tree (O'Keefe 1986, O'Keefe & Roach 1987, Doukidis 1987, Doukidis & Paul 1985, Abdurahimman & Paul. 1994), a neural network model (Curram 1997, Liang et al 1992) or a logistic regression equation (Malhorta et al 1999, Carvalho et al 1998); a separate model can be constructed for each decision maker.

Expert systems software is capable of constructing a decision tree from a set of examples, such as those collected via the VIS. One such method for constructing a decision tree is Quinlan's ID3 algorithm (Quinlan 1979); see, for example, Mingers (1987). The algorithm prioritises the attributes according to the degree to which they match the data set with the correct decisions. Neural network models (McClelland & Rumelhart 1988) can be used to develop a system of equations which can predict decisions once the parameters of the equations have been calibrated, using a set of examples such as those collected via VIS. Finally, logistic regression can also be used to represent decision making strategies and to predict decisions once the parameters of the regression equation have been calibrated.

Taking into account the capabilities of the above techniques, KBI proposes to apply one (or more) of these techniques to model decision making. This process will enable the derivation of the decision making strategy f_j employed by each individual decision maker j.

4.3.4 Stage 4: Determining the consequences of the decision making strategies

Having determined the decision making strategies, that is, a decision model f_j for each expert j, the next step that is proposed as part of KBI is to assess and compare the performance of each strategy. The ultimate performance measure in most manufacturing facilities is the level of throughput. This means that, in the case study used in this research, each expert can be assessed on the basis of the throughput which is achieved in the simulation, when the decision making process is controlled using his/her decision making strategy. To predict the throughput, conditional on each human expert, the VIS can be linked with the expert systems software (or, indeed, the neural network model or the logistic regression equation).

The decision model can be used in place of a decision maker to interact with the simulation (as proposed by Williams 1996, Flitman & Hurrion 1987 and Curram 1997). With this approach, each time that the simulation reaches a decision point, the simulation stops and the software containing the decision model is invoked. The value of each decision attribute is passed from the simulation to the software with the decision model. In turn, the software with the decision model returns the values of the decision variables to the simulation before the simulation run continues.

When the simulation has reached the end of the run, the throughput of the production line provides an indicator of the performance of the expert whose decision model was used during the run. Employing this approach and running the simulation under each expert's decision making strategy for a number of replications, KBI enables the most efficient strategy to be found by comparing the output from each run.

Of course, having identified the most efficient expert does not mean that the most efficient strategy has been found, since there is no guarantee that the best current strategy is the optimal one. Although the best strategy may not be optimal, it can still be used to train less efficient decision makers, providing improvements in overall performance.

4.3.5 Stage 5: Seeking improvements

The last stage that the KBI methodology proposes is to use the decision making strategies of the most efficient experts as a starting-point to search for an improved strategy. The search could be made informally by combining strategies and by making incremental changes. Alternatively, heuristic search methods could be implemented in order to seek for improvements. In each case, the alternative strategies can be tested by running them with the simulation model in order to determine their effectiveness.

4.4 Conclusion

The formation of the KBI methodology, in the previous sections of this chapter, enables the development of a framework to test an approach for understanding and improving decision making. In these earlier sections, as part of the first objective of the research, the methodology that will be tested in this research has been formed based on the theoretical findings of previous research. In the subsequent chapters, as part of the second objective of the research, the above methodology will be applied to a real problem (section 1.3) which forms the case study of the research. As a result of such an application, the methodology will be refined and extended.

Chapter 5

Evaluation of stage 1 of the Knowledge Based Improvement methodology

Chapter 4 described the formation of the KBI methodology. As part of the second objective of the research the purpose of this chapter is to test and evaluate the first stage of the KBI methodology. As discussed in section 4.3.1, this stage involves the definition and formulation of the decision making problem and specification of the data requirements necessary for designing the data collection. This chapter describes the qualitative techniques that were used to define and formulate the decision making process as an OR problem. The chapter ends by summarising the main methodological conclusions from the implementation of the first stage of KBI.

5.1 Steps of stage 1 of the KBI methodology

The stage 1 of the KBI methodology is an iterative process which involves five steps and, as shown in figure 5.1, it combines problem-structuring with pilot simulation-based data collections. During the implementation of the first three steps as explained in sections 5.2, 5.3 & 5.4, qualitative techniques and VIS are applied in order to identify the main elements of the problem. To identify the decision variables $\mathbf{D}_{i,j}$ and attributes $\mathbf{A}_i$ which are required to represent the human decision makers, a conceptual model is developed during the fourth step (section 5.5) of the problem-understanding process. This model is formulated iteratively, based on the scope of the modelling exercise and taking into account the feedback from the interviews with the people involved in the decision making process. During this stage, the level of detail in which the problem is represented in the model is determined by deciding the attributes and the decisions which should be included in the conceptual model.

Based on the attributes and decisions included in the conceptual model, the data requirements are specified in the final step (section 5.6). This step takes into account the understanding of the problem and its constraints, the number of decision makers who should be involved in the data collection and the availability of the experts.

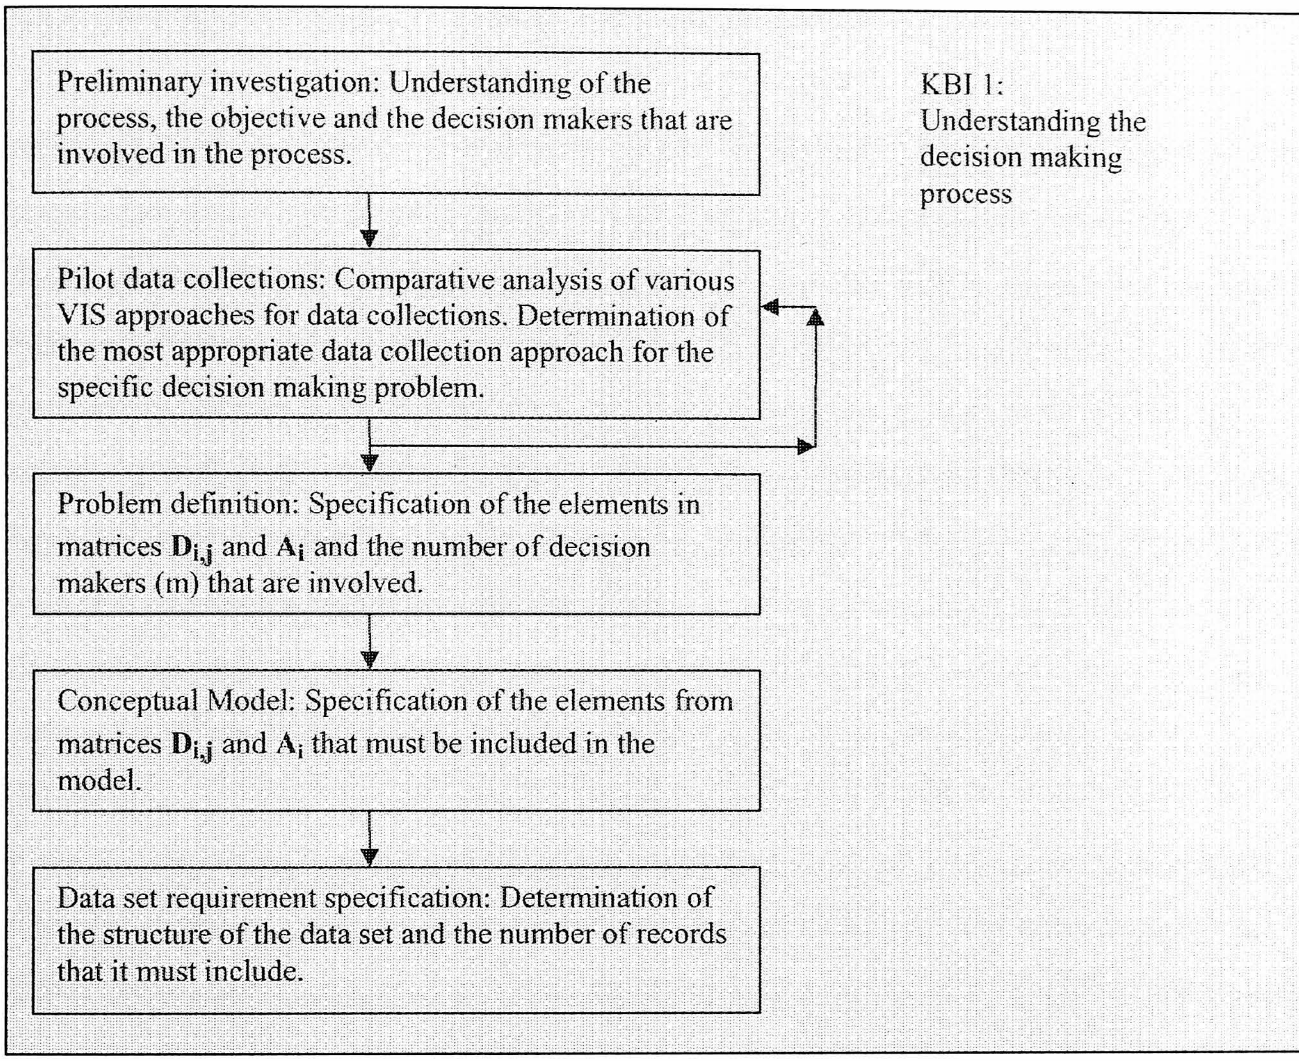

Figure 5.1: Understanding the decision making process (KBI Stage 1)

5.2 Preliminary investigation

In order to gain an understanding of each of the elements of the decision making problem, a series of informal discussions along with some action research took place during the first step of stage 1 of the KBI methodology.

Firstly, a discussion facilitated by the author and his supervisor took place. In this unstructured and relatively informal discussion, all the employees involved in the decision making process were invited and attended. With the aim of building up a rich picture and to gain an understanding of all the possible aspects of the maintenance decision making process, the theme adopted for the discussion was rather general (Checkland 1981 and Robinson & Stanger 1998). The participants were asked to describe the maintenance process in the production line (section 1.3). During this session the participants had the opportunity to express their own views and to debate how maintenance decision making takes place in different areas of the production line.

Although, with the conclusions from the discussion described above, it was not possible to structure the complete problem, we were able to identify the people who make decisions (each j) and the general process that is followed when a decision is required. In order to enhance our understanding of the decision making process, a three day visit to the factory was arranged. During this visit the author had the opportunity to follow and observe two decision makers in action. This gave the opportunity to ask questions and have informal conversations with other decision makers when they were available. In this visit the author gained an initial understanding of the main decision variables, along with the decision options. In addition, some conclusions were derived about the attributes which the decision makers take into account in order to make decisions.

Overall, the preliminary investigation revealed that the Ford engine plant at Bridgend is one of the main production facilities for the 'Zetec' petrol engine. Consistent with our initial information (section 1.3), the preliminary investigation confirmed that the plant consists of a number of transfer lines which feed the main engine assembly line. In the engine assembly, the line that is modelled in the simulation model, head engine blocks are placed on a palette and pass through a series of automated and manual processes. The palette is a base which is used to protect the head engine during its transfer through the conveyors. It can be seen as a container with an open top which contains the head engine block. There are two parallel decision making processes in the engine assembly line. One for planned machine maintenance and one for unplanned and immediate action machine maintenance. Planned maintenance takes place during the weekend, though the decisions about what must be done are taken during the week. Decisions relating to unplanned maintenance are taken when a machine breaks down and the decision maker has to make a decision at short notice.

The production line is broken up into sectors and a Group Leader is assigned to each sector. Given that there are three shifts per day, there are three different human decision makers per sector (DM1, DM2, DM3). The person who is informed immediately when a machine breaks down is the Group Leader of the maintenance team of the specific sector. The Group Leader, who is the decision maker in charge, is responsible for deciding what must be done in order to minimise delays in the production process. Deciding what action must be taken and who should be involved in this action are two aspects of the decision that the Group Leader should address when a machine breaks down. For the action he has at least two options available: the first, known as repair immediately (RI), involves the repair of the machine. The second, known as stand by (SB), involves manual operation of the machine by an operator and delay of the repair at least until the end of the shift. For the decision on who should implement the action there are five engineers, member of the maintenance team of the specific sector. In terms of deciding what actions must be taken, based on the preliminary investigation, it was concluded that a variety of attributes of the system are taken into account. Amongst others these include the type of breakdown and the general status of the production process.

The reporting system includes a pager that reports the machine number and the type of fault in the machine which has broken down. In addition, details about the breakdown are recorded in the central Ford database system known as Plant Operating Systems Monitoring (POSMON). In the specific production line there are six teams and the production line is divided into six sectors. For reasons that will be explained in detail in section 5.5, the research deals with only one sector.

5.3 Pilot data collections

The second step of stage 1 of the KBI methodology described in this section involves a series of VIS-based pilot data collections which increased our understanding of the decision making process and contributed to the definition of the decision making problem (section 5.4). During this process, that is essentially iteration between KBI stage 1 and 2, four methodologically different pilot data collections took place.

The conclusions from each of them confirmed the key decision variables and attributes of the maintenance decision making problem and highlighted the main practical difficulties involved in the process of collecting decisions using a VIS-based data collection approach.

A slightly different data collection method was applied in each experiment. The data collection approach was progressively improved, building on the experience and the conclusions from the previous sessions. In this section, the details and the conclusions from each pilot data collection are described, revealing how the approach evolved to form the one that was used for the main data collection.

5.3.1 The first pilot data collection

The first pilot data collection took place in the early stages of the project, while our understanding of the decision making process was still incomplete. The aim of the first session was to gain an initial understanding of the main attributes that are taken into account by the decision makers when deciding what decision must be taken.

Pilot 1: Strategy

In the first pilot data collection a VIS prototype was used for collecting the decisions. The prototype was constructed using the simulation model that was developed prior to the research by internal Ford analysts (section 1.3). During the first pilot data collection, the simulation model was running in visual interactive mode and the decision maker was watching the progress of the simulation on a visual display until the point when a machine broke down. Once the breakdown had occurred, the simulation halted and the decision maker was asked to input a decision. When a decision action had been taken, the simulation model continued until the next breakdown event, when the above process was repeated.

Pilot 1: System

Given that the aim of the session was to gain a very basic understanding of the elements of the problem, only one decision maker was involved. To record the decision maker's inputs a dialog box was used. The user interface (UI) for the first data collection was developed in Witness, using the limited UI development functionality which the package provides. The decision situations that were presented to the user were selected by sampling from a uniform distribution of historical machine breakdowns.

As a means of communicating the decision situation to the decision maker, the following five attributes of the system were included, along with other information, in the visual display at each decision point: 'current time in the factory', 'parts waiting for processing behind the broken machine', 'resource availability', 'other machines that are broken down at the same time', 'type of fault'. The first four attributes were generated from the simulation model, whilst the type of fault was taken from a database table containing information about specific breakdowns which had happened in the past. As explained in detail in section 6.2.3, this table was used as a trace for generating the details of the breakdowns incidents which occurred during the simulation run. A new record from the table was used each time that a breakdown event occurred in the simulation.

The structure of the prototype system that was used for the first pilot data collection is represented in figure 5.2. In this version of the prototype, the decision maker (5) interacts with the simulation (3) that was running during the data collections through a user interface (4 & 6). This interface consists of a visual display of the simulation (4) and a dialog (6) that is invoked when the simulation reaches a decision point. The interface is used to establish a communication protocol between the simulation and the user. With an initial front end (2) that appears only once in the beginning of the session, the OR analyst is allowed to specify various settings and initial conditions for the simulation run (warm-up period, speed and type of breakdowns that should be reported). The breakdown incidents database (1) is an input in the system which forms a trace of breakdown details that is used during the run (1). The attributes - decisions data set (7) is an output of the system and is used to record the attributes and the decisions that were taken by the decision maker during the pilot data collection.

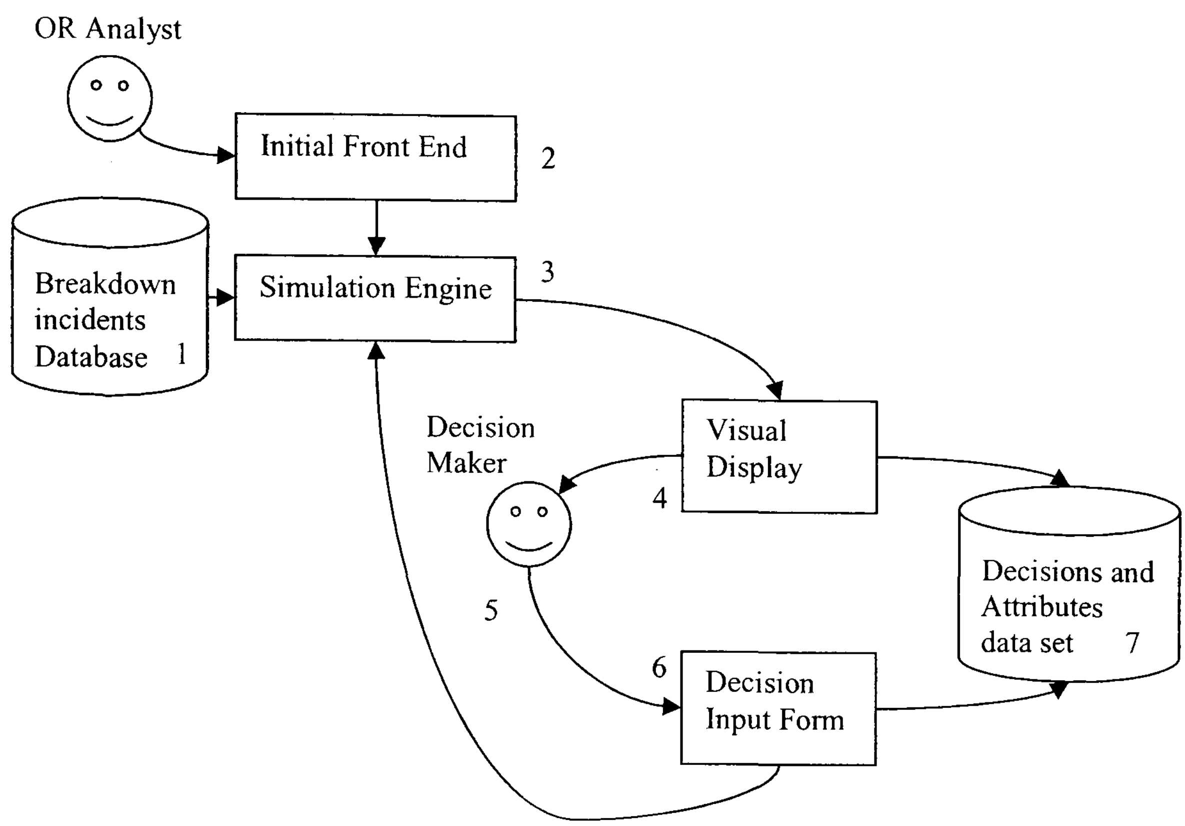

Figure 5.2: Structure of the first pilot data collection

Pilot 1: Outcome

As a result of the discussion facilitated from the decision situations which were generated in the VIS during the first pilot data collection, it was found that in the real system additional attributes are taken into account in order to take maintenance decisions.

One of the attributes taken into account, when deciding what action must be taken in the event of a machine breakdown, is the physical condition of the machine and the physical condition of the part which was last to be processed from the machine. Combining this piece of information with the type of fault reported from the diagnostic system, the decision maker estimates the repair time before he takes his final decision. This estimation is then combined with other attributes in order to take the final decision about what action must be taken. Taking into account that the physical condition of the machine is a visual attribute which cannot be simulated in a model of this scale, to be able to use VIS for knowledge elicitation purposes (which helps to address one of the research questions) it was decided to simplify the representation of the decision making process. This was achieved by directly reporting the estimated repair time to the decision makers. In order to provide this information, it was decided that sampling from an appropriately calibrated statistical distribution should applied.

The simplification of the representation of the decision making process, as a result of not modelling the stage which involves the estimation of the repair time (by the decision makers), has an impact on the outcome of the research, since it does not allow the identification and assessment of the strategies used by decision makers to forecast the repair time. This simplification, as will be referred to in the conclusion, highlights one of the limitations in the use of VIS for knowledge elicitation.

5.3.2 The second pilot data collection

In the second data collection, 25 decision situations were presented visually and interactively to the decision maker using the interface shown in figure 5.3.

Pilot 2: Strategy

Taking into account the conclusions from the first pilot data collection, the version of the prototype used in the second experiment informed the decision maker as to approximately how long it would take to repair each breakdown of each machine which was involved in each decision. In addition, having found from the first pilot data collection that the decision maker had difficulties in understanding the decision situation, due to the richness of the information that the visual display provided, for the second pilot data collection it was decided to change the method for communicating each decision situation to the user. The key information which the visual display revealed at the time of a machine breakdown was represented in a dialog (figure 5.3) that was invoked in the case of a breakdown event. In the same dialog, the decision maker was allowed to input his decision. Once a decision had been recorded, the dialog closed and the simulation process continued until the next breakdown incident, when the above process was repeated. The UI for the second pilot data collection was developed in Visual Basic and was invoked from Witness each time a breakdown event occurred.

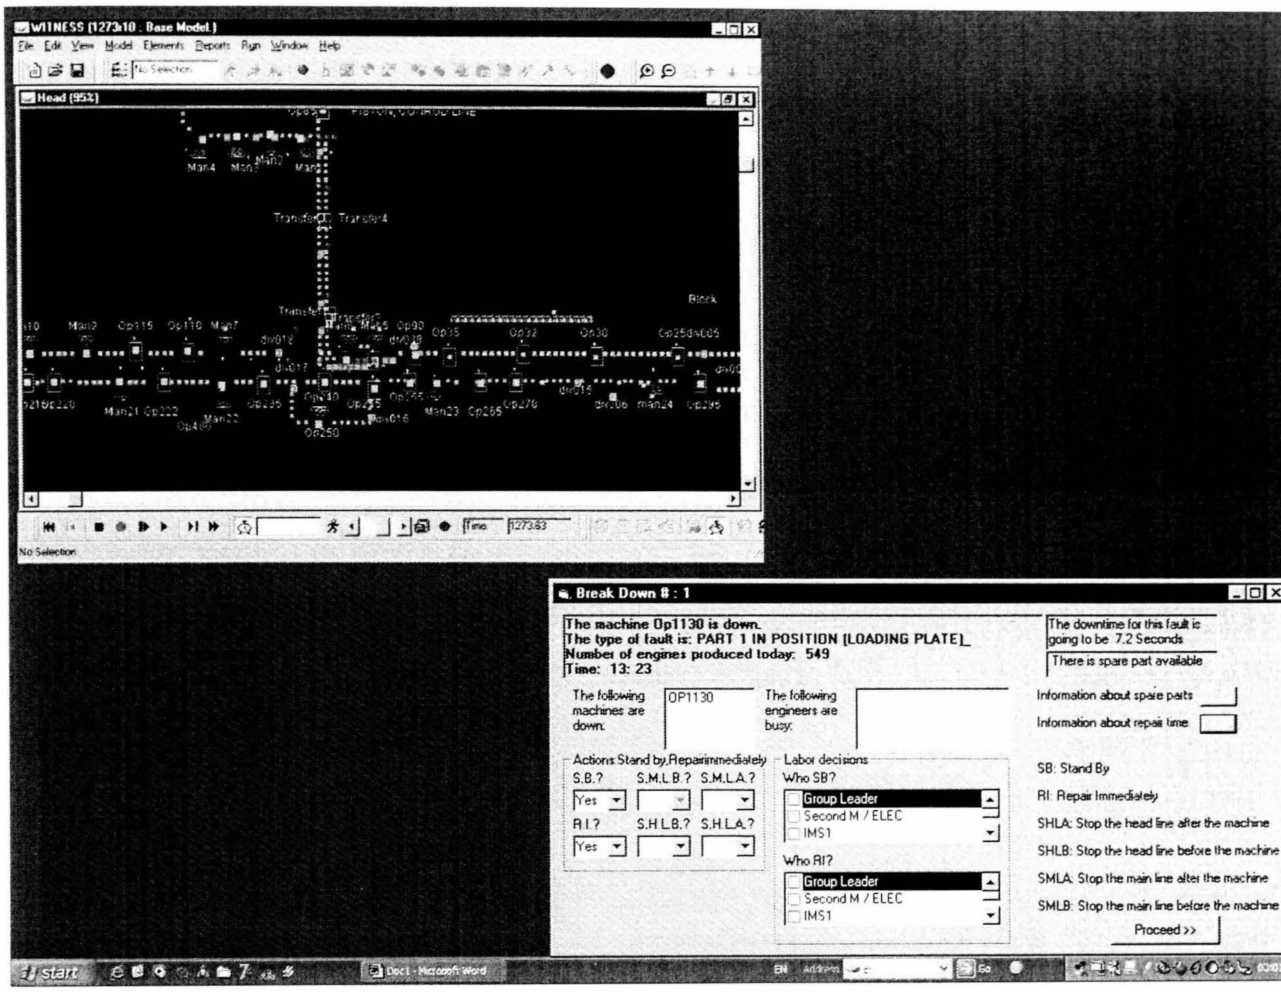

Figure 5.3: Second pilot data collection

As figure 5.3 shows, the user interface in the second pilot data collection reports to the decision maker information about a wide range of attributes (figure 5.3 - bottom right form).

For the second data collection, the scenarios were sampled from a uniform distribution and the simulation was running interactively during the data collection. The decision maker had to watch the simulation progress, waiting for the next breakdown incident to occur. An additional decision variable was included in order to capture the decision about which engineer should implement the desired action. In addition, two extra decision options were included to incorporate the option of the decision makers being forced to stop part of or even all the production line when a breakdown occurs.

Pilot 2: System

The structure of the system was not very different from the configuration of the prototype used for the first pilot data collection and, therefore, the main difference in the second data collection was the representation of additional decision elements and the inclusion of some extra attributes which helped the decision maker to take more informed decisions.

Pilot 2: Outcome

Analysing the data set after the second data collection, it was found that the collected decisions did not represent equally each decision option. 20 out of the 25 times that a decision was required, the decision maker decided to repair the machine immediately. Comparing the attribute data of the situations that were presented to the decision maker with a larger sample of decision situations that were generated by running the simulation for a little longer, it was found that one possible reason to explain the lack of variation in the decisions was that all scenarios presented to the decision maker during the second data collection involved breakdowns with relatively short repair time.

On the basis of the above observation it was concluded that, in order to avoid collecting a data set with stereotype responses, a structure must be imposed on the decision situations that were presented to the decision makers. It was also concluded that when a decision is made, apart from what actions must be taken and who should be involved in taking that action, complementary actions are also considered (for example 'ask production manager' or 'plan off shift repair'). Finally, in this data collection it was found that the decision makers, before they make a decision, might be interested to know about the frequency of a specific breakdown in a specific machine but not about which other machines have also broken down at the same time.

5.3.3 The third pilot data collection

Pilot 3: Strategy

Taking into consideration the above conclusion, for the third pilot data collection a non-randomised decision situation generation design was applied. This means that a filter was used during the run, causing the simulation to stop on a breakdown event only if certain conditions specified in the filter were met. Having identified from the discussion with the decision maker that the estimated repair time is one of the most important attributes when determining what decision must be taken, it was decided to differentiate the situations that were presented to the user on the basis of this attribute. Thus, the rules included in the filter allowed the simulation to stop during the next breakdown incident only if the estimated repair time was at least 50 % different from that reported in the previous incident when the simulation had stopped and user involvement had been required.

As in the previous experiments, the data collection was 'on line'. The simulation was running during the session and the decision maker who was involved in the specific session had to watch the progress of the simulation process, waiting for a decision point to be reached. Using the filter mentioned above, the situations that were presented were systematically sampled, forcing the simulation to generate an amalgam of situations which represented a wide range of estimated repair times.

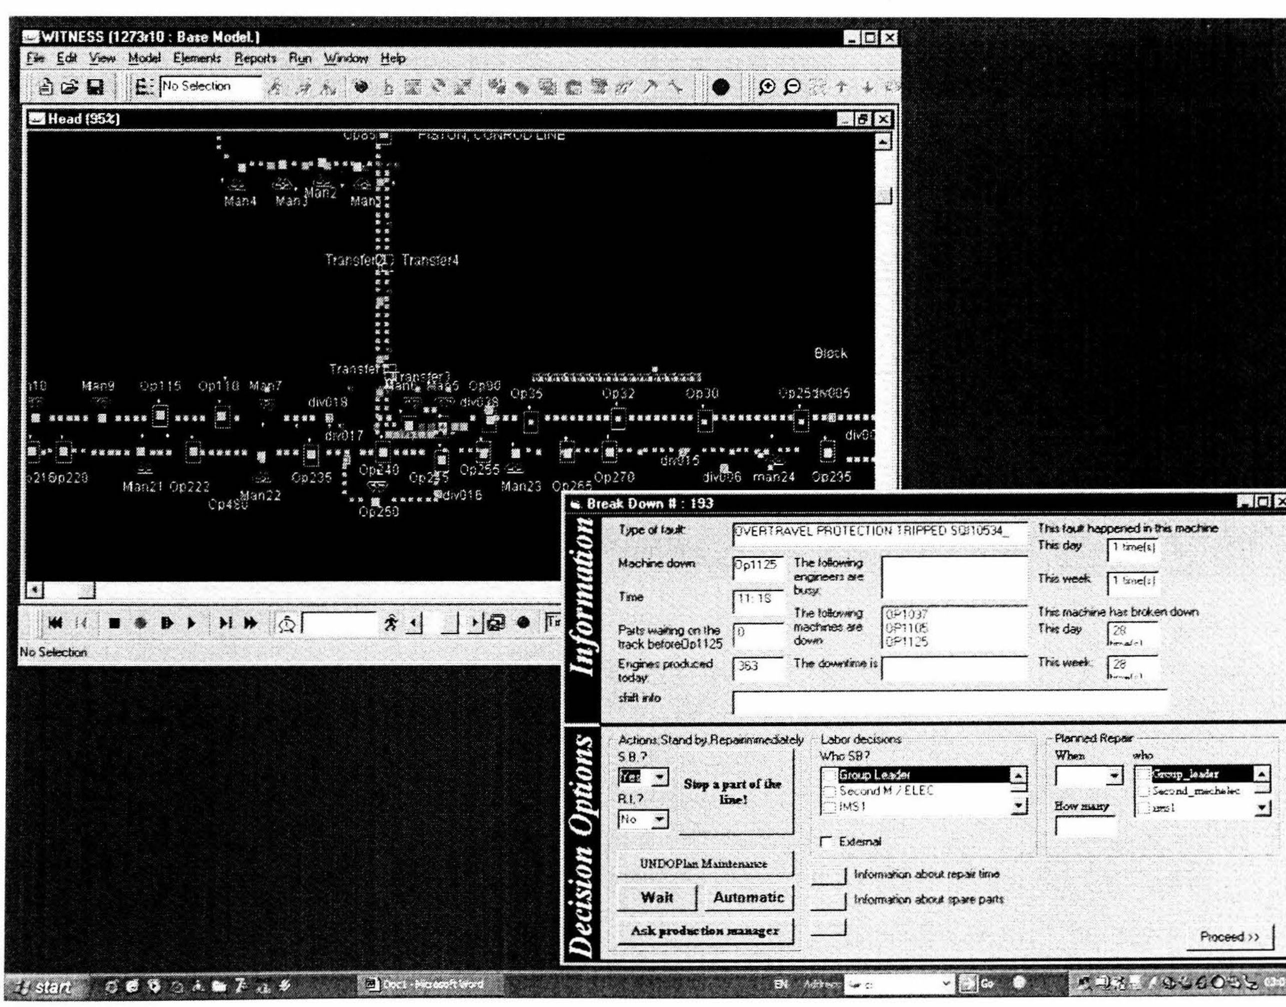

Figure 5.4: Third pilot data collection

Taking account of the conclusions from the second pilot data collection, a number of additional elements of the decision and some additional attributes were included in the system that was used for the third pilot data collection. From the informal discussion that took place during the second pilot data collection it was found that, when a decision is made, complementary actions are at times taken to support the main decision. The options that are available for this complementary aspect of the decision include the request for planned maintenance - which must take place either at the end of the shift or during the weekend - and the option to seek authorisation from the production manager.

Pilot 3: System

In order to allow the decision maker to make more realistic decisions in the prototype (figure 5.4) which was prepared for the third pilot data collection, the above options were included as part of an additional decision variable.

In addition, from the debate that the VIS facilitated during the second pilot data collection, it was found that the decision maker was interested in knowing how many times a specific breakdown had occurred. He was also interested to know how many times the machine had broken down on the specific day and in the specific week. Bearing the above questions in mind, in the prototype that was prepared for the third pilot data collection it was decided to report in the UI information on how many times the machine had broken down in the last week and on the last day. It was further decided to report how many times the specific breakdown had happened during the last week and on the last day in relation to a specific machine. Finally, for completeness, it was decided to report in the user interface the current simulated shift being run by the factory.

The structure of the system that was used in the third pilot data collection is represented in figure 5.5. The main differences from the configuration used in the two previous experiments are the application of a filter and the involvement of more than one decision maker. The filter, as is shown in figure 5.5, checks whether the estimated repair time is appropriate to force the simulation to stop. If it is not, the simulation continues until the next breakdown incident without the involvement of the decision maker.

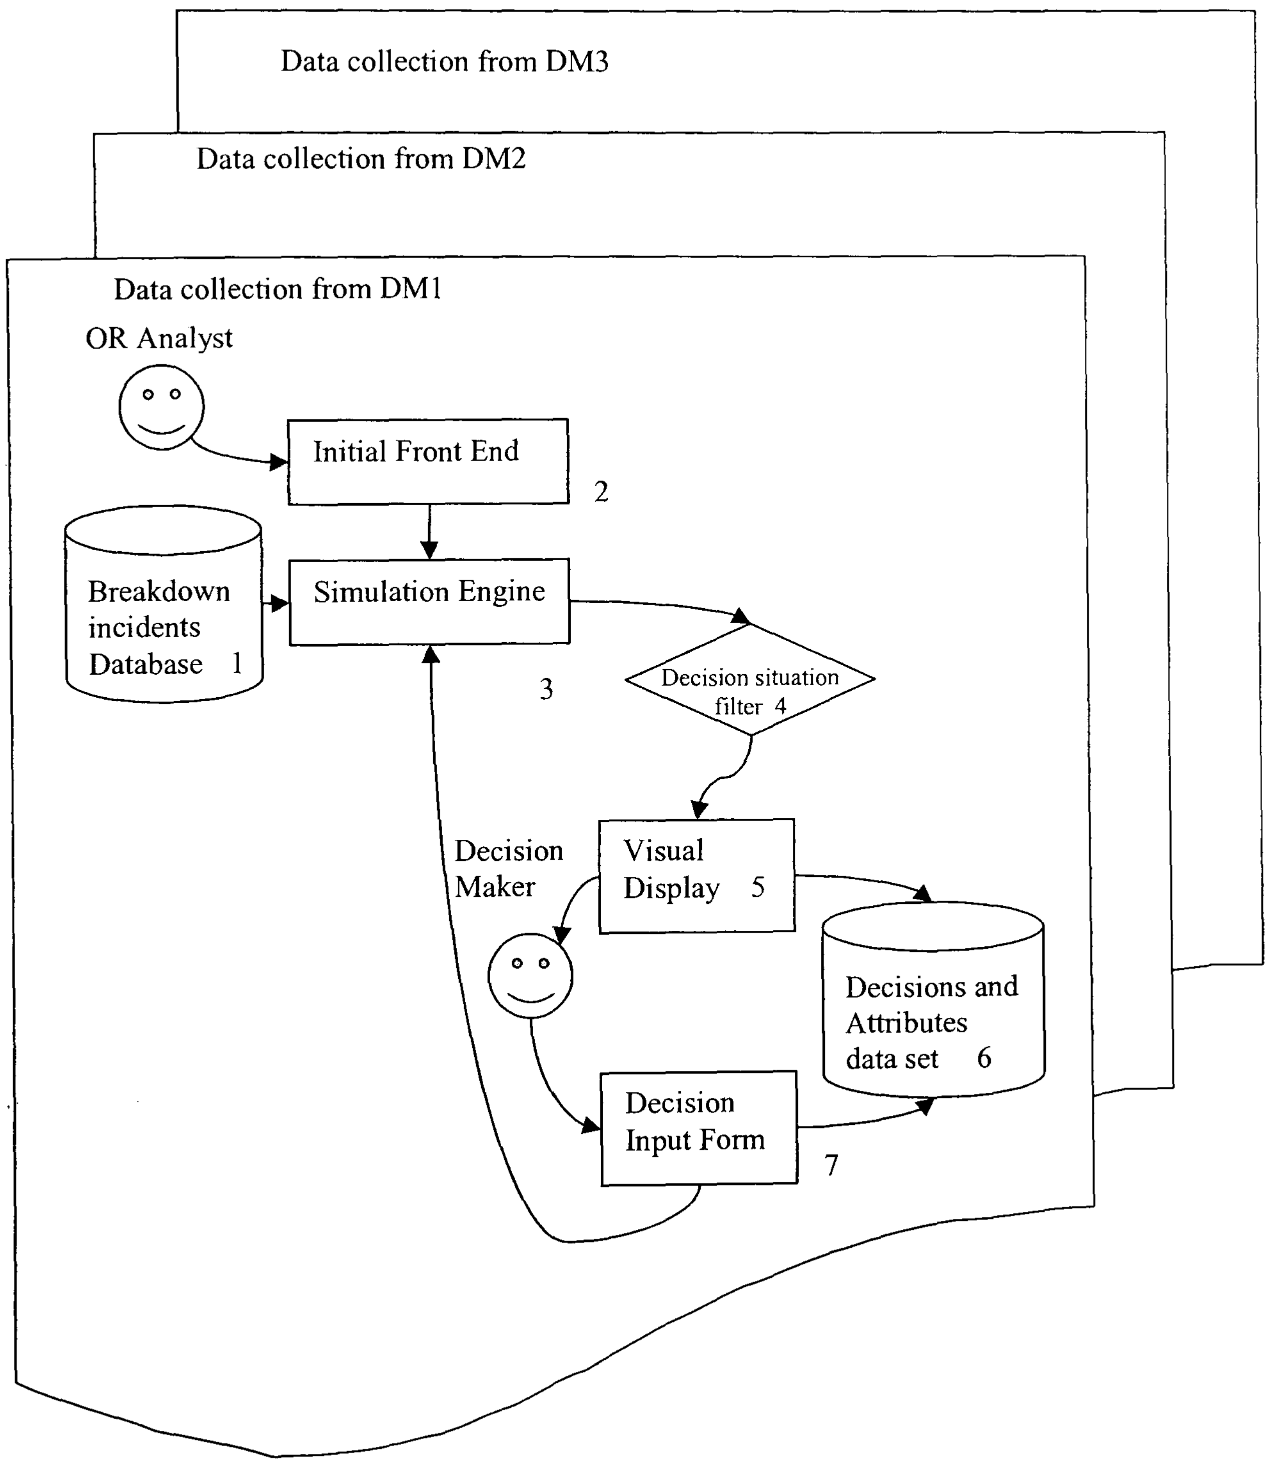

Figure 5.5: Structure of the third pilot data collection

Pilot 3: Outcome

Using the breakdown filtering approach for taking a sample of decision situations, during the third pilot data collection it was possible to collect a set of decisions which include a significant number of SB actions. However, the dominant decision in the data set was still the action RI. With more than 2/3 of the decisions being RI, it was decided that the data set does not represent fully the situations and the actions that the decision making process involves. The above conclusion was reached after interviewing the two production managers in the factory. During the interview they gave reassurances that SB policies are very often applied, since the Group Leaders have received instructions to follow SB process where appropriate and certainly when they think that the repair time will be more than 5 minutes. One of them insisted that SB must account for at least 40 percent of the decisions that are made since 40 percent of the breakdowns that occur in the production line require a repair that, on average, takes at least 10 minutes.

From the discussion which the VIS facilitated during the third pilot data collection, it was also concluded that one of the main attributes that is taken into account but that was missing from the model is the 'number of heads' that are in stock when a machine breaks down. 'Heads' is the main output of the line segment for which the decision makers are responsible. It is the component that feeds the other segments of the line. Clearly, a shortage of this part can cause extensive delays. To minimise the risk of a shortage, a buffer with capacity of storing 200 items is located in-between the two segments of the line and is used to accumulate stock of this component. Indicating this buffer in the visual display of the simulation, the decision makers explained that the stock level in the buffer is key information which helps to predict whether there is going to be any shortage and therefore is a key attribute for deciding whether actions must be taken to accelerate the production process and to increase the stock level. The decision makers said during an informal discussion that there are times when the buffer in which the heads are stored is empty. In these situations they are very likely to decide on SB if a machine breaks down.

In the third pilot data collection it was found that the stock of spare parts is an attribute that is not used frequently since spare parts are usually available. The decision options 'stop the line' and 'wait before decide' are relevant but very rarely considered. The reason for this is that such actions are time-consuming and they have an enormous impact on the supply chain of engines. They do not, therefore, belong to the standard actions that are taken on regular basis.

From the feedback on the data collection session given by the decision makers at the end of the third pilot data collection, it was concluded that watching the progress of the simulation run and waiting for a decision point to come does not support the decision making process, since it is not taken into account when deciding what action must be taken. On the contrary, given the time that the decision makers had available and the time that they had to wait until a decision point was reached in the simulation while watching the simulation progress, it was found that this restricted the number of situations which the decision makers had the chance to see and decide. In the third pilot data collection, three sessions took place and in each session the same set of situations was presented to a different decision maker. 1.5 hours were spent with each decision maker and only 25 instances of decisions were collected from each of them.

5.3.4 The fourth pilot data collection

Pilot 4: Strategy

Taking into consideration the above conclusion, for the fourth data collection, it was decided to review the data collection approach in order to increase the number of decision records that could be collected from the decision makers within the limited time available.

Pilot 4: System

Having concluded from the previous pilot data collection that the dynamic display of the simulation process is not taken into account when the decision makers take decisions, in order to reduce the time that the users had to wait before the next decision point was reached we developed a prototype which does not require running the simulation model during the data collection.

Information about stock levels in the head buffer was added as an extra attribute and the simulation was still used as part of the general data collection approach, but it ran and generated the decision situations prior to the sessions with the decision maker. In order to save the situations of interest, a set of status files was used. This is a functionality provided by Witness and enables the user to save the simulation status (including the visual display) at any given time. It also allows the user to continue the simulation from the point that it was stopped when the status file was saved. Having saved the status of the simulation for each decision situation of interest, the user did not have to wait for the simulation to reach a breakdown event in order to input a decision. During the data collection all that was necessary was to load and present those simulation status files which included the breakdown events of interest.

Although this was a significant reduction of the time that the user had to wait, the process was still time-consuming since Witness can be quite slow in loading the status files for a simulation model of this scale. Figure 5.6 shows the structure of the prototype system that was used for the fourth pilot data collection. As is clear from the diagram, there are multiple copies (5) of the simulation model (3) that are used to present the decision situations to the user. Each copy represents a decision situation that was generated after running the simulation model and saving its status at the point when a machine breakdown of interest had occurred. To collect a sample of decisions and decision situations, the filter used in the third pilot data collection was also applied to generate the status files. This means that the simulation stopped and a file was generated only for a number of selected machine breakdown situations which met the rules applied by the filter.

Pilot 4: Outcome

Despite the use of the simulation status files to present the decision situations, as a result of which we managed to collect more decisions than in any of the previous data collections, once again the sample size did not represent each decision action adequately. RI was by far the most popular decision action for each of the three decision makers who were involved. During the fourth pilot data collection, from the 34 decisions that were collected from the first decision maker only 4 were SB. The second and third decision makers, who also input 34 decisions, did not take any SB action since they thought that the most appropriate decision was to repair immediately for each of the situations that were presented to them. The situations were the same as those that were presented to the first decision maker.

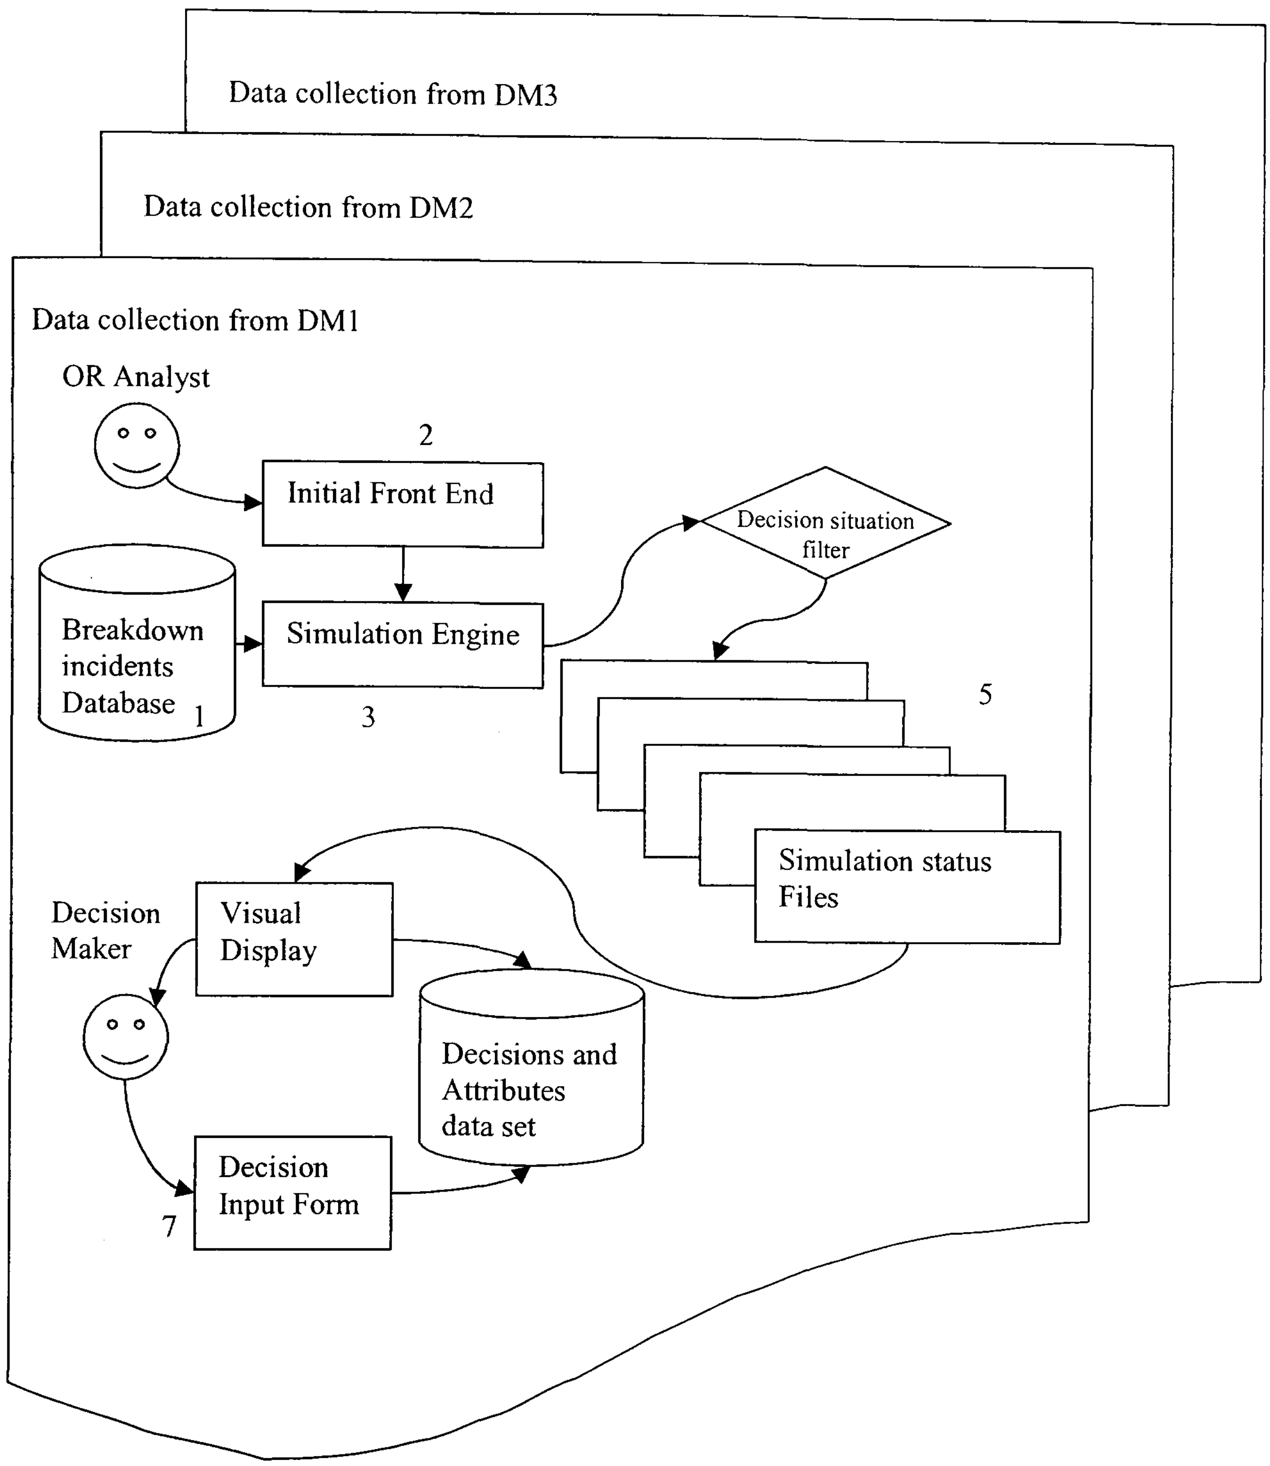

Figure 5.6: Structure of the fourth pilot data collection

When the decision makers were asked to explain why the most of the their decisions were of the same type (RI), they insisted that this is due to the fact that the repair time of the breakdowns that were presented during the data collection was not particularly long, so there was no reason to follow stand by policies. To justify the above statement, they explained that SB is a policy which is usually followed when the time that will be saved is higher than the total time that will be lost. Stand by process, they explained, involves manual processing of the parts as they arrive. Given that manual processing is slower than the automatic processing which the machine performs, deciding on stand by might cause delays due to manual processing. In addition, deciding to stand by means that a stand by station must be set up. According to the decision makers, this requires about 5 to 10 minutes. In order to decide to stand by, the decision makers must be sure that the repair time of the machine - and therefore the delay that this will cause in the production process - is higher than the delay that will be caused by setting up a stand by station and processing the parts manually.

In this data collection, it was concluded that a representative sample of decision situations that occur in a typical day in the production line must include a significant percentage of breakdown incidents whose estimated repair time should be over 20 minutes. This is because a significant number of SB decisions are taken on daily basis and, according to discussions with the decision makers, it appears that the SB-RI trade off becomes practically relevant only in breakdown incidents with an expected repair time of over 20 minutes. Given the high frequency of breakdowns with a relatively short repair time on each machine, the above indicated that stratified sampling was needed for the main data collection. Finally, in this pilot data collection it was concluded that the decision about who should be involved in the action which is decided upon is not necessarily a decision that is determined by the attributes of the system. From the discussion that the data collection facilitated, it was clear that the decision about which engineer should be involved is dependent on the desired action. For instance, an RI action would require the involvement of a skilled engineer whereas an SB action would require the involvement of a semi-skilled engineer.

5.3.5 Conclusion from the pilot data collections

Overall, the pilot data collections contributed to the development of a knowledge elicitation approach that led to the main data collection of the research, as explained in the next chapter. In addition, the conclusion from each pilot contributed to the specification of the requirements and the functionality of the data collection system which was developed to support the main data collection.

5.4 Problem definition

Having undertaken four pilot data collection sessions, our knowledge about the problem situation was significantly improved. From the above discussion of the pilot data collections, it is clear that a maintenance decision making process takes place in the production line. This process is not represented in the simulation model that has been developed in Ford.

A total of at least 18 Group Leaders (one for each sector for each shift) are involved in the decision making process for the whole production line. Each of them is responsible for making maintenance decisions, taking into consideration information collected from the broken-down machine and the surrounding environment.

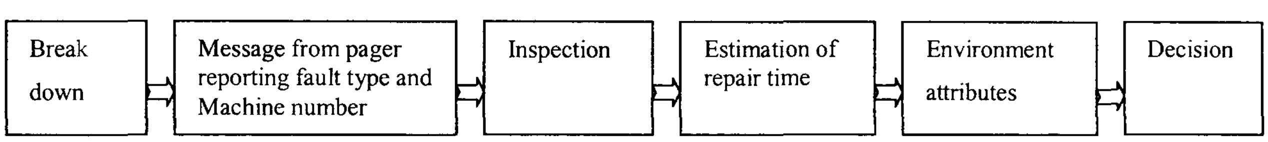

Figure 5.7: Decision-making process

As shown in figure 5.7, in the event of a machine breakdown the relevant decision maker is informed via a pager which indicates a breakdown in the area supervised by him. The decision making process then involves four stages. After the inspection of the breakdown, various attributes of the system are taken into account in order to estimate the repair time and to assess the available decision options. Having decided which option will cause the minimum delay, the decision maker takes the appropriate actions to implement the decision.

Based on our current understanding of the problem, the decision making process for unplanned maintenance involves three main decision elements: actions that must be taken when a machine breakdown occurs (Action Decisions), engineers who should be involved to deal with the situation (Resource Decisions) and complementary decisions that should be taken to support the main action (Complementary Decisions).

Action Decisions

Although the obvious action to take when a machine breaks down is to repair it immediately (RI), this may not always be the most appropriate action for a variety of reasons:

- *Inappropriate*: If there is a long queue of parts downstream from the machine requiring repair, then immediate repair may not be the most appropriate action and the maintenance engineers may be better deployed elsewhere.
- *Insufficient*: Repairing a machine takes time. Meanwhile the rest of the production facility continues to process parts and to move them around. This means that, during the repair of the machine, queues may occur upstream, while downstream the process will be starved of parts. Simply repairing the machine may be insufficient to reach target throughput.
- *Impossible*: It can be assumed that sometimes it may not be possible to repair the machine immediately since all the maintenance engineers are busy. There is always the option of borrowing an engineer from a different part of the production line or to interrupt the repair of another machine and so to release one of the engineers, but this may not be the best course of action. Further to this, it can also be assumed that on occasions spare parts required for the repair of the machine may not be available.

Bearing the above in mind along with the conclusions from the pilot data collections, it is clear that other policies besides immediate repair are considered when a machine breaks down.

Stop the line (SL) is another option which might be considered. In this case, the maintenance supervisor (Group Leader) should decide whether it is useful to stop the whole line or part of it. This might be used, for instance, to avoid a build up of work-in-progress in a section of the line.

Do nothing (DN) is an alternative decision that might be the desired course of action under certain circumstances, for instance, close to the end of a shift. Obviously, this decision must be revised eventually and the machine repaired later.

Stand by (SB) can be considered as an alternative to repair immediately. In this case, an engineer processes the parts manually and pushes them to the next machine through the conveyor. In general it is not possible to repair the machine at the same time as stand by is being operated, because of space restrictions. This creates one of the main trade-offs of the unplanned maintenance decision making problem.

In view of the fact that the objective of the decision making problem is to minimise the delays and the throughput fluctuation caused by machine breakdowns, it is clear that the best choice between SB and RI is always situation specific. This is because SB might improve the flow of parts, but at the same time it requires an engineer to work on the machine bypassing the parts, it requires time to set up an SB bypass station and it operates at a slower rate than running the machine at normal mode (section 5.3.4). Hence if the repair time is 8 minutes and it is the beginning of the early morning shift then SB might not be such a good idea. Assuming that an SB action will operate parts only 1 second slower (a reasonable assumption in the light of what the decision makers explained during the pilot data collections) than the normal time required if the machine is running without problems and if 5 minutes are required to set up an SB station with 800 parts waiting to be processed before the end of the shift, SB will cost 18 minutes (800 * 1 sec /60 secs) while RI will cost only 8. On the other hand, if the repair time was 19 minutes it would clearly be better to follow an SB policy (if the machine can be operated in SB mode and other things being equal) since this would cost 18 minutes delay and it would save 1 minute.

The above example shows clearly that the trade-off between SB and RI depends on the status of the system, which can be defined by a number of attributes, and this is what decision makers take into account in reaching a decision. The example also shows that significant production improvements can be achieved if the right decision is made for each machine breakdown incident which occurs. This reveals that the choice between SB and RI is a significant element of the decision, but it is not the only one. The decision about how to resource the action that is decided and the decision about what else must be done in order to support the main action are still important elements in those decisions which contribute to the minimisation of the delays caused by machine breakdowns.

Resource Decision

Based on the experience from the pilot data collections, the resource decision is not constrained by the availability of resources as much as it is constrained by the suitability of the resources. This means that the decisions about who must implement the action that is decided (usually SB or RI) depend on what action has been decided and not whether the engineers are available. This is because if a specific engineer is not available then an equivalent engineer can be borrowed from a different area of the factory.

In order to resource the actions that have been decided, the decision maker has to choose one of the following engineers who must act when a machine breaks down:

- Group Leader
- Mechelec
- IMS1
- IMS2
- Operator

The Group Leader is the decision maker, a skilled engineer educated at HND or BSc level who is usually involved in 'non-standard' repairs. The Mechelec, a mechanical electrician, is also a skilled engineer who is deployed in machine repairs. He is also the person in charge when the Group Leader is not available. IMS1 and IMS2 (Integrated Manufacturing Specialists) are trainee engineers who are involved mostly in routine repairs. Finally, the operator is the person deployed in more practical tasks which involve supporting the SB process by carrying out and setting up stand by stations.

Complementary decisions

As was concluded during the pilot data collections, a number of other decision variables might sometimes be involved in the decision making process. Decisions about planned maintenance, seeking advice from the production manager and the switching off of a broken machine are some of the additional variables that might be involved in the decision making process when a machine breaks down.

In more detail, whether there is a need for planned maintenance and when this might take place is one of the decision variables that sometimes complements unplanned maintenance decisions. As the decision makers explained during the pilot data collections, when a machine breaks down very frequently, repairing it immediately or setting up stand by stations is not the best solution in the long run. Frequent breakdowns on a specific machine is a matter of concern that requires investigation which can only take place when the production line is not running.

Whether it is appropriate to switch off a machine that has just broken down is another decision variable which sometimes accompanies the decision about what action must be taken. This decision variable is relevant only when the broken down machine is one of those that is used for testing and monitoring the quality of the parts. Switching off one of those machines, if it has broken down, prevents delays that the repair of a machine would cause, but is a highly risky decision since it involves bypassing the quality control practices which are applied in the factory. A decision to switch off such a machine must always be approved by the production manager and, therefore, should always be accompanied by a decision to seek advice from the production manager.

Attribution of decisions

The above discussion of the unplanned maintenance decision making problem has revealed the complexity of the problem and the decision variables and the options that are considered when a machine breakdown occurs. It has also revealed that, in determining what course of action to take, the Group Leaders rely upon their knowledge and experience (tacit knowledge) which is combined with the attributes of the system at the time that the machine breakdown incident occurs. These attributes are an essential part of the problem since, as explained earlier in this section, they support the decision makers in determining the trade off between SB and RI and eventually in deciding what must be done.

According to the discussion with the Group Leaders and the observation of working practices combined with the use of VIS, some or all of the attributes in table 5.1 are taken into account when making a decision:

Attribute	**Description**
Machine number	A code that is used to uniquely identify the machine that has broken down. The machine number is one of the attributes which the decision maker receives in the short message that is sent to his pager once a machine has broken down.
Type of fault	A short description of the problem that is produced from the diagnostic system attached to the machine. This attribute is also sent to the pager that the decision maker carries during the shift.
Number of engines produced so far this shift	It is one of the environment attributes that indicates the current performance of the production line. This attribute is publicly accessible through a display that is located in the middle of the production line.
Availability of spare parts	An attribute that indicates whether a specific part that is required for the repair of the machine is in stock. This attribute is not easily accessible and once it has been decided that a specific part is required, one of the engineers has to check the stock availability in the warehouse. This check might take up to 5 minutes.
Number of other machines down	It is the attribute that indicates whether other machines have also broken down.
Physical condition of the machine	It is the attribute that informs the decision maker about the condition of the machine. The colour of specific components in the machine and the quality of the last part processed by it are two of the characteristics which reveal its physical condition.
Frequency of breakdowns	It indicates how often the machine breaks down. This attribute is not easily accessible since it is not reported when a breakdown occurs. It is available upon request, which involves queries in the central database.
Number of heads in the buffer	It indicates the number of parts that are currently waiting in the buffer for processing. This buffer is located in the middle of the production line and feeds with parts(heads) the second half of the line. Depending on where the machine is located, this attribute can be used to give an indication of the likelihood of a temporary bottleneck due to

Attribute	Description
	the repair of the machine.
Time	It is the attribute that reports the current time. It can be used in combination with the attribute 'engines produced so far this shift' in order to asses the performance to date.
The estimated repair time	It is a meta-attribute which reveals the decision maker's view about the time that will be required to repair the machine. This attribute is estimated by the decision maker in order to reach a decision about what action must be taken. Clearly, the estimation of the repair time is based on some of the attributes that have already been mentioned, such as the type of fault and the physical condition of the machine.

Table 5.1: *Decision attributes*

The above list reveals the dimensionality of the decision making problem and indicates that a number of simplifications will be required.

Having described the decision options and the attributes which determine the trade-off between the various options that are available when a machine breakdown occurs, this section has described our current understanding of the unplanned maintenance decision making problem in the production line modelled in this research. This description also shows that using VIS-based pilot data collections it is possible to identify the objective, the options and the attributes that are taken into account in decision making. However, with the use of the above approach it is not possible to identify the exact relationship that associates the attributes and experience of the decision makers with the decisions that are taken. As was anticipated in section 4.3.3, the determination of this relationship requires the development and calibration of a decision making model. In order to collect the data for developing such models, it is important to identify those elements of the problem which should be included in the model. This process is explained in the following section, which describes the development of the conceptual model.

5.5 Conceptual model

Whilst taking account of the purpose of the research and the implications of simplifying the problem, the conceptual model was developed by determining the scope of the modelling exercise and by making a number of simplifications which determined the level of detail of the model.

5.5.1 Scope and level of detail

In order to define the scope, all the aspects of the decision making problem were considered, but only those directly relevant to the investigation of the research were included. As was explained in greater detail in sections 1.3, 3.1, 3.2 and in Chapter 4, in addressing the questions raised by the research, the purpose of the research is to identify and compare the alternative unplanned maintenance decision making strategies which are employed during the production process and which have a direct impact on the throughput levels. In the previous section, it was shown that unplanned maintenance might sometimes include decisions which involve the request of planned maintenance. Planned maintenance, however, does not have a direct impact on the throughput levels, at least in the short run (daily and weekly). On this basis it was decided that the representation of planned maintenance is beyond the scope of this modelling exercise. Planned maintenance was, however, included as a decision option in order to create a realistic decision making environment.

As explained in the pilots (section 5.3 – Pilot 1), the estimated repair time is an input in the decision model. The action 'stop the line' does not form part of the standard practice and, as discussed in the problem definition (section 5.4), the resource decision depends on the suitability and not the availability of the decision makers. As was pointed out in the previous section, this means that the decision about who should be involved in the action is determined by the desired action.

Furthermore, given the scale of the application, it was decided to model the decision-making process only for a self-contained segment of the production line (the head engine assembly) which employs one decision maker per shift .

5.5.2 Conceptual model expressed

Notwithstanding the scope and the simplifications which were outlined and justified above, the conceptual model that has been developed for the purposes of this research is a simplification of the real problem, yet it contains sufficient level of detail to allow the specification of the data requirements.

The conceptual model consists of two elements: the attribute and the decision. The attribute element represents the various pieces of information which influence the decision making process. The decision element represents some of the decisions that the decision makers take when a machine breaks down.

In the conceptual model a decision consists of three decision variables: Action, Resource Decision and Complementary Decisions. The first variable, the 'Action', can only take a value that represents one of the two actions: SB and RI. The second variable, the Resource Decision, can only take one value that represents one of the following five options: Group Leader, Mechelec, IMS1, IMS2 or Operator. The third decision variable, the Complementary Decisions, which is included in the conceptual model for symbolic purposes, can take values that represent one or more from the following set of actions: ask production manager, plan weekend repair, plan off-shift repair, switch off machine.

The attribute element of the model consists of a set of variables, each of which represents specific information about the status of the system. For reasons that were explained earlier (section 5.3.2 and 5.3.3), the attributes 'number of other machines down' and 'stock availability' have been removed from the list of attributes. The following are the attribute variables which are considered in the conceptual model:

- Type of fault
- Machine number
- Estimated Repair Time
- Time Hours
- Time Minutes
- Engines
- Parts Waiting
- Number of breakdowns on this machine today
- Number of breakdowns on this machine this month
- Number of breakdowns of this type on this machine today
- Number of breakdowns of this type on this machine this month
- Number of Heads in buffer

Having determined the variables that must be included in each element of the model, the conceptual model can be described in KBI terms as a relationship between $\mathbf{D}_{i,j}$ and $\mathbf{A}_i$:

$$\mathbf{D}_{i,j}^{Conceptual} = f_j\left(\mathbf{A}_i^{Conceptual}\right) \qquad i = 1,2,3,\ldots n \text{ and } j = 1,2,\ldots m \qquad (5.1)$$

where

$$\mathbf{D}_{i,j}^{Conceptual} = \begin{bmatrix} \text{Action} \\ \text{Resource Decision} \\ \text{Complementary Decisions} \end{bmatrix}, \qquad \mathbf{A}_i^{Conceptual} = \begin{bmatrix} \text{Type of fault} \\ \text{Machine number} \\ \text{Estimated Repair Time} \\ \text{Current Time} \\ \text{Engines produced} \\ \text{Parts Waiting} \\ \text{Machine breakdowns today} \\ \text{Machine breakdowns month} \\ \text{Breakdown today} \\ \text{Breakdown month} \\ \text{Heads in buffer} \end{bmatrix}$$

n: number of decision situations, m: number of decision makers

Action : {RI,SB}, Resource decision : {Group leader, Mechelec, IMS 1, IMS 2, Operator}

Complementary decision: {Plan repair (and when), Ask production manager, Switch off machine}

As shown in expression 5.1, the elements of matrix $\mathbf{A}_i^{conceptual}$ represent the attributes that have been included in the conceptual model and not the attributes that were identified during the problem-definition process. Similarly, the decisions that are included in matrix $\mathbf{D}_{i,j}^{conceptual}$ reflect the elements of the decision making process which have been considered necessary for the purpose of the knowledge elicitation stage of the research.

5.6 Data requirements

Having designed the conceptual model, it is also possible to identify the data required for modelling the relationship which associates the attributes with the decisions taken by each decision maker. The process of specifying the data requirements is the final step of stage 1 of the KBI methodology and it involves deciding the number of records that must be collected and the number of decision makers who must be involved. In KBI terms this step involves deciding the number of records in matrices $\mathbf{A}$ and $\mathbf{D}_j$ and the value of the numbers n and m.

In order to decide the number of records that each data set must include, the following factors were taken into account:

- The decision makers' availability
- The type of decision situations that must be included in the data set
- The time that the decision makers need in order to make a decision

The decision makers' availability was the principal factor which constrained the size of the data sets. Owing to their commitments as Group Leaders, the availability of the decision makers who were involved in the data collections was very limited. From the preliminary stage they made it clear that they could not afford to be distracted from their duties for more than a total of 30 hours (almost 4 working days) each. Given the time that was required for the pilot data collections and the interviews (which, as explained earlier, contributed to the process of understanding and formulating the problem), each decision maker had no more than two hours available for the main data collection and for the follow-up meetings. Considering our experience from the pilot experiments, in which approximately 25 decisions were collected within a time slot of one hour, it was decided that in the final data collection each decision maker must be asked to input decisions up to a total of no more than 50 decisions.

The decision about the number of decision makers who should be involved in the main data collection was rather obvious. The segment of the production line that is modelled in this research (section 5.5) was operated 24 hours five days a week at the time when the problem-understanding and the data collection stages took place (May 2002). Round the clock operation requires three shifts and as it was pointed out in section 5.2 a group leader is employed for each shift. Based on this shift pattern it was clear that each of the three Group Leaders (DM1, DM2, DM3) from the three shifts of the specific section should be involved in a separate data collection. This indicated that a model should be produced for each of those three individuals.

5.7 Evaluation and analysis

The application of the first stage of the KBI methodology has shown that it is a process which involves five steps. The implementation of each of those five steps has contributed to the identification of the main elements of the problem and to the specification of the data requirements. From the description of each step the main difficulties which stage 1 involves have been identified and addressed. The decision making problem involves a large number of variables and due to limited data, only some of the AI techniques described in section 4.3.3 may be appropriate for the purposes of this research.

It has been found in this chapter that during stage 1 a preliminary investigation is required in order to gain a basic understanding of the problem situation. During the preliminary investigation, it is not expected that the problem and its elements will be fully understood since the decision makers might have difficulties in describing the problem. To gain a clear understanding of the problem, a number of pilot data collections that might involve VIS prototyping are required. The pilot data collections are perhaps one of the most time-consuming steps of stage 1. The number of iterations that are required and the time spent on this step depends very much on the availability of the decision makers and on whether a simulation model is available. Having gained a clear understanding of the problem situation, the problem must be defined by determining its elements: the objective, the decisions, the options that are available for each decision and the attributes which determine the decisions. The problem-definition is a key part of stage 1 since it forms the basis for deciding what must be included in the model of the decision making problem. The determination of the elements of the problem that must be included in the model is implemented during the fourth step of stage 1. This step is known as conceptual model development. As was shown in section 5.5, during this step the main assumptions and the boundaries of the problem are determined, taking account of the scope of the modelling exercise and the technological capabilities of the simulation model. From the description of the implementation of this step it is clear that, although it is not expected that the conceptual model will reflect all the details of the problem, it is expected that the model will incorporate all those elements required for developing decision models which can be used to assess individual decision making policies. The final step of stage 1 of the KBI methodology, as shown in section 5.6, involves the specification of the data requirements, taking into account the elements of the problem which are included in the conceptual model and the availability of the decision makers.

Having used a type of VIS in almost each iteration of stage 1, it is concluded that VIS is an effective technique that can be used for facilitating discussion and as a tool for problem understanding. Using VIS we were able to communicate and express our understanding of the decision making process through the presentation of the user interface which included the decision variables, the decision options and the relevant attributes that it was thought are taken into account by the decision makers. Using VIS it was possible to improve our understanding of the decision-making process and the attributes that are taken into account when a decision is taken. This was achieved through the discussion which the VIS facilitated and through the addition of missing attributes during the pilot data collection experiments that took place. During the pilot data collections we had the opportunity to ask the decision makers how valid our understanding of the decision making process is. The decision makers were able to give us feedback, taking account of the representation of the decision making process in the visual interactive interface.

5.8 Conclusion

In this chapter the first stage of the KBI methodology has been tested and evaluated. This stage involved the process of understanding the decision making problem and the process of specifying the data requirements. The application of KBI stage 1 on the case study that is used in this thesis has shown that the specific decision making problem involves a large number of variables both for the attributes and the decision. As will become clear in subsequent chapters, for the purposes of the analysis some elements of the problem will be simplified and reduced.

Chapter 6

Evaluation of stage 2 of the Knowledge Based Improvement methodology

This chapter tests and evaluates the second stage of KBI. This involves the design and the implementation of the data collection. After a brief description of the steps of the data collection, the implementation of each step is described in detail in the following sections of this chapter. The chapter ends by summarising the main methodological conclusions from the implementation of the data collection approach that is proposed as part of KBI.

6.1 Steps of the data collection process

The data collection stage of KBI involves three steps (figure 6.1) and its purpose is to elicit knowledge from the decision makers by collecting a set of decision instances for a set of decision situations. It builds upon the problem-understanding stage and combines system development with a simulation-based data collection. In KBI terms, this stage involves the population of the matrices $\mathbf{D}_j$ and $\mathbf{A}$.

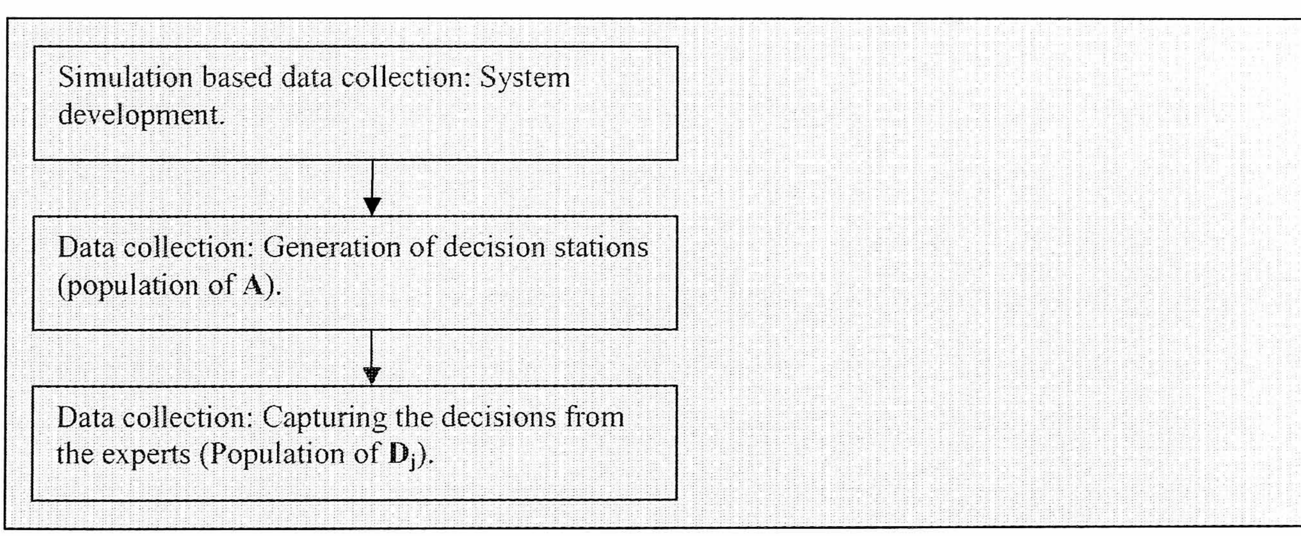

Figure 6.1: Simulation-based data collection process (KBI Stage 2)

6.2 Simulation based data collection: System development

Having determined the data requirements in Chapter 5, this section describes the development of the system used in the main data collection of the research. The system, which is a modular simulation, has been specifically designed to minimise the time required for the involvement of the decision makers during the data collection. The development of the system involves the specification of its requirements (what the system must do) and the design of its functionality (how the data collection system should meet the specified requirements).

6.2.1 Requirements specification

The system requirements which are described in the following paragraphs have been designed with the aim of improving the knowledge elicitation approach by developing a process that resolves the problems identified during the pilot data collections (section 5.3). In order to develop the system for collecting the decisions, the requirements and the functionality of the initial VIS were iteratively re-specified, taking into account the data requirements and the feedback from the decision makers who were involved in the pilot data collections. As with the pilot data collections, simulated decision situations were used to facilitate the knowledge elicitation in the main data collection of the research. The process involved the development of a common platform which acts as a communication interface between the analyst and the decision makers. This combines our experience from the pilot data collection with principles discussed by previous authors (O'Keefe & Pitt 1991, Au & Paul 1996, Williams 1996, Perry & Moffat 1997). To support the process, it was clear that the system must allow the analyst to express and communicate each decision situation in a way which enables the decision makers to apprehend it quickly and easily, enabling them to make a decision. Based on this general principle, a number of specific requirements that the system must have were identified and are described in the following paragraphs.

For each scenario that is presented to the decision maker, it was clear that the system must provide sufficient information about the decision situation. Taking into consideration the conclusions from the pilot data collections, it has been assumed that the attributes in the conceptual model provide sufficient information to specify clearly and objectively a decision situation. Given this assumption it was decided that, for each scenario presented to the decision makers, the decision situation must be communicated to them by reporting the value for each attribute which is included in the conceptual model.

Having identified from the pilot data collections the elements of the decisions which must be modelled, it was decided that the user interface for the main data collection must not allow the decision maker to input non-standard decisions. In view of this requirement, it was decided that the input user interface should consist of a combination of mandatory and non-mandatory questions, for which the user can choose an answer from a drop down list of options. The questions which should be included in the user interface must reflect the decision variables that are included in the conceptual model.

In order to enable the analyst to use the data for the modelling process, it was decided that the information on the attributes and the corresponding inputs should be automatically recorded and stored in a database.

The system must be capable of generating a series of decision situations by using the simulation model. In addition, in order to enable stratified sampling from those situations (section 5.3.4), it was decided that the system must allow the analyst to specify criteria and filter the decision situations which can be generated. To minimise the delays, it was decided that the simulation must not run during the data collection and no simulation status files (more details of this functionality is provided in sections 5.3.3 and 6.2.3) should be used. On the contrary, the system should allow the analyst to generate and store the decision situations in a database or in a data file using the simulation model. In addition, the system should be capable of retrieving and presenting the decision situations to the user during the data collection without re-running the simulation.

Finally, as the pilot data collections revealed that the decision makers do not use the schematic display (since they focus on the information provided to them through the electronic form - logical display), it was decided that the final version of the system should not use the visual schematic display.

6.2.2 Functionality and architecture

Having identified the above requirements to specify how the system should meet those requirements, the functionality of the system was designed taking into account the previously described requirements. Part of this functionality was developed using objects from the prototypes which were implemented for the pilot data collection. The software, which was written in Visual Basic, contains the following four main modules:

- Decision situation generator
- Decision situation filter
- Decision situation data collector
- Database

After a brief overview of the architecture of the system, each of the above four modules is described in detail in the following paragraphs. Figure 6.2 shows the structure of the system that was used for the final data collection. As is shown in figure 6.2, using a front end (2) the system enables the analyst to specify the options for the mode of the simulation run. The simulation model (3) then generates a database of decision situations (4) using information from the simulation process and details from the machine breakdown database (1). At the end of the run the decision situations are filtered, so that only the situations which represent different scenarios are included in the database table with the representative breakdown situations (6). This table is used with an interface (7) to present the decision situations to each decision maker who, using the interface (8), inputs his details and his decisions to the decisions table of the database (10). Difficulties had been experienced with the use of the simulation status files (section 5.3.3). This had mainly been due to the time required to upload a simulation status in Witness. In an attempt to make the data collection less time-consuming, the data collection system used for the main data collection does not use the status file. Instead, the situations that have been stored in the database of the system are presented to the user only through a dialog form (section 6.2.5). This, as was found in the pilot data collection, does not reduce the quality of the decisions since, as explained earlier, the decision makers in the specific case study do not take into account the dynamic visual display of the simulation.

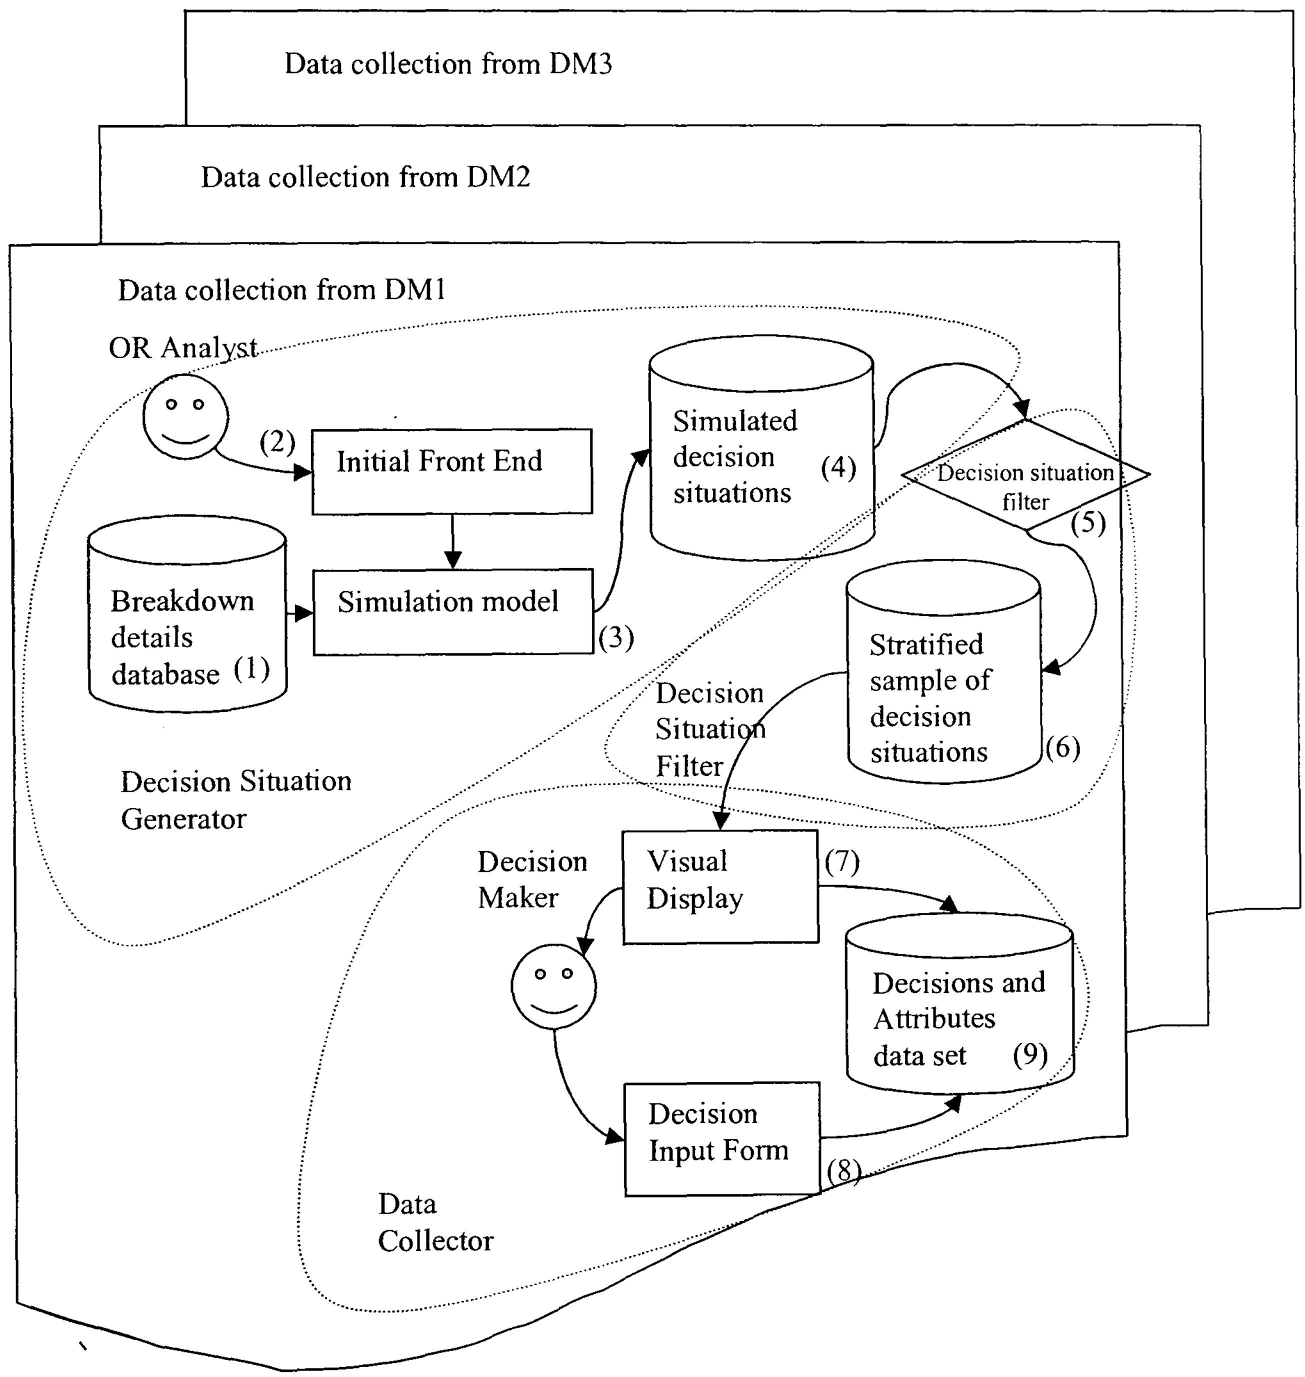

Figure 6.2: Architecture of the data collection system

6.2.3 Decision Situation Generator

The Decision Situation Generator, most of the parts of which were developed after the pilot data collections, is an automated object which is used to generate and store simulated decision situations by taking snapshots of the simulated system at the time when a machine breakdown occurs in the simulation. It consists of a user interface and an enhanced version of the initial simulation model, which is connected with the database of the system where the snapshots are stored in the form of attribute records. The purpose of this module is to allow the analyst to generate automatically a significant number of simulated decision situations. These, amongst other means, can be used to sample the scenarios which must be presented to the decision makers and also to investigate the diversity of situations that the decision makers are likely to encounter in the real system.

Simulation model

The main role of the simulation model in the Decision Situation Generator is to generate breakdown events in the context of the production system which is simulated. The breakdown events that occur during the simulation run contribute to the generation of the decision situations by providing information about various attributes of the simulated system at the time when the breakdown event occurs. For each decision situation that is generated, the simulation model contributes to the composition of the decision situation by providing specific information about the following attributes:

- Type of fault*
- Machine number
- Estimated Repair Time
- Time Hours
- Time Minutes
- Engines
- Parts Waiting
- Number of Heads in buffer
- Number of breakdowns on this machine today
- Number of breakdowns on this machine this month
- Number of breakdowns of this type on this machine today
- Number of breakdowns of this type on this machine this month

** non-simulated attribute*

To complete the generation of the decision situation, each time that a machine breakdown event occurs in the model, non-simulated attributes of the breakdown (the attribute with the asterisk in the above list) are retrieved from a trace with real breakdown incidents which have been recorded in the past in POSMON (section 5.2). The database table that contains the POSMON data in the system is known as the breakdown details table and it was populated with a complete download of breakdowns records covering the period 20-JAN-2000 to 01-AUG-2000.

As well as generating some of the attributes for the decision situations, the simulation run is used for composing each decision situation and for populating the 'decision situations' table of the database. The composition and population of the 'decision situations' table is achieved by the following process: each time that a breakdown event occurs in the simulation run, the simulation process pauses and the database table 'breakdown details' is queried to retrieve information about the 'next in the list' breakdown record for the machine which has broken down in the simulation. The result from the query is a record that consists of a breakdown description and the corresponding breakdown fault code. These attributes are combined with the attributes of the simulated system at the moment when the breakdown occurred. This set of information forms a new decision situation record which is then sent to the database table 'simulated decision situations'. Once the new decision situation is appended in the table, the simulation continues until the next breakdown event.

User interface

To control and enhance the functionality of the Decision Situation Generator, a user interface was developed to form a front end for the decision situation generator module. With the front end, the development of which started in the early stages of the research and which was also used for the pilot data collections, the analyst is able to customise the simulation run to meet the requirements of each data collection session. For example, as is explained in section 5.3.4, the decision situations in the fourth pilot data collection were presented to the decision makers using status files. Introduced in section 5.3.3, a status file is a Witness functionality that enables the analyst to save the status of the simulation at a specific time as a file that can be reloaded, enabling the simulation to continue from where it stopped when the file was saved. As shown in figure 6.3, using the front end the analyst is able to select the option 'Generate Sim Files' (Sim is the filename extension for a status file in Witness). Selecting 'Yes' from the combo box 'Generate Sim files' each time that a breakdown occurs in the simulation, the system automatically generates a file that contains the status of the system at the time when the breakdown occurred in the simulation.

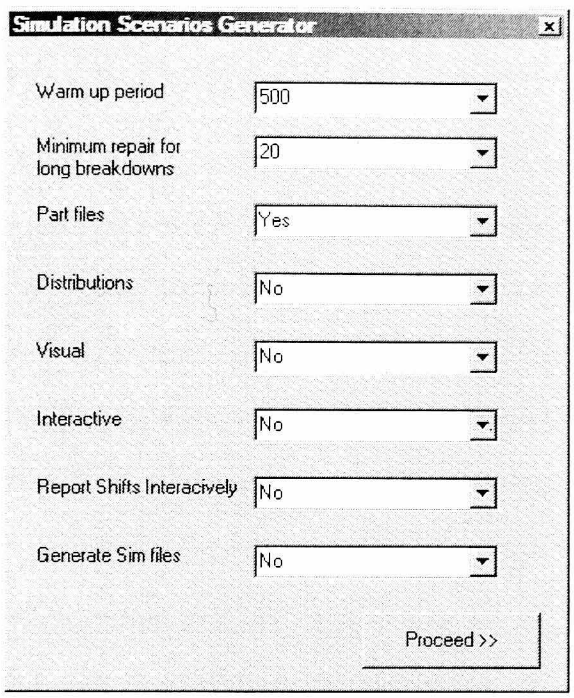

Figure 6.3: Decision situation generator - Front end

6.2.4 Decision Situation Filter

This module was extended and refined after the pilot data collections and, as will be explained in section 6.3.1, it was used to prepare the main data collection, in which the decision situations were generated from the simulation model prior to the data collection session.

The role of this module is to support the analyst in determining which of the decision situations that are generated from the Decision Situation Generator and are stored in the table 'decision situations' must be presented to the decision maker.

The decision situation filter consists of a set decision support tools which allow the analyst to explore the range of decision situations that occurred in the past and were used in the simulation run. These tools include a set of charts similar to the one shown in figure 6.4 and an interactive algorithm, which uses a set of criteria and is applied to the decision situations database (item 4 in figure 6.2) in order to take a stratified sample from the decision situations. It does so by extracting a subset of records which represent a wide range of situations that the decision maker may face (for more details see section 6.3).

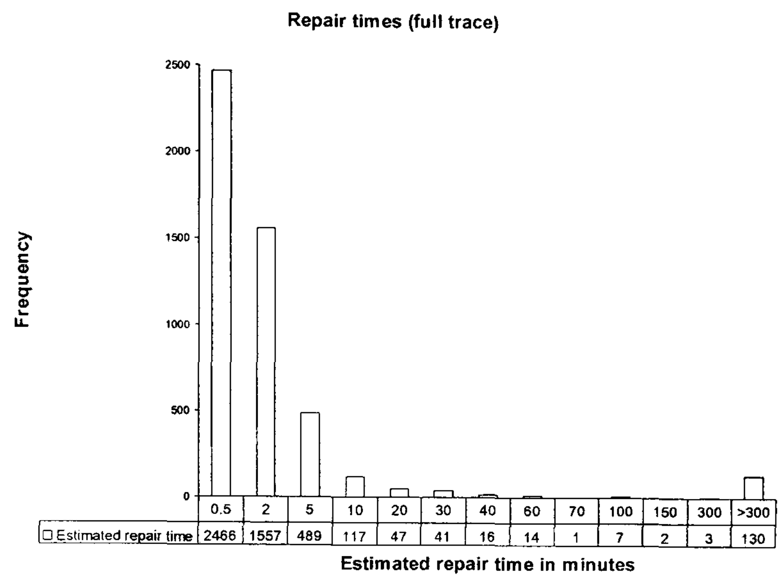

Figure 6.4: Decision situation filter – cluster chart

Having filtered the decision situations using the appropriate criteria, the algorithm populates the 'representative decision situations' table with the decision situations that are used for the data collection.

6.2.5 Data collector

The third module of the system is the data collector. This was developed after the completion of the pilot data collections and it is used in the main data collection to present the decision situations to the decision maker, to collect their decisions and to populate the records of the table 'decisions' (item 10 in figure 6.2). It consists of a set of user interfaces and the 'decision' database table used to store the decisions of each decision maker. The decision maker's ID interface (figure 6.5) is used to record the details of the decision maker who participates in the data collection.

The decision situation user interface is used to present each decision situation to the decision maker. A final version of this interface is shown in figure 6.6 (bottom left form). The decision input interface (top right form in figure 6.6) enables the user to input the decision, taking into account the attributes of the decision situation which are shown while the 'decision' form is active. Once the decision maker has input a decision and has pressed 'proceed', a new record is created in the decision table of the specific decision maker and a new decision situation is presented to him.

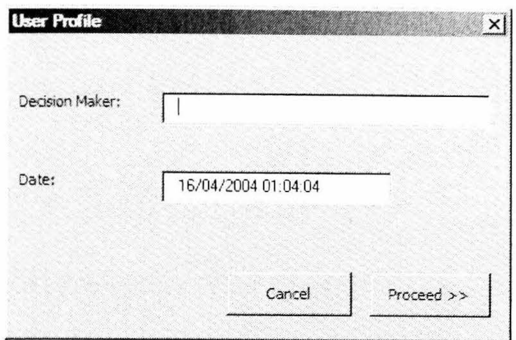

Figure 6.5: Data collector - Decision maker's ID form

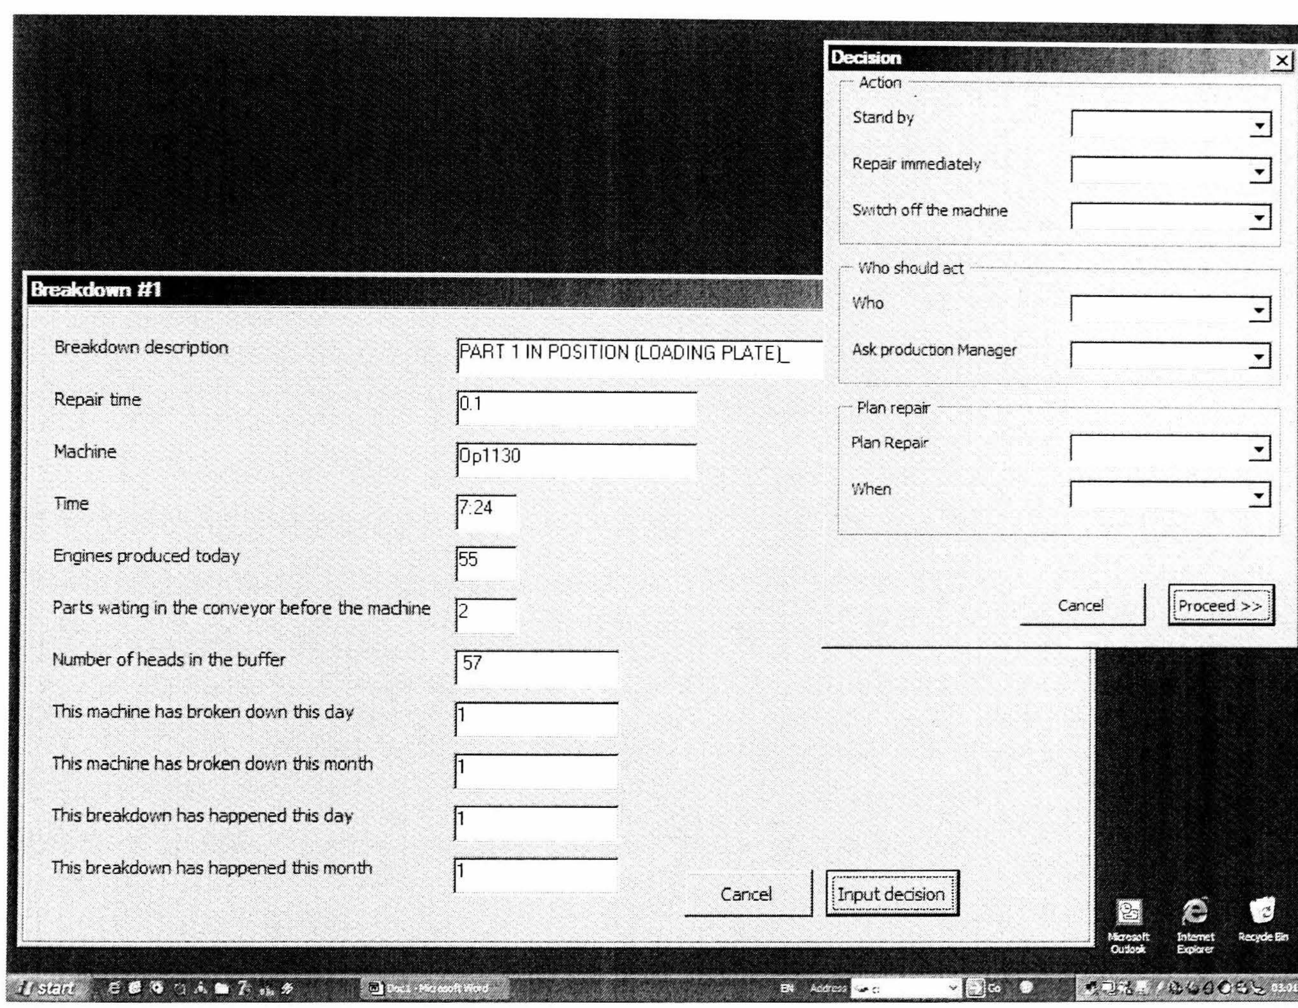

Figure 6.6: Data collector - Decision situation & Decision input forms

6.2.6 Database

The database consists of a set of tables (items 1,4,6,9 in figure 6.2) that are used to store and analyse the data sets. The breakdowns table (item 1 in figure 6.2) includes detailed information about breakdowns which were downloaded from the operational system from the production line (section 6.2.3). The decision situations table (item 4 in figure 6.2), as explained earlier, includes information about the breakdowns that occurred during the simulation run. The representative breakdowns table (item 6 in figure 6.2) is a stratified sample of breakdowns from the table decision situations (item 4 in figure 6.2). The decision and attributes table (item 10 in figure 6.2) includes the decisions that were collected during the main data collection sessions.

6.3 Data collection

Having described the development of the system which is used to support the data collection, this section explains the second and third steps of the data collection, which involve the process of generating the decision situations that were presented to the decision makers and the process of capturing their decisions.

6.3.1 Decision situations generation process

In order to model the strategies that are employed by each decision maker, it is clear that a representative data set should be used. To develop a decision model that is robust, generic and able to provide decisions for a wide range of situations, the data set must reflect each type of situation that might occur in the real system. To ensure that the data sets in this research include most of the situations that often occur in the real system, simulation experiments were used to investigate the range of situations that the decision makers might encounter during their duties. Using the simulation engine of the data collection system, a set of 4773 decision situations (the equivalent of 30 simulation days) was generated and stored in the database of the system. Table 6.1 shows the variation of each attribute that was observed after running the simulation for 30 days.

After the first five simulated days most of the decisions situations had occurred at least once, with very few significantly different new situations to occur after day ten. In addition, the cumulative mean of the estimated repair time did not change more than 1% after the 10^{th} day of the simulation for the additional situations which were generated. In view of the above findings, it was decided that 30 days is sufficient time for the simulation to reveal the variation of each attribute. The set of situations which were used in the data collection was then determined by stratified sampling from the set of situations generated during the 30 days simulation run.

Attribute	**Min**	**Max**
Estimated Repair Time	0.3 Minutes	>5 hours
Time (Hours)	00:00	23
Time Minutes	0	59
Number of engines produced so far this shift	0	914
Parts Waiting	0	10
Machine breakdowns this day	1	25
Machine breakdowns this month	1	320
Breakdown day	1	15
Breakdown month	1	77
Heads in buffer	0	74

Table 6.1: *Range of attribute values*

Analysing the decision situations, a significant variation in the attribute estimated repair time was observed amongst the incidents that were recorded during the simulation run. Figure 6.7 (top left) shows that, for many of the simulated incidents, the estimated repair time is either less than 2 minutes or more than 300 minutes. Based on the knowledge and experience acquired through our involvement in the maintenance process during the preliminary investigation and during the pilot data collection, where a number of incidents with an estimated repair time of less than 30 seconds were presented to the decision makers, it was decided that these extreme situations are not part of the decision making process in which the decision makers are involved and, therefore, must be excluded from the sample of decisions that should be presented to them.

The incidents with a repair time of under 30 seconds reported in the simulation can be explained by the fact that very often a machine fails because the parts are not positioned properly in the machine from the transfer conveyor. The repair of this type of failure is always obvious and requires resetting the machine, a process that takes normally 10 to 20 seconds. Although these incidents are recorded as breakdowns, since the machines become idle and are unable to process parts until an engineer resets them, those situations are not part of the decision making process since the repair action is always the same. In addition, from the conclusion drawn from the discussion with the decision makers during the pilot data collections, it is known that the incidents that are reported to have an estimated repair time of more than 300 minutes usually represent situations in which, for reasons related with imperfections of the diagnostic system, the down time is not properly recorded. This might happen because the glass door located in front of some of the machines as an additional safety measure might not be fully closed and so the diagnostic system assumes that the machine is broken down while it is actually fully operational.

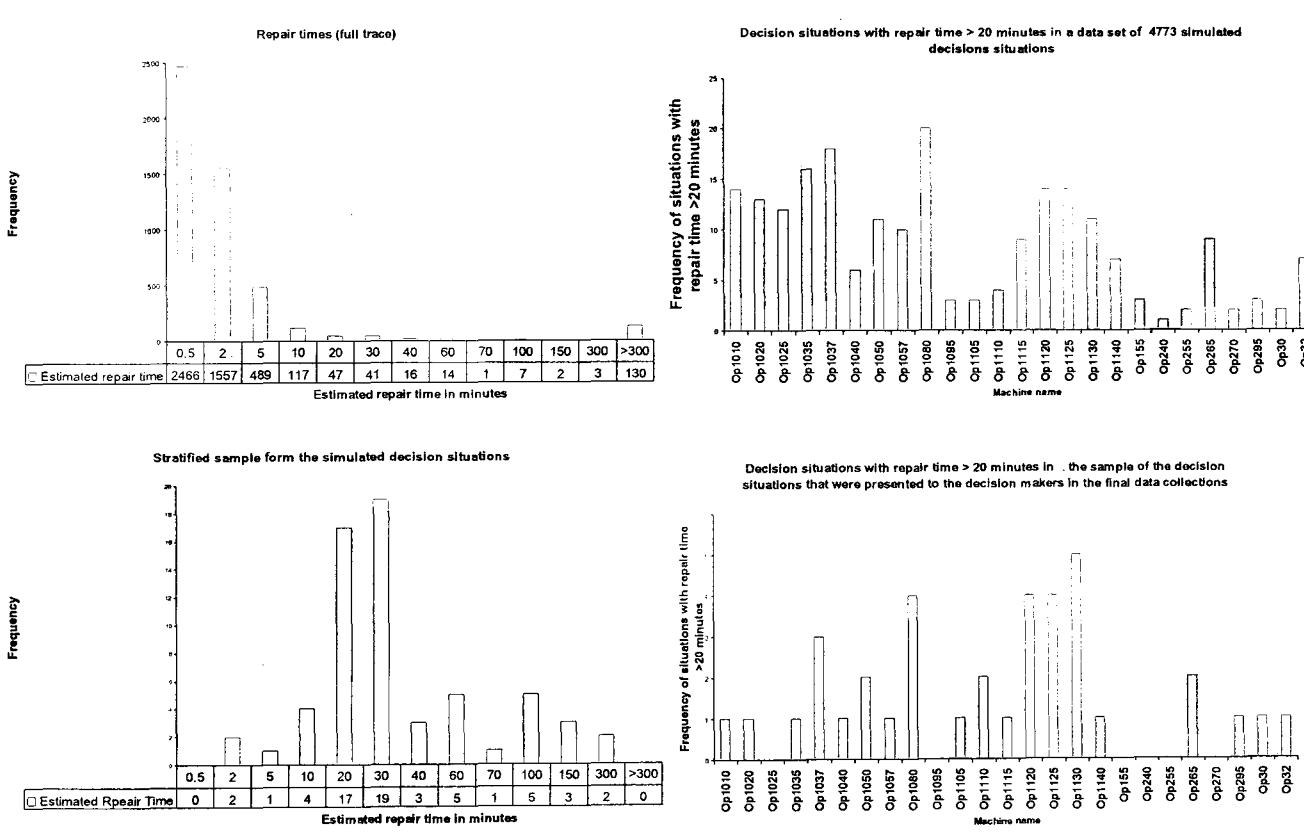

Figure 6.7: Decision situations – simulated population (top) and sample (bottom)

Observing the remaining incidents in figure 6.7 (top left) it was found that, although the repair time for many situations is estimated to be between 2 and 5 minutes, there is a significant percentage of incidents with an estimated repair time of between 30 and 60 minutes. Analysing the incidents that occurred in each machine included in the head engine assembly (the segment of the line that is modelled in this research – section 5.5). it was found that for most of the machines there was at least one breakdown incident with an estimated repair time of more than 20 minutes - figure 6.7 (top right). It was therefore concluded that the wide variation in the estimated repair time is not an isolated incident which occurred only on one machine and, therefore, it must be reflected in the sample of decision situations that should be presented to the decision maker during the data collection. Having decided that the sample of the situations to be presented to the decision makers must reflect the above variation in the repair time, to avoid knowledge gaps it was also decided that the data set must include at least one incident from each machine.

Given the above requirements, 63 decision situations were sampled from the simulated machine breakdown incidents. These were used for the main data collection. The sample shown in figure 6.7 (bottom left & right) contains at least one incident for each machine which is included in the segment of the line that is modelled in this research. In addition, in order to ensure that incidents with a long repair time are reflected in the models for most of the machines, there is at least one incident where the repair time is estimated to be more than 20 minutes (figure 6.7 bottom right). Besides the above criteria, a number of optional criteria were also applied in order to make sure that the data set is a representative sample of a complete set of situations which might occur during the decision making process.

6.3.2 Decision capturing process

For the main data collection of the research, each of the three Group Leaders from the three shifts of the area head engine assembly participated in an interactive data collection session, using the data collection system described in section 6.2. Prior to the data collection, basic training on how to use the system was given to each of the decision makers who participated. After a demonstration of the main features and capabilities of the system, the decision maker was asked to input his details in the introductory electronic form of the system. Once his details were recorded, an interactive dialog with the system was initiated. A decision situation was presented initially in an electronic form. Pressing the button 'input decision', an input dialog box was invoked. The decision maker, taking account of the attributes of the situation still visible on the screen, made a decision and on pressing the button 'Proceed' the forms closed, the input was recorded in the appropriate database table and the next decision situation was presented to him. The above process was repeated 63 times and 63 decisions were collected during the timeslot which he had available (2 hours). The above process was repeated three times, each time with a different decision maker but with the same decision situations. As is shown in the form in figure 6.6, the attributes and the decisions which were included in the dialog match with those which it was decided to include in the conceptual model.

As a result of the completion of the data population process, the three data sets that were collected and are available for the reader of this thesis in appendix 3 contain 63 records of data for each of the fields listed in table 6.2. Each data set contain a wide variety of situations and as will become clear in the subsequent chapters each type of decision action has been captured with a significant frequency.

The $\mathbf{A}_i$ fields of the table represent the attributes of a decision situation i and describe the status of the system at the time when the decision was taken. As the column 'Type' shows, the status of the system is described using a set of quantitative and qualitative fields.

Since the values of five decision variables were recorded for each decision situation which was presented to the decision makers in the final data collection, five data fields have been used to store the recorded decisions in the database. Given that some decision variables represent Yes/No type decisions while others represent multiple choice decisions, a mix of integer and Boolean data types were used to store this information in the database. Table 6.2 shows that a decision made by a decision maker at a decision point consists of selecting one of the alternative options for each of the five decision variables in $\mathbf{D}_{i,j}$.

	Id	Field name	Type	Instance
	1	Type of fault	String	DISTRIBUTOR IN POSITION(FEEDBACK)
	2	Estimated repair time	Real	5.4
	3	Machine number	String	Op1037
	4	Time	Time	7:24
	5	Number of engines produced so far this shift	Integer	1
	6	Parts waiting in the conveyor in front of the machine	Integer	1
A_i	7	Number of heads in the buffer	Integer	56
	8	The machine has broken down this day	Integer	1
	9	The machine has broken down this month	Integer	10
	10	The breakdown has happened this day	Integer	1
	11	The breakdown has happened this month	Integer	1
	12	Action	1:Stand By 0:Repair	1:Stand By
	13	Switch off the machine	1:Yes 0:No	1:Yes
	14	Who	1:Group Leader, 2:Mechelec 3:IMS1,4:IMS2, 5:Operator1	1:Group Leader
$D_{i,j}$	15	Ask the Production Manager	1:Yes 0:No	1:Yes
	16	Plan Repair	1:Yes 0:No	0:No
	17	Plan repair – When?	0 no planned repair 1 End of the shift 2 weekend	1:End of shift

Table 6.2: *Fields in the data set*

As already mentioned the decision making problem involves a large number of variables. It is clear that for the purposes of the analysis the variables listed above will be simplified and reduced.

6.4 Evaluation and analysis: Towards a simulation-based knowledge elicitation approach

Having tested stage 2 of the KBI methodology, in this chapter an emerging knowledge elicitation approach for modelling human decision making in simulation has been evaluated. From the development and the application of the approach, a number of lessons have been learned regarding the practical difficulties involved in a data collection for modelling decision making in simulation. This section summarises the main conclusions about the benefits and the practical difficulties that a simulation-based knowledge elicitation approach involves. In addition, some guidelines are given about how data collections intended for modelling human decision making in simulation should be implemented in practical applications, such as the case study described in this research.

Implications for the practitioner

The simulation-based knowledge elicitation approach which was tested in this chapter has shown that simulation is a useful technique for generating data for experimental designs and data collections. How appropriate simulation is for generating decision situations was shown in this research by the fact that, while in the first pilot data collection session (section 5.3.1) almost every decision was of the same type, having used simulation in the final data collection, we were able to control the decision situations. This created an experimental environment in which a wide variety of situations had a chance to happen and each type of decision was captured with a significant frequency.

Simulation was, therefore, found to be an essential tool for collecting decision making data since it enables the researcher and OR practitioner to control the attribute values and so the decision situations which are presented to the decision maker. Given that, in a data set appropriate for modelling decision making, every possible type of decision option should be represented with a balanced percentage of instances, the decision situations to be included in the data set must be an amalgam of situations representing each scenario that the decision maker faces. Unlike the real world system, where decision situations occur randomly and are not controlled by the modeller, simulation showed that it allows the modeller to create a significant number of decision situations within a few hours by running the simulation model for a sufficient number of simulated days. By selecting the appropriate decision situations from the simulation model at the end of the simulation run, as was done in this research, it is possible to generate each possible decision situation within a few hours, regardless of how rarely it occurs in the real system. Clearly, in order to do so a valid simulation model of the real system is required.

Practical difficulties

From the application of the data collection methodology some practical difficulties were identified and possible ways to resolve them were found. First of all, it was found that an existing simulation model, in which human decision making representation is not modelled, might not be appropriate for data collection intended for modelling human decision making. This is because the attributes which are taken into account might not be simulated in the model.

Due to the general structure and scale of the model and the nature of the attributes, it might not be possible to represent in the model one or more of the attributes that are taken into account in order to make a decision. This problem can be resolved either by retrieving data about the missing attributes which are not able to be simulated from databases that are connected with the simulation model or by replacing those attributes with others that are correlated with them or with meta-attributes. The use of the breakdown details database table for reporting the type of fault to the decision makers and the use of the estimated repair time to replace the physical condition of the machine are two examples that were applied in this research and they illustrate how the problem referred to previously can be resolved.

The need to involve the human decision makers in interactive simulation runs is an additional problem that might prohibit or constrain the use of VIS for practical data collections. If large amounts of data are required, the involvement and commitment of the human decision makers might be a potential problem, since human decision makers are extremely busy and might not be available. This was one of the main challenges during the data collections in this research. While a large data set was required, only a small data set could be collected within the time that decision makers had available (two hours). To resolve this problem, the researcher or OR worker can increase the number of situations presented to the decision maker by not running the simulation model during the data collection. It can instead be run in advance and the decision scenarios saved with the appropriate format, so that they can be retrieved and presented to the decision maker using a system that does not require running the simulation during the data collection. Clearly, this approach assumes that the inputs from the decision maker do not influence the attributes that will be used for the subsequent situations that will be presented to him.

The need to generate a wide range of decision situations in order to avoid collecting a data set of stereotyped decisions is another challenge that the use of simulation for data collection purposes might involve. Identifying the range of variation of the decision situations and filtering out similar situations is perhaps the most efficient way to resolve the problem. Having experienced difficulties in the early stages of collecting a data set of non-stereotyped decisions, we can say that the above difficulty, combined with the limited time that the decision makers might have available, is the most important limitation of the approach. The use of stratified sampling and the use of the filter for excluding short repair time breakdowns illustrates how filtering out similar situations can resolve the problem described above.

Finally, depending on the benefits expected from collecting the decision making data and on whether a pre-existing simulation model can be used, the monetary cost that the simulation-based data collection involves might be an issue which prohibits its practical application. As will be explained in the next chapter, in order to collect a data set that is appropriate for modelling the human decision makers, significant time was spent in each of the iterations of the methodology and some time was required from the decision makers. 1.5 years project time was required, excluding the development of the base simulation model, which in this case pre-existed. Clearly, knowing the steps that must be followed for capturing the data set and using the experience which has been gained and which is reflected in this conclusion, this period could be reduced the next time that the above exercise is repeated.

6.5 Conclusion

In this chapter the second stage of the KBI methodology has been evaluated by attempting to collect a set of maintenance decisions using a simulation model. The approach using a real world case study revealed the main strengths and weakness of a simulation-based data collection. The capability of the simulation to generate a wide range of situations and the use of the visual interface for structuring the problem and identifying the data requirement are amongst the major strengths of the methodology. Whereas the involvement of the human decision makers is necessary, the requirement to represent and simulate a wide range of attributes and the time that is required to collect a data set with a wide range of decisions situations are amongst the main challenges of the approach.

Chapter 7

Evaluation of stage 3 of the Knowledge Based Improvement methodology

This chapter tests and evaluates the third stage of the KBI methodology. As explained in section 4.3.3, this stage involves the process of identifying and modelling the decision making strategies. After a detailed description of the application of the modelling technique, the results of each model are reported and the validation checks that were undertaken are described. The chapter reports the conclusions of the modelling exercise and recommends the modelling structure that was found to be the most appropriate for representing the human decision makers in the simulation. Section one describes the processing of the data sets that were collected during the simulation-based data collection (section 6.3). Section two describes the modelling tools used in this chapter. Section three describes the process of modelling decision making using decision trees. Finally, section four summarises the main methodological conclusions from the modelling process.

7.1 Dimensions reduction

In section 5.8 it was anticipated that the decision making problem includes too many variables for the purposes of this analysis. A number of these variables either cannot be modelled or their representation in the model is not required since they are correlated with alternative variables. This section describes the process of reducing the variables of the problem.

As already explained in sections 5.3.1 and 5.5.1, due to the difficulties that the decision maker faces in making a decision about how long it will take to repair a breakdown, based only on the attributes of the simulated system and without being able to inspect the machines, it was decided to simplify the decision making process that is modelled by revealing to the decision maker the estimated repair time (sampled from the distribution of the repair time for that operation). Using the estimated repair time as a meta–attribute in the situations that were presented during the data collection, the decision makers did not have to inspect the machine. They were not obliged to take into account the type of fault since they did not have to decide how long it would take to repair the machine. Assuming that the type of fault influences only the estimation of the repair time and not the other elements of the decision, it was decided to exclude the variables which represent the types of faults from the attribute set $\mathbf{A}^{\mathrm{Calibration}}$ that was used for the calibration process. Although reducing the dimensions of the data set, by excluding the type of fault from the calibration process, could have an impact on the accuracy and the specification of the models, it was necessary given the number of dimensions that are required to represent qualitative variables such as the type of fault.

In addition, as observed in section 5.4, the resource decision depends on the action that is decided. For immediate repair, expert knowledge is very often required and for this reason it is always undertaken by skilled engineers (the Group Leader or the second Mechelec). The stand by process requires practical skills and it is time-consuming. To enable the skilled engineer to focus on immediate repairs, the stand by process is always resourced with IMSs or operators. The decision about which specific IMS or operator will be involved is based purely on their availability status. The above was concluded during the pilot data collections. In the sessions with the decisions makers, it became evident that the decision about who should be involved is based on what action is required. They explained that, having decided what action must be taken, the decision about who should be involved is determined by the factory rules. This was verified during the main data collection and it is reflected in the data sets that have been collected. On the basis of the above conclusion, it was decided to exclude from the calibration process the variables which represent the resource decisions.

Further to this, the complementary elements of the decision (variables 'switch off the machine', 'ask the production manager', 'plan repair' and 'plan repair when' in table 6.2-Chapter 6) were also excluded from the model calibration process. The decision to exclude the above variables was made taking into account that there were not enough records where complementary actions were taken. Their representation in the model does not affect the assessment of the decision making strategies and these types of decisions do not constrain the simulation in any obvious way. Although planned repair might have an impact on the production process, the representation of planned maintenance was beyond the scope of the research, for reasons given in section 5.5.1. In addition, the authorisation from the production manager is a process that is followed principally to ensure that the production manager is aware of the actions rather than in order to constrain the decision.

Having decided to exclude from the calibration the variables that represent resource, complementary decisions and the attributes that represent the type of fault, each of the three data sets $\mathbf{D}_j^{calibration}$, $\mathbf{A}^{calibration}$ which were used to calibrate the models include 63 records with 1 binary dependent variable and 35 independent variables. The binary dependent variable represents the action decision and the 35 independent variables represent the attributes of the system. Table 7.1 shows the structure of the data set which was used for training the models.

	Attribute Name	Type	Instance
	Repair Time	Real	5.4
	Machine number	Boollean array[25]	[0,0,0,1,0....0]
	Time	Integer Array[2]	[07 :10]
	Number of engines produced so far this shift	Integer	1
	Parts waiting in the conveyor in front of the machine	Integer	1
$\mathbf{A}^{calibration}$	**Number of heads in the buffer**	Integer	56
	The machine has broken down on this day	Integer	1
	The machine has broken down this month	Integer	10
	The breakdown has happened on this day	Integer	1
	The breakdown has happened this month	Integer	1
$\mathbf{D}_j^{calibration}$	**Stand By or Repair Immediately**	Boollean [1/0]	

Table 7.1: *Structure of the calibration data set* $\mathbf{D}_j^{calibration}$, $\mathbf{A}^{calibration}$

7.2 Method of analysis

In order to determine the technique to be used for modelling and understanding the strategy employed by each decision maker, a number of well-established techniques which have been used in the past (section 4.3.3) to represent decision making were considered.

The data sets which became available as a result of the data collection stage of KBI contain data for 25 machines but only a total of 63 decisions. This means that on average there are less than 3 decisions per machine. Statistically this is a very limited number of decisions, so it was concluded that it would not be possible to use the data to model accurate representations of the decision makers. This conclusion was taken on the basis that it would not be possible to test the statistical validity of the relationship that associates the attributes with the decisions.

As the collected decisions are a representative sample of the decision situations faced by the supervisors, it was considered that a more appropriate approach would be to use the collected data to identify differentiations amongst the strategies employed by each individual decision maker.

On the basis of this conclusion and in order to test the third stage of KBI, it was decided to use decision trees as a technique for identifying aspects of the strategy employed by each decision maker. Decision trees were chosen from a wide range of alternative techniques since it is the only one which can be used to understand the differentiations in the strategies employed by the alternative decision makers. It has the advantage of being transparent and it is the least data-demanding technique, as cross validation can be avoided if the trees are tested qualitatively.

Besides decision trees (O'Keefe & Roach 1987, Doukidis 1987, Doukidis & Paul 1985, Abdurahimman & Paul. 1994), other techniques were considered (Neural Networks - Liang 1992, Logistic Regression - Malhorta et al 1999, Case-based Reasoning - Liao 2000) which could have been used to represent and assess decision making. However, given the limited data and given the emphasis of the research on understanding why some strategies perform better than others, it was decided to limit the investigation to the above technique since, ultimately, it was felt that it is the most appropriate for the purpose of the research.

7.3 Rule induction

Using rule induction, the model calibration process involved the derivation of a set of rules that use the attributes from the data set for clustering the dependent variable. The calibration algorithm which is used is known as ID3 (Quinlan 1979) and it was originally developed for data sets with discrete limited dependent variables. The priority rule which is applied, to determine the sequence in which the attributes must be used for clustering the data set, is determined either automatically from the calibration algorithm (automatic ID3) or by the user (semi-automatic ID3). When automatic ID3 is used, the attribute which can classify the most decisions is applied first. When semi-automatic ID3 is used, the priority with which the attributes will be applied is determined by the user, by specifying a list with attributes that are ordered by priority preference. When automatic ID3 is applied, the number of iterations and the attributes which are used are determined by the algorithm and there is little scope for experimentation. However, when semi-automatic ID3 is used, given that the user has the option to decide the priority with which the attributes will be used to classify the decisions, a possible experimentation with the algorithm is to determine which attributes must be used first, based on the knowledge of the decision making situation.

For the purpose of this research the commercially available rule induction software Xpertrule (Attar software 2000) was used. The software allows the calibration with automatic and semi-automatic ID3. For the first experiment, the automatic rule induction algorithm was applied. Allowing the algorithm to decide which attributes must be used first, based on their classification power, the models that were produced classified correctly all the decisions in each data set. In order to improve the robustness of the models and to ensure that theoretically important variables are included in the models, a second set of models was produced using semi-automatic ID3. The following paragraphs describe the two stages of the rule induction process and the analysis of the results.

7.3.1 Model calibration with automatic rule induction

Decision models

The decision trees t_{DM1}^{A}, t_{DM2}^{A}, t_{DM3}^{A} in appendix A1.1-1.3 represent the initial set of models that were derived from the data in matrices $\mathbf{D}_j^{\mathrm{Calibration}}$, $\mathbf{A}^{\mathrm{Calibration}}$ and from allowing the calibration algorithm to determine which attribute should be used first in the tree. As is clear, the rules are organised into trees and each tree represents the model of a specific decision maker. Depending on the situation, there are three alternative decisions that a model may recommend. A node 'SB' at the end of a terminal branch indicates that a stand by decision will be recommended by the model if the input decision situation meets the conditions of the rule represented in the branch. A node 'RI' at the end of a terminal branch indicates that a repair immediately decision will be recommended by the model if the decision situation meets the conditions of the rule. Finally, a node 'Empty' indicates that there is a knowledge gap and there is no recorded decision for the set of situations which the specific rule represents. The frequency next to each outcome indicates the number of decision situations in the data set which meet the conditions of the rule that the branch represents. The probability at the end of each terminal node is a measure of the goodness of fit and it is calculated by the ratio shown in the following expression:

$$P(D_{i,j} = X / A_i^{calibration}) = \frac{[\text{Number of correctly predicted decisions for the node}]}{[\text{Total number of decisions in the data set for the node}]} \qquad 7.1$$

Where:

- $P(D_{i,j}=X/A^{calibration})$ is the probability that a decision $\mathbf{D}_{i,j}$ of type X will be taken when the observed attributes $\mathbf{A}_i^{calibration}$ are within the boundaries that the specific branch of the tree represents.
- [Number of correctly predicted decisions for the node] is a count of the decisions in the data set that the model has predicted correctly and which have attribute values that fall within the range that the specific branch of the tree represents.

- [Total number of decisions for the node] is a count of the decisions in the data set that have attribute values that fall within the range that the specific branch of the tree represents.

Comparison of the decision making strategies

Comparing the structure of each of the three decision trees $t_{DM1}{}^{A}$, $t_{DM2}{}^{A}$, $t_{DM3}{}^{A}$ it can be seen that the structure of the rules and, therefore, the strategy that each decision maker follows are significantly different. In the model of the first decision maker (DM1), the first attribute that is used to split the decisions is the estimated repair time (ERT). The fact that ERT is the first condition in the rules reveals that it is the attribute which can classify more decisions in DM1's data set than any other attribute. Based on the knowledge that the estimated repair time is the most consistent predictor of his decisions, it was initially assumed that DM1 behaves as if he believes that, depending on the estimated repair time, most of the machines can be operated in stand by mode.

The first attribute that is used in the models representing the second and third decision maker is the machine number. This shows that, unlike the first decision maker, DM2 and DM3 behave as if they believe that it is the machine number which mostly constrains their decisions. The above behaviour is consistent with what we expected since, from the discussion during the pilot data collection, it is known that certain machines must not be set in stand by mode and, therefore, it is just not possible to consider a stand by policy for those machines. Although it seems that both DM2 and DM3 look first at the machine number before they make a decision, their strategies about which machines cannot or need not be operated in stand by mode are quite different.

Figure 7.1, which highlights part of the trees represented in appendix A1.2 – 1.3, shows that DM3 seems to behave as if he believes that only nine machines cannot or need not be operated in stand by mode. DM2 behaves as if he believes that 19 machines cannot or need not be operated in stand by mode. Comparing these two sets, it is concluded that most of the machines which DM3 seems to believe are not necessary or possible to be operated in SB mode are included in the set of machines which DM2 seems to believe are not possible or necessary to be operated in SB mode.

The estimated repair time that was reported in most situations where the above machines were involved was much higher than the average repair time regarded as the threshold after which the decision makers opt for SB. This indicates that, although DM2 and DM3 naively should have operated the machines in SB mode, they did not do so because they acknowledged that their decisions are constrained by various factors.

$t_{DM2}{}^{A}$: MA= Op1010,Op1020,Op1025,Op1037,Op1040,Op1050,Op1057 ,Op1095,Op1105,Op1110,Op1115,Op1120,Op155, Op255, Op295, Op30, Op32, Op240, Op270 ——— RI, Freq 36, Prob1

$t_{DM3}{}^{A}$: MA= Op1025,Op1035,Op1095,Op1105,Op15 5,Op255, Op32,Op240,Op270 ——— RI, Freq 12, Prob1

Figure 7.1: Branches from the trees $t_{DM2}{}^{A}$ & $t_{DM3}{}^{A}$

An additional difference between the decision trees which is worth mentioning is that only the decision tree representing DM1 has apparent knowledge gaps (branches with outcome 'empty'). This seems to be due to the fact that the algorithm has given priority to the attribute 'estimated repair time' which is the first to be used in the conditions of the rules in the model. The estimated repair time splits the data set into two subsets. The first includes all the records that have an estimated repair time of less than 20 minutes while the second includes all the records with an estimated repair time of more than 20 minutes. The knowledge gap seems to be caused by the fact that there are insufficient records to represent each machine at each of the two levels of ERT.

Validation

In order to validate the above conclusions, the models were presented to the decision makers and to the production manager of the factory (in separate meetings). Discussing the constraints of the decision making process with them, it was found that each decision maker and the production manager believe that the machine number is the attribute which mostly constrains their decisions, as certain machines must not be operated in SB mode. When this issue was discussed with DM1, he did not agree with the conclusion that his decisions are based mainly on the estimated repair time, but he admitted that the set of the machines which he considers must not be operated in SB mode is less inclusive than the set of the machines that the other two decision makers think that must not be operated in SB mode.

Taking into account the follow-up discussion with DM1, it was concluded that, although the estimated repair time is the first attribute used in the conditions of the rules in his model, it is not the attribute that carries the highest weight in his decisions. It is used first and is therefore the most consistent predictor of the decisions, because the number of machines that DM1 believes cannot be operated in SB mode are not enough to classify as many decisions as the clusters which are created when the estimated repair time is used to classify the decisions.

7.3.2 Model calibration with semi-automatic rule induction

In order to remove the knowledge gaps and to make sure that the machine number is always the first attribute to be taken into account, since it represents the decision making problem qualitatively better, semi-automatic rule induction was used as a second experiment. Taking account of the information collected during the discussion with the decision makers, it was decided that, after the machine number and the estimated repair time, the number of heads in the buffer should be used in the conditions of the rules if a third attribute is required.

Having specified the priority with which the attributes must be used in the rules with the semi-automatic rule induction algorithm that was applied in matrices $\mathbf{D}_j^{\text{Calibration}}$,$\mathbf{A}^{\text{Calibration}}$ the models $t_{DM1}{}^{s}$, $t_{DM2}{}^{s}$, $t_{DM3}{}^{s}$ were produced (appendix A1.4-1.6). Given that each branch of each model has an outcome with a probability of one, it is concluded that each model can classify all the decisions in the data set which is used for its training.

In the model representing DM1, there was an obvious knowledge gap in the branch which represents the decision making strategy for the machine OP1140. However, this time the knowledge gap is not genuine and it has been caused by the requirement to use the estimated repair time as a second attribute. Analysing the situations which involve the machine OP1140, it was found that there were two incidents when the machine OP1140 was involved in a breakdown. Having found that the two decision situations have the same estimated repair time, it was concluded that the knowledge gap has been caused by the algorithm, which is unable to use the repair time for splitting the two decisions. Therefore the knowledge gap was resolved by removing the attribute 'estimated repair time' for the specific branch of the specific decision maker.

Observing the models $t_{DM1}{}^{s}$, $t_{DM2}{}^{s}$, $t_{DM3}{}^{s}$ it is clear that, although most of the rules are consistent with the general RI/SB trade-off described in section 5.4, the strategy represented by each tree is quite different. The set of machines that are always repaired immediately is significantly different across the three decision makers. The splitting values of each numerical attribute vary across the three models and the number of times that the attribute 'number of heads in the buffer' is used in each model is different.

There also rules which indicate that for some machines it is always better to operate them in SB mode once they fail during the shift. These rules are not consistent across the three decision makers, who apparently have very different views about which machines must always be operated in SB mode once they fail during the shift. Analysing the estimated repair time in the incidents in which the 'Always SB' machines are involved, it is concluded that the 'Always SB' rules might not fully reflect the decision making process for these machines. This is because the estimated repair time for the incidents in which the 'Always SB' machines are involved is relatively high. This might have influenced the decisions for these machines.

Discussing the validity of the 'Always SB' rules with the production manager and with an independent senior maintenance engineer at Ford, they were quite happy to accept that some decision makers may decide to operate 'Always SB' policies for specific machines. They insisted that an 'Always SB' policy is feasible and it can be very efficient for machines which cause bottlenecks in the production process. They also explained that this policy is more likely to be efficient for machines located near to other machines which require an operator. This is because there are times when the operator is not busy and so he might be involved in the SB process.

Based on the above discussion and taking into account that 'Always SB' represents incidents involving machines with a generally high estimated repair time, it was concluded that the 'Always SB' rules might, after all, represent a decision maker's strategy for handling and minimising the causes of bottlenecks.

On the basis of the above conclusion and taking into consideration the purpose of the research, it was decided not to modify the above rules, assuming that the 'Always SB' strategies are genuinely complete.

In addition, observing the tree $t_{DM1}{}^{s}$, it is clear that there is a rule (figure 7.2) that recommends SB for breakdowns with an estimated repair time of less than 21.5 minutes while it recommends RI for machine breakdowns with an estimated repair time of more than 21.5.

```
          ┌── ERT<21.5  ──────────────── SB, Freq 1, Prob1
Op1050    │
          └── ERT>=21.5 ──────────────── RI, Freq 1, Prob1
```

Figure 7.2: Inverted rule in $t_{DM1}{}^{s}$

This rule is not consistent with the RI/SB trade off and there were concerns that this inconsistency might have been caused by a specific decision maker who might not have been consistent during the data collection. There were also concerns that it might have been caused by the algorithm that failed to identify the other attributes which explain this behaviour. Analysing the incidents that generated this rule and bearing in mind the discussion with the decision maker, it was concluded that this 'inverted rule' is genuine. This conclusion is based on the fact that the decision maker was able to justify this decision by explaining that, on a specific machine, lengthy repairs off shift must be avoided since they interfere with other activities, such as planned maintenance and safety checks. As the above rule reflects the views of a decision maker on how unplanned maintenance must be performed on a specific machine, it was decided not to amend it.

Having validated qualitatively the above models, it was concluded that the models $t_{DM1}{}^{s}$, $t_{DM2}{}^{s}$, $t_{DM3}{}^{s}$ contain less serious knowledge gaps and the priority with which the attributes are used to classify the decisions is consistent with what the decision makers told us. On this basis it was concluded that they are a more robust representation of the decision making process.

The models could have also been validated using the 'take one out' approach (UrbanHjorth 1994). This validation strategy would involve calibrating 63 models (as many as the data records) by excluding a different record each time. The excluded data record could then be used to test the model. Due to the stratified approach that was used to generate the decision situations and due to the fact that most of the branches in the trees have been derived from unique decision situations (branches with frequency of situation =1), it was decided not apply this calibration approach since this would cause knowledge gaps.

In KBI terms, each decision tree t_j^s represents the relationship f_j that is a model of the real relationship Φ_j which associates the attributes of the system with the decisions taken by the decision maker j.

7.4 Evaluation and analysis

Having described the process of modelling human decision making using rule induction, as part of the evaluation of the third stage of the KBI methodology the following paragraphs explain the lessons that have been learned. Using the collected data, the interpolation capability of the rule-based models has been assessed by examining their tree structure. Some of the rules which have been derived have shown that the interpolation capability of the decision trees is highly dependent on the completeness of the input data. There are concerns that some of the 'Always SB' rules may not be genuine and may have been caused by the limited capability of the algorithm to interpolate. Although it has been decided not to amend those rules, since they have been qualitatively accepted with the decision makers, it is believed that further investigation with additional data could improve the analysis.

Given that the collected data is a representative but not a complete set of scenarios, the conclusion concerning limited interpolation indicates that the AI models produced in this research represent only some aspects of the decision making strategy employed by each decision maker.

The fact that the model uses only as many attributes as it needs to classify the decisions in the data set correctly is an important restriction. It does not allow the user to see how the attributers which are not included in the model affect the decision maker and it reduces the robustness and generality of the model.

Despite the above limitations, having compared the decision making strategies represented by the models with the verbal explanations given by decision makers, it has been concluded that the models represent some of the key aspects of the decision making process. The machine number is used as a first rule to filter out the machines which cannot be operated in stand by mode. This shows that the model is capable of preventing unrealistic decisions from being taken. The 'estimated repair time' which, according to the decision makers, is the main variable that is taken into account when deciding what action must be taken is the second attribute used in the models.

Taking into account that each of the six models developed using rule induction predicted correctly all the decision in the data set which was used for their calibration, it is concluded that with rule induction it is possible to capture the individuality of specific decision makers in the decision making problem analysed in this research. Although the quality of the predictions was not tested using data excluded in the calibration set, the qualitative validation in which the decision makers were involved has shown that there is evidence to indicate that the models represent generically a significant percentage of the strategies followed by the decision makers.

Using the tree representation of the rules, it is relatively easy to understand how the value of an attribute influences the decisions in the models. With the tree representation of the rules it is possible to identify the differences in the strategies followed by the alternative decision makers. This is a particularly useful feature when the efficiency of each model is assessed with simulation, since it enables the attribution of why some strategies are more efficient than others.

The hierarchical structure with which the rules are organised in the tree reveals the classification power of each attribute in the model when this is derived using automatic rule induction. Observing the hierarchical structure, it was possible to identify that, in one of the models, the classification power of the attribute 'estimated repair time' is higher than the classification power of the attribute machine number.

The validation process involved modification of the initial models using semi-automatic rule induction, incorporating knowledge that was gained during the presentation of the initial models to the decision makers. From the above it can be concluded that rule induction can be used to analyse and understand the mental process which links the attributes with the decisions and, therefore, can be used for external as well as internal analysis (section 2.5.3) and for understanding how the decision makers take decisions.

Amongst other benefits, this contributes to the identification of aspects of the quality of the strategies which cannot be assessed using quantitative techniques such as simulation. For example, as discussed in the previous section, analysing the tree structures in this research it was possible to conclude that DM1 takes safety and production quality risks when he decides SB (SB for some machines means that testing is not undertaken) for machines about which other decision makers are reluctant to take similar decisions, due to the risks that are associated with this type of decision for the specific machines. It is also concluded, despite the difficulties in assessing the statistical importance of each attribute in the tree, that the relative importance of each attribute can be assessed by noting its position in the tree. The participation of the analyst in the decision about the sequence with which the attributes will be used when semi-automatic rule induction was used supports the inclusion of such theoretically important variables as the machine number.

The modification of the tree of DM1, with the aim of removing the knowledge gap, shows that with the tree structure of the models it is possible to modify and extend the models by adding or removing branches without the need for recalibration. Finally, given that the calibration time for each model was less than one minute, it is concluded that computational effort and time does not appear to be a significant constraint when rule induction is used.

7.5 Conclusion

In this chapter the process of modelling the human decision makers using decision trees and data sets which come from simulation-based data collections has been described. This has tested the third stage of KBI.

The models which have been developed have been presented to the decision makers and have been qualitatively validated. The structure and the sequence with which the attributes are used in the models have been discussed with the decision makers and it has been concluded that they are an adequate representation of the decision making process.

The models that have been developed do not form a complete representation of the decision making strategies, but they provide a tool for understanding differentiations in the strategies employed by alternative experts. Using the decision trees, the individual characteristics of each decision making strategy can be represented and, as will be explained in detail in the next chapter, their efficiency can be assessed using simulation.

Chapter 8

Evaluation of stages 4 & 5 of the Knowledge Based Improvement methodology

This chapter tests and evaluates the final two stages of KBI. These stages involve the process of representing and assessing the decision makers by linking the rule-based models with the simulation. Section one describes the process of representing the decision actions in the simulation. Section two describes the design of the simulation experiments that were undertaken in order to assess the decision making strategies and section three reports the results from these experiments. Section four deals with the statistical tests that were used to assess the validly of the findings. Section five provides an overview of KBI stage 5 and shows how the conclusions from the assessment of the decision making strategies can be used to improve the decision making process in the operations system. Finally, section six evaluates the application of the final part of the KBI methodology.

8.1 Representation of the decision making process in the simulation model

In order to represent the decision making process, it was decided to control the logic of the actions that are executed when a machine breaks down by linking the simulation model with XpertRule, which contains the decision strategies. Although it was possible to hard code the rules in Witness, it was decided to separate the logic of the decision making process from the rest of the simulation in order to improve the flexibility of the integrated knowledge-based simulation model. Representing and controlling the logic of the decision makers externally, using a third party application, enables experimentation with alternative decision makers and with alternative decision making strategies. In addition, it makes validation and verification of the model easier and enables the modification of an existing decision rule without the need for editing the simulation model.

Implementation of decision actions in Witness

In order to use the decision making model to control the logic of the actions that are executed when a machine breaks down, the first step was to implement in the simulation every alternative action that the decision making models can predict. Given that the rule-based model predicts decisions using a binary decision variable (repair immediately is denoted as =0 and stand by as=1), these two actions were implemented in the simulation.

The 'repair immediately' action was implemented using conventional machine breakdown functionality, which is provided with the simulation software. When a machine breaks down, if the decision is to repair immediately then the state of the machine changes to breakdown. The machine is unable to process the parts, causing the length of queue to increase as more parts are arriving and require processing. The time for which the machine is in the breakdown state (the 'repair time') is an input in the model and, as will be explained in detail in section 8.2.1, it is determined by sampling from an empirical distribution. The resource which, if available, is engaged in order to fix the machine is one of the qualified engineers (Mechelec). If this type of resource is not available the machine remains in breakdown mode until the appropriate resource becomes idle. This representation is consistent with the rule that when a machine is repaired immediately it is the qualified engineers who are mainly involved in the repair.

If the decision is stand by, a dummy machine is used to represent the process of by-passing the parts by asking a trainee engineer (IMS) to process the parts that are arriving. A set of dummy machines has been used to represent the whole by-passing process. Each element of the set corresponds to each of the machines in the model. The representation of the by-passing process, using a different dummy machine for each 'actual' machine in the model, enables the representation of the fact that the time required for the by-passing process depends on the type of the machine.

Figure 8.1 provides an example of the logic that is used to represent the by-passing process. This logic has been implemented in each machine of the model (in the segment of the production line that is modelled in this research). As is shown in figure 8.1, if a machine $\mathbf{m}_i$ has broken down and if the decision is SB, then the part is sent to the dummy machine $\mathbf{m}^d{}_i$ and not to the actual machine $\mathbf{m}_i$. Once the part is processed, it is forwarded to the conveyor C which then pushes it to the next machine $\mathbf{m}_{i+1}$ for the next stage of processing. If the next machine $\mathbf{m}_{i+1}$ has not broken down, the parts are processed by that machine. If the next machine has broken down and is in stand by mode then the parts are sent to the dummy machine $\mathbf{m}^d{}_{i+1}$.

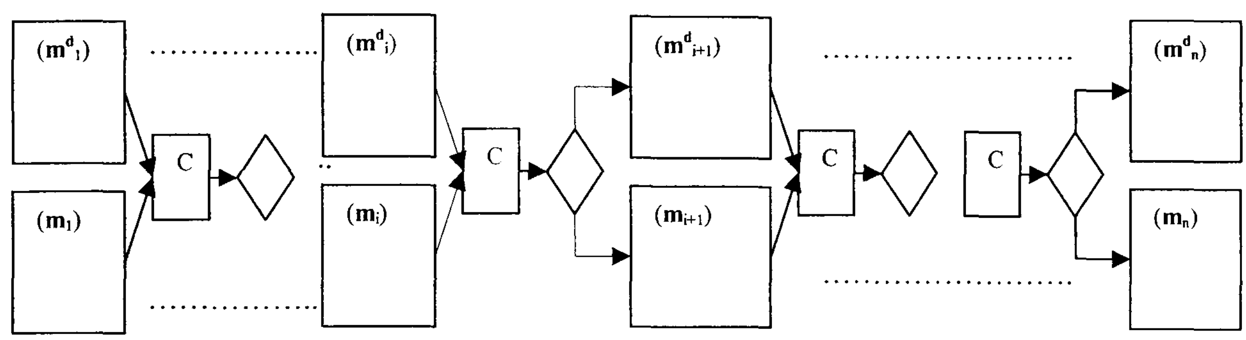

Figure 8.1 Model representation of the by-passing process

In the by-passing process, while the parts are processed manually by the trainee engineer, the status of the actual machine is broken down. Due to safety rules it is not possible to repair a machine while parts are being processed. The repair cannot occur, therefore, until the end of the shift. Once the machine repair starts the 'by- passing' process is terminated and the dummy machine becomes idle. In order to represent the fact that manual processing is usually more time-consuming, the cycle time of the dummy machine is greater than the cycle time of the standard machine (by 10%).

Linking XpertRule with Witness

Having implemented the modelling of the action that each possible decision option represents then, in order to control which action will be executed when a machine breaks down, the simulation model was linked with Xpertrule. Since the OLE2 automation functionality which Witness supports allows third party applications to invoke and control the simulation, although it does not allow the simulation itself to invoke and control third party applications, a communication interface was developed to link Witness with Xpertrule. This interface was developed in Visual Basic and it is an extension of the approach described by Robinson et al (2003b). The interface (figure 8.2) has the form of an executable (exe) application and it is invoked by the simulation software when a machine breakdown occurs. This approach makes the simulation more efficient, resolving problems described by previous authors (Robinson et al 2003b). The Visual Basic application does not run parallel with simulation and so there is no need to check whether a machine has broken down every time that a simulation event is executed.

When a machine breaks down in the simulation, the execution of the simulation events stops. The interface application described above is invoked using the Witness 'Appexec' function. This function enables the opening of an external application from Witness but it does not allow control of the external application. For this reason, once the communication interface is launched it takes control of the simulation events. It retrieves the data about the value of each attribute required by the rule-based model in order to make a decision and it passes them to XpertRule (Appendix 2 - Figure A2.1). Using the appropriate function, it requests a decision from XpertRule (Appendix 2 - Figure A2.2). Once a decision is taken from Xpertule, the communication interface passes this decision to Witness by assigning the appropriate value to the variable that is used to represent the decision (Appendix 2 - Figure A2.3).

The last instruction that the interface sends to the simulation before the Visual Basic executable shuts down is to continue the simulation process (Appendix 2 - Figure A2.4). Once this instruction has been sent to the simulation, the action which must be executed is decided by checking the value of the local (Witness) decision variable. If this is =1, the parts are sent to the dummy machine for processing and this represents the stand by process. If the decision variable is =0, the parts are waiting for the machine to be repaired and this represents the immediate repair.

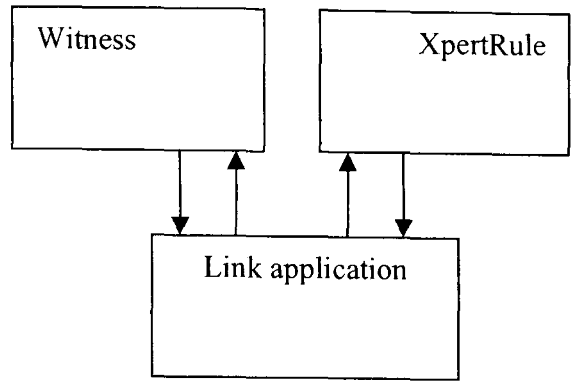

Figure 8.2: Witness XpertRule link

Verification

Having developed the interface as part of the verification process - the process of testing the software implementation (Pidd 1998, Robinson 1999) - the system was verified by comparing the simulation output from the base simulation model (section 1.3) with the output from a simulation run of the knowledge based simulation with an XpertRule decision strategy, reflecting the strategy which is hard coded in the base simulation model i.e. always repair immediately. Taking into account that the output from these simulation runs was identical, it was concluded that the knowledge based simulation does what is expected.

8.2 Experimental design

Having verified the new model, with the aim of testing the use of simulation to identify efficient decision making, a number of simulation experiments were undertaken. This section provides an overview of these experiments. It outlines the policies that were simulated, the implicit assumptions that were made and the decision scenarios that were used. It explains how the decision scenarios were generated, how the simulation run length was determined and which attributes were used for the scenarios.

8.2.1 Policies compared in the simulation experiments

To set a benchmark for comparison and to compare the base decision making process (section 1.3) against the representation that takes into account the decision making strategies of each individual, it was decided that the policies which should be assessed during the experiment must include the decision making process that assumes that repair immediately is the only option available when a machine breaks down. On this basis, in order to compare the efficiency of the strategies identified in section 7.3.1, it was decided to simulate each of the following four decision making policies:

- Base model: Current representation of decision making
- Strategy DM1: Decision tree $t_{DM1}{}^{s}$ derived from DM1's data
- Strategy DM2: Decision tree $t_{DM2}{}^{s}$ derived from DM2's data
- Strategy DM3: Decision tree $t_{DM3}{}^{s}$ derived from DM3's data

Given that the manufacturing process is represented in the model as a non-terminating system, it was decided that for each simulation run a warm up period must be used in order to enable the simulation to reach a steady state. Inspecting the time series (visually - Pidd 1998) from pilot runs, it was found that 500 minutes were required for the simulation to reach a steady state. On this basis, it was decided that the throughput from the first simulated day must not be used as a performance indicator for assessing the decision making strategies. Having excluded the first day, the simulation length for each run was determined by taking into account time and hardware considerations. Considering also the scale of the simulation model, it was decided that an overnight simulation run should be undertaken for each decision making strategy. This was the equivalent of 112 simulation days (replications) and, from the repeated pattern in the time series of the observed throughput, it was decided that it was sufficient for the purpose of the experiment.

8.2.2 Decision scenarios and their inter-arrival times during the simulation run

The decision scenarios (machine breakdowns) during the simulation runs were generated by combining attributes of the simulated system. The inter-arrival time between two decision scenarios was sampled using a set of negative exponential distributions. The negative exponential distribution was chosen because, compared with alternative distributions, it is theoretically the most appropriate representation of inter-arrival times

(Robinson 2004, Pidd 1998). Assuming that the inter-arrival time of machine breakdowns on a specific machine is independent of the state of the other machines, a different distribution parameter and a different random number seed was used in each machine in the model. The parameter in each distribution was determined by fitting a negative exponential curve in data collected from inter-arrival times of breakdowns for each specific machine.

8.2.3 Decision attributes

In order to simulate the decision strategies $t_{DM1}{}^{s}$, $t_{DM2}{}^{s}$, $t_{DM3}{}^{s}$ during the simulation run for each decision scenario, the simulation provided information about each of the following attributes:

- Estimated Repair time
- Number of heads in the buffer
- Number of machine (Machine ID)
- Current Simulated time

Estimated Repair time

The 'estimated repair time' for each breakdown for each machine is an input in the model and it was determined by sampling from a machine-specific empirical distribution that was developed using historical data. Using recently observed machine-specific repair times in each machine's empirical distribution, it was assumed that the repair time for each machine is stochastically distributed, and it depends on the age of the machine and on the complexity of the operation which the machine performs. The empirical distributions are also used in the existing simulation model (section 1.3) and these were developed prior to the research by a simulation specialist in Ford. According to Ladbrook (1998) these distributions have been calibrated during the validation of the existing model and, compared with statistical distributions, they lead to more accurate simulation prediction. Although the accuracy of the prediction is not the only criterion for assessing the validity of the model, in order to achieve a like for like comparison the same empirical distributions were also used in the simulation models that are compared in this research.

Number of heads in the buffer

The attribute 'number of heads in the buffer' is determined by the simulation using a Witness object known as 'buffer'. This object keeps a record of the number of heads that have been produced and the number of heads that have been used during the simulation run. The 'buffer' object is dynamically updated during the simulation and it can provide information on the fly about the current availability of heads in the model.

Machine ID and Current time

As discussed earlier in this section, the inter-arrival times of machine breakdowns are sampled from a set of negative exponential distributions which determine where and when the decision scenarios will occur during the simulation run. This means that the attributes 'machine id' and 'current simulated time' are determined within the simulation from the machine breakdown inter-arrival times. As has already been noted, a different distribution and different random seed is used in each machine. This ensures that the above attributes occur independently in the decision scenarios.

8.3 Simulation results

Running the simulation four times for 112 simulated days with a different decision making configuration each time, the results that are shown in figure 8.3 were collected. It must be noted that, for confidentiality purposes, the time series of the numbers on the Y axis have been multiplied by a random multiplier to avoid disclosure of sensitive information.

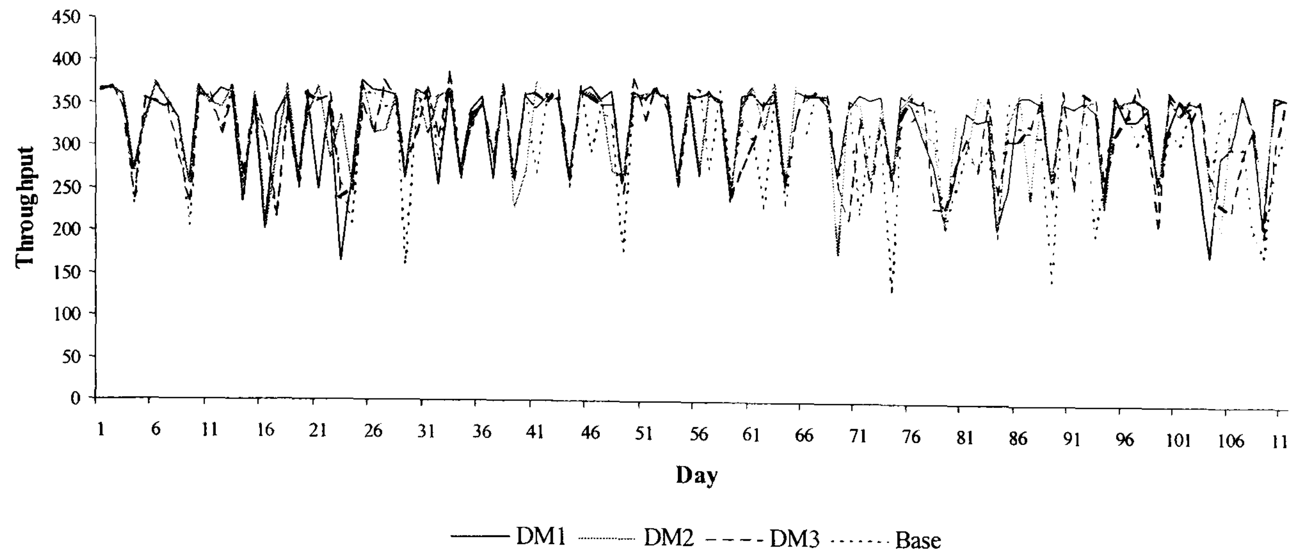

Figure 8.3: Simulation prediction - throughput levels under different decision making policies

Figure 8.3 shows the daily throughput for each decision making policy. Although none of the four strategies seems to perform markedly better than the others, it is clear that there are days when the strategy employed by the base model performs worse than the strategies that model the policies followed by the human decision makers.

The misalignment of the time series of the throughput of each decision making model reveals the degree to which representing the decision making policies in the simulation has an impact on the prediction of the model. This is also confirmed from the mean daily throughput figures that are reported in table 8.1.

	DM1	DM2	DM3	Base
Mean daily throughput	325.53	318.87	325.15	312.06
Standard deviation	36.28	34.34	28.54	39.58

Table 8.1: *Simulation prediction - mean daily throughput and standard deviation*

According to the mean throughput, the most efficient policies are those employed by DM1 and DM3. Based on the simulation results, the strategy employed by DM2 leads to lower daily throughput levels and the least efficient strategy is the one that is represented in the base model. The reduced performance of the strategy employed by DM2 is attributable to the relatively lower number of SB decisions that are recommended from the rules used in his model. Table 8.2 shows that the percentage of machine breakdowns that are handled with stand by decisions is relatively lower for the decisions taken by DM2. Taking into account that DM1 and DM3 perform better than DM2, this shows that the decision making approaches which involve strategies for minimising the delays caused by machine breakdowns (i.e. strategies with a relatively higher percentage of stand by decisions) lead to higher throughput compared to strategies in which the delays caused by machine breakdowns are not handled.

	DM1	DM2	DM3	Base
SB decisions as % of all decisions	51%	16%	46%	0%
RI decisions as % of all decisions	49%	84%	54%	100%

Table 8.2: *% of Stand by versus Repair Immediately decisions*

8.4 Statistical validation

In the previous section it was concluded that the representation of individual decision making strategies may have an impact on the prediction of the model. This indicates that DM1 and DM3 have the most efficient strategies. To assess the validity and the statistical significance of the above conclusions, two types of hypothesis were tested. First of all, in order to assess the statistical significance of the variation in efficiency that was observed across the three alternative human decision making strategies, the confidence intervals of the differences between the results for the three strategies were calculated.

To ensure that the observations are not auto-correlated, Fishman's procedure (Fishman 1978) was applied and the mean daily throughput and the corresponding standard deviations were calculated using batches of two observations. Calculating the confidence intervals for the difference in the mean daily throughput that was predicted from each policy, the hypothesis that there is no difference between each pair of the decision making strategies was tested by examining whether the confidence intervals include the value zero. To calculate the 95% confidence intervals, a significance level of 1.67(5%/3) was used in order to account for the Bonferroni inequality and for the fact that three confidence intervals were being calculated (Law & Kelton 2000). Normalising the standard error using a standard normal value of 1.67 (and not 1.96 for 95 % Confidence Intervals), the Bonferroni inequality ensures that the confidence level is not reduced. Had the adjustment of the standard error remained at 1.96, the significance level would have deteriorated. It would give less than or equal to a 5% chance for one of the true differences to fall outside the confidence intervals, but less than or equal to a 15% chance for the three true differences to fall outside the intervals.

From table 8.3, which shows the confidence intervals of the difference in mean daily throughput between each pair of decision makers, it is concluded that the variation in the throughput due to the application of a different human decision making policy is not significant since zero is included in the confidence intervals. Although this conclusion indicates that the difference between the policies is not significant, it is of limited validity. This is because, in this research, the decision making strategies were only modelled in a specific segment of the production line, while the throughput which was predicted and which was used for the statistical assessment represents the performance of the whole production line. Had we modelled the decision makers in the whole production line, the differences in mean throughput would have been higher and could have been statistically significant.

	DM3	**DM2**
DM1	-8.43, 9.19 (0.38, 26.45, 1.67,55) (no sig. difference)	-3.80, 17.11 (6.66, 31.37, 1.67,55) (no sig. difference)
DM3		-2.11, 14.66 (6.28, 25.17, 1.67,55) (no sig. difference)

Table 8.3: *98.33% paired-t confidence intervals for differences in daily throughput between each decision making strategy. Values in brackets represent mean throughput difference, standard deviation, standard normal value used for calculating the confidence intervals and sample size.*

In order to assess the statistical significance of the impact of representing the decision making in simulation, confidence intervals were also calculated for the difference between the mean daily throughput predicted for each human decision making strategy and the throughput that was predicted from the base model.

	Base
DM1	0.84, 26.10
	(13.47, 37.91, 1.67,55)
	(DM1>Base)
DM2	-4.04, 17.66
	(6.81, 32.57, 1.67,55)
	(no sig. difference)
DM3	1.21, 24.98
	(13.09, 35.67, 1.67,55)
	(DM3>Base)

Table 8.4: *98.33% paired-t confidence intervals for differences in daily throughput between each decision making strategy and the base model. Values in brackets represent mean throughput difference, standard deviation, standard normal value used for calculating the confidence intervals and sample size.*

Calculating the above confidence intervals, it was possible to test the hypothesis that repair immediately is always the best option when a machine breaks down. From the results that are reported in table 8.4, it is clear that we can reject the null hypothesis. It is possible to confirm statistically that repair immediately is not always the best option. For two out of the three human decision making strategies, the output that was produced is significantly different from the output that was produced employing the base decision making strategy. Clearly, the above conclusion shows that representing human decision making in the simulation model has a statistically significant impact on the prediction. This suggests that, in a simulation model of the engine assembly plant which aims to assess plant performance in the face of machine failures, an appropriate representation of maintenance strategies is important.

8.5 KBI Stage 5: Seeking improvements

Although, due to the limited time scales, full implementation of stage five is beyond the scope of this thesis, from the above conclusions it is clear that a number of improvements can be achieved in the operations system modelled in this research. The simulation assessment indicates that a statistically significant improvement can be achieved by employing a strategy that involves stand by policies. Taking into account that the throughputs predicted by applying the policies DM1 and DM3 appear to be the most desirable of the four that were assessed, it is clear that throughput improvements can be achieved by making one of these policies standard practice. In addition, although it is beyond the scope of this thesis, having identified that DM1 and DM3 are the most efficient strategies, it might be possible to achieve throughput improvements by designing and assessing a strategy that combines elements of these two policies. Finally, given that DM1 and DM3 are currently the best strategies that have been identified but not necessarily the best strategies that exist, it is proposed that throughput improvements can be achieved by using one of these strategies as a starting point for an incremental heuristic search. This might lead to a policy with higher utilisation rates that would allow the system to reach its full potential. The heuristic that could be implemented using a local search would involve the maximisation of the following production function:

$$Y = \left[\sum_k Y(t_l)_k \right] / K \qquad 8.1$$

In expression 8.1, Y is the mean throughput from a simulation run of k days that uses a tree t_l. The heuristic in a KBI context would involve the identification of the tree t_l^* that, in an iterative search which starts from the tree t_{DM1} or t_{DM3}, locally maximises Y. The search would involve a number of iterations. At each iteration an incremental alteration could be applied in the tree t_{DM1} or t_{DM3} and simulation runs would be required in order to assess the performance of the new tree. If the new tree leads to a higher simulated throughput, then this could form the current best solution and the next move would involve alteration and assessment of this tree. In the context of the case study, each heuristic move would involve modification of the values in the rules which determine the length of the estimated repair time after which a specific machine must be operated in stand by mode. Given that lengthy simulation runs would be required for each move, only a tiny percentage of the alternative possible solutions can be assessed. To make sure that the local optimum is identified, a rejection method such as simulated annealing (Eglese 1990) or taboo search (Wright 1996) could be applied in order to allow diversification in the search. Applying rejection techniques in the search and by so doing allowing iterations with relatively bad solutions, it is expected that the search will identify a near-optimal solution (Reeves 1995).

8.6 Evaluation and analysis

With the discussion on the potential improvements that can be achieved, the previous section has completed the description of the implementation of KBI. As part of the evaluation process in this section, the approaches that can be used for validating the conclusions are described and form the basis for the discussion and the conclusion of this thesis.

As explained in section 3.2 and again later, during the implementation of the stages of the research, an iterative and incremental validation strategy was employed in this research. The decision making models were validated using qualitative knowledge that was gained during the interviews and discussion with the decision makers. In addition, these models were presented to the decision makers in order to see to what extent the strategies that are represented in the models are consistent with the decisions that the decision makers take.

The impact of representing the decision making strategies in the simulation, as described in section 8.4, was initially validated by assessing the statistical significance of the difference between the throughput of the models. A more straightforward validation strategy would have been to compare the output from the simulation models with the decision making representation against real observed data, as is suggested by Moffat et al (2004). This kind of validation was not possible for a number of reasons. Firstly, daily throughput data is rather sensitive information which Ford would not share with us. Secondly, specific roster information about the personnel who were involved in each shift is even more sensitive information which is safeguarded by the Data Protection Act, due to implications which this might have for the performance assessment of the engineers. Thirdly, even if we could get hold of this information it would not be particularly useful, since the research focussed on only one segment of the production line.

Taking into account the above limitations, the simulation model was validated by considering to what extent it is sufficiently accurate for the purpose of the research (Robinson 2001). Using the simulation model, it was possible to compare the efficiency of the alternative decision making strategies (section 8.4). The decision trees provided a tool for understanding why some strategies perform better than others and the simulation assessment enabled the identification of the most efficient decision making policies. This has enabled the recommendation of improvements in the decision making process (section 8.5). Considering that the purpose of the research was to test the use of simulation for understanding and improving the design of decision making, the above findings have shown that the simulation model is sufficiently accurate for the purpose of the research.

8.7 Conclusion

In this chapter, the process of linking the representation of the human decision making strategies with the simulation model has been described. Using the simulation model, the various decision making strategies have been compared and it has been found that none of the human decision makers is statistically more efficient than the others, yet the strategies that involve SB policies are statistically more efficient than the strategy that does not use SB. Comparing the output of the base simulation model with the output of the models with the representation of the decision making strategies, it has been found that representing decision making in simulation has a significant impact on the prediction of the model. Finally, adopting the validation approach, which suggests that a model is valid if it serves the purpose for which it was developed, it has been possible to evaluate the findings of the final two stages of KBI.

Chapter 9

Discussion and refinements of the KBI methodology

The last four chapters have tested and evaluated the stages of the KBI methodology in an industrial application. As part of the third research objective (section 3.1), this chapter contributes to the refinement of the methodology by discussing the main conclusions and by summarising the lessons that have been learned. Section one discusses the conclusions from the evaluation of the first two stages of the KBI methodology (problem understanding and the data collection process). Section two discusses the findings from the evaluation of the third stage (the use of AI) and section three discusses the conclusions from the evaluation of the final two stages of the methodology (simulation assessment of the decision making strategies). Finally, in section four the benefits and the general limitations are discussed.

9.1 KBI stages 1 and 2

As explained in section 3.1, as part of the second objective of the research, the research questions 1 and 2 were set out in order to investigate the feasibility and the practical difficulties of a VIS-based data collection in an industrial environment. Using a deductive case study, part of this objective was also to explore what can be learned from a VIS-based data collection and what are the limitations of this process in an industrial environment (section 3.2.1). As will be explained in section 10.2, the conclusions from the implementation of the stages of KBI have addressed the above research questions. They have also contributed to the third objective of the research and to the refinement of the KBI methodology. The following paragraphs discuss the findings related with the first two stages of KBI and the section concludes by explaining the lessons that have been learned and the refinements of the first two stages of the KBI methodology.

9.1.1 Testing KBI stage 1

During the problem-understanding process (stage 1 of the KBI methodology), the main elements of the decision making problem were determined and the practical difficulties that may arise during the problem-understanding process in an industrial environment were identified and addressed. The main difficulties that the problem-understanding process involves were revealed from the initial discussion with the people who are involved in the maintenance process. During this discussion it was not possible to identify the decision making problem. It was not possible to identify any obvious process which involves human beings deciding from a set of alternative options by taking into account the status of the system. The initial response to our question about what do you do when a machine breaks down was '.... we fix the machine as soon as we can'. The same response was also given when they were asked what if more than one machine needs to be repaired urgently. To identify the decision making problem we had to go through an iterative process which progressively improved our understanding of the problem. As has been explained in detail in Chapter 5 (section 5.2), the main decision variables (What and Who) were identified only when a three day visit to the factory took place. The options that are considered for each decision variable and the attributes that are taken into account were later identified when a series of iterative simulation-based pilot data collections were conducted. Four pilot data collections, supported with the use of a VIS that was iteratively re-specified, have been required in order to gain a clear understanding of the problem. From the problem-understanding process a number of lessons have been learned and these have contributed to the development of a detailed approach towards structuring decision making problems.

From the problem-understanding process that was described in Chapter 5, it is clear that the decision makers were unable to explain the decision making problem to us verbally, although it is clear that they were actively involved in it. This illustrates that human beings sometimes do things that they have not questioned. This might be due to the fact that their decision making duty is part of a more generic role and, therefore, their role as decision makers was thought to be a rather trivial one which, in their minds, was underestimated and essentially removed from their conscious awareness of the list of things which their role involves (tacit knowledge).

The difficulties that the decision makers had with describing the problem revealed the challenge involved in the process of gaining a detailed understanding of the decision-making problem through directly asking the decision makers. The process of re-specifying the VIS iteratively, to include or exclude certain attributes in order to identify those that the decision makers need in order to make decisions, shows that a simulation model can support the problem understanding process. Given the uncertainty about which attributes the decision makers take into account when making decisions, the model must allow the user to receive additional information upon request. Recording the frequency of times that information about a certain attribute is requested might contribute to the identification of the attributes which influence the decisions.

9.1.2 Testing KBI stage 2

The collection of the decisions data (stage 2 of the KBI methodology) was also an area in which practical difficulties were identified. Unlike most of the previous research in modelling decision making in simulation, where the decisions data were collected from hypothetical decision makers who have the time or incentives to participate in lengthy data collections, in this research the decisions were collected by asking real decision makers to interact with the data collection system. Clearly, the involvement of the real decision makers in the research imposed many constraints on the number of sessions that were able to be conducted and, therefore, on the quality and quantity of the data sets which were collected. However, it was the involvement of the real decision makers and the limitations on the size of the data sets which contributed to the development of the data collection approach that is reflected in KBI stage 2. The limitations resulting from the involvement of the real decision makers provided the motivation for designing a data collection approach in which a representative set of decisions is inferred from a limited data collection session. This data collection approach addresses the issues of involving real experts raised by Liang et al (1992) and shows that it is possible to use simulated decision situations for collecting decisions from real experts.

The practical difficulties of collecting decisions data from real decision makers using VIS were identified shortly after the second pilot data collection in which 24 decisions were collected within the available time, all of which were of the same type (repair immediately). At that stage of the research, it was clear that the lack of diversification was due to the limited size of the data set and the nature of the decision situations which were presented to the decision maker. The manner of resolving this problem was not obvious, given that it was not possible to ask the decision maker to be involved in lengthy data collections. The solution to the problem was to use a non-visual simulation and to design an experiment containing significantly different decision situations in the hope that this would lead to a more representative data set. In order to obtain the so-called significantly different decision situations, as explained in Chapter 6, a large number of simulated decision situations were analysed. Based on this analysis, it was possible to generate decision situations that enabled the collection of decisions which contained sufficient variability to develop sensible models of decision making. This shows that, in the specific application, simulation provided the facilities to design controlled experiments which enabled the analyst to collect representative data sets within realistic time scales.

VIS was used to present the decision situations to the decision makers in some of the pilot data collection experiments. According to what the decision makers told us, the schematic visual display of the simulation did not add any further information that helped them to make decisions. Running the simulation during the data collection session imposed many constraints and restrictions about the type of decision situations that were presented to the decision makers during those experimental data collections. Clearly, the most appropriate data collection approach was the one that involved the use of a logical display that did not require running the VIS during the data collection session. The main reason is that when we had the opportunity to run the simulation in advance, we had the chance to select a subset of situations which represented a wider range of decision scenarios. From the above it is concluded that, for the specific case study, the dynamic schematic display was not particularly appropriate for presenting the decision situations to the decision makers, principally because this type of display provides details about the general status of the system and this was not taken into account by the decision makers when they made decisions.

9.1.3 Evaluation of stage 1 and 2

Having discussed case specific conclusions about the implementation of the problem-understanding and data collection stages of KBI, the following paragraphs describe the general refinements, the limitations and the lessons which can be learned as a result of the implementation of the first two stages of the methodology.

Data collection process

From the process that was described in Chapters 5 and 6 it is clear that, although the problem-understanding and data collection in KBI are treated as two distinct stages, the empirical evidence from the implementation of the methodology has shown that these two are actually parts of the same process, which has emerged from the implementation of the initial two stages of KBI. This is because the problem-understanding process involves iterative pilot data collections. Addressing the need (highlighted by Curram 1997) for a model building methodology, this process involves eight steps and has led to a new approach to problem structuring and data collection, known as KBI stages 1&2.

From the development and the application of KBI stages 1&2 it has been learned that, depending on the nature of the problem, sometimes it is more appropriate to use non-visual simulation. In order to proceed to the data collection the problem must be defined and a conceptual model should be developed, reflecting the scope of the modelling exercise and any simplifications that have been made. The conceptual model can be used to specify the requirements of the system which needs to be developed to support the data collection. After the development of the system, the data collection should take place by involving each decision maker.

Data collection system specification

It has also been learned that the functionality of the data collection system proposed as part of KBI stages 1&2 must allow the analyst to design the data collection by using simulation for generating and analysing decision situations. This contributes to address the need for an integrated intelligent simulation package, as highlighted by Williams (1996). It shows that when the availability of the experts is limited, the system must allow the analyst to elicit the decisions by simulating the system prior to the data collection, generating a trace of decision situations which can be stored and presented to the decision makers without the need to run the simulation during the data collection. To avoid presenting duplicated decision situations the system must use a filter, which combined with the use of stratified sampling should filter out situations similar to those that have already been presented to the decision makers. To enable the analyst to identify the attributes which the decision makers need to know in order to make decisions, the system must provide the facilities that allow the decision makers to retrieve information upon request.

Data collection limitations

The size of the data sets that can be collected using KBI stages 1 and 2 is perhaps the most serious limitation. Due to the nature of the approach, requiring the intensive involvement of the decision makers, it is unrealistic to assume that it is possible to collect data sets of any size. The size of the data set is constrained not only by the availability of the decision makers but also by the intellectual capacity of the human brain. In this research, it has been found that the continuous involvement of each decision maker in the task of inputting decisions in the system for more than two hours per day can bore the human decision maker, leading to a degree of fatigue which can have a direct impact on the quality of the inputs. Clearly, it is not realistic to assume that the decision makers will be able to be involved consistently for more than a couple of hours per day, even if they are willing to do so.

As has been seen in section 6.4, the nature and the number of attributes required by the decision makers in order to take decisions is also an important constraint limiting the scope of KBI stages 1 and 2. In order to use simulation to collect decision making data, the decision making situations which the system should be able to generate must include all the key attributes that the decision makers need in order to make decisions. Based on the empirical evidence from the implementation of KBI in this research, it is clear that there are situations in which this might not be possible. This is more likely to be the case when the decisions are influenced by attributes with a low level of detail, such as colour or temperature of parts of the equipment of the system which is being simulated.

The number of options that are considered in the decision making process is also an area that imposes limitations on the applicability of KBI stages 1 and 2. Given that the data set must include each option that is considered when a decision is required, an increase in the number of options that are considered significantly increases the size of the data set which must be collected. Taking into account the empirical evidence on the limitations in the number of decisions that can be collected using simulation, it is clear that KBI stages 1 and 2 are more appropriate for small to medium size decision making problems (problems represented using decision variables which reflect decisions with a small number of options two, three or possibly four).

9.2 KBI stage 3

As part of the second objective of the research (section 3.1), research question 3 was set out to investigate the feasibility of the third stage of the KBI methodology (the use of AI to model individual decision making in simulation). Using a deductive approach, the purpose of this research question was to test the theoretical developments related to the use of AI in simulation. As will be explained in section 10.2, the research has addressed this research question and, as part of the third research objective, it has contributed to the refinement of the third stage of the KBI methodology. The following paragraphs summarise the findings in relation to the third stage of KBI (the use of AI) and the section concludes by explaining the lessons that have been learned and the relevant refinements in the KBI methodology.

9.2.1 Testing KBI stage 3

Based on the conclusions in section 7.5, the research has shown that in the specific application, with certain limitations and assumptions, it is possible to use rule induction in an industrial environment to identify aspects of the decision making strategies which industrial experts employ during the production process. As explained in section 7.3, the strategies that were identified are not complete representations of the decision makers, however they do model the differences in the way in which they take decisions. This has enabled the simulation assessment of the decision makers and it has shown that, for the aim of the research, the aspects of the strategies that have been identified are appropriate to represent individuals in simulation.

Clearly, the above conclusions are accompanied by certain assumptions and limitations. In order to reduce the dimensionality of the data sets which have enabled the calibration of the models, it has been assumed that the repair time reflects all the required information that the decision maker collects from the physical inspection of the machine and from the type of fault. An estimation of the repair time has been used to represent the physical inspection of the machine and the type of fault that is reported by the diagnostic system during a machine breakdown. On this basis, it has been assumed that the decision makers are equally capable of predicting the repair time for any given situation. From the dimensionality reduction, it is also concluded that the limited data sets might also impose restrictions in the number of decision variables which can be modelled.

It is evident from the validation strategies that have been employed to test the validity of the decision trees that the conclusion on the capabilities of rule induction for identifying the decision making strategies are based on some assumptions. The qualitative validation involving the presentation of the models to the decision makers (and to other stakeholders in the problem) has provided some confidence that the models are generic. However, the limited interpolation capabilities and the knowledge gaps that might exist show that the models have their limitations. They are founded on the assumption that the sample collected during the data collection process is a representative set of the population of the decision situations which might occur during the production process.

9.2.2 Evaluation of KBI stage 3

Based on the above discussion, a number of methodological lessons associated with the process of identifying the decision making strategies in KBI can be learned by considering the implications of the above case-specific conclusions for the potential user of the approach.

The conclusions in sections 7.3 and 8.3 have shown that rule induction is appropriate for understanding the decision making strategies and for gaining insights into why certain individuals perform better than others during the simulation assessment. This is key information for the KBI approach proposed in this research. Considering that a fundamental objective in the KBI approach is to identify improved practices by comparing the alternative strategies and understanding why some of them perform better than others in the simulation assessment, it is clear that the tree structure gives a comparative advantage to rule-based models. Due to the algorithm that is used to induce the strategies, rule induction models can be qualitatively validated by facilitating a discussion with the decision maker who is involved in the data collection.

As in the case of the lessons learned from the implementation of the data collection and the problem-understanding stages, the lessons which have been learned from the process of identifying the decision making strategies have their limitations. Overfitting and knowledge gaps caused by missing decision situations in the data sets are an important limitation of the technique which might limit its application. If the data set has serious knowledge gaps, then the strategies will not be complete and qualitative validation should be used to test the extent to which the rules represent the strategies employed by the decision makers. There is clearly little scope to use rule induction if it is not possible to use the models to represent the rules which are implicitly included in the data sets.

9.3 KBI stages 4 and 5

As a final part of the second objective (section 3.1), research question 4 was set out to investigate the use of simulation to identify efficient decision making strategies. The research has implemented and tested the above process in a real context. The following paragraphs discuss the relevant conclusions, the lessons that have been learned and the relevant refinements of the KBI methodology.

9.3.1 Testing KBI stages 4 and 5

Chapter 8 described the process of linking the simulation with a rule-based expert system that was used to store the decision making strategies. In addition, it described the main conclusions from the implementation of the fourth stage of KBI which involved the process of comparing the performance of each decision making strategy using simulation. The implementation of the simulation assessment has contributed to the validation of KBI and it has provided empirical evidence on the feasibility and the benefits of the approach.

The conclusion from the simulation runs has revealed that, in this research, it has been possible to assess the performance of various decision making practices related to unplanned maintenance. From the conclusion of the comparison of the simulation results (throughput prediction), it has been found that the strategy apparently employed by most of the decision makers is significantly better than the base decision making strategy which involves only the policy RI. This demonstrates that, in the specific application using KBI, it has been possible to make recommendations regarding the nature of the policy that should be applied in order to improve efficiency. From the decision making strategies that have been identified and from the assessment of their performance, it has been found that the decision makers appear to follow quite different decision making strategies and that their performance may vary. This indicates that performance improvement may be achieved by spreading the knowledge held by specific decision makers.

9.3.2 Evaluation of KBI stages 4 and 5

Based on the above case-specific conclusions, it can be concluded that under certain circumstances commercial expert systems packages, such as the one used in this research, can be used to store, assess and maintain complex decision rules which are applied in simulations. Having applied the above idea in a commercial simulation model, the research has provided empirical evidence to support the feasibility of an approach for separating the logic of the simulation from the rest of the model. This empirical evidence addresses the need for testing the above approach in an actual industrial environment, as suggested by Flitman (1986).

Generalizing the case-specific conclusions described earlier in this section, it can also be concluded that the implementation of KBI in this research, using a simulation-based data collection and rule induction has shown that, under certain circumstances, simulation and AI can be used to identify, model and assess decision making practices in an industrial environment. In addition, from the conclusion arrived at from the comparison of the decision making strategies it is clear that, with an implementation of KBI which involves rule induction, it is possible to identify good practices that enable the analyst to make recommendations to the owner of the problem i.e. practices which are relevant and can improve the overall performance of the system.

The above conclusion highlights the fact that KBI is the only way to isolate and assess the impact of individual decision making strategies. This is because the operation of systems such as that investigated in this research involves many decision makers who act in parallel fashion. As a result, the performance of each decision maker cannot be fully assessed on the basis of the strategy that he/she follows since observed performance measures are affected by the strategies followed by other decision makers. Simulating the system under one individual decision making strategy appears to be the only way to assess the performance of specific strategies.

9.4 Scope and limitations of KBI

Taking into account the difficulties that have been described elsewhere in this thesis, it is clear that there are certain constraints and assumptions in the above discussion which might limit the use of simulation-AI methodologies in general and KBI in particular for identifying and assessing decision making practices.

It is clear that KBI is not yet another performance assessment alternative approach to well-established performance assessment techniques, such as Data Envelopment Analysis or Stochastic frontiers (Dyson et al 2001). KBI is an approach that is intended for the solution of different kinds of problems. Unlike established performance assessment techniques, KBI focuses on problems which involve assessment of non-homogeneous decision making units, such as decision makers who have their own individual characteristics and distinctive attitude and act within an environment with loosely pre-specified rules.

In contrast to all the other performance assessment approaches, with KBI the assessment of a strategy does not involve the use of historical performance data which might not reveal the full potential of a specific strategy. When KBI is employed, the assessment of the performance of a specific strategy involves the use of simulation where extreme case situations are bound to occur at some stage. This eliminates the risk of bias due to non-representative sampling of historical performance data.

KBI enables explicit derivation of the strategy that is assessed. This enables a better understanding of why a specific strategy performs better than others. KBI does not treat individual decision makers as production systems, whose performance is assessed based on the rate with which they can transform inputs into outputs. With KBI, the decision makers are treated as individual human beings whose interaction with the system is investigated based on the emerging behaviour that this can generate. KBI supports the assessment of qualitative performance measures such as the safety risk that the decision makers take. For instance in section 7.4 it was found that DM1 takes safety and production quality risks when he decides SB for machines. KBI facilitates a process of understanding key constraints (for example, in the case study of the research, the machines which cannot be set in SB mode) and it can be used to form strategic propositions and to provide initial solutions to a heuristic search.

KBI uses simulation to predict and assess the performance of human decision making strategies. It is clear that the methodology is more appropriate for situations where the decision makers have a single common objective directly related to the decisions that they make. Further to this, it should be possible to predict the performance of that objective by using the simulation model. Based on empirical evidence, there are concerns that this might not always be the case. This could be because the decision makers might have quality controls and moral-related objectives which cannot be assessed by the simulation.

From the conclusions in Chapter 8, it is clear that the nature of improvements which can be achieved using KBI come from the identification of good practices that can be communicated, facilitating a learning environment through knowledge management practices. Despite the fact that the recommendations which can be made and the improvements which can be achieved with KBI are less ambitious than those that are expected from optimisation approaches involving exhaustive search, the recommendations from KBI are more pragmatic and realistic. With the continuous involvement of the decision makers during the process of identifying the decision making strategies, with KBI it is more likely to identify key constraints (section 7.3) and to improve learning. These are two issues which are very often overlooked when optimisation techniques are applied.

From the description and the implementation of KBI, it is clear that the knowledge elicitation method which is proposed as part of the methodology is an experimental process, in which the decision makers who participate are involved in interactive data collection sessions that are facilitated by a system that generates, stores and presents simulated decisions situations. Using experimental data to identify the decision-making strategies, it is clear that KBI is based on the assumption that the decision makers who are involved during the experiments behave as if they were taking decisions that have real effects on the production system. It is also assumed that the decision makers during the experiments take the same risks and are equally responsible for their actions as when they are involved in real decisions. From the experience of implementing KBI, as was anticipated in Chapter 3 (section 3.3), it has been found that there is a risk of collecting unrealistic decisions, owing to the fact that the decision makers interact with a simulation rather than with the real system. This risk, known as the gaming effect, shows that the above assumption is an important one and so care should be taken by the potential user of KBI to make sure that is not violated. Constraint violation checks of the input decisions (such as those described in sections 7.3.1 & 7.3.2) are one way to eliminate this problem. Increasing the awareness of the decision makers about the consequence of taking risks during the experiments and cross validation of a sample of decisions with production managers, as was done in this research (sections 7.3.1 & 7.3.2), can also contribute to a minimisation of the gaming effect.

Knowledge gaps in the data sets is a limitation which can constrain the implementation of KBI, especially if decision trees alone are used to identify and represent decision making strategies. Having found that knowledge gaps can be identified by comparing the decision rules with qualitative knowledge from the experts, it is concluded that qualitative validation and perhaps quantitative cross validation (if sufficient data is available) is the most efficient current approach for identifying and resolving those problems.

In the version of KBI that has been developed in this research, consistent decision making has been assumed. The decision models predict expected decisions $E(\mathbf{D}_j/\mathbf{A})$ and these decisions are used in the simulation. With this approach it has been assumed that the quality of the decisions of a specific decision maker is not affected by the working conditions, such as the temperature and the noise level. There are concerns that this might not always be the case since there is preliminary evidence (Baines & Kay 2002) which suggests that the quality of inspections from specific groups of decision makers might vary from time to time. A parallel research project (Baines & Kay 2002) is under way to investigate what effect the working conditions might have on the performance of individuals.

9.5 Conclusion

The previous sections of this chapter have discussed the lessons that have been learned from the implementation of KBI. From this discussion it is clear that KBI is an area of great potential. The fact that the conclusions have been based on the application of KBI on a specific case study reveals that the approach can be refined and that it would benefit from future research, which would provide additional empirical evidence on its applicability. The thesis concludes in the next chapter by summarising what has been achieved in terms of the research question and by outlining areas for future research.

Chapter 10

Conclusion

Having discussed the implementation and the refinements of the KBI methodology, this chapter concludes by providing a reflection of what has been achieved and what can be achieved by future research. Section one summarises the research which has been described in this thesis. Section two outlines what has been achieved measured against the objectives and the aim of the research. Section three summarises the key findings and section four outlines the contribution of the research. Section five explains general limitations and section six outlines areas for future research.

10.1 Summary of the research

Chapter 1 introduced the subject of the research by explaining the challenge involved in the use of simulation for understanding and improving decision making. Chapter 2 provided an overview of the relevant literature. The main research gaps were identified and the main techniques which can potentially be used to model decision making in simulation were described. In Chapter 3 the aim and the objectives of the research were set out and the strategy that was used to address the research questions was explained. In Chapter 4 a methodology for capturing efficient decision making using simulation was formulated. Applying this methodology (known as KBI), the subsequent chapters (5, 6, 7 and 8) tested and evaluated the use of simulation for understanding and improving the design of decision making. Chapter 5 evaluated the first stage of the KBI methodology. This involved the use of VIS for understanding and formulating the decision making problem. Chapter 6 evaluated the second stage of KBI. For this stage a data collection was required. In Chapter 7 the third stage of KBI was tested and the use of AI for modelling decision making was evaluated. Chapter 8 evaluated the remaining two stages of the KBI methodology, which involved the use of simulation to identify efficient decision making policies. Finally, Chapter 9 discussed the conclusions, the lessons learned and the refinements in the conceptual version of the KBI methodology outlined in Chapter 4.

Based on these conclusions the research has provided the required evidence to address the objectives and the research questions. These are discussed in the following section.

10.2 Aim, objectives and research questions - what has been achieved

This section restates and discusses the findings in relation to the objectives of the research. It refers to relevant sections of the discussion in Chapter 9 and it provides a reflection on what has been achieved in terms of the aim, the objectives and the research questions.

The first objective of the research was to form a methodology to capture efficient decision making. The research has achieved this objective by formulating the KBI methodology in Chapter 4.

The second research objective was to test the above methodology by addressing the research questions which were set out in section 3.1. The research has tested and evaluated the stages of the methodology in Chapters 5, 6, 7 and 8. The following paragraphs summarise what has been achieved in relation to the research questions.

In relation to research question 1, the research has shown in Chapters 5 and 6 that, subject to limitations in the size of the data set, it is possible to use simulated decision situations to elicit the knowledge and to collect the data required for modelling (with AI) decision makers in simulation. Using a case study, the research has tested thoroughly the use of VIS as a tool for problem understanding. This has been achieved by using various types of VIS for collecting decision making data (section 5.3).

For research question 2, the research has shown that there are practical difficulties and limitations involved in the process of collecting decision making data using VIS. As has been discussed in section 9.1.3, these include the time that is required, the attributes which cannot be represented in the simulation and the limited availability of the decision makers. Amongst the lessons that have been learned, it has been found that sometimes, in order to elicit the knowledge, it is necessary to present decision scenarios without running the simulation during the data collection. This is a direct implication for the OR analyst and it is reflected in the requirements of the system which has been described in section 6.2.

As has been discussed in Chapter 7 in relation to research question 3, the research has shown that it is possible to identify aspects of decision making strategies using AI and a sample of decision situations. The strategies that can be identified are not complete representations of the decision makers, but they can be used to identify differentiations in the decision making policies employed. These are appropriate for representing individuals in simulation.

In relation to research question 4, the research has shown in Chapter 8 that, under certain circumstances (discussed in section 9.3.2) and subject to limitations (discussed in section 9.4), it is possible to identify efficient decision strategies by representing and assessing the decision making in a simulation. In this research this has been achieved by providing a real world example where three strategies were assessed using simulation (section 8.3).

The third objective of the research was to refine the methodology and to discuss the lessons that have been learned as a result of its application in an industrial problem. The research has met this objective by discussing the general lessons learned as a result of the implementation of the KBI methodology. These were discussed in Chapter 9 and are reflected in the key findings which will be outlined in section 10.3.

Based on the above discussion, it is clear that the research has addressed the research questions and it has met its objectives and its aim. The research has developed and tested the use of simulation for understanding and improving decision making in a real context. This has been achieved by identifying difficulties and limitations and by using a real world example to test the KBI methodology. Based on this example a number of lessons have been learned and have been discussed in Chapter 9. Amongst the lessons there are key findings which highlight the degree of success of the approach. These are summarised in the following section.

10.3 Key findings

In this research it has been found that it is possible to use simulation to improve human decision making. The approach requires the representation of decision making in simulation and this involves a process that is reflected in the stages of the KBI methodology. This has been tested on an industrial problem and the findings have been discussed in detail in Chapter 9.

In order to use simulation to collect decision making data, it is necessary to represent the attributes which the decision makers take into account during the decision making process. This may not always be possible and, in order to resolve the problem, the inclusion of meta-attributes must be considered (section 5.3.1).

The involvement of the real experts limits significantly the size of the data set that can be collected using simulated decision situations. When real experts are involved, in order to collect a representative sample of decisions, a stratified approach with the use of a decision situations database (section 6.2) should be considered.

The use of decision trees enables the identification of differences in the decision making policies employed by different decision makers. However, the limited availability of data may not enable a complete representation of each decision making strategy.

The representation of decision making may have an impact on the prediction of the simulation (section 9.3.2). In this research this impact was statistically significant and it enabled efficient decision making practices to be identified.

10.4 Contribution to knowledge

Having addressed the research questions, the research has addressed a number of research issues outlined in section 2.6. The research has been applied to an industrial problem. This addresses the need for testing the use of simulation for understanding and improving decision making in a real context (Flitman 1986). It shows that the interface between simulation and AI, which has been proposed in the literature over the last 20 years, can be applied in industrial problems and can be used to improve the design of decision making.

By forming and testing the KBI approach, the research has also addressed the need for a model building methodology to represent decision making in simulation (Curram 1997). The involvement of real decision makers has tested the use of VIS as a data collection tool. It has enabled the identification of practical limitations and it has addressed the need for testing the use of VIS with the involvement of real experts (Liang et al 1992).

The functional specification of the data collection system described and discussed in sections 6.2.2 and 9.1.3 has contributed to the identification of the requirements of an intelligent simulation package (Williams 1996). This highlights a promising area for software development and indicates some of the areas for future research. These are described in section 10.6.

10.5 Limitations

The limitations of specific findings and the limitations of the KBI approach have been discussed extensively in Chapter 9. This section summarises the key limitations of the research in general.

As anticipated in section 3.2.3, the main risk in the validly of the findings is the use of a single case study. Clearly, the validity of the findings is based on the assumption that the case study which has been used in this research encapsulates the key challenges involved in the use of simulation for improving decision making. Based on this assumption, the research has recommended the feasibility of the use of simulation to improve decision making in a real context.

The limited availability of data is also a general limitation of the research. The strategies that have been identified are based on a very limited data set. As has been pointed out in section 7.4, this means that the decision making models may not form a complete representation of the decision makers. Furthermore, it must be remembered that the data set that was used in this research was collected from a series of simulation experiments. The research is based on the assumption that the decision makers who were involved in the experiments were not influenced by the fact that the decisions that they made did not have a real impact on the real system.

10.6 Future research

A VIS-based data collection presents a number of challenges. The use of three dimensional displays is an approach which could improve the knowledge elicitation process since it may create a more realistic environment in which to present decision scenarios. Further research is required to investigate the potential benefits to be gained from the use of 3D displays.

The application of rule induction as part of KBI is also an area that presents challenges and would benefit from future research. The limited interpolation in the decision trees indicates that there are unanswered research questions. For example, what is the best approach in order to avoid knowledge gaps in decision trees? Also, can the use of neural networks or logistic regression improve the process of filling the knowledge gaps by extending the sample data sets?

Due to the limited data that could be collected, a number of difficulties were identified in the process of validating the models that were produced. These difficulties show that an area of potential research interest is the use of bootstrapping (UrbanHjorth 1994) to support the model validation process in the KBI methodology. The capability of the bootstrapping approach to generate multiple samples from a data set is promising. It may be able to support the process of generating decision situations and the process of validating the models by generating samples for cross validation.

As explained in Chapter 7, the research has focussed on binary decision making problems in which only one dependent variable is used in the decision making models which have been calibrated. The investigation of the feasibility of the modelling techniques to identify strategies in decision making models with multiple non-binary decision variables is a challenging area of significant potential interest in which future research must also be undertaken.

From the discussion of the conclusions from Chapter 8, it is clear that the version of KBI which has been implemented in this research has focused on recommending improvements that can be achieved by identifying and sharing good decision making practices. The fifth stage of KBI has not been fully tested and an area of potential interest that can extend the frontiers of KBI is the investigation of the feasibility of a process for combining and assessing good decision making practices.

References

Abdurahimman V. & Paul R.J. (1994), Machine learning and simulation model specification. Journal of simulation practice and theory 2, 1, 1-15.

Alifantis T., Robinson S., Edwards J., Hurrion R.D., Ladbrook J. & Waller T. (2001), Knowledge Based Improvement: Simulation and AI for Improving Unplanned Maintenance Operations. 2nd European Conference - Operations. Intelligent Management Systems in Operations (2001).

Alifantis T. & Robinson S. (2001), Using simulation and neural networks to develop a scheduling advisor. Proceedings of the 2001 Winter Simulation Conference.

Angelides M.C. & Paul R.J. (1999), A methodology for specific, total, enterprise, role-playing, intelligent gaming-simulation environment development. Decision support systems 25, 89-108.

Attar software (2000), Xpertrule. KBS reference manual. Attar software: Leigh, UK.

Au G. & Paul R.J. (1996), Visual Interactive Modelling: A pictorial simulation specification system. European Journal of Operational Research 91, 14-26.

Baines T.S. & Kay J.M. (2002), Human performance modelling as an aid in the process of manufacturing systems design: A pilot study. International Journal of Production Research 40,10, 2321-2334.

Baines T.S., Kay J.M., Oyarbide A. & Ladbrook J. (2000), Modelling the human in manufacture. First world conference on production and Operations Management, Sevilla.

Bell P.C. & O'Keefe R.M. (1994), Visual interactive simulation: A methodological perspective. Annals of Operations Research,53, 321-342.

Bell P.C. & O'Keefe R.M. (1995), An experimental investigation into efficacy of Visual Interactive Simulation. Management Science 41,6 1018-1037.

Bryman A. & Bell E. (2003), Business research methods. Oxford: Oxford U.P.

Carvalho M.D. & Dougherty M.S., Fowkes A.S. (1998), Forecasting travel demand: A comparison of logit and artificial neural network methods. Journal of the Operational Research Society 49, 717-722.

Checkland P.B. (1981), Systems Thinking Systems Practice. Wiley Chichester, UK.

Clark D.N. (1999), Strategic level MS/OR tool usage in the United Kingdom: a comparative survey Asia-pacific. Journal of the Operational Research Society 16(1), 35-51.

Cornnell N.A.D., Klein J.H. & Powel P.L. (2003), It is tacit knowledge but not as we know it: Redirecting the search for knowledge. Journal of the Operational Research Society 54, 2, 140-152.

Curram S. (1997), Representing intelligent decision making in discrete event simulation: a stochastic neural network approach. PhD thesis. University of Warwick, England.

Davies R.M. & O'Keefe R.M. (1988), Simulation modelling with Pascal. Prentice-Hall, Upper Saddle River, NJ.

Deslandres V. & Pierreval H. (1997), Knowledge acquisition issues in the design of decision support systems for quality control. European Journal of Operational Research 103, 296-311.

Doukidis G.I (1987), An anthology on the homology of simulation with Artificial Intelligence. Journal of the Operational Research Society 38, 701-712.

Doukidis G.I. & Paul R.J. (1985), Research into expert systems to aid simulation model formulation. Journal of the Operational Research Society 36, 319-325.

Doukidis G.I. & Paul R.J. (1990) A survey of the application of Artificial intelligence techniques within the OR society. Journal of the Operational Research Society 41, 5, 363-375.

Dyson R.G., Allen R., Camanho A.S., Podinovski V.V., Sarrico C.S. & Shale E.A. (2001), Pitfalls and protocols in DEA. European Journal of Operational Research 132, 2, 245-259.

Eden C., Jones S. & Sims D. (1979), Thinking in Organisations. MacMillan Press: London.

Edwards J.S. (1991), Building knowledge-based systems, towards a methodology. Pitman, London.

Edwards JS, Alifantis T, Hurrion RD, Ladbrook J, Robinson S &Waller T, (2004) Using a Simulation Model for Knowledge Elicitation and Knowledge Management. Simulation Modelling Practice and Theory 12 527-540.

Eglese R.W. (1990), Simulated annealing: A tool for operational research. European Journal of Operational Research, 46.

Fiddy E., Bright J. G. & Hurrion R. D. (1981), See why: Interactive simulation on the screen. Proceedings of the Institute of Mechanical Engineers, London.

Fildes R. & Ranyard J.C. (1997), Success and survival of operational research groups – a review. Journal of the Operational Research Society 48, 4, 336-360.

Fishman G.S. (1978), Grouping observations in digital simulation. Management Science 24,5,510-521.

Flitman A.M. (1986), Towards the application of artificial intelligence techniques for discrete event simulation. PhD thesis. University of Warwick, England.

Flitman A.M. & Hurrion R.D. (1987), Linking discrete event simulation models with expert systems. Journal of the Operational Research Society 38,8,723-733.

Haykin S (1999), Neural Networks: A comprehensive foundation. Prentice Hall International, Upper Saddle River, N.J.

Hindle T., Checkland P., Mumford M. Worthington D.(1995) Developing a methodology for multidisciplinary action research Journal of the Operational Research Society 46,453-464.

Hurrion R.D. & Secker R.J.R. (1978), Visual interactive simulation. An aid to decision making. Omega. The International Journal of Management Science 6, 4, 5419-426.

Hurrion R.D. (1976), The design use and required facilities of an interactive visual computer simulation language to explore production planning problems. PhD Thesis, University of London.

Hurrion R.D. (1978), An investigation of visual interactive simulation methods using the job-shop scheduling problem. Journal of the Operational Research Society 29, 11, 1085-1093.

Hurrion R.D. (1980), An interactive visual system for industrial management. European Journal of Operational Research, 5, 86-93.

Hurrion R.D. (1985), Implementation of a visual interactive consensus decision support system. European Journal of Operational Research 20, 138-144.

Hurrion R.D. (1991), Intelligent Visual Interactive Simulation. European Journal of Operational Research 54, 349-356.

Hurrion R.D. (1986), Visual interactive modelling. European Journal of Operational Research 23, 281-287.

Jackson P. (1998), Introduction to expert systems. Harlow: Addison-Wesley.

Jeffrey P. & Seaton R. (1995), The use of operational research tools: A survey of operational research practitioners in the UK. Journal of the Operational Research Society 46(7), 797-808.

Johnston J. & DiNardo J. (1997), Econometric Methods. MCGraw-Hill, London.

Kolodner J. (1993), Case-Based reasoning. Morgan Kaufman, San Mateo: CA.

Ladbrook J. (1998), Modelling breakdowns: An inquest. MPhil Thesis. University of Birmingham

Lanner Group (2003), Witness 2003. Lanner Group: Redditch, UK.

Law A.M. & Kelton W.D. (2000), Simulation modelling and Analysis 3rd edn. McGraw-Hill, London.

Liang T.P., Moskowitz H. & Yuehwern Y. (1992), Integrating neural networks and semi-Markov processes for automated knowledge acquisition: An application to real scheduling. Decision Sciences, 23, 1297-1311.

Liao S. (2000), Case based decision support system: Architecture for simulating military command and control. European Journal of Operational Research 123, 558-567.

Lightfoot M.J. (1999), Expert knowledge acquisition and the unwilling expert: A knowledge engineering perspective. Expert Systems, 16, 3.

Malhorta M.K., Sharma S. & Nair S.S. (1999), Decision making using multiple models. European Journal of Operational Research, 114, 1, 1-14.

Mason C.R. & Moffat J. (2000) Representing the C2 Process in Simulations: Modelling the Human Decision-Maker. Proceedings of the 2000 Winter Simulation Conference.

MathWorks (1999), Matlab 6. Reference manual. Natick, MA. US.

McClelland J.L. & Rumelhart D. E. (1988), Explorations in parallel distributed processing: A handbook of models programs and exercises. MIT Press, Cambridge: MA.

Mingers J. (1986), Expert Systems - Experiments with rule induction. Journal of the Operational Research Society 37, 11, 1031-1086.

Mingers J. (1987), Rule induction with statistical data - A comparison with multiple regression. Journal of the Operational Research Society 38, 4, 347-351.

Moffat J, Campbell I, Glover P. (2004), Validation of the mission-based approach to representing command and control in simulation models of conflict. Journal of the Operational Research Society 55, 340-349.

Moffat J. & Witty S. (2002), Bayesian decision making and military command and control. Journal of the Operational Research Society 53, 709-718.

Moffat J. (2000), Representing the command and control process in simulation models of conflict. Journal of the Operational Research Society 51,431-439.

Mundermann F.(1993), Fuzzy concepts for predicting the behaviour of other drivers on a highway. In E.D. Klement & w. Slanny (Eds) Fuzzy logic in artificial intelligence conference, Linz, Austria, June 1993. Stringer-Verlag, London.

O'Keefe R. (1986), Simulation and expert systems: A taxonomy and some approaches. Simulation 46, 1, 10-16.

O'Keefe R.M & Roach J.W. (1987), Artificial intelligence approaches to simulation. Journal of the Operational Research Society 38, 8, 713-722.

O'Keefe R. & Pitt I.L. (1991), Interaction with a visual interactive simulation and the effect of cognitive style. European Journal of Operational Research 54, 339-348.

Paul R.J., Macredie R.D. & Baldwin L.P. (1997), Special issue: Interdisciplinary Approaches to Neural Computing. Neural Computing & Applications 6, 4, 187-192.

Paul R.J. & Doukidis G.I. (1990), Further developments in the use of artificial intelligence techniques which formulate simulation problems. Journal of the Operational Research Society 37, 8, 587-810.

Perry W.L. & Moffat J. (1997), Developing models of decision making. Journal of the Operational Research Society 48 357- 470.

Pidd M. (1996), Object orientation discrete simulation and the three phase approach. Journal of the Operational Research Society 46, 3, 362-374

Pidd M. & Lewis P. (2001), Innovative research in OR/MS. European Journal of Operational Research 128, 1-13.

Pidd M. (1998), Computer Simulation in Management Science, 4th edition Wiley: Chichester, UK.

Pierreval H. & Ralambondrainy H. (1990), A simulation and learning technique for generating knowledge about manufacturing systems behaviour. Journal of the Operational Research Society 41, 6, 461-474.

Proudlove NC, Vadera S. & Kobbacy K.A.H. (1998), Intelligent management systems in operations: A review. Journal of the Operational Research Society 49, 682-699.

Quinlan J. R. (1979), Discovering rules by induction from large collections of samples. Expert system in microelectronic age. Edinburgh University Press, Edinburgh, UK, 168-201.

Quinlan J. R. (1983), Learning efficient classification features and their application to chess end games. In R.S Michalski, J.G. Carbonell & T.M. Michell (Eds) Machine learning: An artificial intelligence approach, 463-482 Tsioga Publishing.

Reeves C.R. (1995), Modern heuristic techniques for combinatorial problems. McGraw-Hill, London.

Robinson S, Alifantis T., Hurrion R.D., Edwards J., Ladbrook J. & Waller T. (2003), Modelling and Improving Maintenance Decisions: Having Foresight with Simulation and Artificial Intelligence. Transactions Journal of Materials & Manufacturing, Warrendale, PA, 256-264 (2003).

Robinson S, Alifantis T, Hurrion R. D., Edwards J., T, Ladbrook J. & Waller T.(2002) Modelling and Improving Maintenance Decisions: Having Foresight with Simulation and Artificial Intelligence. Foresight Vehicle Technology: Consumer Driven Design, Manufacturing, Supply Chain, and Purchasing, International, SP-1694, Warrendale, PA, 201-208 (2002).

Robinson S, Alifantis T., Edwards J.S., Ladbrook J. & Waller T., (2005), Knowledge Based Improvement: Simulation and Artificial Intelligence for Identifying and Improving Human-Decision Making in Operations Systems. Journal of the Operational Research Society 56, 7.

Robinson S. (2001), Soft with a Hard Centre: Discrete-Event Simulation in Facilitation. Journal of the Operational Research Society 52 905-915.

Robinson S. (1999), Simulation verification, validation and confidence: A tutorial. Transactions of the Society for Computer Simulation International. 16, 63-69

Robinson S. (2004), Simulation: The Practice of Model Development and Use. Chichester.

Robinson S. (2005), Discrete-Event Simulation: From the Pioneers to the Present, What Next? Journal of the Operational Research Society 56.

Robinson S., Edwards J.S. & Yongfa W. (1998), An Expert Systems Approach to Simulating the Human Decision Maker. Proceedings of the 1998 Winter Simulation Conference.

Robinson S., Edwards J.S. & Yongfa W. (2003b), Linking the Witness simulation software to an expert system to represent a decision making process. Journal of Computing and Information Technology - CIT, 11, 2, 123-133.

Robinson S. & Stanger M. (1998), Developing a simulation strategy for British Airways. Proceedings of the 1998 Winter Simulation Conference.

Schimidt B. (2000), The modelling of human behaviour. SCS Europe publishing house, Ghent.

Standridge C.R. & Steward D. (2000), Using expert systems for simulation modelling of patient scheduling. Simulation, 75, 5, 148-156.

Tocher K.D. (1963), The art of simulation. The English Universities Press. London.

UrbanHjorth J.S. (1994), Computer intensive statistical methods. Chapman & Hall, London.

Verbos P.J. (1994), The roots of Backpropagation. From ordered derivatives to neural networks and political forecasting. Wiley, New York.

Welbank M. (1983), A review of knowledge acquisition techniques for expert systems. British Telecom research laboratories, Ipswich.

Wielinga B.J. & Breuker J.A. (1988), Interpretation of verbal data for knowledge acquisition. Principles of Expert Systems, IEEE press 214-219, New York.

Williams T. M. (1996), Simulating the man-in-the-loop. OR Insight 9, 4, 17-21.

Williams T. M., Gittins R. P., Burke D. M., (1989), Replenishment at Sea. Journal of the Operational Research Society 40, 10, 881-887.

Wong B.K., Lai V.S. & Lam J. (2000), A bibliography of neural network business applications research 1994-1998, Computers & operations research 27,1025-1076.

Wright M.B. (1996), School timetabling using heuristic search, Journal of the Operational Research Society 47(3), 347-357.

Yin R.K. (2003), Case study research: Design and methods Thousand Oaks CA: Stage publications.

Zadeh L.A. (1973), Outline of a new approach to the analysis of complex systems and decision processes I.E.E.E. Trans on Systems Man and cybernetics, 3,28-44.

Appendix 1: Decision Trees

Decision trees with automatic rule induction

A1.1 Initial calibration: t_{DM1}^{A}: *Decision tree generated with automatic ID3 that represents strategy DM1*

A1.2 Initial calibration: t_{DM2}^{A}: *Decision tree generated with automatic ID3 that represents strategy DM2*

A1.3 Initial calibration: t_{DM3}^{A}: *Decision tree generated with automatic ID3 that represents strategy DM3*

Decision trees with semi automatic rule induction

A1.4: Final calibration: t_{DM1}^{s}: *Decision tree generated with semi-automatic ID3 that represents strategy DM1*

A1.5: Final calibration: t_{DM2}^{s}: *Decision tree generated with semi-automatic ID3 that represents strategy DM2*

A1.6: Final calibration: t_{DM3}^{s}: *Decision tree generated with semi-automatic ID3 that represents strategy DM3*

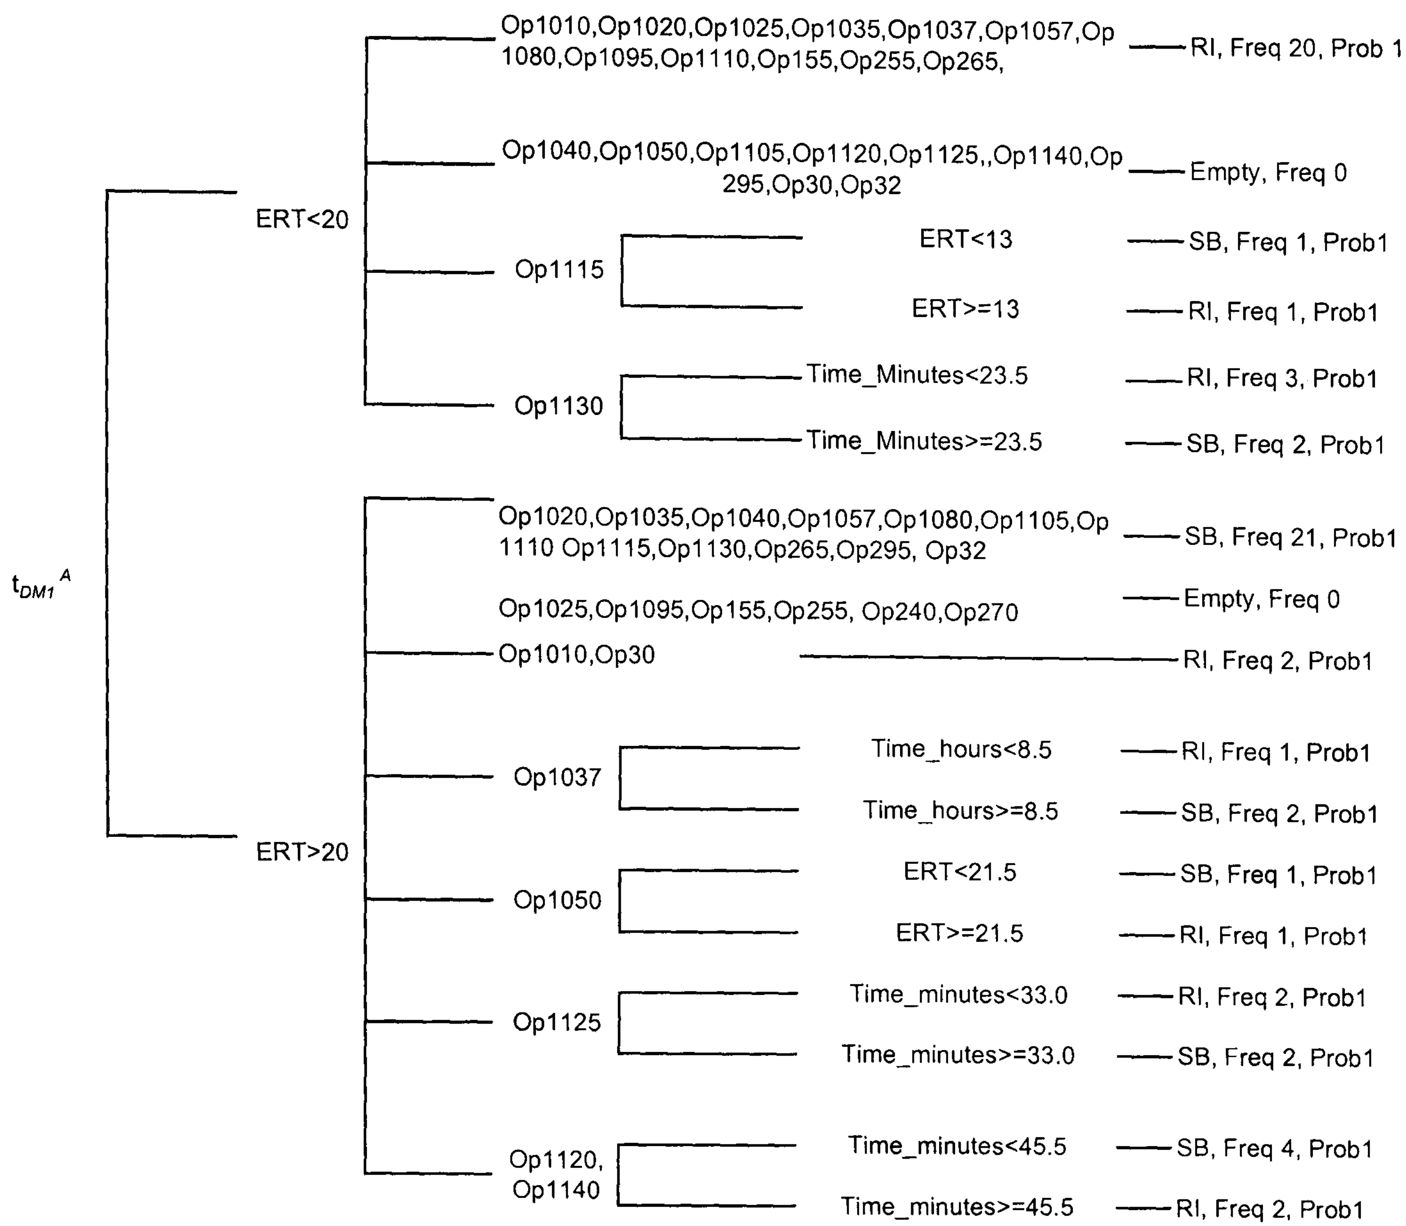

A1.1 $t_{DM1}{}^{A}$: *Decision tree generated with automatic ID3 that represents strategy DM1.*

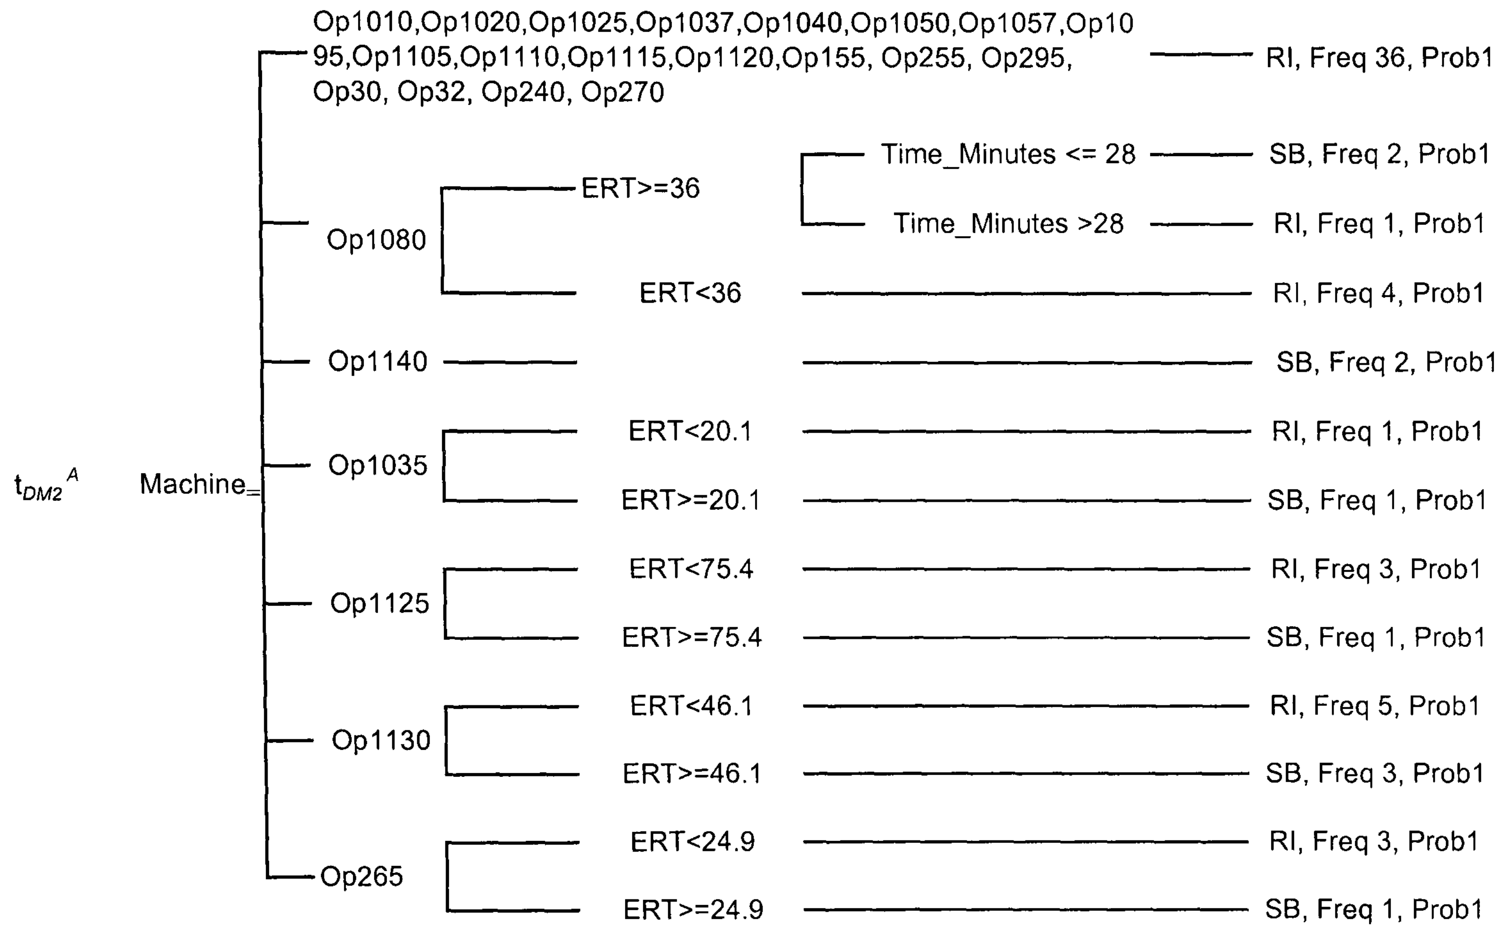

A1.2 t_{DM2}^{A}: *Decision tree generated with automatic ID3 that represents strategy DM2.*

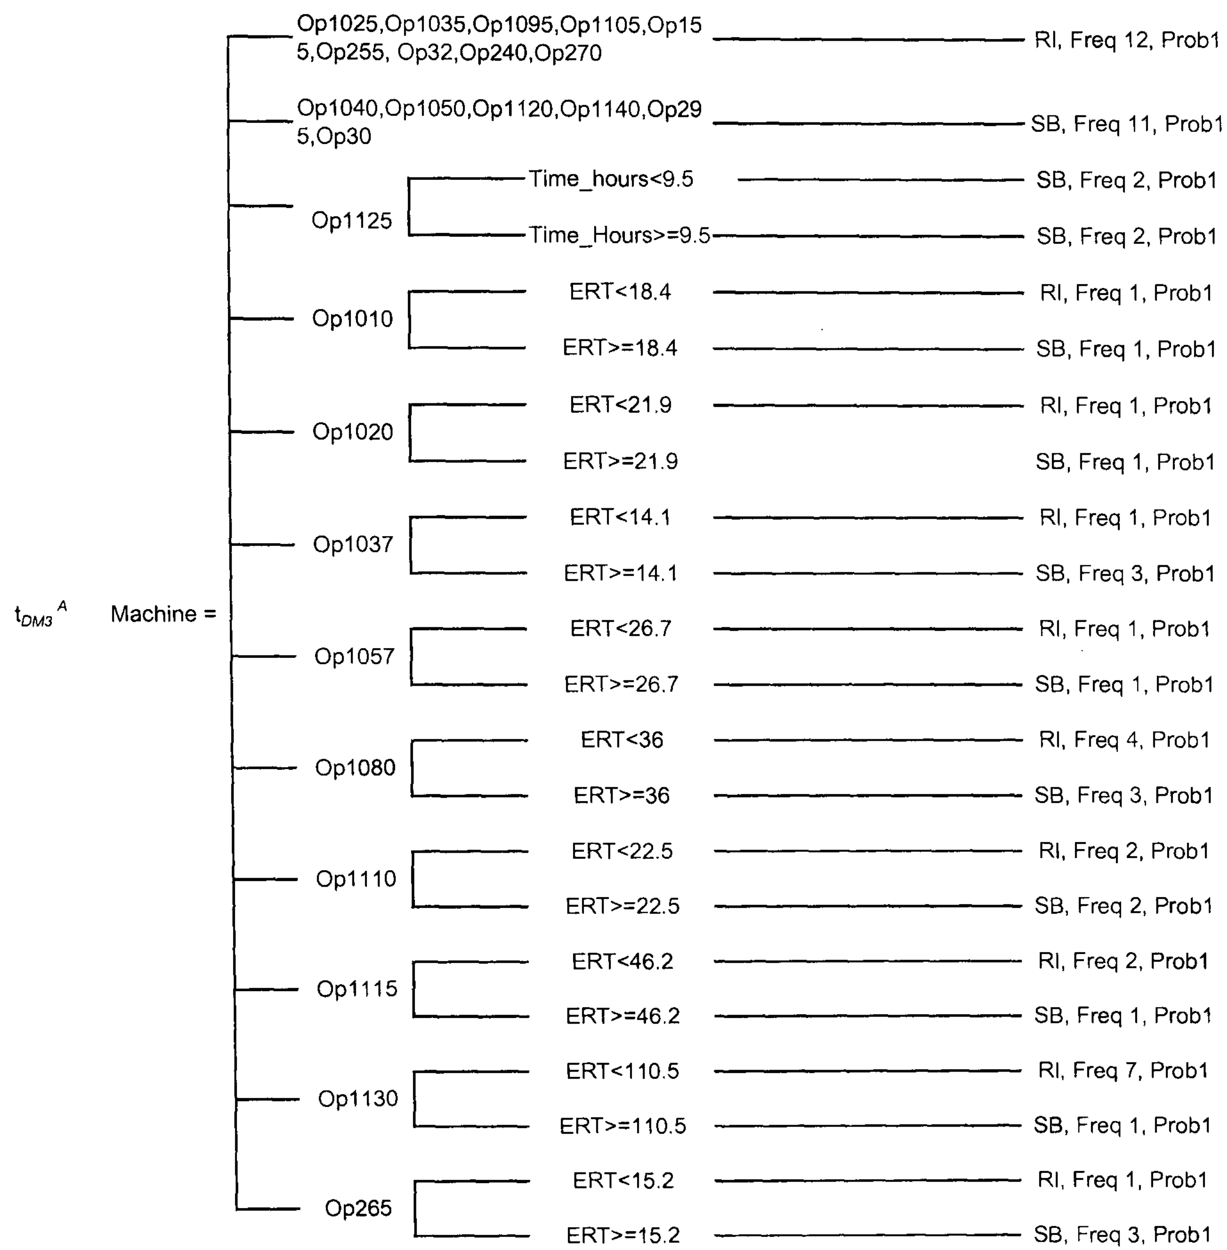

A1.3 $t_{DM3}{}^{A}$: *Decision tree generated with automatic ID3 that represents strategy DM3.*

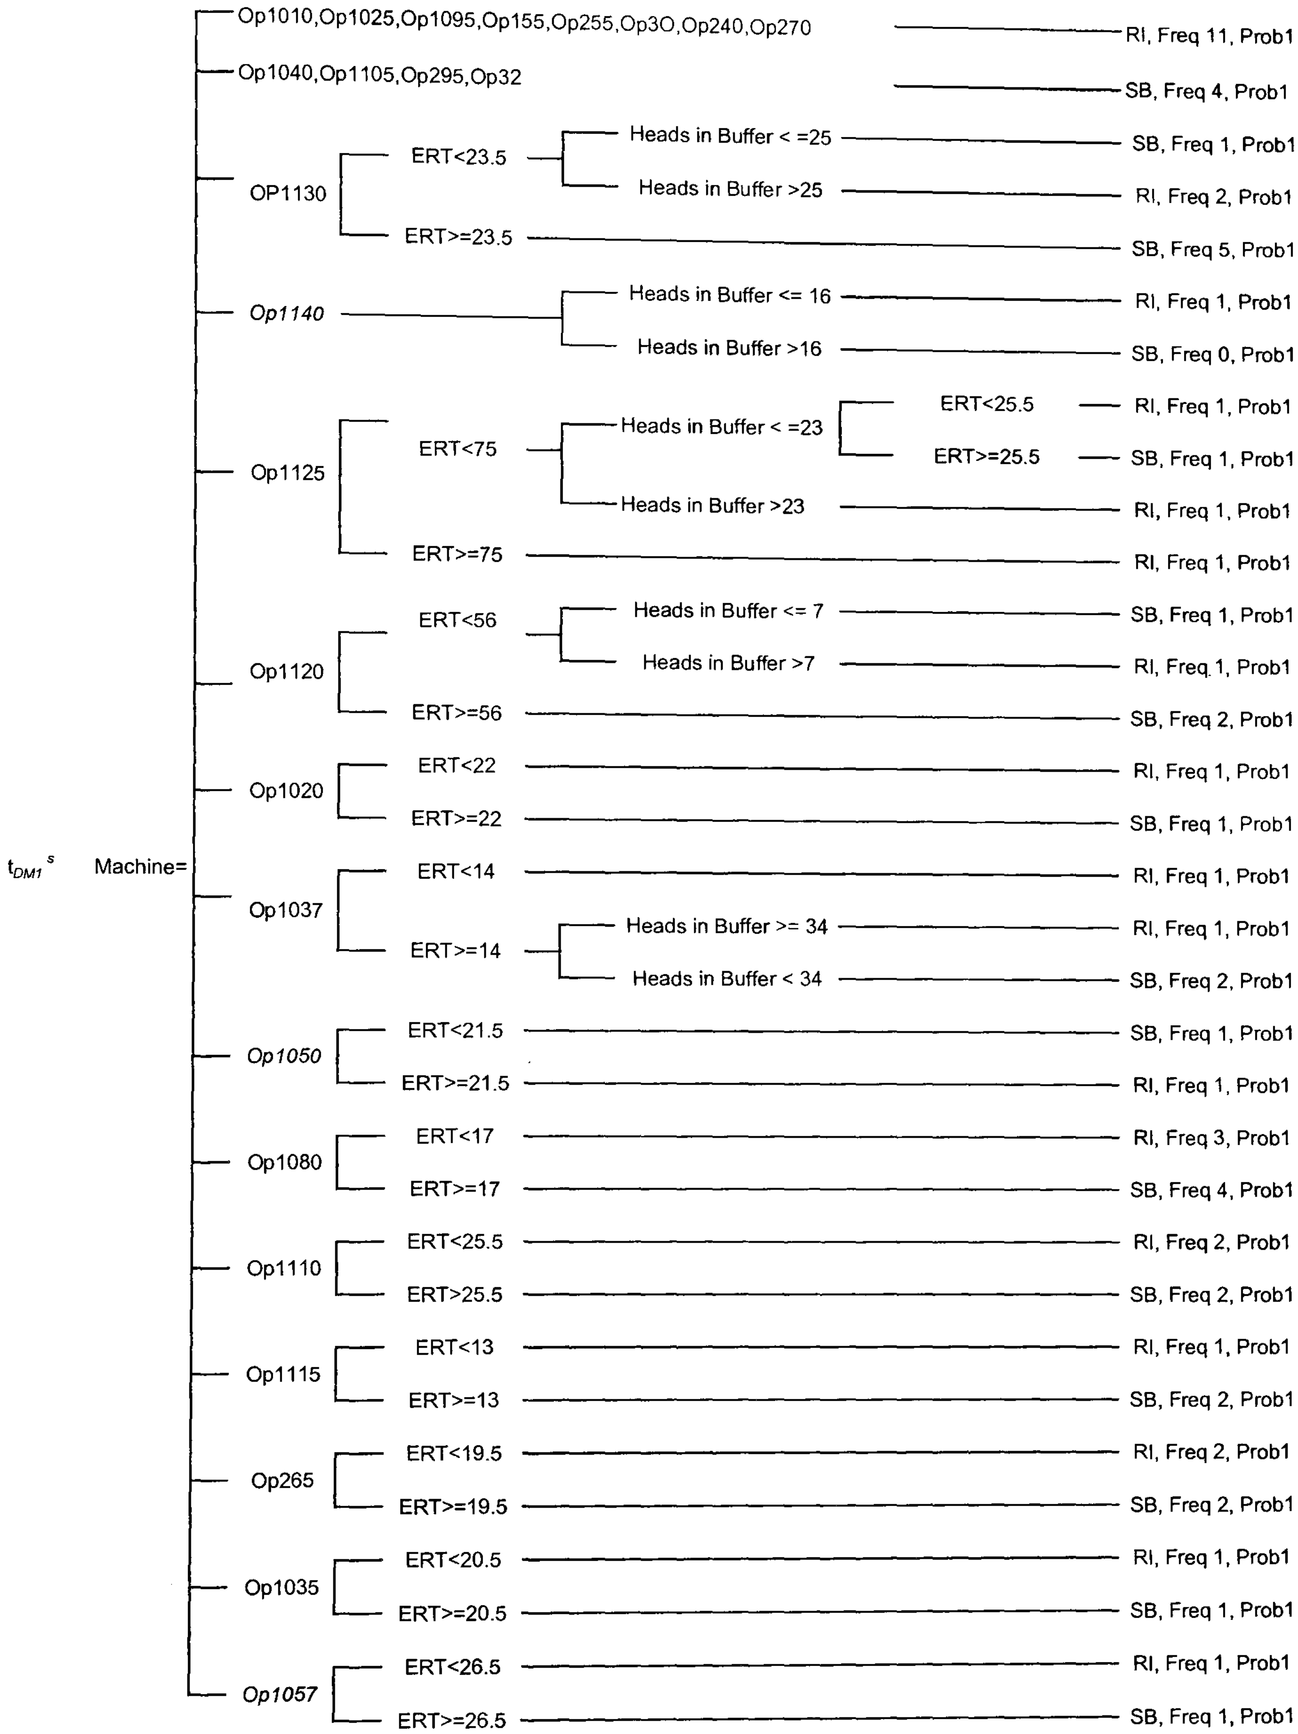

A1.4 *t_{DM1}^{s}: Decision tree generated with semi-automatic ID3 that represents strategy DM1.*

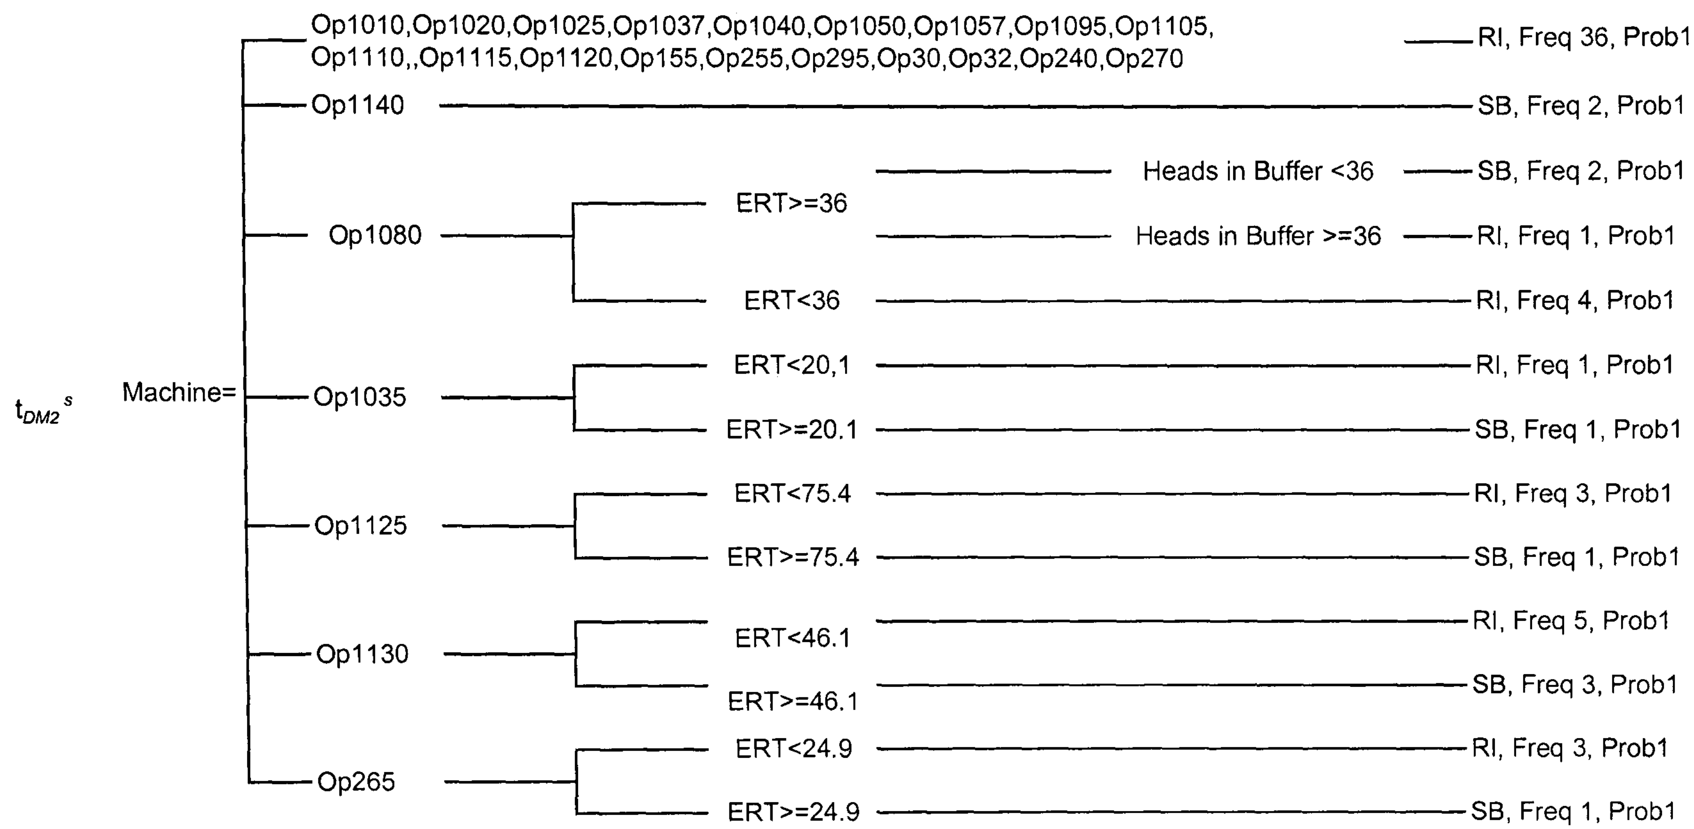

A1.5 $t_{DM2}{}^{s}$: *Decision tree generated with semi-automatic ID3 that represents strategy DM2.*

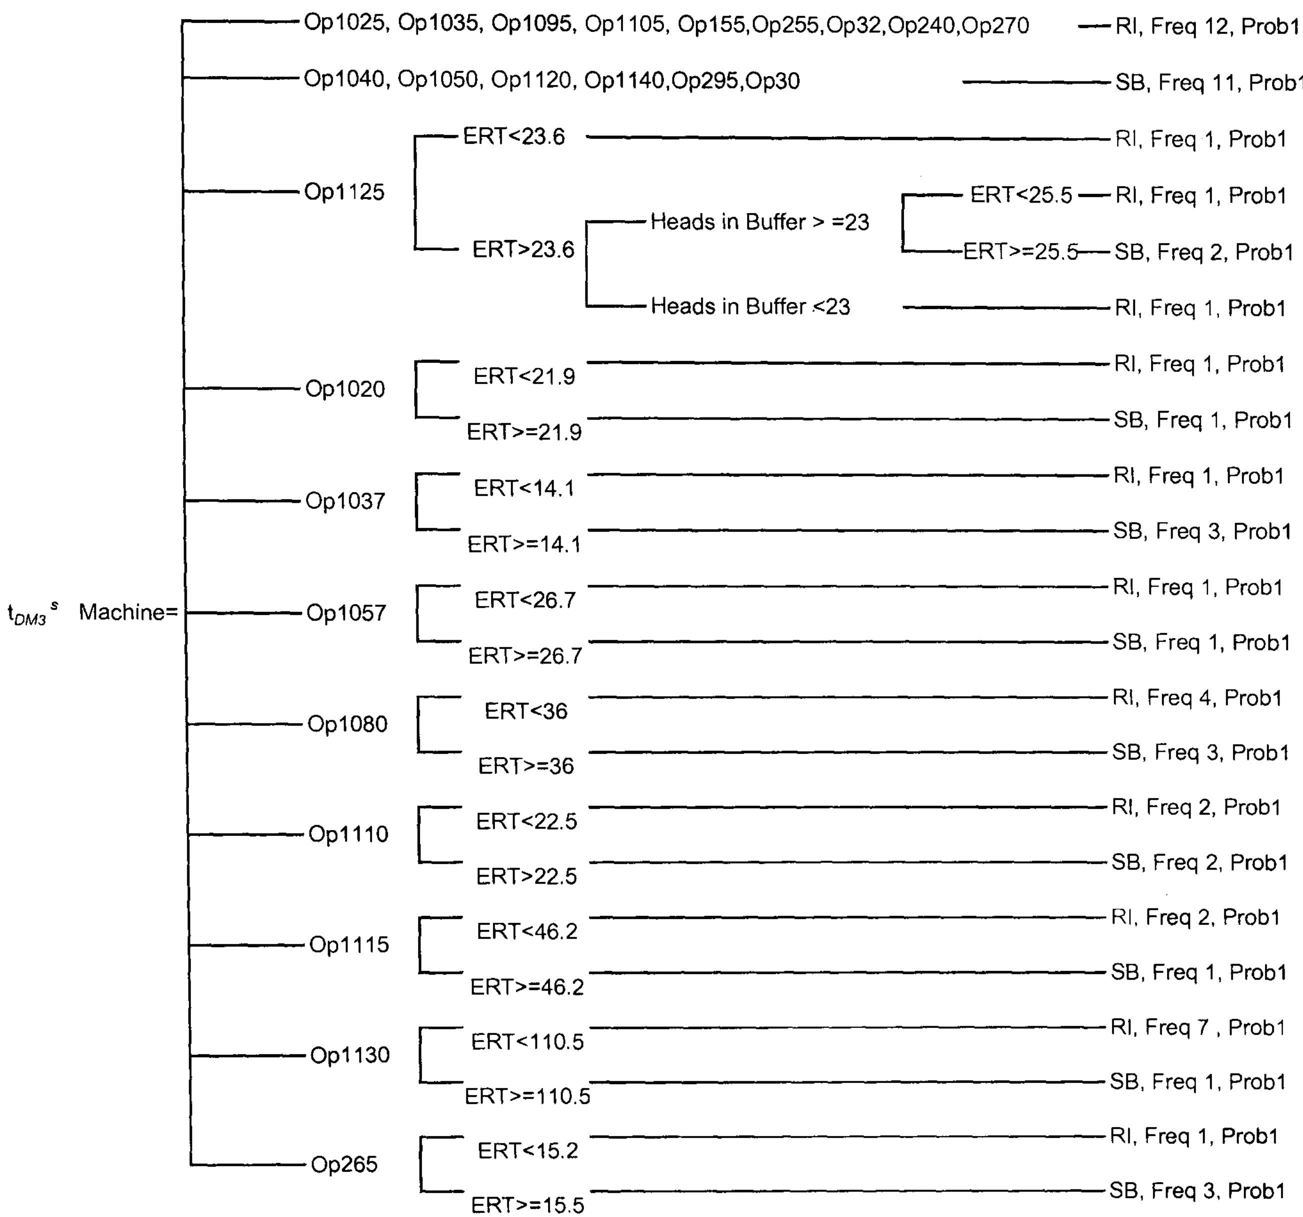

A1.6 t_{DM3}^{s}: *Decision tree generated with semi-automatic ID3 that represents strategy DM3.*

Appendix 2: Visual Basic interface for linking Witness with XpertRule

```
'Collect attribute vallues from the simulation model
    Dim witobj
    Dim rt
    Dim ma
    Dim hb
    Dim flag As Integer
    Dim decision_counter As Integer
    Set witobj = GetObject(, "witness.wcl")
    DoEvents
    rt = witobj.variable("dmaking.ma_rt")
    ma = witobj.variable("dmaking.madown")
    hb = witobj.variable("dmaking.hb")
    decision_counter = witobj.variable("dmaking.decision_counter")
```

Figure A2.1: ***Communication interface to collect attribute information***

```
'invoke expertrule pass the attribute vallues and request a decision
    Dim XpertruleObject As Object
    Dim i
    Dim a As String
    Set XpertruleObject = CreateObject("xrclient.xr32run")
    Call XpertruleObject.StartApp("d:\warwick phd\sigma line\dialog\MODELS\rm.xra -server
-hidden")
    DoEvents
    Call XpertruleObject.Poke("RT", rt)
    Call XpertruleObject.Poke("MA", ma)
    Call XpertruleObject.Poke("HEADINBUFFER", hb)
    Call XpertruleObject.Command("CONTINUE")



'Wait untill xpertule make a decisions
    Do
      Call XpertruleObject.Peek("?", a$)
    Loop Until a = "READY"
```

Figure A2.2: ***Communication interface to invoke and control XpertRule***

```
'Cpllect the decision from xpertrule and close it
    Call XpertruleObject.Peek("PETERHARIS", a$)
    Call XpertruleObject.Command("EXIT")

'Pass the decision to witness
witobj.variable("Dmaking.xr_outcome", 1) = a
```

Figure A2.3: ***Communication interface for passing the decision to Witness***

```
witobj.batch
witobj.endole
End
```

Figure A2.4: ***Communication interface for passing control to the simulation***

Appendix 3: Attributes - Decisions data sets

A3.1 Decision maker DM1

A3.2 Decision maker DM2

A3.3 Decision maker DM3

A7. 1: *Attributes decisions data set - Decision maker DM1*

A_I D_{DM1}

Record	1.Type of fault	2. Repair Time	3. Machine number	4. Time		5. Engines Produced so far this shift	6. Parts waiting in the conveyor before the machine	7. Number of heads in the buffer	8. The machine has broken down this day	9. The machine has broken down this month	10. The breakdown has happened this day	11. The breakdown has happened this month	12 Action	13. Switch off the machine	14. Who	15. Ask the production Manager	16. Plan Repair	17. Plan repair – when?
1	CENTERING ADVANCED(SQI150)_	26.3	Op32	9	44	184	8	57	1	1	1	1	SB	No	IMS1	No	Yes	End off shift
2	GRIPPER 2 CLAMPED(FEEDBACK)_	22.6	Op1035	10	30	200	1	66	1	1	1	1	SB	No	Operator	No	Yes	End off shift
3	VALVE COLLET 1.2 LOST._	171.8	Op1080	10	30	200	0	66	1	1	1	1	SB	No	IMS2	No	Yes	
4	GUARD DOOR OPENED_	70.7	Op1105	10	30	200	0	66	1	1	1	1	SB	No	IMS1	No	Yes	End off shift
5	OVERTRAVEL PROTECTION TRIPPED_	97.8	Op1130	13	10	595	0	0	23	23	1	1	SB	No	IMS1	No	Yes	
6	INDEX(SQI151_	15	Op255	10	11	306	3	0	2	6	1	1	RI	Yes	Second mechelec	Yes	No	
7	INDEX LOWERED;EXIT 1 EMPTY, SQI141 NSQI124 NSPI170, SB027/A_	12.2	Op1095	11	32	353	0	17	1	4	1	1	RI	No	Group Leader	No	No	
8	PART 2 IN POSITION (LOADING PLATE)_	23.5	Op1125	11	59	412	0	0	16	27	8	15	SB	No	IMS2	No	No	
9	FEEDING EMPTY 1_	12	Op1080	11	13	436	2	0	6	8	1	1	RI	No	Group Leader	No	No	
10	OP1025>FAULT(DATA TRANSFER,ASM-HOST)_	18.5	Op1025	11	18	436	5	0	5	8	3	4	RI	No	Group Leader	No	No	
11	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	77.8	Op1115	12	8	525	1	2	1	2	1	1	SB	No	Operator	No	No	
12	LIMIT SWITCH MONITORING BLOCKER 2_	219.2	Op1080	14	11	583	0	0	8	10	1	1	SB	No	IMS1	No	Yes	

13	PART 1 IN POSITION (LOADING PLATE)_	32.4	Op1130	15	16	667	1	4	22	49	12	25	SB	No	IMS1	No	Yes	
14	PART 2 IN POSITION (LOADING PLATE)_	59.7	Op1130	8	56	70	2	5	5	54	2	18	SB	No	IMS1	No	Yes	End off shift
15	PASSWORD INPUT(NOK)_	134.2	Op1120	10	33	136	0	19	1	7	1	3	SB	No	IMS2	No	Yes	
16	PART 2 IN POSITION (LOADING PLATE)_	27.7	Op1130	11	34	270	0	10	10	59	3	19	SB	No	IMS2	No	Yes	End off shift
17	DISTRIBUTOR IN POSITION(FEEDBACK)_	22.7	Op1037	12	6	451	1	7	8	36	5	29	SB	No	Operator	No	Yes	End off shift
18	PART 1 IN POSITION (LOADING PLATE)_	18.5	Op1130	12	43	472	1	51	23	103	13	50	RI	No	Group Leader	No	No	
19	READ IDS(FINISHED MESSAGE), FEEDBACK(F9.1), SB022/F_	33.2	Op1110	8	57	69	1	12	2	7	1	2	SB	No	IMS1	No	Yes	End off shift
20	PART 1 IN POSITION (LOADING PLATE)_	23.8	Op1125	7	16	914	0	47	2	54	1	21	RI	No	Group Leader	No	No	
21	LIMIT SWITCH MONITORING BLOCKER 2_	50.3	Op1080	8	27	132	1	6	1	22	1	5	SB	No	Operator	No	Yes	End off shift
22	PART IN POSITION(FEEDBACK)_	28.9	Op1037	9	2	212	1	0	3	41	1	2	SB	No	IMS1	No	Yes	End off shift
23	START DISPENSER(NOT ENABLED)_	40.6	Op1020	15	29	9	5	0	1	6	1	3	SB	No	IMS2	No	Yes	End off shift
24	TAGGING SYSTEM(FAULT,CONNECTIONS)_	44.6	Op1120	18	58	303	2	0	2	18	1	1	RI	Yes	Second mechelec	Yes	Yes	End off shift
25	GRIPPER CLAMPED(FEEDBACK)_	89.3	Op1140	20	55	556	0	0	2	16	1	10	RI	No	Second mechelec	No	No	
26	GUARD DOOR OPENED_	27.4	Op1125	13	20	665	0	0	20	82	1	1	RI	No	Group Leader	No	No	
27	GRIPPER 1 CLAMPED(FEEDBACK)_	17.5	Op1035	15	0	883	0	29	1	8	1	1	RI	No	Second mechelec	No	No	
28	CHECK TYPE(FEEDBACK)_	17.7	Op265	15	6	896	2	28	1	51	1	2	RI	No	Second mechelec	No	No	
29	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	11	Op1115	16	42	32	0	30	1	6	1	2	RI	No	Group Leader	No	No	
30	CONTROLLER ENABLE MISSING(RCM)_	11	Op1080	19	12	376	0	14	2	28	1	3	RI	No	Group Leader	No	No	
31	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	14.5	Op1115	20	39	442	1	0	2	7	2	3	SB	No	Operator	No	Yes	End off shift
32	FEEDING EMPTY 1_	21.8	Op1080	16	20	120	2	0	3	35	3	9	SB	No	IMS1	No	Yes	End off shift
33	LOWER INDEX(NOT ENABLED)_	20.8	Op265	17	14	250	2	0	5	69	1	13	SB	No	IMS1	No	Yes	End off shift
34	GUARD DOOR 1 OPENED_	15	Op1110	18	48	297	0	12	2	15	2	5	RI	No	Group Leader	No	No	
35	SEPARATING 1 LOADED_	20.8	Op1050	19	55	455	1	1	1	12	1	2	SB	No	IMS1	No	Yes	End off shift
36	PART 2 IN POSITION (LOADING PLATE)_	19	Op1130	20	4	573	0	0	25	207	11	77	SB	No	IMS1	No	Yes	End off shift
37	READ IDS(FINISHED MESSAGE), SB022/A_	29.9	Op1110	20	22	613	8	0	3	16	1	2	SB	No	Operator	No	Yes	End off shift
38	WRITE IDS 1+2(FINISHED MESSAGE)_	21.4	Op295	21	24	703	0	0	1	4	1	1	SB	No	Operator	No	Yes	End off shift
39	PLATEN IN POSITION SB015/F_	22.1	Op1040	16	15	74	1	60	1	2	1	1	SB	No	Operator	No	Yes	End off shift
40	GRIPPER 2 TRANSFER POSITION(FEEDBACK)_	23	Op1037	8	18	74	1	61	2	61	1	1	RI	No	Group Leader	No	No	

41	SEPARATING 2 LOADED_	22	Op1050	18	50	297	1	48	1	14	1	11	RI	No	Second mechelec	No	No	
42	GRIPPER CLAMPED(FEEDBACK)_	89.3	Op1140	20	36	531	0	33	1	16	1	10	SB	No	IMS2	Yes	Yes	End off shift
43	TAGGING SYSTEM SIM MISSING_	24.1	Op1120	6	9	737	0	15	5	30	2	6	SB	No	Operator	No	Yes	End off shift
44	GRIPPER CLAMPED(FEEDBACK)_	29	Op265	0	38	31	2	17	2	101	2	48	SB	No	Operator	No	Yes	End off shift
45	FAULT(BAR CODE SCANNER)_	13.3	Op1057	5	3	618	0	48	4	13	1	1	RI	No	Group Leader	No	No	
46	CLAMP GRIPPER(CRANKSHAFT)(YVQ182)_	26	Op30	1	58	219	10	43	1	4	1	2	RI	No	Second mechelec	No	No	
47	PASSWORD INPUT(NOK)_	67.1	Op1120	1	2	229	0	44	3	33	1	10	SB	No	Operator	No	Yes	End off shift
48	PASSWORD INPUT(NOK)_	40	Op1057	3	28	506	1	49	2	15	1	4	SB	No	IMS2	No	Yes	End off shift
49	PLATEN IN STATION(SQI122_	35.7	Op1010	4	52	559	5	64	3	53	1	2	RI	No	Group Leader	No	No	
50	FEEDING EMPTY 1_	11.3	Op1080	0	19	128	0	46	2	50	2	14	RI	No	Second mechelec	No	No	
51	PLATEN ENTERING(FEEDBACK)_	12	Op155	0	54	64	4	72	1	4	1	2	RI	No	Group Leader	No	No	
52	OVERTRAVEL PROTECTION TRIPPED_	17.1	Op1130	1	55	194	0	52	5	308	3	38	RI	No	Second mechelec	No	No	
53	OVERTRAVEL PROTECTION TRIPPED_	123.3	Op1130	5	1	623	2	41	17	320	15	50	SB	No	IMS2	No	Yes	
54	PART 2 LOST (TOOL)_	123.4	Op1125	8	46	45	0	74	1	224	1	5	SB	No	Operator	No	Yes	End off shift
55	GRIPPER CLAMPED(FEEDBACK)_	12.6	Op265	4	2	539	2	64	7	142	5	63	RI	No	Group Leader	No	No	
56	DISTRIBUTOR IN POSITION(FEEDBACK)_	5.4	Op1037	7	26	1	1	56	1	1	1	1	RI	No	Group Leader	No	No	
57	GUARD DOOR 1 OPENED_	7.9	Op1110	2	30	200	0	66	1	1	1	1	RI	No	Second mechelec	No	No	
58	OP1025>FAULT(DATA TRANSFER,ASM-HOST)_	8.6	Op1025	2	43	234	5	51	2	2	1	1	RI	No	Group Leader	No	No	
59	NO SHORTAGE SIGNAL EXIT 1, SQI227, SB091/F_	5.2	Op1095	3	21	413	1	4	2	2	1	1	RI	No	Group Leader	No	No	
60	RETURN SEPARATING(NOT ENABLED)_	3.2	Op1020	0	26	1	5	56	1	1	1	1	RI	No	Second mechelec	No	No	
61	GRIPPER CLAMPED(YVQ062)_	1.1	Op1010	1	40	173	5	51	1	1	1	1	RI	No	Second mechelec	No	No	
62	CLAMP CLAMPING UNIT 1(FEEDBACK)_	1.6	Op240	1	48	193	0	57	1	1	1	1	RI	No	Group Leader	No	No	
63	RETURN ADJUSTING UNIT_	1.6	Op270	1	57	200	2	66	1	1	1	1	RI	No	Group Leader	No	No	

A7. 2: *Attributes decisions data set - Decision maker DM2*

A_i $D_{i,DM2}$

Record	1.Type of fault	2. Repair Time	3. Machine number	4. Time	5. Engines Produced so far this shift	6. Parts waiting in the conveyor before the machine	7. Number of heads in the buffer	8. The machine has broken down this day	9. The machine has broken down this month	10. The breakdown has happened this day	11. The breakdown has happened this month	12 Action	13. Switch off the machine	14. Who	15. Ask the production Manager	16. Plan Repair	17. Plan repair – when?
1	CENTERING ADVANCED(SQI150)_	26.3	Op32	9 44	184	8	57	1	1	1	1	RI	No	Second mechelec	No	No	
2	GRIPPER 2 CLAMPED(FEEDBACK)_	22.6	Op1035	10 30	200	1	66	1	1	1	1	SB	No	IMS1	Yes	Yes	End off shift
3	VALVE COLLET 1.2 LOST._	171.8	Op1080	10 30	200	0	66	1	1	1	1	RI	No	Group Leader	Yes	No	
4	GUARD DOOR OPENED_	70.7	Op1105	10 30	200	0	66	1	1	1	1	RI	No	Group Leader	Yes	No	
5	OVERTRAVEL PROTECTION TRIPPED_	97.8	Op1130	13 10	595	0	0	23	23	1	1	SB	No	Operator	Yes	Yes	End off shift
6	INDEX(SQI151_	15	Op255	10 11	306	3	0	2	6	1	1	RI	No	Second mechelec	No	No	
7	INDEX LOWERED;EXIT 1 EMPTY, SQI141 NSQI124 NSPI170, SB027/A_	12.2	Op1095	11 32	353	0	17	1	4	1	1	RI	No	Second mechelec	No	No	
8	PART 2 IN POSITION (LOADING PLATE)_	23.5	Op1125	11 59	412	0	0	16	27	8	15	RI	No	Second mechelec	No	No	
9	FEEDING EMPTY 1_	12	Op1080	11 13	436	2	0	6	8	1	1	RI	No	Second mechelec	No	No	
10	OP1025>FAULT(DATA TRANSFER,ASM-HOST)_	18.5	Op1025	11 18	436	5	0	5	8	3	4	RI	No	Second mechelec	No	No	
11	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	77.8	Op1115	12 8	525	1	2	1	2	1	1	RI	No	Group Leader	Yes	No	
12	LIMIT SWITCH MONITORING BLOCKER 2_	219.2	Op1080	14 11	583	0	0	8	10	1	1	SB	No	IMS1	Yes	No	

13	PART 1 IN POSITION (LOADING PLATE)_	32.4	Op1130	15	16	667	1	4	22	49	12	25	RI	No	Second mechelec	No	No	
14	PART 2 IN POSITION (LOADING PLATE)_	59.7	Op1130	8	56	70	2	5	5	54	2	18	SB	No	IMS1	Yes	Yes	End off shift
15	PASSWORD INPUT(NOK)_	134.2	Op1120	10	33	136	0	19	1	7	1	3	RI	No	Group Leader	Yes	No	
16	PART 2 IN POSITION (LOADING PLATE)_	27.7	Op1130	11	34	270	0	10	10	59	3	19	RI	No	Second mechelec	No	No	
17	DISTRIBUTOR IN POSITION(FEEDBACK)_	22.7	Op1037	12	6	451	1	7	8	36	5	29	RI	No	Group Leader	No	No	
18	PART 1 IN POSITION (LOADING PLATE)_	18.5	Op1130	12	43	472	1	51	23	103	13	50	RI	No	Group Leader	No	No	
19	READ IDS(FINISHED MESSAGE), FEEDBACK(F9.1), SB022/F_	33.2	Op1110	8	57	69	1	12	2	7	1	2	RI	No	Second mechelec	No	No	
20	PART 1 IN POSITION (LOADING PLATE)_	23.8	Op1125	7	16	914	0	47	2	54	1	21	RI	No	Second mechelec	No	No	
21	LIMIT SWITCH MONITORING BLOCKER 2_	50.3	Op1080	8	27	132	1	6	1	22	1	5	SB	No	IMS1	Yes	No	
22	PART IN POSITION(FEEDBACK)_	28.9	Op1037	9	2	212	1	0	3	41	1	2	RI	No	Group Leader	No	No	
23	START DISPENSER(NOT ENABLED)_	40.6	Op1020	15	29	9	5	0	1	6	1	3	RI	No	Group Leader	No	No	
24	TAGGING SYSTEM(FAULT,CONNECTIONS)_	44.6	Op1120	18	58	303	2	0	2	18	1	1	RI	No	Group Leader	Yes	No	
25	GRIPPER CLAMPED(FEEDBACK)_	89.3	Op1140	20	55	556	0	0	2	16	1	10	SB	No	IMS2	Yes	Yes	End off shift
26	GUARD DOOR OPENED_	27.4	Op1125	13	20	665	0	0	20	82	1	1	RI	No	Second mechelec	Yes	No	
27	GRIPPER 1 CLAMPED(FEEDBACK)_	17.5	Op1035	15	0	883	0	29	1	8	1	1	RI	No	Second mechelec	No	No	
28	CHECK TYPE(FEEDBACK)_	17.7	Op265	15	6	896	2	28	1	51	1	2	RI	No	Second mechelec	No	No	
29	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	11	Op1115	16	42	32	0	30	1	6	1	2	RI	No	Second mechelec	No	No	
30	CONTROLLER ENABLE MISSING(RCM)_	11	Op1080	19	12	376	0	14	2	28	1	3	RI	No	Group Leader	No	No	
31	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	14.5	Op1115	20	39	442	1	0	2	7	2	3	RI	No	Group Leader	No	No	
32	FEEDING EMPTY 1_	21.8	Op1080	16	20	120	2	0	3	35	3	9	RI	No	Group Leader	No	No	
33	LOWER INDEX(NOT ENABLED)_	20.8	Op265	17	14	250	2	0	5	69	1	13	RI	No	Second mechelec	No	No	
34	GUARD DOOR 1 OPENED_	15	Op1110	18	48	297	0	12	2	15	2	5	RI	No	Group Leader	No	No	
35	SEPARATING 1 LOADED_	20.8	Op1050	19	55	455	1	1	1	12	1	2	RI	No	Group Leader	No	No	
36	PART 2 IN POSITION (LOADING PLATE)_	19	Op1130	20	4	573	0	0	25	207	11	77	RI	No	Second mechelec	No	No	
37	READ IDS(FINISHED MESSAGE), SB022/A_	29.9	Op1110	20	22	613	8	0	3	16	1	2	RI	No	Group Leader	Yes	No	
38	WRITE IDS 1+2(FINISHED MESSAGE)_	21.4	Op295	21	24	703	0	0	1	4	1	1	RI	No	Group Leader	No	No	
39	PLATEN IN POSITION SB015/F_	22.1	Op1040	16	15	74	1	60	1	2	1	1	RI	No	Group Leader	No	No	
40	GRIPPER 2 TRANSFER POSITION(FEEDBACK)_	23	Op1037	8	18	74	1	61	2	61	1	1	RI	No	Second mechelec	Yes	No	
41	SEPARATING 2 LOADED_	22	Op1050	18	50	297	1	48	1	14	1	11	RI	No	Group Leader	No	No	
42	GRIPPER CLAMPED(FEEDBACK)_	89.3	Op1140	20	36	531	0	33	1	16	1	10	SB	No	IMS2	Yes	No	
43	TAGGING SYSTEM SIM MISSING_	24.1	Op1120	6	9	737	0	15	5	30	2	6	RI	No	Second mechelec	Yes	No	
44	GRIPPER CLAMPED(FEEDBACK)_	29	Op265	0	38	31	2	17	2	101	2	48	SB	No	Operator	Yes	No	
45	FAULT(BAR CODE SCANNER)_	13.3	Op1057	5	3	618	0	48	4	13	1	1	RI	No	Second mechelec	No	No	

46	CLAMP GRIPPER(CRANKSHAFT)(YVQ182)_	26	Op30	1	58	219	10	43	1	4	1	2	RI	No	Group Leader	Yes	No
47	PASSWORD INPUT(NOK)_	67.1	Op1120	1	2	229	0	44	3	33	1	10	RI	No	Group Leader	Yes	No
48	PASSWORD INPUT(NOK)_	40	Op1057	3	28	506	1	49	2	15	1	4	RI	No	Second mechelec	No	No
49	PLATEN IN STATION(SQI122_	35.7	Op1010	4	52	559	5	64	3	53	1	2	RI	No	Group Leader	No	No
50	FEEDING EMPTY 1_	11.3	Op1080	0	19	128	0	46	2	50	2	14	RI	No	Group Leader	No	No
51	PLATEN ENTERING(FEEDBACK)_	12	Op155	0	54	64	4	72	1	4	1	2	RI	No	Second mechelec	No	No
52	OVERTRAVEL PROTECTION TRIPPED_	17.1	Op1130	1	55	194	0	52	5	308	3	38	RI	No	Second mechelec	No	No
53	OVERTRAVEL PROTECTION TRIPPED_	123.3	Op1130	5	1	623	2	41	17	320	15	50	SB	No	IMS1	Yes	No
54	PART 2 LOST (TOOL)_	123.4	Op1125	8	46	45	0	74	1	224	1	5	SB	No	IMS2	Yes	Yes
55	GRIPPER CLAMPED(FEEDBACK)_	12.6	Op265	4	2	539	2	64	7	142	5	63	RI	No	Second mechelec	No	No
56	DISTRIBUTOR IN POSITION(FEEDBACK)_	5.4	Op1037	7	26	1	1	56	1	1	1	1	RI	No	Group Leader	No	No
57	GUARD DOOR 1 OPENED_	7.9	Op1110	2	30	200	0	66	1	1	1	1	RI	No	Group Leader	No	No
58	OP1025>FAULT(DATA TRANSFER,ASM-HOST)_	8.6	Op1025	2	43	234	5	51	2	2	1	1	RI	No	Group Leader	No	No
59	NO SHORTAGE SIGNAL EXIT 1, SQI227, SB091/F_	5.2	Op1095	3	21	413	1	4	2	2	1	1	RI	No	Group Leader	No	No
60	RETURN SEPARATING(NOT ENABLED)_	3.2	Op1020	0	26	1	5	56	1	1	1	1	RI	No	Group Leader	No	No
61	GRIPPER CLAMPED(YVQ062)_	1.1	Op1010	1	40	173	5	51	1	1	1	1	RI	No	Second mechelec	No	No
62	CLAMP CLAMPING UNIT 1(FEEDBACK)_	1.6	Op240	1	48	193	0	57	1	1	1	1	RI	No	Second mechelec	No	No
63	RETURN ADJUSTING UNIT_	1.6	Op270	1	57	200	2	66	1	1	1	1	RI	No	Second mechelec	No	No

A7. 3: *Attributes decisions data set - Decision maker DM3*

A_i $D_{i,DM3}$

Record	1.Type of fault	2. Repair Time	3. Machine number	4. Time		5. Engines Produced Today	6. Parts waiting in the conveyor before the machine	7. Number of heads in the buffer	8. The machine has broken down this day	9. The machine has broken down this month	10. The breakdown has happened this day	11. The breakdown has happened this month	12 Action	13. Switch off the machine	14. Who	15. Ask the production Manager	16. Plan Repair	17. Plan repair – when?
1	CENTERING ADVANCED(SQI150)_	26.3	Op32	9	44	184	8	57	1	1	1	1	RI	No	Second mechelec	No	No	
2	GRIPPER 2 CLAMPED(FEEDBACK)_	22.6	Op1035	10	30	200	1	66	1	1	1	1	RI	No	Group Leader	No	No	
3	VALVE COLLET 1.2 LOST._	171.8	Op1080	10	30	200	0	66	1	1	1	1	SB	No	Operator	No	Yes	
4	GUARD DOOR OPENED_	70.7	Op1105	10	30	200	0	66	1	1	1	1	RI	No	Group Leader	No	No	
5	OVERTRAVEL PROTECTION TRIPPED_	97.8	Op1130	13	10	595	0	0	23	23	1	1	RI	No	Second mechelec	No	No	
6	INDEX(SQI151_	15	Op255	10	11	306	3	0	2	6	1	1	RI	No	Second mechelec	No	No	
7	INDEX LOWERED;EXIT 1 EMPTY, SQI141 NSQI124 NSPI170, SB027/A_	12.2	Op1095	11	32	353	0	17	1	4	1	1	RI	No	Second mechelec	No	No	
8	PART 2 IN POSITION (LOADING PLATE)_	23.5	Op1125	11	59	412	0	0	16	27	8	15	RI	No	Second mechelec	No	No	
9	FEEDING EMPTY 1_	12	Op1080	11	13	436	2	0	6	8	1	1	RI	No	Group Leader	No	No	
10	OP1025>FAULT(DATA TRANSFER,ASM-HOST)_	18.5	Op1025	11	18	436	5	0	5	8	3	4	RI	No	Second mechelec	No	No	
11	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	77.8	Op1115	12	8	525	1	2	1	2	1	1	SB	No	Operator	No	Yes	

12	LIMIT SWITCH MONITORING BLOCKER 2_	219.2	Op1080	14	11	583	0	0	8	10	1	1	SB	No	Operator	No	Yes	
13	PART 1 IN POSITION (LOADING PLATE)_	32.4	Op1130	15	16	667	1	4	22	49	12	25	RI	No	Second mechelec	No	No	
14	PART 2 IN POSITION (LOADING PLATE)_	59.7	Op1130	8	56	70	2	5	5	54	2	18	RI	No	Group Leader	No	No	
15	PASSWORD INPUT(NOK)_	134.2	Op1120	10	33	136	0	19	1	7	1	3	SB	No	Operator	No	Yes	
16	PART 2 IN POSITION (LOADING PLATE)_	27.7	Op1130	11	34	270	0	10	10	59	3	19	RI	No	Second mechelec	No	No	
17	DISTRIBUTOR IN POSITION(FEEDBACK)_	22.7	Op1037	12	6	451	1	7	8	36	5	29	SB	No	IMS1	No	Yes	End off shift
18	PART 1 IN POSITION (LOADING PLATE)_	18.5	Op1130	12	43	472	1	51	23	103	13	50	RI	No	Second mechelec	No	No	
19	READ IDS(FINISHED MESSAGE), FEEDBACK(F9.1), SB022/F_	33.2	Op1110	8	57	69	1	12	2	7	1	2	SB	No	IMS1	No	Yes	End off shift
20	PART 1 IN POSITION (LOADING PLATE)_	23.8	Op1125	7	16	914	0	47	2	54	1	21	SB	No	IMS2	No	Yes	End off shift
21	LIMIT SWITCH MONITORING BLOCKER 2_	50.3	Op1080	8	27	132	1	6	1	22	1	5	SB	No	IMS1	No	Yes	
22	PART IN POSITION(FEEDBACK)_	28.9	Op1037	9	2	212	1	0	3	41	1	2	SB	No	IMS1	No	Yes	
23	START DISPENSER(NOT ENABLED)_	40.6	Op1020	15	29	9	5	0	1	6	1	3	SB	No	Operator	No	Yes	End off shift
24	TAGGING SYSTEM(FAULT,CONNECTIONS)_	44.6	Op1120	18	58	303	2	0	2	18	1	1	SB	No	IMS1	No	Yes	
25	GRIPPER CLAMPED(FEEDBACK)_	89.3	Op1140	20	55	556	0	0	2	16	1	10	SB	No	Operator	No	Yes	
26	GUARD DOOR OPENED_	27.4	Op1125	13	20	665	0	0	20	82	1	1	RI	No	Group Leader	No	No	
27	GRIPPER 1 CLAMPED(FEEDBACK)_	17.5	Op1035	15	0	883	0	29	1	8	1	1	RI	No	Second mechelec	No	No	
28	CHECK TYPE(FEEDBACK)_	17.7	Op265	15	6	896	2	28	1	51	1	2	SB	No	IMS1	No	Yes	End off shift
29	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	11	Op1115	16	42	32	0	30	1	6	1	2	RI	No	Second mechelec	No	No	
30	CONTROLLER ENABLE MISSING(RCM)_	11	Op1080	19	12	376	0	14	2	28	1	3	RI	No	Second mechelec	No	No	
31	INDEX RAISED, FEEDBACK(SQI136) SB021/F_	14.5	Op1115	20	39	442	1	0	2	7	2	3	RI	No	Group Leader	No	No	
32	FEEDING EMPTY 1_	21.8	Op1080	16	20	120	2	0	3	35	3	9	RI	No	Group Leader	No	No	
33	LOWER INDEX(NOT ENABLED)_	20.8	Op265	17	14	250	2	0	5	69	1	13	SB	No	IMS1	No	Yes	End off shift
34	GUARD DOOR 1 OPENED_	15	Op1110	18	48	297	0	12	2	15	2	5	RI	No	Second mechelec	No	No	
35	SEPARATING 1 LOADED_	20.8	Op1050	19	55	455	1	1	1	12	1	2	SB	No	IMS1	No	Yes	End off shift
36	PART 2 IN POSITION (LOADING PLATE)_	19	Op1130	20	4	573	0	0	25	207	11	77	RI	No	Second mechelec	No	No	
37	READ IDS(FINISHED MESSAGE), SB022/A_	29.9	Op1110	20	22	613	8	0	3	16	1	2	SB	No	IMS1	No	Yes	End off shift
38	WRITE IDS 1+2(FINISHED MESSAGE)_	21.4	Op295	21	24	703	0	0	1	4	1	1	SB	No	IMS1	No	Yes	End off shift
39	PLATEN IN POSITION SB015/F_	22.1	Op1040	16	15	74	1	60	1	2	1	1	SB	No	Operator	No	Yes	End off shift
40	GRIPPER 2 TRANSFER POSITION(FEEDBACK)_	23	Op1037	8	18	74	1	61	2	61	1	1	SB	No	IMS2	No	Yes	End off shift
41	SEPARATING 2 LOADED_	22	Op1050	18	50	297	1	48	1	14	1	11	SB	No	IMS2	No	Yes	End off shift
42	GRIPPER CLAMPED(FEEDBACK)_	89.3	Op1140	20	36	531	0	33	1	16	1	10	SB	No	IMS2	No	Yes	
43	TAGGING SYSTEM SIM MISSING_	24.1	Op1120	6	9	737	0	15	5	30	2	6	SB	No	IMS1	No	Yes	End off shift

44	GRIPPER CLAMPED(FEEDBACK)_	29	Op265	0	38	31	2	17	2	101	2	48	SB	No	IMS1	No	Yes	End off shift
45	FAULT(BAR CODE SCANNER)_	13.3	Op1057	5	3	618	0	48	4	13	1	1	RI	No	Second mechelec	No	No	
46	CLAMP GRIPPER(CRANKSHAFT)(YVQ182)_	26	Op30	1	58	219	10	43	1	4	1	2	SB	No	IMS1	No	Yes	End off shift
47	PASSWORD INPUT(NOK)_	67.1	Op1120	1	2	229	0	44	3	33	1	10	SB	No	Operator	No	Yes	
48	PASSWORD INPUT(NOK)_	40	Op1057	3	28	506	1	49	2	15	1	4	SB	No	IMS1	No	Yes	End off shift
49	PLATEN IN STATION(SQI122_	35.7	Op1010	4	52	559	5	64	3	53	1	2	SB	No	IMS1	No	Yes	End off shift
50	FEEDING EMPTY 1_	11.3	Op1080	0	19	128	0	46	2	50	2	14	RI	No	Second mechelec	No	No	
51	PLATEN ENTERING(FEEDBACK)_	12	Op155	0	54	64	4	72	1	4	1	2	RI	No	Group Leader	No	No	
52	OVERTRAVEL PROTECTION TRIPPED_	17.1	Op1130	1	55	194	0	52	5	308	3	38	RI	No	Second mechelec	No	No	
53	OVERTRAVEL PROTECTION TRIPPED_	123.3	Op1130	5	1	623	2	41	17	320	15	50	SB	No	IMS1	Yes	No	
54	PART 2 LOST (TOOL)_	123.4	Op1125	8	46	45	0	74	1	224	1	5	SB	No	Operator	No	Yes	
55	GRIPPER CLAMPED(FEEDBACK)_	12.6	Op265	4	2	539	2	64	7	142	5	63	RI	No	Group Leader	No	No	
56	DISTRIBUTOR IN POSITION(FEEDBACK)_	5.4	Op1037	7	26	1	1	56	1	1	1	1	RI	No	Second mechelec	No	No	
57	GUARD DOOR 1 OPENED_	7.9	Op1110	2	30	200	0	66	1	1	1	1	RI	No	Group Leader	No	No	
58	OP1025>FAULT(DATA TRANSFER,ASM-HOST)_	8.6	Op1025	2	43	234	5	51	2	2	1	1	RI	No	Second mechelec	No	No	
59	NO SHORTAGE SIGNAL EXIT 1, SQI227, SB091/F_	5.2	Op1095	3	21	413	1	4	2	2	1	1	RI	No	Second mechelec	No	No	
60	RETURN SEPARATING(NOT ENABLED)_	3.2	Op1020	0	26	1	5	56	1	1	1	1	RI	No	Second mechelec	No	No	
61	GRIPPER CLAMPED(YVQ062)_	1.1	Op1010	1	40	173	5	51	1	1	1	1	RI	No	Group Leader	No	No	
62	CLAMP CLAMPING UNIT 1(FEEDBACK)_	1.6	Op240	1	48	193	0	57	1	1	1	1	RI	No	Group Leader	No	No	
63	RETURN ADJUSTING UNIT_	1.6	Op270	1	57	200	2	66	1	1	1	1	RI	No	Group Leader	No	No	